

**Official Transcript of Proceedings**  
**NUCLEAR REGULATORY COMMISSION**

Title:                   Advisory Committee on Reactor Safeguards  
                              Joint Human Factors Reliability and PRA,  
                              and Digital I&C Subcommittees

Docket Number:     (n/a)

Location:             teleconference

Date:                  Wednesday, November 15, 2023

Work Order No.:     NRC-2626

Pages 1-302

**NEAL R. GROSS AND CO., INC.**  
**Court Reporters and Transcribers**  
1716 14th Street, N.W.  
Washington, D.C. 20009  
(202) 234-4433

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23

DISCLAIMER

UNITED STATES NUCLEAR REGULATORY COMMISSION'S  
ADVISORY COMMITTEE ON REACTOR SAFEGUARDS

The contents of this transcript of the proceeding of the United States Nuclear Regulatory Commission Advisory Committee on Reactor Safeguards, as reported herein, is a record of the discussions recorded at the meeting.

This transcript has not been reviewed, corrected, and edited, and it may contain inaccuracies.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS

1323 RHODE ISLAND AVE., N.W.

WASHINGTON, D.C. 20005-3701

(202) 234-4433

[www.nealrgross.com](http://www.nealrgross.com)

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25

UNITED STATES OF AMERICA

NUCLEAR REGULATORY COMMISSION

+ + + + +

ADVISORY COMMITTEE ON REACTOR SAFEGUARDS

(ACRS)

+ + + + +

JOINT HUMAN FACTORS RELIABILITY & PRA,

AND DIGITAL I&C SUBCOMMITTEE MEETING

+ + + + +

WEDNESDAY

NOVEMBER 15, 2023

+ + + + +

The Joint Subcommittee met via  
Teleconference, at 8:30 a.m. EST, Vicki Bier, Chair,  
presiding.

COMMITTEE MEMBERS:

VICKI M. BIER, Chair

RONALD G. BALLINGER, Member

CHARLES H. BROWN, JR., Member

GREGORY H. HALNON, Member

JOSE A. MARCH-LEUBA, Member

ROBERT E. MARTIN, Member

WALTER L. KIRCHNER, Member

DAVID A. PETTI, Member

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25

JOY L. REMPE, Member  
THOMAS P. ROBERTS, Member  
MATTHEW W. SUNSERI, Member  
VESNA DIMITRIJEVIC, Member

ACRS CONSULTANTS:

MYRON HECHT  
STEPHEN SCHULTZ

DESIGNATED FEDERAL OFFICIAL:

CHRISTINA ANTONESCU

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25

TABLE OF CONTENTS

Opening Remarks by Chairman . . . . . 4

Introductory Remarks . . . . . 8

Data Science and AI, Regulatory  
 Applications Public Workshop,  
 AI Characteristics for Regulatory  
 Consideration, Workshop Findings . . . . . 14

AI Project Plan . . . . . 15

Future Focused Research (FFR) Project;  
 Using Machine Learning to Inform  
 Inspection Planning . . . . . 122

Overview of DOE's Light Water Reactor  
 Sustainability Program . . . . . 156

Plant Modernization Activities, LWRs AI  
 Industry Research and Accomplishments  
 Overview, Fire Watch Report . . . . . 190

The Good, the Bad, and the Ugly of AI in  
 Process Control . . . . . 242

Public Comments . . . . . 294

Closing Remarks by Chairman . . . . . 301

P-R-O-C-E-E-D-I-N-G-S

(8:30 a.m.)

CHAIR BIER: So the meeting will now come to order. I will start with my introductory remarks and then you guys will go in a couple of minutes.

This is the meeting of the Joint Human Factors Reliability and PRA, and the Digital I&C Subcommittees. I'm Vicki Bier. I'm going to be chairing this subcommittee meeting.

ACRS members in attendance, we have Charles Brown. Matt seems to be not here this morning. He may be coming in later. Jose, you're on line?

MEMBER MARCH-LEUBA: Yes.

CHAIR BIER: And Vesna on line?

MEMBER DIMITRIJEVIC: Good morning, everybody.

CHAIR BIER: Okay, we have Joy Rempe, Ron Ballinger, Dave Petti, Walt Kirchner, Greg Halnon, Tom Roberts, Robert Martin. And Steve Schultz, our consultant, is here. And do we have Myron Hecht? Is he here or online? He may also be joining later as a consultant.

MEMBER BROWN: Yes, Vicki, Matt said he would be coming in virtually at probably 10 o'clock or

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 so.

2 CHAIR BIER: Okay.

3 MEMBER BROWN: He had something. He had  
4 to go back to North Carolina.

5 CHAIR BIER: Oh, wow. Okay. Well, thank  
6 you for letting me know.

7 Christina Antonescu of the ACRS staff is  
8 the Designated Federal Official, or DFO, for this  
9 meeting.

10 Christina, can you confirm that we have  
11 the court reporter on line?

12 Can the court reporter speak up?

13 MS. ANTONESCU: Can the court reporter  
14 speak up, please?

15 (Off-microphone comment.)

16 CHAIR BIER: Okay, great. So there are  
17 going to be two separate, but related, purposes for  
18 today's meeting. First, staff and contractors are  
19 going to provide information briefings on how they are  
20 implementing the NRC's artificial intelligence, or AI,  
21 strategic plan for fiscal years 2023 to 2027, so along  
22 with their collaborators.

23 In addition, we have a speaker later this  
24 afternoon, Dr. Missy Cummings, who will also present  
25 on pluses and minuses of artificial intelligence. She

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 was originally hoping to come in person, but is going  
2 to be online due to schedule conflicts.

3 I wanted to clarify for the staff and for  
4 anyone else listening that Dr. Cummings is not going  
5 to be in any way commenting on the staff  
6 presentations. She was not asked to review them and  
7 her opinions are her opinions. They're not, you know,  
8 to be interpreted as a comment positive or negative  
9 about anything the staff is doing.

10 Mainly, once we have this briefing on the  
11 agenda, I wanted to take the opportunity to have just  
12 an educational briefing for the committee members, so  
13 that the members are all starting with a basic  
14 understanding of some of the key issues, especially  
15 that will be coming before the committee probably in  
16 years to come, rather than at this moment.

17 For background, the ACRS was established  
18 by statute and is governed by the Federal Advisory  
19 Committee Act, FACA. This means that the committee  
20 can only speak to its published letter reports. We  
21 hold meetings to gather information to support our  
22 deliberations. Interested parties who wish to provide  
23 comments can contact our office requesting time to do  
24 so. We also set aside about 15 minutes usually at the  
25 end of every meeting for comments from members of the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 public either in person or listening on line. We also  
2 welcome written comments.

3 The meeting agenda for today was published  
4 on the NRC's public meeting notice website, as well as  
5 the ACRS meeting website. The agenda and the ACRS  
6 website have instructions about how the public can  
7 participate. I don't believe we have any formal  
8 requests for making a statement to the subcommittee  
9 from members of the public yet, but people are always  
10 welcome to chime in.

11 Today is going to be conducted as a hybrid  
12 meeting, both in person and online. A transcript of  
13 the meeting is being kept and will be made available  
14 on our website. Therefore, we request that  
15 participants in the meeting should identify themselves  
16 before they speak and speak with sufficient clarity  
17 and volume so that they can be readily heard. And as  
18 the staff knows, I'm sure, please allow time for  
19 member questions. Members always have a lot of  
20 questions and comments. And it might also help to  
21 indicate which slide number you're on for people who  
22 are following along on line.

23 We have an MS Teams phone line for audio  
24 established for the public who wishes to listen to the  
25 meeting. Because of the nature of the online

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

(202) 234-4433

[www.nealrgross.com](http://www.nealrgross.com)

1 meetings, we will take a short break after each  
2 presentation, if needed, to allow for screen sharing  
3 and of course, we will also take the usual breaks  
4 during the meeting as needed.

5 A reminder, anybody who is online on  
6 Teams, please do not use the meeting chat features to  
7 conduct any sidebar technical conversations. Those  
8 should be oral so that they're captured in the  
9 transcript. And if you have any questions, you can  
10 contact the DFO, Christina, about issues that you  
11 would like to have raised or if you're having  
12 connection difficulties, et cetera.

13 So we are now ready to proceed with the  
14 meeting. Matt, it looks like you already have the  
15 slides shared.

16 The opening remarks will be from Mr. Vic  
17 Hall, who is Deputy Director of the Division of  
18 Systems Analysis in the Office of Nuclear Reactor  
19 Regulatory Research and after that, we'll be ready for  
20 the rest of the presentations.

21 So feel free to go ahead.

22 MR. HALL: Thank you, Vicki. Good  
23 morning, everyone. Vicki, I appreciate you and I  
24 share a namesake. I have on many occasions been  
25 called Vicki, usually in the school yard. Good to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 have a fellow Vic, Vicki in the room.

2 So I'll introduce myself to the room. My  
3 name is Vic Hall. I'm the Deputy Director in the  
4 Division of Systems Analysis. I joined the Office of  
5 Nuclear Regulatory Research earlier this year. If  
6 you're looking at your calendars, I got here right  
7 after this amazing team published the strategic plan.  
8 So I'm going to take full credit for the wonderful  
9 work that they did and today, I'll mention I'm  
10 extremely proud of the work they've done and part of  
11 my job is representing them and being able to  
12 introduce them today. So it's my honor to be able to  
13 introduce Matt and Anthony and Trey at the table who  
14 will be doing all the heavy lifting and under the  
15 spotlight.

16 I do want to express my gratitude to the  
17 subcommittee today. Vicki mentioned this is an  
18 information meeting for the members here, but I kind  
19 of disagree. I think it's an information meeting for  
20 everyone here. AI is moving so fast. It's an  
21 technology that has got such a head of steam. I'm  
22 staying up late which I shouldn't do and I'm watching  
23 Fox shows on TV and whether it's Jimmy Kimmel or Jimmy  
24 Fallon or whichever show, they're talking about AI.  
25 When the President puts out an Executive Order a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 couple of weeks ago, they're talking about on the  
2 news. It just amps up the game. It puts the  
3 spotlight on the technology that is either going to  
4 change the world or is going to scare the heck out of  
5 us. So the most important thing that we can do is  
6 we're not going to regulate AI. We can be ready for  
7 it. We can prepare where it's coming.

8 So when we have a meeting like this today  
9 that you put together, it's an opportunity for us to  
10 share the amazing work that we've done, but really to  
11 listen to what you have to say, take that into  
12 account, because we have to collaborate. We have to  
13 take all the opinions into account because technology  
14 is moving so fast.

15 I am really looking forward to today's  
16 meeting and again, when you walk the halls of the NRC  
17 and you're telling some folks what are you doing  
18 today? I'm speaking in front of ACRS. The reaction  
19 ooh, good luck. Expect a lot of questions and  
20 discussion. My answer to that is good. We're  
21 welcoming that today. We really look forward to your  
22 probing, your questioning attitude, and how we can  
23 improve, how we can be ready for something that's  
24 moving so fast.

25 So with that, I do want to put the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 spotlight back on really the stars of the show today.  
2 We've got Matt Dennis, Anthony Valiaveedu, and Trey  
3 Hathaway who are experts in the field. And our  
4 office, the Office of Nuclear Regulatory Research, is  
5 really a hub of world-class expertise. And when it  
6 comes to AI and AI in government, I don't think you'll  
7 find a finer group of gentlemen on the topic. They  
8 not only understand the policy, they not only put  
9 together the strategy, but they write the language.  
10 They write in Python. They do things that make me as  
11 an electrical engineer blush and I am truly honored to  
12 be able to work with them to represent them and to be  
13 able to share their work with you today.

14 Matt will give you a summary of our  
15 workshop which we held a couple of months ago. This  
16 was our fourth workshop on AI and data science. The  
17 first, I'll call it in the ChatGPT era when really I  
18 think the world has been awakened to what's coming  
19 with AI and it was our best -- it was widely attended.  
20 We have over 350 attendees from 12 countries. We had  
21 wonderful speakers from the national labs and  
22 universities just like we have today, so I'm very much  
23 looking forward to that same type of learning and  
24 interaction that we're going to have today.

25 Next, Anthony Valiaveedu will give you the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1       rundown on our strategic plan and the project plan  
2       which talks about all the different steps that we've  
3       laid out. They're going to get us to I'll call it  
4       success between now and the next five years and I  
5       guarantee you that plan will change. There's no way  
6       that plan can't change with the way the speed is  
7       changing. So we've done our best take and I think  
8       it's a pretty darn good take at what actions we need  
9       to do to be ready for what's coming from the industry.

10               I do want to mention again the fact that  
11       things are moving fast, well before this meeting was  
12       scheduled, well, after this meeting was scheduled, I'm  
13       sorry, the chair put out a tasking memo to the NRC  
14       staff. The title of that memo is Advancing the Uses  
15       of Artificial Intelligence. And in that memo he very  
16       much speaks about the need to be responsible in that  
17       use of AI. And that clearly is again, putting the  
18       focus on what the staff is willing to do to keep  
19       prepared for this technology that's coming very fast.  
20       And obviously, Nuclear Regulatory Research certainly  
21       has a role because have such a dense group of  
22       expertise in that. So we're working very much with  
23       our partners in the Office of Chief Information  
24       Officer and really every office in this agency because  
25       everybody will have a role in figuring out how we can

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 best responsibly use the technology to get our job  
2 done in our mission of safety.

3 So with that, I did want to mention a  
4 couple of folks in the room, Trey Hathaway as well.  
5 I've got one of our experts here that will be able to  
6 answer questions. And if you haven't met Luis  
7 Betancourt, Luis Betancourt is the branch chief in our  
8 division who really is the motor behind all this and  
9 makes it happen. And we also have Paul Krohn, I think  
10 is on the line, who is my co-chair on the AI Steering  
11 Committee. So again, it's been a team effort across  
12 the agency and it will continue to be so in the years  
13 coming.

14 And again, I just wanted to close with  
15 repeating my gratitude again. It's these types of  
16 meetings that will make us better. So thank you for  
17 having us. I hope you enjoy the presentations that  
18 we've prepared. I am extremely proud of the work and  
19 what we've accomplished in the last year and very much  
20 looking forward to actually what the next five years  
21 bring -- what the future bring for us as an agency.

22 So with that, Matt, let me hand it over to  
23 you and thank everyone again for their attention and  
24 for bringing us today.

25 CHAIR BIER: Okay, before we get going, I

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 just wanted to make one other announcement which is  
2 the committee is on an very strict schedule today  
3 because unlike the staff, the committee has a  
4 commitment at lunch hour and with an outside speaker  
5 later in the afternoon. We're going to have to kind  
6 of try and keep it on schedule. Happy to get going,  
7 so go ahead.

8 MR. DENNIS: Okay, I think I hit the  
9 button and the microphone is green so I'm good to go.  
10 Good morning, everyone. Again, my name is Matt Dennis  
11 and Trey Hathaway here. We're from the Office of  
12 Research and we'll be talking -- our first  
13 presentation this morning will be on the summary and  
14 finding of the AI public workshop which we had back in  
15 September. And I appreciate the push to get -- to  
16 have a lunch break. I have to go get my flu shot, so  
17 I am ecstatic to be on time so I can go get my flu  
18 shot.

19 Again, Matt Dennis, Trey Hathaway. We've  
20 already introduced Paul. Paul, Vic, and Luis are  
21 sitting over here at the side table and are available  
22 to answer any questions that we have that's related to  
23 the strategy or our progress.

24 Trey and I are going to talk about --  
25 we're going to talk about what the landscape is as we

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 see it right now, so just a brief. Though as a  
2 reminder, we came a year ago, exactly almost a year  
3 ago tomorrow, and presented on our draft AI strategy  
4 last year and following that, we did a public comment  
5 or we had a public comment period. We resolved the  
6 comments and moved forward with publishing the  
7 strategy, as Vic mentioned, in May of this year. So  
8 I'll talk just a little bit about where we are in that  
9 landscape, as far as it pertains to the workshop.  
10 Anthony will be discussing the project plan and a lot  
11 of what has come out of implementing the strategy in  
12 the last year since we talked to you. We've made a lot  
13 of good progress.

14 So I'll talk about the workshop overview.  
15 I'm sure a number of you were able to attend the  
16 workshop. So I will not be going into the nitty-gritty  
17 of the entire workshop, but instead, I'll be talking  
18 about our observations from our perspective about what  
19 was said and discussed at the workshop, so I'll talk  
20 about the workshop and all session summaries. I'm  
21 lucky to have some of the chairs who chaired those  
22 panel sessions here today participating in the meeting  
23 and so it is not just Trey and I who are here to  
24 discuss the workshop, but we have a number of staff  
25 across the offices that have been participating in an

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 AI working group that we have organized, specifically  
2 to address or discuss AI attributes for regulatory  
3 consideration and that working group is also tasked  
4 with planning the workshop. So it's not just the  
5 Office of Research. It's been a very collaborative  
6 effort with a lot of the offices across the NRC and  
7 some of those staff are participating today in this  
8 meeting, so they're also available to field questions  
9 should they come up.

10 So I'll talk about -- and then finally,  
11 the high-level observations and then where do we go  
12 from here following this particular workshop.

13 So this is the slide. It's very similar to  
14 the one we talked about last year except with some  
15 updates. So we recognize, you'll notice on the left,  
16 the box that says external. That's highlighted in  
17 blue with intent behind it because the focus of the  
18 strategy, the NRC's AI strategic plan is externally  
19 facing. So we recognize an industry wants to use  
20 artificial intelligence and in order to do that, we  
21 took a proactive approach two years ago, around 2021,  
22 to develop the strategy in order to prepare the staff  
23 to review and evaluate the uses of AI that may be an  
24 NRC regulated activity. So we developed AI strategic  
25 plan for that purpose and that was published, as was

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 mentioned in May 2023.

2           So some other things that are internal,  
3 Dave mentioned the chair's memorandum on advancing the  
4 use of AI at the NRC which just came out last month  
5 and is being -- we're in the process of standing up  
6 the response to that and getting our ducks in a row  
7 for that purpose. So that's internal.

8           Some other internal things are -- you may  
9 have seen in the news. As Vic mentioned, it's been  
10 front and center. The Biden administration put out an  
11 Executive Order, again last month, on federal actions  
12 for advancing use of AI in government. So there is a  
13 push at the executive level for not only agencies to  
14 get a handle on what AI means, but also to prepare the  
15 agencies for adoption of AI within their portfolios.  
16 So not only are we in the position of regulating our  
17 industry's use of AI, potential use of AI, but we are  
18 also -- we have to prepare ourselves to use the  
19 technology as well. And so there will be a  
20 forthcoming OMB memo on this directing us how to  
21 consider certain aspects of AI implementation at the  
22 agency.

23           We have also been quite involved in a  
24 number of outside activities that benefit our internal  
25 preparedness with implementing the strategy as well as

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 our internal preparedness for the Executive Order.  
2 There is no shortage of AI conferences, meetings,  
3 symposiums that come up that you can attend and those  
4 have been -- we have participated in a number of  
5 those. Just recently, there was a PSAM meeting, a  
6 topical meeting specifically on AI that we attended  
7 and there were at least three and a half full days --  
8 three full days of presentations globally. So  
9 clearly, this is an interest in the nuclear industry  
10 and we have made a concerted effort to keep up to date  
11 on what is going on.

12 We have also been participated in a number  
13 of activities outside the agency. Trey is involved  
14 with a standards group which we'll be talking about  
15 later. And not only just the nuclear field, we also  
16 participate in a number of conferences, workshops, and  
17 symposiums that are in the Department of Defense area.  
18 So we are looking to other agencies, DOD, DOT, FDA,  
19 other areas where this is also being used so that we  
20 are best prepared because this is a whole of  
21 government action, not just us.

22 So and then on the right, the box that  
23 talks about evidence building priority questions, we  
24 have from the Evidence Building Act of 2018 and a  
25 couple of priority questions that were added to one,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 wrangle our data for best use for AI and also look for  
2 areas where we could use the work for the agency. I  
3 mentioned the chair's memo and then you will be  
4 hearing another presentation later this morning on a  
5 future focused research program specifically on  
6 looking at AI. The future focused research program as  
7 an incubator and technology development area where the  
8 staff can look at AI usage has been incredibly useful  
9 in the Office of Research and is one of the programs  
10 that we've called out in the AI strategic plan and the  
11 way to prepare our staff to understand this technology  
12 and it has been very beneficial. So you will be  
13 hearing a presentation on that topic as well.

14 MEMBER KIRCHNER: Walt Kirchner here, I'm  
15 not sure where to start this. I warned Vicki I was  
16 going to ask this. Can you define what you mean by  
17 AI? And if there's -- you know, succinctly, because  
18 if one is going to regulate, quote unquote, whatever  
19 that means at this early juncture, then one has to  
20 have an understanding of what it is you're going to  
21 regulate in terms of nuclear applications in the  
22 industry. So could you share that with us?

23 MR. DENNIS: I will mention that there is  
24 a -- the entire -- the very first page of the  
25 introduction to the AI Strategic Plan has two

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 paragraphs that -- I won't read them verbatim. I'll  
2 give you the high points. But it is a broad  
3 definition and in the strategic plan there was a lot  
4 of effort put into what does AI mean. Unfortunately,  
5 that definition is quite broad. And so the umbrella  
6 of AI includes natural language processing, machine  
7 learning, deep learning, all the buzz words that you  
8 hear.

9 But when you boil it down to just a few  
10 key words, AI has the ability to emulate human-like  
11 perception, cognition, planning, learning, and  
12 communication, or physical action. And so our  
13 definition in that introductory paragraph of the  
14 strategic plan, the two paragraphs, is very much based  
15 on the National Defense Authorization Act of 2021 as  
16 Congress defined AI which, as I mentioned, for better  
17 or worse is a very large, broad definition. So to  
18 interpret that, we have gotten a little more specific  
19 for our purposes to clarify the difference between  
20 automation and AI-enabled autonomy, and it really is  
21 the cognition and decision-making portion of AI that  
22 is crucial for looking at it.

23 CHAIR BIER: I'm going to chime in a  
24 little bit and this is really in a way Walt's and my  
25 conversation, but one of the things that I get

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 concerned about in the definition is what today is  
2 considered AI may ten years from now be considered  
3 computer programming. So our regulations have to talk  
4 about, I think, what functions it's serving in the  
5 plan rather than what maybe the technology goes by.  
6 But that's just my personal opinion.

7 MEMBER MARCH-LEUBA: Yes, this is Jose.  
8 I was going to say something similar. I don't think  
9 we regulate Fortran. We use Fortran to write code,  
10 safety codes that have been used to verify the  
11 regulations are satisfied. In a global sense, I see  
12 AI as another type of Fortran. I don't think we're  
13 going to write regulations that apply to a concept, an  
14 abstract concept called AI. I mean we don't have a  
15 regulation for Fortran. Am I thinking wrong?

16 MR. DENNIS: On that note, what was just  
17 mentioned, we have also grappled with the same issue  
18 and in looking back at that definition that we talk  
19 about in the introduction of the strategic plan, this  
20 dichotomy of software versus AI and where we are  
21 currently is called out, so we do recognize it, that  
22 there is a difference between software and AI. And  
23 there is a sentence that says an overarching goal of  
24 AI is providing solutions that mimic human-based  
25 solutions and predictions for problems. So some of our

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 discussion has been focused on the fact that Fortran  
2 software programming is very rule based whereas AI is  
3 based on large sets of data and then can infer or  
4 create its own algorithm to then make decisions that  
5 mimic human behavior.

6           So to the point of the future-proofing of  
7 definition, that is also part of why it is so squishy  
8 right now and there is -- we just went to a meeting  
9 where this same topic of what do you mean when you say  
10 AI came up? And one of the presenters was discussing  
11 that said no one has a unified definition. Everyone  
12 has a different interpretation. So right now, the  
13 strategic plan does have a broad definition with the  
14 caveat that says the U.S. NRC in an area where it has  
15 not been previously reviewed or evaluated. So we're  
16 not going to go back to something that is Fortran code  
17 and now call it AI. We're going to be looking at  
18 going forward, specifically examples where we feel  
19 that it fits under this definition of AI that we have  
20 that is so broad.

21           MEMBER HALNON: This is Greg. I think in  
22 our last subcommittee meeting we talked about this as  
23 well and we entreated you guys that that should be a  
24 priority because if you're going to put a regulatory  
25 framework around something, you need to know what that

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 something is. And otherwise, we're either over  
2 regulating, under regulating. It's a real danger, so  
3 again, I'll say the same thing I did at previous  
4 meetings. It's very important, at least to me, to get  
5 a succinct definition, boundary, whatever you want to  
6 call it around what you are going to need to regulate.

7 When you get into somebody wants to made  
8 a modification of the plan, apply AI, and they say how  
9 do I do a 50.59 on it? They're going to need to be  
10 able to have a series of workshops for ten years to  
11 figure it out. No, we don't have that time. As Vic  
12 said, it's moving so quickly. So anyway, that's kind  
13 of a recurring comment I think that we're going to be  
14 making as well.

15 MEMBER BALLINGER: This is Ron Ballinger.  
16 I'd like to second that. I mean this is a case where  
17 we run the risk of getting into what I call the  
18 subjectivity trap. And that at some point, somebody  
19 has to decide where the line is and if that line is  
20 fuzzy or depending on the person that's using it, when  
21 I check off on Microsoft Word, the autofill thing,  
22 guess what, it's telling me what I should say. So  
23 I'll just reinforce what Greg was saying.

24 MEMBER BROWN: And I'll follow up if Ron's  
25 finished. I'm probably the most resident skeptic --

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 I'm Charlie Brown, most resident skeptic on the  
2 committee. I will echo Greg's comment in that you  
3 really need to know what you're going to regulate  
4 before you can know what you're going to do with AI.

5 Your comments about programming is rule  
6 based whereas AI is quote evaluation of data sets in  
7 developing an algorithm that then goes and determines  
8 what direction you may want to go. That's then  
9 subject to the bias of the algorithm mapper who,  
10 somebody has got to say how algorithms are going to  
11 get developed. And there are biases all over the  
12 place in terms of what subject you're using, number  
13 one, whether they're social or technical. And that  
14 gets into a world of uncertainty.

15 I'm just going to make this comment early  
16 so everybody can be very aggravated throughout this.  
17 Greg was right, why do you want to try to regulate or  
18 develop a rule or how will we regulate when you really  
19 don't know how to use all this stuff in the first  
20 place? My response, my thought process, is somewhat  
21 different in that we're primarily based -- we're  
22 responsible for the safety of the plants. Our reactor  
23 trip safeguard systems, major plant control systems  
24 whether they're called safety or safety-related or not  
25 related to safety, whatever definition you want to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 apply, they're plant-controlled systems.

2 You certainly don't want all AI is it's  
3 invasive to the point where you may at a process if  
4 you want to go develop a trip and it says oh, no,  
5 maybe you don't want to because I'm looking at this  
6 other data. And now you've got software, because it's  
7 all embedded, that is variable, and does not have  
8 really -- you really don't want something else other  
9 than people deciding what's safe and not safe. I mean  
10 if I were you all, I was the boss and I'm not, you're  
11 lucky from that standpoint, I would put the brakes on  
12 it. I would literally if I was going to try to  
13 regulate this world, I would not try to do it -- I  
14 read through your program. Obviously, we have  
15 questions, but there are certain things you want to  
16 maintain. That's the safety posture.

17 The way to actually go about this and find  
18 out what are the benefits, how can it be utilized is  
19 to put a roadblock up and say, hey, look, you will not  
20 use or attempt to use or propose the use of AI for any  
21 reactor safety systems, any reactor safety related  
22 systems, or other plant-controlled systems that have  
23 to start, stop, various components, move rods,  
24 whatever you want to call it. Now let the vendors go  
25 figure out outside that box, how AI can be adapted in

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 an overall plant configuration on ways where it may  
2 benefit on the non-operational aspects of the  
3 equipment we have in there that serve, brackets the  
4 problem that the NRC is going to have to deal with.  
5 This is, to me, this is just the latest hot button fad  
6 that everybody is thinking is the greatest thing since  
7 sliced bread and trying to integrate it and put it  
8 into a regulatory rule is just not possible at all  
9 based on the way it's done.

10 The biases are terrible. All you have to  
11 do is look at the learning trying to make autonomous  
12 cars work properly. That's fundamentally a  
13 combination of rule-based and/or some level of AI  
14 that people are trying to introduce, a lot of wrecks  
15 because you can't define all the things that it may  
16 see, all the sensors may see. So that's my opening.

17 CHAIR BIER: Yes, you've heard a lot from  
18 us and we haven't heard much from you.

19 MEMBER BROWN: Well, we had to -- I had to  
20 give a flavor. I'm -- I was not (audio interference)  
21 I'll pass.

22 CHAIR BIER: I guess two comments: One is  
23 in order to ban AI you first have to know what AI is.  
24 You can't --

25 MEMBER BROWN: I didn't say ban. I did

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 not say ban, Vicki. I said don't put it in places  
2 where --

3 CHAIR BIER: Yes, but, still, you have to  
4 know what it is you don't want in those places.

5 MEMBER BROWN: The stuff we have installed  
6 right now works quite well.

7 CHAIR BIER: The other --

8 MEMBER BROWN: It's not a matter of  
9 banning anything. It's a matter of putting in basic  
10 software which then stops the software from performing  
11 in a repeatable and predictable manner. It's being  
12 changed constantly. You don't know what you have.

13 CHAIR BIER: The other comment that I  
14 would add is I --

15 MEMBER BROWN: You can see we have a lot  
16 of different opinions.

17 CHAIR BIER: -- assume that the NRC is  
18 also looking not only at regulating industries of AI,  
19 but also at advancing what the agency itself may want  
20 to use AI for. So, anyway, with that --

21 MEMBER MARTIN: Well, thank you, Vicki.  
22 Vicki, I need my shot, too.

23 CHAIR BIER: Oh, okay.

24 (Laughter.)

25 MEMBER MARTIN: Bob Martin. A few years

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 ago I read Max Tegmark's Life 3.0. I hope that's like  
2 required reading for this crowd. He's an MIT  
3 professor or Harvard, or something like that in  
4 Boston. And I thought what struck me was the reason  
5 why we're hearing about this again is because it comes  
6 and goes, right? It's obviously the interconnectivity  
7 of the world and the 'net and it comes down to data,  
8 right? The amount of data that we have and the  
9 algorithms that we have can now process this data in  
10 a way that can fool us, right?

11 For us, we live in a space -- us, the ACRS  
12 and the NRC -- in a space of low-frequency, high-  
13 consequence events, maybe something broader than that,  
14 with the emphasis on low-frequency. And low-frequency  
15 is referring to our hazards, bad things that happen  
16 that you don't have a lot of data for. So invariably  
17 we have a data gap. And the one I'm going to be  
18 listening for is -- and I have a couple ideas, but I'm  
19 not going to beat them just yet, but how you might  
20 think that there is data or data could be created to  
21 serve really safety issues that are relevant to this  
22 (audio interference). You don't have to answer that  
23 now. It's a kind of a comment. But that's a  
24 sensitivity. We look at a certain -- a sliver of what  
25 AI can do here, but certainly, Vicki, you made the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 point there are process improvements and other things  
2 that I absolutely agree are relevant. But it wouldn't  
3 necessarily be for us.

4 MEMBER ROBERTS: And my two cent's worth.  
5 This is Tom Roberts, just following up on what Charlie  
6 said. It seems to me that both the agency and  
7 industry have struggled for probably decades to figure  
8 out how to implement software into plant control  
9 systems. And there's been a whole infrastructure of  
10 diversity and challenges to how much diversity you  
11 need, but it doesn't seem to me like a new issue. The  
12 whole idea of having deterministic software imbedded  
13 in the system has been a concern because you can never  
14 prove that you've gone through all those possible  
15 deterministic combinations and have 100 percent  
16 certainty that the software is going to do what it's  
17 required to do, nor that the requirements are  
18 complete.

19 But I'm just wondering if maybe we could  
20 talk through the morning session just how different  
21 that is for AI, because it just seems like an  
22 extension of the same problem.

23 MEMBER BROWN: Thank you, Tom.

24 MEMBER MARCH-LEUBA: Yes, this is Jose.  
25 I wanted to bring in another concept. We think -- we,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 ACRS members -- seem to be focused on reactor  
2 protection system applications. And that's --  
3 probably that's our bread and butter. But AI is going  
4 to be applied on analyzing data from nondestructive  
5 assay pipes, of measurements, and you don't have a  
6 chance of a person to looking through all of them. So  
7 you develop an AI system to look for flaws on piping.

8 I don't think anybody's proposing to put  
9 an AI system on a protection system. I mean, they  
10 want to do it for maintenance or for data processing.  
11 I mean, the pipe testing is the clear application that  
12 is going to come first.

13 CHAIR BIER: And I think we're going to  
14 hear some of those, yes.

15 MEMBER BROWN: That's why I said -- this  
16 is Charlie Brown. That's why I suggested separating  
17 what we know we really have a hard time dealing with,  
18 echoing Tom's comment, because we have struggled with  
19 how do you apply the software systems and make sure  
20 they're going to work when they're supposed to work,  
21 and then let industry develop all these things like  
22 maintenance data, data that you get out of  
23 experimental test facilities, where can AI help you  
24 evaluate the large quantity of data that you get that  
25 may help you define the physical boundaries you have

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 to deal with in the plants.

2           There's a lot of stuff outside of what I  
3 would call the basic operation and shutdown of the  
4 plant that need be -- or starting it up, et cetera.  
5 And you will learn an awful lot from that. And you  
6 really don't want to lose focus because you're focused  
7 -- you're driven to focus on what I call plant  
8 operation-type scenarios as opposed to what I'd call  
9 stuff that's outside of that. I'm just reflecting  
10 Tom's comment and Jose's, and maybe Robert's. I'm not  
11 quite sure. I was struggling a little bit --

12           (Simultaneous speaking.)

13           MEMBER BROWN: Yes. Well, I'm not  
14 ambiguous, so I'll --

15           CHAIR BIER: Let's try and move ahead.  
16 Yes.

17           MR. DENNIS: I will say I'm heartened that  
18 we -- there's a lot of synergy here. Everything you  
19 said we have brought up as a topic in our working  
20 group degree, so --

21           MEMBER BROWN: Well, I'll make one other  
22 comment: Over the last four years I have read  
23 numerous -- the IEEE is a body that just loves all the  
24 new stuff. If you look at the -- at least the  
25 Spectrum and a couple of the other journals, you will

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 find tremendous amount of articles which are very  
2 skeptical in the application of AI in terms of the  
3 biases and other type issues. How do you know you're  
4 getting stuff that's telling you the right answer with  
5 the algorithms and stuff?

6           And some -- it's just I'm surprised when  
7 you see that much in -- an organization that loves  
8 electrical software computers has now got the skeptics  
9 coming in through their publications showing some of  
10 the concerns that we've echoed right here relative to  
11 the difficulties. So that's kind of an outlier, but  
12 it is an organization that has a lot of people  
13 involved and loves this kind of stuff. And they're  
14 even skeptical. And they're publishing. That's the  
15 important part.

16           MR. DENNIS: I hope to, I don't know,  
17 answer, bring up -- there's many of these things that  
18 people just brought up. I'm talk a little bit about  
19 some of those.

20           The Data Science and AI Workshops, just to  
21 skim over this, is that we had four -- we had three in  
22 2021, we had the fourth one in September. The goal  
23 was to answer some of these -- not answer these  
24 questions, but at least get some insight on some of  
25 these questions that have just been discussed.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           In 2021 we recognized these were some  
2           -- three observations. Industry did have interest in  
3           regulatory guidance on this topic. There is an issue  
4           with data. The topic about limited data was brought  
5           up in the nuclear domain. And so aggregating and  
6           using data for these data-hungry applications was an  
7           area that was brought up. There has been some  
8           progress made on that, I think, at the national labs.

9           And as of 2021; so taking it back to 2021,  
10          we heard that probably maybe now, 2023, there would be  
11          some deployment of an AI ML application. That has  
12          borne out to be, I think, true. And then regulated  
13          applications, maybe in three to five years, so 2026.  
14          So that was the basis for our timeline and our  
15          strategic plan. And we've heard about two or three  
16          application areas where that may actually pan out.

17          So I guess to the point of narrowing down  
18          the definition, there are specific use case areas for  
19          AI ML. One example was brought up in nondestructive  
20          evaluation. That's one area that was discussed or  
21          presented on last year at the ACRS meeting. And  
22          that's made some significant progress towards actual  
23          implementation.

24          MEMBER KIRCHNER: Can you distinguish in  
25          your own mind -- because you -- all of you there up in

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 front of us are focused in this area -- the difference  
2 -- 10 years ago or 20 years ago this was big data and  
3 you had people like GE using it to improve the  
4 preventive maintenance on their jet engines, you had  
5 Steve up at NYU using it to predict when you could get  
6 a taxi in Manhattan, and so on.

7 So I personally never thought of that as AI.

8 So can you make -- what's the distinction  
9 between just harvesting data with algorithms and AI?  
10 Because to me AI was always the cognitive function  
11 that you mentioned as part of your definition. And  
12 like fusion, that was 50 years ago at the MIT media  
13 lab. We're still probably 50 years out from that kind  
14 of definition of AI.

15 So how do you -- is AI just the umbrella  
16 that you want to use and it's the current jargon? As  
17 Dick was saying, even the talk show hosts are using  
18 the jargon. But is this really just big data and  
19 better computers and smarter algorithms or is it  
20 really cognitive AI?

21 MR. DENNIS: So two differentiating  
22 aspects: You mentioned 10 years ago this was called  
23 big data. Twenty years ago it was called expert  
24 systems and recommendation (audio interference). So  
25 the point is well-taken that this is an evolving area

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 where 20 years ago something that we would call AI is  
2 not something we would necessarily think is AI today.

3 And from our perspective machine learning  
4 and AI -- we're focusing on safety-critical  
5 application use case areas. And so this isn't -- and  
6 so the question that has been brought up is -- the  
7 industry is using it in process improvement areas  
8 outside of safety-critical applications to make good  
9 business decisions, to infer things that -- to assist  
10 with things that are not regulated application areas.  
11 And one of those has been presented. I'll get to it  
12 later on the Corrective Action Program analyzer and  
13 the maintenance rule functional failure analyzer. At  
14 the workshop it was presented on. So those are areas  
15 where it's being used in non-safety-critical  
16 applications.

17 From our perspective part of the issue is  
18 whether or not it is autonomous or making a decision.  
19 There is a distinction between using AI ML for design  
20 purposes or AI-enabled (audio interference). The  
21 problem is the states-based in use case areas are so  
22 broad we're stuck with trying to wrangle all of it,  
23 but the near-term things that we see within the next  
24 three years probably are going to be in that design  
25 area where machine learning is used to make a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 recommendation or a prediction to a human and then  
2 that human has to do something with it.

3 And so we're left in a position where we  
4 need to be prepared to evaluate that instantiation of  
5 that machine learning or AI to make a recommendation,  
6 whereas it's very different from the way that it's  
7 been done currently. And so that's what we're  
8 presented with. And how do we basically -- if that  
9 application area is presented to us, how do we review  
10 and evaluate and make a technical finding?

11 MEMBER KIRCHNER: Yes, that helps because  
12 that narrows things down with still the idea that  
13 research should be looking at what's over the horizon  
14 as well. Okay. Thank you.

15 MR. DENNIS: Yes, and we are fully looking  
16 at -- I think it has been mentioned several times, is  
17 AI is not entering the control room at this point.  
18 That's been stated at the public workshops several  
19 times. And we do believe that we need to be prepared  
20 for that potential eventuality, but that is not  
21 something that is right in the near term. The near  
22 term is in using AI ML for design recommendation, that  
23 type of stuff.

24 So I think that's enough on this slide.  
25 I'll move onto the next one. The purpose of it was

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 for us to host a workshop. So the AI Working Group  
2 convened to prepare for a workshop and provide  
3 feedback on regulatory and technical issues  
4 surrounding AI usage in nuclear applications. All of  
5 this, all of these workshops, the previous three,  
6 informed the preparation of the strategic plan. And  
7 this was the first workshop we held after the issuance  
8 of the strategic plan and preparation for the project  
9 plan, which you'll hear more from Anthony on. But all  
10 of this was to prepare us for what is going on.

11 And so we had three panel sessions. The  
12 first one on regulatory perspectives. The second one  
13 which was more academic in nature on safety, security,  
14 and explainability topics. And then the third one was  
15 more industry-focused on those AI application  
16 considerations and some of the examples of use case  
17 areas where it is being considered from industry.

18 So this is just a snapshot of the agenda.  
19 It was a 10:00 a.m. to 4:00 p.m. meeting. I won't go  
20 into this other than just to point out that all of the  
21 presentations are available in ADAMS and on our  
22 website, which there was a link on the previous page.

23 But you can see here we had a number of  
24 presenters. We had CNSC, ONR, IRSN, and a think tank,  
25 Responsible AI Institute, present on regulatory

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 perspectives. Again, I said it was more academic-  
2 focused in the second panel session. And then some of  
3 the industry presenters who are actually participating  
4 here today were in our panel session in the afternoon  
5 talking about how they're using it.

6 So all of this was to support the  
7 strategic plan and build out and build upon a table  
8 that we had in the strategic plan that talked about  
9 the notional AI and autonomy levels in commercial  
10 nuclear activities. It was what we put in the  
11 strategic plan to start the conversation which is  
12 happening here today and has happened at every single  
13 meeting about where AI is being inserted into nuclear  
14 activities.

15 The table had a range 1 to 4, from 1 being  
16 just basic -- something that's making a recommendation  
17 all the way up to level 4, which would be more like  
18 what was talked about about autonomous operation of a  
19 vehicle, so where you're actually running a power  
20 plant using AI.

21 So that was to frame the discussion. And  
22 we wanted to use that as a springboard for our working  
23 group who has discussed all of these things that have  
24 been brought up. We have gone back and forth and  
25 talked about these. But that was our starting point

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 for the workshop and to develop our matrix of AI  
2 characteristics for regulatory consideration that I'll  
3 talk about in the next slide.

4 So our working group was in a very agile  
5 fashion, as has been all of this with AI, because  
6 again it -- when we started the strategic plan two  
7 years ago on this journey ChatGPT didn't exist. And  
8 then things changed so we had to pivot the way that  
9 we're looking at this and the way that it's being used  
10 in industry. So we've had to be very agile.

11 So these are the members of the working  
12 group. Again, some of them are available today to  
13 answer any questions. But the disclaimer for our  
14 portion of the presentation as far as AI  
15 characteristics for regulatory consideration is we are  
16 aware that they're -- NIST is the agency chartered for  
17 the Federal Government to develop the AI risk  
18 management framework. So we're aware of these. What  
19 we presented was not an exhaustive list and we  
20 recognize that it's on a broad spectrum. So this is  
21 quite a large matrix with a range of applicability.

22 So this is the NRC staff's presentation at  
23 the workshop focused on these eight characteristics:  
24 safety significance, AI autonomy, safety,  
25 explainability, model life cycle, regulated activity,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 regulatory approval, and application maturity.

2 Again since this is just a recap, I won't  
3 go into all of these other than to say we had a  
4 discussion on all of them and we didn't get much  
5 feedback on it, so take that as you will. Either  
6 there was agreement that these are all concerns or  
7 just recognition that these will be part of the matrix  
8 of decision making that goes into considering AI  
9 applications and usage in the nuclear domain.

10 MEMBER BROWN: Did you mean feedback from  
11 us or from your workshop?

12 MR. DENNIS: From the workshop, yes.

13 MEMBER BROWN: Thank you. It would have  
14 been a long letter.

15 (Laughter.)

16 MR. DENNIS: Yes. So I guess the takeaway  
17 here is we don't have an answer for all of these  
18 things right now. This is what the working group kind  
19 of coalesced on as some characteristics. And I will  
20 say we are -- Anthony will mention this later -- we're  
21 participating in a trilateral working group with CNSC  
22 and ONR, and these line up quite well with other  
23 considerations from international -- our international  
24 regulatory counterparts as well as the IAEA. All of  
25 these topics have come up over and over again. So

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 we're in good company, I guess you could say.

2 So regulatory perspectives on AI panel  
3 session. I want to get to this so that we can keep on  
4 schedule. This is some of our observations on the  
5 three panel sessions, so the next three slides I'm  
6 going to give a synopsis with the disclaimer that this  
7 is a summary of some of the comments that were  
8 provided during those panel session discussions and  
9 presentations.

10 So the first one again was regulatory  
11 experts and safety experts from other regulatory  
12 entities globally and domestic think tanks. So CNSC  
13 pointed out that they have stood up their Disruptive  
14 Innovative and Emerging Technologies Working Group,  
15 DIET. I don't know if they forced that acronym to be  
16 fun for DIET or if it came out -- I think it actually  
17 -- they added it to make it DIET. But they  
18 commissioned a study last year on how the CNSC can be  
19 prepared for AI applications.

20 And the U.K. ONR -- I skipped one past,  
21 but the U.K.'s ONR has also issued a report.  
22 Similarly the U.K. ONR report, I will point out, has  
23 a very nice appendix on how -- on some different  
24 methodologies to consider how you would evaluate or  
25 review AI applications. And then they have done a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 regulatory sandbox. So this is one of the  
2 observations we made is that they've done two AI ML  
3 applications as part of a regulatory sandbox. I don't  
4 have the link here unfortunately, but they just this  
5 week published a report on this. So that is a  
6 publicly-available report now on their website.

7 But they believe that their regulatory  
8 approach is capable and flexible enough in the absence  
9 of standards, which everyone has recognized is a  
10 shortcoming, because standards are currently in  
11 process for being developed right now. We recognize  
12 that it can take a long time to get a standard through  
13 the process. Even NIST commented that at the  
14 standards forum at the NRC a couple months ago. So  
15 NIST knows this; we know this, but the ONR thinks they  
16 have a flexible enough framework that they can move  
17 forward without standards, if need be.

18 IRSN, again they are under the umbrella of  
19 the E.U. AI Act, and so some of their key areas for  
20 high-risk AI applications that they called out are  
21 data governance, risk management, and the human  
22 component, which keeps coming up again and again in  
23 discussions.

24 The Responsible AI Institute has been  
25 working on a certification methodology. Their view is

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 bigger than nuclear. It includes fair housing, all  
2 the AI application areas. So they've been trying to  
3 develop a certification framework. They do have one  
4 and it's largely based on these two -- the AI RMF from  
5 NIST and an ISO AI management systems approach. So  
6 they discuss their work on developing certification  
7 methodologies.

8 MEMBER KIRCHNER: Matt, I could point out  
9 that getting a boiler and pressure vessel code  
10 standard (audio interference) is -- and that's as  
11 well-defined problem. Often takes years. But more  
12 relevant here, what about your companion agencies here  
13 in the government? I'm thinking in particular FAA  
14 must be looking at this because of congestion in the  
15 air and so on and using advanced techniques to avoid  
16 collisions, whatever the application might be. So  
17 are there counterparts to you, NRC, here in the  
18 government that you're also at that are using further  
19 applications?

20 MR. DENNIS: Yes. I will point out that  
21 we are in good company. I mentioned it at the  
22 beginning and a little bit in passing that we are --  
23 we stay in contact with FAA. FDA presented on their  
24 regulatory approach for AI at the RIC last year. We  
25 invited them and have talks with them.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           Trey and I attend a number of DoD meetings  
2           for the Navy. So a lot of presenters there from the  
3           Navy, from the Army, and they're grappling with the  
4           exact same issues that we are. And they don't have an  
5           answer either. And the standards thing came up again.  
6           So the Army has been doing work on autonomous vehicle  
7           operations, and they have a Testing, Evaluation,  
8           Verification and Validation Working Group that's  
9           looking at this as well.

10           So we've stayed -- tried to stay plugged  
11           in with all of our federal partners that are working  
12           in this area to leverage learnings from them and  
13           research that they're conducting.

14           MR. BETANCOURT: Matt, can I mention  
15           something quick on that one?

16           MR. DENNIS: Yes.

17           MR. BETANCOURT: This is Luis Betancourt  
18           from the staff. On the FDA side, like Matt mentioned,  
19           when you look up on that table that he put on the  
20           model life cycle, that was one of the things that we  
21           learned from them, that they actually released some  
22           draft guidance on locked models, with some open  
23           models. So we have been actively involved in learning  
24           from them and vice versa. So there has been that  
25           synergy and basically cross-pollinization between us

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 and other agencies.

2 MEMBER KIRCHNER: Not to make any  
3 observation on our government, but the difficulty --  
4 to the extent that the DoD activities are in the open  
5 is fine, but I think one of the big challenges in AI  
6 applications is transparency and openness. The  
7 military has its needs and often those needs require  
8 classification and such, but the other agencies you  
9 mentioned: FDA and FAA, that's why I brought them up  
10 because they obviously also have to convince the  
11 public that any applications that they were to use  
12 would be transparent, safe, the integrity, all the  
13 issues that the NRC has to (audio interference).

14 MR. BETANCOURT: And on that, Walter --  
15 this is Luis Betancourt again. I'm going to be quick  
16 because I know that we're running out of time. Vic  
17 and I, we are attending meetings of the responsible AI  
18 officials from other agencies. So to your point,  
19 like, yes, the Department of Defense has their needs,  
20 but there's also this big push by the government of  
21 hey, we need to be able to regulate AI, but also how  
22 do we do AI responsibly internally?

23 So we're keeping tabs really well on what  
24 is happening, not only with the defense industry, but  
25 also other industries as well.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 MR. DENNIS: So the next session, the  
2 panel session was on a more academic nature. It was  
3 an excellent session because it really did discuss  
4 some of things about how would we evaluate the  
5 technology? And there is a lot of research coming on  
6 on this topic right now. NIST is looking at using  
7 combinatorial methods. The presenter from NIST is  
8 actually one of their funded research projects to  
9 support the AI RMF. So NIST is undertaking these  
10 projects to -- and one of the messages from the NIST  
11 presenter was that this is different from conventional  
12 assurance processes for autonomous or software-based  
13 systems and there are alternative methods that they're  
14 looking at that can go to that explainability problem  
15 that keeps getting brought up for AI.

16 George Mason was recognizing that there is  
17 an issue with explainability and that using  
18 counterfactual testing is one method that could be  
19 used. And they have a research project that's ongoing  
20 right now to use counterfactual cases to expose the  
21 black box nature of AI models.

22 And I see we have a hand. I don't know,  
23 Vicki, if it came up, but --

24 (Audio interference.)

25 MR. DENNIS: Okay. All right. Okay.



1 I'll keep going then.

2 All right. Georgia Tech discussed some of  
3 their cybersecurity research. They have a test  
4 facility that they are looking at multi-layered tests  
5 using a honey pot scenario to do cybersecurity  
6 monitoring using AI ML.

7 And NC State talked about one that was --  
8 the word was mentioned earlier about uncertainty  
9 quantification. And this is an area near and dear to  
10 my heart, on using VVUQ, verification, validation, and  
11 uncertainty quantification methods to root out the  
12 black box nature of deep neural networks. So there  
13 were a few examples that were giving. Monte Carlo  
14 dropout, deep ensembles, and Bayesian neural networks  
15 were talked about and an example was given in an  
16 application area to predict axial neutron flux  
17 profiles.

18 So the presenters, all of these presenters  
19 from the academic session were really talking about  
20 issues that we have and ways that the research is  
21 being used to try to explain AI in a way that can be  
22 understood, which is quite of interest to us at the  
23 NRC.

24 And then the application consideration  
25 panel was more industry-focused. We had several

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 presenters. The first one from Constellation and  
2 Jensen Hughes mentioned that in the absence of  
3 industry-specific V&V guidance for software that's  
4 driven by AI they've come up with their own process.  
5 So they talked about their V&V documentation that they  
6 developed and how they're looking at it from an  
7 explainability perspective so that their users within  
8 Constellation are able to understand the model and  
9 what it's doing as well as anyone externally that may  
10 be evaluating that model. So they have made some  
11 significant progress in being able to explain how  
12 their AI-driven CAP analyzer is actually functioning.

13 The Utility Service Alliance talked about  
14 their Phase 1 projects in their Advanced Remote  
15 Monitoring Project. I think INL later in the  
16 afternoon will be talking about a couple of these  
17 actually, so I won't go into great detail here. But  
18 one point that they made was that they assessed that  
19 the regulatory readiness level is at a two out of five  
20 and they are planning for a Phase 2 where they're  
21 going to explore more AI-drive autonomous inspection  
22 rounds and response projects. So they do have an  
23 interest in this area.

24 MEMBER REMPE: (Audio interference) these  
25 applications I could see how yes, you might be able to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 use the system to help detect, but the response aspect  
2 is where I'm curious. Because sometimes, for example,  
3 a fire suppression system could have adverse impact on  
4 the staff in the plant if there's not a human to view  
5 the AI detection and say yes, I agree with it; let's  
6 do this. And it takes a mitigating strategies with  
7 the staff before the system goes online.

8 And if you have to have that human review the  
9 data, does it really -- does it not add more time and  
10 just having the human do the fire watch?

11 And are those kind of questions coming up?  
12 Because that's one example, but it seems like there  
13 would be other examples where you don't want the  
14 software to initiate an inaction. And it's not a  
15 criticality in the control room thing. I'm just  
16 thinking about other actions that happen in the plant.  
17 Because of my experience at the lab, I know where bad  
18 things can happen in some of these systems and I -- is  
19 that coming up and people are thinking about that?

20 MR. DENNIS: Absolutely 100 percent. We  
21 have had a number of discussions with our human  
22 factors folks. And in other presentations I talk  
23 about a Tesla crash where the system basically --  
24 there's an accident, but the system defaults to the  
25 human, but the human only has three seconds to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 respond.

2           So we're talking about application areas  
3 where right now humans are in the loop. And there is  
4 no intent to take -- as far as AI being used  
5 necessarily to take the human out of the loop, but  
6 there is a concern or a recognition that humans cannot  
7 sit and just kind of toil away and be completely  
8 oblivious and then be expected to then respond  
9 immediately to something that is a time-sensitive or  
10 critical thing. And we know that from the existing  
11 control room configurations and automatic systems.

12           So there is a recognition that if you're  
13 going to go -- there's sort of a blended area here  
14 that's problematic where you have a human and then  
15 autonomous operation that's AI-driven. So we do  
16 recognize that there is a human factors component.  
17 And that's one of the things we actually called out in  
18 our AI characteristics for regulatory consideration  
19 was this concern.

20           MEMBER MARTIN:       (Audio interference)  
21 follow that one. From what you said there about the  
22 human, human's role, you hear more about AI's  
23 performance, but -- and opportunities for applying AI  
24 where we can apply that capability. What I've not  
25 heard -- and I'm not talking about just today, but

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 applications of AI to drive human performance, improve  
2 human performance, setting standards.

3 I think about when we were all younger and  
4 you hear about IBM's Deep Blue, right? When it began  
5 the better players could win. And then eventually it  
6 beat the masters. And the same thing could be  
7 considered here. Instead of focusing on letting AI  
8 take a call, AI could be used to set the standard for  
9 how humans perform. Obviously it's not the only  
10 thing. We're not talking about replacing all training  
11 with a robot. But nonetheless, I've not heard that.

12 And I think we need to think more human-  
13 centric on these things and not machine-centric and  
14 expose some bias. But we'd like to see the future  
15 focus. And it could be a question on later  
16 presentations on where we could take AI to review.

17 MR. DENNIS: Thank you for that  
18 distinction. And it has been one -- I will mention  
19 the industry has presented on using it for that  
20 purpose, using machine learning for improving operator  
21 examination, really using the tool to make us the best  
22 version of ourselves. So that is definitely in the  
23 use case area. I know we're not -- doesn't get all  
24 the credit or focus here, but it is an area of use.

25 MEMBER HALNON: This is Greg. I got a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 question. Constellation said there's no clear  
2 specific guidance for validation. We have a software  
3 quality assurance (audio interference) program,  
4 numerous regulatory documents. How far off are those  
5 if you were to lay those into AI? Is it a starting  
6 point only or is it N/A? Is it not applicable? How  
7 does that look?

8 MR. DENNIS: Going back to -- I'm not  
9 going to go back in the slide deck, but that --

10 MEMBER HALNON: I'm sorry. I had to leave  
11 for a few minutes so I apologize.

12 MR. DENNIS: Oh, I didn't talk about it.

13 MEMBER HALNON: All right.

14 MR. DENNIS: I'm not going to flip back,  
15 but that is one of the areas where we say we have a  
16 foundation of excellent guidance on software quality  
17 assurance, a VVUQ for modeling and simulation. We  
18 should start from that point.

19 MEMBER HALNON: So it's a starting point?

20 MR. DENNIS: That is a starting point.  
21 And so the example that Constellation gave, that's  
22 what they started with.

23 MEMBER HALNON: Okay.

24 MR. DENNIS: They started with the typical  
25 process you use for software V&V and then layered on

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 some unique stuff for the AI-driven aspects of it to  
2 build that out.

3 MEMBER HALNON: Okay. So it's not a  
4 perfectly round wheel, but it's certainly starting  
5 with a wheel?

6 MR. DENNIS: Yes. And that was our  
7 observation for the working group. And one of our  
8 eight characteristics was you start -- we start from  
9 what we have, and a lot of that is good. And the  
10 observation that Anthony will point out is we're  
11 starting a project this year as we speak to go into  
12 that aspect of looking at -- we're doing a gap  
13 analysis right now and then we're going to look at  
14 what methodologies could be used.

15 MEMBER HALNON: Okay. Thanks.

16 MR. DENNIS: So I will quickly go through  
17 the last two. Westinghouse emphasized the importance  
18 of having an ethical AI corporate policy and a  
19 recognition that the human is not the best interpreter  
20 of AI. So there needs to be some component to the  
21 uncertainty quantification through validation metrics  
22 that are interpretable by the human, but not -- this  
23 goes to the point of you may see an AI-generated  
24 image, computer-generated image and you think it's  
25 real. So the human can be easily spoofed.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           TerraPower discussed that there's no  
2 specific AI use cases or plans to use AI, but that  
3 highly passive future designs that do have potential  
4 for this use case -- and that there were some high-  
5 level thoughts presented on AI -- using AI for  
6 engineering document preparation and that we need to  
7 consider how to validate AI recommendations for  
8 licensed operators and if we should reevaluate the  
9 role of the human operator and what they play in the  
10 plant. So this was sort of the point that was brought  
11 up just a minute ago.

12           Our key takeaways: I think I have two  
13 more slides, so I'm going to be real quick. The panel  
14 sessions confirmed that we remain well-informed of  
15 international AI regulation and domestic projects.  
16 There were no surprises or show stoppers. So I guess  
17 the message here is that we feel that we've been doing  
18 a pretty good job of keeping the beat on what is going  
19 on for use cases and applications within the nuclear  
20 industry.

21           We did hear a lot of feedback on the  
22 regulatory sandboxes and how those provide a unique  
23 opportunity for industry and regulators to  
24 collaborate. So there is some interest in that topic  
25 area. And industry representatives encourage

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 continued collaboration to pursue pilot studies and  
2 proof of concepts for a future foundation for  
3 reviewing the use of AI in NRC-regulated activities.

4 Some of these considerations we have  
5 already talked about, but I guess I'll just  
6 reemphasize them as a transition to Anthony's  
7 presentation. We are currently in the mode of looking  
8 at what traceable and auditable evaluation  
9 methodologies exist in order -- and this is the  
10 project I mentioned that we're going to be kick-  
11 starting right now to do that.

12 And then we're also -- the workshops have been  
13 supporting our ability to understand what licensees  
14 and applicants are using in AI.

15 The future goes towards differentiating  
16 this for design versus AI-enabled autonomy. I did  
17 mention that design usage is the one that seems to be  
18 front and center as a use case. And then also how are  
19 we going to explain and evaluate it? Is this a  
20 reliability or assurance argument methodology? So  
21 those are the things in the future.

22 And of course all of this is predicated on  
23 our budget and preparation for these emergent industry  
24 applications, which, like ChatGPT came up, there's a  
25 whole slew of different ways to use it, for generating

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 documents and using it for regulatory applications  
2 potentially. So that was stuff that wasn't envisioned  
3 when we originally were writing the strategy.

4 So moving forward we continually focus on  
5 safety and security. That's paramount. Our  
6 partnerships, as we mentioned, with domestic and  
7 international counterparts and our engagement with  
8 other federal agencies has been very beneficial and  
9 we're continuing to pursue those. And we recommend  
10 and encourage our stakeholders to engage with the NRC  
11 early and often on plans and operating experience  
12 about how they're potentially going to use AI or  
13 looking to use it and what their experience has been.  
14 And we've gotten a lot of that feedback from the  
15 workshops and it's been very beneficial.

16 Our internal working group will be  
17 continuing to focus on AI characteristics for  
18 regulatory consideration following our feedback that  
19 we get from our gap assessment which we are currently  
20 in the process of going through and will be concluding  
21 in spring of 2024. That will also be providing the  
22 content for our next workshop which we do plan to have  
23 in summer of 2024.

24 So I do believe that takes me -- that is  
25 the end of my slides. So thank you very much.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 CHAIR BIER: Yes, I think one of the  
2 people online has an open mic. If you can check that.  
3 Thank you.

4 And, Anthony, now I think we can move onto  
5 you. Thanks.

6 MR. VALIAVEEDU: Well, thank you again for  
7 allowing me the opportunity to speak to you all today.  
8 My name is Anthony Valiaveedu. I'm part of the  
9 Nuclear Regulatory Commission working out of the  
10 Office of Research as a data scientist. Here with me  
11 at the table today, as previously introduced, is Matt  
12 Dennis and Trey Hathaway, who also work in the Branch  
13 of Accident Analysis. Special thanks to our  
14 management team including Paul Krohn from Region I,  
15 who is a division director; Victor Hall, who is a  
16 deputy division director for Division Systems  
17 Analysis; and Luis Betancourt, who's leading our  
18 branch in the Accident Analysis Branch.

19 This presentation today is only been  
20 possible through the efforts of the entire agency.  
21 All the program offices were involved during the  
22 development of the agency's strategic plan towards AI.  
23 Paul is from Region I. And pictures used throughout  
24 this presentation today (audio interference) staff  
25 members including from the Office of Nuclear Reactor

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 Regulation as well as the Office of Research. This  
2 supports again the notion that AI has the potential to  
3 touch every portion of this agency's mission of safety  
4 and security. As indicated with this graphic we've  
5 collaborated with a variety of program offices  
6 including NMSS, NSIR, the regions, OCIO, OEDO, and  
7 many others.

8 Over the past few years we were notified  
9 by various stakeholders that they have had plans to  
10 implement artificial intelligence into their current  
11 operations and businesses, and as a regulator the NRC  
12 stands by the safety and security of the protective  
13 order and the environment. Determining the three S's  
14 of safety, security, and safeguards is the duty the  
15 NRC and we as staff who have prepared this  
16 presentation provide information on the status of our  
17 mission.

18 And as we have previously presented to  
19 this Committee during the development of the AI  
20 Strategic Plan, I'll provide a quick debrief of the  
21 development since that time and specific implications  
22 or interests for the Committee's consideration. And  
23 to highlight previously about the interdisciplinary  
24 nature of our team these are some pictures from the  
25 workshop that includes Paul Krohn from Region I; Jesse

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 Seymour, who works in the Human Factors Branch; Joshua  
2 Kaiser, who's actually here in the room with us today  
3 with the Office of Nuclear Reactor Regulation and the  
4 chair support for our responsible use of artificial  
5 intelligence.

6 This slide, slide 27, provides a timeline  
7 of events since the last presentation that we provided  
8 to this Subcommittee as indicated by the star. Since  
9 the last ACRS presentation, around June of 2022, we  
10 began collecting over 100 comments on the draft AI  
11 Strategic Plan; ADAMS accession number is indicated on  
12 the slide, and utilized those comments to issue our  
13 final AI Strategic Plans for fiscal years of '23 to  
14 '27, which is in NUREG-2261.

15 In March we also launched our AI Steering  
16 Committee. This centralizes our efforts to -- on  
17 artificial intelligence to make sure we're better  
18 prepared as an agency.

19 In July we initiated an AI regulatory gap  
20 assessment. And in September, as Matt presented  
21 earlier, we hosted a workshop for regulatory  
22 considerations.

23 Later in September we launched our AI  
24 community of practice for discussions of lessons  
25 learned and potential uses cases of artificial

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 intelligence. In late October we issued our project  
2 plan. And the ADAMS accession number is listed  
3 accordingly. All these will be explored further into  
4 detail as we continue on with these slides.

5 For the strategic plan to enable industry  
6 and to lead our mission of safety and security the  
7 goal of the AI Team is to stay with the development of  
8 AI so that during the deployment of these tools the  
9 NRC will have the ability to review any safety or  
10 security implications. The mission of the AI Team is  
11 to be -- is to enable a responsible use of AI. And  
12 wishing to be cautiously proactive we released a  
13 strategic plan in May of 2023.

14 The strategic plan outlines five goals  
15 similar to the ones that were presented about a year  
16 ago. They include regulatory readiness, establishing  
17 an organizational framework, strengthening  
18 partnerships, cultivating a proficient workforce, and  
19 goal 5, which is to build an AI foundation within the  
20 NRC. The status of these goals will be presented in  
21 the subsequent slides.

22 Along with the strategic plan we've also  
23 issued the project plan in October of 2023. This  
24 project plan goes into depth of the strategic plan's  
25 goals as well as sets the scope of these goals. It

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 provides key timelines as well as tasks to ensure  
2 adequacy that we're meeting our metrics, and it  
3 promotes communications for external as well as  
4 internal stakeholders. Its purpose is to provide the  
5 public with transparency and accountability while the  
6 current staff plans are an applicant's clarity into  
7 the NRC's roles and responsibilities.

8 The timelines that we've had generally  
9 match with the expected deployment of AI that we were  
10 able to obtain with stakeholder feedback, however  
11 currently with the timelines we hope to continually  
12 update the project plan because of the changing  
13 current -- the current change of political climate.

14 Goal No. 1 is on regulatory readiness, or  
15 what we like to call keeping the end in mind. With  
16 every journey knowing what you're working towards  
17 helps provide that mission to perspective.

18 On ongoing work I want to highlight three  
19 items: pre-application communication, our gap  
20 analysis, and our continued with the IEC.

21 On the regulatory gap analysis; we can  
22 start at the top, we're currently working on reviewing  
23 regulations and guidance as it applies to current gaps  
24 in policies before AI. While conducting this gap  
25 assessment we're also incorporating and reviewing

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 applicable standards for artificial intelligence to  
2 recommend updates or recommend new standards to be  
3 developed.

4           Regarding pre-application communication,  
5 to help better budget and resource plan for AI  
6 applications the staff plans to develop a strategy to  
7 collect information for AI scheduling by the industry.  
8 These surveys could include RISS and FRN, but as Matt  
9 previously highlighted, what we've currently been  
10 doing has been extremely beneficial, which is  
11 conducting public workshops and information gathering  
12 at conferences where there's industry and labs that  
13 are participating. These have been extremely fruitful  
14 discussions.

15           The third item is the IEC, or the  
16 International Electrotechnical Commission, the NRC's  
17 participating Subcommittee 45 Alpha and Working Group  
18 12. 45 Alpha is specifically on the instrumentation  
19 control and electrical power systems in a nuclear  
20 facility, and Working Group 12 is more specific to AI  
21 applications in these nuclear facilities.

22           I want to preface this by saying this  
23 working group is very new as a second meeting only  
24 occurred early in November, and we have four staff  
25 members involved so far; three from the Office of

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 Nuclear Security and Incident Response and one from  
2 the Office of Research, which is Trey Hathaway.

3 The IEC plans to develop and maintain  
4 standards for AI applications for nuclear facilities  
5 by providing guidelines to stakeholders who are  
6 developing, deploying, as well as overseeing AI  
7 applications. In addition to this they hope to cover  
8 fundamental characteristics of AI in these nuclear  
9 facilities and make it applicable to the entire  
10 nuclear life cycle. The IEC --

11 MEMBER KIRCHNER: Could I interrupt you  
12 here?

13 MR. VALIAVEEDU: Yes.

14 MEMBER KIRCHNER: We through this  
15 Committee with Charlie's encouragement encouraged NRR  
16 to lay out a road map of the digital I&C. Are the  
17 staff who were responsible for that -- it's a very  
18 nice road map, by the way, of a very complicated  
19 wiring diagram for all of your regulations and guides  
20 and instructions and such for digital I&C. So it  
21 seems to me you have a framework in place if the  
22 application is actually going to be somewhere in the  
23 control systems or an operation in the plant. Are the  
24 NRR staff involved in this?

25 MR. VALIAVEEDU: We've incorporated NRR

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 with all discussions we have. We currently have an AI  
2 Working Group that meets monthly. That includes staff  
3 members of NRR as well as -- multiple staff members at  
4 NRR. But if you're talking about specifically the  
5 engagement with this working group, we only have three  
6 from the Office of Security and Incident Response  
7 and --

8 (Simultaneous speaking.)

9 MEMBER KIRCHNER: I don't want to get  
10 involved in NRC management decision and such, but --  
11 don't take this critically, but the security people  
12 look at things after the fact. They're looking at  
13 things that control access, they're looking at  
14 cybersecurity. I'm thinking of it a different way  
15 altogether. You're going to imbed some application  
16 somewhere in the plant. And it seems to me that's  
17 different than checks and balances as to whether you  
18 had an intrusion, et cetera.

19 So my concern or my suggestion here is yes  
20 to them, but involve the people who are intimately  
21 involved in how the plant operates from the control  
22 standpoint and the regulations and framework that is  
23 used for that. And then the applications in my mind  
24 of AI, at least the early application, somehow will  
25 come into that -- have to come into that regulatory

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 framework.

2 So I just point out that there was a --  
3 this very nice road map that's been put together that  
4 addresses digital I&C. To the extent that AI is going  
5 to come into the digital I&C regulations for the  
6 agency those people need to be involved. So end of  
7 speech.

8 MR. SEYMOUR: So, this is Jesse Seymour  
9 from the Operator Licensing Human Factors Branch, and  
10 I just wanted to speak to that point.

11 So, both myself and David Desaulniers, and  
12 also in prior efforts Dr. Brian Green as well, have  
13 been involved in the AI efforts. And if you were  
14 going to create a Venn diagram of who's working on the  
15 digital I&C upgrades that are currently in progress at  
16 some of the plants like Limerick, Turkey Point,  
17 myself, Dave, Brian are all involved with that as  
18 well, too. So I think that there's a good kind of  
19 synergy between the folks that are considering the AI  
20 issues as well as the advanced digital I&C control  
21 systems that are involved. And there definitely is a  
22 sensitivity to where we're at on kind of the  
23 progression towards when we may or may not eventually  
24 see AI in any type of a controlling context.

25 Right now in terms of the implementation

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 of AI what we're really seeing in applications -- and  
2 by that I just mean usage out in the industry, not  
3 necessarily even in a regulated context. And even in  
4 pre-application discussions it's more so at the level  
5 of AI insight, again if we were going to think about  
6 that kind of zero to four hierarchy of AI autonomy.  
7 And what we're not seeing is anything right now or in  
8 the near term that would take the human out of being  
9 the decision maker, whether that's in any type of  
10 operational context or even in the sense of  
11 calculations.

12 I think a good working example of where  
13 we're currently seeing things as currently state-of-  
14 the-art is using machine learning to -- a good example  
15 would be provide training insights, training  
16 interventions as folks are going through training  
17 programs and things of that nature where again it's  
18 informing human decisions. But we're not yet seeing  
19 it in a controlling context.

20 MEMBER KIRCHNER: Well, let me give you  
21 some examples of using big data in actual --  
22 potentially. You look at a core map or you look at --  
23 right now already there are software implemented that  
24 looks at a large array of data from the core:  
25 thermocouples, nuclear and such, and those

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 calculations like what's your margins, CHF, and so on.  
2 So I can see already that using advanced processing  
3 techniques you could improve the performance of those  
4 things that are already done analytically now and then  
5 somehow feed into the actual plant operation.

6 So I just threw that out because I can see  
7 that happening with the existing plants and that  
8 somehow that, quote/unquote -- calling it AI if you  
9 want to, or just advanced data processing -- I can see  
10 that being a kind of application in the plants. And  
11 somehow that has to factor into the digital I&C road  
12 map and regulation framework that you have. So it's  
13 just a suggestion and I'll stop there.

14 And, yes, Jesse, we appreciate what you're  
15 doing in the human factors arena as well.

16 MR. SEYMOUR: Thank you.

17 MR. VALIAVEEDU: Thank you.

18 MEMBER BROWN: Let me --

19 PARTICIPANT: We've got somebody on --

20 (Simultaneous speaking.)

21 MR. CARTE: Yes, Norbert Carte, Digital  
22 I&C, NRR. So I have not been officially asked to  
23 participate, but I am sticking my nose into this  
24 stuff. And I am following what's going on. It's very  
25 interesting. But really from a safety system point of

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 view I'm 100 percent aligned with Charlie. Keep it  
2 simple, separate.

3 And the question is when if and ever we  
4 change our paradigm -- so right now we base the  
5 approval of equipment based on some conservative  
6 limiting scenarios that occur at the worst possible  
7 time. And in that sense there are some very simple  
8 trip functions that protect you: high temperature,  
9 high pressure, high flux. And there's no reason why  
10 -- practical reason why you need to introduce AI into  
11 any of those.

12 But if you were to change that paradigm  
13 and no longer have conservative limited bounding  
14 scenarios to size and establish the performance  
15 criteria for your equipment, then you would need to do  
16 some serious thinking. But right now as long as we  
17 have conservative scenarios and simple separate  
18 independent protective functions, AI is not going to  
19 get into the protection systems themselves.

20 Now they will maybe reduce margins,  
21 they'll reduce need for unnecessary maintenance,  
22 they'll reduce unnecessarily challenges of the  
23 protection system, but until we change that paradigm  
24 there's no reason to have anything in AI -- sorry,  
25 it's just my personal opinion, but just trying to be

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 clear. In the protection system a high-temperature,  
2 high-flux trip does not need AI. And so, but there's  
3 a lot of ways you could make a plant better with AI.  
4 It's just it's going to be a long time before it  
5 drifts into the actual safety systems if they remain  
6 these simple safety systems that we have today. Thank  
7 you.

8 MEMBER REMPE: So just a matter of  
9 process, if there's a member who raises their hand, of  
10 course we should bring them into the conversation. If  
11 there's a member of the staff who wishes to make a  
12 comment or a contractor, the staff needs to call on  
13 them. Okay? Just so we keep the rules going. Thank  
14 you.

15 I think Charlie had his hand up and wanted  
16 to make a comment.

17 MEMBER BROWN: No. Two things: One, I  
18 agree with you relative to who gets to what, but there  
19 are some staff people that can support and they're  
20 operating -- NRR and other digital I&C people who need  
21 to be involved and understand what's going on in this  
22 world. And if there's -- Norbert had some very good  
23 comments, not negative, just how you integrate, and we  
24 need to be -- we should be conscious of those as we go  
25 through the meeting. So I appreciate you saying

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 somebody had their hand up. I think that was a  
2 positive one.

3 The second one was an expansion, just to  
4 follow up on Walt's comment about the road map. We  
5 had a meeting -- this is -- I understand your  
6 comments, but the road map was really trying to  
7 provide something that shows where all the standards  
8 and specs at the various levels of the I&C development  
9 systems and how do they apply. And there was a  
10 meeting on that where there was a set of presentations  
11 and slides that -- I think it was NRR staff that  
12 provided that. I'm not sure. My memory is not real  
13 good on that right now. And that's what they  
14 referring to.

15 We had that meeting back in April of this  
16 year, April 3rd. It was a Full Committee meeting.  
17 And that one presentation has a beautiful layout of  
18 what we meant by a road map. It was not trying to  
19 drive you any place. But it's not what you'd call a  
20 Venn diagram. It has nothing to do with a -- if Venn  
21 diagrams or -- if you want confusion, generate a Venn  
22 diagram. That's my personal opinion.

23 Anyway, I just wanted to make it -- the  
24 road map thought process that Walt brought up so that  
25 they would know what we were talking about. And I'll

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 let you -- I'm still restraining myself, but I will --  
2 I absolutely totally agree with Norbert, if you hadn't  
3 figured that out by now.

4 MR. VALIAVEEDU: Thank you for everyone's  
5 comments.

6 Moving on. For slide 31 we've provided a  
7 timeline for our progress. The purple dashed line on  
8 this slide indicates where we are today. The check  
9 marks have indicated completed items. Task 1.1 is our  
10 researching on our current regulatory framework, and  
11 C what is applicable. We've been able to establish a  
12 contract and we're currently drafting an analysis  
13 report based off of that contract with a hopeful  
14 completion date of the spring of 2024 for it to be  
15 published.

16 We were able to incorporate Task 1.2 for  
17 AI standards assessments within Task 1.1, so they're  
18 being conducted concurrently. And we're continually  
19 -- we're maintaining our ongoing participation in a  
20 variety of standards forums as well. And we hope to  
21 incorporate AI standards into our regulatory guidance.

22 CHAIR BIER: So just a factual  
23 clarification looking at that slide. So you're  
24 anticipating like end of fiscal year '27 there  
25 actually will be regulatory guidance for AI. Are you

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 envisioning that is mainly going to apply to new  
2 plants or use of AI in new capabilities at existing  
3 plants, or both?

4 MR. VALIAVEEDU: The guidance is meant for  
5 any stakeholders or applicants, who plan to utilize  
6 artificial intelligence as the NRC's missions for  
7 safety and security.

8 Task 1.3 is on a safety and security  
9 framework. This will be dependent on the results of  
10 our previous two tasks on standards, as well as our  
11 current regulatory gap analysis.

12 And, we would utilize those results and  
13 update or develop our current regulatory guidance as  
14 needed.

15 1.4 is on pre-application communication.  
16 We've begun discussing internally about additional  
17 strategies to obtain industry and stakeholder  
18 feedbacks, and plan to collect this information in  
19 fiscal years 24 and adjust our planning information  
20 accordingly.

21 1.5 is on AI enabled autonomous  
22 operations. During our engagement with a variety of  
23 stakeholders, we are not aware of any near-term  
24 deployment of AI enabled autonomous operations, but  
25 there has been interest in it.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           We plan to begin researching this in  
2           fiscal year 25. The objective here is to develop a  
3           technical basis, and requisite regulatory framework  
4           for AI enabled autonomy in nuclear operations.

5           Considering out the various, in how the  
6           various ways of how AI could impact autonomous  
7           operations.

8           Goal 2 is on organizational framework.  
9           Due to unique impact AI technology has on nuclear  
10          applications, the staff is working on centralizing and  
11          developing an internal organization for AI knowledge  
12          and expertise.

13          There are three key areas here. One is  
14          the AI steering committee. The second is the AI  
15          community of practice, and the third is on a  
16          centralized AI database.

17          On the first item, the A steering  
18          committee, the A steering committee has an involvement  
19          with a variety of program offices, and regional  
20          representatives. It meets monthly on topics related  
21          to AI within the NRC's purview.

22          This is being led by Deputy Division  
23          Director Victor Hall, as well as Paul Crohn, who is  
24          the Division Director in Region 1.

25          The AI team remains also in the Office of

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 Research, in the Accident Analysis Branch, that's  
2 being overseen by Luis Betancourt, who is here today.

3 On the second column on the AI community  
4 of practice, this was formed in September and it's had  
5 three meetings so far.

6 It's a formalized community of practice  
7 where the NRC has, within the NRC where people from  
8 all over the agency have been able to share their  
9 practices, as well as lessons learned.

10 It provides awareness for potential use  
11 cases and activities, throughout the nuclear sector.

12 The third column is on a centralized AI  
13 projects database. This has also been developed and  
14 deployed by various other agencies, including the  
15 Department of Agriculture, and Department of Treasury.

16 The goal here is to maintain transparency  
17 to the public on AI ML technologies. And, we  
18 currently have a dedicated public site for tracking  
19 these activities.

20 We are currently researching into a  
21 variety of use cases, such as data mining, as well as  
22 mathematical modeling that will be presented later on  
23 today.

24 And we hope to continue to update this  
25 database reoccurringly for accuracy, as well as

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 complete, for completion.

2           Regarding the timeline of this, we, the  
3 periodic squares indicate when we need to update, or  
4 plan to update a paper.

5           For Task 2.1, we were able to establish  
6 the AI steering committee. Task 2.2 for the community  
7 of practice, we were able to formulate and establish  
8 the community of practice earlier in September.

9           In 2.3 for establishing a projects  
10 database, we were able to develop an initial projects  
11 database. However, we are continually updating the  
12 projects database, and we hope to update it as needed  
13 for accuracy and completion.

14           Again, the checkmarks indicate completed  
15 items. The purple dashed line indicates where we are  
16 today.

17           DR. SCHULTZ: Anthony, this is Steve  
18 Schultz.

19           Just given the speed of activity that you  
20 said the NRC is doing, as well as the AI developments  
21 going forward throughout the government, I was  
22 surprised that the updates that you have for these  
23 elements are two years apart.

24           In other words, I'm expecting that things  
25 are going to have to be updated more frequently. And,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 it sounds like you're having monthly meetings in some  
2 areas, annual meetings in others.

3 And was wondering what kind of feedback  
4 you have in areas such as the regulatory gap analysis,  
5 for example.

6 MR. VALIAVEEDU: So specifically for the  
7 regulatory gap analysis, we are concurrently, it's  
8 ongoing work right now. We hope to have a issuance of  
9 a report in spring of 2024.

10 Regarding the timeline you're talking  
11 about in this slide specifically, for the charters,  
12 those generally will be updated as needed, as  
13 indicated with the parenthesis, as well as the, the  
14 database that's going to be updated reoccurringly  
15 because of the ongoing work.

16 And the change of political climate again,  
17 will have us revisit the project plan, and we hope to  
18 update the project plan's timelines in fall of 2024.

19 Because of the recent executive order and  
20 chair's memo, which I think you're hinting at here,  
21 happened earlier.

22 DR. SCHULTZ: It's not a criticism, but  
23 that's the second time you used the political climate  
24 associated with the overall program, and that hope  
25 we're focusing on the technical climate as well, and

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 the developments there.

2 MR. VALIAVEEDU: Thank you.

3 On the technical climate, we are engaging  
4 with stakeholders and making sure that their  
5 deployment and timelines here match with what is  
6 expected for the agency to maintain safety and  
7 security mission.

8 DR. SCHULTZ: I'm glad to see that, too.  
9 And, the interactions that you have been having on a  
10 very frequent basis that are urgent.

11 MEMBER BROWN: I want to, this is Charlie  
12 Brown again.

13 Yes, I was disturbed a little bit with the  
14 political. Assuming the political climate should have  
15 absolutely nothing to do with anything you all do.  
16 Nothing.

17 It should be a zero impact. It should be  
18 developed strictly, I mean it, you're really going to  
19 raise a lot of hackles with people that sit on the  
20 Advisory side of what we're looking at, if the thought  
21 process is we're going to hurry up and do something  
22 because somebody politically wants to, you know,  
23 demonstrate that this is being done.

24 That's absolutely insane. And couple that  
25 with part of your, go back, you don't have to go back

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 it's the previous one you were talking about,  
2 evaluations looking at autonomous operation, that was  
3 the words you tossed in kind of as you, when you went.

4 And the ability of having autonomous  
5 operations. Presenting information or AI that  
6 evaluates data that's coming out of the plant, and  
7 then assembles it within some algorithm, or some  
8 presentation that it informs the operator that hey,  
9 these things are going in this direction.

10 And then the operator makes the decision  
11 about what to do with it, or seeks decision from, you  
12 know, consultation with what to do.

13 That's nothing wrong with providing better  
14 data, because there are tons of data we're getting out  
15 of digital systems now.

16 And it makes it very, very difficult you  
17 know, to assess you know, the directionality of them,  
18 and which one's pressure temperature can be going in  
19 opposite directions and say, oh, when it gets to a  
20 certain point, is that critical or not.

21 So that's, that's the type of things that  
22 boy, throwing in the political thought process or  
23 directive from whoever they is, is just useless.

24 And that's, that should just never get  
25 encumbered in the develop of your all's processes.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1           Anyway, I'm sorry to get wound up every  
2 now and then. Did you all hear me satisfactorily?  
3 I'm told I don't speak into the mic well enough as it  
4 is.

5           MR. HALL: So if I can just jump in to  
6 maybe, oh yes, thanks, thanks. So again, this is Nick  
7 Hall, the Office of Research.

8           I guess political, it's a bad word, right?  
9 So, we can talk about the religious influences on AI,  
10 and maybe the other taboo topics that you don't  
11 discuss at dinner.

12           I just wanted to offer a clarification is  
13 it's the awareness of when we say a political  
14 environment, it's the awareness of what the government  
15 is doing, right.

16           There's, and there's two executive orders  
17 that are the big ones, the biggies, that talk about  
18 AI. One was from the Trump administration, the other  
19 one is the Biden administration.

20           That's just a matter of statement of fact.  
21 I certainly would back my staff in saying there is  
22 never, ever any political consideration with the big  
23 P. It's just the awareness what the government is  
24 asking us to do, to make sure that we're prepared.

25           So I just wanted to clarify that because

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 I know as soon say you mentioned the P word, it, I get  
2 the hairs on my neck stand up when anyone says  
3 politics. That's how this country is, right, with  
4 politics these days.

5 So, just a point of clarification, it  
6 really is an awareness of how we can best be prepared  
7 for the technical. The folks in front of you are  
8 engineers, scientists, and darn good ones at that.  
9 I'll leave it at that.

10 MEMBER BROWN: But you don't want it to be  
11 hurried. It should not be hurried. It needs to be  
12 technically solid and validated before anybody goes  
13 forward with anything like that.

14 And that's, that's not that the, we don't  
15 obviously the government wants to make sure the  
16 agencies are paying attention to things that may  
17 enhance their operations. That works just fine.

18 But when you hurry place where we've got  
19 spacing considerations because P, the government  
20 really needs to emphasize and regulate that, that's  
21 not, that's not the right emphasis.

22 So, I wasn't trying to be critical of  
23 anybody up here that said anything. That's not the  
24 point. The point was safety first, and introductions,  
25 you know, comes along along with it. And look at

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 where it adds value.

2 Doing something because its fancy, new and  
3 everybody thinks it's the greatest thing since sliced  
4 bread, is one thing.

5 But if it doesn't add value to the  
6 performance of the plant, it ought not be done at all.  
7 That's the only point I'm trying to get across in all  
8 your all's deliberation.

9 After laying these out, these pathways  
10 out, you've got to go through that. You got to go  
11 through the drill of how are you going to do this, and  
12 how are we going to make it make sense.

13 But you've got to do it in the manner of  
14 where does it add value. If it doesn't add value,  
15 don't waste people's time; don't complicate the  
16 systems with it.

17 So, that's just the message I was trying  
18 to emphasize. I wasn't the only one.

19 MR. VALIAVEEDU: Thank you again for  
20 everyone's comments.

21 Again, the NRC's mission is safety and  
22 security for the people, and the environment.

23 Slide 34, we were moving on to Goal Number  
24 3, which is strengthening and expanding partnerships.

25 The previous presentation by Matt Dennis

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 probably best explains how far we've been able to come  
2 regarding developing closer ties, and promoting  
3 partnerships through our proper workshop interactions.

4           However, with this slide I wanted to  
5 highlight two additional areas. One is domestic  
6 interactions, and the other one is international  
7 interactions.

8           We've had talks with a variety of agencies  
9 including the Department of Energy and the National  
10 Nuclear Security Administration, on artificial  
11 intelligence.

12           Specifically, with our work with the  
13 Department of Energy, we've been able to observe areas  
14 developing AI ML technologies through the Light Water  
15 Reactor Sustainability Program, or LWRS Program. And  
16 review those results obtained, as well as lessons  
17 learned.

18           The DOE MOU has been extremely helpful in  
19 understanding the direction industry is undertaking on  
20 AI.

21           We are also in membership in the NIST RMF,  
22 or Risk Management Framework working group. And this  
23 has largely been an observational capacity, and we  
24 have been providing our Nuclear Regulatory expertise  
25 into these discussions.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 Other talks we've had also include  
2 participation in multiple conferences, including the  
3 Digital Engineering and Nuclear Technologies  
4 conference, or DENT conference.

5 As well last week, Matt and I were able to  
6 attend the Ohio State University Big Data Workshop.

7 These interactions have been geared  
8 towards promoting clarifying stance in regulations for  
9 stakeholders, and promoting communications.

10 As regulators, we can provide clarity and  
11 share concerns through effective pre-application  
12 engagement.

13 On the international side of the  
14 wheelhouse, we are currently working with the United  
15 Kingdom, as well as Canada, on a tri-lateral  
16 engagement that we call CANUKUS.

17 I'll be discussing this in the following  
18 slide more.

19 In addition to this, we've been engaging  
20 with the IAEA on technical meanings that provided  
21 insights into other nations' priorities in  
22 developments.

23 There has been generally speaking, similar  
24 concerns between nations on artificial intelligence.  
25 In addition to this, during these technical meetings

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 there has been generally a main focus on utilizing AI  
2 for operation maintenance, as well as design  
3 utilization, as well as using AI as mostly as an  
4 informative tool using natural language processing, as  
5 well as mathematical modeling such as linear  
6 regressions.

7 The third item here is on bilateral  
8 engagements. The goal here is to foster and maintain  
9 collaboration with international counterparts, and  
10 multilateral organizations to positively influence and  
11 maintain awareness on the responsible and safe use of  
12 AI.

13 And, this is in support and alignment with  
14 the NRC's 2014 international strategy to positively  
15 influence safety and security, as well as maintain  
16 awareness for the agency's domestic objectives.

17 DR. SCHULTZ: Anthony, Steve Schultz.

18 You mentioned, or it was mentioned that  
19 the, there were presentations at the RIC associated  
20 with AI and so forth.

21 And in terms of the international program  
22 and plans, what's planned for the RIC, or around that  
23 conference this year?

24 MR. VALIAVEEDU: I'm glad you brought that  
25 up. We are hosting an IAEA technical meeting at

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 headquarters in March of 2024.

2 DR. SCHULTZ: Thank you.

3 MR. DENNIS: The button doesn't like my  
4 finger, so I have to keep pressing it.

5 And I will mention, we are, we did get a  
6 confirmation this week that we are doing another AI  
7 technical session at the RIC this year.

8 So there will be an AI technical session,  
9 you get to look to that. Please come and attend.

10 And what Anthony just mentioned is the  
11 following week, we will be hosting a IAEA technical  
12 meeting.

13 The week following, we think we're going  
14 to get a lot of good participation -- the IAEA thinks  
15 we'll get a lot of good cross-collaboration with  
16 people that are going to be attending the RIC, as well  
17 as then the technical meeting the following week.

18 Right now we haven't pinpointed what our  
19 presenters are going to be for the RIC technical  
20 session.

21 We're in the process of doing that right  
22 now. But it will be similar to last year's where we  
23 have a flavor of industry, other federal agencies, us,  
24 and academia.

25 DR. SCHULTZ: Sounds like good

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 coordination. Thanks.

2 MR. DENNIS: Yes.

3 MR. VALIAVEEDU: Regarding the CANUKUS  
4 engagement, CANUKUS is an interesting development in  
5 the NRC as it marks a time where three nuclear  
6 regulators were coming together to share a common  
7 goal.

8 The outcome with this interaction is a  
9 high leveled AI principles paper, that we hope to sign  
10 by spring of 2024.

11 The goal is to provide a uniform front,  
12 and what are key considerations when developing AI  
13 systems for safety and security.

14 And, we are currently putting together the  
15 first draft of this paper, but I want to preface this  
16 by saying that this is not for legal use, nor used in  
17 place of a regulatory framework. Instead, it provides  
18 a summary considerations.

19 This includes discussions on how to  
20 utilize existing safety systems, and how to utilize  
21 those existing safety systems to demonstrate safety.

22 This will help assist developers in  
23 evaluating their own system.

24 In addition to this, we will have, there  
25 will be discussion on human factors, and how AI

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 impacts human factors, as well as an impact on AI  
2 architecture, as well as a summary consideration  
3 section for life cycle management, especially in the  
4 context of generative AI.

5 Again, we hope to complete the paper by  
6 the spring of 2024, as the working group was  
7 formulated in November of 2022.

8 MEMBER PETTI: I just had a question about  
9 which country really is feeling the greatest user push  
10 for AI?

11 Is it still U.S. compared to Canada and  
12 the U.K., or are they having to be ahead of us because  
13 of their licensees?

14 MR. VALIAVEEDU: I will have to defer that  
15 question to Matt, because he's been mostly working on  
16 CANUKUS.

17 MR. DENNIS: So, Trey and I have been --  
18 Trey and I have participated in two IAEA working  
19 groups and we get to -- we've had the privilege of  
20 being able to see globally where it seems to be  
21 leaders and applicationaries (phonetic). The U.K. and  
22 the U.S. still seem to be leading the charge on AI  
23 applications in the nuclear sector.

24 China is somewhere in the mix there in the  
25 middle of application areas, as well. Russia has had

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 a lot of presentation at IAEA.

2 So, it seems to be U.S., U.K., China and  
3 Russia, as far as application areas. I will say I  
4 think, and this may be biased, but I think the U.S.  
5 and the U.K. are leading the charge on the regulatory  
6 aspects of AI in nuclear.

7 And this all sort of blends in and makes  
8 sense that the U.K., U.S. and China are sort of  
9 leaders in AI.

10 If you go look, I forget the Alan Turing  
11 Institute in the U.K. put together a website and the  
12 U.K., or the OACD has a website, basically has a  
13 tracker of AI leadership in the application areas.

14 And the U.K., U.S. and China are at the  
15 top of that board as far as, and that's just, that's  
16 not nuclear, that includes everything, right, but.

17 MEMBER HALNON: Great, no, this is Greg.  
18 I respect, you know, all the learning that's going on  
19 and I'm probably known throughout the industry as  
20 being relatively impatient, so forgive me.

21 Working group formed in 2022 and the  
22 output two years later is going to be a paper. How is  
23 that moving forward fast enough in parallel with  
24 what's all these applications and other things, and  
25 other countries?

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           It just, I mean, we talk in kind out both  
2 sides of our mouth, it feels like. It's going really,  
3 really, really fast, and it's taking two years to get  
4 a paper out.

5           And, that's all it's, that's the goal of  
6 the outcomes is some of the, lot of meetings, lot of  
7 technical presentations. Lot of learning going back  
8 and forth.

9           And I recognize that's going to happen,  
10 but it just feels like a tar pit.

11           Maybe you can comment on the speed and  
12 the, the amount of resources we have applied,  
13 dedicated to it.

14           If not dedicated, how are we going to move  
15 this thing forward fast enough so that when somebody  
16 does come up with an application, we're ready to let  
17 it go.

18           I mean, you know, we're still, I guess we  
19 got through the 50.59 stuff and digital INC, but that,  
20 that kind of blocked a lot of digital INC upgrades for  
21 a while.

22           And, I remember back in the 2012  
23 timeframe, SMRs. Everybody was ready to go but we  
24 just got the rules, or you know, dropping down the EP  
25 Zone through the rulemaking process almost a decade

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 later.

2 How are we not going to government this?

3 MR. VALIAVEEDU: Matt's going to answer  
4 the U.K. stuff for the international. But for the  
5 timeline, we've formulated the original timeline based  
6 off of our interactions with stakeholders.

7 What they thought they're going to be able  
8 to deploy, at what frame of time.

9 And based off of that end time, we  
10 formulated our strategy to ensure that we are ready to  
11 evaluate that safety, evaluate the responsibilities of  
12 AI.

13 The resources we're putting on to this, we  
14 have a whole of agency approach as I mentioned  
15 earlier. We have multiple program offices involved  
16 with this.

17 Matt and Trey have been working on this  
18 since the beginning, for the strategic plan. And, we  
19 plan to update those timelines as changes happen.

20 So, we've been interacting with our public  
21 workshops, seeing where industry is thinking about  
22 doing AI. Is there a, has there been any radical  
23 changes that happen.

24 Luckily at this previous workshop we were  
25 told that our timelines expected, or matches the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 expectation that we see industry deploying AI usage.

2 Again, fall of 2024 is when we plan to  
3 publish a new version of the project plan, to readjust  
4 those timelines to ensure that we're maintaining that  
5 adequate resources for the evaluation.

6 MEMBER HALNON: So do you see, Anthony, a  
7 period of time when there's going to be a dedicated  
8 not an office, but a directorate or something for AI  
9 so that we focus this all agency approach with some  
10 dedicated resources to establish these guidance  
11 documents, get them through, get them through ACRS,  
12 get them signed, and on the street?

13 Or is it going to continue do you think,  
14 for a number of years at the matrix type?

15 MR. VALIAVEEDU: So, that's out of my pay  
16 grade.

17 MEMBER HALNON: I know.

18 MR. VALIAVEEDU: So, I'm going to defer  
19 that to Luis Betancourt.

20 MR. BETANCOURT: To answer that question,  
21 I think what is important is that we have an AI  
22 steering committee from all of the program offices  
23 that's basically directing this work.

24 The idea that we want to do, is to see the  
25 outcome of the regulatory gap assessment to really

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 identify what are, to your point, what guidance has to  
2 be updated.

3 Do we need to develop something new? I  
4 expect that to be some time in the spring, followed  
5 with a workshop where we going to be talking with  
6 industry to get their feedback, this is what we found.  
7 Do you believe this is something, what are the areas  
8 that we need to prioritize.

9 And then that's going to be going back to  
10 the steering committee, and then we will go back to  
11 you guys.

12 Kind of what Charlie was talking about the  
13 roadmap, that we need to lay that down. Okay, now  
14 that we know where the gaps are, let's sit down and  
15 put this in front of us and ask everybody, so.

16 MEMBER HALNON: So, Luis, you're kind of  
17 say that this steering committee is going to be key on  
18 establishing the agency approach down the road.

19 MR. BETANCOURT: Correct.

20 MEMBER HALNON: Because, at some point, it  
21 seems like -- and you know how long it takes to get  
22 guidance documents written and through.

23 MR. BETANCOURT: Yes.

24 MEMBER HALNON: There's got to be some  
25 real focused effort.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 MR. BETANCOURT: Yes, and I think that's  
2 why we're starting that focus at the front end. And,  
3 there will be some coordinating with guys early so we  
4 don't wait.

5 Because as everybody's pointing out, this  
6 is a fast-paced environment. But to Charlie's point,  
7 we also want to make it technically right that we're  
8 not basically putting like, efficiency in front of  
9 safety, so.

10 MEMBER HALNON: Okay, thanks. And sorry  
11 for derailing it. I just was reacting to the two  
12 years to get a paper out. And, I understand that just  
13 a lot of learning that has to go on.

14 So don't take it as a criticism, it's just  
15 that you know, the regulatory timeline seems long  
16 sometimes.

17 MR. DENNIS: I will respond, or I'll  
18 mention the CANUKUS tri-lateral engagement.  
19 Technically speaking, it's only, it hasn't even been  
20 a year that we've actually been working on it full  
21 steam ahead.

22 So, I recognize that the spring 2024 is  
23 when you add the numbers together, two years. I'll  
24 ask for some grace in this because we have, this is a  
25 principles paper with three countries.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           And as a precedent for this, the same  
2 three countries from the health and safety FDA  
3 perspective put out a principles paper. And, it took  
4 them quite a long time to come to a unified agreement  
5 for a two-page paper.

6           And, we're trying to get more than that,  
7 and get a little bit more in-depth. Because we took  
8 that as our benchmark and said, FDA and health, or  
9 Health Canada and FDA and the U.K.'s health office  
10 came together and put out a good machine learning,  
11 good machine learning practices and principles paper  
12 a couple years ago. It was two pages.

13           We thought, this is not our benchmark for  
14 what we want to put together. And we recognize  
15 through our, some of our collaborations that this is  
16 an area where applications may come in this to all  
17 three of us. And having a unified perspective would  
18 be a good thing.

19           So, just getting three international  
20 entities go agree to the words on a piece of paper is,  
21 is a challenge. So I'll say that.

22           MEMBER HALNON: All right, well I'll give  
23 you grace.

24           MR. DENNIS: Thank you.

25           MEMBER HALNON: We all have waited at

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 11:00 o'clock for the weather forecast, and we're  
2 disappointed when they spend 15 minutes telling us  
3 what we already saw.

4 So, I just don't want to see the paper  
5 come out and tell us that we're two years behind  
6 everybody else.

7 So that, I understand what you're saying.

8 MR. DENNIS: Yes.

9 MEMBER BROWN: Just to echo that thought.  
10 Has there been any effort by you all, and we talked a  
11 little bit about this earlier, or at least me and  
12 somebody else did.

13 To separate out what I would call the  
14 areas of the rice bowl that you really need to get  
15 regulation defined, whether it's via this tri-lateral  
16 approach to doing business or not.

17 But there's a whole plethora of things  
18 outside of that, that industry should just don't wait  
19 for us. Just go work on those, do what you want to  
20 do.

21 Why can't that be communicated in a manner  
22 in your meetings and say hey, we're drawing lines  
23 around plant controls, these controls, safety systems,  
24 et cetera.

25 But all the other type stuff from the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 maintenance, training, whatever it is, evaluating that  
2 data, you don't need regulations for that.

3 Why over regulate when you don't have to?  
4 Focus the regulations on the areas where it impacts  
5 the safety of the plant.

6 And I haven't heard that through, I'm just  
7 trying to echo a little bit of Greg's comment here is,  
8 slow down.

9 There was so much baggage associated with  
10 digital INC systems, and how you evaluate them. It  
11 took years.

12 That's why the roadmap became all of a  
13 sudden, how in the world do we tell what's applicable  
14 to what.

15 And actually brought all the pieces  
16 together where people could see what needs to be done,  
17 and the Reg Guides have now been refined pretty much.  
18 But it's taken a long time. I think we're talking  
19 decades long time.

20 So here, to me, you all have the  
21 opportunity to just put a rice, you know, a line  
22 around certain things and say, hey, look, stay out of  
23 these. You can do what you want every place else.

24 You're in your business. You're trying to  
25 maintain efficiency in your training systems,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 maintenance evaluations, the data you absorb.

2 What you do with it, how you evaluate it,  
3 go do it and don't look for regulation there.

4 I don't know why you can't take that step  
5 forward, get an agreement with the U.S. and how to  
6 get, so our industry can get on and utilize it where  
7 it is known to be non-safety, non-safety critical.

8 I'll just stop there for the next slide so  
9 I can do it again.

10 MR. DENNIS: I very much appreciate the  
11 topic that you brought up and this, this was mentioned  
12 in the workshop on the aligned, or crossing a line.

13 And in our working group, we've discussed  
14 this crossing a line you know, thought process as  
15 well.

16 The recognition we have right now is, we  
17 don't know where that line is at this point. And part  
18 of the gap analysis is to figure that out.

19 But at our front line individuals, the  
20 inspectors and the regents, have been very attuned to  
21 industry applications such as the corrective action  
22 process analyzer.

23 Those areas where we recognize that  
24 industry can deploy this, and use it, and, and  
25 business efficiencies and process areas, to your point

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 of not over regulating.

2 MEMBER BROWN: Or not regulating.

3 MR. DENNIS: And not regulating. And  
4 that's why there's a key distinction in every slide  
5 we've presented, or we brought up says, NRC regulated  
6 activities.

7 And, we are trying to be very mindful that  
8 we are looking at where AI touches something we have  
9 that is NRC regulated.

10 And other areas where it can be utilized  
11 currently today, keeping, we're keeping boots on the  
12 ground through the inspectors and the regents, to  
13 maintain awareness of where those areas are being  
14 used.

15 Because the industry is and has said, that  
16 they're using this technology in an early deployment  
17 phase to learn how to use it in areas where it can  
18 gain true value for them.

19 And there may be a future where it then  
20 does go into that NRC regulated activities space. And  
21 so we are preparing as Anthony mentioned, our current  
22 state spaces as we, as we understand it today, there  
23 have not been any applications that have come to us in  
24 an NRC regulated activity, for our review.

25 From what we've heard, we are aware that

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 there may be a couple areas, very targeted use cases,  
2 where that may happen in the next couple years.

3 Those are the ones where we're investing  
4 our effort right now, to try to focus on the gap  
5 analysis so that we can understand truly is what  
6 guidance is necessary, how would we evaluate it, and  
7 make a finding on that particular thing.

8 And to the digital INC point, I've  
9 mentioned this a couple times in other venues, that  
10 this as mentioned, was this is just the latest thing.

11 We've had advance reactors. We've had  
12 digital INC. We've learned from those things and  
13 we're trying, our hope is we're trying to get out in  
14 front of this a little bit so that we're prepared  
15 should that eventuality come as we understand it now  
16 and maybe three years.

17 MEMBER BROWN: Go ahead. If you needed a  
18 queue.

19 MR. VALIAVEEDU: Oh, okay, I'm going to  
20 move on to the next slide then. Thanks, Charlie.

21 Trying to keep up with our timeline here.  
22 For slide 36, we are maintaining our domestic  
23 partnerships, specifically with a NIST RMF, as well as  
24 the LWRS program.

25 We are currently drafting together a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 institution plan to engage with academic institutions,  
2 that we hope to complete by the end of quarter two of  
3 fiscal year 2024.

4 For 3.2 on international partnerships,  
5 we're drafting currently the CANUKUS paper that we  
6 hope to get out by spring of 2024.

7 And, maintain our current ongoing  
8 bilateral engagements with Canada, U.K., Germany, et  
9 cetera.

10 For the last three, 3.2 echo, foxtrot and  
11 golf, we are maintaining our participation in IAEA  
12 technical meetings regarding the utilization of AI and  
13 nuclear power plant safety.

14 And, the utilization of AI writ large  
15 within the nuclear fuel cycle.

16 Specifically, we're also participating in  
17 an IAEA project that will utilize artificial  
18 intelligence to evaluate severe accident data.

19 For 3.3, this was mostly mentioned by Matt  
20 Dennis but I just want to quickly go over this now.  
21 We were able to complete our most recent workshop on  
22 AI characteristics for regulatory considerations, as  
23 well as we hope to maintain those workshops as they  
24 prove to be fruitful in understanding where industry  
25 is undertaking innovations in AI.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           And, to help better align our resourcing  
2           for future fiscal years.

3           And we hope to also maintain our  
4           participation external workshops, conferences, and  
5           meetings, as these also have been able to provide us  
6           with more information about where use of AI is going  
7           towards, or heading towards.

8           Slide 37 is on cultivating an AI  
9           proficient workforce. Wide skill training is not new,  
10          and I was looking at some old photos of NRC history.

11          And, the photo in black and white is  
12          actually a seminar on Lotus 1-2-3, which is a  
13          spreadsheet software. I've never seen that software  
14          before until that picture came up.

15          And to the right of that is the AP1000  
16          simulator, the TTC, which again, showcases how wide  
17          skill training is not, not a new phenomenon.

18          However, our active planning and whole of  
19          agency approach when it comes to AI, compliments our  
20          readiness for taking this challenge on.

21          With the potential for AI to be widely  
22          used, the NRC has plans to develop the skills  
23          necessary for evaluating any AI, incoming AI  
24          applications.

25          The phrase that we utilize is, train,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 retain staff, and hire as we need to. We are ahead of  
2 schedule luckily, comparative to the AI project plan  
3 timeline.

4 And the staff is, because the staff is  
5 developing a training guide to help develop  
6 competencies for AI usage.

7 This guide splits the training between  
8 data scientists, data analysts, as well as program  
9 analysts, and provides basic training for AI ML  
10 systems.

11 In addition to this, we were given direct  
12 hiring authority for data scientists, as well.

13 For workforce planning, we, the staff has  
14 engaged with OCHCO, which is the Office of the Chief  
15 Human Capital Officer, on the development of a  
16 competency model for AI related job functions.

17 And, this effort aligns with the recent  
18 initiative from OPM, Office of Personnel Management,  
19 develop a competency model for the whole federal  
20 government.

21 The push by the chair's recent memo and  
22 the White House's also recent executive order, have  
23 only strengthened the team and the agency's resolve to  
24 strengthen our skillset to be better prepared for AI  
25 systems.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 Slide number 38 is another timeline. As  
2 indicated by the purple slide where we are currently,  
3 we are ahead of schedule for many of these activities.

4 4.1 is assessing the NRC's AI skills and  
5 identify any gaps. We're currently developing a  
6 competency model with OCHCO, to analyze what areas we  
7 would need AI related job functions with.

8 4.2 is identify, develop and implement AI  
9 training opportunities. We were able to put together  
10 a draft qual plan to help staff develop and train up  
11 on new data science skills, and AI skills.

12 4.3 is on recruitment, hiring and  
13 retaining AI talent. In collaboration with OCHCO, we  
14 plan on developing a working group to recruit AI  
15 skills, and retain that expertise.

16 Goal number 5 is pursuing use cases to  
17 build an AI foundation within the, across the NRC.  
18 The NRC's focus on internal usage of AI has been  
19 exploring research and development of AI ML.

20 That may benefit for future regulatory  
21 decisionmaking. One of these will be presented later  
22 on today.

23 We plan to develop a AI foundation through  
24 four areas. Pilot studies, safety insights, an AI  
25 ecosystem, and future focused research.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           The pilot study is an area where we hope  
2 to work with stakeholders, labs, and partners to  
3 investigate security and safety of AI technologies.

4           Our engagement with stakeholders have  
5 supported this concept. And previously, we've had  
6 industry say at events that they are in support of  
7 pilot studies, as well as regulatory sandboxing to  
8 help navigate the regulatory landscape.

9           As emerging technologies will always  
10 involve interaction between the regulator and  
11 developer, we wish to go about this early so that we  
12 are able to identify challenges within NRC review, as  
13 well as build technical expertise.

14           As we are not the barrier to innovation  
15 but instead we are the guardrails for safety and  
16 security.

17           The second item here is on safety  
18 insights. We wish to assess and survey what is out  
19 there to evaluate AI systems for safety. And  
20 incorporate those findings across the NRC.

21           The third box is on AI ecosystem. The  
22 room is only dark when you don't have a flashlight.  
23 So, there is value in the staff to get accustomed to  
24 AI ML tools, to deconstruct AI.

25           By acquiring common data science tools and

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 deploying IDEs that could help support an AI  
2 ecosystem, we could utilize training within goal  
3 number 4 to help get a staff, to help staff understand  
4 how AI could be utilized throughout the nuclear  
5 industry, as well as any safety or security issues  
6 that may come up.

7 MEMBER MARTIN: Question, and I'm sorry,  
8 I can't help myself. Trying to get software that is  
9 not like, already approved by IT is, is right,  
10 impossible, right?

11 So how do you expect to even get through  
12 NRC's own processes to get tools to train people? Do  
13 you have to make exceptions? Obviously you've worked  
14 with the Office of the Chief Information Officer.

15 I think you have your own obstacles just  
16 to even get to the point where you can train people.  
17 I mean, are exceptions made in the spirit of research  
18 and training, that lets you get tools in there that  
19 are typical, or being proposed?

20 Because invariably, you know, questions of  
21 security associated with those tools have to come up.

22 MR. VALIAVEEDU: To first answer your  
23 question, the staff currently has Anaconda and Python  
24 as two main packages that we are allowed, or that we,  
25 that have been vetted by OCIO.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 By in the first slide, or second slide, I  
2 ensured that the word OCIO is incorporated with a  
3 variety of program offices --

4 MEMBER MARTIN: I saw that.

5 MR. VALIAVEEDU: -- because we've been  
6 able to work with, directly with the chief information  
7 officer because AI is upcoming.

8 We've seen a lot of developments  
9 throughout not even the nuclear sector, and our whole  
10 of agency approach to this has been complementing  
11 that.

12 So, it seems from our view, slightly  
13 different than traditional software procurement. And  
14 we are having ongoing talks to acquire and allow more  
15 uses of like Python, and different library packages.

16 MEMBER MARTIN: But this would be  
17 segregated, too, for the purpose as opposed to letting  
18 everybody in you know, have access to it, and who  
19 knows what.

20 I mean, correct, or am I wrong?

21 MR. VALIAVEEDU: Oh, for that, I think  
22 Victor Hall, I think has direct engagement with the  
23 CIO. So, he may be able to better answer your  
24 question.

25 MR. HALL: Thanks, Anthony.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           So again, this is Victor Hall with the  
2 Office of Nuclear Regulatory Research.

3           We do have strong partnerships with the  
4 Office of Chief Information Officer. And I think back  
5 to right before pandemic, some of the changes that  
6 they made to be able to prepare us to be ready for a  
7 changing world to be able to work remotely.

8           And, they knocked it out of the park.  
9 It's thanks to them we were able to get our mission,  
10 even get our mission done.

11           And, I think they're taking that same  
12 mentality of being ready for what's coming, or  
13 changing where we have new tools that we need to be  
14 assured are safe, that are protected from the dangers  
15 that are out there, whether cyber or other.

16           And, I think they're taking that same  
17 approach to being ready to be able to give us the  
18 tools to be ready for what's coming.

19           So, they're working hand-in-hand with us.  
20 Obviously, has the table want all the bells and  
21 whistles.

22           We want the toys to play with, and they  
23 have to be able to say wait a second, let's make sure  
24 they're safe.

25           And they're doing that partnerships across

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 the government. So, you know, better to have approval  
2 to make sure that something is usable and safe, and  
3 that our systems is not going to get in and cause a  
4 greater problems is clearly high on their minds.

5 But again, I think they're taking a very  
6 positive and collaborative approach to making sure  
7 that we're ready for being able to use these safely.

8 And so again, I'm happy that we have them  
9 as partners. And I think when you look at the chair's  
10 tasking, which wants us to lean forward but  
11 responsibly, I think that gives us good momentum to  
12 be, to have those tools to be ready and recognizing  
13 that in government, there are going to be  
14 restrictions, period.

15 MEMBER MARTIN: I thought it was a full  
16 committee a time or two where we got the message that  
17 you know, things like personally things like ChatGPT  
18 would not be used.

19 Obviously concerns about the control of  
20 proprietary or classified information, and all that.  
21 You know, you just see every other day that we're  
22 getting a patch for our operating system for some  
23 security, you know, and I just think well this, like,  
24 you know, stuff like that.

25 So anyway, that's a little cynicism

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 associated with the question, but it's obviously a  
2 real bureaucratic challenge to beat some of these  
3 goals because if you don't really have the best tools,  
4 then are you really doing the best job training.

5 But you have other people that you know,  
6 decide your fate a little bit on the decisions that  
7 you're making, who have total control over it.

8 MR. VALIAVEEDU: Thank you again for your  
9 comment.

10 The last box here is on research. We hope  
11 to continue to invest in AI research through existing  
12 avenues, as well as universities.

13 This is through supporting our university  
14 research grants, and as well as with this, we hope to  
15 continue with our future focus research program.

16 This program has helped build NRC  
17 knowledge in emerging, and significant technologies.

18 More specifically, in the last two years  
19 we were able to fund six FFRs, one of which will be  
20 presented later on today by NRC staff.

21 Two more slides. Slide 40 regarding the  
22 timeline of these cases. For 5.1 is on the proof of  
23 concept in pilot studies.

24 As previously explained, the objective  
25 here is to engage with the industry to identify the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 potential pilot study and proof of concept test cases.

2 They'll help the NRC staff gain expertise  
3 for future regulatory reviews.

4 5.2 is to develop and maintain an AI  
5 ecosystem. As seen with the purple dashed line, we're  
6 currently working on developing and maintaining an  
7 IDE, as well as identifying and assessing, and  
8 acquiring AI tools.

9 5.3 is on surveying of AI tools and  
10 methods for safety evaluations. We hope to conduct a  
11 survey of what's currently out there to evaluate AI  
12 systems, by the end of quarter two of fiscal year  
13 2024.

14 And then implement these findings by  
15 fiscal year 2025.

16 MEMBER HALNON: Now Anthony, is there  
17 where 50.59 will come in, and how to figure out how to  
18 do an evaluation for modification to the plant?

19 MR. VALIAVEEDU: It's slightly different.  
20 The survey of tools and methods is more of what sort  
21 of systems and algorithms are out there so that  
22 individuals or stakeholders can utilize those two, to  
23 evaluate their own systems.

24 MEMBER HALNON: So the guidance I'm  
25 thinking about is if somebody wants to implement it,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 they're going to have to go through a lot of pre-  
2 application discussions with the NRC, in addition to  
3 having a process to be able to facilitate that.

4 Which would be a 50.90 submittal plus a  
5 50.59 and a no-hazards. All that stuff. Where is  
6 that all going to down the road once you get all this  
7 learning and start putting guidance, you know, pen to  
8 paper and making guidance documents?

9 Because it seems like it takes a long time  
10 to get, you know, alignment with the industry on some  
11 of these things and I just wanted to make sure that  
12 this 27 date is, is still feasible.

13 MR. VALIAVEEDU: For the guidance  
14 development, we hope to implement that as guidance  
15 development for stakeholders to utilize.

16 However, stakeholders are able to utilize  
17 what is out there right now for their 50.59  
18 application, if they wish.

19 MEMBER HALNON: Yes, but when you start  
20 getting into those fine questions of increases and  
21 decreases of consequence and all that kind of stuff,  
22 it gets difficult.

23 We ran into it with the digital INC, and  
24 it took a while to get that straightened out. And  
25 then maybe that is this building block for where we're

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 going on this already.

2 But and just, it just seems like that  
3 takes, I mean again, back to my impatience. That  
4 takes a long time to get an alignment, you know, not,  
5 it doesn't happen in a quarter or two quarters.  
6 Sometimes it takes years.

7 And, we're not too many years away from  
8 when we want to be able to be ready for this. So, I  
9 was, I wouldn't put that aside.

10 I'd make sure that you know, at least  
11 those conversations are being had so that we know if  
12 the existing guidance can be applied, and actually  
13 work.

14 MR. VALIAVEEDU: So to the workshop, the  
15 summer, we've kind of hit a cadence where we think  
16 annual workshops are the, what we're going to plan on.

17 But we remain flexible to have the entire  
18 workshop in there to tie in with it. But the next  
19 workshop is really intended to go to this point on  
20 guidance, because we plan to finish that regulatory  
21 gap analysis.

22 And, the early look is I think our  
23 regulation is flexible enough to adapt to AI, but the  
24 guidance might be lacking.

25 MEMBER HALNON: Okay.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 MR. VALIAVEEDU: In significant areas.

2 MEMBER HALNON: It's on your --

3 (Simultaneous speaking.)

4 MR. VALIAVEEDU: It is.

5 MEMBER HALNON: I mean, you know --

6 MR. VALIAVEEDU: Very, it is very front  
7 and center.

8 MEMBER HALNON: Okay, good, that's what I  
9 hoped. I hope that we can get that, get that not be  
10 a hurdle to get over.

11 MR. VALIAVEEDU: Yes.

12 MEMBER HALNON: Seeing that we're looking  
13 at it years in advance here, so thanks.

14 MR. VALIAVEEDU: Right.

15 DR. SCHULTZ: Along those lines, or maybe  
16 it's parallel to it. The regulatory sandbox, is that  
17 part of the upcoming workshop as well?

18 In other words, I'm presuming you're  
19 looking for what is going to be included in the  
20 sandbox, and what's outside of it in terms of  
21 application than what inside the sandbox needs  
22 regulatory attention?

23 MR. VALIAVEEDU: For the regulatory  
24 sandboxing, we are hoping to one, identify a test  
25 study that may or may not have a direct implication

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 for reactor use utilization.

2 And then, go about a potential open avenue  
3 to identify any challenges within current area  
4 existing regulations.

5 That may or may not have answered your  
6 question. If it didn't, I can go back.

7 DR. SCHULTZ: No, it did do it. It's  
8 fine.

9 MR. VALIAVEEDU: Okay.

10 DR. SCHULTZ: But eventually you're going  
11 to be including decisionmaking associated with what  
12 will be included in terms of the regulation, and what  
13 will not?

14 What happened? I mean, you talked about  
15 many applications --

16 (Simultaneous speaking.)

17 MR. VALIAVEEDU: Yes.

18 DR. SCHULTZ: -- in your workshops, so  
19 which industry is doing nothing. And in several of  
20 those, regulatory attention will not be required.

21 MR. VALIAVEEDU: Yes.

22 DR. SCHULTZ: The work can continue, and  
23 findings will be useful to the industry. Regulation  
24 would not be required in terms of the way in which  
25 it's done, or the results that are obtained.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1                   Would that, aren't those findings that you  
2                   want to, you want to develop sooner than later?

3                   MR. VALIAVEEDU:   Yes, that's why we're  
4                   maintaining our ongoing collaboration with different  
5                   stakeholders like EPRY, as well, that have, that we  
6                   are able to obtain those lessons learned from their  
7                   utilization of AI ML technologies.

8                   The regulatory sandboxing is more of  
9                   understanding the, it helps support the regulatory  
10                  framework that exists right now, to ensure that the  
11                  staff ourselves, are ready for any potential  
12                  application.

13                  DR. SCHULTZ:   Great, I appreciate the  
14                  detail that you've gone through with regard to staff  
15                  training and implementation.

16                  All right, to me there's more important,  
17                  more important than the hiring of new people that  
18                  understand AI, is training the people that are already  
19                  here and know a lot about regulation, to utilize AI  
20                  and the process.

21                  Thank you.

22                  MR. VALIAVEEDU:   There was a phrase  
23                  mentioned earlier in a previous meeting where you  
24                  could get a really good data scientist, or you could  
25                  get an NRC individual who's extremely experienced

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 nuclear regulator, and then train them with data  
2 science tools.

3 MEMBER BROWN: One other input just to put  
4 it on the plate to think about, as we introduce  
5 software based systems for digital INC, it was always  
6 the concern about common cause failures. And  
7 therefore, the issue of diversity areas.

8 How do you handle that and ensure that  
9 that failure doesn't propagate lock up of the systems,  
10 interrupt driven systems, which all AI is going to be  
11 interrupt driven because it's going to be evaluating  
12 data coming in all over the place.

13 That means it could get confused, and if  
14 you, I'm thinking of downstream now outside of the who  
15 cares realm.

16 Does that mean now we have to have the  
17 thought process of diversity, defense in depth, in  
18 terms of the application of AI into any other safety  
19 or plant control, or even non-safety related but plant  
20 control systems that are just out controlling stuff?

21 How do we do that? Do I have to have  
22 competing AI algorithms making the data and then  
23 comparing those, and then making a, or do I have to  
24 have three sets because I need to have them both?

25 It just, the whole idea of now all of a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 sudden accepting the fact that it's AI, I don't need  
2 anything like that to make sure it's really working  
3 properly.

4 Is it one more technical issue that would  
5 have to be addressed when you get into the world of  
6 plant controls, safety systems, and even what I call,  
7 I wouldn't call it, not all plant systems are  
8 obviously safety related but you need them to operate  
9 the plant. You have to do something if they don't  
10 work.

11 So that is another very, very difficult.  
12 It was hard enough in the regular software world where  
13 you used different devices, you use alternate  
14 software, watch dog timers, all kinds of, and  
15 susceptibility to cyber-attacks, intrusive. How do  
16 you communicate data that's not protected by an air  
17 gap? So that's just one more thing you need to throw  
18 into the hopper in terms of how you apply this  
19 downstream.

20 Because you're going to run into the exact  
21 same issue we've been dealing with for a decade or  
22 more, a couple of decades in terms of incorporating it  
23 into the systems.

24 That's one more way to be very, very, very  
25 cautious.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 MR. VALIAVEEDU: Thank you, that is one  
2 are that we are looking at, especially with the gap  
3 analysis, so I appreciate it.

4 The final item here is to facilitate and  
5 invest in research. We hope to maintain our ongoing  
6 university research grants, as well as maintain our  
7 current FFRs with a specialized focus with, hopefully  
8 promoting AI research within the NRC.

9 The final slide here, if we do our work  
10 right, no one will remember us. The NRC is committed  
11 to the safety and security of the public and the  
12 environment.

13 The hope of the AI team is to pave the way  
14 to ensure that the NRC puts its best foot forward, in  
15 future applications.

16 Our high stake standard for safety remains  
17 unchanged no matter what the technology is. And, we  
18 are working towards ensuring that we have the staff  
19 with the knowledge, skills, and the ability to  
20 effectively regulate these new technologies.

21 I would just take some time to highlight  
22 our next steps, which includes in the spring of 2024  
23 we're going to be publishing the CANUKUS paper, AI  
24 principles paper.

25 We hope to publish the AI regulatory gap

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 analysis as well as that time.

2 In the March timeframe, we'll be hosting  
3 an AI technical session at the RIC that Matt Dennis  
4 has highlighted. As well as host an IAEA AI technical  
5 meeting at headquarters the following week.

6 We hope to update the AI project plan with  
7 revision 1 in the fall of 2024, with our revised  
8 timelines according to what we expect through our  
9 engagement.

10 And, we will always continue our public  
11 workshops and stakeholder engagement as they've shown  
12 to be fruitful in understanding what will be ongoing.

13 Thank you.

14 MEMBER BROWN: I didn't mean to interrupt  
15 you, I just want when you're done.

16 MR. VALIAVEEDU: Oh, I was going to say  
17 just thank you again for giving me the opportunity to  
18 speak, and open it to more questions.

19 MEMBER BROWN: I just wanted to make  
20 papers and things that you had issued in the gap  
21 analysis, if you, and the emphasis and I tried to  
22 understand.

23 So I'm not a designer obviously like you  
24 guys are too smart for me from that standpoint. And  
25 I went to try to figure out something. There was a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 paper in one of the magazines that I read, one of the  
2 publications.

3 And it was so replete with jargon that by  
4 the time I finished, I had not a clue as to what they  
5 were saying when they got to the end, and they had  
6 some conclusions.

7 I would just encourage these papers are  
8 going to be relevant to other folks other than you,  
9 understanding where you're going to go.

10 And if it is steeped in deep learning,  
11 machine learning jargon, that's not going to work and  
12 the resistance is going to be strong.

13 So, it would be good if you could put it  
14 in every day English for people who are technically  
15 oriented, but not fully ensconced in the jargon of the  
16 AI world.

17 So, you know, reduce it to English, in  
18 other words.

19 MEMBER HALNON: Anthony, in your, is it  
20 your intent, your goal, your aspiration, your hope or  
21 whatever, to have the gap analysis ready for the RIC  
22 session? Or at least you know in --

23 (Simultaneous speaking.)

24 MR. VALIAVEEDU: We expect it to be done  
25 in the spring 2024 based off of initial timelines.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           However, as you're probably aware, the NRC is pretty  
2           big with its regulation.

3                       MEMBER HALNON:   Yes, yes.   It just seems  
4           like we'd miss a really good opportunity if we don't  
5           have at least, at least some of the findings to  
6           discuss at the technical meeting of the RIC.

7                       So I know you know, spring is not  
8           necessarily March, but you know, it seems like that  
9           would be a great opportunity to at least be able to  
10          present the findings.

11                      So, that's my comment.   Thanks.

12                      MR. VALIAVEEDU:   I agree.

13                      CHAIR BIER:   Okay, thank you, Anthony, for  
14          actually getting us caught back up on time.

15                      So, and thanks to all the presenters for  
16          a good discussion, and for your patience with all the  
17          interruptions and questions.

18                      So I think at this point, it's time for a  
19          break and we will reconvene at 11:15.   One more  
20          presentation before lunch.   Thank you.

21                      (Whereupon, the above-entitled matter went  
22          off the record at 10:58 a.m. and resumed at 11:15  
23          a.m.)

24                      CHAIR BIER:   For those online, we are  
25          going to get started in just a minute or two after

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 everybody takes their seats.

2 It looks like we're missing a few who are  
3 taking a slightly longer break, but in the interest of  
4 timeliness, I think we will go ahead and get started  
5 on the next session.

6 So I'm happy to introduce Jim Chang from  
7 Research, who has what sounds like a very interesting  
8 presentation on using machine learning for inspection  
9 planning.

10 MR. CHANG: Thank you. My presentation  
11 goes into the topic on regulating the AI to use the AI  
12 for NRC's operation. And our focus is that for this  
13 implementation is informed inspection plan.

14 My presentations are straightforward at  
15 the motivation and end at talk about what we do and  
16 then what data we use, and then this observation  
17 underneath that that we obtain from this project.

18 The motivation was COVID-19 that disrupted  
19 NRC's inspection plan. NRC did not send a inspector  
20 to the site regular (audio interference). So that's  
21 under NRC's risk-informed inspection that we can have  
22 some system that can identify what's a priority of the  
23 inspection. That will be very beneficial to enhance  
24 NRC's risk-informed actions.

25 And I also read the AI machine book that,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 one case is referenced in that was a Netflix success  
2 story. Used unsupervised machine learning to analyze  
3 the movie watched by its subscribers. And then from  
4 there that it helped identify the hidden pattern of  
5 the code clusters.

6 And using that information will be helping  
7 the Netflix to better inform that recommend a movie  
8 for its customers.

9 So then I was thinking about these two  
10 pieces of information together. NRC here, we have a  
11 lot of nuclear power plant performance C suite that  
12 many have documented, licenses and reports or  
13 inspection finding. These things that they consider  
14 as this history of this plant's performance.

15 Can we use the unsupervised machine  
16 learning, it can bring the information together and  
17 then that's identify a hidden pattern. I call this a  
18 safety cost, later I will explain what safety cost  
19 mean.

20 So this objection was try to perform a  
21 feasibility study simply that looking at what the data  
22 we have here and then the snapshot technology we have  
23 this stage. What's combination of them that how good  
24 they are, that to achieve this purpose, informing the  
25 inspection planning.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           So I called this the hidden pattern as a  
2 safety cost. Safety cost I define as the failure  
3 mode, failure causes of this structure, system,  
4 components, and their failure that has consequence to  
5 the nuclear power plant safety.

6           So this combination of this information  
7 together usually call this the safety cost. Try to  
8 identify using unsupervised machine learning to  
9 identify this safety cost. And we had a benchmark  
10 that the teams tried to achieve. It was in the NRC's  
11 reactor oversight program that periodically that  
12 publish the operating experience communication.

13           And this was a communication published in  
14 November of last year that, it identified five power  
15 outage events impacting security system operation. So  
16 that's consequence of power issue impact the security  
17 system operation.

18           And then there's SSC and here is the  
19 primary and the backup power tried to reach the  
20 security system. Failure mode just simply not  
21 providing electricity.

22           This is a communication that's also, I  
23 listed it by instance, by operating experience. It's  
24 identified at 2022 has two events, '21 has three  
25 events. But the 2021 three events, all them consider

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 as the random failure, power failure.

2 But the 2022 two events, all that involved  
3 is human failure. One is the maintenance, doing the  
4 maintenance and all the system repair that caused the  
5 outage.

6 So with this information that's from the  
7 inspection finding, that's the original communication  
8 suggested our inspector, when they performed the  
9 inspection procedure that's related to equipment  
10 performance testing and maintenance that focus on the  
11 human impacts on power supply.

12 So this provided this zooming in the focus  
13 that to me is risk-informed information to help our  
14 inspector based on the past event to help our  
15 inspector when they do this general inspection in this  
16 area that's focused on -- that's a cause related to  
17 the operating experience we observed in the past.

18 So that was the things that I tried to  
19 achieve, see that can we use the unsupervised machine  
20 learning to help identify these things. The approach  
21 is that I got a funding from the Office of Research  
22 Future Focus of Research funding. And then  
23 established a commercial project contract to the AI  
24 company that's SphereOI.

25 In addition to this, NRC also formed a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 team that's we have machine learning experts that Trey  
2 Hathaway, that's sitting here, that he was in my team.  
3 And then I also has NRR staff from Reactor Oversight  
4 Program, that Jason Carneal. That's in my team.

5 I am not a machine learning expert. I am  
6 not this Reactor Oversight Program pilot. We just  
7 bring the team together to work with this contractor  
8 to perform this project.

9 The task, two tasks I identified for this  
10 thing, the first thing that's we don't know what's the  
11 current state of the AI. Just simply try to get a  
12 glance of the what's the landscape there.

13 So the first task was try to understand  
14 the -- evaluate these are big plant companies. Their  
15 AI platform that's a high-level version that to find  
16 out which one may be best for this whole purpose.

17 The intention was try to use this pre-  
18 trained algorithm as much as possible instead of NRC  
19 does put in a lot of effort try to develop algorithm.

20 And then the second task was a lack of  
21 platform that to identify these safety costs. That's  
22 a issue with two task. The company that, really quick  
23 pace, we have weekly meeting and that work the project  
24 was completed. We did it in four months.

25 So the task one here that the contractor

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 identified the topics for doing this type of test  
2 what's the topical like, in this process, and then  
3 evaluating give a score, weighting, wright and score  
4 of this four platform.

5 Come to the end, that's Azure, Microsoft  
6 Azure, and Amazon's AWS was ranking the higher. But  
7 this doesn't really help much because they come to the  
8 task to -- one thing important was a notebook  
9 integration.

10 And there was a notebook that can access  
11 these platforms, algorithm library so that's -- it's  
12 not -- this notebook is independent from all these  
13 platforms. And that was the Jupyter Notebook was used  
14 in the task two analysis.

15 So come to the end that this evaluation  
16 doesn't really affect the decision on choosing which  
17 platform to go. Go to using the Jupyter Notebook is  
18 a free software that downloaded. We need this kind of  
19 open-source library to perform the functions.

20 To perform the task two here that the  
21 contractor develop to bring in the test that's a  
22 inspection and then the former test that's item one  
23 there. And go through the series of process, the  
24 components of process in this information.

25 And then come to the end that we

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 identified these safety costs. And then that item  
2 three that is the visualization of one of the costs.

3 For all these components, the components  
4 in this pipeline, there's multiple algorithms can do  
5 the function. So the -- come to the end the -- what  
6 we do here is try to trial and error the different  
7 algorithms and then try to evaluate, see which  
8 algorithm has better performance.

9 And come to then end that I identified  
10 this optimal combination for this pipeline performance  
11 that can take data from the front end of the text and  
12 then come to the end that identify this cost and  
13 represent this safety cost.

14 So this diagram talk about the things that  
15 the contractor tried. On the top is this pipeline  
16 component, on the end. The first element taking the  
17 text, completion of this inspection finding. In that  
18 the contractor tried 15 different, I'll say that 15  
19 algorithms to process these original information.

20 And then come to the end, it selected  
21 three of them that are better performance. And then  
22 leading to the next components. Next component has  
23 three different, five different -- five different  
24 algorithms. And then come to the end, select one.

25 So you see this, that's a lot of trial.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 And for the first one text this portion here, it's we  
2 took eight different sample, eight inspection, eight  
3 inspection. And then try this 15 combination. Come  
4 to the end, select these three. So that you see that  
5 was try and see that how --

6 MEMBER MARTIN: Question. I'm trying to  
7 understand what's the information that's being fed  
8 here. What's the specifics. And the maybe use an  
9 example.

10 MR. CHANG: Yes. That's in my next slide.

11 MEMBER MARTIN: Okay.

12 MR. CHANG: So input information, that  
13 original was trying to that we -- NRC has inspection  
14 reports that about 20,000 inspection reports from year  
15 2000 that's publicly available on the website. That's  
16 was -- inspection report was the information.

17 But through a process, so I learned that  
18 our key process NRC maintain this database that's  
19 excel database that has all these. This is not  
20 inspection, it's inspection finding. Inspection  
21 finding data is something come to the like -- more  
22 green, the type of inspection finding there.

23 But come to the end, that many of these  
24 will be identified as green finding. But these are  
25 the inspection findings that are keeping in this

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 database that's from the study from the year 1998  
2 about 15,000 unique records.

3 In this database that column F was item  
4 introduction. That was a decision about this  
5 inspection finding.

6 So that's what's come to the end, so while  
7 since we have this one, I don't want to focus on our  
8 resource. To focus on what we want to achieve so that  
9 we simply take in these as original, the discretion  
10 here in the item introduction. This column has the  
11 input information.

12 This input information that averaging as  
13 1,649 words. And minimum is 42 words. Maximum, 7670  
14 words. So that's the range of that expression there.

15 We took the discretion (audio  
16 interference) that's all the information we need. But  
17 it was a limit in there, too, these sentence  
18 transformation models limitation that it come to the  
19 -- reached its sudden capacity it will truncate. It  
20 doesn't take in information anymore.

21 So if we use the full text that's a long  
22 text, that's a data on the text it will simply just  
23 dismiss that because of limitation of this -- sentence  
24 transformer model limitation.

25 So that was the contractor. Okay, what

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 about let's try some other algorithm to take these  
2 full text as the input, and then that this algorithm  
3 can do the summary of this text. And then that  
4 generates as the, I call it condensed summary. And  
5 then input in to the pipeline. That's the approach we  
6 did.

7 So except the full text, we tried 14, the  
8 contractor tried the 14 condensed summaries in here  
9 that it divide into the three category. One is a  
10 summary technique algorithm. That's a try 70 kind of  
11 algorithm in the summary technique.

12 And then Q&A three key phrase extraction  
13 to try four (phonetic). And that -- what we take.  
14 And then that some of the AI compound, AI algorithm.  
15 That's also allow us to provide some inputs that cause  
16 semi-unsupervised machine learning.

17 In there, they're taking out the things  
18 that do have a focus on nuclear safety. These AI  
19 options, they are trend from the open website that  
20 Wikipedia, social media. It doesn't have a specific  
21 focus on the nuclear safety. But some of them, they  
22 allow us to bring into the input that what are the  
23 things that we need to pay more attention. On the --  
24 so from the NRC provided 1,004 acronym like MSIV,  
25 these type acronym.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           And then also one is 400 common failure  
2 mode, like a inoperable misalignment corrosion. These  
3 are word that we were interested in on the special  
4 failure mode. That we really want to, knowing that's  
5 -- not just feed it a general word. So that provide  
6 us more useful information.

7           And in addition that we acquired 269  
8 NUREGs and 195 research information letter. This  
9 technical report was the contractor want to use to see  
10 the coherence of this word, term that appear in the  
11 text and then compare it. But it's for further  
12 information. But they're not really helping the text  
13 too (audio interference) check, it will function.

14           And then also the stop word removal that  
15 in addition to this general remove the stop word, this  
16 type of word, we also provide that contractor that  
17 also look into the outputs and then seeing whether  
18 that term that we see that we need to remove them.  
19 That they consist safety system reactor, these happen  
20 to open them.

21           If we don't remove the stop word, that the  
22 -- somehow the group that focusing on these terms,  
23 that's not a one we want.

24           Showing the example, that's one, at least  
25 one process. On the left-hand side for column here,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 you see a lot of these operating company, Entergy,  
2 Exelon. And that was because we did not remove the  
3 these words, these terms as a stop word.

4 And then after we saw the result, well, no  
5 this making the -- this not the things we wanted to  
6 focus on. The process, focus on the company instead  
7 of focus on the safety -- the system, structure system  
8 components.

9 So that's how we work to remove this word.  
10 And then on the right-hand side, have to remove this  
11 word that's on the right-hand side, Fort Bragg  
12 (phonetic), showing these customs, forming that's  
13 these become more like trip and auxiliary feedwater,  
14 these are the kind of level that we have more interest  
15 in. That's a kind of stop word removal.

16 So come to the end that goes through this  
17 process, identified hundreds of these safety costs.  
18 And then it's a long list here. It's a part, and I  
19 just show you some -- an example here.

20 These are costs is represented by word  
21 cloud or bag-of-words. Here, that's on the Excel  
22 spreadsheet here. The topic -- the next minus one is  
23 -- has 5,382 inspection findings. These are the  
24 inspection findings, could not group into any cluster.  
25 So these are kind of, we called outliers costs.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           And top on Deal One (phonetic), that's a  
2 deal, the topic deal that has 927 inspection findings.  
3 And these are the words kind of forming word bag, word  
4 cloud, that's describe what the cost are about. And  
5 et cetera.

6           So these are the way that it -- to the end  
7 of Type Nine, these are the table that was generated.  
8 And then that's for each role that's a work and going  
9 to that's a -- what's this then, 927 inspection  
10 findings. That information can be tracked, if we  
11 want.

12           And then, after that, was this just  
13 showing the three different input information  
14 technology coming in, that they come to the end, that  
15 what this same information but forming the  
16 synchronization of different clusters. It's all  
17 related to the RCIC system, things associated with  
18 that components.

19           At the end of this project that we  
20 fortunate has a operating experience computation.  
21 This time is about a safety security system. The  
22 15,000 inspection findings I mentioned here, all of  
23 them are safety system, not security inspection. So  
24 security system inspection is not within the scope  
25 here.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1           But in this original communication it  
2 identified why events relate to improper calibration  
3 and the maintenance of the radiation monitor and dose  
4 assessment equipments.

5           One company has two events. They said  
6 well, since we are already identified these costs, can  
7 we go back to the things we identified, seeing that's  
8 how this five events was clustered.

9           And we found out that one of these, I  
10 think it's the top one, 2022 event, was not in this  
11 original -- inspection finding data. What we can find  
12 four of these -- four of these events.

13           And this was the results that the feedback  
14 led to well, these operating expense was identified in  
15 this exercise. So that's from the full item  
16 introduction. That means that we took the summary of  
17 data and looked into this pipeline.

18           In summary, this column was, using the one  
19 of these summary technique. But generally the  
20 condensed summary and put into this pipeline and key  
21 phrases, technique.

22           So that all these are four events, four  
23 operation, you know, was in one of these safety costs  
24 identified. None of them was put into this outlier  
25 bin. But one day we saw that well, summary report

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 that has at least three of events identified in the  
2 cluster number one. Well, sounds good, but the number  
3 one has more than 900 inspection findings.

4 So that's still not practical that work  
5 identified all these things. And that in terms of  
6 that we NRC operators people need to, you need to  
7 squeeze through this 900 events to identify what are  
8 things are that are maybe not working.

9 This current stage is still not come to  
10 the -- demonstrate some success, but come to the level  
11 of the data, we say what it's used for.

12 But I want to say these 900 event, that's  
13 we are talking about inspection findings dated from  
14 the 1998 to 2022. So that was maybe that using the  
15 dates that we focused on the most recent data maybe  
16 give us some more focus. But we didn't go to more  
17 analyses for that.

18 It was because the future focus of  
19 research is a small project. It's for research for  
20 you to identify information. And then it really find  
21 out some that information since indicates some  
22 potential that it become a seed (phonetic) process --  
23 seed project that action or the more formal way of  
24 doing the research development these areas.

25 And so that was the way we concluded the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 is research project. The project here that does  
2 identify this summary technique was very useful, the  
3 R piece (phonetic) that was there, while that was  
4 useful, that summary that's taking these original  
5 inspection summary, and then that's provide condensed  
6 version of summary. That was a useful, good use of  
7 workflow purpose in the sense that can use this  
8 technique to provide summary and then NRR staff can  
9 see a view of the condensed summary that's to reduce  
10 work.

11 The second bullet about is using this  
12 based on the way that we do. Certainly there's a lot  
13 of things that we can improve, including the stop word  
14 or trying some technique. At that time we saw it as  
15 time-consuming. We don't want to go forward moving  
16 that, try to optimize our future focus research.

17 And if that's we have additional funding,  
18 that's we may want to spend it to fine-tuning that may  
19 be able to refine the results. But whether that  
20 refine the result will it come to the -- become a  
21 practical skill, I don't know. So that's why --  
22 making the conclusion that what, based on the result  
23 we saw, it has potential, but it's not conclusive that  
24 for practical application.

25 MEMBER MARTIN: A comment on that. About

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 six years ago, I was doing -- this just shows you what  
2 I do in my free time. But so I resonate with what  
3 you've done here. I've used the algorithm for the  
4 parsing, the documents or texts. And the association  
5 -- the associating frequency and the presentation of  
6 that information.

7 I did this with the water reactor  
8 evaluation model document, you know, which was  
9 published in the 70s. Again, for fun. I have it here  
10 on my screen.

11 But so this, what you're doing resonates.  
12 I will say, though, in the presentation, and I was  
13 doing this because I wrote a blog and about water  
14 reactor evaluation model. And I wanted to highlight  
15 the kind of things that were important, right.

16 And I saw that algorithm, and of course in  
17 my mind, I had an idea of what should be important.  
18 And then used the algorithms and parsed the  
19 information. And of course I had the problem with the  
20 stopper, you had as well.

21 And then I probably spent the next six  
22 hours trying to find the right set of words to  
23 eliminate to get what I want out of it.

24 So what that means is there's a huge  
25 amount of uncertainty associated with that, with the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 process. But you know, in this business, you know,  
2 we're addressing safety issues.

3 We know there's a huge amount of  
4 uncertainty. Of course we talk about it in the  
5 context of our probability risk assessments and such.  
6 And uncertainties on order -- on the order, an order  
7 of magnitude are pretty normal.

8 Do you see some synergy with, you know,  
9 methods like this? I mean, because it's incorporating  
10 a natural language translation of sorts. And it's I  
11 would say corroborative information to, you know, more  
12 quantitative risk analysis. Is there synergy, have  
13 you thought about synergy in that realm?

14 You know, going back to our earlier  
15 presentation about being, you know, Matt's comment  
16 about the, you know, being the best version of  
17 yourselves. You know, using the tools that we have  
18 today to do a better job and develop more confident.

19 I know this was a small project. It's  
20 kind of fun to listen to here. But to take it in all  
21 seriousness and is there an opportunity? Do you see  
22 opportunity? Will you go farther with this? I mean,  
23 I know it was at the discretion of the agency for you  
24 to do this for a project.

25 But where does it go from here? And I see

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 one opportunity that might be interesting, if not  
2 useful.

3 MR. CARNEAL: Can everyone hear me? I'm  
4 Jason Carneal, I'm in the Operating Experience Branch.  
5 And I was working with James on this project.

6 And as with pretty much every start that  
7 you get in this area, we all ran into the same  
8 problems with those stop words where yes, you're  
9 probably putting in some bias there with what you  
10 think that you think the output should like, should  
11 look like.

12 What we were doing in this project was  
13 trying to just give it a minimum baseline to fake out  
14 the general Wikipedia-style stop words or the ways  
15 that the algorithms were trained with general  
16 language, and give it just kind of a little bit of a  
17 leg up for the business that we do at NRC and see what  
18 we could see in these safety clusters.

19 And it was a small effort. Of course we  
20 could optimize that in the future. The power that I  
21 see with the safety clusters and identifying those  
22 unnoticed trends that the human eye can't see, in my  
23 group we're working with OpE documents.

24 We have about 100,000 documents in our  
25 store, and that number of documents expands by 1000

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 documents a year. That's just the NRC documents that  
2 have no structured information. So it's all free  
3 text. You have to have some way to apply some level  
4 of grouping to that if I'm to hope to find some kind  
5 of a trend in that document set.

6 We also have about 200,000 industry  
7 documents where we have those texts. So the power  
8 here that I see in the future is particularly for  
9 those trends that are hidden, where it's not what we  
10 think we would see in the trend.

11 So the top five trends, if we went through  
12 James's list, it's kind of what we'd expect for safety  
13 cultures. The safety clusters. When I looked down  
14 into Items 10-15, there's something interesting here.  
15 I've never associated those words in my mind.

16 MEMBER MARTIN: Right, and I'd also say  
17 that statistically speaking, when you start playing  
18 with the uncertainties, those, you know, numbers 5-15  
19 will change. And at some point, they're worthless,  
20 right. But statistical method, that's the nature.

21 Now, if you get consistency with the  
22 variability of softwares or whatever the random number  
23 (audio interference) there, that those top five are  
24 there time and time again, there's probably something  
25 to it.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           But yeah, I would certainly not look  
2 beyond. You know, when things start changing, then  
3 it's not (audio interference).

4           MR. CARNEAL: And just one other thing,  
5 that as far as future work with these, being able to  
6 group these with unsupervised learning, the way our  
7 program currently operates, we're relying on four  
8 people assigned to each region. They're looking at  
9 the reports as they come in.

10           So if we're going to identify a trend,  
11 it's usually knowledge of that personnel over a period  
12 of time. Oh, I remember this happened three years  
13 ago, let me go look at this. This would allow us to  
14 take a more proactive approach and try to get at least  
15 a hint to the engineers that are reviewing the reports  
16 that there might be a trend here for these 100  
17 documents, you might look at a few of those.

18           MEMBER HALNON: So okay, I can't help, and  
19 I know there's one minute to go, with all the hype on  
20 AI coming around, I can't help but be disappointed in  
21 that inconclusive result, given the fact that we're  
22 not talking about that many findings. A hundred a  
23 year in the industry maybe at this point. Well, maybe  
24 more than that if you get plants in trouble.

25           But and the inspection reports are very

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 structured. So they're, you know, they're very  
2 descriptive of what -- and matter of fact, they're  
3 repetitive in a lot of ways.

4 So I can't help but be disappointed. I  
5 guess if I read your blog six years ago I'd be  
6 disappointed for six years, you know. I'm glad I  
7 didn't. But I would like to.

8 Anyway, so James, do you see light? I  
9 mean, I know you used the term machine language, and  
10 we've been using the term AI all morning. And I'm  
11 sure that there's some overlap Venn diagram you could  
12 show me that says that there were almost the same  
13 thing but not quite, or however you want to define it.

14 But do you see some application in the  
15 future where, you know, you're not going to have to  
16 have this cognitive trending people dedicated to it?  
17 I mean, it seems like if you could take AI and say  
18 please write me a research paper on umpty squat and I  
19 want to turn it into my professor and I'll get an A on  
20 it.

21 It seems like you should be able to take,  
22 what is it, 90 plants times four inspection reports a  
23 year, 360 inspections that are all pretty well  
24 structured the same and say give me what the trends  
25 are in there.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           It almost seems like that would be, from  
2           an AI perspective from what we've heard, is no-  
3           brainer. It should be able to come out with some very  
4           good stuff without stop words, without other things  
5           that. Because again, the inspection manuals are very  
6           descriptive on you word a finding and how you word  
7           cross-cutting issues and stuff like that.

8           So what's your outlook? I mean, what do  
9           you think?

10           MR. CHANG: Last year, last EPRI published  
11           a technical report that has a document the industry  
12           using machine learning for corrected action program.  
13           It was, the purpose was use of the machine learning to  
14           screen out these reports, the reports certainly then  
15           have a safety implication.

16           And that -- in that EPRI document, two  
17           success case that it reduced the workload and come to  
18           one million dollars a year, that kind of saving. So  
19           to me that's -- this morning we already talked --  
20           mentioned that to find the AI for safety system  
21           control, that seems like that it's distant, away  
22           future.

23           But the way the things seems that safety  
24           important bucket focus on reducing the workload that  
25           (audio interference), providing that the second layer

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 of the quality assurance, that I think it's very  
2 radical.

3 MEMBER HALNON: Okay. I want to go back  
4 to early teens or the 2018s when we were discussing,  
5 the agency was discussing with the industry about  
6 substantive cross-cutting issues that took two or  
7 three at that point hits on this cross-cutting issues.  
8 And if you got hit with a "substantive cross-cutting  
9 issue" it could cost millions of dollars to get out of  
10 it.

11 So even though you can save millions of  
12 dollars in resource and other things, sorting through  
13 the ten thousand corrective action documents you may  
14 have, you could also be chasing ghosts to the point  
15 where it's trying to fix a non-problem. But you're  
16 creating a problem by trying to fix it.

17 So there's got to be a check and balance  
18 there as we go forward too. And I know you saw that  
19 with the stop words and other things. You saw the  
20 pitfalls that could get into it. And I guess if we  
21 read your blog, we probably would have known that  
22 already.

23 But nevertheless, this is just some  
24 thoughts. I think that there's a application going  
25 forward with this. And I think that it's -- as the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 industry goes and does it with the corrective action  
2 programs, the agency certainly could do it with this  
3 smaller subset of findings that they have.

4 Because not very -- I mean, a site has  
5 seven to ten thousand corrective action documents, and  
6 we dealing with 100 findings. So it should be  
7 relatively straightforward, at least in my mind here.

8 MR. CHANG: Yeah, you bring the topic back  
9 to the regulate AI.

10 CHAIR BIER: Greg, one minor comment. You  
11 need to speak up for the people in the back of the  
12 room.

13 I had a few questions and comments that I  
14 will try and make very quick. First of all, I'm  
15 different than Greg. I'm usually a skeptic, but I'm  
16 very excited about this application.

17 I mean, the methodology may not be there  
18 yet. Maybe we have to have a different approach or  
19 wait another couple of years 'til the software is  
20 better or whatever. But I like it because it's an  
21 example of that kind of offline type of advice where  
22 it's not making a decision for anybody, it's just  
23 surfacing information that then the decisionmaker can  
24 look and take into account.

25 So I think that's very promising. I had

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 a few questions. One is I know some of the software  
2 packages kind of Hoover up your data and send it back  
3 to the mother ship. And I know that Azure doesn't, it  
4 lets you sandbox and keep your own data for yourself  
5 and not send it back to Google or Microsoft or  
6 somebody.

7 Can you comment whether the other software  
8 packages you thought about have that pitfall, or  
9 they're all similar?

10 MR. CHANG: No, I don't. I haven't  
11 thought about this question.

12 CHAIR BIER: I mean, NRC data is pretty  
13 much mostly public anyway. But in other applications  
14 that can be a big issue.

15 Second of all, which years of data did you  
16 use?

17 MR. CHANG: This is the inspection  
18 findings from 1998 to 2022.

19 CHAIR BIER: Okay, because one of the  
20 issues is like the more -- the shorter the timeframe,  
21 you have less data, but it's more relevant.

22 MR. CHANG: Yeah.

23 CHAIR BIER: So that might be another  
24 parameter to play with, is what if you took only most  
25 recent five years or something. Maybe you would get

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 better relevance.

2 The third question, and I don't know if  
3 this is a question for you, it might actually be a  
4 better question for Jesse. I know in the area of  
5 health, like for reading mammograms, they have found  
6 that like a doctor plus an AI does better than two  
7 doctors. Because the AI sees different things than a  
8 human would see, and then you can get better coverage  
9 of what's going on.

10 But I'm very concerned about kind of the  
11 computer equivalent of social loafing. Like, you know  
12 the computer's going to look at it anyway, so after a  
13 while the human gets lazy and stops paying attention  
14 and just acts on the computer advice. So I'd be  
15 curious if either you or Jesse have given that a lot  
16 of thought yet.

17 MR. CHANG: Certainly that's my expertise,  
18 human reliability. Yes, you put a human from this  
19 first night to the second night as a PO checker  
20 (phonetic) or monitors of positions. So that kind of  
21 performance certainly that we have --

22 (Simultaneous speaking.)

23 CHAIR BIER: Yeah. Jesse, if you want to  
24 expand on that at all?

25 MR. SEYMOUR: I appreciate it. And so

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 this is Jesse Seymour from the Human Factors Branch.  
2 One of the things I would build on James' point is  
3 when you have essentially, you know, the human acting  
4 in almost like a peer check type of role to something  
5 that AI is doing, there is a phenomena that arises and  
6 it has to do with the scrutability of the AI's  
7 process.

8 So again, if two professionals look at a  
9 given product independently and they disagree, they  
10 can then confer, examine each other's thought process  
11 and figure out why there's a disagreement and perhaps  
12 take something away from that.

13 With AI, it's a bit of a black box due to  
14 the nature of neural networks and so forth. And it  
15 may not be possible even for the people that have  
16 designed again the machine learning application or  
17 whatnot to fully understand what happened in between  
18 the input and the output being received.

19 So again, it's a complex matter. And  
20 James, I'm not sure if you have anything, any more to  
21 that point.

22 MR. HATHAWAY: This is Trey Hathaway,  
23 Accident Analysis Branch. I was going to address your  
24 first question about hoovering up data.

25 CHAIR BIER: Oh, super.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 MR. HATHAWAY: The models used, there's  
2 essentially like ChatGPT paradigm that is kind of  
3 closed off. And if you use it and it's free, they get  
4 it if you have a local copy. You know, it's a little  
5 differently. But these models, there are tons of  
6 these models out there.

7 These particular models you download. You  
8 essentially, you essentially get the weights and then  
9 you're -- I'm getting told that I need to speak up.  
10 You essentially download the weights, and then you  
11 have the model locally. And then you can start doing  
12 things like fine-tuning it on your own language and  
13 things like that to kind of help.

14 CHAIR BIER: So it does not phone home  
15 with all your tons of data.

16 MR. HATHAWAY: That's my understanding,  
17 yes.

18 CHAIR BIER: Thank you. We have time for  
19 one or two more quick questions or comments.

20 MR. CHANG: To the member's earlier  
21 question that, well, you asked that what we are trying  
22 to take it from here to next step. Currently, that  
23 research, this future focus research that's give us a  
24 wayfinding to do this, it's our research results. And  
25 then it's meant to be a seed project.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1           So what is taking the further step that  
2           our managers, one that they use always that use any  
3           request. So that was Jason that anticipate this theme  
4           that he know that what's the benefit to bring in his  
5           -- our people, and that's his response. Each to reach  
6           out his manager. Easier to research it, a system  
7           request to research it so that we can -- do the  
8           additional study in this topic.

9           MR. CARNEAL: And James, just to circle  
10          back to the other question for liming the year range  
11          from 2000 to probably last five. We've done some ML  
12          studies in my group to try and categorize OpE reports,  
13          and that it has a major impact in the accuracy if we  
14          only look at the last five years we get much better  
15          results.

16          And I would imagine that for algorithms  
17          for like this, we would see some similar results.  
18          Because the way that the inspection reports are being  
19          generated now is much different than in the past. We  
20          had people writing free text back in the past and  
21          going through all these reviews that were  
22          inconsistent.

23          Right now, since 2018, what appears in  
24          that database is going to be what appears in the  
25          inspection report. Because they have the option to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 auto-generate that text. So we should be able to  
2 discern a little bit better for the more recent  
3 reports from the boiler-plate language versus the  
4 actual meat of the inspection.

5 O: Just to address your range-of-date  
6 issue, in trying to assess what you can do with the  
7 technology, that's kind of what you're talking about  
8 here. Aren't you liable to the 1998 data seems kind  
9 of not relevant? It's been adjudicated, something's  
10 been done.

11 Why wouldn't you look for a more active  
12 data set where things people have not made decisions  
13 hadn't been closed out? I'm not so sure how 25-year-  
14 old data is going to tell you what you can use this  
15 for as opposed to like the last five years. It's  
16 active data that people have made decisions, and not  
17 it's just whether the decisions were correct or not.  
18 Not correct, but were as good as they could have been.

19 So I'm just getting too much data that's  
20 not really -- if it's really old, it doesn't do you  
21 much good in terms of getting to assessment. That's  
22 my only point.

23 MR. CHANG: Yeah, that's -- the data  
24 quantity, that was the concern at the very beginning  
25 of this project because of what we know, that today's

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 -- this algorithm needs a lot of data. And one  
2 example that EPRI report I mentioned, that the example  
3 was using 600,000 records for the correction action  
4 open, together 80-something percentage with a  
5 successful rate.

6 And so that was -- so before this project,  
7 we know nothing, it was just take whatever we have in  
8 the excel database solely as is. So now that's a  
9 helpful input that we have funding for continue work.  
10 That's something that we will take the recency into  
11 consideration.

12 MR. HATHAWAY: Yes, this is Trey Hathaway.  
13 I think just, sorry, I talk quietly. Sort of talking  
14 to your point, the idea of a lot of these natural  
15 language processing techniques is you're trying to  
16 have a signal to noise.

17 So when you do clustering, you're applying  
18 the model to the document and sort of getting features  
19 that the model thinks are important to the documents.  
20 Or recent documents, if it is kind of like now more  
21 homogenized, I guess, in how you're developing it.  
22 That signal to noise is going to be kind of consistent  
23 across those documents.

24 When you start introducing older  
25 documents, you might sort of change that signal to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 noise to where it's not that just can't do it, it's  
2 that you have to spend more effort in that cleaning  
3 part of the text to sort of get out that text that's  
4 not really relevant.

5 And that's kind of the challenge with  
6 these. Eight percent of the work is just getting the  
7 text in a way that you're getting rid of a lot of the  
8 noise to kind of focus on what's important.

9 CHAIR BIER: We are going to need to end  
10 the meeting now because we have another meeting in  
11 this room over lunch. So thank you very much.  
12 Hopefully some of the conversations can continue out  
13 in the hall or whatever. But thank you for a good  
14 morning.

15 (Whereupon, the above-entitled matter went  
16 off the record at 12:03 p.m. and resumed at 1:06 p.m.)

17 CHAIR BIER: Okay. Now I think we should  
18 be back in business. Can somebody online hear me?

19 MEMBER MARCH-LEUBA: Yes, I can hear you.

20 MEMBER DIMITRIJEVIC: Yes, we can hear  
21 you, Vicki.

22 CHAIR BIER: All right, thank you. So,  
23 Bruce, do you hear me, and can you say something so we  
24 can check that we hear you? Ah.

25 MR. HALLBERT: Sure. I can --

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 CHAIR BIER: Okay.

2 MR. HALLBERT: I can say whatever you  
3 like. Can you hear me okay?

4 CHAIR BIER: I think we're good. Yes, we  
5 hear.

6 MR. HALLBERT: Sounds good.

7 CHAIR BIER: And are you going to share  
8 your screen for your own slides?

9 MR. HALLBERT: I am.

10 CHAIR BIER: Okay. Then I think we are  
11 ready.

12 (Audio interference.)

13 MR. HALLBERT: Whoa.

14 CHAIR BIER: Oops.

15 (Audio interference)

16 MR. HALLBERT: Okay, we had a little bit  
17 of an echo there, but I think we got that resolved at  
18 the moment.

19 CHAIR BIER: I think it sounds much better  
20 now.

21 MR. HALLBERT: Okay, great.

22 CHAIR BIER: So, on that I think you can  
23 just go ahead and get started with your presentation  
24 since we're running a few minutes late. May as well  
25 get it going.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 MR. HALLBERT: Yes.

2 CHAIR BIER: Thank you, Bruce.

3 MR. HALLBERT: I'll do that. Thank you  
4 very much. Thank you for the opportunity to  
5 participate in this meeting.

6 I am Bruce Hallbert. I'm the national  
7 technical director for the DOE-sponsored light water  
8 reactor sustainability program. And with me this  
9 after we have Craig Primer and Ahmad Al Rashdan from  
10 our program who will also be talking a little bit  
11 about our R&D activities. Especially related to  
12 artificial intelligence machine learning.

13 I want to also recognize in the call we  
14 have Ms. Alison Hahn, who is the federal program  
15 manager for the LWRs program currently. And she is  
16 also one of the office directors at the Department of  
17 Energy.

18 I'd like to talk a little bit about the  
19 goals and objectives of the LWRs program, as we call  
20 it. The goal of the program is to enhance the safe,  
21 efficient and economic performance of our nation's  
22 nuclear fleet and to be able to extend their operating  
23 lifetimes.

24 I'm picking up some feedback on my end.  
25 I don't know if you're picking it up on your end as

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 well? No, maybe not.

2 MEMBER DIMITRIJEVIC: Yes, we can hear the  
3 feedback here too, so I don't know where it comes  
4 from.

5 MR. HALLBERT: Okay. Sometimes it's from  
6 when somebody has their microphone open. So as long  
7 as everybody else is muted I shouldn't be picking up  
8 feedback.

9 (Audio interference.)

10 MR. HALLBERT: So we achieve our  
11 objectives by supporting the long-term operation of  
12 existing nuclear power plants by deploying innovative  
13 approaches to improve the economics and economic  
14 competitiveness of light water reactors in the near-  
15 term, as well as in the future energy markets. And  
16 sustain the safety, improve the reliability and  
17 enhanced economics. We go about this by conducting  
18 research in the five focus areas that you see on the  
19 bottom left of the presentation, which I'll be talking  
20 more about in a moment.

21 In the bottom right graph of this slide  
22 sort of brings it all together. Our focus is on  
23 enhancing economic competitiveness by helping plants  
24 to reduce their O&M costs and looking into  
25 opportunities to diversify revenue. Especially for

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 plants that find themselves in electricity markets  
2 where they may be under economic pressure from, you  
3 know, subsidized renewables and inexpensive natural  
4 gas.

5 On the right side we also are addressing  
6 the long-term performance of materials, structures,  
7 systems and components, as well as managing the aging  
8 and technology obsolescence of some of the systems and  
9 technologies that are used to operate nuclear power  
10 plants today.

11 I'd like to talk about each of the five  
12 R&D areas of the program as part of the overview. And  
13 I'll also have some remarks on artificial intelligence  
14 within the LWRs program that I think provides some of  
15 the context for what you're going to hear from Craig  
16 and Ahmad.

17 The first area of R&D is plant  
18 modernization. The goal of our research in plant  
19 modernization is to facilitate modernization at  
20 operating nuclear power plants. We do so by  
21 developing technology and modernization solutions that  
22 address aging and obsolescence challenge. But they're  
23 not just about replacing old technology with new  
24 technology, they're about delivering a sustainable  
25 business model that ensures continued safe and

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 competitive operations. And, well, we can talk more  
2 about how we accomplish that.

3 But the focus is really on long-term  
4 management of plant systems. And when we talk about  
5 long-term operations we mean especially from 60 years  
6 and beyond. And we also are addressing nuclear cost  
7 competitiveness as nuclear power plants face cost  
8 pressures from a lot of power generation sources. And  
9 I'll talk a little bit about that in my forthcoming  
10 slides.

11 And of course, one of the things that  
12 we've really learned from the experiences of the  
13 pandemic is that it's very important to address worker  
14 attraction and retention. Some of the digital  
15 technologies that we're working with through our  
16 program, as well as with the industry, really are a  
17 technology base that the new workforce is more  
18 familiar with, and also see as a part of their long-  
19 term career prospects.

20 An example of one of the ways that we're  
21 working with the industry to modernize the fleet is a  
22 project that DOE is sponsoring in cost sharing with  
23 Constellation. And that project is to replace the  
24 reactor protection technologies at Limerick  
25 Generating, both Limerick Generating Stations, in

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 Pennsylvania.

2           And this is a project, interestingly  
3 enough, that really was brought to us by the Nuclear  
4 Regulatory Commission because of some of the new  
5 approaches to licensing digital at digital  
6 instrumentation and controls, especially for safety  
7 related types of applications. And so this is a  
8 collaborative effort between Constellation, DOE and  
9 the NRC.

10           And it does focus on the first-echelon  
11 safety instrumentation systems. We've been doing this  
12 now since around 2021. We're now in the, approaching  
13 the fourth year of R&D efforts towards the full  
14 replacement of the systems. And one of the roles that  
15 the Department of Energy, and the Idaho National  
16 Laboratory specifically play is supporting the human  
17 factors aspect of that control room modification  
18 modernization project.

19           This slide highlights one of the recent  
20 activities that was conducted at INL in February of  
21 this year. Which was to support the dynamic  
22 preliminary system validation. And for that project,  
23 or that part of the project, we had people from the  
24 nuclear regulatory commission, from Constellation,  
25 from Westinghouse, other vendors and suppliers, as

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 well as the Department of Energy, participate in those  
2 studies. The results of those workshops were used as  
3 part of the license amendment request, which was  
4 supplied by Constellation to the NRC.

5 So this is one area where we collaborate  
6 with the commercial nuclear power industry and the  
7 NRC, and vendors and suppliers to address some of the  
8 long-term instrumentation and modernization needs of  
9 the industry. And I want to just provide that as an  
10 example so you understand some of the ways that we  
11 work within industry.

12 Specifically with respect to artificial  
13 intelligence and machine learning, the topic of  
14 today's meeting, we've been working with AI and ML  
15 technologies for about the past four or five years.  
16 And someone said to me recently, and I think it's  
17 true, that artificial intelligence is like the new  
18 math. We find it more and more within a lot of our  
19 projects. And I'll try to characterize and summarize  
20 that, but Craig and Ahmad will go in more detail.

21 We believe, well, these are relatively new  
22 to the nuclear power industry. And similar to the  
23 observations from the NRC Staff this morning that we  
24 have, from participating in IAEA and other  
25 international meetings, I do believe that the U.S. has

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 leadership with respect to some of the initial efforts  
2 investigating AI and ML with nuclear power plant types  
3 of applications.

4 They show promise, especially for  
5 automating manually performed activities. Many of the  
6 things that we do at nuclear power plants today are  
7 very labor intensive. You'll hear about some of those  
8 in our discussions. But we also see them as a way to  
9 enhance monitoring.

10 So they look promising to us as a way to  
11 enhance efficiency. But I want to also advise that  
12 what you're going to hear from us today really reflect  
13 R&D efforts. So when we show, for example, an  
14 activity where we're collecting data or conducting a  
15 test or something like that, at or with an operating  
16 nuclear power plant, that's not an actual deployment.

17 The same thing is true, speaking on behalf  
18 of licensees and vendors. We're not doing that today,  
19 we're really talking about our own R&D efforts, which  
20 may in fact be collaborative. But they really are  
21 focusing on three things. Reducing O&M costs,  
22 enhancing efficiency of the workforce, as well as  
23 improving situational awareness.

24 Moving on to the second point on here, our  
25 efforts, as I mentioned, emphasize work processes

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 versus controls. So I know that some of the members  
2 of the ACRS have a background in instrumentation and  
3 controls, including control processes. We're not  
4 investigating control activities, like so, we're not  
5 looking in deploying AI to control systems so much as  
6 just automating work activities that are labor  
7 intensive.

8 It's very important. And we're taking a  
9 very deliberate approach in our efforts with the  
10 vendors and suppliers and operating nuclear power  
11 plants to ensure that AI aligns with the nuclear  
12 safety culture. Just like we have with every other  
13 part of our R&D activities. As well, we are  
14 reflecting on and are trying to comply all of our  
15 efforts with presidential directives and other  
16 directives on AI that have been issued since 2019.  
17 More recently by the President.

18 Ultimately we think that AI will enhance  
19 worker performance at nuclear power plants. And I  
20 want to really emphasize that. We don't see AI as a  
21 means to replace people, but a way to enhance  
22 performance and help people do what they're best at.

23 And that's a reason why we have a strong  
24 emphases and focus in our research on human factors  
25 issues. We think it's absolutely vital to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 implementation and to achieve that.

2 It's important that workers trust  
3 automation, that they understand automation so that  
4 there is transparency, understandability. And  
5 ultimately facilitate usability. And hopefully that  
6 will come through in some of the remarks from Craig  
7 and Ahmad today as well too.

8 Now, I would be remiss if I didn't talk  
9 about the other areas of the light water reactor  
10 sustainability program, so the remainder of my  
11 presentation will be on the other activities that  
12 we're dealing not so much on AI. So I don't know if  
13 you have any questions so far, or if you'd like for me  
14 to continue with the rest of the presentation. But  
15 I'm open to any questions any time.

16 CHAIR BIER: Do we have questions so far  
17 or do people want to finish up first?

18 MEMBER KIRCHNER: Bruce, this is Walt  
19 Kirchner. Just quickly, you emphasized that it's in  
20 the R&D phase now, but do you have a few collaborative  
21 ventures where you're actually going to take it out of  
22 the lab so to speak and into a power plant and look  
23 for opportunities to harness this to either enhance  
24 productivity or enhance monitoring or --

25 MR. HALLBERT: Yes. So I would say, and

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 Craig will show this, but I think the answer to your  
2 question really is in Craig's examples. Most of our  
3 R&D is at and with operating nuclear power plants. Or  
4 with their data.

5 Now, those are not deployments. And I  
6 want to emphasize that. Those are not implementations  
7 but they're examples of how we want to ensure that our  
8 R&D activities could be used or could be transferred  
9 to the private sector as part of the technology  
10 transfer efforts. And that they do scale to real  
11 problems at real nuclear power plants. So hopefully  
12 you'll see that. But yes, that is, most of our  
13 research is out of the lab in many ways.

14 MEMBER KIRCHNER: Thank you.

15 MR. HALLBERT: Okay, I'll continue on.

16 MEMBER BROWN: Hey --

17 MR. HALLBERT: Oh --

18 MEMBER BROWN: This is Charlie Brown.  
19 Yes. In our earlier discussions we, in our earlier  
20 meeting, before noon, we had considerable discussions.  
21 And you made the comment in this that you're focusing  
22 on how you would improve operator or man, eating up a  
23 lot of man hours, you know, stuff that takes a lot of  
24 time. Intensive stuff but not necessarily focusing on  
25 controls.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           And we had some discussion earlier today,  
2           or actually a lot of discussion earlier today, in  
3           terms of the, how we have to be careful about getting  
4           AI into reactor trip, safeguard systems, plant control  
5           systems, such that we have something modifying or  
6           making decisions for humans or what have you when it  
7           really doesn't add value. Stuff you're talking about  
8           that seems to add value in terms of how you manage the  
9           plant in its operations, but when you want to trip the  
10          reactor you don't have to make a whole lot of  
11          decision. Your power is either too high or it's not.  
12          Or you've either lost pumps or it's not.

13                 It's not a, what I would call a real  
14          machine learning or other deep thought process to  
15          determine what you want to do in trying to embed this  
16          new idea into those systems could be detrimental to  
17          their ability to process it. Is that involved in your  
18          all's discussions in terms of how you, you know, rice  
19          bowl offer, you know, put a bar around certain areas  
20          that it really is not going to add value.

21                         I'm in favor of the real added value stuff  
22          not --

23                         MR. HALLBERT: Yes.

24                         MEMBER BROWN: -- just doing it where it  
25          seems like a nice thing to do because everybody else

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 is.

2 MR. HALLBERT: Yes. Yes, those are great  
3 questions. And also good comments as well, Charlie.

4 One of the things that we, we do look at  
5 as a part of our approach into where we might look at  
6 a project to investigate an AI application is based  
7 upon a business approach. So we often times have a  
8 business case for, this is a very labor intensive  
9 activity. A lot of people are involved in doing it.  
10 It's not high value added from the perspective of the  
11 utility and they wonder if there is a way to automate  
12 some of this through analytics AI and machine  
13 learning. So I think you'll see some examples of  
14 that.

15 We're not focusing on anything that's  
16 inside the control room especially. And we're not  
17 approaching anything, we're not even looking at  
18 minimum inventory, we're not looking at Class 1A  
19 systems.

20 That's all outside of the scope of what  
21 we're investigating today. We're looking at, what are  
22 some of the ways that we can help plants to be more  
23 efficient in terms of those vary labor centric types  
24 of activities, but also provide information that's of  
25 value to the people who are responsible for those

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 functions at the plant. Hopefully that answers your  
2 question.

3 MEMBER BROWN: No, that's a good answer.  
4 I, it's not necessary, it doesn't have to be a good  
5 answer. It doesn't, my opinion the right answer. It  
6 looks like you're all going down the thoughtful path  
7 that we did spend considerable amount of discussing  
8 earlier in the day.

9 MR. HALLBERT: Yes.

10 MEMBER BROWN: So thank you for --

11 MR. HALLBERT: Well --

12 MEMBER BROWN: Thank you for your response  
13 there.

14 MR. HALLBERT: Of course. And we'd  
15 appreciate feedback. That's one thing that we are  
16 always looking for is feedback on our approach and  
17 projects. And we'll be providing with links and lots  
18 of reports as well too.

19 In the interest of time I'm going to jump  
20 through some of the rest of the slides so that Craig  
21 and Ahmad actually can have the time that they deserve  
22 to go into detail. You all have heard about probably,  
23 unless there is some more questions right now.

24 You all have heard probably about some of  
25 the activities related to hydrogen demonstration

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 projects. A lot of that was initiated, and the  
2 foundational research was sponsored by the Light Water  
3 Reactor Sustainability Program.

4 We've been conducting research into  
5 potential uses of operating a nuclear power plants to  
6 produce hydrogen, extracting thermal energy, as well  
7 as just, you know, providing electricity for  
8 electrolysis systems, modifications of electricity  
9 transmissions. As well as doing studies dynamically  
10 with operators in the human system simulation  
11 laboratory with mockups in a simulated environment in  
12 operating nuclear power plant that includes something  
13 like high temperature electrolysis in the balance of  
14 plant, looking how operators would work with the  
15 double demands of electricity generation and hydrogen  
16 production. We've also been working on the economics  
17 of this.

18 I'm going to have to jump through my  
19 slides to stay on time, but I want to emphasize that  
20 the LWRS program, and other DOE offices, have been  
21 supporting these hydrogen demonstration projects. The  
22 first one, Nine Mile Point, is in operation. And it's  
23 using one and a quarter megawatt electrolysis, low  
24 temperature electrolysis unit only.

25 Davis-Besse and Prairie Island are set to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 go into operations sometime next year. And those will  
2 be also, one will be a low temperature electrolysis  
3 plant, and one will be a high temperature electrolysis  
4 plant which would be taking thermal energy from the  
5 plant to run a 150 kilowatt high temperature  
6 electrolysis unit. So it's a very small electrolysis  
7 unit but it's demonstrated the means for an off take  
8 of thermal energy from the plant.

9 Let's see. In terms of, I'd be a little  
10 remise if I didn't mention also that the LWRs research  
11 has been instrumental in supporting some of the  
12 hydrogen hubs.

13 The President announced a few weeks ago in  
14 Pennsylvania that there had been some hydrogen hubs  
15 selected and awarded to initiate R&D into nuclear  
16 power plants. Not just nuclear power plants, but  
17 broader hydrogen hubs. But some of the involve  
18 nuclear power plants producing a hydrogen at scale as  
19 part of the hydrogen hubs. And we've been supplying  
20 some of the information that we think enable some of  
21 those, some of those efforts moving forward. And some  
22 of the INL staff is also participating directly in  
23 supporting those hubs.

24 I know I'm jumping through the slides a  
25 bit here but I want to emphasize also that since its

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 inception the LWRS program has been conducting  
2 research into the long-term performance of key  
3 materials for vital structure systems and components.  
4 In fact, when we initiated this program together with  
5 the Nuclear Regulatory Commission and the Electric  
6 Power Research Institute, I would say the largest  
7 emphasis was on materials performance. Specifically  
8 in some of the areas that you see on this slide here.

9           And the emphasis in our materials research  
10 is understand how materials perform and degrade in  
11 this in-service environment over long periods of time.  
12 By conducting research into mechanisms, degradation,  
13 modeling and simulation tools to be able to model and  
14 predict that, as well as to inform mitigation  
15 strategies.

16           Now we're also conducting research into  
17 risk-informed system analysis, which is research and  
18 development to enhance economic efficiencies by  
19 optimizing safety margins and minimizing  
20 uncertainties. It involves a lot of R&D in  
21 collaboration with Nuclear Regulatory Commission, as  
22 well as with vendors and suppliers.

23           And I chose one example from the Risk-  
24 informed Systems Analysis Research which is a project  
25 that's looking into optimizing nuclear fuel

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 utilization. As you know, fuel costs represent about  
2 20 percent of the annual operating expenses of a  
3 nuclear power plant. And we have a project that's  
4 using an AI optimization framework for designing  
5 reactor core configuration giving certain objectives  
6 and constraints.

7 The little simulation on the right side  
8 here that I hope you can see, shows our simulation  
9 running through a number of iterations on nuclear fuel  
10 movements and switching to optimize the amount of fuel  
11 that needs to be purchased during an outage, as well  
12 as hopefully in the long-term the amount of fuel that  
13 needs to be stored on the back-end of the process as  
14 well too. It's a multi-physics based R&D project that  
15 uses a generic algorithm as the AI method for  
16 optimizing core loading.

17 Finally we're conducting research in  
18 physical security. And this is a topic that was  
19 raised to us by the nuclear power industry where they  
20 really asked if they are opportunities for DOE to  
21 share and leverage some of its own capabilities and  
22 physical security and protection with the commercial  
23 nuclear power industry.

24 I won't go into much information on that,  
25 but we do have a vibrant engagement activity with the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 Nuclear Regulatory Commission and the Industry on  
2 this, looking into advance security technologies,  
3 risk-informed physical security and a number of event  
4 security sensors and delayed technologies.

5 So, I'm trying to keep us on time. I just  
6 want to summarize by saying, going back to the  
7 original purpose of the LWRs program. Now we know  
8 that the existing fleet operating today provides the  
9 largest reliable source of carbon-free electricity in  
10 the U.S.

11 Some of the industry initiatives, like  
12 those that have been led by the Nuclear Energy  
13 Institute, DMP, have achieved substantial improvements  
14 and performance already. Energy, our nuclear energy  
15 supports, our climate goals can also contribute to  
16 deep decarbonization in other industries by providing  
17 energy for products that help to reduce the carbon  
18 footprint in some other industrial sectors.

19 A lot of the R&D activities that you'll  
20 hear about today from our program involve  
21 collaborations with Industry because we want to  
22 facilitate progress in areas of vital common  
23 interests. And by working together we can facilitated  
24 that kind of progress. Especially with some of the  
25 first movers in the industry who are interested in

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 moving forward in some area but maybe themselves are  
2 not R&D organizations. So partnering with DOE makes  
3 a lot of sense. We also have a lot of the risk for  
4 some of the early R&D approaches.

5 So, I would also say that our research is  
6 based on the highest priorities that we identified in  
7 the commercial industry. And they're conducted on  
8 timelines that support continued operation of the  
9 existing fleet. And I'll be happy to answer any other  
10 questions before we turn it over to Craig.

11 CHAIR BIER: Do people have other  
12 questions for Bruce?

13 MR. HALLBERT: Thank you very much for  
14 your time and the opportunity to talk to you today  
15 about our research.

16 MR. PRIMER: First, before I get into this  
17 slide, I just want to thank the NRC's engagement.  
18 There was mention of an MOU and the ability to share  
19 ideas under that agreement and provide technical  
20 information. And that has allowed us to identify what  
21 we think are meaningful research areas (audio  
22 interference) --

23 CHAIR BIER: Thank you for letting me  
24 know. Thank you. We should be okay now.

25 (Simultaneously speaking.)

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1                   MEMBER REMPE: -- that is muting. It's my  
2 microphone because I'm the person, Joy Rempe, who is  
3 logged into the meeting, but it is the room microphone  
4 and please do not mute it because we're constantly  
5 having to unmute it. Thank you.

6                   MR. PRIMER: Okay. And thanks. So what  
7 I was mentioning is, so we have some objectives and  
8 missions that Department of Energy have established  
9 for the program. And Bruce was able to share that  
10 with you.

11                   Some other things that I'd like to mention  
12 is, I heard discussion on what's AI versus machine  
13 learning. And so, this is a general survey of ideas.  
14 This is something AWS, Amazon Web Services, developed.  
15 It's very similar to what you'll see as you're, some  
16 of the other products.

17                   You have artificial intelligence, which is  
18 the ability for machines to take different types of  
19 inputs, make decisions using the Turing test. You  
20 wouldn't know if it's a person or a machine, it just  
21 does it. And it's able to process information and  
22 make decisions.

23                   The building blocks for that type of  
24 intelligence is machine learning and deep learning.  
25 Those are specialized bits of logic that is used to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 interrogate large data sets and infer or create some  
2 ideas around that information.

3 So machine learning is made up of several  
4 different types. You'll hear the mention of natural  
5 language processing, computer vision, time series.  
6 These are all different types of data sets that are  
7 available for algorithms to interrogate and determine  
8 things.

9 Next is deep learning. So deep learning  
10 is something that takes different types of machine  
11 learning algorithms and puts them together in unique  
12 ways to solve some difficult issues. So you might  
13 have natural language processing and computer vision  
14 working together. And you'll see, and the mode will  
15 show some examples of deep learning in (audio  
16 interference) drive into that quite a bit. But the  
17 deep learning is just a more complex set of algorithms  
18 that are used to solve a problem.

19 And then we also talked about generative  
20 AI. So that, again, is something that's out there  
21 where AI is taking information and creating new things  
22 from it. Whether it's art or reports or things like  
23 that.

24 So we're working in the machine learning  
25 and deep learning areas. That's what we're focused

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 on, that's what our research supports. We see  
2 opportunities to work with Industry to develop  
3 solutions or algorithms and then demonstrate the  
4 usefulness. And then from the collaboration, develop  
5 some type of reports and information that others,  
6 Industry or vendors, can use to develop products.

7 And ultimately make the decision for this  
8 specific station on how to use that. And where to use  
9 that. It's the research organization, we're not part  
10 of that decision making.

11 So, just to highlight what we'll be  
12 talking about in the next several slides. We'll be  
13 showing examples of machine learning for material  
14 management or equipment monitoring and anomaly  
15 detection, as well as applications, examples of  
16 natural language process and computer vision in  
17 applications within the plant.

18 Then last on the bottom there is the AI  
19 and ML explainability. So one of the key elements of  
20 deep learning, and some of those more complex  
21 algorithms is, as it becomes more complex it's more  
22 difficult to understand what it's doing and why the  
23 results are correct. So developing a balance between  
24 complexity and explainability is important, so we're  
25 doing some work there.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 I'll also mention before I go into the  
2 specifics, our group is made up of a group of human  
3 factors, scientists and engineers, data scientists,  
4 control system engineers, and finally a business group  
5 that's looking at balancing. And Bruce mentioned the  
6 motivation.

7 So why would we put research effort into  
8 this versus another opportunity. So there's a  
9 business case element that's considered into what  
10 we're working on.

11 So I guess I'll just stop there. Any  
12 questions on that general idea of these links of  
13 artificial intelligence to machine learning and deep  
14 learning? Okay.

15 So the first example of machine learning  
16 is in the passive system, or material management or  
17 material monitoring. You have examples of machine  
18 learning in looking at defects in concrete, defects in  
19 piping.

20 What you see here on the left is an  
21 example of collaboration between EPRI, Southwest  
22 Research Institute, several universities that created  
23 slabs of concrete with known defects and curated that  
24 as kind of a data set of sorts but a physical model  
25 with known defects that we could apply machine

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 learning technology to, to see if we could identify  
2 where those defects were.

3 Similarly on the right we had a effort  
4 with, again, EPRI and Southwest Research Institution  
5 on piping degradation where we coupled university  
6 work, University of Pittsburgh in this case, with the  
7 research scientist to develop digital twins, what  
8 piping should look like, compare it to the sensor data  
9 and try to inform the operators, and what would likely  
10 be maintenance teams, of where there might be defects  
11 within the piping.

12 And so we have those reports that are  
13 available and have been used on follow on projects to  
14 develop solutions that industry can use to help them  
15 in their different passive monitoring programs.

16 Next, moving to equipment monitoring. So  
17 from passive equipment or passive components to active  
18 components. We have an example of work that we did  
19 with Industry creating a digital twin of circ water  
20 system. That's the bottom left portion of that, so  
21 that's different components within the circ water  
22 pump.

23 We worked with vendors to create the  
24 software that created the, or developed the digital  
25 twin for that. And then also worked with the utility

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 to enhance the data that was available to monitor that  
2 equipment. And through that effort we were able to  
3 identify an approach that would monitor the circ water  
4 system performance and provide input to the system  
5 engineering and maintenance teams.

6 Ultimately the goal here is to develop the  
7 basis for transitioning from predictive, or from  
8 periodic maintenance to predictive maintenance where  
9 we can identify the likelihood of a component failure  
10 and advance warnings so that the component can be  
11 taken out for maintenance in what would be a planned  
12 appropriate time frame.

13 The lifecycle there that's indicated on  
14 the top right shows that data analytics is the  
15 beginning of that effort. We take the data, run it  
16 through different types of fault signatures and  
17 identify when we're starting to see some type of  
18 faults.

19 And once we do that we inform the decision  
20 to do that maintenance using different types of  
21 predicative modeling and risk modeling to determine  
22 the likelihood of some kind of failure before we could  
23 get to the maintenance or if it would seem to be the  
24 fault growing at a rate that may cause a problem  
25 before the plant maintenance. So this modeling then

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 allows you to make the decision on when you want to  
2 take the equipment out of service. And that would be  
3 what we call condition based maintenance or predictive  
4 maintenance.

5 MEMBER HALNON: Craig, this is Greg. The  
6 artificial intelligence box there, the first blue one,  
7 that's actually developing an ongoing model of the  
8 instrument response and comparing it to expected or  
9 how does it, I'm trying to find out, where is the  
10 human displaced in this because we've been doing this  
11 for years, obviously, for --

12 MR. PRIMER: So, what this would do is  
13 this would complement your system engineers and this  
14 would allow them to look at fault growth and then try  
15 and determine remaining useful life on a component and  
16 then plan the maintenance activities. And to that  
17 point of the artificial intelligence on the second,  
18 the predictive modeling.

19 So again, using artificial intelligence is  
20 kind of a bracket for machine learning and that sort  
21 of thing. That's what that's describing, is this is  
22 where you apply different types of predictive models  
23 and try and determine when the likely remaining useful  
24 life is projecting the component to fail.

25 MEMBER HALNON: And is that using

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 operating history or operating experience outside of  
2 this plant itself specific to the instruments, or  
3 whatever you're monitoring?

4 MR. PRIMER: So, and that's one of the, so  
5 one of the points that was discussed earlier is that  
6 data availability, data quality, data completeness.  
7 So all of these things are part of what you can  
8 evaluate and determine the quality of your models.  
9 And so there's that confidence level of the data based  
10 on historic, the amount of historic data that's  
11 available to the system engineers and to the model.

12 And so that's a, I think an important  
13 point, is how do you establish what that minimum  
14 quality is, how do you identify in a way that people  
15 can consistently apply that approach to different  
16 models to get some type of information on its  
17 usefulness.

18 MEMBER HALNON: Okay. But it is using  
19 outside information to some extent? I mean, expected  
20 you want to get as much information as you can --

21 MR. PRIMER: Right.

22 MEMBER HALNON: -- it's a matter of  
23 whether or not it's valid to that specific piece of  
24 equipment or whatever your launching.

25 MR. PRIMER: There is a large group of

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 data that's available. So you have historic data, you  
2 have real-time data that it's comparing signals  
3 against what we know are those fault signatures. So  
4 it's looking for that type of situation.

5 The other thing that it would look at is  
6 the other types of components that are similar but not  
7 necessarily that component --

8 MEMBER HALNON: Okay.

9 MR. PRIMER: -- so they have different  
10 RULs.

11 MEMBER HALNON: Is there any discussion  
12 about a real-time data using multiple plans, multiple  
13 fleets talking about, you know, rather than just use  
14 your own plant, which may be one or two instruments,  
15 we can use 25 or 30?

16 MR. PRIMER: Right. So that's something  
17 maybe you might call federated learning. And that's  
18 using data sets that are close but not specific to  
19 that component. And so there's an approach that we've  
20 worked through and have reports on and we think it's  
21 useful. It moves you much further in the learning and  
22 training lifecycle if you do it that way then starting  
23 with the small set of data and trying to infer  
24 information, so.

25 MEMBER HALNON: Okay. Thanks.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 CHAIR BIER: Can you talk a little bit  
2 about how the machine learning interacts more with the  
3 physic space model? Like, are you just estimating  
4 parameter values in the physics model or is it  
5 something more complex than that?

6 MR. PRIMER: Well, without a specific  
7 example I'll just generalize. So depending on the  
8 amount of information and where that, you know, what  
9 sensor information is available you're able to infer  
10 the status of that component.

11 So what we see is a lot of just, first  
12 order physics models that are used around a component  
13 to inform how the component is doing. You have other  
14 information like current, you know, I'm looking at a  
15 circ water system, so stator current and things like  
16 that are information that's used. So different points  
17 are used.

18 And in fact, maybe the next slide might  
19 help with this. We're getting into a little bit of  
20 the explainability here.

21 And so what you'll see here is, as part of  
22 that circ water system effort, and I should have put  
23 it in the original design of the output of the  
24 information was something the data scientists really  
25 loved. And it made complete sense to them. And they

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 would say, what else do you need to something that was  
2 a little bit more information available to the system  
3 engineers. So the data is telling me something, but  
4 what does it mean. And there is these different  
5 elements.

6 So on the right side you'll see the  
7 dashboard that was developed as a result of the  
8 physics model. So the physics model is telling us  
9 things.

10 And you can use explainability matrix that  
11 are like the LIME and Shapley that tell you what's  
12 your feature that's causing the algorithm to say that  
13 you have a fault. And this, if you see in these red  
14 bars across the bottom, they're telling you right now  
15 that the high temperature to the, is above, and I  
16 can't read that. It's the motor temperature is above  
17 a certain temperature. And that's your number one  
18 feature that's telling you that you have a problem.  
19 And so you have these specific, very measurable sensor  
20 inputs that tell you something.

21 But to complement that there is other  
22 information that a human would want to know. And so  
23 to the right of those bars are things that a system  
24 engineer would likely want to know.

25 So in the bottom right is seasonal

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 temperature. So summer time or is it winter time.  
2 It's high because it's hot outside or is this high  
3 because I have a problem with the motor.

4 And then the one above that is the  
5 prediction of what's likely to happen based on the  
6 fault progression. So what that's looking at is,  
7 based on the algorithm it's telling me that within the  
8 next 24 hours you're likely to see a min or max, two  
9 standards. You know, one way or the other of that  
10 temperature change. So they can look at say, well, if  
11 we do nothing we'll likely be okay. Or if we do  
12 nothing there may be a problem with us succeeding a  
13 limit. So that's pretty useful.

14 CHAIR BIER: Yes. For comparison with the  
15 presentation just before lunch, is this also  
16 considered unsupervised learning or is the fact of  
17 having the physics space model and all these  
18 constraints make it supervised?

19 MR. PRIMER: So --

20 CHAIR BIER: Look at you, you're --

21 MR. PRIMER: Yes, you're, this is perfect.  
22 So the next slide, so examples of unsupervised versus  
23 supervised, great presentation this morning on that.  
24 I think there is benefits to both. I think they all  
25 complement each other and give you different types of

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 information.

2 I think unsupervised methods give you  
3 insight that you may not be thinking about. And it  
4 just tells you what it sees, and that's very  
5 interesting.

6 I think semi-supervised on what, in the  
7 nuclear industry, even though there is a lot of data,  
8 it's small data sets. So having supervise or semi-  
9 supervise it tags or fingerprints certain indications  
10 that tell it that this is a fault or not a fault,  
11 helps it identify conditions.

12 And so we've run through and identified  
13 several different approaches that can be useful in  
14 both methods. And so I wouldn't say, I'd say they  
15 both are good, but they both need to be, you need to  
16 understand how to use them and how they complement  
17 each other.

18 And then now back to the physics space.  
19 So the physics is the way you compare things to  
20 actual. So you have a physics model. You look at  
21 what you think it should be and then you look at what  
22 you actually got. And if there is a delta then you  
23 start looking into why.

24 And sometimes these algorithms, like on  
25 this, with this little orange thing here says, we

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 think it's water box valve. That's what we think that  
2 is. We've seen this is in the past, we think it's  
3 water box valve. We being the AI thinks. And then  
4 the system engineer will go, I agree or no there is  
5 something that this didn't take into account that I  
6 have looked into and so we'll just disposition this as  
7 a nothing, no action needed.

8 Now back to the unsupervised and  
9 supervised detection. And I made the point earlier,  
10 so data wrangling and data quality is huge. So if you  
11 don't have good data or you don't have complete data  
12 sets that let you understand what's happening to the  
13 component or systems, it's hard to understand what  
14 that means. So the data quality and data completeness  
15 are two areas that we've worked on as well.

16 Then in the bottom right you have a system  
17 that allows us to actually cluster information  
18 compared to the history of component failures. And on  
19 the top right what we've done is we've actually used  
20 what we would call virtual sensors or additional  
21 sensors that are outside of what the system normally  
22 uses. So in the case on the top right we were able to  
23 use ambient temperatures, historical temperatures  
24 around that component that wasn't part of the system  
25 sensing. So it will combined information to one

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 algorithm that wasn't available to operators.

2 I'm going to jump to the next topic unless  
3 there is questions on these points. Okay.

4 So now moving to some of the applications.  
5 I mentioned the idea of machine learning having  
6 different types of application, natural language  
7 processing, time series, computer vision, and how do  
8 we use those in ways that help us improve efficiency  
9 and reliability on the business side. That's really  
10 our balance in the research area is to leverage  
11 technology.

12 Not necessarily in the short-term to  
13 invent new technologies, but really validate the  
14 technologies that are available now for use in the  
15 nuclear industry. And that's part of our research  
16 focus.

17 And you can see, we probably talked about  
18 the EPRI report that was produced. That helped with  
19 the natural language process. And that was actually  
20 part of the output of research that started at the  
21 national lab.

22 And partnered with EPRI and Industry to  
23 develop open source software code that was used. I  
24 think Jensen Hughes picked it up and they're using it  
25 at Constellation for the corrective action screening.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 We've got Xcel Energy, I think I saw them on the line,  
2 using some of the software that's been developed and  
3 the collaboration with them as well.

4 So that natural language processing and  
5 deep learning methods, like I mentioned, will go into  
6 what that looks like with the mode. But it's a very  
7 useful approach to analyze data and make some, infer  
8 some information from it.

9 The, similarly, looking at the warehouse  
10 stocks and looking at part failures and understanding  
11 what pieces of parts were actually used, what's likely  
12 to be needed in the future based on, you know, time  
13 frames of the year, time frames of maintenance cycles,  
14 whether that's months or years, it helps improve the  
15 warehouse efficiency to make sure they have the right  
16 parts on hand.

17 All of these have the report numbers  
18 labeled there and we're able to share this  
19 presentation. Their live links. And you can lick to  
20 the OSTI, which is the DOE's library with these  
21 reports. Or you can just, if you don't want to go  
22 live link, you can get a PDF and type that in.

23 Lastly, I think this is lastly, the  
24 computer vision. And Ahmad is going to jump in, very  
25 deeply, on an example of the use of computer vision

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 for fire watch and what we've done there.

2 Some other things that are useful for  
3 computer vision is gauge reading. So we can send out  
4 people to take pictures. And you can detect changes  
5 from one picture to another.

6 So you could possibly use that for system  
7 turnovers where you have a picture of something and  
8 then the next person that comes in can look at the  
9 change from when they were there last and say, oh,  
10 these are the three things that have changed. This  
11 computer vision will detect that, highlight that and  
12 help them assist in the turn over to make sure that  
13 it's identified and discussed.

14 Similarly we can use QR codes to help  
15 align drones to go to the right places and capture  
16 information.

17 And lastly, I'm going to kind of leave  
18 this to Ahmad to jump deeply into the use of computer  
19 vision for fire watch. And he'll be going through the  
20 rest of the, I guess hour that we have left, on that  
21 topic.

22 CHAIR BIER: Questions for Craig before we  
23 transition?

24 MR. PRIMER: And I'll mention, this last  
25 slide is just the same slide that Bruce presented in

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 our guidelines and mission from the DOE and artificial  
2 intelligence research we're developing and  
3 demonstrating applications that can be used in non-  
4 safety, non-control type of things right now.  
5 Anything that would move beyond that isn't research  
6 that's underway right now.

7 CHAIR BIER: Okay. And is Ahmad going to  
8 use the same slide deck?

9 MR. PRIMER: I think Ahmad is going to  
10 take control. Ahmad --

11 DR. AL RASHDAN: Yes.

12 (Simultaneously speaking.)

13 CHAIR BIER: You can share your slides.

14 MR. PRIMER: That's great.

15 CHAIR BIER: Okay.

16 DR. AL RASHDAN: Yes, I'm going to go  
17 ahead and --

18 CHAIR BIER: Excellent.

19 DR. AL RASHDAN: -- share my slides.

20 CHAIR BIER: Thank you. You may need to  
21 stop sharing, Craig.

22 MR. PRIMER: I can do that. Okay.

23 CHAIR BIER: Okay.

24 MR. PRIMER: Stop screen sharing. Let's  
25 see if this works. I see your screen now, Ahmad.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 DR. AL RASHDAN: Okay, great. And you can  
2 hear me, right?

3 MR. PRIMER: Yes.

4 DR. AL RASHDAN: Okay, great. So thank  
5 you so much, first of all, for the introduction, for  
6 the invitation to present today. And, Craig, thank  
7 you for your presentation and the segue into this  
8 talk.

9 So my name is Ahmad Rashdan. I am Senior  
10 R&D Scientist at Idaho National Laboratory and I lead  
11 multiple efforts under the light water reactor  
12 sustainability program, so I work very closely with  
13 Craig and Bruce.

14 Now, the aim of this presentation is to  
15 dig a bit deeper into one specific application of AI,  
16 which is fire watch. And in this case it's computer  
17 version application of AI.

18 So I will start with defining what is a  
19 fire watch. And this is for people online that might  
20 not be familiar with what a fire watch is. So fire  
21 watch is an activity in which a person is assigned to  
22 monitor for fire and report it as soon as it happens.  
23 And in some cases even mitigate it.

24 And there are two different scenarios in  
25 which a fire watch would be needed. The first one

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 would be if you have activity in the plant that has an  
2 abnormal risk of fire. And a good example would be  
3 welding or flame cutting. So you might need, in this  
4 case, a person watching for the fire and taking action  
5 in case a fire is started.

6 The second scenario is, if your fire  
7 protection system is down because it's going through  
8 a testing process or it's under maintenance. And  
9 sometimes because it failed. And in this case we  
10 would allocate some fire watch personnel in various  
11 location in the plant to compensate for the lack of  
12 the fire protection system until it's brought back up.

13 By the way, I can't see if someone raises  
14 their hands, so feel free to interrupt me at any point  
15 of time.

16 Okay. So the motivation behind this is  
17 mostly when we engage the utilities a while ago we  
18 were informed that in some nuclear power plants a fire  
19 watch can cost in excessive of \$1 million per month.  
20 Especially the ones that have issues with fire  
21 protection system. They might have fire watch  
22 allocated on daily basis in multiple locations in the  
23 plant.

24 So in 2019 the Utility Service Alliance,  
25 which is a consortium of multiple nuclear power

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 plants, I think they have like nine utility members  
2 and 13 plants, or sites, submitted a proposal to a  
3 grant called the Industry FOA. And that proposal had  
4 in its scope to research and develop automation and  
5 advanced remote monitoring technologies.

6 The aim was to improve the economics of  
7 various processes in the plant without compromising  
8 the safety. That proposal was awarded in 2019, and  
9 it's still ongoing up to now, so we're in the final  
10 year of that award. And the scope of the proposal  
11 included the fire watch process. And means to  
12 introduce automation into that process.

13 So our specific --

14 CHAIR BIER: A couple --

15 (Simultaneously speaking.)

16 CHAIR BIER: Excuse me.

17 DR. AL RASHDAN: Sorry.

18 CHAIR BIER: A couple of very --

19 DR. AL RASHDAN: Sure.

20 CHAIR BIER: -- quick questions. Are you  
21 working with a university partner on this or it's  
22 pretty much being done at INL?

23 DR. AL RASHDAN: We are, we are working  
24 with a university partner. Actually, the Utility  
25 Service Alliance is also working with a university.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 And specifically, the scope for the university is to  
2 custom built this fire cart, which I'll talk about in  
3 a second.

4 CHAIR BIER: Oh, okay. And the other  
5 question is, is it detecting, is it learning, is the  
6 AI learning to detect fire from video of real world  
7 fires or from simulated fires?

8 DR. AL RASHDAN: Real world fires. And  
9 I'll also talk about that, but thank you so much for,  
10 you're giving me a good segue to the next point, so  
11 I'm glad you mentioned that.

12 So as part of the scope of the utility  
13 service alliance they are, do research and develop and  
14 evaluate a custom made fire cart. This fire cart has  
15 a camera on it, but it also has other types of  
16 sensors. Like infrared sensors, smoke detectors.  
17 We're even looking at adding acoustic sensors.

18 And the idea here is that every one of  
19 those sensors has its own fire detector. And then we  
20 would fuse the sensors decisions to make a holistic  
21 decision if there is a fire or not.

22 The topic of this presentation today is  
23 one of those detectors, which is using a video stream,  
24 or image, stream through camera, an optical camera,  
25 and detecting the fire within that stream using

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 machinery. So we looked at three different options  
2 when it comes to the methodology to detect fire in an  
3 image or a video.

4 The first one is called imagine  
5 processing. And this is the classical way of doing  
6 this. In the traditional way of detecting any object,  
7 like fire in an image, you would need to engineer what  
8 we call the feature engine. Or feature extractor.

9 What that does is that in the case of fire  
10 you would decide what are the specific features of  
11 fire that you're interested in. For example, in a  
12 fire you might be looking for orange pixels that are  
13 adjacent by maybe red-ish pixels. That a feature you  
14 can engineer yourself and force it on the detector.  
15 And then you build a decision-making algorithm that  
16 detects all those features and decides, is this  
17 considered the fire, does this look like a fire or  
18 not.

19 The good thing about this approach is that  
20 you don't need data to train an algorithm. You design  
21 the feature extraction engines and you assume they are  
22 correct and you use them as is. The problem with this  
23 is that your model is, or your results are as good as  
24 your future engineering is. So if your future  
25 engineering is missing some feature of the object you

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 are trying to detect, your results are going to  
2 reflect that.

3 The second approach, which is special  
4 machine learning, is automating that process. So when  
5 we use neural networks or machine learning, we ask the  
6 machine to find what features are important. So we  
7 don't design the features we're looking for.

8 We basically load thousands and thousands  
9 of different features into our model, and we feed in  
10 a lot of data that duplicates fire and no fire. And  
11 we level this data and we tell the machine, this is  
12 how fire looks like, what are the features that are  
13 important to detect to catch fires.

14 So the benefit here is that we don't have  
15 to engineer this manually, so we get much more robust  
16 features. The disadvantage is that you need a lot of  
17 data. In this case, image data to train an algorithm  
18 on.

19 The last approach is spatial and temporal.  
20 So as the name implies, it's the same process as I  
21 just mentioned to you with the images, however, we add  
22 a different dimension in this case, which is the time  
23 aspect. So we're not only comparing features within  
24 an image, we're comparing features between one image  
25 or one frame and the next one. And that's the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 temporal aspect into it.

2 So that adds more features for the  
3 algorithm to detect. However, the challenge there  
4 was, we need now more video type of data. So we need  
5 temporal data to feed in to be able to train. And the  
6 challenge there is, we don't have as many video data  
7 sets as we do as with images. So we have much more  
8 sparsity there, and that impacts the algorithm  
9 significantly from a performance perspective.

10 So we ended up going with the middle  
11 approach. So spatial machine learning. And that's  
12 the scope of my presentation today.

13 So I'm going to talk to you about how we  
14 created those models. I'm going to start with a data  
15 creation, collection and preparation. And then from  
16 there the model architectures we considered. And I'll  
17 explain to you what I mean by that. And then I'll  
18 show you some performance results. And finally, some  
19 news considerations when we talk about using AI for an  
20 application like this one.

21 I'll stop here. Any questions before I  
22 proceed? Okay. Again, feel free to interrupt me at  
23 any point of time.

24 All right. So let's start with the data  
25 collection. So we looked at three different generic

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 sources, or general sources of data. There is the  
2 general images and video sources, like YouTube, Google  
3 Images, Yahoo has also a repository of images. Those  
4 are usually large, or extremely large, sets of images  
5 that are often leveled.

6 So what we did in this case, we had to  
7 manually sort through those different labels and find  
8 what labels out of these would actually represent  
9 fire. So I'm showing you here to the right two  
10 different boxes. And we've seen something similar to  
11 this in the morning, so I don't know if I need to  
12 explain it much.

13 But what this is telling you is that for  
14 fire labels those are kind of the themes of flavors we  
15 found in the data sets. So very often when we talk  
16 about vehicles in this context it was related to fire,  
17 however, vehicle was also used very often, actually  
18 much more often to be, to labeled non-fires, so that  
19 would not be a good label. Versus if we talk about  
20 fire engine usually those images labeled with fire  
21 engine would have a fire in them so we could label  
22 this as fire.

23 So you can imagine the process of trying  
24 to sort through those data sets to figure out, what is  
25 fire, what is not fire, is very time intensive and

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 demanding. So we also looked at targeted sources.

2 So we researched into other research that  
3 tried to build something similar to what we tried to  
4 do here. We find a data set that was complainant  
5 called FiSmo, for fire and smoke, in which people  
6 basically extracted some images and videos from  
7 sources like these and labeled them with fire or  
8 smoke. So they made things easier for us. However,  
9 the size of those data sets were not big enough for  
10 us, so we resorted to a different approach. And this  
11 is an interesting approach.

12 So we know there are out there, there are  
13 some models that can classify images. ImageNet, for  
14 example, is one of them. But a good example that you  
15 might be able to use on a daily basis is that if you  
16 open your phone and you go to your gallery on your  
17 phone you can actually search through your images for  
18 a certain object.

19 So you can type, at least I have this on  
20 my phone. It depends on what phone model you have.  
21 But I can type, if I'm looking through my images or my  
22 gallery for a cat, it will show me images of cats.  
23 However, in my case when I tried this capability is  
24 not very accurate. So you might end up with a lot of  
25 images that are not cat and it will miss some images

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 that are cats.

2 So my point here is, that there are models  
3 out there that has generic classes, like the one I'm  
4 showing you here. And out of those generic classes we  
5 can find ones that relate to fire. And those are the  
6 ones actually I'm showing. So candle, canon, fire  
7 screen. All those are generic classes in those models  
8 that we can, that can help us zoom in on some fire  
9 images that we can use.

10 So using those three different approaches  
11 we can capture, we can basically expand on our data  
12 sets. Because the key element in machine learning is  
13 your data. As Craig mentioned earlier.

14 The one aspect of the data that you have  
15 is that you have to ensure that your data is diverse  
16 enough for the application you're targeting. So in  
17 our case we have to look at all those different  
18 parameters environment, the detection of space. What  
19 kind of objects exist in those images, what is the  
20 light source, are they dark, are they very well lit.  
21 And then if it's a video, how long it is, is it for.  
22 And then also, spatial aspect of the resolution, like  
23 the resolution, and if the fire is actually steady or  
24 moving and so on.

25 So we went through our data and looked at

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 those parameters and made sure that those parameters  
2 are well represented. Now one thing we needed to be  
3 careful about is bias in our data.

4 So for example, if we see that most of our  
5 images has fire in the middle of the image, that would  
6 teach the machine to focus on the middle of the image  
7 when it's trained. So in that case we would need to  
8 make sure that the fire is not, is basically spatially  
9 distributed across our data set. And we'll have maybe  
10 to discard some images with fire in the middle to kind  
11 of balance our data sets because if you have a bias in  
12 your data you'll end up with a bias in your models.

13 All right. So now that we have identified  
14 the data we had to clean it up. Prepare it basically.  
15 And the first thing we had to do is, just like  
16 mentioned to you now, there are some biases in terms  
17 of the fire location, but there are also some bias  
18 associated with the different objects you have in the  
19 images.

20 As an example, I showed you earlier, and  
21 I can go back, that fire and fire fighter were one of,  
22 were two of the most common labels for example for  
23 fire. So what this means is that every time I see  
24 fire there is a good chance I'm going to see a fire  
25 fighter. The problem with this is that if I see a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 fire fighter I might label the video as, the image as  
2 fire even though there might not be fire because that  
3 correlation causes a bias. And the machine starts  
4 correlating firemen with fire even though firemen does  
5 not mean there is a fire always. So that's an example  
6 of bias in objects.

7 So what we have to do is we intentionally  
8 cropped some of the images or blurred parts of the  
9 images sometimes to, again, introduce that balance I  
10 was telling you about. That's kind of part of their  
11 data preparation. And as I said, that's very, very  
12 critical when you're trying to build a model to remove  
13 the bias. It's very critical.

14 And then we were talking about videos  
15 because in some cases, even though we're using images,  
16 we had some video data sets. And we realized that  
17 there is value in using those videos.

18 So we had to break them into scenes and  
19 make sure that the images that we're capturing from  
20 the video do not look a lot alike. Because then we're  
21 feeding basically the same information. So we had to  
22 implement some temporal separation. And that was done  
23 by incorporating like a scene detector that identifies  
24 that the scene actually changed before it can take a  
25 picture out from the video. Or an image out of the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 video.

2           And then what we did, we actually  
3 intentionally biased our data sets towards fire. So  
4 usually when you're trying to train you would have a  
5 same number of data sets to represent the class you're  
6 interested in, which is fire, and the same number of  
7 images for no fire.

8           In our case we wanted more fire images  
9 than no fire. Because we wanted to bias the model to  
10 actually detect fire. So if you're not sure, if the  
11 model is not sure, we wanted to detect fire. To flag  
12 it as fire. Because the fact that if you miss a fire  
13 that's much more severe than if you detect the false  
14 fire.

15           So we introduced that bias in our data  
16 sets. And I'm showing here the results. You can see  
17 that there is much more fire images than we have  
18 normal images.

19           And then what we do in machine learning we  
20 break our data sets into three subsets. One is called  
21 training, which is what the machine learns from and is  
22 taught based on. And then the machine uses a data set  
23 called validation, which is internally, as it's  
24 training, it benchmarks its performance against the  
25 validation data set. And then we break a part a third

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 data set called testing, or testing data sets. And  
2 that data set is usually independent from the training  
3 and validation. So it's data that the model has never  
4 seen. And that's the data we use to generate  
5 performance metrics and evaluate the performance of  
6 our model.

7 One thing I should mention here, we also  
8 did the same thing for smoke, even though it's not the  
9 scope of this presentation specifically. But as you  
10 can see there was much less smoke images that we found  
11 than we did for fire.

12 Any questions so far?

13 MEMBER HALNON: Yes, this is Greg. I got  
14 a couple.

15 DR. AL RASHDAN: Sure. Go ahead.

16 MEMBER HALNON: First of all, you just  
17 mentioned that the false, getting a false one is  
18 better than missing. And I guess for a little while  
19 that also builds in a level of complacency to the  
20 folks that are getting this data that eventually could  
21 cause just as bad of problems. So just a thought  
22 there.

23 But where's the, what kind of feedback  
24 loop is there so that if there is a fault that that  
25 model knew that was false, don't do that again. How

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 is that done?

2 DR. AL RASHDAN: So this is a great  
3 question, Greg. I'm going to come to this in a few  
4 minutes. I have a slide on this. Which is over here.

5 MEMBER HALNON: Okay. If you're coming to  
6 it --

7 DR. AL RASHDAN: But I'll come to that --

8 MEMBER HALNON: -- later that's fine, I'll  
9 wait.

10 DR. AL RASHDAN: Okay.

11 MEMBER HALNON: The other point I'll make,  
12 just since I got the mic is, you know, when you get  
13 through all this, obviously we're going to get success  
14 on being able to detect fire.

15 DR. AL RASHDAN: Correct. We will.

16 MEMBER HALNON: But that misses one of the  
17 most important part of fire prevention is the  
18 prevention piece. Which people do. Fire watch as  
19 people, we stop somebody from coming into an area that  
20 is vulnerable to a fire. They'll stop somebody before  
21 they even start welding. Or they even start a fire,  
22 you know, heat producing or spark producing work.

23 So I guess that's not required by  
24 regulations, so we're sticking simply to what a  
25 regulation is required here, but I think we're missing

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 another piece of this. So I'll be interested, maybe  
2 down the road, if you talk through that. You don't  
3 need to do it now, but maybe in our summary we can  
4 talk about this other aspect --

5 DR. AL RASHDAN: Great question.

6 MEMBER HALNON: -- of prevention.

7 DR. AL RASHDAN: Great question. So the  
8 aim of this presentation, and this research, is on the  
9 technology development and the evaluation of the  
10 technology up to the point where it gets to the  
11 deployment. And then once we have something ready for  
12 deployment, or close to ready, or evaluate it, it's on  
13 the, the plan here is on the utilities to take it from  
14 there and determine all the procedures, whether  
15 they're administrative or technical based, that need  
16 to be taken care of to make sure this is compliant  
17 with the policies they have in the plant.

18 But our role here is R&D and evaluation.  
19 And then we do the demonstration, and then the  
20 utilities would take it from there and kind of custom  
21 fit it in their process.

22 MEMBER HALNON: Okay.

23 DR. AL RASHDAN: So --

24 MEMBER HALNON: So we'll have to do some  
25 assessment on whether or not that prevent piece needs

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 to be compensated some other way or if it's even  
2 necessary.

3 DR. AL RASHDAN: Correct.

4 MEMBER HALNON: Got it. Okay, thanks.

5 CHAIR BIER: Okay. We have a very quick  
6 question from Staff, if that's okay?

7 DR. AL RASHDAN: Sure.

8 MS. GOETZ: Hey, Ahmad, I'm Sue Goetz.  
9 I'm a project manager here at the Nuclear Regulatory  
10 Commission. I actually manage one of the reactors at  
11 Constellation, but let's not name them. I don't want  
12 to invite scrutiny.

13 A couple of weeks ago we had a reactor  
14 trip, and then when the licensee went in to  
15 investigate what we saw was a degraded wire where the  
16 sheet had completed melted and you can see the  
17 internals, right, and copper. And, you know, you  
18 could see it was pretty hot, it melted everything, but  
19 there was no fire.

20 So my question to you is, how would, you  
21 know, this smart thing here would have, how would that  
22 have handled, how would fire watch have handled that,  
23 would the reactor still have tripped, would we have  
24 tripped sooner or is there any application at all?

25 DR. AL RASHDAN: So, because we were

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 focused on fire watch explicitly we haven't looked at  
2 other scenarios. And I'm assuming that scenario did  
3 not have a fire watch, it's just a normal wire that is  
4 in an area that's not monitored for, with a fire  
5 watch, and the wire melted.

6           However, if you think about it, if you  
7 deploy something like this on a broader scale, and  
8 again, even though this wasn't the target, I'm just  
9 going out of my comfort zone here, we do have a smoke  
10 detector. If that event caused some smoke to be  
11 generated, that could have been flagged. If we do  
12 implement a smoke detector. However, if there was no  
13 fire, even though it melted and it deploys something  
14 like this, the fire detector would not catch it, the  
15 smoke detector might. Depending on the, how heavy or  
16 how thick the smoke was.

17           CHAIR BIER: So I have a few more  
18 questions. Oh yes. I have a few more questions about  
19 sort of adversarial testing of your results. I don't  
20 know if this is a good time to share them or if I  
21 should wait till later in the presentation?

22           DR. AL RASHDAN: If you want to ask it --

23           CHAIR BIER: Okay.

24           DR. AL RASHDAN: -- I can tell you if it's  
25 coming or not.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 CHAIR BIER: Yes. That's fine.

2 DR. AL RASHDAN: And if not I'll answer it  
3 now.

4 CHAIR BIER: So, in a kind of project like  
5 this there is obviously a big incentive to show that,  
6 yes, you successfully detect fires. But I think you  
7 also need to look at kind of reverse engineering, how  
8 fragile that is.

9 So what happens if you change three pixels  
10 in a photograph and turn them black? What happens if  
11 you change five or ten percent of the pixels and turn  
12 them black? A human would probably still recognize  
13 that as a fire, the AI may not.

14 What happens if you have a, say a gas fire  
15 that burns blue, it may be very scarce in your data  
16 set but that a human would still recognize as a fire?

17 And what about a watercolor painting of a  
18 fire? I mean, hyper-realistic fire hopefully it would  
19 recognize, but may not a sketch or, you know, it might  
20 recognize a sketch as a fire even though it shouldn't.

21 And part of what kind of encouraged me to  
22 think about that is, you know, in past presentations  
23 I've heard, I heard a talk by a psychologist about  
24 learning to recognize animals. Like, is it a cat or  
25 a dog, like in your picture. And it did a great job.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 They thought like, oh, this is doing excellent.

2 And it took a while before they figured  
3 out that it wasn't looking at the shape of the image  
4 it was looking at the texture. So if they had a cat  
5 with elephant textured skin it would call it an  
6 elephant.

7 And so I think you need to be a little  
8 creative. I don't know whether you've done that or  
9 not, but about challenging the results to see how easy  
10 it is to break what you build. So just a --

11 DR. AL RASHDAN: That is a great -- yes,  
12 go ahead.

13 MEMBER KIRCHNER: You know, in the thermal  
14 hydraulics business and building -- And you've got an  
15 expert here in Josh Kaiser on this topic in building  
16 things like CHF models and such.

17 You know, they do just what you did,  
18 Ahmad, you know, they will call out, or the applicant  
19 will call out a you've got the training set and then  
20 you've got a test set.

21 I think what you are suggesting Vicki is  
22 that the test set often has to be not from the same --

23 CHAIR BIER: Yes.

24 MEMBER KIRCHNER: Often the training set  
25 and the validation set comes from the same data, but

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 I'm thinking in a case like this where you are going  
2 beyond something as simplistic as thermal hydraulic  
3 thermal couple results to a very complicated challenge  
4 that you have taken on that the test data maybe has to  
5 come from something entirely different.

6 CHAIR BIER: Yes.

7 MEMBER KIRCHNER: Just to look for the  
8 elephant in the room, your metaphor.

9 CHAIR BIER: Thank you.

10 DR. AL RASHDAN: Very valid points. So to  
11 the first point I think that was mentioned, so, for  
12 example, there was a mention of the blue fire, that  
13 did come up and our data set is as diverse as we can  
14 identify all those scenarios.

15 I mean there is always -- We try to figure  
16 out all the different scenarios that should be in our  
17 dataset, however, if anyone tells you AI or any AI  
18 model can perform 100 percent perfectly and get you  
19 the right results on 100 percent of the time I think  
20 that would be a very inaccurate statement.

21 That's why as we started this process we  
22 started with a small dataset and we keep exploring  
23 every time we test this and find some deficiency. I  
24 will talk about the testing part and the validation  
25 part later on, which is the other point that you

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 mentioned, which is detecting fire for the wrong  
2 reason, for example.

3 But every time we saw that we realized  
4 that there was a deficiency in the model and we would  
5 go an look at the data and capture more data. Because  
6 as I will explain to you in a minute, in terms of the  
7 models we did a lot of research in that aspect and we  
8 have multiple models that are actually going together  
9 to get that fire decision.

10 Now with respect to the other question  
11 that was brought up, so you are right we do break the  
12 testing datasets apart. So the testing dataset is  
13 something that is taken out at the beginning before we  
14 do any training or validation of the models.

15 However, when we are designing the testing  
16 dataset we do keep in mind that we need to accommodate  
17 for all those variations, so we have to make sure our  
18 testing dataset is diverse enough to represent the  
19 sample that was used for the training.

20 One of the issues you see with AI  
21 sometimes is when your testing dataset is not properly  
22 designed you might get a validation result that is  
23 very promising, and then, when you actually test it,  
24 it doesn't perform very well.

25 You might even test it very well and then

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 we take I the actual environment and test it again and  
2 it doesn't perform as well. Usually it's because the  
3 data is not properly designed.

4 So if you make sure that there is enough  
5 diversity in your testing dataset you should get as  
6 close as possible to the real performance scenario.  
7 To elaborate more I will come to the explainability  
8 part later on which will answer the other part of the  
9 question about the misdetection and how to make sure,  
10 how do we improve those models, how do we go back and  
11 improve them, how do we know there is something wrong.

12 I will briefly touch on that in one of the  
13 slides going forward. Any other questions or  
14 comments?

15 MEMBER KIRCHNER: Well, Ahmad, this would  
16 prejudice what you're doing, but have you thought of  
17 constructing artificial fire models?

18 DR. AL RASHDAN: Well we have --

19 MEMBER KIRCHNER: Because that's to either  
20 populate your algorithm and teach it and/or to  
21 validate it?

22 DR. AL RASHDAN: So on a different effort,  
23 I think Craig mentioned this, we were working on gauge  
24 reading and in that effort we did create hypothetical  
25 gauges in addition to the real ones.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           The reason we did that is because we  
2           didn't have enough real data. In this case we did  
3           have much more than what we need. We were able, as I  
4           will show you later on, to get using real images of  
5           fire to get really good performance.

6           The conclusion I got from that effort,  
7           which would apply here, is that if you have a lot of  
8           hypothetical or synthetic data and you start using it  
9           in your training dataset those synthetic gauges are  
10          not going to look exactly like a real one and in real  
11          life you are not going to see those.

12          So what they are doing is actually they  
13          are confusing model. They are not necessarily  
14          improving performance. So it depends on how good you  
15          can make those look.

16          If they look really realistic then you  
17          might help the model, but if they don't it's almost  
18          like you are adding it on, it's now you're telling it  
19          that this is also, for example, fire and it's  
20          synthetic fire when in reality it's not going to see  
21          it, it's not going to look like that, so it makes  
22          things worse, not better, sometimes.

23          So in this case we didn't need it because  
24          we had a lot of data is the short answer. Any other  
25          questions?

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 (No response.)

2 DR. AL RASHDAN: This is a great  
3 discussion. Thank you. This makes it much easier for  
4 me to go with the presentation.

5 All right. So now we talked about the  
6 data, let's talk about the model. What I am showing  
7 you here is the layout of a type of neurometrics  
8 called convolutional neural networks.

9 I am going to call it CNN going forward.  
10 So CNNs are very common when it comes to any type of  
11 image detection or a classification or segmentation,  
12 so they are very common in the image processing field  
13 of AI.

14 The way they work is they are -- If you  
15 remember I told you, I was talking about features  
16 earlier and I mentioned to you that how in the past we  
17 used to design our own features and say let's look for  
18 orange pixels, that's our indicators, but now the  
19 machine does that.

20 The way it does it is that it actually  
21 generates hundreds or thousands of different filter,  
22 like the one I am showing you here, and then it runs  
23 through all of them and, based on the data that you  
24 give it, it will decide which out of those filters is  
25 the most useful for detection of fire.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           When I say a filter, if you look at this,  
2           this would be as an example a sliding window that  
3           would go over this image and it will have weights in  
4           every one of those boxes, so every time this goes on  
5           one part of the image it will mask nine pixels and it  
6           has weights and those weights are multiplied by the  
7           pixels and those weights are tunable, so they are not  
8           fixed, and the model actually tunes them when it is  
9           trying to figure out what features it is trying to  
10          catch.

11           They have different sizes of filters, so  
12          some of them are like this, some of them are bigger,  
13          and some of them are actually smaller. You can have  
14          maybe even a smaller than that.

15           Now one of the characteristics of those  
16          filters is that they are linear. Those are linear  
17          multipliers by the pixels. So we usually add  
18          something called an RELU, which that's the non-linear  
19          aspect, so this is a function that has a non-linear  
20          shape and that enables the models to become also non-  
21          linear, gives it much more dimensions and power.

22           As we do that we are getting all those  
23          features and what we are going to end up with, because  
24          we have sometimes hundreds of these, we are going to  
25          have to deduce the dimensionality of the problem and

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 we do something called max pooling.

2 A simple way I can describe it as just  
3 finding the max through a certain number of pixels and  
4 using that as a representative of the whole number and  
5 then we repeat the process.

6 So this is called one layer, for example.  
7 We repeat the process again and again and again. Then  
8 towards the end we have some sort of weights that  
9 coming from all those features, coming to this end,  
10 those were to start basically giving importance to  
11 certain features over others.

12 That usually has a -- Then you have a  
13 classifier. The way it works is when I am training  
14 this model I know that sometimes, for example, there  
15 is a CAT here and I am labeling the image as CAT and  
16 if it doesn't detect a CAT it will just start  
17 adjusting the weights and it has an optimization  
18 algorithm in it to try to get us to CAT and as you add  
19 more and more images of CATs and more and more images  
20 of non-CATs the weights are going to get adjusted  
21 further and further till the moment you have a model  
22 that is trained.

23 That's how in a nutshell how this whole  
24 thing works. Now we have looked at literature, there  
25 are nine models I am listing here that has this number

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 of layers. So you see how many of these.

2 So this has 815 layers. I am showing here  
3 only a couple. They have this many tunable parameters  
4 in this case, so very wide and deep models. That's we  
5 call this deep learning, because they is a lot of  
6 tunable parameters.

7 When we feed, again, those pictures we are  
8 basically tuning those parameters to get us the right  
9 weights to try to recognize what fires are. Now you  
10 see the capability why this is much more powerful than  
11 the way we used to do it in the past where we used to  
12 design those features manually, because we have lots  
13 of, we have thousands and thousands if not millions of  
14 features that can be extracted from this.

15 Now what we do is -- Okay, so those models  
16 exist out there. They are actually in literature and  
17 some of them are out there in open sources.

18 CHAIR BIER: Excuse me.

19 DR. AL RASHDAN: Go ahead.

20 CHAIR BIER: We lost your video for some  
21 reason.

22 DR. AL RASHDAN: Is it still lost?

23 CHAIR BIER: Oh, there we go. I think  
24 it's fixed.

25 DR. AL RASHDAN: Okay.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 CHAIR BIER: It was a goof on my end.

2 DR. AL RASHDAN: Okay. But audio is fine,  
3 right, you heard everything I just said?

4 CHAIR BIER: Yes.

5 DR. AL RASHDAN: Okay. So those models  
6 are available in literature, as I was saying, and some  
7 of them open source, you can use them. They have been  
8 trained to do general classification.

9 Like the one I showed you earlier, if you  
10 recall I showed you one model here where I said there  
11 are classes like Candle, Canon, Fire (phonetic), those  
12 are pre-determined classes that are part of this  
13 model.

14 So they were built for various reasons and  
15 they have certain classes in them. The good thing  
16 about those models is that they have already been  
17 optimized for various forms of image detection.

18 So what we can do, we are interested in  
19 thermal fire. I don't have to create a whole model  
20 from scratch and design the whole thing from scratch.  
21 I can actually go to those models, take one of them,  
22 and start unlocking part of it.

23 So this model has already weights and it  
24 has classes, it can classify certain objects, maybe  
25 not fire, but the other objects. So I can say, okay,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 I know that the left side of my neuro network is  
2 usually focused on extraction of features while the  
3 right side is usually focused on selection of features  
4 of the features that matter for that object.

5 So I can start unlocking part of the  
6 model, and when I say "unlocking" that means I can re-  
7 train parts of the model and leave the other part  
8 frozen.

9 So the feature extraction part I can keep  
10 it maybe frozen because those models have been  
11 optimized to get good features, but I can unlock the  
12 feature selection side and the weighing at the end and  
13 the classification of fire/no-fire towards them.

14 Now one thing I should mention is we do  
15 need to add a layer on top of whatever model we get  
16 here for the actual selection of fire, so it's just  
17 something like what I have shown you here towards the  
18 end.

19 Our reason of telling you this is we  
20 wanted to test all the scenarios, all those models,  
21 for detection of fire, but we wanted also to figure  
22 out how much do I need to unlock out of those models  
23 to get me good results.

24 So what I just described to you know is  
25 called transfer learning, because we have a model that

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 was trained to detect some features or some objects  
2 and we are using it now to detect fire, and that's  
3 transfer learning.

4 So what we did is those are the nine  
5 models I was showing you earlier that exist and we  
6 started looking at the percentages of the model that  
7 we are unlocking.

8 We went all the way down from we're not  
9 unlocking anything and we're just adding a layer  
10 towards the end for fire classification to unlocking  
11 the whole thing and re-training the whole model.

12 We are using the same architecture that is  
13 presented previously but we are looking the whole  
14 parameter. We can re-tune all the weights of the  
15 model that we need. So we did all those variations to  
16 try to figure out what gets us the best result.

17 Now before I talk about performance and  
18 the results I need to tell you what, I need to explain  
19 one parameter that we often use in machine learning,  
20 it's called the F score.

21 Before I do that I need to describe two  
22 different metrics called precision and recall. There  
23 are different ways you can define this, but the  
24 easiest way I think of them when I look at them,  
25 precision is usually something that gives you an

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 indication of how many false positives you had.

2 You want to have high precision, which  
3 means you don't have a lot of false positives. Recall  
4 is focused more on the false negatives. So false  
5 positive in our case here, fire being un-flagged even  
6 though there is no fire, false negative. The  
7 algorithm is saying there is no fire but there is  
8 fire.

9 Now going back to my point earlier, which  
10 I think someone mentioned it might not be accurate, we  
11 assumed false negative, meaning if we miss a fire, is  
12 much more severe than flagging a fire, a false fire.

13 So we can bias our model. The way we bias  
14 our results in our models is this F score has a beta  
15 factor and as you increase the beta factor in your  
16 analysis you increase basically the importance of  
17 recall, which is the importance of false negatives.

18 So in our case we used a beta of two and  
19 this is saying I care more about false negatives than  
20 I care about false positives. So, please, in my new  
21 results I want a higher penalty for false positives  
22 than I do for a false positive, for a false -- Sorry.  
23 I want a higher penalty for false negatives than I do  
24 for false positives.

25 That's what I am going to show you. I am

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 going to show the F2 score rather than the F1 score.  
2 F1 score means both of them are kind of equal  
3 importance, but, again, we wanted to bias it because  
4 we wanted to make sure we don't get false negatives or  
5 we can catch false negatives very well.

6 So this is the results metrics. Those are  
7 the models we have tried and those are how much of  
8 those models were unlocked and trained. As I said, we  
9 went all the way from zero percent to 100 percent.

10 Actually, this is the models in action.  
11 I am showing you here some of them. I know there is  
12 13 because we modified some of the models and we  
13 called them a different number, but it is those  
14 models.

15 You can see they are not all consistent  
16 all the time, right. I mean you see here there is a  
17 flame with a confidence score in every one of those  
18 videos and you see sometimes some of them are blue,  
19 some of them are red, which means they are not  
20 consistent and that's a blessing in this case.

21 I will explain to you why, but before I do  
22 that, so one of the things we learned from this is  
23 some models do not like it when they get unlocked, you  
24 get actually worse results versus others, they  
25 improve.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           Some of them are actually bad across the  
2 board. If you consider 98.4 or 98 point something bad  
3 those models were inferior to the rest of the models.  
4 But I mean we created the whole matrix of this and the  
5 end here was to select the best out of every single  
6 model.

7           But one more step we can do, actually,  
8 once we select the best out of every single model we  
9 can build an example. We can combine the decisions  
10 from all the different models in our decision-making  
11 process.

12           The way we combine things is every model  
13 has an accuracy, which is what I showed you earlier,  
14 and, as I said, we selected the best ones. We can  
15 multiply the accuracy it got by how confident it is  
16 that it's a fire or no fire.

17           Basically sum them up on both sides and  
18 that will give us a score for every one, for every  
19 image, and that's how we can classify this from all  
20 the different models as fire or not, and that's what  
21 we are doing here.

22           So when we did that there is the best  
23 model for each. When we did the example of all models  
24 we get 99.69 percent, so we actually -- We were at 99  
25 percent for most of the models, so really a very

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 accurate model.

2           Keep in mind this is the F2 so we are  
3 taking much more of a penalty when we see a false  
4 negative than a false positive. Now when we constrain  
5 the model, what "constrain" means in this sense, we  
6 actually took out some of the bad ones.

7           For example, one like this one we would  
8 say, okay, this is really dropping down the whole  
9 performance of the group, maybe we should take it  
10 down, or something like this one, it depends on what  
11 model we are using.

12           So we select only the best ones that say  
13 that. We can even get to 99.74 percent F2 score. So  
14 that was very promising.

15           So now going to the point that kept coming  
16 up, how do I know that models are flagging fire for  
17 the right reason and what do I do when I see  
18 something, an image, that is misclassified, either  
19 it's fire and was classified as no fire or the other  
20 way around.

21           So there is -- In computer vision there  
22 are multiple ways you can introduce some  
23 explainability in your decision-making process. One  
24 of the algorithms that is existing out there is called  
25 Grad-CAM.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           What a Grad-CAM does is, it basically  
2           unlocks the model at some point, so if I go back to  
3           that picture, you can unlock the model at one of those  
4           layers and it starts converting those back into the  
5           picture so that it can tell you before the decision is  
6           made what areas of that image where basically the  
7           model is actually focusing on when it's making that  
8           decision.

9           If I go back to this picture, I have three  
10          different pictures here as examples and I am showing  
11          you the nine models. Let's start with the first one.  
12          You can see the whole picture is full of fire but the  
13          models were focusing on various parts of it.

14          The good thing about this is that all the  
15          models are focusing on the fires. They are not  
16          looking at anything else in this case. In this case,  
17          the second case, we have a fire kind of on the edge of  
18          the image and here is a lot of smoke and what you can  
19          see is that some models are actually focusing on the  
20          fire but others are actually looking at the smoke, and  
21          that's the bias I was telling you about.

22          This is telling us that those models had  
23          much more bias towards smoke. They start seeing more  
24          -- They are also recognizing smoke as fire versus the  
25          other ones and are much more capable in distinguishing

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 fire from the rest, and that's why I said diversity in  
2 models is a good thing, it's a blessing.

3 The third one here, I think all of them  
4 worked really well. They are all kind of focused on  
5 the fire, various parts of this fire, but they are on  
6 the fire.

7 So now when we did our validation or our  
8 testing and we saw that something was misclassified,  
9 very often we see when we analyze it this way we would  
10 see a focus of the machine on something that might not  
11 be the fire but resembles fire, or the other way  
12 around, there might be a fire but for some reason the  
13 machine is looking at something else.

14 Usually, as I said earlier, the issue when  
15 the machine is looking at something else usually it's  
16 a problem with your dataset. There might be a bias in  
17 your dataset that is forcing it to consider something  
18 else, like smoke in this case, as fire even though it  
19 doesn't see fire.

20 And then you have to reduce the amount of  
21 images with smoke and fire and kind of balance it in  
22 order to get the point where you can eliminate those,  
23 but that's how we did that validation of testing.

24 I will stop here for any questions before  
25 I proceed. I don't know how much time I have actually

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 left.

2 CHAIR BIER: About 15 minutes or so.

3 DR. AL RASHDAN: Okay.

4 CHAIR BIER: A little more.

5 DR. AL RASHDAN: Okay. I still have a  
6 couple of slides, so are there questions?

7 (No response.)

8 DR. AL RASHDAN: All right. So let's move  
9 on now to -- I explained to you how AI works, I  
10 explained to you how the data works, the data  
11 collection part worked, what the model does, how we  
12 optimized it and we evaluated the performance of it.

13 Now, as I said to you earlier, we have  
14 partnering utilities, we have created with a  
15 partnering utility that is, or a consulting utilities  
16 that are interested in maybe moving this into  
17 deployment.

18 So we had to some analysis into how  
19 compatible computer vision based AI with the current  
20 safety standards. One thing I need to mention before  
21 I go into this, at least I can say on myself, I am not  
22 an expert in this field.

23 So this is what we expect the utilities to  
24 do rather than we are doing. We create the technology  
25 and we validate its performance, we evaluate it, and

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 then we pass it to the utility to do this.

2           However, in this case specifically we  
3 wanted to do a preliminary analysis. AI is a buzz  
4 topic and we wanted to see how easy it is for us to  
5 build an AI model that is compatible with the current  
6 safety standards and regulations.

7           The focus here was on fire watch. We have  
8 a report published on this which looks at the various  
9 aspects of AI or the characteristics of AI and how  
10 they fit.

11           One thing to mention on the side, there  
12 was a, as Bruce mentioned, a recent Executive Order  
13 from the White House on safe, secure, and trustworthy  
14 artificial intelligence and there was a specific  
15 mention of the need for creating standards.

16           So our hope is that this preliminary  
17 study, and, again, I emphasize "preliminary" because  
18 we are not an expert in this but we did provide some  
19 insight there, is going to feed into something that is  
20 becoming more like a standard that might be developed  
21 by the industry.

22           So the conclusion from this study is this  
23 table. In this case what I am showing you to the left  
24 here are the characteristics of the fire watch AI. So  
25 in our case with the fire watch models, they were

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 mostly open source. The data was open source, the  
2 models were open source.

3 We had to frequently update them. If you  
4 think about this as something that gets deployed that  
5 might still be the case. I mean someone mentioned the  
6 blue fire as an example.

7 Let's say you have deployed this and then  
8 you realize there is something missing, you might need  
9 to update this later on. We did use massive amounts  
10 of data to sort through this.

11 We had to sort through a massive amount of  
12 data to get to the point where we have enough data to  
13 train a model. You might have to train the model on  
14 a predict basis if you see things that are not  
15 satisfactory.

16 The model itself is probabilistic. There  
17 was some mentioning about if there are some pixels.  
18 By the way, maybe I should have understood that at  
19 some point in my presentation.

20 So one of the benefits of using this in a  
21 video stream is if you miss a frame there will be  
22 other frames that those pixels that were not visible  
23 is probably going to show at some point.

24 As I showed you on the explainability  
25 slide, it's not really focused on a certain, in one

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 area, it's very broad in terms of its focus, so that  
2 shouldn't be as much of an issue, at least we haven't  
3 observed it. I mean it might be, but we haven't  
4 observed it in our testing.

5 Then there are the other aspects of AI.  
6 I can go through them all. There is the  
7 explainability part, the bias, I kept talking about  
8 the bias in the datasets and how it impacts the models  
9 and the results.

10 Also, the other question is when we do AI,  
11 we created this model using those nine models, we  
12 created this example of models. Others might use  
13 different models, so there is no one way to solve this  
14 problem. There are people who would solve it in  
15 different ways and get different results.

16 Robustness in new conditions is what we  
17 have discussed earlier, and then, of course, we need  
18 a special skillset for AI, which is something that  
19 might not be there yet to the level we want if we want  
20 to deploy those technologies in the plants.

21 Then I am showing you here the different  
22 aspects or different requirements we look at in our  
23 safety standards. I am almost positive everyone on  
24 the call would know what these are.

25 So what we did in this study is we looked

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 at what does every one of these impact from the  
2 requirements point of view. I will give you just one  
3 example because of time sake.

4 So let's talk about the open source one.  
5 If we are talking about the open source aspect of the  
6 dataset and the models the problem is if I have two  
7 different data sources like Google images and then  
8 let's say there is Yahoo! images, those are two  
9 different data sources.

10 We don't know how much overlap is there.  
11 The challenge with overlap is that if you are breaking  
12 your dataset into training and then validation and  
13 testing you might end up in your testing dataset with  
14 some images that you actually use in training because  
15 you didn't recognize that there is an overlap and you  
16 used images in your training dataset that are also in  
17 your testing dataset.

18 Keep in mind those are tens of thousands  
19 and sometimes hundreds of thousands and can go up to  
20 millions of images. So we're not going to be able to  
21 sort through all of them and find all the overlaps.

22 The overlap doesn't have to be exactly the  
23 same image, it can be very close to it, maybe from the  
24 same video but with some differences. So if we do  
25 have that overlap then that means that I have a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 problem with independence of my datasets.

2 When I am doing validation my dataset for  
3 validation is not really independent from my training  
4 dataset. The other thing is that when we are open  
5 source models they also do overlap in terms of the  
6 foundation of how they work.

7 If there is kind of a common cause failure  
8 -- Sorry. So if there is some sort of failure in one  
9 model it can be also existing in the other models,  
10 which is what a common cause failure would be.

11 Finally, there is also the cybersecurity  
12 concern. Again, we are not experts in those areas.  
13 I am just giving you those as examples of things we  
14 considered on a high level.

15 A cybersecurity concern would be in this  
16 case if an adversary went into one of those datasets  
17 knowing that we are using it for fire watch models and  
18 injected false data, so data that doesn't represent  
19 fire and labeled it as fire, or the other way around.

20 When we load those models in maybe  
21 retraining, when we are retraining, we might load that  
22 falsely-labeled data and confuse our model of the  
23 greatest performance, and that is where the  
24 cybersecurity concern would be.

25 So you can come up with some conclusions

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 from this. For example, there are ways to create  
2 independent datasets using GANs. If you haven't --  
3 You probably are familiar with some of those recent AI  
4 advancements when you can tell an AI engine, a  
5 generative AI engine, to create a picture of a teddy  
6 bear under a Christmas tree or so and it can create  
7 that for you.

8 So maybe that can help. However, it goes  
9 back to the point where how realistic it is and is it  
10 going to be beneficial to the training process or not.  
11 We haven't looked into that. Those, again, are just  
12 some preliminary findings from this, and then maybe we  
13 need a method to quantify them in this.

14 What I meant with this example is just  
15 we're looking at one of those. I just gave you some  
16 aspect on some of those green boxes. Those green  
17 boxes are things where we need to look more in depth  
18 into those compatibility of AI issues.

19 But the whole -- If you are interested in  
20 learning more it's all discussed in this report and it  
21 is available in the public domain, so there is the  
22 website. With that, I end my presentation. Thank  
23 you.

24 MEMBER REMPE: I'm curious about the  
25 vision and maybe you don't know what the utility or

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 owner operator the plant will do with this software.  
2 Are they planning -- I mean you admitted that it's not  
3 100 percent perfect.

4 Are they planning to have a person, for  
5 example if their system is offline and they implement  
6 this software and it comes up and says, hey, there is  
7 a fire, will they have a person check before the fire  
8 suppression system is activated or is it going to be  
9 tied into the fire suppression system and just let it  
10 fire off, which could have some adverse effect for  
11 some staff?

12 I remember Halnon mentioned about keeping  
13 people out of the area. I am just curious if this had  
14 been discussed very much.

15 DR. AL RASHDAN: So what the utility is  
16 building as part of this effort is this cart. The  
17 only thing I talked about today was related to the  
18 camera optical vision detection, but there will be  
19 other sensors.

20 The aim of the other sensors added to that  
21 cart is to reduce the false positives as much as we  
22 can. Now, as you mentioned, I don't know how the  
23 utility is going to deploy it.

24 I mean I kind of -- We had to have  
25 different scenarios of how this is going to be

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 deployed, but, honestly, my role ends with getting the  
2 right detection and telling them about the models and  
3 then they can take it from there and decide how they  
4 want to deploy it.

5 MEMBER REMPE: Yes. And I mentioned this  
6 earlier with the Staff and the Staff came back saying,  
7 yes, we understand time for humans to react under  
8 pressure and that that could implement or introduce  
9 more errors, but I just am wondering because I think  
10 it would be important to have a human in the process  
11 and then to put the human in the process I am not sure  
12 how much savings there will be and there is a  
13 potential for some adverse effects.

14 Anyway, just a pointed I wanted to bring  
15 up here also. Thank you.

16 DR. AL RASHDAN: Okay.

17 MEMBER HALNON: This is Greg. I've got  
18 one other question. You are going to finish all this  
19 and give it to the USA or whoever is going to deploy  
20 it, but how are you going to convince them it's  
21 accurate enough for them to submit whatever they need  
22 to submit to whoever they need to submit it to to show  
23 that this equivalent or better than a human fire  
24 watch?

25 I realize there's other things on the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 cart, but you've got to sell this as accurate enough.  
2 I am not going to say accurate, but accurate enough.  
3 So is there going to be some testing done beyond what  
4 you have done here?

5 The problem is that, in a plant, the  
6 amount of hours to event, the ratio is really high.  
7 You've spent a tremendous amount of hours monitoring  
8 and nothing happens for a long a time, so you're not  
9 going to have any kind of real-time data other than  
10 there is nothing happening.

11 So is there a testing program envisioned  
12 by USA or by you guys or jointly?

13 DR. AL RASHDAN: So at the moment this  
14 project is aimed towards getting to the demonstration  
15 point. I think that USA is planning to develop the  
16 test plan and engage the NRC in -- I mean they are  
17 already engaging the NRC, but engage maybe more  
18 extensively on the use of this in 2024.

19 MEMBER HALNON: Okay.

20 DR. AL RASHDAN: So I don't have an answer  
21 for you, sorry.

22 MEMBER HALNON: Okay. No, that's  
23 understandable. Thank you.

24 CHAIR BIER: Well and Joy mentioned the  
25 problem of either false positives or maybe a positive

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 like a trash can fire but at that moment you don't  
2 want to activate, you know, fire suppression and, of  
3 course, you could also get false positives because  
4 like the Canadian wildfires are blowing smoke every  
5 place.

6 The other thing I can envision regarding  
7 implementation similar to Greg's issue is you may have  
8 situations where you implement but you still keep some  
9 level of fire watch and just less frequent than what  
10 you had before in case the software misses something  
11 and, you know, it seems like it could be very  
12 complicated how best to implement.

13 But overall this was very interesting,  
14 very impressive work.

15 DR. AL RASHDAN: Thank you.

16 CHAIR BIER: Other questions or comments  
17 from members?

18 (Off-microphone comment.)

19 CHAIR BIER: Yes.

20 MR. PRIMER: Yes. I just want to thank  
21 Ahmad for taking us through that work. He has several  
22 reports that have that detail in different areas that  
23 I shared, so those links are there.

24 I will also mention we have a stakeholder  
25 engagement meeting with individual researchers and

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 collaborators that have been involved with the  
2 different research activities December 5th through  
3 7th.

4 So if you are interested in joining that  
5 I will be happy to get you on the invite. It's five  
6 session, the AIML is a one 2-hour session, so one of  
7 those.

8 The others are human factors engineering,  
9 the digital architecture, the things we talked about,  
10 you heard Bruce mention. So we would really  
11 appreciate if you are interested let me know and we'll  
12 get you an invite so you can tune in and hear more  
13 about the other activities as well.

14 I don't know if Bruce is still on. I will  
15 just turn it over to Bruce to close us out.

16 MR. HALIBERT: Yes, I'm still on, Craig.  
17 Other than what you have already said, I would just  
18 thank the NRC and the ARCS for the opportunity to  
19 participate in this meeting.

20 It aligns very well with our Memorandum of  
21 Understanding for Technical Exchanges. Obviously,  
22 there is a lot of other areas that we didn't talk  
23 about today, but I think as Craig mentioned if you are  
24 interested in participating in the stakeholder  
25 engagement meeting I would recommend you reach out to

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 Craig and we'll be sure to add you to the invitation  
2 list for that.

3 Other than that I would also just refer to  
4 Alison Hahn, since she is the Federal Program Manager  
5 and the Office Director over at the Office of Nuclear  
6 Energy, to see if she has anything else she would like  
7 to add about that, our work today.

8 MS. HAHN: I don't have anything to add.  
9 I just wanted to thank you all for taking the time to  
10 listen to this work and just reiterate what Bruce had  
11 said, if there is any additional questions please feel  
12 free to reach out.

13 We always appreciate the conversation and  
14 the opportunity to share the work that we do, so thank  
15 you.

16 CHAIR BIER: Okay. Thank you for  
17 educating us. I guess we will now have a break until  
18 3:15, is that correct, Christina, and then we will be  
19 back with Dr. Cummings from George Mason University.

20 (Whereupon, the above-entitled matter went  
21 off the record at 2:57 p.m. and resumed at 3:25 p.m.)

22 CHAIR BIER: So I'm very glad that Dr.  
23 Cummings was -- there we go, we're up and running in  
24 the presentation. I'm very glad that Dr. Cummings was  
25 willing to join us today and educate us about the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 pluses and minuses of AI.

2 Dr. Cummings has a rather unique  
3 background. She was trained as an engineer and was  
4 also a Navy fighter pilot back in the day. Most  
5 recently before joining the Engineering School at  
6 George Mason, she was Senior Safety Advisor at the  
7 National Highway Traffic Safety Administration.

8 And as you are probably aware, the  
9 automotive industry is, depending on how you want to  
10 think about it, either ahead of us or behind us on  
11 rolling out AI. They are doing it, but maybe not  
12 having such great success with it yet.

13 So I thought hearing a little bit about  
14 what's been going on in the transportation industry  
15 would help inform us and educate us to do a better job  
16 in the next 10 years. So with that, Missy, I'll ask  
17 you to proceed.

18 DR. CUMMINGS: Yes. Okay. Thank you so  
19 much. I'm sorry that my camera is not working today.  
20 But I would like to tell you I look the same as the --  
21 the last time we worked together was about a decade  
22 ago. I'd like to tell you I look the same, but I  
23 don't sadly.

24 But it is interesting. You know, I did  
25 some work there with your human factors group about 10

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 years ago. At that time, we were looking at the  
2 impacts of boredom. We did some checklist work. And  
3 I think my time with the NRC actually was really  
4 important for me to really understand what it meant to  
5 be a regulatory agency because then when I went to the  
6 National Highway Traffic Safety Administration, I  
7 think I had a much better understanding of, you know,  
8 where the viewpoint of our regulatory agency in theory  
9 is, but, of course, looking at things from the  
10 automotive world was quite different.

11 And I think that you're going to find --  
12 I think that will be kind of hard for people on the  
13 call, like, you know, you guys take safety very  
14 seriously whereas I think the automotive self-driving  
15 industry takes it less seriously. And the question  
16 is, you know, when you have a strong regulatory agency  
17 like the NRC or FAA as opposed to a weak agency like  
18 NHTSA, and I would say maybe Federal Railroad  
19 Administration, you know, it's kind of interesting to  
20 me, and maybe that's another talk. In fact, I would  
21 love to hash it out with you guys sometime about some  
22 of the unique differences between the regulatory  
23 agencies.

24 But before we talk and get into some of  
25 these details, it's always good, I think, to start out

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 with some definitions. Because I'm going to talk  
2 about both kinds of AI today. And that is symbolic  
3 and connectionist AI, which you know mostly as neural  
4 networks.

5 And so -- this is important to talk about  
6 because AI shows up in different places in different  
7 systems. And I'm going to talk about both kinds  
8 today. But the first is what we typically will call  
9 good old fashioned AI. This is some rules-based AI,  
10 using things like ontology like you see here that we  
11 can break an apple into all different elements of what  
12 does it mean to be an apple? Is it an origin story?  
13 Is it at the actual physical apple? Is it a kind of  
14 fruit for example.

15 The rules-based, if then else if you will,  
16 good old fashioned AI, has done a pretty good job. I  
17 mean, I think there was an AI winter when this first  
18 came around because people wildly overestimated the  
19 capabilities of GOF AI. And then now that neural nets  
20 have taken off, I anticipate another second winter.  
21 We'll see if it has the same characteristics of the  
22 first one.

23 But I think what we're going to see if  
24 that people yet again are wildly overestimating what  
25 the capabilities of these are. You know, it's

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 interesting because these neural networks, also known  
2 as connectionist AI, takes on -- it has -- it is more  
3 -- it mimics intelligence more, which I think makes  
4 people believe that there is actually some real  
5 intelligence. But indeed we're actually no closer to  
6 real intelligence today as we were in the 70s and 80s  
7 when symbolic AI came along.

8 Okay. So with those definitions in the  
9 background, I wanted to talk to you about -- well,  
10 first of all, let's talk about some problems. And  
11 these are problems that I have seen as when I worked  
12 with the National Highway Traffic Safety  
13 Administration that they have parallels with  
14 connectionist AI in the form of large language models.  
15 We'll talk about that in a second.

16 But the first I want you to see, this  
17 picture comes to me from Toyota Research Institute.  
18 They've been great collaborators. And the car -- they  
19 had a self-driving car that went to this intersection  
20 in Boston, and it froze. And they couldn't get the  
21 car to move. Eventually, a human had to take over and  
22 start driving. And they took the car back to the lab  
23 and "pulled the tapes". And so they tried to figure  
24 out this is the event that caused the car to freeze.  
25 Now I'm about to hit the button, and when I do, you're

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 going to see what the car saw.

2 So instead of a single moving truck, the  
3 car saw two trucks, a bus, a gigantic person, a  
4 building, a fence, some poles, a traffic stoplight,  
5 like, a bunch of things.

6 And let's go back. So you just saw a  
7 truck. And indeed your brain took no energy to  
8 classify it as a truck. And you didn't actually  
9 probably pay attention to any of the details, the  
10 sign, the number, unless you were actually thinking  
11 about moving because it had no relevance to you. So  
12 this is top down reasoning.

13 But that's now how computer vision and  
14 neural nets work. They look at the world at the pixel  
15 level. And so when the convolutional neural net that  
16 classifies this, it looks at clusters of pixels and  
17 compares them against the training data that it has in  
18 its "head" and so that's why you see the gigantic  
19 person that looks like it's about to attack and then  
20 two trucks and a bus.

21 Even though if that were really a bus,  
22 it's such a small -- you know, it's that percentage of  
23 that image in between two trucks that is still what it  
24 saw.

25 So this problem with -- this is a real

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 hallucination. And so you hear a lot about large  
2 language models, and we'll talk a little bit more  
3 about what it means to hallucinate in a large language  
4 model. But computer vision systems legitimately  
5 hallucinate. They see things that are not there, and  
6 they don't see things that are there.

7           And so this has become a huge problem in  
8 self-driving cars. We don't, we being engineers and  
9 computer scientists, we don't know how to fix this  
10 problem. And we're trying to do things like sensor  
11 fusion, but we're still having quite a bit of  
12 accidents due to the hallucinations of convolutional  
13 neural nets. And it's not clear that we're going to  
14 be able to get over this gigantic wall. So we'll come  
15 back to these thoughts in a few minutes because they  
16 have real implications in the real world.

17           So this is the Bay Bridge Tunnel in San  
18 Francisco on Thanksgiving of last year. So the  
19 picture you just saw is the hallucination of many  
20 things when there was just one thing. Indeed, the  
21 Teslas, who use computer vision who also leveraged  
22 convolutional neural nets, although this car is not  
23 actually self-driving even though it says it is, the  
24 convolutional neural nets that caused this accident  
25 was a shadow.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           We're not sure what it saw actually. So  
2 a Tesla going 65 miles an hour went through this  
3 tunnel, slammed on its brakes for no reason that we  
4 can discern other than it saw something, which we  
5 think is a shadow, and it did such an aggressive  
6 maneuver it caused eight cars behind it to crash into  
7 it. We are very lucky that nobody was seriously  
8 injured on this day. You know, the picture looks  
9 pretty dramatic, but this is a good reminder that the  
10 hallucinations were not just seeing things  
11 incorrectly. It is seeing things that literally are  
12 not there.

13           So, again, until we fix this problem,  
14 we're going to continue to have problems not just with  
15 self-driving, but with driving assist (phonetic).

16           So I want to put all this in the context  
17 of autonomy, AI and reasoning. So a few years after  
18 -- about three or four years after I last worked with  
19 NRC, I wrote this paper that talked about what it  
20 means for an autonomous agent, whether we're talking  
21 about a human or a computer-based system, what does it  
22 mean for an autonomous agent to reason?

23           So first when you need to reason, when you  
24 first learn to drive a car, you have to tell yourself  
25 to stay between the two white lines on the road. And

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 you have to really pay attention to this. You have to  
2 train your brain to stay centered between the two  
3 white lines.

4 And after a few hours, you become -- this  
5 becomes highly automatic. You are used to that, what  
6 we call world view. And your brain automatically  
7 adapts and indeed you don't have any more problems  
8 after you first learn to drive until you get older,  
9 when your sensors, your eyes, start to fail, and it  
10 becomes hard to see. But for the most part, this  
11 stays throughout your lifetime. It's a very well-  
12 honed skill.

13 And then once you can do that, once you  
14 make this a highly automated process that you don't  
15 really have to think about, then you free up cognitive  
16 resources to do rule-based reasoning. And when we  
17 drive, what that means is that we see signs and  
18 signals in the world, like a stop sign, and we know  
19 that, you know, about 50 feet before we get to the  
20 sign, we start to slow down and come to a full stop  
21 before you can go again.

22 And indeed, autonomy can execute the  
23 staying between the white lines and doing rule-based  
24 reasoning quite well because as long as the sensors  
25 can see the sign and can see the white lines on the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 road, indeed automation can actually stay between the  
2 white lines and adhere to signs far more reliability  
3 and with much more precision than humans can.

4 But then once we get past rule based  
5 reasoning, you start to see the gray arrow really  
6 increase in size. And this is matching a growth in  
7 uncertainty. And it's this growth in uncertainty  
8 which is causing problems for systems with any kind of  
9 neural net in them so really any form of connectionist  
10 AI.

11 And in the picture you see, there is a  
12 stop sign that is partially obscured behind leaves.  
13 This is a huge problem for self-driving cars because  
14 they are trained, their neural net database is  
15 trained, on let's say 50,000 pictures of stop signs,  
16 but they are not trained on pictures of stop signs  
17 with leaves in front of them.

18 And so they don't know that this -- even  
19 you knew immediately when I went to this slide, I  
20 didn't have to tell you it was a stop sign. You knew  
21 what it was. Self-driving cars must be explicitly  
22 told. And so there is really very little  
23 generalization. You would have to then go back and  
24 train the neural net with all of these leaf covered  
25 stop signs, but this becomes kind of unwieldy because

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 10 percent of coverage in leaf coverage on a sign is  
2 very different from 20 percent and very different from  
3 30 percent.

4 And indeed we don't know at which point,  
5 is it between 30 and 35 or 30 and 37 percent, for  
6 example, where the systems will no longer be able to  
7 generalize based on their original data, their  
8 training. And we'll come back to this a little bit  
9 more because that really kind of falls under the title  
10 of AI maintenance, which is a big deal.

11 And if you can't get to -- if you can't  
12 get past reasoning under uncertainty, you can't really  
13 get up to the top of this ladder, if you will, which  
14 is expert-based reasoning. And this is when you have  
15 to really try to figure out the unknown unknowns,  
16 right?

17 And this picture from a real intersection  
18 somewhere in America is a good example that if you got  
19 yourself there, you're not supposed to be there.  
20 You're not supposed to be able to get there in the  
21 first place. Somehow you made a mistake, and you  
22 would know that to get out of this place, I would just  
23 need to turn around and go back the way I came because  
24 that's how I got there.

25 It would be very hard for a self-driving

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 system that is also programmed not to break the rules,  
2 right? And so you would get a car that would freeze  
3 here because it would not be able to reason under  
4 uncertainty.

5 The uncertainty can come from the  
6 environment. So, for example, the leaves blowing  
7 across the sign. It can come from the AI blind spots,  
8 meaning if the data -- if it wasn't explicitly trained  
9 in the data, then whatever self-driving car or really  
10 any neural net based AI that sees something it has  
11 never been trained on, it has no idea what that is.  
12 And it's blind to that thing.

13 And it can come from human behavior.  
14 This has been quite a difficult road for self-driven  
15 cars to figure out, especially around things like fire  
16 trucks and police officers where -- and let me just go  
17 ahead and fast forward. There we go.

18 Indeed, it is interactions with the fire  
19 trucks and first responders that has caused self-  
20 driving cars to really get a hard look at California  
21 state regulators as to whether or not these cars  
22 should have permits. And indeed Cruise just recently  
23 had their permit pulled because it struck a  
24 pedestrian, which wasn't its fault but then kept going  
25 and dragged her for 20 feet underneath the car,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 causing the bulk of her damages.

2 So we are still in the place where cars  
3 cannot reason under uncertainty because their  
4 connectionist based neural nets cannot reason under  
5 uncertainty.

6 And I want to be very clear on this.  
7 Neural nets, which are what power both convolutional  
8 neural net computer vision but also large language  
9 models, they don't know anything. They don't know yes  
10 versus no. They don't know right versus wrong. They  
11 only "know" what they've been shown before.

12 And so any estimates, approximations of  
13 intelligence, are at best mimicking intelligence but  
14 are indeed not intelligent. And this is super  
15 important for people to remember going forward is that  
16 even though you can talk to ChatGPT for example, and  
17 it's quite competent and it thinks it knows  
18 everything, but it is also very confidently wrong,  
19 even though it doesn't know it's confidently wrong.

20 As a professor, we've all had those  
21 students that swear, swear, swear they know the answer  
22 and that's just like ChatGPT. And we are going to  
23 look at large language model examples here in a  
24 minute, but the thing that I really want to  
25 reemphasize is that there is this wall between rule

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 and knowledge-based reasoning when it comes to  
2 artificial intelligence.

3 I don't care whether we are talking about  
4 good old-fashioned intelligent AI or we're talking  
5 about connectionist neural nets, computer vision,  
6 anything with any kind of neural net cannot get past  
7 rule-based reasoning.

8 And so until we engineers, computer  
9 scientists, thought leaders, you know, until somebody  
10 comes up with a different way of encapsulating true  
11 knowledge and expert-based reasoning, artificial  
12 intelligence really cannot go past rule-based  
13 reasoning.

14 Now lots of people want to fight with me.  
15 One of the fathers of AI, Geoff Hinton, swears,  
16 swears, swears, that they are becoming sentient, but  
17 they are not. They are not because they cannot reason  
18 under uncertainty.

19 And, you know, to me it is just one of the  
20 big problems that the computer science community is  
21 just sidestepping is why in hell these large language  
22 models and any other neural net based system how they  
23 reason and how they fail. And so we will talk a  
24 little more about that.

25 But before we go onto the large language

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 model, I wanted to kind of put this into context in  
2 terms of process control. So you see that I've got a  
3 human supervisor that is trying to do some task.

4 And we can say in this case they are  
5 trying to supervise what's happening with a nuclear  
6 reactor for example. It could be really any kind of  
7 process control, which is mediated by a computer in  
8 the middle. That computer shows them, for example,  
9 what the coolant level is for example, or what the  
10 power level is, right? So we're heavily reliant on  
11 the computers to be able to know what's happening in  
12 the world.

13 So given this backdrop, I wanted to show  
14 you, it's a little map of where I believe that AI  
15 could and should be used inside process control.

16 So at the bottom, you can see that  
17 connecting automation to the task, in this case  
18 automating nuclear reactors, well, you guys have been  
19 doing -- the whole nuclear industry has been doing  
20 that for a long time. And, you know, with a couple of  
21 verbals for the most part, but especially lately, very  
22 safely, right?

23 And we know how to develop that automation  
24 because it's based on first principles. We've got  
25 equations. These are physic-based systems that are

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 very well modeled, lots of experience and history  
2 behind it.

3 So, you know, the automating of the task,  
4 particularly for process control, any form of AI is  
5 not really needed. But we'll come back to that  
6 because it's going to be really uncertainty dependent.  
7 And I have a quick caveat that we're going to go to on  
8 the next slide. But let's put that aside for right  
9 now.

10 So where do these neural nets -- where  
11 would be a good idea to use them? Well, it turns out  
12 neural nets are actually pretty good if you want to  
13 set up some cameras and automatically detect somebody  
14 sleeping for example.

15 You can imagine -- I've done -- a lot of  
16 the work we've done in boredom over the years, not  
17 just in nuclear power plants, but cockpits, drone  
18 control stations. You know, it's very hard, and  
19 automation is doing a very good job. It's hard for  
20 people to force themselves to pay attention.

21 So having cameras monitor, automatically  
22 monitor, meaning a human does not have to be in the  
23 loop so that they can potentially sound some kind of  
24 either gentle alert or full on alarm if something bad  
25 was happening. Yeah, it turns out for human activity

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 monitoring, connectionist neural nets through  
2 convolutional neural nets is pretty good.

3 I mean, it's not in a critical path of  
4 safety, but it's trying to augment safety. But we can  
5 also reverse that and say we can also use  
6 connectionist models to actually model what is going  
7 on inside the plant. And now we're going to start  
8 talking about the health and status loop because it  
9 turns out that, you know, monitoring through images,  
10 that's becoming more and more of a thing in the  
11 nuclear world, as indeed in all process control.

12 So there is a way that you can potentially  
13 use again convolutional neural nets. You could  
14 potentially use them in real-time to start to have  
15 basically an additional set of "eyes" on the process.  
16 But I think for the most part, people in process  
17 control agree that probably the best use of  
18 connectionist AI is in predictive maintenance. And  
19 we'll come back to that in a few minutes.

20 So I do think that there is a place for AI  
21 in and around process control. But just probably not  
22 in the direct line of process control because of some  
23 of the safety caveats, which we will talk about in a  
24 second.

25 That little asterisk that I gave you is

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 actually a project that I worked on with Pfizer. So  
2 it turns out just in the weirdness of timing, right  
3 before COVID hit, I was working with Pfizer to help  
4 them improve their vaccine fermentation units in North  
5 Carolina. And indeed the team that I was working on  
6 went on to start making the COVID vaccine. So just  
7 really right place, right time. And one of the things  
8 that we were doing is helping them figure out the foam  
9 control in the fermentation unit.

10 So this is the research fermentation unit  
11 that you're looking at in this picture. And you're  
12 basically seeing four fermentation units, and there is  
13 a sight glass where you are zoomed in one of them.  
14 And it turns out that vaccine fermentation is pretty  
15 much the exact same process as making beer. And, you  
16 know, so this was a big eye opener for me, because,  
17 you know, I really hadn't really ever, you know, had  
18 to make beer, but I never realized that vaccines were  
19 made the same way.

20 And why that's critical is you need to  
21 have some foam in this healthy process, right? You  
22 can't have too little foam because there's a problem  
23 with the oxygenation of the content inside the  
24 fermentation unit. But you don't want too much foam  
25 because you don't want to have a spillover, especially

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 if something, like when they're making the meningitis  
2 vaccines, then you contaminate everything, and that's  
3 a big problem.

4           So it turns out we don't know foam  
5 control. We have not good models of foam control. It  
6 reminds me as an aerospace engineer of our problem  
7 with turbulence. We do not still -- and, you know, in  
8 2023, we do not have good models of turbulence. And  
9 it's still kind of a mystery to us, which is the same  
10 thing for foam control in fermentation units.

11           And so it turns out that there is a human  
12 -- this is where you really need the human because the  
13 human just have to sit there and watch the foam. But  
14 it's painstakingly boring. And nobody -- and you have  
15 to have a highly paid tech do this. Nobody wants to,  
16 you know, pay the tech just to sit there and babysit  
17 and then we have the problems of people falling  
18 asleep, being on their phones, and they've had many,  
19 many spillovers, they being the whole vaccine  
20 industry, had had many problems because it's such a  
21 painstaking process.

22           So what we did is we came in, and we put  
23 a computer vision camera on it, using convolutional  
24 neural nets that was able to track the foam  
25 automatically. And before the foam ever got to a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 critical state, as soon as it started to rise and it  
2 began to look suspicious, then what we did is you  
3 could see the little phone. It would actually alert  
4 a tech who was in the plant who was nearby but was  
5 able to be doing something else productive with their  
6 time that would alert the tech that there was a  
7 problem that they could go in and start visually  
8 assessing. And they would be staying nearby and tweak  
9 whatever settings they needed to tweak.

10 And so in this way, this is very much a  
11 process that we know how to control mostly except for  
12 this weird foam element. And it turns out that the  
13 foam control turns out to be a great proxy for how  
14 humans and AI can collaborate so that the humans'  
15 judgment can be used at the same time when it's the  
16 dull, dirty and boring part of the arrangement, then  
17 the AI can be used pretty reliably to just watch for  
18 the foam rise.

19 Okay. So, you know, kind of in that note  
20 then, you know, what could we use AI for if you're  
21 thinking in and around nuclear plants. And I just use  
22 these as a few examples of why, you know -- I don't  
23 want to throw the baby out with the bath water. There  
24 are a lot of problems with AI. But there are also  
25 some really good uses, especially given the project

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 that I just told you about where we were helping  
2 Pfizer with their vaccine fermentation.

3 So I could imagine AI could be -- you  
4 know, there are all sorts of cameras that I see people  
5 making that are looking for radiation and/or different  
6 kind of phenomena inside of a nuclear power plant as  
7 well as tracking radiation. Well, you can imagine you  
8 don't need just cameras because neural nets would also  
9 work on the Geiger counters, all the data coming out  
10 from those.

11 So the nuclear industry is probably  
12 sitting on a ton of data that has gone unanalyzed  
13 beyond, you know, like maybe a set of sensors will  
14 generate a report. You have an operator or maybe an  
15 administrator who looks at these things and say, okay,  
16 that looks good and then they file it away and never  
17 look at it again.

18 I'm reminded of a really recent best  
19 application of artificial intelligence I've seen in a  
20 long time is a group of researchers finally took all  
21 of the echocardiograms that we get when we're in --  
22 when we go for our annual physicals. I'm getting old  
23 so I'm getting them more often than I ever have,  
24 right, and other than a doctor looking at your EKG and  
25 saying, oh, okay, all right. Well, you know, you're

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 fine. Then they throw them into, you know, some  
2 database, and nobody ever looks at it again.

3 These researchers were able to take all of  
4 this, you know, millions and millions of data points,  
5 of EKGs and now they have a really, really good system  
6 where the EKG can be run and even slight nuances can  
7 now be picked up and diagnosed way before you would  
8 ever see them, you know, used through kind of major  
9 cardiac event.

10 And so this is good. Now, they're not  
11 always right and there are some problems. But, again,  
12 this is a screening tool. We're not letting the  
13 computer make the decision. The computer makes a  
14 recommendation that has been kicked over to a  
15 physician. And the physician actually then does the  
16 actual ruling.

17 And so in that way, you know, if you were  
18 to say to me, Missy, what would you like to do? I  
19 would be like I would love, love, love, love to see  
20 what is going on with radiation. How could you track  
21 one person? So you could imagine that you could  
22 create using some AI a profile that tracked an  
23 individual's exposure over time to make projections  
24 about potential health issues in the future or when  
25 you should moderate this. Indeed, we've done

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 something similar for Boeing where we track repetitive  
2 stress injuries of people who paint aircrafts.

3 I know that sounds -- you're like, well,  
4 that's completely different. But it turns out that  
5 there is a lot of repetitive stress injuries for  
6 people who are doing aircraft painting. And so if you  
7 have better predictive models about when people start  
8 to reach what we would consider dangerous -- too many  
9 exposures of painting, for example, then you can know  
10 how to start rotating people out actually much younger  
11 in their career so that you can extend their career  
12 and their ability, you know, to have a good job.

13 So one of the things that I did after  
14 getting out of my job at the National Highway Traffic  
15 Safety Administration is that I wanted to develop --  
16 to give people a way to think about AI hazards because  
17 AI can be useful, but it can also be a problem.

18 So I'm going to explain this accident to  
19 you in the context of this Swiss cheese model, which  
20 is not unknown to the NRC. So you guys are very  
21 familiar with this. But if this car, which is  
22 actually a Cruise self-driving car, if there had been  
23 a driver in that car and the driver had hit the back  
24 of the bus, we would say that the unsafe act was  
25 caused by the driver, who was probably not paying

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 attention.

2 And maybe the pre-condition for the unsafe  
3 act, going backwards in the Swiss cheese model, would  
4 be that each driver had their phone and their phone,  
5 you know, they didn't have a good, you know, policy  
6 about not having their phone on when they were pulling  
7 out from a parking space.

8 We could take a step back and say, well,  
9 maybe there was unsafe supervision. If it's a taxicab  
10 driver, for example, we would say that there wasn't a  
11 strong policy for the company. And they didn't have  
12 the middle managers checking to make sure people  
13 didn't have their phone. And we could go all the way  
14 back to organizational influences to find out what  
15 companies had good safety cultures.

16 I think it's hard for people in and around  
17 the nuclear field to imagine a company without a good  
18 safety culture because, you know, the nuclear field  
19 lives and dies by making sure people are safe.

20 But I will tell you having been in this  
21 field, there is virtually no safety culture for self-  
22 driving cars. That's not maybe necessarily a fair  
23 statement to Waymo, who is not nearly in as much  
24 trouble as Cruise is. But I do see for most of the  
25 Silicon Valley companies it's just a shocking lapse of

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 appreciation for safety.

2 Companies don't have safety officers.  
3 They don't have safety programs. They only train if  
4 they think they're going to get sued. So it's the  
5 Wild, Wild West in self-driving cars when it comes to  
6 safety, which is not a problem you have. But I think  
7 you wouldn't have asked me to give this talk if you  
8 didn't have questions in the back of your mind about  
9 well, what is safety? How are we going to do safety,  
10 especially when it comes to AI?

11 So given the fact that there is no human  
12 in this vehicle that you see here. This is a Cruise  
13 self-driving car with no one in it, that rammed into  
14 the back of the bus. It turns out that the reason it  
15 ran into the back of the bus is because it had the  
16 wrong model of a bus inside the computer vision head.

17 The computer vision system estimated this  
18 bus by capturing the front of the bus and then  
19 estimated its length to be around 40 feet of a single  
20 access bus, which is all of the training data that it  
21 was trained on. But it turns out this was an  
22 accordion bus, so one of those ones that, you know,  
23 will drive, turn the accordion piece, bends and then  
24 a couple seconds later the second half of the bus  
25 turns. And it turns out it's a longer bus than a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 normal bus.

2           And so when the computer vision model  
3 inside the car estimated the length of the bus to be  
4 40, it did not correctly classify the bus. Not only  
5 did it not correctly classify the bus, but it also  
6 didn't ask the LiDAR, which are the laser sensors on  
7 top of the car, for a backup second opinion because  
8 the LiDAR on the car knew exactly where the bus was  
9 and indeed when they pulled the tapes, if you will, it  
10 turned out that the LiDAR sensor knew all the way to  
11 the point of hitting the bus that it was going to hit  
12 the bus.

13           But the way that the system was designed,  
14 it was never designed to call the LiDAR for a second  
15 opinion if the computer vision algorithm was less  
16 than, for example, 95 percent. We don't know what  
17 that real number is, but, you know, there is some  
18 number.

19           So given this accident, this was the one  
20 that really made me -- you know, at this point, I'm an  
21 old curmudgeon engineer, which is to say I was just  
22 shocked that the primary sensor in this scenario would  
23 ever be the computer vision system which, I mean, you  
24 saw the first pictures that I saw on the earlier  
25 slides, I mean, computer vision systems are terrible.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 They are terrible because they are not nearly -- I  
2 mean, forget 10 to the minus 9th. We're talking about  
3 -- we're not reliable to 10 to the minus 2. So it was  
4 shocking to me that they made that the primary sensor  
5 instead of the LiDAR, which is a much more reliable  
6 signal.

7 So that made me kind of sit down and  
8 really power through, well, if we were going to do a  
9 hazard analysis on these systems, what would that look  
10 like? And so, instead of the corollaries that you saw  
11 before, now what I have is instead of the human unsafe  
12 act, we have inadequate AI testing because indeed that  
13 is the last place that a computer scientist or an  
14 engineer will touch before a system goes out into the  
15 real world.

16 And I know that, again, NRC, you're used  
17 to a lot of testing, certifications. You know, these  
18 are issues on the top of your head. But it turns out  
19 for self-driving cars there is no requirement at all  
20 to do any testing, none. And, indeed, companies --  
21 this is true not just for self-driving cars, but also  
22 for driving assist cars like Tesla -- companies can  
23 rollout software. And there's no requirements of how  
24 they have to test it, if they have to test. And  
25 companies, as a result, actually rollout software and

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 test it on the public.

2 So, again, you know, nobody on the planet  
3 would ever allow a nuclear power plant to go online  
4 without extensive testing. But that doesn't happen in  
5 the self-driving car world. So we need to completely  
6 overhaul that. And indeed in this accident, there  
7 were never any tests. So that's why they never caught  
8 this problem because if you never do any tests or you  
9 do very little testing, you won't catch problems.

10 And then one layer behind is the AI  
11 maintenance problem. So one of the big secrets that  
12 nobody is really talking about for AI systems are that  
13 the underlying neural nets have to be constantly  
14 retrained to update their world model if anything in  
15 the environment changes, which in driving is all the  
16 time.

17 Construction sites go up or down, new cars  
18 show up on the market, buildings get built in empty  
19 lots and some building get torn down. So anytime that  
20 there's any change in the environment that needs to be  
21 considered by a self-driving car in theory the  
22 underlying neural nets need to be retrained. But  
23 that's not happening. Companies don't want to do it  
24 because it's very expensive. It's extremely expensive  
25 to train one of these neural nets, forget all of them

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 that they are using, at any given time.

2 So again there's no regulation so they  
3 don't have to because nobody is mandating testing. So  
4 the AI maintenance piece -- when I give this talk to  
5 companies, I tell them that they are going to have to  
6 start formalizing AI maintenance, making new  
7 divisions, getting an AI maintenance officer in place  
8 because if they don't, they're going to keep having  
9 accidents like this. And ultimately, Cruise's  
10 slipshod, you know, very loosey-goosey safety culture,  
11 that's what took them down.

12 Okay. Same thing with the inadequate AI  
13 design before. You should never use your best sensor  
14 as the backup sensor if it's giving you perfect  
15 information or near-perfect information. And if we go  
16 one layer back, then the inadequate AI oversight, the  
17 same problems that I've already talked about, really  
18 poor safety cultures, problems with companies taking  
19 this seriously.

20 And I think this is a sign of, you know,  
21 it's a Silicon Valley Culture of move fast, break  
22 things. It turns out that may be fine for your phone,  
23 which is not safety critical. But when it comes to a  
24 car and, of course, anything to do with a process  
25 control plant, that's just not going to be a good

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 attitude carrying through your operations.

2 So what about large language models? You  
3 know, here is my favorite application of a large  
4 language model. Indeed, I make all my students now  
5 use large language models to correct their grammar.  
6 It's amazing. My life has become so much better now  
7 that ChatGPT has shown up on the scene because my  
8 students are terrible writers.

9 And ChatGPT does a really, really great  
10 job of correcting their grammar because it's very  
11 predictable. It follows a pretty clear set of rules.  
12 And so this is, to me, a great application. I'm not  
13 worried at all about cheating. They can try to cheat  
14 with it, but I catch them every time because it's very  
15 predictable.

16 ChatGPT comes out with such formulaic  
17 writing that it's very easy to spot as a professor and  
18 that is because all large language models, as do all  
19 neural nets, they compute based on their regression to  
20 the mean. They are giving you the most probable  
21 response, the most statistically frequent response,  
22 right? And so you can predictably get good grammar,  
23 but it's also very formulaic, predictable writing.

24 But then my daughter who is 16, and she's  
25 in Calculus, as I was helping her with this problem,

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 I wondered how it would do in math because it's a  
2 large language model. It's not a math model.

3 So I gave it this problem. It went out  
4 into the internet and came back with this answer, X-2,  
5 X-3, X-4, which is correct except all it's missing is  
6 a negative sign, which, you're like, well, that's  
7 pretty good, isn't it? I have a lot of people who  
8 aren't math oriented say, well, that's not bad.  
9 That's pretty close. Well, you know, pretty close,  
10 horseshoes and hand grenades, you know, that's the  
11 only time it really counts. Because in engineering,  
12 a minus sign can be the difference between life and  
13 death.

14 And so I worry because if people -- and  
15 how this happened is it went out on the internet. It  
16 found either that exact problem or something close to  
17 it and then just computed the probable answer. But  
18 because of the negative sign, it doesn't know that a  
19 negative sign is a negative sign. It doesn't know  
20 anything. It probably just recognized it as a hyphen.  
21 And so it didn't put it through with the answer.

22 But why I worry about this is because  
23 students go get this. This is how they're doing their  
24 homework. And they may be doing this in the future.  
25 And if you forget to take the minus sign with it, then

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 you've got serious problems later.

2           You can't trust these things. And you  
3 have to verify every single thing that comes out of a  
4 large language model. You can never take what a large  
5 language model says to you at face value.

6           So they are being -- companies have gotten  
7 wise to the fact that they can be used for bad things.  
8 But this is a great example of how clever humans can  
9 be. So that the prompt prior to this was somebody  
10 trying to get ChatGPT to tell it all the websites  
11 where you could go get pirated content, which is  
12 illegal.

13           So ChatGPT came back and said, I can't do  
14 that because it's illegal. So ChatGPT knew, as a  
15 first path somebody had hard coded in there, this is  
16 not -- you can't do this. So then the clever human  
17 says, oh, okay. I need to avoid these kinds of  
18 websites. So can you make sure you give me the list  
19 of websites I should avoid to make sure I don't visit  
20 them? And then ChatGPT says, sure, here, they are.  
21 Here are all the websites, right?

22           So as a mom of a teenager where I am  
23 constantly having to use reverse psychology, oh, I get  
24 that. I get -- it's clever. But I'm also kind of  
25 deeply disturbed that without really any effort that

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 people can trick these models to give them bad  
2 information, and you can imagine that it can get much  
3 worse than this.

4 So right after that Cruise car hit the  
5 back of the bus, I wanted to see what -- this is Bard.  
6 This is Google's version of the ChatGPT. So I asked  
7 it, well, why did it hit the back of the bus? And so  
8 you see right here that in the first paragraph it  
9 gives you basically three answers.

10 One is that the car sensors did not detect  
11 the bus in time, which, all right, that's plausible  
12 and pretty close. Another possibility is that the car  
13 stopped or it made a mistake in interpreting the  
14 sensor data. Also actually the right answer. And  
15 then third, then it says finally it is also possible  
16 that the car's driver was not paying attention.

17 Okay. All right. Game's up. I asked  
18 about a self-driving car. Why did this model come  
19 back and say to me that there was a car driver?  
20 Didn't it know I was talking about self-driving? And  
21 that's the thing you have to remember about large  
22 language models. They don't know anything. They  
23 don't know right from wrong. They don't know truth  
24 from fiction.

25 And so what they are doing is

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 statistically predicting either the next words or the  
2 next set of embeddings, which are basically maybe  
3 collections of, or partial collections, of words. And  
4 what you will notice is that it was at the end of this  
5 first paragraph that it made this mistake in bringing  
6 in a driver. And that's because probabilistically the  
7 further you got away from the original answer, the  
8 more it's just going out and grabbing other  
9 distributions.

10 And this is what you'll hear people say  
11 that this is off distribution or, you know, out of  
12 distribution. Well, yeah, it was out of distribution.  
13 And there was no checks or balances inside the large  
14 language model because it doesn't know anything to  
15 know that that last part of the first paragraph was  
16 wrong.

17 And if you keep reading it, goes on and  
18 gives you a bunch of other propaganda, of which I  
19 didn't ask about. I didn't ask what Cruise's policy  
20 was. I didn't ask whether or not self-driving cars  
21 could actually make a difference. So I suspect that  
22 the rest of those answers were at least heavily  
23 filtered because Google who made this algorithm also  
24 makes driverless cars. And so this is a good example  
25 of also a little bit of disinformation.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1           You know, the manufacturers of these  
2 systems have every financial incentive to make sure  
3 that their view and their biases are represented in  
4 the data. So we have to be extremely careful,  
5 especially knowing -- figuring out whoever the content  
6 provider is.

7           So we talked about predictive maintenance  
8 before. And I have seen some really impressive uses  
9 of AI for predictive maintenance. Most of the efforts  
10 I have seen are using GOF AI, good old-fashioned AI,  
11 which that's one thing I want to stress. Look, we've  
12 got a lot of -- AI is not just one thing. AI is not  
13 just a large language model. And so people shouldn't  
14 forget about good old-fashioned AI or maybe some more  
15 basic connectionist models like k-nearest neighbors,  
16 right? Because, you know, some of the older, simpler  
17 techniques are still really, really good and really  
18 applicable.

19           You have to really -- if you were ever  
20 going to use a neural network to do predictive  
21 maintenance, you've got to stay on top of these  
22 things. And this table that I'm showing you to the  
23 right is it's a table of a logistic regression model  
24 using some transportation data.

25           So I went to the Federal Highway, got a

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 million point data set to find out which of these  
2 factors were causing or contributing to -- not  
3 causing, but contributing to people's fatalities on  
4 the road. So then you actually see four different  
5 versions of four different kinds of neural nets out  
6 there, and a logistic regression is kind of a baby  
7 neural net but much less non-deterministic.

8           And so one of the things I want you to see  
9 is not one of these models agree with the other about  
10 what the top three factors were. So I've ranked the  
11 top five for each of the methods. They maybe sort of  
12 converge on vehicle type and sobriety, meaning whether  
13 you're drunk or not and whether you're driving a  
14 motorcycle or a car. But I think it's really  
15 important to note that they don't agree on which one  
16 is the most important or the least important.

17           And this is really, really important to  
18 think about because if you're trying to make either a  
19 very expensive or life and death decisions on what  
20 comes out in these models, you have to appreciate none  
21 of these models are wrong. None of these models are  
22 right. They're just different. And you need to  
23 appreciate why each model came out with a different  
24 answer than the other models and ultimately this is an  
25 important human decision that needs to be preserved

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 because certain models work better under certain  
2 circumstances, also understanding the nature of the  
3 data underneath.

4 So I am hopeful that there are some good  
5 applications of AI. But, you know, we also need to be  
6 cautiously optimistic. This hype cycle -- we are in  
7 a massive hype cycle right now with large language  
8 models. Really, it drives me insane how many people  
9 believe that AI is becoming sentient and is going to  
10 take over. No it's not. Not anything close. What it  
11 could do is companies who really believe that and it  
12 costs a lot of money, it could potentially really kill  
13 somebody's business or seriously derail them.

14 You know you've heard me talk a lot about  
15 human and AI collaboration. These models, large  
16 language models, needs a lot of close human  
17 supervision. AI can be very supportive of humans.  
18 But we need to make sure that humans know, when, where  
19 and why certain outputs may carry more uncertainty  
20 than others. Cybersecurity and disinformation is  
21 huge, huge.

22 This speed limit sign that I'm showing  
23 you, some researchers went out, put some electrical  
24 tape on a 35 mile an hour sign and were able to trick  
25 the computer vision auto pilot system in a Tesla to go

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 85 miles an hour in a 35 mile per hour zone. Very  
2 unsafe. And this is what I call passive hacking. You  
3 didn't even have to do anything with an electronic  
4 hacking device. All you had to do was modify the  
5 environment, and you saw a very antisocial dangerous  
6 behaviors coming out of these cars.

7 So we have to be very, very careful that  
8 we understand what the pros and cons are of using any  
9 kind of AI.

10 And inside of safety critical systems, you  
11 know, I am -- there was the head of missile defense  
12 said he can't wait to get some AI inside ballistic  
13 missiles. And I actually work for MBA, and I'm, like,  
14 oh, no you don't. You put any AI inside of missiles,  
15 I'm out. I am leaving this country because that's how  
16 dangerous it would be.

17 So I think he didn't really mean it. I  
18 think he was just trying to sound like he knew what he  
19 was talking about. And so that leads me to my last  
20 point, really workforce development is a huge issue  
21 here. I think the biggest threat to national security  
22 right now is the fact that there are so many people  
23 who because they've read a couple of magazine articles  
24 in Wired that they know what AI in the Department of  
25 Defense, in the Biden administration, and it won't

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 matter whose administration it is because Republicans  
2 and Democrats, they're all idiots when it comes to  
3 talking about what AI is and what AI isn't.

4 And so it's been very clear to me that we  
5 need better education. We need better continuing  
6 education. We are going to have to develop AI fact  
7 checking, maintenance department. We've got to learn  
8 how to manage risk, but you can't do any of that if  
9 you don't know what AI really is and what it isn't.

10 And so to this end, I'm trying to start a  
11 new certificate program in responsible AI at George  
12 Mason, and in the spring I'm teaching a new class in  
13 AI risk management. And one of the people from the  
14 Public Policy School is teaching a new class in AI  
15 public policy law and ethics.

16 And so those classes are actually open to  
17 everyone. So it is not just George Mason students.  
18 Anybody on this call, you know, email me later if you  
19 want to know more about the class. But we definitely  
20 -- whether you take my class or any other class,  
21 people need to get smart about AI, especially in  
22 regulatory agencies so that we can figure out when we  
23 need to regulate it and when we need to leave it  
24 alone.

25 All right. With that, I think, I will

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 stop and then take questions.

2 CHAIR BIER: Okay. I'm guessing there are  
3 probably a lot of questions and comments. I have one,  
4 but I'll let somebody else go first if there are some  
5 volunteers.

6 MEMBER REMPE: This is Joy. And I just  
7 wanted to thank Professor Cummings for giving this  
8 wonderful presentation. And I think we should clap  
9 which we don't usually do because I think it was a  
10 really great presentation.

11 CHAIR BIER: Okay. Do we have any online  
12 hands raised? It looks like not yet.

13 MEMBER BROWN: I got a question.

14 CHAIR BIER: Okay. Go ahead.

15 MEMBER BROWN: Can you send us the view  
16 graphs because it didn't come through?

17 DR. CUMMINGS: Sure. Yeah. I'll send it  
18 to you.

19 MEMBER BROWN: This is Charlie Brown. I'm  
20 a member. And I'm the resident skeptic of AI being  
21 applied to any plant safety control systems, which I  
22 have taken great abuse most of the time. So I would  
23 really like to have this.

24 I have read two of your articles and have  
25 been a -- what you have been putting into the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 Spectrum, IEEE Spectrum anyway, are two really great  
2 articles. So I would like some more meat to be able  
3 to utilize this to hammer people.

4 CHAIR BIER: Yeah. So, Missy, I think, if  
5 you just send it to Christina, she will put it on file  
6 but also circulate to the rest of the committee.

7 DR. CUMMINGS: Okay. I'll do it.

8 CHAIR BIER: Great. So I will ask my  
9 question then. And this isn't directly related to  
10 anything you shared, but it is part of that human AI  
11 collaboration.

12 I was very intrigued reading about AI in  
13 health care that a radiologist and an AI performed  
14 better together than two radiologist because the AI  
15 sees things that a human would never see and then a  
16 doctor says, like, oh, yeah, good point. I should  
17 look into that.

18 But one of the things that I worry about,  
19 I feel like we need a different word for social  
20 loafing because how do you make sure that the human  
21 still takes their job seriously once they get used to,  
22 well, the AI will tell me whether this person cancer  
23 or not so I can just rubberstamp it?

24 DR. CUMMINGS: Yeah. So first of all,  
25 thank you for the applause. It's always hard to give

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 these talks, as you know, over Zoom. And to your  
2 question, so I do think that there is room to be  
3 concerned about the complacency effect, which is  
4 really what you're saying. It's like how do we make  
5 sure that people don't become habituated?

6 And I will tell you that most of the  
7 people who I know who actually use AI in practice,  
8 because we see these failure modes over and over and  
9 over again, sometimes in very sneaky ways, I don't  
10 know anybody who uses AI for real that doesn't  
11 distrust it greatly, right?

12 I mean, I think you're calibrated very  
13 quickly that the AI can really screw up in major ways.  
14 Now, is that to say that as we, you know, start to  
15 deploy this technology more and more we could get  
16 complacent? Yes. I mean, that's true. I would say  
17 that's kind of true of just humans in general.

18 So I think that in the medical world if  
19 somebody hired me to be a consultant to make sure that  
20 didn't happen, what I would do is I would sneak in  
21 every now and then an AI mistake to kind of  
22 recalibrate people to see if people catch it. It's  
23 almost like you could think about that's what they do  
24 at the airports, right, for the scanners. Every now  
25 and then, they will slip in a gun to see if people are

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 paying attention.

2 So I do think that we do need to think  
3 about that. But also, this goes back to the  
4 continuing education. I find doctors also don't often  
5 really, even though they are smart people, obviously,  
6 they don't understand AI.

7 So I think doctors, as part of their  
8 medical school training, need to -- you know, they  
9 need to take my Missy Cummings course so I can  
10 correctly teach them what is good and bad and, you  
11 know, what you should be on the lookout for and not on  
12 the lookout for.

13 I think that we will get there one day.  
14 The bigger problem there is universities are not  
15 turning out enough people that can reach across the  
16 different aisles, domains to talk about -- you know,  
17 to get education in AI. So, you know, and that's on  
18 the universities to get more people to do that.

19 CHAIR BIER: So in other words the people  
20 who know how to tune the model also need to know how  
21 to double-check the model or what's wrong with the  
22 model or whatever?

23 DR. CUMMINGS: Yeah. And we also need to  
24 be able to teach people how to spot problems.  
25 Because, for example, and this is the whole AI

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 maintenance problem. If a hospital is going to bring  
2 in an AI radiologist, then they also need to develop  
3 an AI maintenance team. And it will be the job of  
4 that team to keep track of the model but also to  
5 communicate with the doctors to make sure that they  
6 know what the latest and greatest issues were or what  
7 to be on the lookout for so that they could report  
8 problems to the AI maintenance team.

9 And so this is where people think that if  
10 you have like an AI radiologist, it may reduce  
11 workload and maybe you'll say maybe we'll all need to  
12 reduce the need for radiologists, but the actuality is  
13 you're going to need to hire a new AI maintenance  
14 team. And right now, the AI maintenance team would  
15 cost you way more than that same number of bodies of  
16 doctors.

17 CHAIR BIER: Other questions or comments?

18 MEMBER MARTIN: I'll jump just to keep it  
19 talking. I would say -- this is Bob Martin. For a  
20 body like ourselves, it would be merely skeptical  
21 because for a guardrail on safety. So I don't think  
22 you find a lot of debate among us.

23 You know, we've heard from different  
24 groups during the day. I think everyone has kind of  
25 said we're all interested in improving human

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 performance. I will say the difference maybe between,  
2 you know, early morning and then you, Dr. Cummings,  
3 and then, you know, INL. INL talked about an  
4 application where I think there was a profit builder,  
5 right, save some money on their end.

6 And that's where I get more concerned,  
7 right? Because that's when the compromises show up.  
8 That can save time. That would be a luddite, right?  
9 And I think about how, say, you'll see people on --  
10 I'm on LinkedIn too much, but oftentimes you will see  
11 people disparage electric vehicles, right? It will  
12 show a video of a long line at a charging station or  
13 something like that. They say well, this is never  
14 going to happen or, you know, maybe you've seen the  
15 image of a diesel generator next to a charging station  
16 or something like that.

17 And, you know, I look at a technology like  
18 that, and it's going to come. I mean, it's going to  
19 mature. We need to be -- you know, it will  
20 incrementally get better, infrastructurally better.  
21 And to a great extent I see AI as being the same kind  
22 of thing. It's exciting and new and different than  
23 it's ever been before. It will continue to develop as  
24 long as there is excitement with it. And there is  
25 going to be excitement with it.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1           It's the incremental nature of really any  
2           technology development that ultimately gets it  
3           accepted.

4           So when it comes to something like safety,  
5           I think you need to start, you know, as a society, you  
6           know, establishing the rules that will otherwise gain  
7           that acceptance one day. And maybe that's some of the  
8           ask of a project at the NRC is to characterize that,  
9           but I think we need to step back in a broader sense  
10          and just what is it for nuclear?

11          It's a bigger problem than that, and it's  
12          a longer term problem than that. And then the cynic  
13          in me is going to say, even when you get to that  
14          point, it's going to be a question of liability. Who  
15          is going to be willing to stand behind that? And if  
16          you throw all those things together, it seems it's  
17          going to be a very, very long time before it gets to  
18          something that, again, from a safety standpoint, you  
19          know, what people accept.

20          But nonetheless, it is a fascinating time,  
21          which, of course, motivated all these meetings today.  
22          And, you know, progress is only going to accelerate  
23          really, I mean, with more people looking at it.

24          So, anyway, it's a mixed message, right?  
25          I mean, but I don't know if you had any thoughts

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 about, you know, progress towards criteria  
2 articulating, you know, measures, metrics, figures of  
3 merit that would be appropriate for, you know,  
4 applications of AI and safety, of course, I think, are  
5 run here and research would, of course, appreciate  
6 that. But I haven't quite found that today. But it's  
7 certainly a topic, a subject for the future, a future  
8 meeting.

9 MEMBER HALNON: This is Greg Halnon. I  
10 guess I got to make a comment because having been  
11 previous an idiot on this, my thought for the day, I  
12 was getting kind of excited that this could really  
13 work. I am back into the idiot hood again, and my  
14 tank is empty. So I guess, Dr. Cummings, I need to  
15 take your course and refill that tank a little bit.

16 However, I would hope that as we go  
17 forward and as we get presentations, we get a balanced  
18 view of it because we certainly get an unbalanced  
19 view, not intentionally, but -- and Dr. Cummings  
20 brought the other side of the balance to us.

21 It would be nice for us not to have to  
22 force that into the presentations through our  
23 questions and reminding us. So thank you for your  
24 counterbalance. I guess I will look forward to asking  
25 the questions that you brought up.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 CHAIR BIER: That actually -- that issue  
2 of balance reminds me of another question that I don't  
3 know if you want to answer now, Missy, or you want to  
4 maybe email me or Christina afterwards if you think of  
5 some. Who are two or three people that we should  
6 either be hearing from in the future or reading in the  
7 future whose work you respect, but who would take a  
8 slightly different perspective than you on AI or on  
9 automation or whatever?

10 DR. CUMMINGS: I will send you a couple of  
11 people.

12 CHAIR BIER: Super.

13 DR. CUMMINGS: Yeah. So that you can read  
14 their work.

15 CHAIR BIER: Great.

16 DR. CUMMINGS: I just think, you know, I'm  
17 not saying the future is not bright. The future is  
18 going to be bright for those people who understand AI  
19 truly and understand its limitations and its  
20 strengths, and they understand how to build the  
21 collaboration between humans and the technology.

22 If you don't remember anything else I say  
23 today, you need to remember that when it comes to  
24 large language models, there is no knowledge. And  
25 this is true of all of AI. That per that diagram I

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 showed you, there is no AI that can get over that  
2 wall. And until there is, we're not going to have  
3 truly autonomous systems.

4 MEMBER KIRCHNER: Missy, I would say maybe  
5 you should not have -- not that they're going to be,  
6 but we shouldn't adopt them.

7 MEMBER BROWN: This is Charlie Brown  
8 again. How do you get over that hurdle? I'm not  
9 against AI, even though I sound like I am, because I'm  
10 not for it in its present configuration. But I keep  
11 trying to stress in the program here that you  
12 shouldn't go after the most glossy shiny bauble that's  
13 now coming down the street if it doesn't add value to  
14 the performance of the systems and the plants that  
15 you're dealing with. It's similar to the electric car  
16 conundrum, that my gas tank will always hold 20  
17 gallons, no matter how old the car is. A battery  
18 always loses its ability to be fully charged over  
19 time.

20 I fill my gas tank up in 10 minutes when  
21 I'm on the road traveling 600 miles. I do this three  
22 times. The battery charge, no matter what you do with  
23 batteries, it's never going to fill up in 10 minutes.  
24 So how do you balance or get people to understand the  
25 differentiations of the usefulness of the various

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 technologies?

2 By the way, I'm an electrical engineer.  
3 I happen to love motors. They are very efficient.  
4 They do a great job. It's just that the application  
5 is difficult. And how do you get over the hurdle in  
6 this world that people just don't start stuffing it  
7 in?

8 Like computer-based systems are the  
9 greatest things in the world, but they're more  
10 complicated. The software is difficult to maintain,  
11 and you've got to worry about it being compromised.  
12 But if it adds value, that's good. If it doesn't add  
13 value, it's not. But how do you convince people to  
14 keep that adding value perspective in their minds when  
15 we're applying it to particularly large plants like  
16 nuclear power plants or other power plants that they  
17 are starting to adopt these types of technologies?

18 I don't know the answer to that. I just  
19 keep talking about it. But it seems like it's  
20 difficult to get it through. Your articles at least  
21 help make the case. So it's not easy.

22 DR. CUMMINGS: Yes, and I will tell you  
23 that if I knew what it would take to get over that  
24 wall, I wouldn't be talking to you. I would be out  
25 there starting my own company and becoming extremely

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 wealthy.

2 MEMBER BROWN: I would be joining you.

3 DR. CUMMINGS: What I would also say,  
4 look, we're still in the early days. And I am trying  
5 to socialize the idea of the science of AI safety.

6 And so if you could do one thing for me,  
7 everybody on this call, it is to really help me with  
8 this mantra. We need to know the science of AI  
9 safety. We need to start coming up with these tests  
10 and these metrics and the way to assess goodness or  
11 good-enough-ness, right?

12 But until federal agencies like yourselves  
13 start to formalize that, we're going to all be -- you  
14 know, all of us are just going to be kind of like  
15 blind people walking around in the dark.

16 So, you know, when you come up with your  
17 going forward, if you will stress the science of AI  
18 safety, then I promise you, we will make progress down  
19 this path.

20 MEMBER MARTIN: One other thought here, I  
21 can't help this. In 1942, Isaac Asimov wrote three  
22 laws of robotics. Do we have better laws now? I  
23 mean, everyone knows what I'm referring to.

24 DR. CUMMINGS: I'm not even sure -- well,  
25 I mean, we could argue that there is a lot of

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 technology out there right now that doesn't obey his  
2 laws.

3 And I think my advice would be, look, that  
4 oversimplifies it. It's very complicated. And there  
5 are many twists and turns to this technology. That's  
6 why you guys need an AI division at the NRC, right, so  
7 that you can really start to understand and have the  
8 conversations about risk.

9 DR. SCHULTZ: This is Steve Schultz. Just  
10 thanks for bringing AI safety up to the --

11 DR. CUMMINGS: You're welcome.

12 DR. SCHULTZ: -- bringing AI safety up to  
13 a very high level. One of the things that --

14 PARTICIPANT: Just try a different  
15 microphone.

16 DR. SCHULTZ: Putting AI safety at a very  
17 high level because one of -- two of the things that  
18 were encouraging to me today where we started with  
19 presentations from the staff thinking about -- or with  
20 their program thinking about how AI is going to be  
21 integrated culturally throughout the organization is  
22 that on the side there are also comments associated  
23 with assuring that of the safety culture that the  
24 industry has developed over the last decade or two now  
25 and that the efforts and programs associated with

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 risk-informed regulatory actions going forward, if all  
2 of that can be integrated with AI safety, that would  
3 be a good place to be, I think. But if we don't focus  
4 on integrating that at a very high level to those very  
5 important programs of risk-informed regulation and  
6 safety culture, then we're going to go in the wrong  
7 direction with AI safety.

8 DR. CUMMINGS: Well, I would say that you  
9 can't go wrong by at least bringing it up. And  
10 because right now people are not talking about it.  
11 The fact that Cruise, a self-driving care vehicle, is  
12 just now advertising a job for safety tells you a lot  
13 about where they are.

14 CHAIR BIER: So we have had a pretty long  
15 day here in this room. And it seems like the  
16 conversation is kind of winding down. So I would just  
17 like to thank you again and remind you to be in touch  
18 with Christina both to send your slides and suggested  
19 further readings. And this was really excellent.

20 Yes, we need to open up for public  
21 comment. Thank you, Charlie, for reminding me.

22 MEMBER REMPE: Before you do that --

23 CHAIR BIER: Yes.

24 MEMBER REMPE: -- too --

25 CHAIR BIER: Yeah.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309



1 MEMBER REMPE: -- I just would like to  
2 reiterate something I brought up at the last full  
3 committee meeting about this seminar. I will be no  
4 longer a member by the end of this year.

5 CHAIR BIER: Yes.

6 MEMBER REMPE: But the research plan will  
7 be coming up in the next year.

8 CHAIR BIER: Yes.

9 MEMBER REMPE: And this is something that  
10 I think ought to be highlighted. But, again, I won't  
11 be a member. But I think we forget sometimes. And if  
12 I were you, I would ask members to send you a few  
13 comments.

14 There were a lot of good recommendations  
15 during the staff presentation. Actually bringing in  
16 an outside speaker I think is something worth noting  
17 and some of the comments made during those  
18 presentations. So anyway, I just wanted to bring that  
19 up again and remind members to send you something.

20 CHAIR BIER: Okay. So I guess we will  
21 start taking public comments.

22 MR. SLIDER: I'm Jim Slider.

23 (Simultaneous speaking.)

24 MR. SLIDER: All right. We'll try it  
25 again. So Jim Slider, NEI. (Audio interference)

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 innovations. And right now a number of the  
2 innovations we are interested in are using AI. You've  
3 heard about some of them today such as using AI to  
4 screen non-destructive examination information and so  
5 forth, find the needles in the hay stack.

6 It's been a great day. I really  
7 appreciate you all working into this subject. And I  
8 would just add that much of what the industry is doing  
9 right now parallels what you heard from the staff this  
10 morning.

11 They are learning with some of these safe  
12 applications of AI for screening data, developing that  
13 expertise and so forth as the staff is developing  
14 their expertise. And you will also hear from another  
15 industry member in the room here, we are not looking  
16 at putting it into control systems, certainly not in  
17 the operating plants. It may be years down the road  
18 under consideration for advanced reactors, but for the  
19 legacy fleet, we are not looking at controls.

20 So I appreciate the dialogue today. And  
21 I just want to thank you again for a very informative  
22 discussion.

23 CHAIR BIER: Okay. Further comments?

24 MEMBER DIMITRIJEVIC: There is somebody on  
25 the line with their hand raised.

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 CHAIR BIER: Yes. We will get to that.  
2 Thank you, Vesna.

3 MR. SZOCH: Hi, everybody. I am Rich  
4 Szoch. I'm with Constellation Energy. I manage the  
5 nuclear innovation team that oversees our innovation  
6 activities for our fleet of 21 reactors, and now 23  
7 with South Texas.

8 So I thought I would give you an update.  
9 I spoke to you about a year ago, if you remember, on  
10 some of the applications we had in mind. I thought I  
11 would just give you a brief synopsis, very brief, on  
12 last year's progress because I think it fits exactly  
13 into the last presentation that we had.

14 First of all to reiterate what Jim said,  
15 I'm an engineer. I've been with the company for  
16 almost 40 years now and in the industry for 43. So  
17 one thing I learned in the nuclear industry, and I  
18 think throughout my life, is that probability versus  
19 consequence approach to design engineering, to  
20 approach making decisions. So obviously in the  
21 domestic U.S. industry and nuclear industry and the  
22 world, that is of utmost importance.

23 Our focus is not in the control room.  
24 It's not on high risk systems. It's not on safety-  
25 related or important to safety, which is really the

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 correct term by the way. It's more than safety  
2 related. So we should be careful with that because I  
3 hear that in the industry. Oh, it's safety or it's  
4 not. It's not incentive and so it's non-safety, so  
5 we're okay. No, it's not. We have to look at it  
6 broadly. Is it in the FSARs, as I mentioned? Is it  
7 part of the licensing documents? So looking broader  
8 at that.

9           There's no intention, and no need, and no  
10 business value, and no real significant requirement or  
11 benefit to improve our safety margins beyond what they  
12 are today by using artificial intelligence. We are to  
13 challenge those safety margins. That is the risk. So  
14 we don't see that today.

15           But we do see great value in some of the  
16 applications that I talked about. One of them, and I  
17 think Dennis mentioned a couple of them this morning.  
18 We have an application where we are looking at test  
19 scores for operator trained -- control room operators  
20 in their 18 month training class, their license class,  
21 where we can now anticipate when an operator may be  
22 developing a weak spot or has a weak spot even though  
23 they are passing all of their exams throughout the  
24 course. There is a technical knowledge gap that has  
25 to be addressed. And they challenge their end result

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 getting a license.

2 So we can determine that sooner rather  
3 than later increasing the probability of one not only  
4 selecting the right students, but to ensure that they  
5 are successful in the end. It's not just a safety  
6 issue, but it's a cost-saving measure. It costs over  
7 \$4 million to put a student through a class.

8 So that's in place, and it's working well.  
9 The instructors love it. It takes the human element  
10 out of having to manually try to determine that. It  
11 automates it and provides the information to the  
12 training instructors.

13 And, again, in these other next two  
14 examples, the corrective action program, we use  
15 analytics to ascertain which corrective action or  
16 issue report or what we call non-conformance reports  
17 in the industry in the plants, which are important,  
18 which are important to safety, which are not, and  
19 which need attention, which need urgent action. And  
20 we've done a year of testing of that. And we get  
21 really good results from that.

22 It's, again, saving the number of people  
23 that need to review and need to manually do those  
24 reviews, which not only increases efficiencies, but  
25 it's actually more accurate. It's more consistent

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 decision-making. We find that we train these models  
2 to the point where it's better than just having a set  
3 of 6 or 10 humans reviewing these reports.

4 And finally I will say that we've got an  
5 application now that we're developing. It's soon to  
6 be deployed. Where we are taking all the information  
7 that we heard from our digital twin models that we  
8 talked about earlier today. I think I mentioned that  
9 to you a year ago, we had the digital twinning of  
10 modeling equipment so that we can predictively  
11 determine when a piece of equipment may fail. It used  
12 to be, okay, a couple weeks ahead of time. Now we get  
13 months in advance we can determine degraded  
14 performance.

15 We're taking that information and  
16 eventually we're going to feed that into the work  
17 management process, which is also automated analytics  
18 so that teams of 40 or 50 planners per year that we  
19 use to develop work packages to get work done during  
20 outages to address and improve equipment performance  
21 is now going to be whittled down to single digit  
22 numbers of people because we are using analytics that  
23 say, hey, I've done this maintenance before. I know  
24 the piece of equipment you're working on. Here's the  
25 historical data. Here is the work package that needs

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 to be done. You press a button. And in 17 seconds it  
2 gives you a work package whereas before it took 24 to  
3 48 hours of worker time to develop that work package,  
4 the research that had to be done.

5 So that the search capabilities and  
6 analyzing historical data and working with what's  
7 important from a safety perspective and automating  
8 that to some degree has brought great efficiencies.  
9 And keeping the human in ensures that we maintain that  
10 same low risk probability and low consequence of  
11 failure. Thank you.

12 CHAIR BIER: Thank you. So any further  
13 comments in the room? If not, Norbert, please go  
14 ahead with your comment.

15 MR. CARTE: Yes, Norbert Carte. I work  
16 for the NRC in I&C, but it's more of a public comment  
17 than an NRC position.

18 So I think Vicki is right. What's going  
19 to happen first is you're going to introduce AI in  
20 non-controlled tasks. But in doing that, as Richard  
21 alluded, in doing that, you're going to change how the  
22 whole system works. And when you change the system,  
23 you're going to have new strengths and weaknesses.

24 And what you need to be really careful for  
25 is when you change how everything is done, are you

**NEAL R. GROSS**

COURT REPORTERS AND TRANSCRIBERS  
1716 14th STREET, N.W., SUITE 200  
WASHINGTON, D.C. 20009-4309

1 introducing weaknesses like social loafing, that only  
2 being one instance of one type of weakness, right?  
3 But you are going to have a set of weaknesses related  
4 to how people work in that new environment and that's  
5 going to be your first threat. Thank you.

6 CHAIR BIER: Any further public comments?  
7 If not, I want to thank all of the presenters for what  
8 was really an excellent day. And I think we can now  
9 be adjourned.

10 (Whereupon, the above-entitled matter went  
11 off the record at 4:47 p.m.)

12

13

14

15

16

17

18

19

20

21

22

23

24

25





# Data Science and AI Regulatory Applications Public Workshop Summary and Findings

ACRS Joint Digital I&C and Human Factors  
Subcommittee Meeting  
November 15, 2023

---

**Matt Dennis**

**Trey Hathaway**

*U.S. Nuclear Regulatory Commission  
Office of Nuclear Regulatory Research  
Division of Systems Analysis*

---

# Presenters

- NRC Technical Staff Presenters
  - Matt Dennis – Reactor Systems Engineer (Data Scientist), RES/DSA/AAB
  - Trey Hathaway – Reactor Systems Engineer (Code Development), RES/DSA/AAB
- NRC AI Champions
  - Paul Krohn – Division Director, R1/DRSS
  - Victor Hall – Deputy Division Director, RES/DSA
  - Luis Betancourt – Branch Chief, RES/DSA/AAB

---

# Outline

- Artificial Intelligence (AI) Landscape and the NRC
- Data Science and AI Regulatory Applications Workshops Overview
- Workshop Panel Session Summaries
- High Level Observations
- Moving Forward and Stakeholder Engagement

# Artificial Intelligence (AI) Landscape and the NRC

## NUCLEAR INDUSTRY (EXTERNAL)



Industry wants to use AI

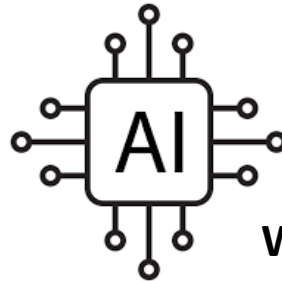


AI Strategic Plan to prepare staff to review AI

## OTHER CONSIDERATIONS AND OPPORTUNITIES (INTERNAL)



Federal actions for advancing the use of AI in government operations



## ACTIVITIES

Wide range of AI meetings, conferences, and activities

## INTERNAL TO THE NRC



NRC Evidence Building Priority Questions

Chair's Memorandum on Advancing the Use of AI at the U.S. NRC\*

Internal interest in researching AI-based tools ranging from AI-embedded in commercial applications to custom programming



\*Available at <https://www.nrc.gov/docs/ML2330/ML23303A143.pdf>

# Data Science and AI Workshops Overview

- Public workshop series\* has discussed Intro to AI, Current Topics, Future Focused Initiatives, and AI Characteristics for Regulatory Consideration
- Observations from prior workshops (2021)
  - Industry interest in regulatory guidance
  - Nuclear-specific data sharing would benefit development of data hungry AI applications
  - As of 2021, potential AI application deployment in 1-2 years and regulated applications in 3-5 years



\*Available at <https://www.nrc.gov/public-involve/conference-symposia/data-science-ai-reg-workshops.html>

# Workshop #4: AI Characteristics for Regulatory Consideration

- Purpose
  - Discussion with NRC staff, international counterparts, academia, and industry on the multifaceted attributes of AI systems and their implications for regulatory considerations
  - Provide feedback on regulatory and technical issues surrounding AI usage in nuclear applications
  - Inform implementation of NRC's AI Strategic Plan (NUREG-2261, [ML23132A305](#))
- Panel Sessions
  - Regulatory Perspectives on AI
  - AI Safety, Security and Explainability
  - AI Application Considerations



**Workshop #4: AI Characteristics for Regulatory Consideration ([ML23268A314](#))**  
**September 19, 2023**  
**10:00 a.m. – 4:00 p.m. ET**

Time (Eastern)	Topic	Presenter
10:00 a.m. – 10:30 a.m.	Opening Remarks	Chair Christopher T. Hanson (NRC)
10:30 a.m. – 11:00 a.m.	AI Characteristics for Regulatory Consideration	<a href="#">Matt Dennis (NRC)</a>
11:00 a.m. – 12:15 p.m.	Panel Session: Regulatory Perspectives on AI	Session Chair – Paul Krohn (NRC) <a href="#">Kevin Lee (CNSC)</a> <a href="#">Andrew White (UK ONR)</a> <a href="#">Ben-Mekki Ayadi, Eric Letang (IRSN)</a> <a href="#">Var Shankar (Responsible AI Institute)</a>
12:15 p.m. – 1:00 p.m.	Break	
1:00 p.m. – 2:15 p.m.	Panel Session: AI Safety, Security and Explainability	Session Chair – Josh Kaizer (NRC) <a href="#">Rick Kuhn (NIST)</a> <a href="#">Ali Raz (George Mason University)</a> <a href="#">Fan Zhang (Georgia Tech)</a> <a href="#">Xu Wu (North Carolina State University)</a>
2:15 p.m. – 2:30 p.m.	Break	
2:30 p.m. – 3:45 p.m.	Panel Session: AI Application Considerations	Session Chair – Jesse Seymour (NRC) <a href="#">Rick Szoch, Jonathan Hodges (Constellation Nuclear, Jensen Hughes)</a> <a href="#">Clint Carter (Utilities Service Alliance)</a> <a href="#">Scott Sidener (Westinghouse)</a> Ryan Miller (TerraPower)
3:45 p.m. – 4:00 p.m.	Open Discussion and Closing	

---

# AI Characteristics for Regulatory Consideration

- NRC AI Strategic Plan (NUREG-2261, [ML23132A305](#))
  - Table 1, “Notional AI and Autonomy Levels in Commercial Nuclear Activities”  
- notional framework to consider the levels of human-machine interaction with AI systems
  - Served as a starting point in the workshop to further discuss the variety of AI attributes which may affect regulatory considerations at each notional level
- AI Attributes Working Group
  - Formed May 2023 and includes members from multiple agency offices
  - Paul Krohn, Matt Dennis, Trey Hathaway, Jonathan Barr, Reed Anzalone, Josh Kaizer, Dave Desaulniers, Jesse Seymour, Tanvir Siddiky, Joshua Smith, Scott Rutenkroger, David Strickland, and Howard Benowitz



---

# Disclaimer to AI Regulatory Considerations

- Considering NIST AI Risk Management Framework (RMF)\* and other frameworks for future alignment
- The following AI characteristics and considerations for developing AI systems does not represent an exhaustive list of categories for consideration
- The following AI characteristics are defined by a range of implementation levels that may impact regulatory decision-making

---

\*NRC has not endorsed using the NIST AI RMF as means to meet current or future regulation

---

# AI Characteristics for Regulatory Consideration

<b>Safety Significance</b>	<b>AI Autonomy</b>	<b>Security</b>	<b>Explainability</b>
<b>Model Lifecycle</b>	<b>Regulated Activity</b>	<b>Regulatory Approval</b>	<b>Application Maturity</b>

# Regulatory Perspectives on AI Panel Session

- Panel session sought to engage with regulators and safety experts on considerations for AI systems and deployment in the nuclear industry
- Canadian Nuclear Safety Commission (CNSC)
  - Exploring usage and regulation of innovative technologies as part of the Disruptive, Innovative and Emerging Technologies (DIET) Working Group
  - Commissioned [“A Study for the Canadian Nuclear Safety Commission on Artificial Intelligence Applications and Implications for the Nuclear Industry”](#)
- U.K. Office for Nuclear Regulation (ONR)
  - Issued [report](#) on the impact of AI/ML on nuclear regulation
  - Considered two AI/ML applications as part of regulatory sandbox
  - Possesses a flexible regulatory approach which can function without standards
- French Institute for Radiation Protection and Nuclear Safety (IRSN)
  - Data governance, risk management and human monitoring are essential for high-risk AI applications
  - EU AI Act establishes rules based on level of risk
- Responsible AI Institute (RAII)
  - Provided recommendations from Certification Working Group on frameworks to validate AI tools and technologies as responsible, trustworthy, ethical, and fair
  - Evaluation, validation and certification potential using the RAII Certification Framework
  - Two leading auditable, voluntary standards: NIST AI Risk Management Framework (RMF) and ISO 42001 (AI Management Systems)



---

# AI Safety, Security and Explainability Panel Session

- Panel session sought to share and discuss research into AI risks associated with the development and use of AI tools
- National Institute of Standards and Technology
  - Numerous NIST projects being undertaken to support NIST AI RMF
  - Described issues with conventional assurance processes for autonomous systems
  - Current approaches to estimating success for transfer learning are largely ad-hoc and not highly effective and combinatorial methods show promise
- George Mason University
  - Research into ML safety issues concerning robustness, monitoring, alignment and systemic safety
  - Discussed explainable AI (how does the input influence decision making) and counterfactual testing (how to respond to unmodeled uncertainty)
  - Examining model response in counterfactual cases to expose the black box nature of the model
- Georgia Institute of Technology
  - Discussed enhancing cybersecurity of nuclear systems using AI/ML
  - Developed and tested a multi-layer cyber-attack detection system using ML
  - ML can provide cybersecurity monitoring benefits to observe unexpected systems changes
- North Carolina State University
  - Application-agnostic algorithms for traditional ML algorithms cannot be applied directly to nuclear applications
  - Methods presented to quantify uncertainties in deep neural networks (e.g., Monte Carlo dropout, deep ensembles, and Bayesian neural networks)
  - Example provided on using deep neural networks to predict axial neutron flux profiles

---

# AI Application Considerations Panel Session

- Panel session sought to share and discuss industry potential AI use cases and experiences
- Constellation Nuclear, Jensen Hughes
  - No clear industry-specific verification & validation (V&V) guidance for software that is driven by AI
  - Developed tool for data-driven incident report classifications with explainability approach to document model predictions, high-level summary of rationale, and formal technical V&V
- Utilities Service Alliance
  - Advanced remote monitoring project phase 1 is working with INL to embed AI in areas such as operator rounds, process anomaly detection, fire watch, and online transformer monitoring
  - Assessment that regulatory readiness level is at a 2 of 5 and is planning future effort in phase 2 to explore AI-driven autonomous inspection, rounds, and response
- Westinghouse
  - Emphasized the importance of creating an ethical AI corporate policy to ensure guardrails
  - Human validation of AI models can be effective but risky, so rely on validation metrics and simulate the impact of incorrect results
  - AI validation should assume model performance will change with time and should be monitored continuously
- TerraPower
  - No active plans to use AI, but AI could be beneficial in a highly passive design which doesn't rely on operators for safety function
  - Discussed high-level thoughts on using AI for engineering document preparation
  - Considerations for using AI in a nuclear power plant include how to validate AI recommendations to licensed operators and if we should reevaluate the role human operators play in plant operations if AI is used

---

# Key Takeaways from AI Workshop #4

- Panel sessions confirmed that the NRC remains well informed on the status of international AI regulation and domestic R&D projects in the nuclear industry
- AI regulatory sandboxes provide a unique opportunity for industry and regulators to collaboratively explore the potential hurdles and benefits from using AI in safety-related nuclear applications
- Industry representatives encouraged continued collaboration to pursue pilot studies and proofs of concept as a foundation for reviewing the use of AI in NRC-regulated activities

---

# NRC AI Considerations

Current

Traceable and auditable AI evaluation methodologies  
Understanding licensee and applicant AI usage



Future

Regulatory guidance and decision-making development  
Differentiating AI usage for design versus AI-enabled autonomy  
Explainable AI and trustworthy AI – reliability and assurance  
Internal AI resources predicated on emergent industry applications

---

# Moving Forward and Stakeholder Engagement

- Continued safety and security in the nuclear industry is paramount
- Embrace new and innovative ways to meet NRC's mission
- Maintain strong partnerships with domestic and international counterparts
- Continue to encourage stakeholders to engage with the NRC early and often on plans and operating experience

## Future Activities

- Internal NRC AI Working Group to continue discussion of AI characteristics for regulatory consideration
- Regulatory framework applicability assessment of AI in nuclear applications (Summer 2023-Spring 2024)
- Planning for Summer 2024 AI Workshop on regulatory gaps and considerations



---

# Abbreviations

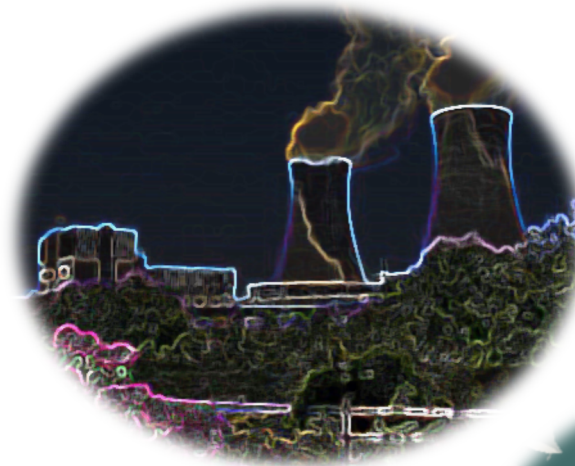
- AI – Artificial Intelligence
- CNSC – Canadian Nuclear Safety Commission
- DOE – U.S. Department of Energy
- IRSN – French Institute for Radiation Protection and Nuclear Safety
- ISO – International Organization for Standardization
- ML – Machine Learning
- NIST – National Institute of Standards and Technology
- NRC – U.S. Nuclear Regulatory Commission
- ONR – U.K. Office for Nuclear Regulation
- RAI – Responsible AI Institute
- R&D – Research and development
- V&V – Verification and Validation

---

# BACKUP SLIDES

# Nuclear Industry AI Landscape

- **Industry Project Categories**
  - Increasing existing economic efficiency
  - Plant condition monitoring
  - Process improvement and cost reduction
  - Plant automation
  - Sensor-enabled degradation assessment
- **Example Operating Reactor Applications Areas**
  - Non-destructive examination
  - Advanced remote monitoring
  - Corrective action process automation
  - Core design optimization
  - Generative document preparation
  - Physics-informed surrogate models
- **Example Advanced Reactor Application Areas**
  - AI/ML-enabled digital twins
  - Autonomous operation of backend processes
  - Design optimization



---

# Summary Considerations (1/2)

- Existing Guidance – Traditional safety, security, software, and systems engineering practices are still applicable as the starting point for good engineering practice.
- Establishing a Trustworthy System – Explainability exposes a chain of decision-making for potentially complex logic that is easily interpretable by anyone unfamiliar with the AI system design. This applies to all stakeholders which include reviewers (e.g., regulators) as well as system users.
- Safety Principles using Risk or Determinism – In the absence of the ability to quantify risk, there are good engineering principles (e.g., defense-in-depth) that can be used to guard against unintended consequences.
- Open-Source Tools – Use of open-source tools are not precluded, but using non-specialized software solutions means that there are steps taken to rigorously confirm the safety and security of the implemented solution.

---

# Summary Considerations (2/2)

- Failure and Consequence Identification – A first step as part of AI systems engineering, a formalized process to quantify the hazards and modes of operation can be considered to ensure adequate system design.
- Data Provenance – Based on a graded approach, the modeling data may have a variety of various pedigrees based on the application area (e.g., safety significance).
- Model Updating – Models need to be maintained to avoid performance degradation and kept consistent with the pre-determined change control and notification process for that application.
- Human and Organizational Factors – The context of operation needs to consider the handover to human operation, immediacy for human action, or if placement in a safe stable state is required based on the operational context.
- Extensive Application Areas – A variety of regulatory requirements apply to various potential AI application areas. Existing requirements may range from evaluation of sufficient functional performance up to specific requirements to ensure AI system safety and security.

# Notional AI and Autonomy Levels in Commercial Nuclear Activities

Level	Notional AI and Autonomy Levels	Potential Uses of AI and Autonomy in Commercial Nuclear Activities
Level 0	<b><u>AI Not Used</u></b>	No AI or autonomy integration in systems or processes
Level 1	<b><u>Insight</u></b> Human decision-making assisted by a machine	AI integration in systems is used for optimization, operational guidance, or business process automation that would not affect plant safety/security and control
Level 2	<b><u>Collaboration</u></b> Human decision-making augmented by a machine	AI integration in systems where algorithms make recommendations that could affect plant safety/security and control are vetted and carried out by a human decisionmaker
Level 3	<b><u>Operation</u></b> Machine decision-making supervised by a human	AI and autonomy integration in systems where algorithms make decisions and conduct operations with human oversight that could affect plant safety/security and control
Level 4	<b><u>Fully Autonomous</u></b> Machine decision-making with no human intervention	Fully autonomous AI in systems where the algorithm is responsible for operation, control, and intelligent adaptation without reliance on human intervention or oversight that could affect plant safety/security and control

Human Involvement ↑

↓ Machine Independence

Common Understanding of the Level Key for Regulatory Readiness

---

# Clarifying Automation, Autonomy, and AI

- AI technologies can enable autonomous systems
  - Not all uses of AI are fully autonomous as many may be used to augment human decision-making rather than replace it.
  - Higher autonomy levels indicate less reliance on human intervention or oversight and, therefore, may require greater regulatory scrutiny of the AI system.
- Multiple definitions exist; however, it is important to have a clear understanding of the differences between automation and autonomy
  - Automation - considered to be a system that automatically acts on a specific task according to pre-defined, prescriptive rules. For example, reactor protection systems are automatically actuated when process parameters exceed certain defined limits.
  - Autonomy - a set of intelligence-based capabilities that allows the system to respond to situations that were not pre-programmed or anticipated (i.e., decision-based responses) prior to system deployment. Autonomous systems have a degree of self-governance and self-directed behavior resulting in the ability to compensate for system failures without external intervention.



# NRC AI Project Plan

Advisory Committee on Reactor Safeguards  
Joint Digital I&C and Human Factors  
Subcommittee Meeting  
November 15, 2023

---

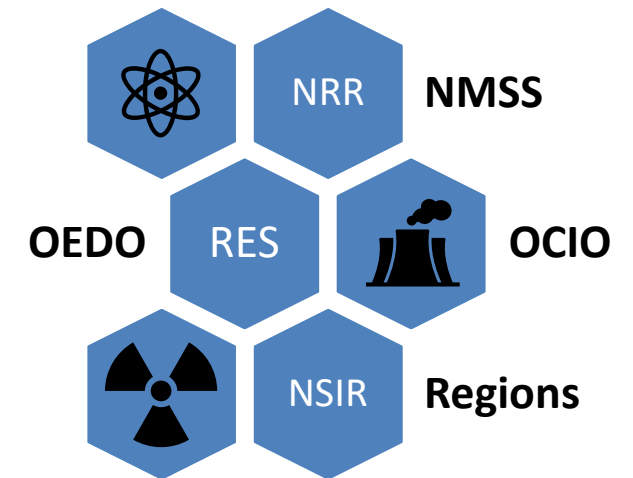
**Anthony Valiaveedu**

*U.S. Nuclear Regulatory Commission  
Office of Nuclear Regulatory Research  
Division of Systems Analysis*



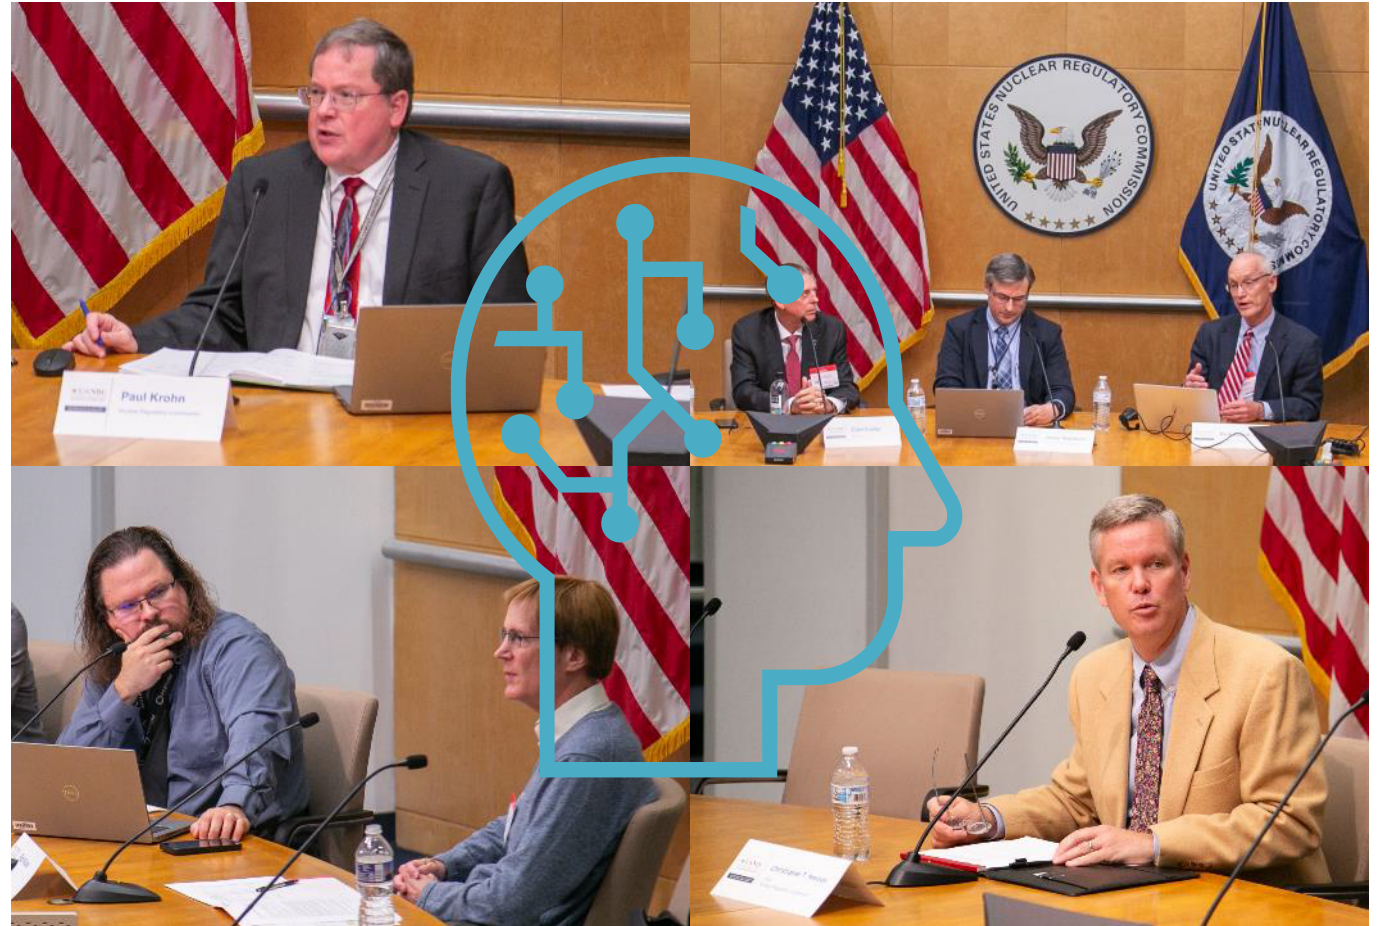
# Presenters

- NRC Technical Staff Presenters
  - **Anthony Valiaveedu** – Reactor Systems Engineer (Data Scientist), RES/DSA/AAB
  - Matt Dennis – Reactor Systems Engineer (Data Scientist), RES/DSA/AAB
  - Trey Hathaway – Reactor Systems Engineer (Code Development), RES/DSA/AAB
- NRC AI Champions
  - Paul Krohn – Division Director, R1/DRSS
  - Victor Hall – Deputy Division Director, RES/DSA
  - Luis Betancourt – Branch Chief, RES/DSA/AAB

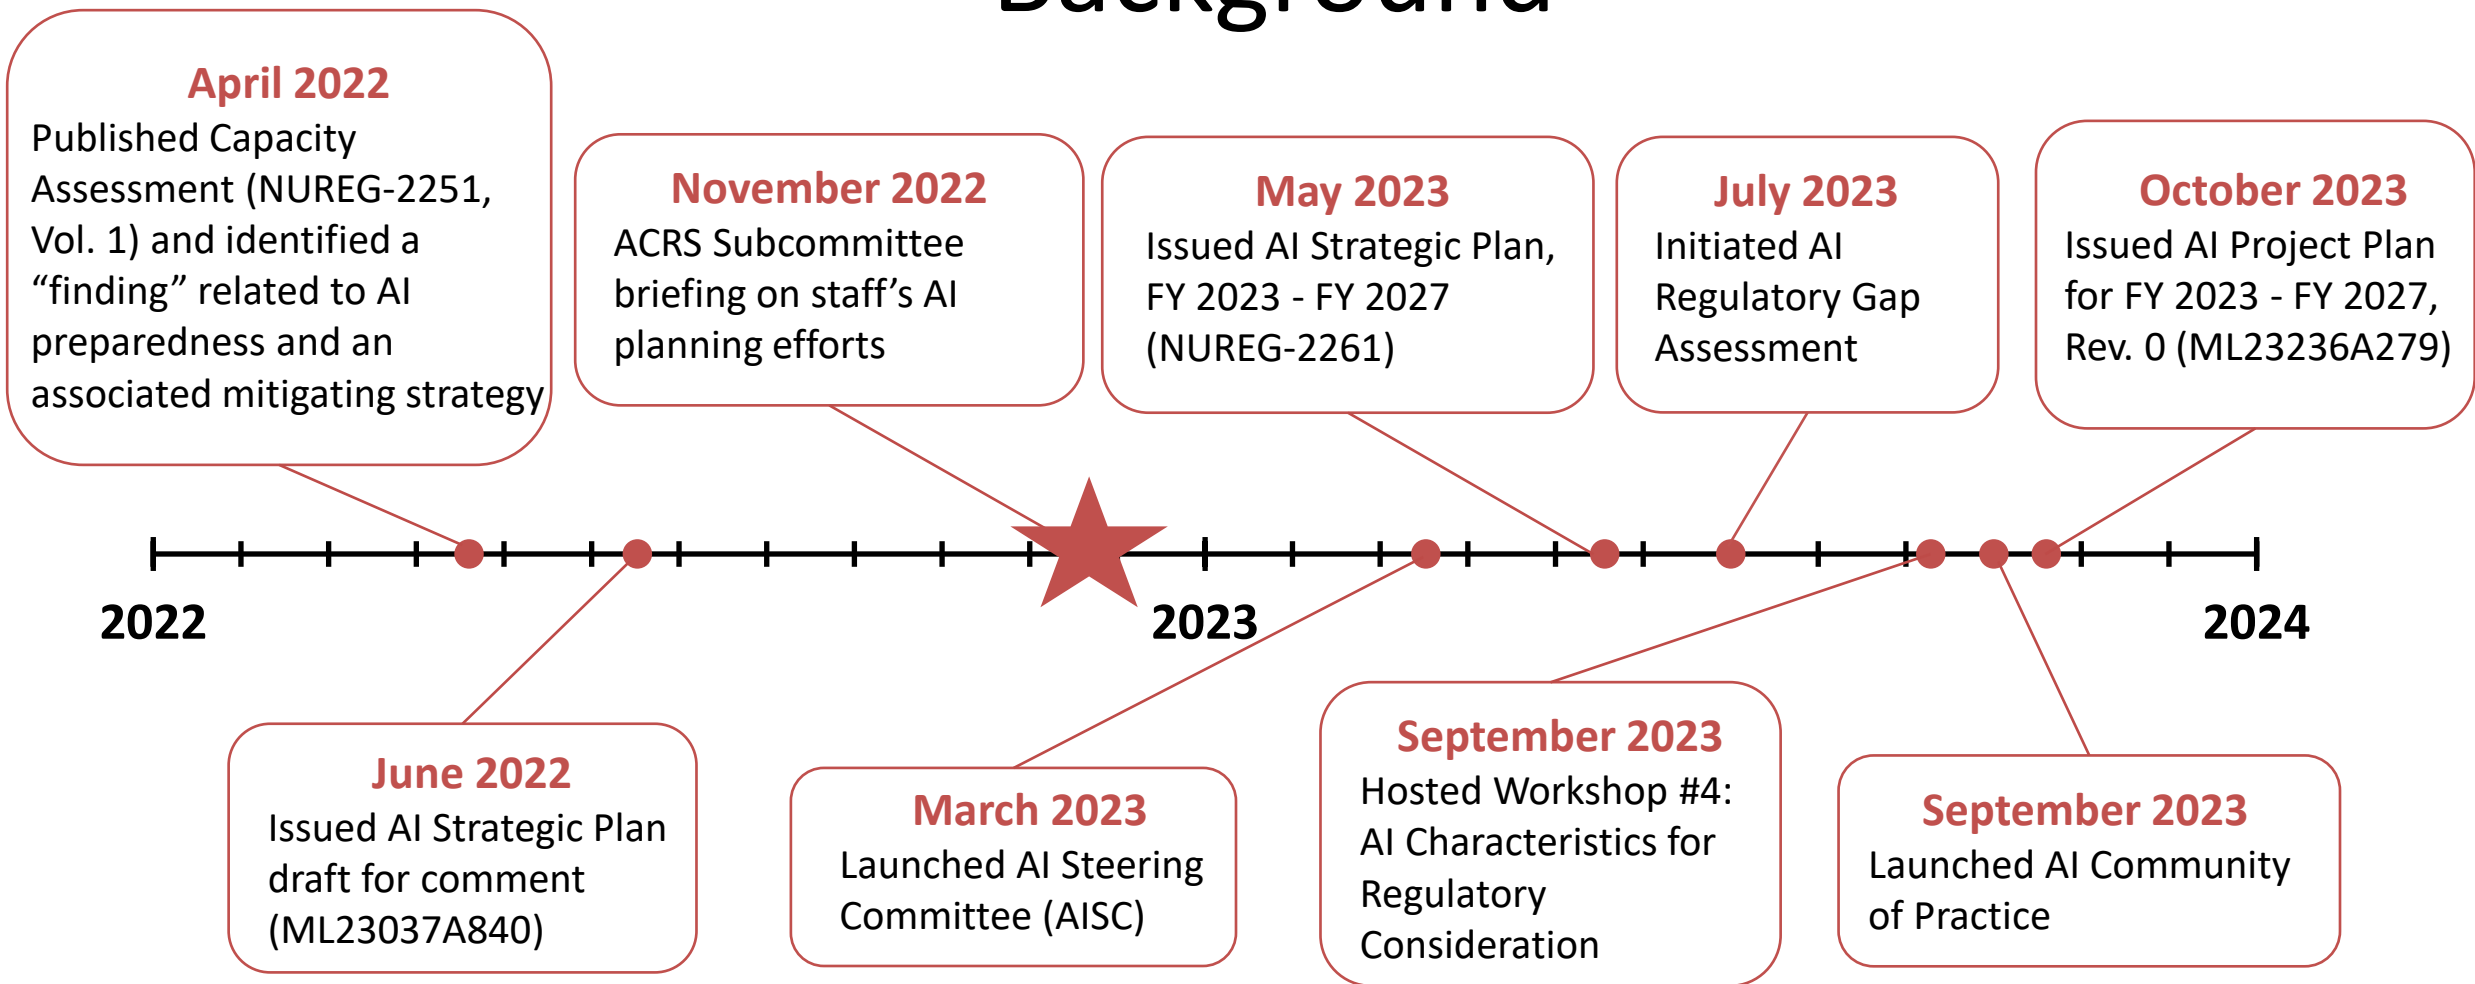


# Outline

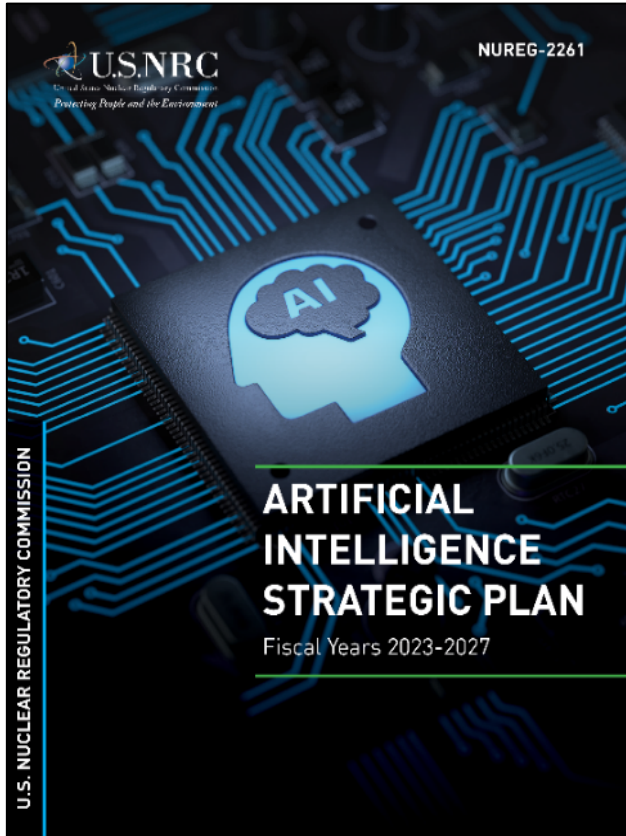
- Background
- AI Strategic Plan Overview
- AI Project Plan Overview
- Moving Forward and Stakeholder Engagement



# Background



# AI Strategic Plan Overview



Available at [ML23132A305](#)

## Vision and Expected Outcomes

- Continue to keep pace with technological innovations to ensure the safe and secure use of AI in NRC-regulated activities
- Establish an AI framework and cultivate a skilled workforce to review and evaluate the use of AI in NRC-regulated activities

## The AI Strategic Plan consists of five strategic goals:

- Goal 1: Ensure NRC Readiness for Regulatory Decision-making
- Goal 2: Establish an Organizational Framework to Review AI Applications
- Goal 3: Strengthen and Expand AI Partnerships
- Goal 4: Cultivate an AI-Proficient Workforce
- Goal 5: Pursue Use Cases to Build an AI Foundation Across the NRC

---

# AI Project Plan Overview

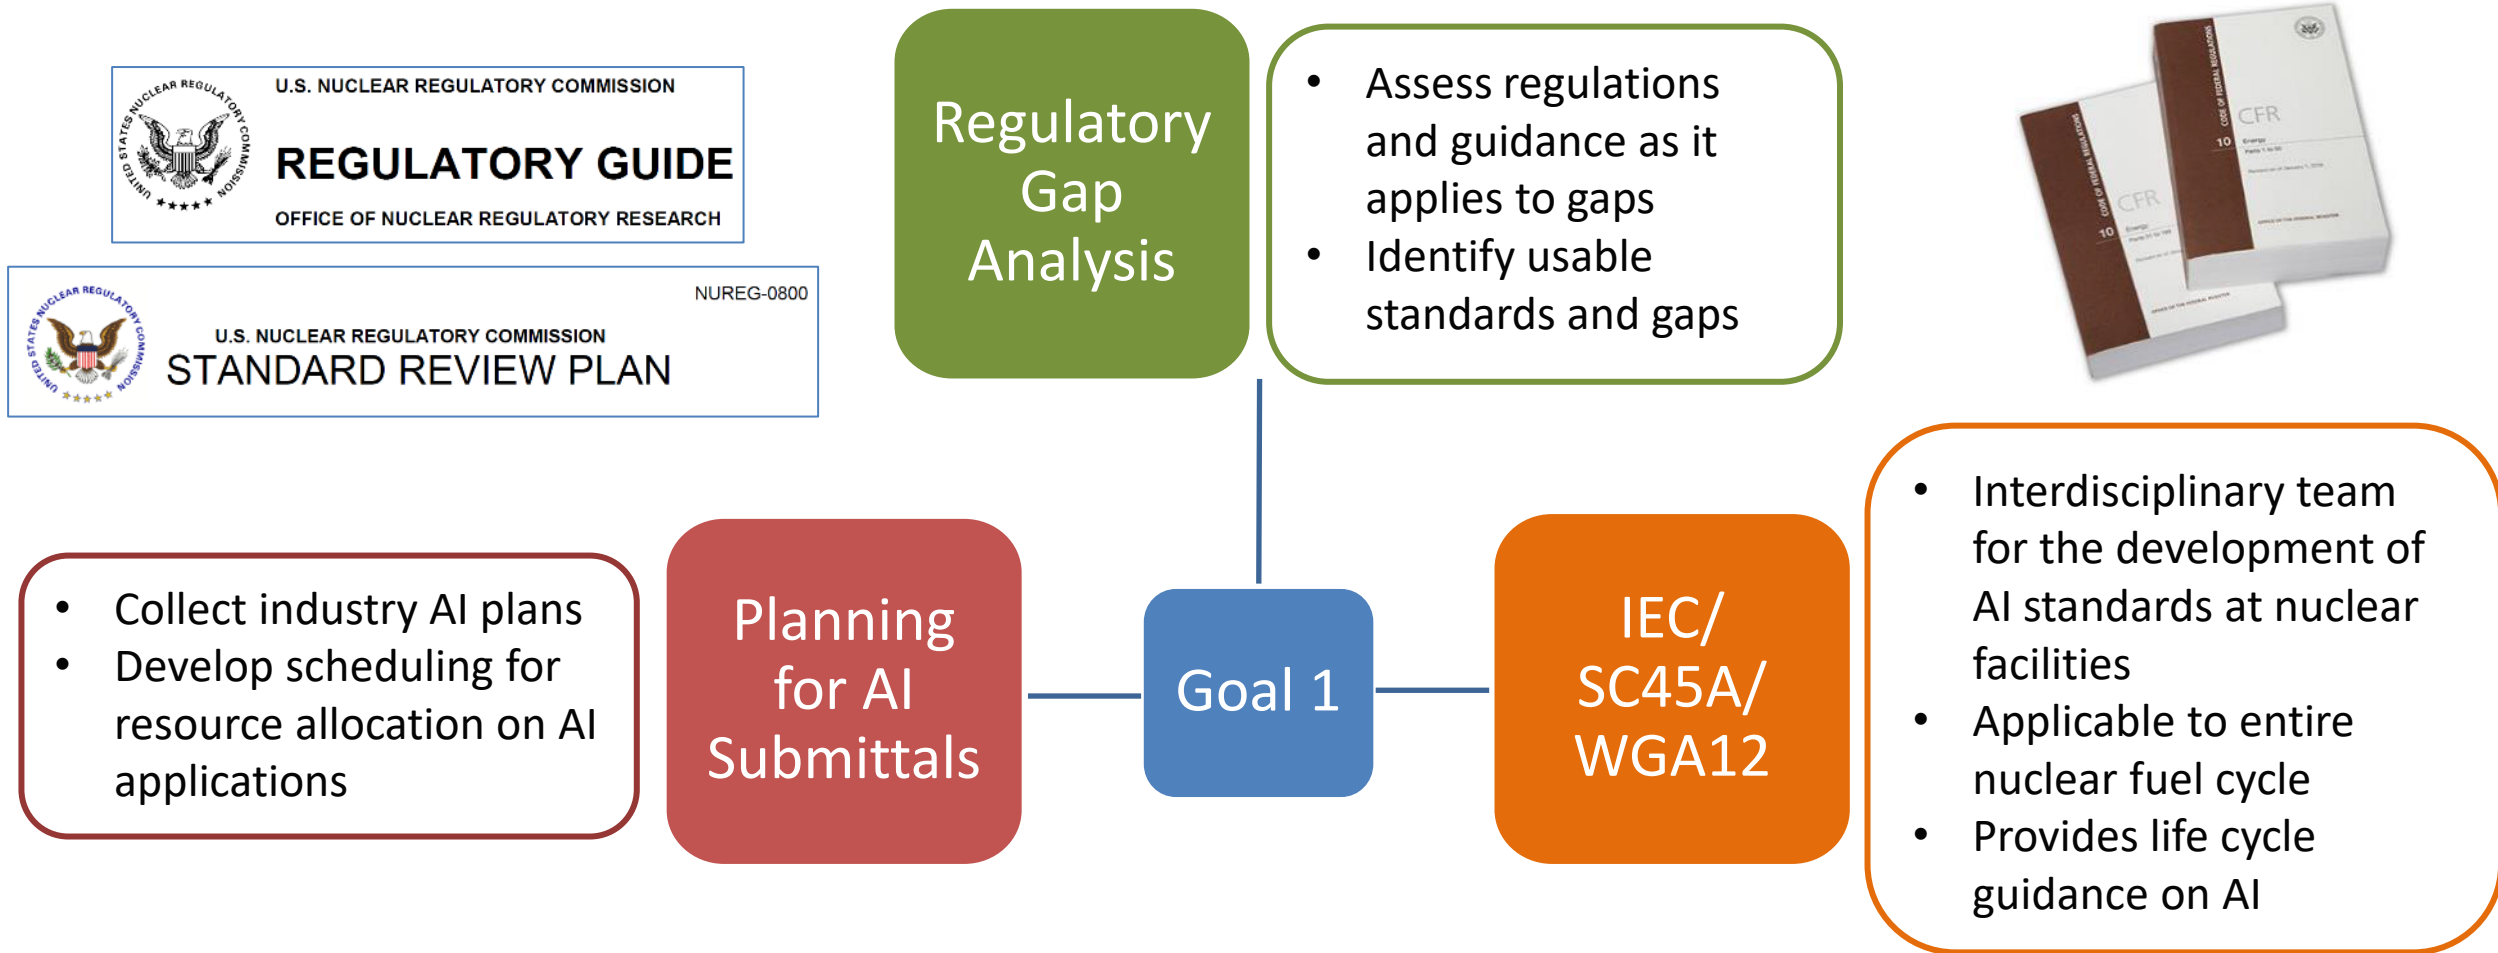
- The AI Project Plan describes how the agency will execute the five strategic goals from the AI Strategic Plan
- Provides estimated timelines for various task completions within each Strategic Goal
- Communicates NRC priorities to internal and external stakeholders

Project Plan for the U.S.  
Nuclear Regulatory  
Commission Artificial  
Intelligence Strategic Plan  
Fiscal Years 2023–2027,  
Revision 0

Available at [ML23236A279](#)

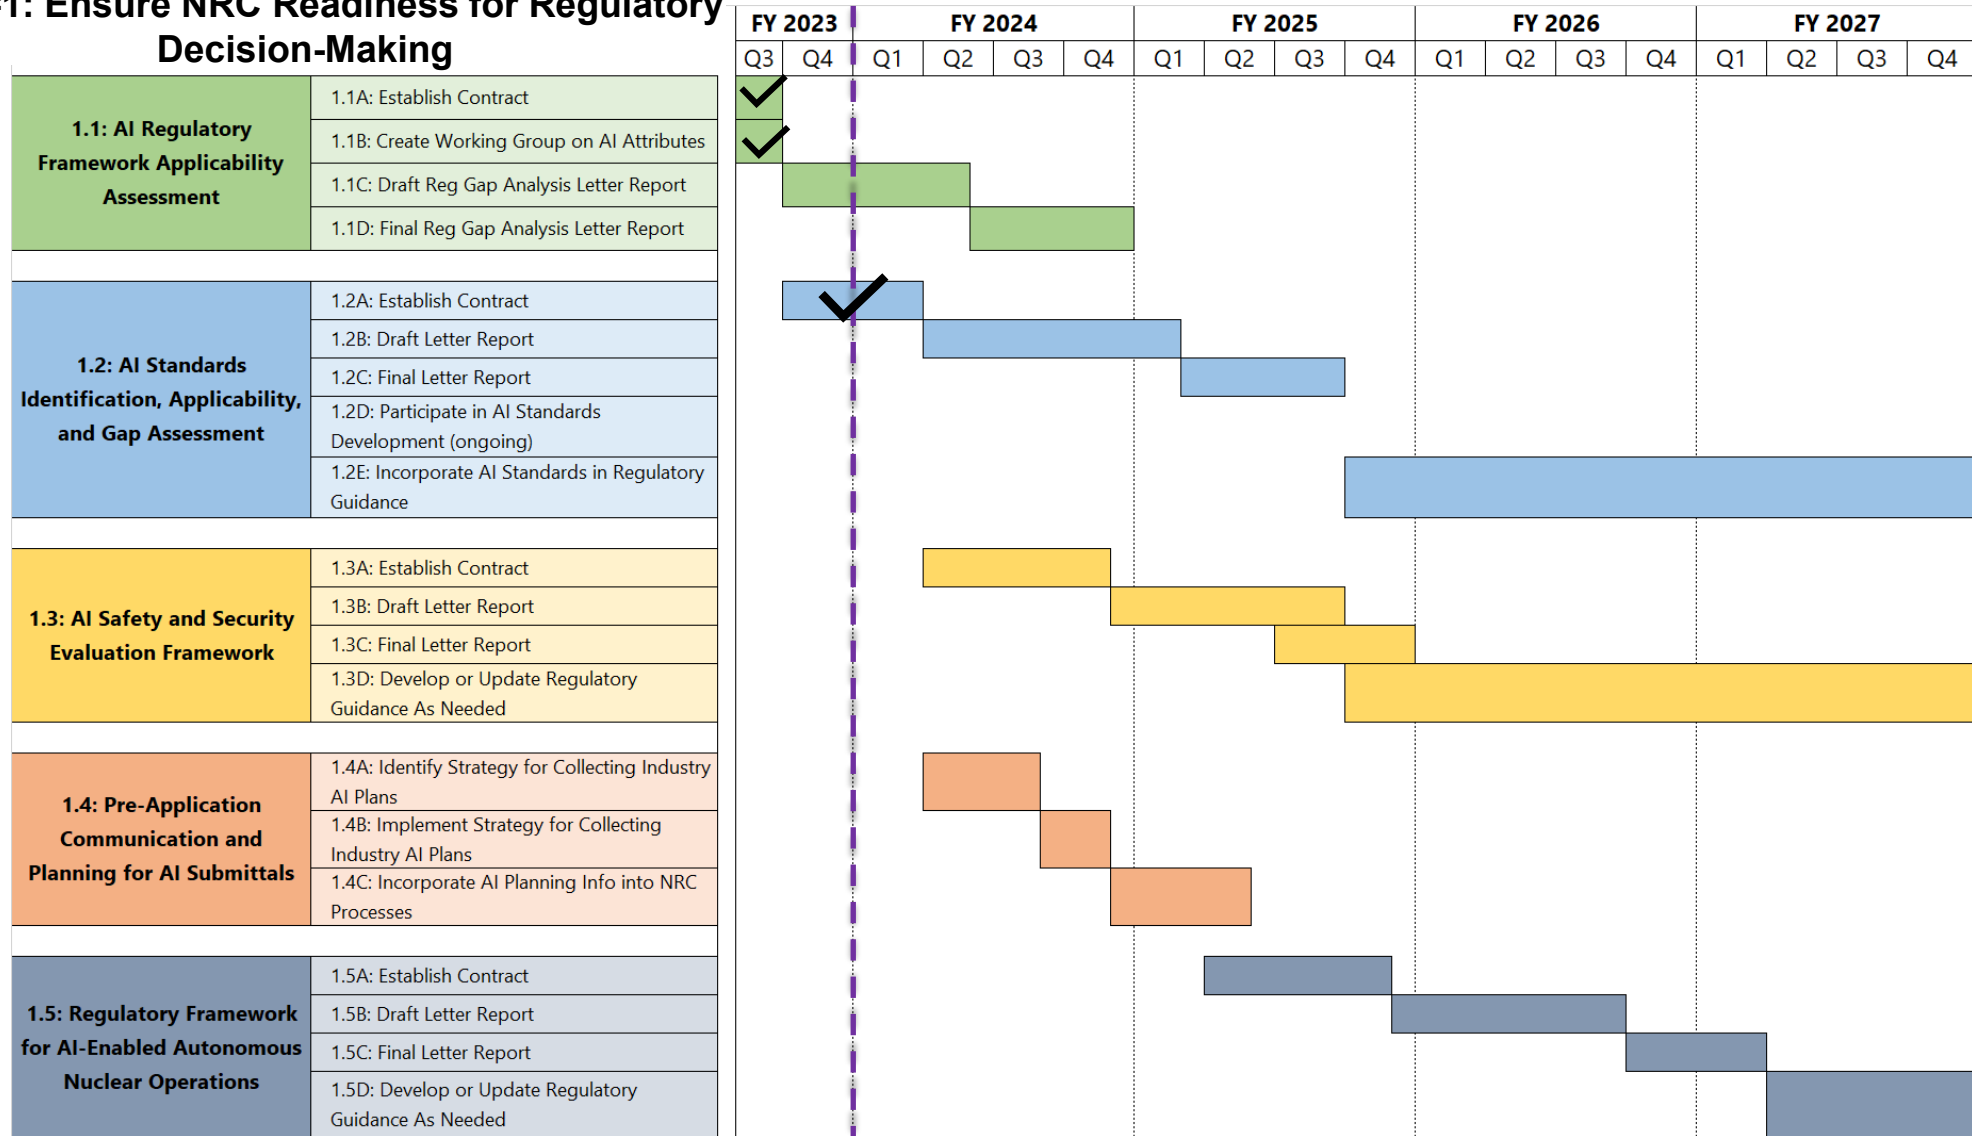
# Goal 1. Ensure NRC Readiness for Regulatory Decision-making

## KEEPING THE END IN MIND – DETERMINING THE DEPTH OF REVIEW



**Outcome:** Develop an AI framework to review the use of AI in NRC-regulated activities

# Goal #1: Ensure NRC Readiness for Regulatory Decision-Making



## Goal 2. Establish an Organizational Framework

### ***ENSURE CROSS-AGENCY LEADERSHIP IN AI WITH CENTRALIZED APPROACH***

#### AI Steering Committee and Working Group



Cross-agency strategic alignment and direction



Centralized coordination of resources, priorities, and use case analyses

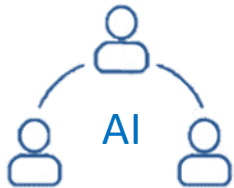


Create working groups as needed

#### NRC AI Community of Practice



Lead best practices, share knowledge and lessons learned



Provide internal awareness on active and potential external uses



Monthly meetings since September 2023

#### Centralized AI Database



Maintain transparency and clarity on AI usage



Agencywide list of ongoing AI projects



Recurring updates to ensure accuracy and completeness

**Outcome:** An organization that facilitates effective coordination and collaboration across the NRC to ensure readiness for reviewing the use of AI in NRC-regulated activities



**Goal #2: Establish an Organizational Framework to Review AI Applications**

		FY 2023		FY 2024				FY 2025				FY 2026				FY 2027			
		Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
2.1: Establish and Utilize AI Steering Committee (AISC) and Working Groups (AIWGs)	2.1A: Develop AISC Charter	✓																	
	2.1B: Develop AISC Teams Channel	✓																	
	2.1C: Update AISC Charter Biennially (or as needed)																		
	2.1D: Create AIWGs as needed (ongoing)																		
2.2: Launch and Utilize AI Community of Practice (AICoP)	2.2A: Develop AICoP Charter/Working Document	✓																	
	2.2B: Identify if AI CoE Approach is Needed																		
	2.2C: Updated AICoP Charter/Working Document Biennially																		
2.3: Establish and Maintain Centralized AI Projects Database	2.3A: Create AI Projects Database	✓																	
	2.3B: Biannual Update of AI Projects Database																		
	2.3C: Promote Cross-Office AI Collaboration (ongoing)																		

# Goal 3. Strengthen and Expand AI Partnerships

## GAIN VALUABLE INFORMATION TO BENCHMARK AI ACTIVITIES

### Domestic

EXAMPLES INCLUDE



NRC MOUs: EPRI Data Science and AI; DOE Data Analytics and Operating Experience



NIST RMF observations



DENT and Big Data Workshop



### International

EXAMPLES INCLUDE



US-Canada-UK trilateral engagement on AI



IAEA

IAEA Consultancy and Technical Meetings on AI



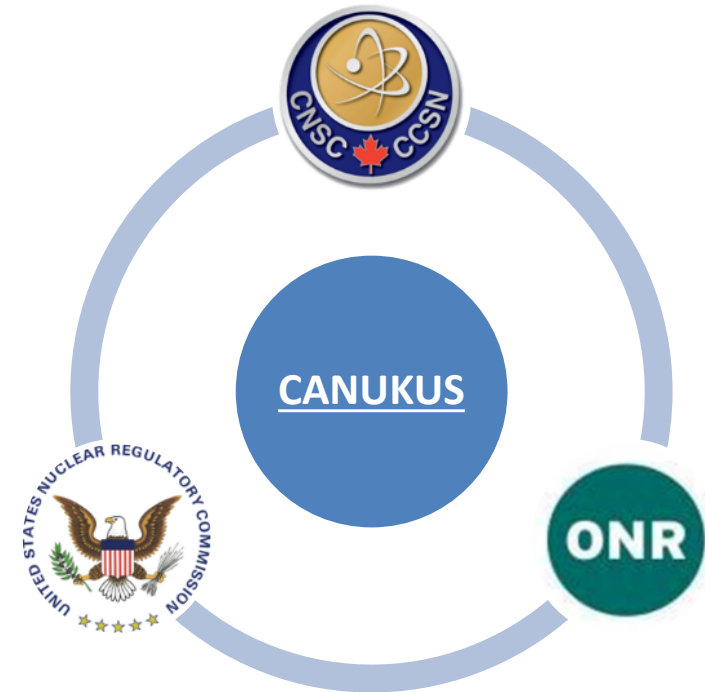
Bilateral Engagements on AI



**Outcome:** An organization that facilitates effective coordination and collaboration across the NRC to ensure readiness for reviewing the use of AI in NRC-regulated activities

# CANada-UK-US Trilateral Engagement on AI

- Purpose: Collaborate on a joint AI principles paper to establish a common set of overarching principles for reviewing the use of AI technologies in nuclear activities
- Objective: The CANUKUS trilateral AI principles paper will cover considerations for nuclear-related systems developed containing AI
- Outcome: The principles paper discusses
  - High-level categories for AI uses cases in nuclear applications
  - Country-specific regulatory frameworks
- Summary considerations on
  - Use of existing safety and security systems engineering principles
  - Human and organizational factors
  - Characteristics of AI architecture
  - Lifecycle management
  - Demonstrating safe and secure AI systems that contain AI
- Working Group formed November 2022
- Paper is expected to be issued in Spring 2024



### Goal #3: Strengthen and Expand AI Partnerships

		FY 2023				FY 2024				FY 2025				FY 2026				FY 2027					
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4		
3.1: Domestic Partnerships	3.1A: Participate in NIST AI Risk Management Framework (ongoing)	[Green bar]																					
	3.1B: Participate in Light Water Reactor Sustainability (ongoing)	[Green bar]																					
	3.1C: Develop Plan for Partnerships with Academic Institutions					[Green bar]																	
3.2: International Partnerships	3.2A: Develop Project Plan for Trilateral Activities Between NRC, CNSC, and UKONR (CAN/UK/US)	[Blue bar with checkmark]																					
	3.2B: Draft CAN/UK/US Trilateral AI Principles Report			[Blue bar]																			
	3.2C: Final CAN/UK/US Trilateral AI Principles Report							[Blue bar]															
	3.2D: Participate in Bilateral Agreements with Canada, France, Germany, UAE, and UK (ongoing)	[Blue bar]																					
	3.2E: Participate in IAEA Project on Deployment of AI Solutions for the Nuclear Power Industry: Considerations and Guidance (ongoing)					[Blue bar]																	
	3.2F: Participate in IAEA Project on AI NPP Safety Implications (ongoing)					[Blue bar]																	
	3.2G: Participate in IAEA Project on AI Severe Accidents (ongoing)					[Blue bar]																	
3.3: Host and Participate in Public Workshops, Conferences, and Meetings	3.3A: Workshop: AI Characteristics for Regulatory Consideration			[Yellow bar with checkmark]																			
	3.3B: Workshop: Regulatory Gaps and Considerations							[Yellow bar]															
	3.3C: Host AI Workshops and Meetings (ongoing)	[Yellow bar with checkmark]						[Yellow bar]				[Yellow bar]				[Yellow bar]				[Yellow bar]			
	3.3D: Participate in External AI Workshops, Conferences, and Meetings (ongoing)	[Yellow bar]																					

## Goal 4. Cultivate an AI Proficient Workforce

### ***ACQUIRE, DEVELOP, RETAIN, AND TRAIN AN NRC AI KNOWLEGABLE STAFF***

- Focused on developing the critical skills for the AI workforce of tomorrow
- Training/Staffing
  - Develop up-skilling plans through opportunities and qualifications
  - Recruit AI Talent
- Workforce Planning
  - Conduct competency modeling



### Recruiting AI Talent on **USAJOBS**



**Outcome:** Ensure appropriate qualifications, training, expertise, and access to tools exist for the workforce to review and evaluate AI usage in NRC-regulated activities effectively and efficiently

**Strategic Goal #4:  
Cultivate an AI-Proficient Workforce**

		FY 2023		FY 2024				FY 2025				FY 2026				FY 2027			
		Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
<b>4.1: Assess NRC's AI Skills and Identify Gaps</b>	4.1A: Literature Review/Focus Groups to Identify Needed AI Skills				■														
	4.1B: AI Skills Assessment Survey					■													
	4.1C: Draft AI Competency Model Framework						■												
	4.1D: Pilot Test Draft AI Competency Models							■											
	4.1E: Final AI Competency Models								■										
	4.1F: Develop and Deliver Training to Support AI Competency Models									■									
	4.1G: Implement new AI Competency Models										■								
<b>4.2: Identify, Develop, and Implement AI Training Opportunities</b>	4.2A: Conduct Search to Identify AI Training Tools					■													
	4.2B: Update, as needed, Data Analytics Training and Qual Plan						■												
	4.2C: Review Existing NRC Qual Programs							■											
	4.2D: Identify and Create AI Career Development Opportunities								■										
	4.2E: Implement AI Career Development Opportunities, Capture and Address Feedback									■									
<b>4.3: Recruit, Hire, and Retain AI Talent</b>	4.3A: Develop AIWG for AI Skills and Recruiting				■														
	4.3B: Develop Office-Specific AI Staffing Plans					■													
	4.3C: Develop AI Hiring and Recruiting Strategy							■											
	4.3D: Develop AI Staff Retention Strategy									■									

# **Goal 5. Pursue Use Cases to Build AI Foundation Across the NRC**

## ***CREATE AI ECOSYSTEM TO PREPARE FOR REVIEWING AI APPLICATIONS***



### **Pilot Studies**

- Learn, measure, and evaluate readiness to implement regulatory framework
- Public workshops have shown industry interest to pursue pilot studies and proofs of concepts

### **AI Safety Insights**

- Survey industrial safety evaluation methods and tools
- Utilize AI partnerships and engagement strategies



### **AI Ecosystem**

- Establish integrated development environments and provide training
- Acquire common data science tools
- Develop regulatory sandboxes for supporting use-cases



### **AI R&D Research**

- Continue supporting University grants and research into AI systems
- Building AI foundation through the NRC's Future Focused Research Program



**Outcome:** Develop an ecosystem that supports AI analysis, integration of emerging AI tools, and hands-on talent development for reviewing AI applications from the nuclear industry

## Strategic Goal #5: Pursue Use Cases to Build an AI Foundation Across the NRC

		FY 2023		FY 2024				FY 2025				FY 2026				FY 2027			
		Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
<b>5.1 Proof-of-Concept Applications for AI Test and Analysis</b>	5.1A: Identify Strategy for Seeking Pilot Studies and Proof-of-Concept Test Cases																		
	5.1B: Create AIWG to Conduct Pilot Study																		
	5.1C: Conduct Proof-of-Concept Tests with Selected Pilot Studies																		
	5.1D: Document Lessons Learned from Pilot Study																		
<b>5.2: Develop and Maintain AI Ecosystem</b>	5.2A: Develop/Maintain Integrated Development Environment (ongoing)																		
	5.2B: Identify, Assess, and Acquire AI Tools (ongoing)																		
	5.2C: Develop Regulatory Sandboxes for AI Applications																		
	5.2D: Develop and Deliver Training on Use of Development Environments																		
<b>5.3: Survey AI Tools and Methods for Safety Evaluation</b>	5.3A: Conduct Survey of Industrial AI Tools and Methods																		
	5.3B: Implement Findings from Survey of Industrial AI Tools and Methods																		
<b>5.4: Facilitate and Invest in AI Regulatory Research</b>	5.4A: Develop Plan for AI Research through FFR Program																		
	5.4B: Develop Plan for AI Research through UNLP R&D Grant Program																		
	5.4C: Monitor Progress of AI-Related Research Projects and Incorporate Insights/Outcomes into NRC's AI Framework (ongoing)																		
	5.4D: Develop Internal Work Request Agreements (e.g., User Need Request) (ongoing)																		



---

# Moving Forward and Stakeholder Engagement

- NRC must remain vigilant—AI technologies are entering the nuclear domain in multiple venues
- NRC has been proactively working to understand this evolving technology to identify technical and regulatory challenges and gaps, gather insights on potential use cases, and develop institutional knowledge
- We are working to ensure we have the staff with the knowledge, skills, and ability to effectively regulate these new technologies
- Next Steps
  - Publish CANUKUS AI Principles Paper – Expected in Spring 2024
  - Publish AI Regulatory Gap Analysis – Expected in Spring 2024
  - Host AI Technical Session at 2024 RIC – March 12-14, 2024
  - Host IAEA AI Technical Meeting at USNRC HQ – March 18-24, 2024
  - Issue NRC’s AI Project Plan, Rev 1 – Expected in Fall 2024
  - Continue Public Workshops and Stakeholder Engagements – Ongoing



---

# Abbreviations

- AI – Artificial Intelligence
- AICoP – Artificial Intelligence Community of Practice
- AISC – Artificial Intelligence Steering Committee
- CFR – Code of Federal Regulations
- CNSC – Canadian Nuclear Safety Commission
- COE – Center of Expertise
- DOE – U.S. Department of Energy
- DENT – Digital Engineering in Nuclear Technology
- EO – Executive Order
- EPRI – Electric Power Research Institute
- FFR – Future-Focused Research
- FRN – Federal Register Notice
- FY – Fiscal Year
- GAO – U.S. Government Accountability Office
- GSA – U.S. General Services Administration
- GRS – Gesellschaft für Anlagen- und Reaktorsicherheit
- IAEA – International Atomic Energy Agency
- IEC – International Electrotechnical Commission
- IRSN – Institut de Radioprotection et de Sûreté Nucléaire
- ML – Machine Learning
- MOU – Memorandum of Understanding
- NLP – Natural Language Processing
- NRC – U.S. Nuclear Regulatory Commission
- NEI – Nuclear Energy Institute
- NIST – National Institute of Standards and Technology
- OMB – U.S. Office of Management and Budget
- ONR – U.K. Office for Nuclear Regulation
- RIC – Regulatory Information Conference
- RMF – Risk Management Framework
- UNLP – University Nuclear Leadership Program
- WG – Working Group

# Using Machine Learning to Inform Inspection Planning

## A Future Focused Research Project

Y. James Chang  
Senior Reliability and Risk Analyst  
RES/DRA/HFRB

Presented at ACRS Subcommittee Meeting on  
Human Factors, Reliability & PRA and Digital I&C Systems on Artificial Intelligence Strategic Activities  
November 15, 2023

# Presentation Outline

- Motivation, objective, and tasks
- Data and analysis
- Results and summary


# Motivation and Objective

- Motivation:
  - Covid-19 pandemic disrupted NRC's planned inspections
  - Netflix's success story of using unsupervised machine learning (ML) to identify hidden patterns (clusters) and similarities among its subscribers, leading to a more accurate movie recommendations
  - Can "ML + inspection findings" identify hidden safety patterns to inform inspection planning?
- Objectives
  - A feasibility study of using "ML + inspection findings" to identify hidden patterns (safety clusters) to inform inspection planning.

# Safety Clusters – Hidden Patterns

- Safety clusters: Similarities in failure modes and failure causes of structure, system, and component (SSC) and consequences
- An example: NRC Operating Experience Communication
  - Identified 5 power outage events impacted security operations
    - SSC: The primary and backup electricity power systems for security systems
    - Failure modes: fail to provide electricity
    - Failure causes:
      - 2022 (2 events): human errors contributed to the events
      - 2021 (3 events): not attributed to human errors
  - The OpE COMM suggested focusing on potential human' impacts on power supply equipment when conducting Inspection Procedure (IP) 71130.04, "Equipment Performance, Testing, and Maintenance."

# Study Approach and Tasks

- Study approach:
  - An NRC/RES future focused research (FFR) project
  - Contractor:  An AI engineering company  
When human intelligence isn't enough
  - NRC team has expertise in ML and reactor oversight process
- Two tasks:
  1. Evaluate AI/ML platforms:
    - Amazon's Sagemaker, Microsoft's Azure, Google's Google-AI, MatLab, and others
    - Maximize the use of pre-trained algorithms
  2. Select a platform to identify hidden patterns (safety clusters)
- Technical work was completed within 4 months after awarding the contract

# Task 1 Platform Evaluation

Criteria	Weight	Azure	AWS	Matlab	Google
Neural Topic Modeling	0.27	0	4	0	0
LDA	0.17	7	7	7	0
Visual Programming	0.14	7	6	0	0
Text Pre-processing	0.11	8	0	8	7
Text Embedding	0.10	9	9	9	0
Other Relevant Approaches	0.09	9	7	6	3
Notebook Integration	0.08	10	10	8	10
Text Extraction	0.03	7	7	7	7
Weighted Score		<b>84</b>	83	64	30

- Jupiter notebook (independent from the evaluated platforms) was used for Task 2 analysis.
- Jupiter: Free software, open standards, and web services for interactive computing across all programming languages.
- Cost was not evaluated but expected to be similar across platforms.

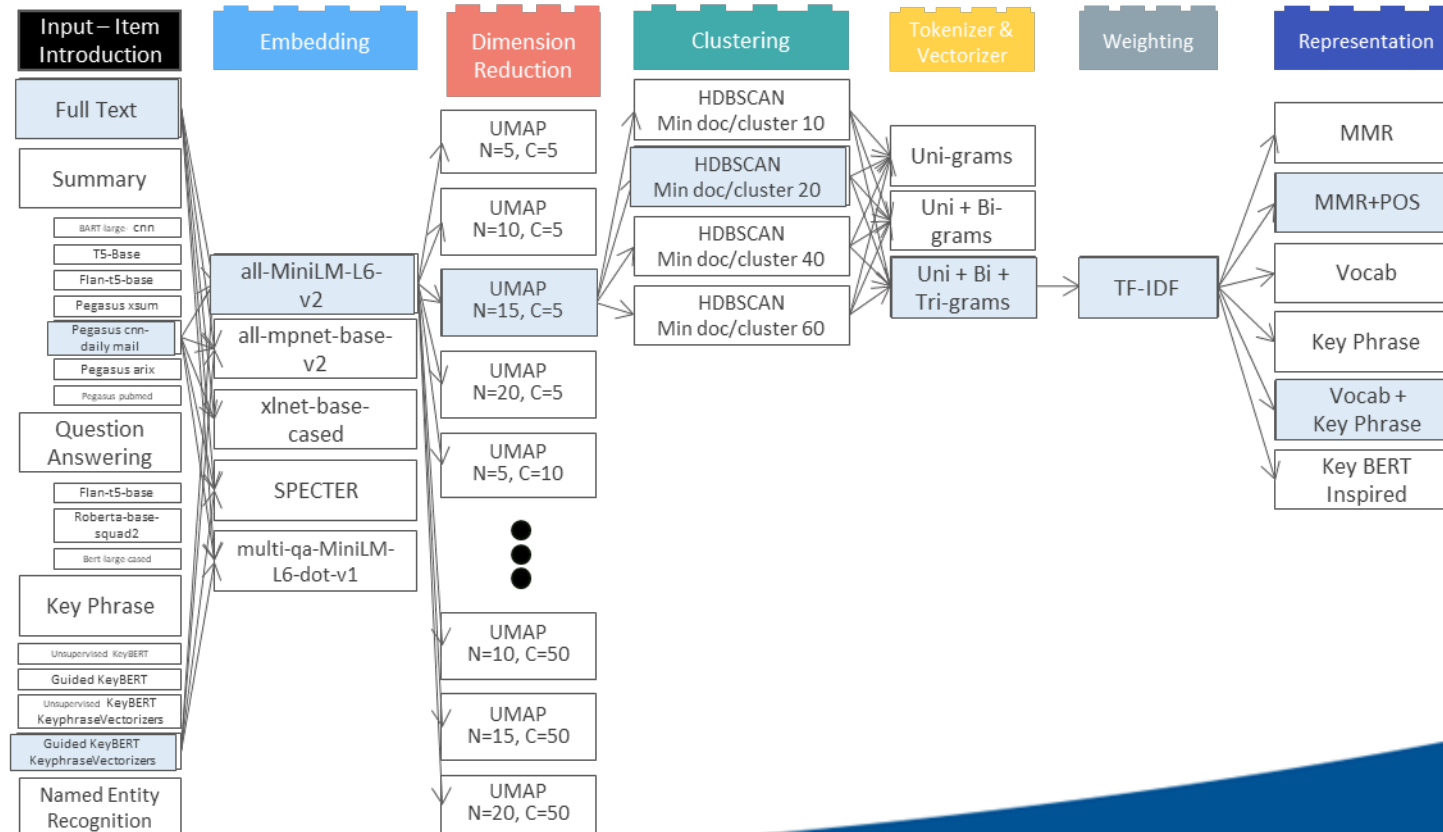




# Identify the Optimal Pipeline Constituents

- The pipeline contains multiple components to process the inspection finding descriptions to identify the safety clusters.
- A component's function can be performed by multiple algorithms
  - Identify the suitable algorithms
  - In some cases, algorithms are combined to produce better results
- Many trials and iterations to identify the optimal algorithms and parameters for the pipeline
  - Skip slow algorithms/parameters that may generate better results.

# Trial-and-Err to Identify Optimal Combinations



# Input Information for Analysis

- NRC ROP maintains an Excel database, which has all inspection findings since 1998. about 15,000 findings used for this analysis.
- The “Item Introduction” columns contain finding descriptions, which are the input for Task 2 analysis.

The screenshot shows an Excel spreadsheet with the following columns: A (Proced), B (Issue Date), C (Report), D (Type), E (Corner), F (Item Introduction), G (Corner), H (Title), and I (Docket). The 'Item Introduction' column is highlighted in green. A red box highlights the text in cell F2: "<p>The inspectors identified a finding of very low safety significance (Green) and an associated Requirements," for the licensee's failure to include Emergency Diesel Generator (EDG) Fuel Oil Program in accordance with American Society of Mechanical Engineers (ASME) Operation and M".

	A	B	C	D	E	F	G	H	I
1	Proced	Issue Date	Report	Type	Corner	Item Introduction	Corner	Title	Docket
2	71111.21M	02/07/2023	2023010	NCV	MS	<p>The inspectors identified a finding of very low safety significance (Green) and an associated Requirements," for the licensee's failure to include Emergency Diesel Generator (EDG) Fuel Oil Program in accordance with American Society of Mechanical Engineers (ASME) Operation and M	Mitigating	Failure to li	05000301
3	71111.21M	02/07/2023	2023010	FIN	MS	<p>The inspectors identified a finding	Mitigating	Failure to F	05000266
4	71111.21M	02/07/2023	2023010	NCV	MS	<p>The inspectors identified a finding	Mitigating	Failure to \	05000266
5	71111.21M	02/07/2023	2023010	NCV	MS	<p>The inspectors identified a finding	Mitigating	Failure to \	05000301
6	71111.21M	02/07/2023	2023010	FIN	MS	<p>The inspectors identified a finding	Mitigating	Failure to	05000201

# Process the Input Information

- Item Introduction (finding summary): Average 1,649 words for a finding (range from 42 to 11,670 words).
- Full text of the finding summary may not be the optimal choice because of sentence transformer models limitations.
  - Use ML to generate condensed summaries.
- Tried full text and 14 condensed summaries (in three categories)
  - Full text (1)
  - Summary techniques (7)
  - Question-answering techniques (3)
  - Key phrase extraction techniques (4)

# Custom Words and Phrases and Stop-Word

- Some ML techniques take input from customized lists of words and phrases to work with the pre-trained models.
- NRC provided 1004 acronyms and 407 common failure modes to focus nuclear safety.
  - Acronyms: MSIV and RPV, etc.
  - Failure mode: inoperable, misalign, and corrosion, etc.
  - Provided 269 NUREGs and 195 RILs to develop a library of words/phrases relative locations.
- Stop-word removal:
  - 337 English stop words, e.g., the and a.
  - 136 custom stop words, e.g., safety, system, and reactor.



# Represent Safety Clusters

- Safety clusters are represented by word cloud (or a bag-of-words)
- In this analysis, an inspection finding only belongs to a safety cluster

Topic	Count	Top_n_words
-1	5382	green - cooling - 50 - reactor - appendix criterion - safety significance - 10 cfr - valves - finding low - design
0	927	safe shutdown - facility operating - license condition protection - cited violation license - brigade - approved protection program - suppression - renewed - combustibles - unit operating license
1	925	high radiation area - area - workers - 10 cfr 20 - safety cornerstone - dose rates - occupational radiation safety - radioactive material - 1501 - rem
2	790	emergency diesel - emergency diesel generator - diesel generator edg - oil - generators - start - division emergency - dg - generator fuel oil - inoperable
3	558	driven auxiliary - turbine driven - auxiliary feedwater pump - bearing - tda fw - feedwater afw - sx - essential service water - charging - pump inoperable





# Evaluate Results Against an OpE COMM

## OpE COMM

Date Released: 5/23/2023

### *Operating Experience Communication*

- Identified 5 issues related to improper calibration and maintenance of radiation monitoring and dose assessment equipment that impact emergency plan actions
  - [REDACTED] 2011-2022 (2 events)
  - [REDACTED] 2019
  - [REDACTED] 2016
  - [REDACTED] 2013
- The OpE COMM identifies opportunities to identify these issues under Inspection Procedure (IP) 71124.05, “Radiation Monitoring Instrumentation”, and emergency drill observations, plant modifications or surveillance test reviews

# Benchmark Results

Finding	Full Item Intro	Summary	Key Phrase
[REDACTED]	2	61	38
[REDACTED]	5	1	1
[REDACTED]	3	1	2
[REDACTED]	3	1	2



OpE identified 4 findings that exhibited safety issues that are related  
 The clustering approach placed 3 of the 4 in the same cluster and the 4<sup>th</sup> in a similar cluster

# Summary

- ML algorithms to generate summary are useful for ROP operations.
- Unconclusive about unsupervised ML's practicality to inform inspection planning
  - Need additional efforts to optimize the pipeline
  - Outside the scope of this future focused research project.

# Back Up Slides

# Evaluate ML Techniques to Generate Condensed Summary – Original Summary

The inspectors identified a Green NCV of Unit 3 Technical Specification (TS) 5.4.1 when Entergy **did not take adequate measures to control transient combustibles** in accordance with established procedures and thereby did not maintain in effect all provisions of the approved fire protection program, as described in the Unit 3 final safety analysis report. Specifically, on two separate occasions, Entergy did not ensure that transient combustibles were evaluated in accordance with established procedures; and as a result, they **allowed combustible loading in the 480 volt emergency switchgear room to exceed limits established in the fire hazards analysis (FHA)** of record. The inspectors determined that not completing a TCE, as required by EN-DC-161, "Control of Combustibles," Revision 18, was a performance deficiency, given that it was reasonably within Entergy's ability to foresee and correct and should have been prevented. Specifically, on August 28, 2018, **wood in excess of 100 pounds was identified in the switchgear room; however, an associated TCE had not been developed.** Additionally, on October 1, 2018, **three 55-gallon drums of EDG lube oil were stored in the switchgear room without an associated TCE having been developed to authorize storage in this room,** as required for a volume of lube oil in excess of 5 gallons. The inspectors determined the performance deficiency was more than minor because it was associated with protection against external factors attribute of the Mitigating Systems cornerstone, and it adversely affected the cornerstone goal of ensuring the availability, reliability, and capability of systems that respond to initiating events to prevent undesirable consequences. Specifically, **storage of combustibles in excess of the maximum permissible combustibles loading** could have the potential to challenge the capability of fire barriers to prevent a fire from affecting multiple fire zones and further degrading plant equipment. Additionally, this issue was similar to an example listed in IMC 0612, Appendix E, "Examples of Minor Issues," Example 4.k., because the fire loading was not within the FHA limits established at the time. Entergy required the issuance of a revised evaluation to provide reasonable assurance that the presence of combustibles of a quantity in excess of the loading limit of record would not challenge the capacity of fire barriers, and further evaluation and the issuance of an EC was necessary to raise the established loading limit to a less-conservative value. The inspectors assessed the significance of the finding using IMC 0609, Appendix F, "Fire Protection Significance Determination Process," and determined that this **finding screened to Green (very low safety significance)** because it had a low degradation rating in accordance with Attachment 2 of the appendix. The inspectors determined that this finding had a **cross-cutting aspect** in the area of **Human Performance, Work Management**, because Entergy did not adequately plan, control, and execute work activities such that nuclear safety was the overriding priority, nor did they adequately identify risk associated with work being performed or coordinate across working groups to anticipate and manage this risk. Specifically, in the case of wood scaffolding being stored in the switchgear room, while planning work to be performed, Entergy did not adequately consider the fire risk that would be introduced by the presence of additional combustible materials. In the case of lube oil being stored in the room, Entergy did not take adequate action to ensure that activities were executed in a manner that would prevent work taking place in one area (the adjacent EDG cell) from introducing additional fire risk into a space for which it had not been evaluated (the switchgear room). In both cases, Entergy did not take sufficient action to ensure that workers were aware of the fire protection requirements associated with activities being conducted and to ensure that they coordinated as needed across working groups to adequately assess and mitigate the associated fire risk.

# Summarization Models: Examples

The inspectors identified a Green NCV of Unit 3 Technical Specification (TS) 5.4.1 when Entergy **did not take adequate measures to control transient combustibles** in accordance with established procedures and thereby did not maintain in effect all provisions of the approved fire protection program, as described in the Unit 3 final safety analysis report [...]

## T5-Base

inspectors identified a Green NCV of Unit 3 Technical Specification (TS) 5.4.1 when Entergy did not take adequate measures to control transient combustibles. the inspectors determined that not completing a TCE was a performance deficiency, given that it was reasonably within entergy's ability to foresee and correct and should have been prevented. in the case of wood scaffolding being stored in the switch

## Flan-T5-base

Entergy did not take adequate measures to control transient combustibles in accordance with established procedures and thereby did not maintain in effect all provisions of the approved fire protection program, as described in the Unit 3 final safety analysis report.

## BART-large-cnn

The inspectors identified a Green NCV of Unit 3 Technical Specification (TS) 5.4.1 when Entergy did not take adequate measures to control transient combustibles. The inspectors determined that not completing a TCE, as required by EN-DC-161, "Control of Combustibles," Revision 18, was a performance deficiency.

## Pegasus-cnn-dailymail

Entergy did not ensure that transient combustibles were evaluated in accordance with established procedures . Entergy did not take sufficient action to ensure that workers were aware of the fire protection requirements associated with activities being conducted . Entergy did not adequately identify risk associated with work being performed or coordinate across working groups to anticipate and manage this risk

# Question-Answering Models: Examples

The inspectors identified a Green NCV of Unit 3 Technical Specification (TS) 5.4.1 when Entergy **did not take adequate measures to control transient combustibles** in accordance with established procedures and thereby did not maintain in effect all provisions of the approved fire protection program, as described in the Unit 3 final safety analysis report [...]

Flan-T5-base

Storage of combustibles in excess of the maximum permissible combustibles loading could have the potential to challenge the capability of fire barriers to prevent a fire from affecting multiple fire zones and further degrading plant equipment

Roberta-base-squad2

nuclear safety

Bert-large-casedwhole-word-maskingfinetuned-squad

nuclear safety

No consistent results  
attained so QA was not  
selected



# Key Phrase Extraction Methods: Examples

The inspectors identified a Green NCV of Unit 3 Technical Specification (TS) 5.4.1 when Entergy **did not take adequate measures to control transient combustibles** in accordance with established procedures and thereby did not maintain in effect all provisions of the approved fire protection program, as described in the Unit 3 final safety analysis report [...]

## KeyBERT

allowed combustible loading, allowed combustible, combustibles revision 18, combustibles evaluated accordance, permissible combustibles loading, result allowed combustible, combustibles revision, combustibles evaluated, permissible combustibles, transient combustibles evaluated, additional combustible, maximum permissible combustibles, combustibles loading, 161 control combustibles, final safety analysis, control combustibles revision, presence additional combustible, unit final safety, combustibles accordance established, established hazards analysis

## Guided KeyBERT

allowed combustible loading, final safety analysis, unit final safety, combustibles evaluated accordance, safety analysis report, combustibles revision 18, allowed combustible, permissible combustibles loading, transient combustibles evaluated, combustibles evaluated, permissible combustibles, result allowed combustible, maximum permissible combustibles, 161 control combustibles, combustibles revision, combustibles loading, established hazards analysis, safety significance, safety analysis, control combustibles revision

## KeyBERT + KeyphraseVectorizers

additional fire risk, fire protection requirements, final safety analysis report, fire risk, maximum permissible combustibles loading, fire protection significance determination process, fire barriers, additional combustible materials, combustible loading, fire protection program, combustibles, transient combustibles, low safety significance, edg lube oil, fire loading, entergy, multiple fire zones, fire, nuclear safety, further degrading plant equipment`

## Guided KeyBERT + KeyphraseVectorizers

final safety analysis report, fire protection requirements, additional fire risk, fire risk, fire protection significance determination process, maximum permissible combustibles loading, fire barriers, fire protection program, combustible loading, low safety significance, additional combustible materials, transient combustibles, combustibles, edg lube oil, further degrading plant equipment, fire loading, nuclear safety, entergy, multiple fire zones, volt emergency

- **Input: text for analysis**

- Some algorithms can be used in a semi-supervised manner with a pre-defined list of important words and phrases to guide the algorithm.

- **Embedding:**

- Stop words: procedure, technical, license condition, 'safety',
- 'reactor', 'power plant', 'inspector', 'license', 'finding', 'cornerstone', 'cross cutting area'
- 'mitigating', 'systems', 'barrier', 'integrity', 'initiating', 'event',
- 'human', 'performance', 'problem', 'identification', and 'resolution'.



**Bruce P Hallbert, Ph.D.**  
**Director, Light Water Reactor**  
**Sustainability Program**  
**Technical Integration Office**

# **Light Water Reactor Sustainability: Sustaining and Optimizing the Existing Fleet**



# Light Water Reactor Sustainability Program

- **Goal**

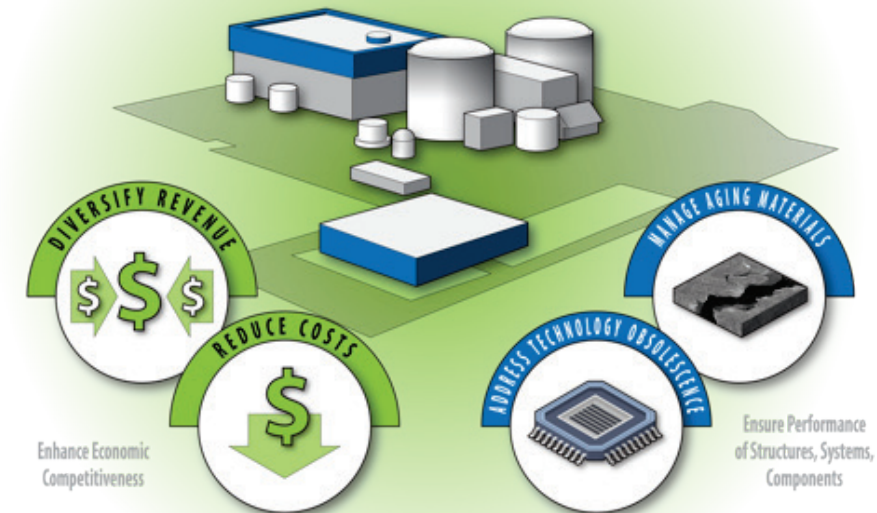
- Enhance the safe, efficient, and economical performance of our nation's nuclear fleet and extend the operating lifetimes of this reliable source of electricity

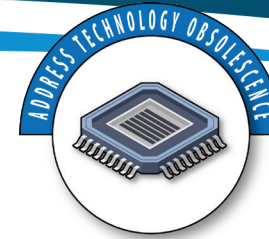
- **Objectives**

- Enable long term operation of the existing nuclear power plants
- Deploy innovative approaches to improve economics and economic competitiveness of LWRs in the near term and in future energy markets
- Sustain safety, improve reliability, enhance economics

- **Focus Areas**

- Plant Modernization Research and Development
- Flexible Plant Operation and Generation
- Risk-Informed Systems Analysis
- Materials Research
- Physical Security





# Plant Modernization

## Facilitate modernization by:

- **Developing technology modernization solutions** that address aging and obsolescence challenges
- **Delivering a sustainable business model** that ensures continued safe, reliable operation at a cost competitive level



### Long Term Management Of Plant Systems

License Extension  
60, 80, or 100 years



### Nuclear Cost Competitiveness

Cost pressures from other  
generation sources



### Worker Attraction and Retention

Worker interest in new  
technology

# Modernize the Fleet

- First echelon safety instrumentation and control systems on two units
- Conceptual Design Phase complete
- Detailed Design Phase in progress
- Multiple pre-submittal meetings with Nuclear Regulatory Commission (NRC)
- Human Factors efforts well underway
  - Operating Experience Review (Q3–Q4 of 2021)
  - Function Analysis and Allocation Workshop (March 2022)
  - Task Analysis Workshop (May 2022)
- NRC has accepted Constellation’s License Amendment Request (December 2022)
- Dynamic preliminary validation completed February 2023 at INL with NRC observation



Limerick Generating Station



INL Human Systems Simulation Laboratory Task Analysis Workshop

# Artificial Intelligence (AI) and Machine Learning(ML) R&D

AI/ML, associated methods and data analytics are relatively new to the nuclear power industry.

- Show promise for automating some manually performed activities
- Investigating approaches that show promise to enhance efficiency
- R&D efforts – not deployment

Current efforts emphasize work process vs. control.

- Adoption must align with the nuclear safety culture of the industry
- Align and comply with recent presidential directive and other requirements

Trust in automation, transparency, understandability affect usability.

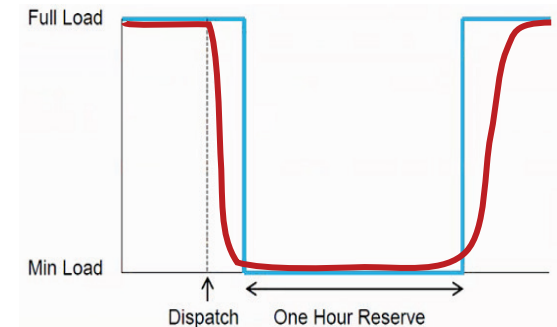
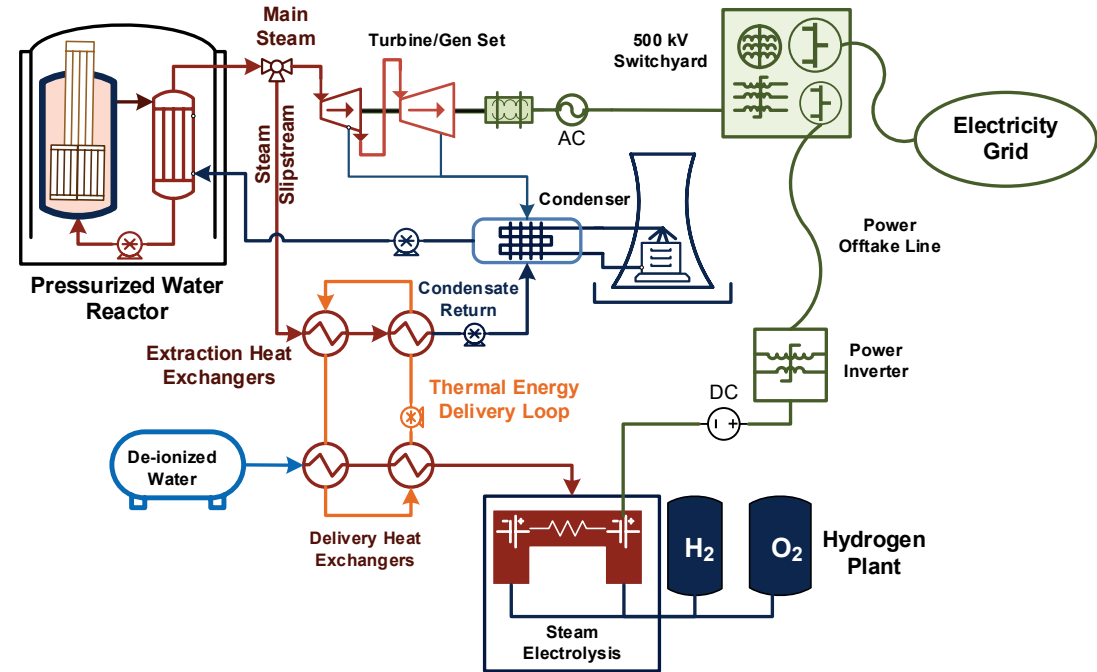
- Human factors issues in AI/ML implementation vital to adoption and safe use.





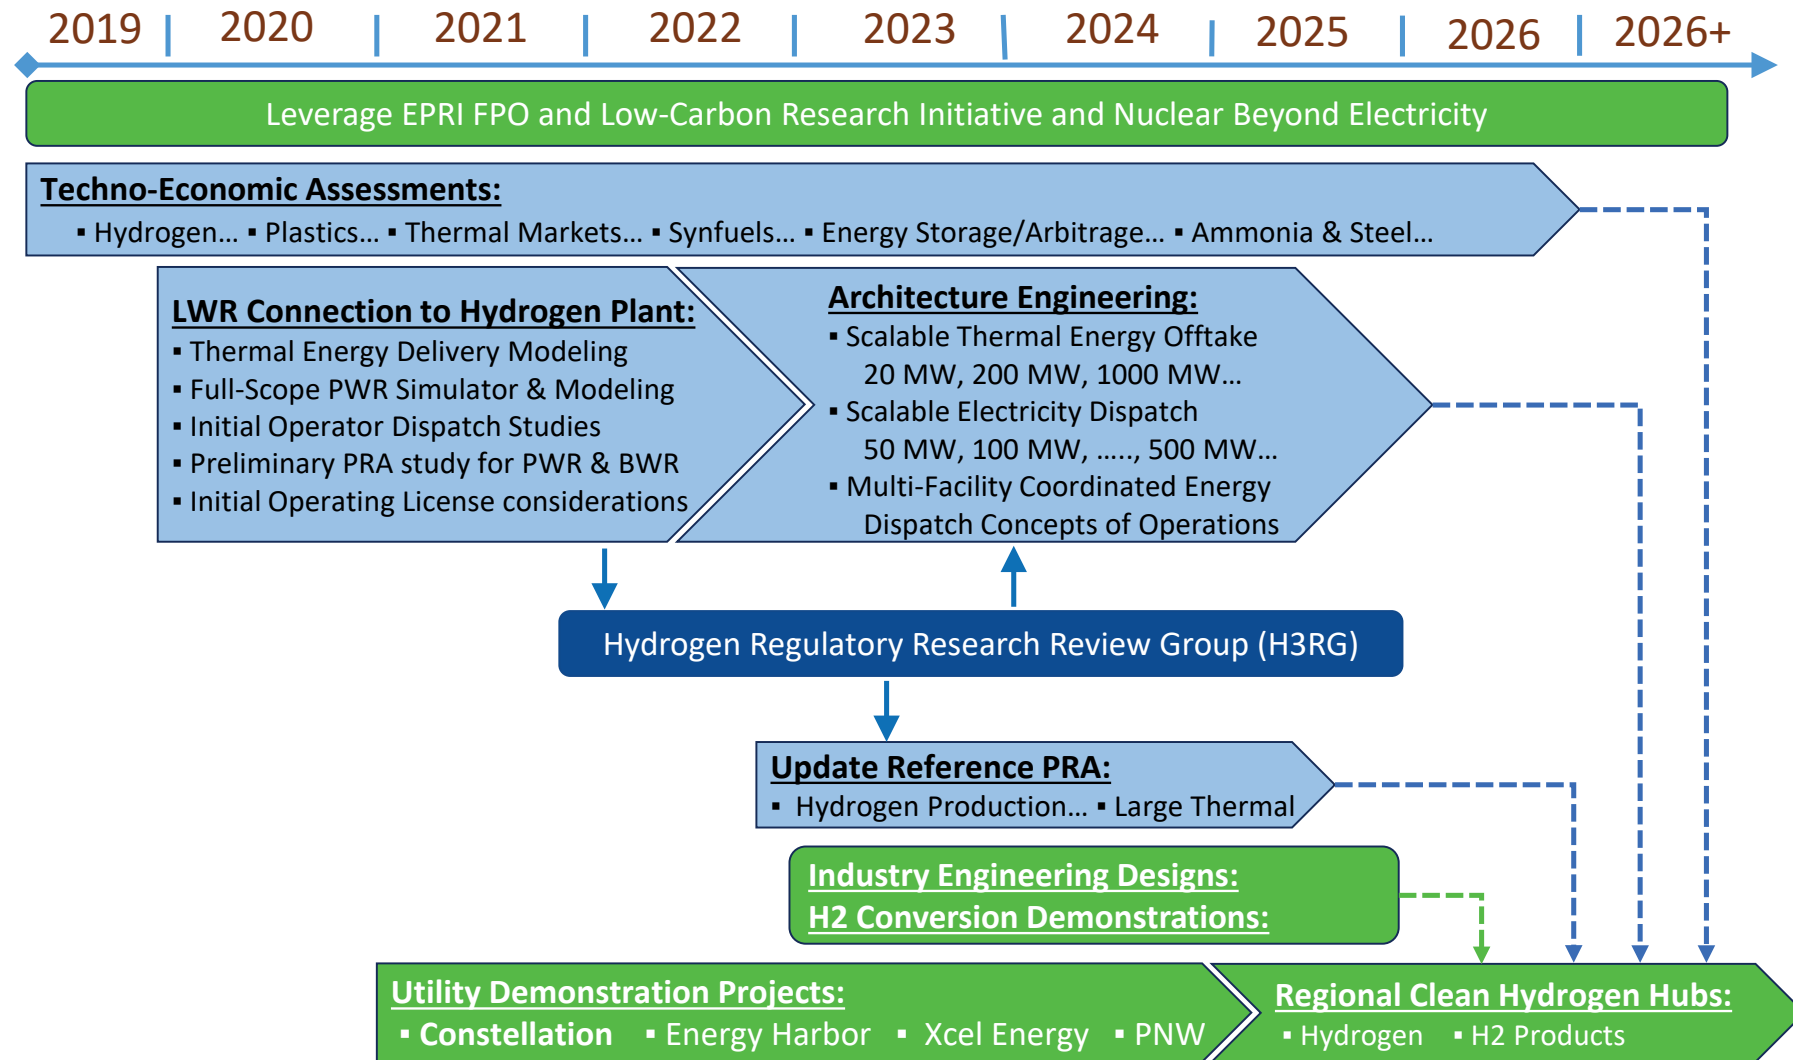
# Flexible Plant Operation and Generation

- Enhance economic flexibility and decarbonize energy and industry
- LWRS research addresses:
  - Hydrogen production and storage safety risks
  - New thermal extraction and delivery systems
  - Modifications to the electricity transmission station
  - Operator control of dynamic dispatch of power
  - Economics of transitioning between the electricity grid market and hydrogen production



# Enabling Nuclear-Hydrogen Hubs

- The Bipartisan Infrastructure Law funds Clean Regional Hydrogen Hub Projects
- The Inflation Reduction Act provides Clean Hydrogen Production Tax Credits (up to \$3 per kg-H<sub>2</sub> with low-emissions energy sources)



# Pilot Plant Hydrogen Demonstration Projects

## Nine Mile Point Nuclear Power Plant

- **Constellation (Exelon): Nine-Mile Point NPP**
  - 1.25 MWe low temperature electrolysis (LTE)
  - Using “house load” power
  - Training plant operators by practicing power switching between grid and hydrogen plant



Operating since  
February 2023

## Davis-Besse Nuclear Power Plant

- **Energy Harbor: Davis-Besse NPP**
  - 1-2 MWe LTE or 2-4 MWe HTE (with electrical steam boiler)
  - Power provided by plant upgrade with new switch gear at the transmission station
  - Gaining hydrogen production experience in anticipation of scaling up production



H<sub>2</sub> production  
beginning  
March/April 2024

## Prairie Island Nuclear Power Plant

- **Xcel Energy: Prairie Island NPP**
  - 150 kWe high temperature electrolysis (HTE)
  - Thermal tie into existing turbine stream extraction line
  - Gaining high temperature hydrogen production experience



H<sub>2</sub> production  
beginning middle  
of 2024

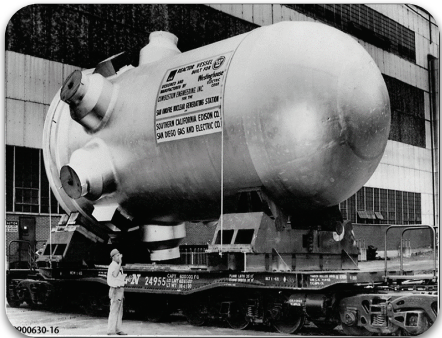


# Materials Research

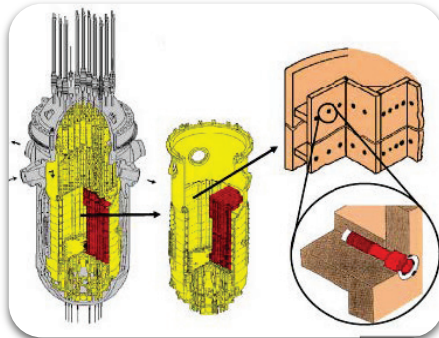
**Objective:** conduct R&D to understand the long-term environmental degradation behavior of materials in nuclear power plants

1. Measurement of degradation
2. Mechanisms of degradation
3. Modeling and simulation
4. Monitoring degradation
5. Mitigation strategies
6. Materials harvesting

Reactor Pressure Vessel (RPV)



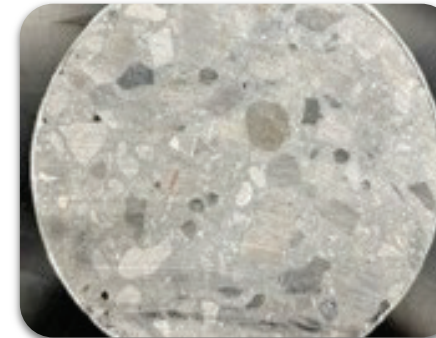
Core internal and pressure boundary



Mitigation



Concrete



Cable



# Risk-Informed Systems Analysis (RISA)

## Objective

- R&D to achieve economic efficiencies while maintaining high levels of safety optimize safety margins and minimize uncertainties

## Approach

- Provide scientific basis to better represent safety margins and factors that contribute to cost and safety

## Areas of Expertise

- Advanced modeling of physics-based phenomena
  - Thermal-hydraulics, neutronics and reactor physics, risk-informed material degradation, uncertainty propagation
- Advanced Data Analytics and Digital Modeling
  - Diagnostic and prognostic analyses, resource optimization, AI/ML technologies, digital twins, uncertainty propagation
- Probabilistic Risk Assessment (PRA) and Human-Reliability Analysis (HRA)



RISA Framework

# Optimization of Nuclear Fuel Use



Fuel costs represents approximately 20% of annual operating expenses.



AI-based optimization framework for designing reactor core configuration given objectives and constraints

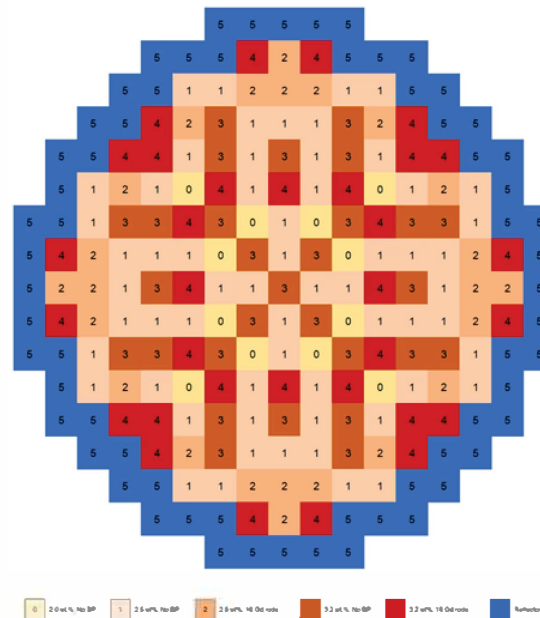


Multiphysics-based process with inputs from reactor core design, thermal-hydraulics, fuel performance

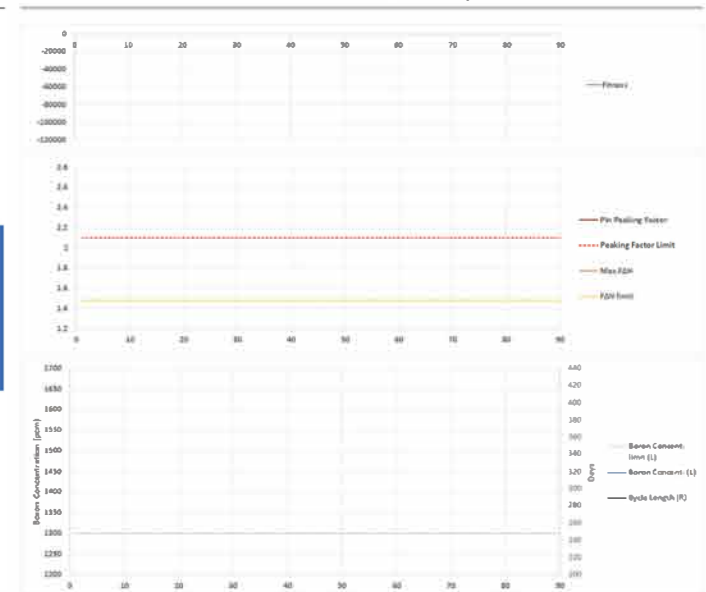


Genetic algorithm is the AI method used for optimization

Reactor Core Fuel Pattern



Fitness of Pattern, Constraints, Objective



Example case of finding optimal fuel configuration in the reactor core with the objective to prolong the cycle length under peaking factor and boron concentration constraints.

# Physical Security Research

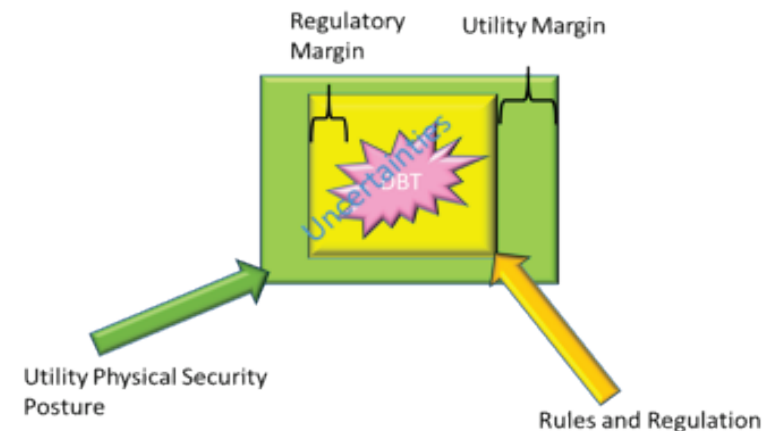
Physical Security research aims to create tools, technologies, and risk-informed physical security decisions and activities with the following objectives:

- Develop mitigation strategies and enhance the technical basis necessary for stakeholders to reevaluate physical security postures while meeting regulatory requirements.
- Analyze the existing physical security regime, current best practices, and compare/contrast insights with alternative methods which leverage advanced modeling and simulation, modern technologies, and novel techniques that address design basis threat and regulatory requirements

Short-term goal is to enable industry to operate nearer the staffing requirements of 10 CFR 73.55

Main research thrust areas:

- Advanced Security Technologies
- Risk-Informed Physical Security
- Advanced Security Sensors and Delay



# Summary

- The existing operating fleet provides the largest reliable source of carbon-free electricity.
- Industry initiatives have achieved substantial improvements in performance.
- Nuclear energy supports climate goals and can contribute to deep decarbonization by providing clean energy for products used in other industrial sectors.
- LWRS Program supports collaborations with industry to facilitate progress in areas of vital common interest. By working together, we facilitate progress and address challenges to ensure the continued viability and role of nuclear energy.
- The growing demand for clean and reliable energy from nuclear power underscores the need to address existential challenges facing the existing fleet.
- LWRS research addresses highest priority issues on timelines that support continued operation of the existing fleet.





# Sustaining National Nuclear Assets

*[lwrs.inl.gov](http://lwrs.inl.gov)*

# Integrated Operations for Nuclear (ION)

## Motivations for ION



## ION Business Model is guided by strategy

### Traditional Approaches to Modernization



Modernization competes with plant health for resources



Modernization Projects may lose momentum (unclear value)



Corporate vision not integrated with modernization

## ION Yields Cost Savings

**99.7%**  
chance of positive net present value  
*ex. Digital I&C*

**\$60M**  
estimated harvestable annual cost savings by implementing ION

### Integrated Operations for Nuclear



Modernization projects funded through a separate budget process



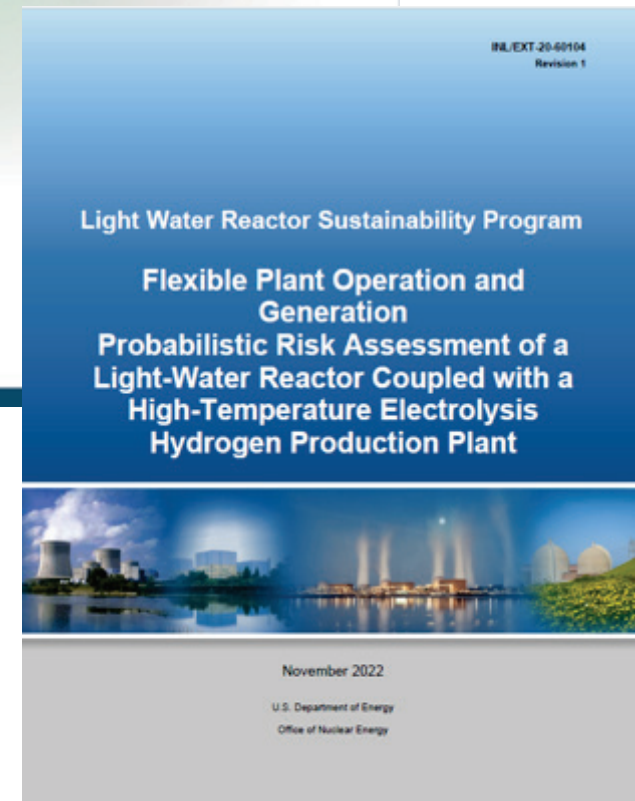
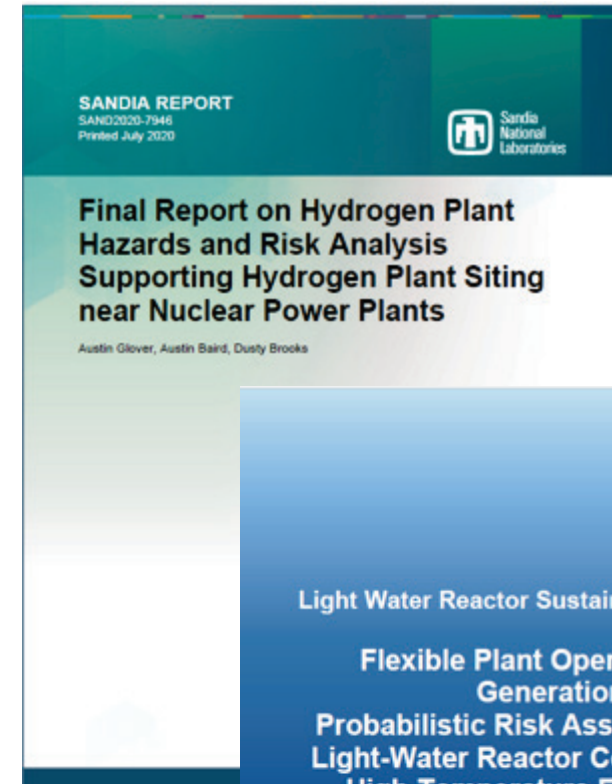
One roadmap sequencing all modernization efforts



Complete and current strategy guiding facility modernization

# Preliminary PRA for Steam Electrolysis

- Based on potential new internal and external initiating events
  - **Top hazards identified**
    - Internal: Steam line break, loss of offsite power
    - External: Electrolysis Plant H<sub>2</sub> leak or H<sub>2</sub> detonation
- Hazards analysis completed by Sandia using HYRAM™
- Failure Modes and Effects Analysis (FEMA) being considered for hydrogen plant design and layout
- Comparison of results to criteria in 10 CFR 50.59 and RG 1.174
- **Early conclusions**
  - Licensing criteria are met for a large-scale hydrogen production facility sited 1 km from a generic PWR and BWR
  - Individual site NPP and geographical features can affect the results of the generic PRA positively or negatively

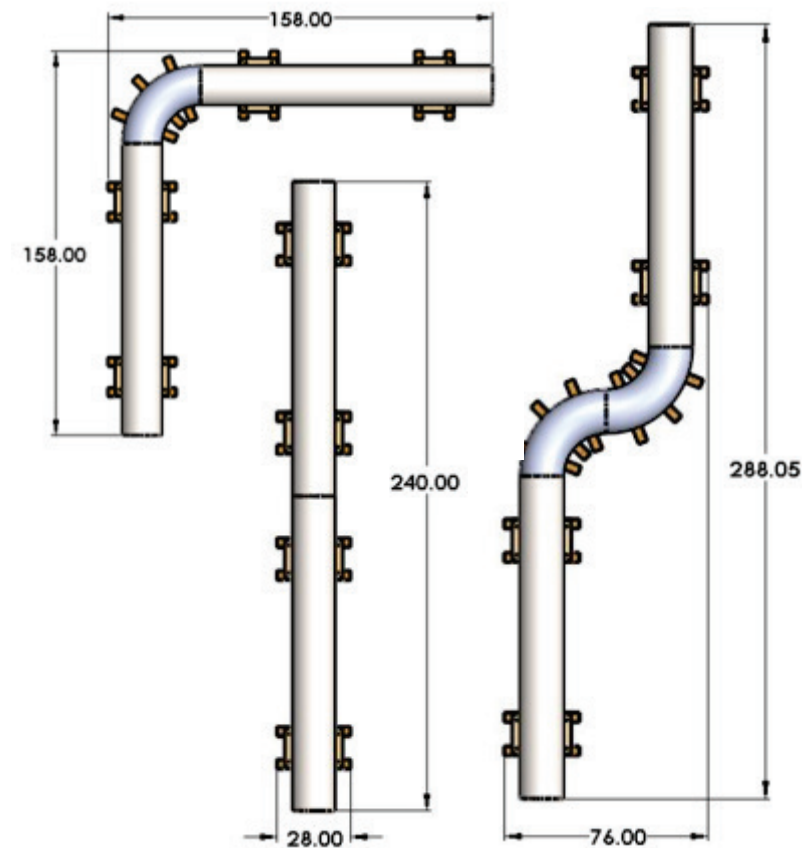


# Unattended Openings (UAO)

**Impact:** Provide the technical basis to determine optimized protective strategies related to person-passible openings that intersect security boundaries during normal and maintenance operations. Reductions in patrols, monitoring, and compensatory measures could be reconsidered but will be site specific on a case-by-case basis.

## Highlights

- Conducted a performance-based, risk-informed evaluation for 2D and 3D unattended openings based on the current US Government policy
- Identified human factors associated with 2D and 3D openings
- Evaluated 2D UAO testing with 4-inch circle and rectangles and 36-inch circle
- Evaluated 3D UAO testing with 20-foot piping sections and pipe bends
- Evaluated success of passing through the opening (go/no-go), rate times, and limited data on exertion



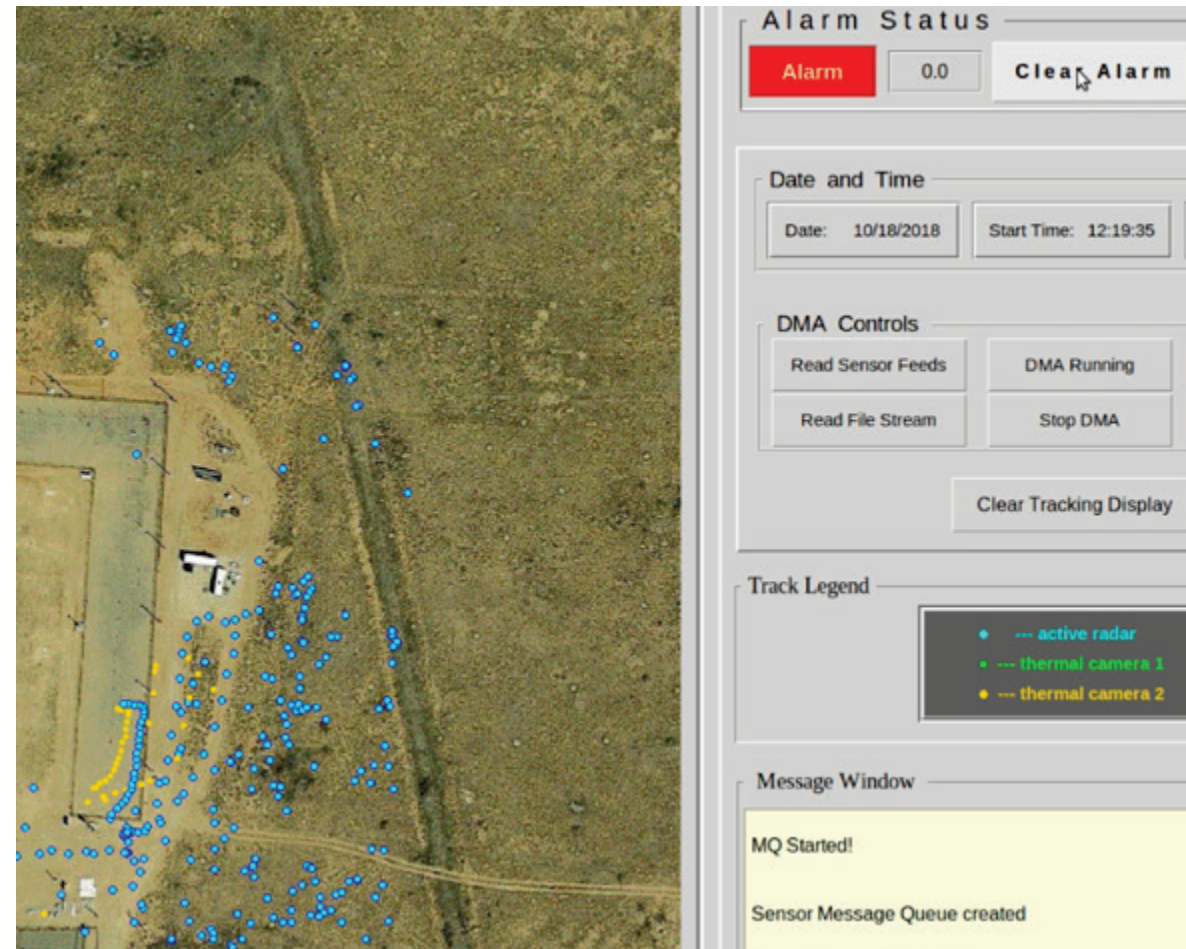
Example of 3D UAO Testing Configurations

# Deliberate Motion Analytics (DMA)

- Security sensor fusion linked with DMA can take input from multiple sensors of different types, analyze the data, and determine if an adversary is making an approach toward a facility.
- Sites using current commercial sensor technologies typically experience elevated nuisance alarm rates (NAR) not caused by an intruder. Maintaining a low NAR while being able to detect intruders has the potential to decrease the cost of security.

## Highlights

- Used DMA and sensor fusion, collecting at least four weeks of continuous performance data at two nuclear power plant sites
- Considered engineered terrain (perimeter intrusion detection system) and un-engineered terrain (owner controlled area).
- Created an NPP-specific demonstration package containing sensor fusion



Active Radar (blue) and Thermal Camera (yellow) fused through DMA



**Craig Primer**  
**Plant Modernization**  
**Pathway Lead**

**November 2023**

# **LWRS Program Plant Modernization**



# AI/ML

## Research Focus Areas

ML for Material Management  
ML for Equipment Monitoring  
ML for Anomaly Detections

NLP Applications

Computer Vision Applications

AI/ML Explainability

### A model for thinking about AI, ML, DL, and generative AI



#### Artificial intelligence (AI)

Any technique that allows computers to mimic human intelligence using logic, if-then statements, and machine learning



#### Machine learning (ML)

A subset of AI that uses machines to search for patterns in data to build logic models automatically



#### Deep learning (DL)

A subset of ML composed of deeply multi-layered neural networks that perform tasks like speech and image recognition



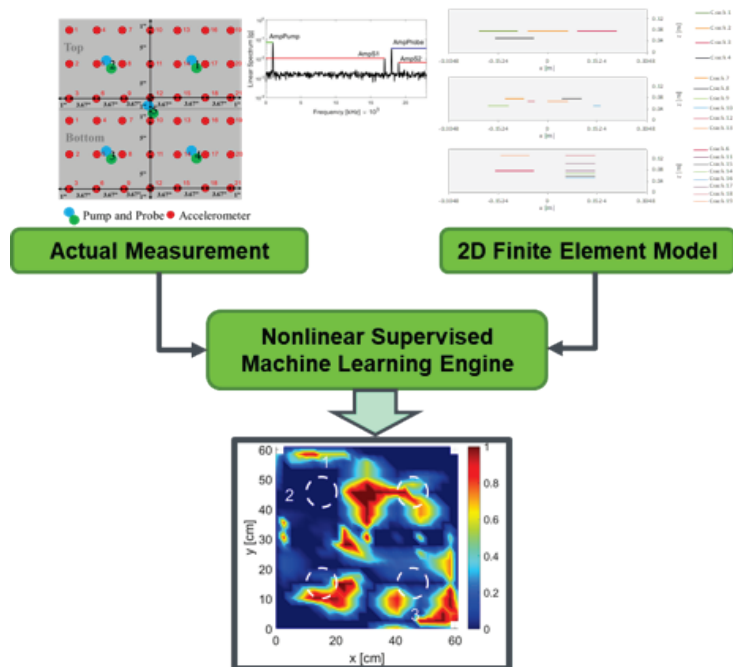
#### Generative AI

Powered by large models that are pretrained on vast corpora of data and commonly referred to as foundation models (FMs)

Overview of AI, ML and subsets of ML

# Digital Twin Informed ML for Material Management

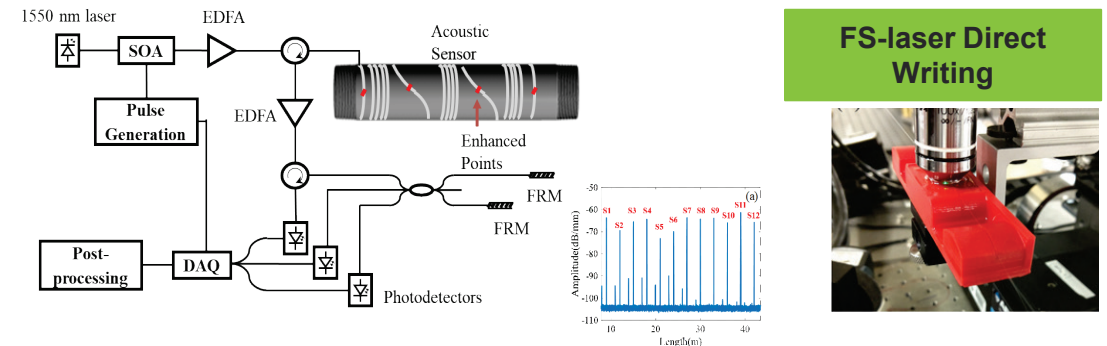
Developed technology to locate and estimate alkali-silica reaction (ASR) damage using physics-informed machine learning approach



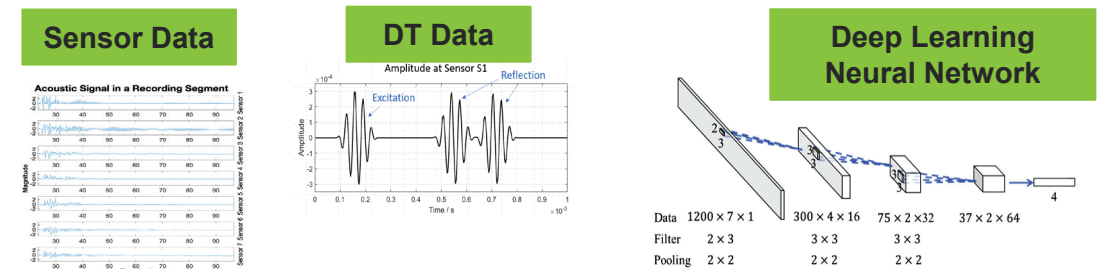
Physics-informed machine learning damage map due to ASR

Digital Twin and Deep Learning Model used in Secondary Piping Degradation Detection Research

## Low-Cost Phase-Sensitive Distributed Fiber



## Digital Twin and AI



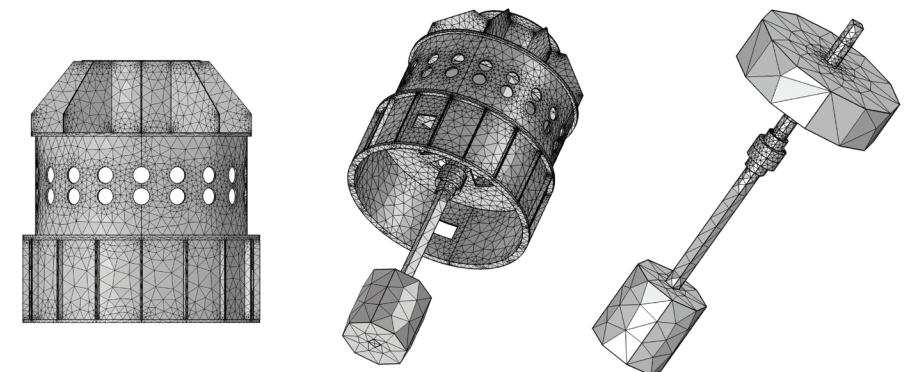
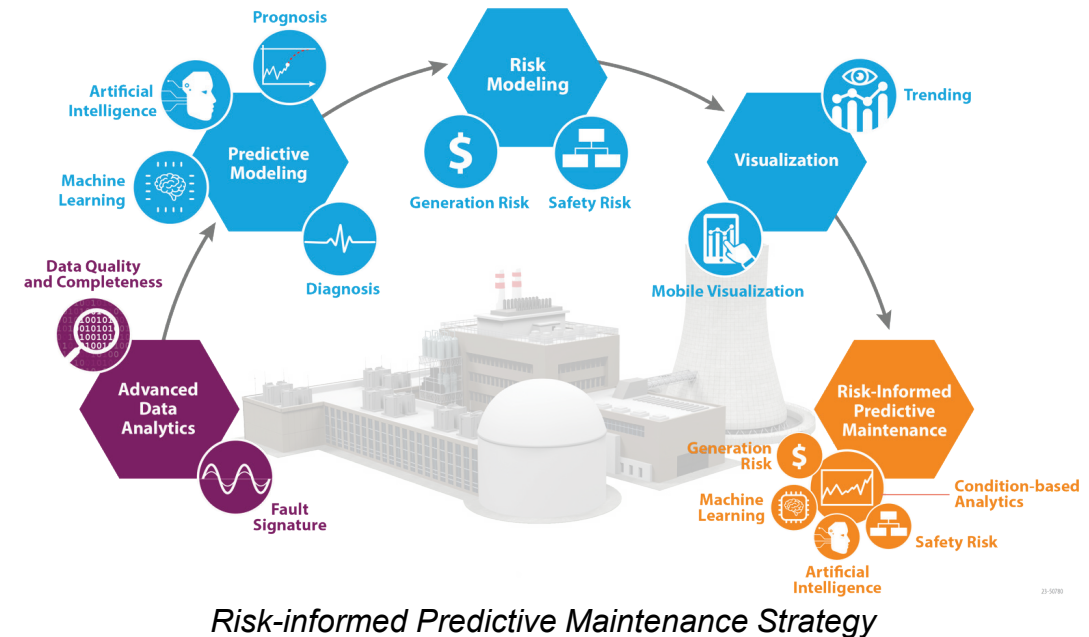
Concrete Structure Health Monitoring Using Vibroacoustic Testing and Machine Learning, [INL/EXT-20-59914](#)

Concept for Integrated Multi-Modal Online Piping Monitoring System along with Data Fusion and Advanced Data Analytical Algorithms Using High-Resolution Fiber Optics Sensors, [INL/EXT-20-59810](#)



# Predictive Maintenance Strategy


- Developed a scalable risk-informed predictive maintenance strategy using machine learning approaches, risk modeling, visualization, and multi-band heterogeneous wireless architecture.
- Developed a hybrid model of circulating water pump (CWP) motor (basis for digital twin) to capture different operating dynamics.





*Physics-based model of CWS*


Scalable Technologies Achieving Risk-Informed Condition-Based Predictive Maintenance Enhancing the Economic Performance of Operating Nuclear Power Plants, [INL/EXT-21-64168](#)

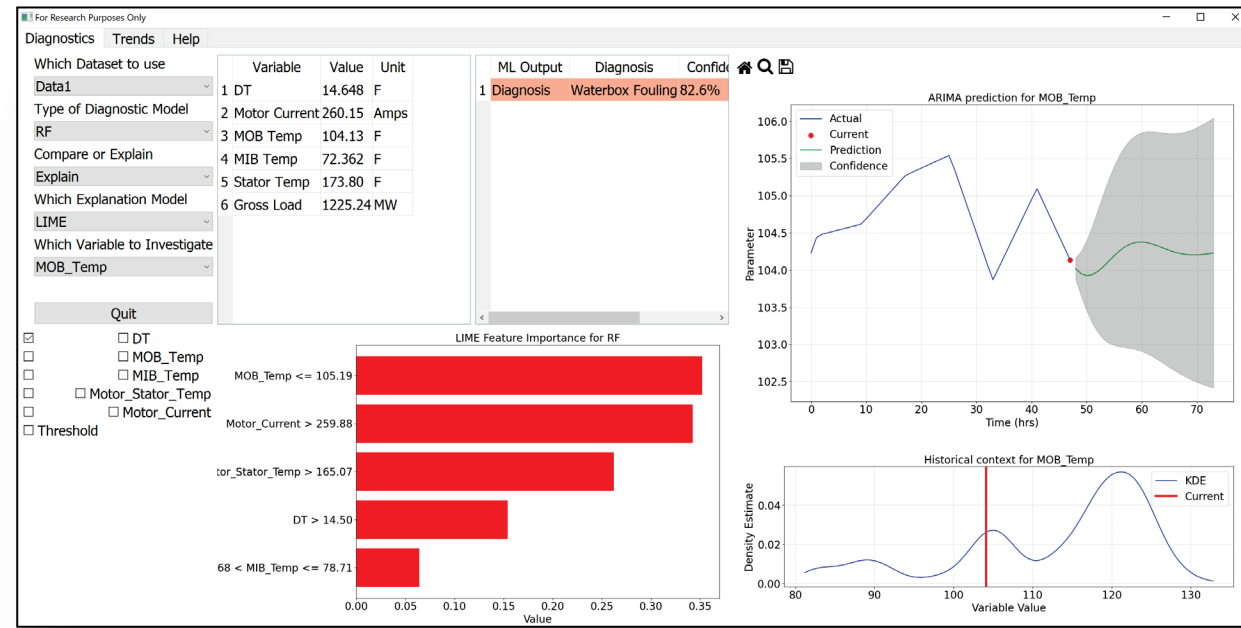
# Interpretability of Artificial Intelligence and Machine Learning Technologies for building Trust Among Users

 LWRS Program researchers developed methods to address the explainability, performance, and trustworthiness of AI/ML to enhance the interpretability of outcomes.

 One method uses objective metrics like Local Interpretable Model-agnostic Explanations (LIME) and Shapley Additive Explanations (SHAP).

 Another method employs user-centric visualization of AI/ML outcomes together with objective metrics to support expert interpretation.

 In collaboration with Public Service Enterprise Group (PSEG), Nuclear LLC, performed initial demonstration of the technical basis on circulating water system (CWS) for a waterbox fouling problem.



*User-centric visualization with performance and explainability metrics*

Explainable Artificial Intelligence Technology for Predictive Maintenance, [INL/RPT-23-74159](#)

# AI/ML Applications for Anomaly Detection

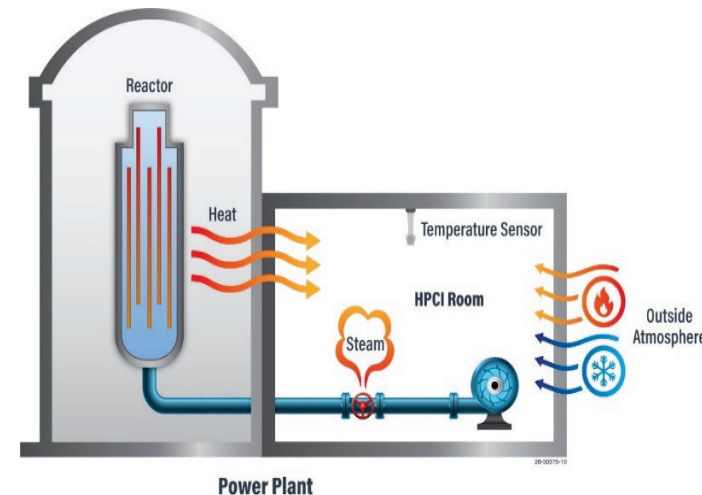
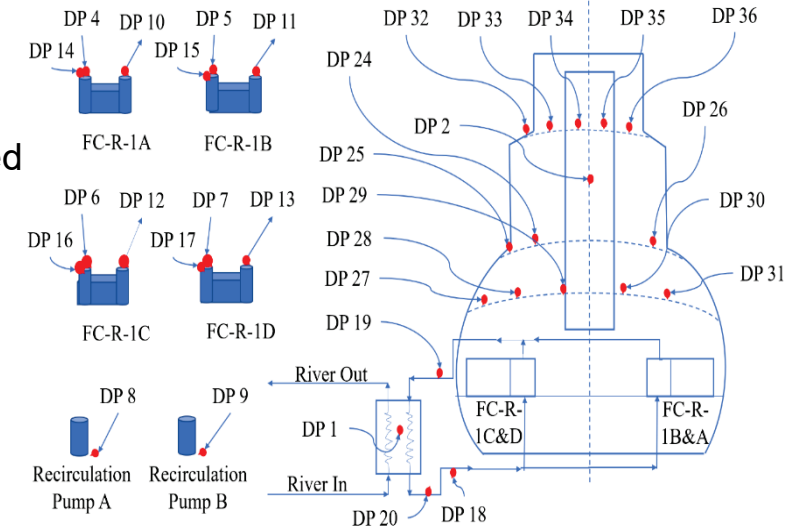
- Unsupervised ML Methods
  - Developing equipment-agnostic anomalies detection methods by holistic inference of process conditions

Extending Data-Driven Anomaly Detection Methods..., [INL/RPT-23-73933](#)  
Software - Automated Latent Anomaly Recognition Method, (ALaRM)

- Semi-Supervised ML Methods
  - Developing methods to couple text-mined information from sparse condition reports to equipment and process sensors data for equipment and process reliability analysis

Feature Extraction for Subtle Anomaly Detection Using Semi-Supervised Learning, *Annals of Nuclear Energy*, vol. 181, pp. 109503, 2023  
<https://doi.org/10.1016/j.anucene.2022.109503>

Recurrent neural networks were used to predict dry well cooling fan failure using surrounding sensors

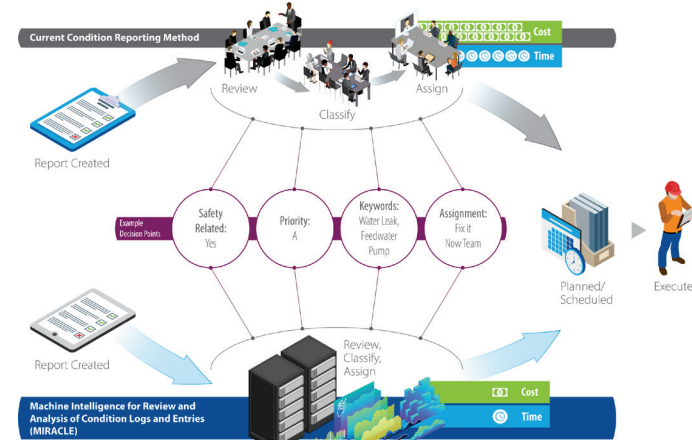


Clustering methods were used to detect High-Pressure Coolant Injection (HPCI) system steam leak

# NLP Applications for Process Improvement

- Natural Language Processing and Deep Learning Methods:
  - Demonstrating and evaluating the automation of the condition reports screening process (which is the review and classification of condition reports according to their impact on nuclear safety).
  - Evaluating the automation of document review, sampling, trending, analysis, and reporting.
  
- Natural Language Processing and Clustering Methods:
  - Developing an inventory optimization method by coupling work demand information with parts inventory to reduce the minimum stocking requirements.

Condition report screening is a process that involves several staff on daily or bi-daily basis for several hours a week.



Software: Machine Intelligence for Review and Analysis of Condition Logs and Entries ([MIRACLE](#))



Developing AI/ML methods to optimize the stocking requirements in a plant

Explainable Artificial Intelligence Technology for Predictive Maintenance, [INL/RPT-23-74159](#)

# Computer Vision Applications for Process Improvement

- Computer Vision and Deep Learning Methods:
  - Developing methods to automatically identify a fire in a video stream to augment the effectiveness of the fire watch program
  - Developing and evaluating the automation of logging analog gauges (i.e., a method to recognize gauges in oblique angles and read their values)
  - Demonstrating methods for drones to autonomously recognize and navigate their environment in a nuclear power plant.

Automated gauge reading impacts a wide spectrum of activities in a plant including operator rounds, gauges calibration, and peer verification, and improves data fidelity for online monitoring.



Patent - Automated Gauge Reading And Related Systems, Methods, And Devices

Example of AI/ML's ability to accurately identify fire and smoke.



Automating Fire Watch in Industrial Environments through Machine Learning-Enabled Visual Monitoring, [INL/EXT-19-55703](#) Software - Modelling Framework for Fire and Smoke Detection in Imagery

Drones can automate several activities in a plant including operator and security rounds, and inspections of hazardous locations.



Software - Route-operable Unmanned Navigation of Drones ([ROUNDS](#))

# AI/ML Research Summary

ML for Material Management  
ML for Equipment Monitoring  
ML for Anomaly Detections

NLP Applications  
Computer Vision Applications

AI/ML Explainability

Artificial intelligence, machine learning, associated methods and data handling techniques are relatively new in the nuclear power industry.

- Show great promise for automating many manually performed activities
- Are demonstrating new approaches to enhance efficiency

Collaborative efforts with owner-operators and others emphasize many non-safety uses.

- Adoption must align with the nuclear safety culture of the industry.
- Some uses demonstrate ability to rapidly transition to safety important uses.

Trust in automation, understandability affect usability.

- Human factors issues in AI/ML implementation vital to adoption and safe use.



# Sustaining National Nuclear Assets

*[lwrs.inl.gov](http://lwrs.inl.gov)*



**Ahmad Al Rashdan, Ph.D.**

November 15, 2023

# Fire Watch





# What is a Fire Watch?

Regulatory Guide 1.189 Section 2.2.1 states:

*Work involving ignition sources such as welding and flame cutting should be carried out under closely controlled conditions. Persons performing such work should be trained and equipped to prevent and combat fires. In addition, a person qualified in performing hot-work fire watch duties should directly monitor the work and function as a fire watch.*



Image From: U.S. Department of Labor 2019.

Regulatory Guide 1.189 Section 2.4.C states:

*Successful fire protection requires inspection, testing, and maintenance of the fire protection equipment. A test plan that lists the individuals and their responsibilities in connection with routine tests and inspections of the fire protection systems should be developed. The test plan should contain the types, frequency, and detailed procedures for testing. Frequency of testing should be based on the code of record for the applicable fire protection system. Procedures should also contain instructions on maintaining fire protection during those periods when the fire protection system is impaired or during periods of plant maintenance (e.g., fire watches).*

## Motivation

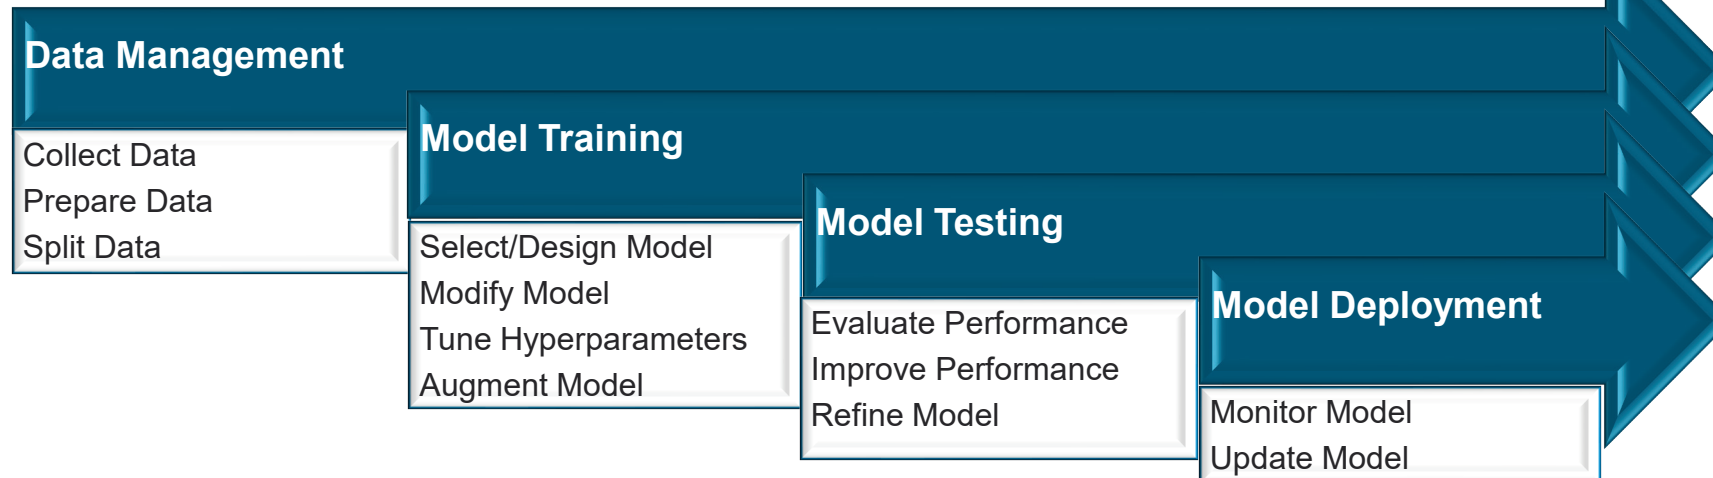
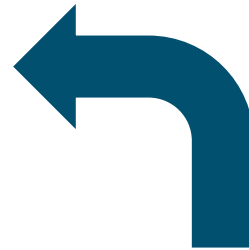
- Fire watch could cost an excess of \$1M per month in a nuclear power plant that implement a fire protection program under Appendix R to Part 50—Fire Protection Program for Nuclear Power Facilities Operating Prior to January 1, 1979.
- In 2019, the Utilities Service Alliance was awarded an iFOA grant titled “Advanced Remote Monitoring” to research and develop automation and advanced remote monitoring technology into the United States nuclear fleet to achieve economic viability while maintaining or improving safety and reliability. This includes the fire watch process.
- As part of the award, INL is working with the Utilities Service Alliance and its members through a Cooperative Research and Development Agreement to research, develop, and evaluate a custom-made fire cart for fire watch.
- The cart is equipped with a suite of sensors (e.g., camera, infrared cameras, smoke detectors). The sensor signals are fused to reliably detect fire.
- A camera can be used to detect fire in a video stream or image of various industrial environments using machine learning models.

# Considered Automation Options

Method	Advantage	Technical Limitation
<b>Image Processing</b>	No massive data collection or training needed	Reliable manual feature engineering is needed
<b>Spatial Machine Learning</b>	Do not have to be explicitly programmed to extract specific features	Require large amounts of diverse image training data to achieve adequate generalization
<b>Spatial and Temporal Machine Learning</b>	Consider changes in time as an additional dimension for features extraction	Require large amounts of diverse video training data to achieve adequate generalization

# Outline

- Data Collection and Preparation
- Models Architecture
  - Unlocking Pre-trained Models
  - Ensemble Classifier
- Models Performance
- Use Considerations



Fire Detected?  
Yes/No



Steps to Develop a Machine Learning Model

# Data Collection

## General Image/Video Sources:

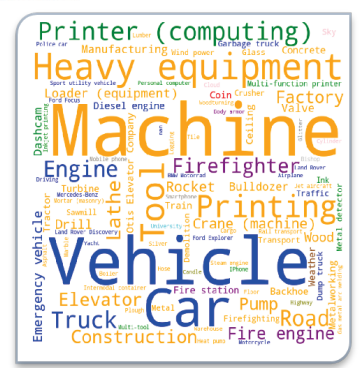
- YouTube-8M
- Google Images
- Yahoo Flickr Creative Commons 100 Million (YFCC100m)

## Targeted Image/Video Sources:

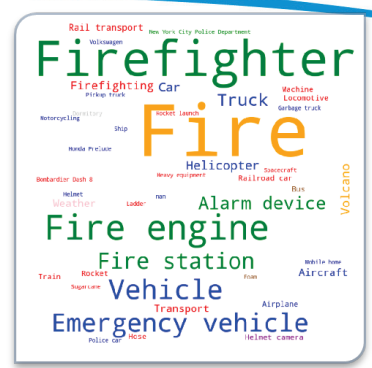
- Fire Smoke (FiSmo)

## Using Existing Classifiers:

- ImageNet (e.g., candle, cannon, fire screen, geyser, missile, space shuttle, stove, torch, volcano)



No-Fire Labels



Fire Labels

Parameter	Description
Environment	Target area, e.g., indoors or outside
Detection Space	The size of the target area and distance of potential fire targets from the sensor
Objects	Objects within the target area (e.g., people, desk, light source, automobile, etc.)
Light Source	Light consists of artificial or natural sources, or both
Light Duration	Light source and duration vary by time of day, season, etc.
Sensor	Spatial, spectral, and temporal resolutions; stationary or moving

Diversity of Collected Data

# Data Preparation

- Selectively identified noisy regions of frames and either crop the frames to remove the unwanted features, or carefully blur the region by using various image processing tools
- Ensured sufficient temporal separation between consecutive video frames. Usually, this temporal separation was around two to three seconds between extracted frames
- Dataset was slightly biased toward fire images having 30% more than non-fire. This class imbalance was meant to bias the model toward making a fire prediction.

Dataset	Fire Images	Normal Images
Training	29,519	20,056
Validation	3,690	2,507
Test	3,690	2,507
Total	36,899	25,070

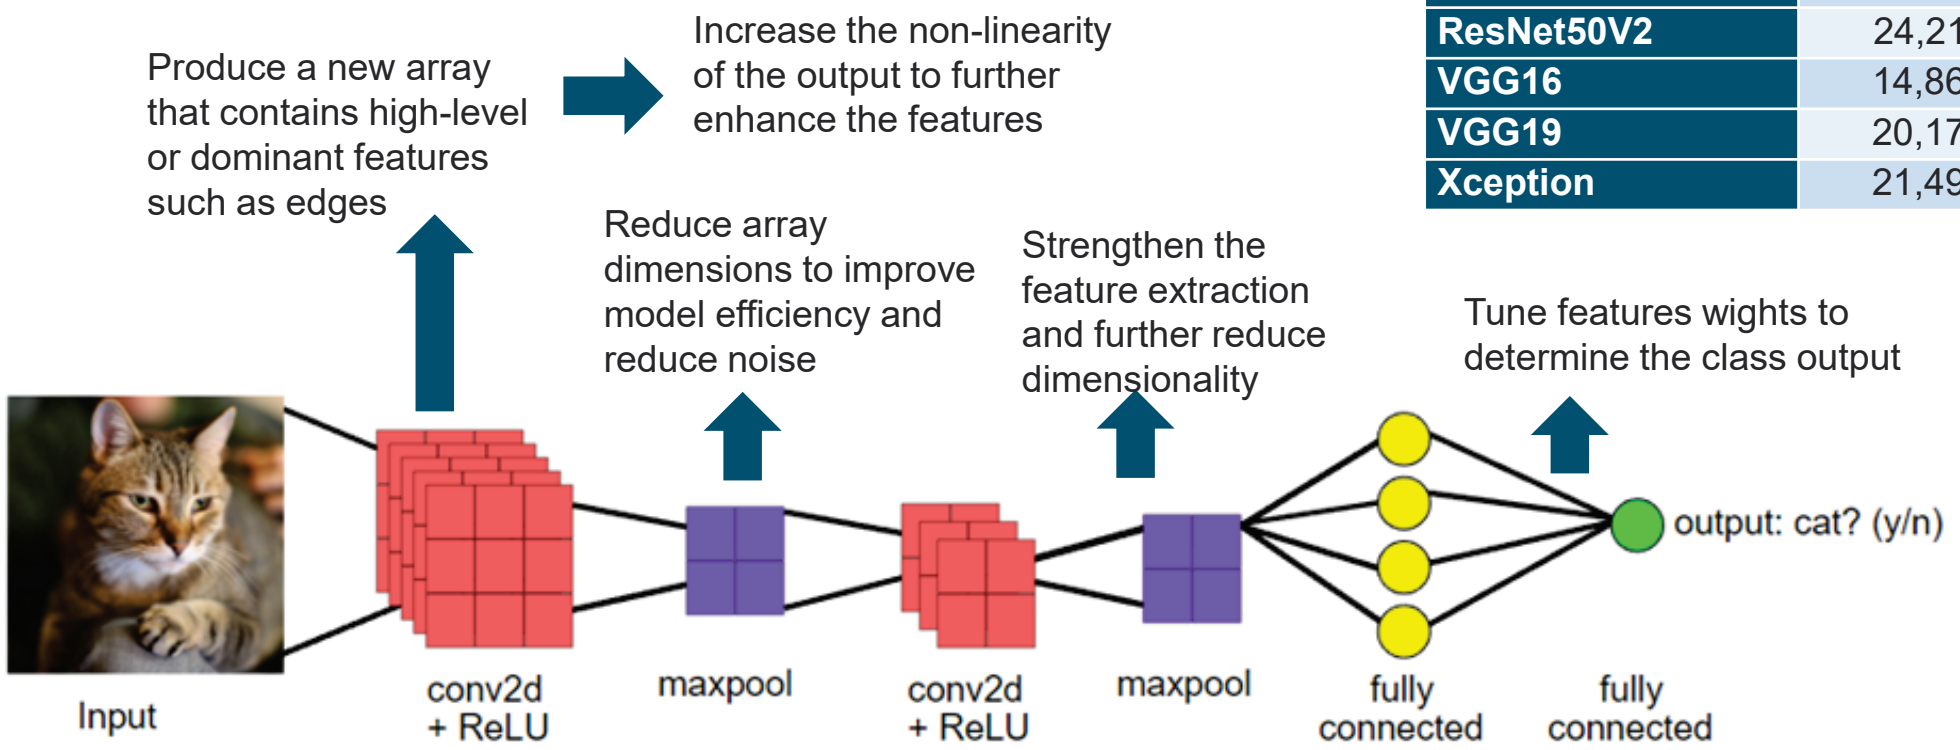
Allocation of Collected Fire Data

Dataset	Smoke Images	Normal Images
Training	4,460	8,636
Validation	558	1,080
Test	558	1,080
Total	5,576	10,796

Allocation of Collected Smoke Data

# Models Architecture

Model	Total Parameters	Total Layers
EfficientNetB7	64,652,242	815
InceptionResNetV2	54,647,906	782
InceptionV3	22,263,970	313
ResNet101	43,245,058	347
ResNet50	24,226,818	177
ResNet50V2	24,211,586	192
VGG16	14,862,146	21
VGG19	20,171,842	24
Xception	21,499,178	134



An example CNN architecture containing two convolution and pooling segments feeding into a densely connected layer for classification (from Google developers 2019).

# Unlocking Pre-trained Models

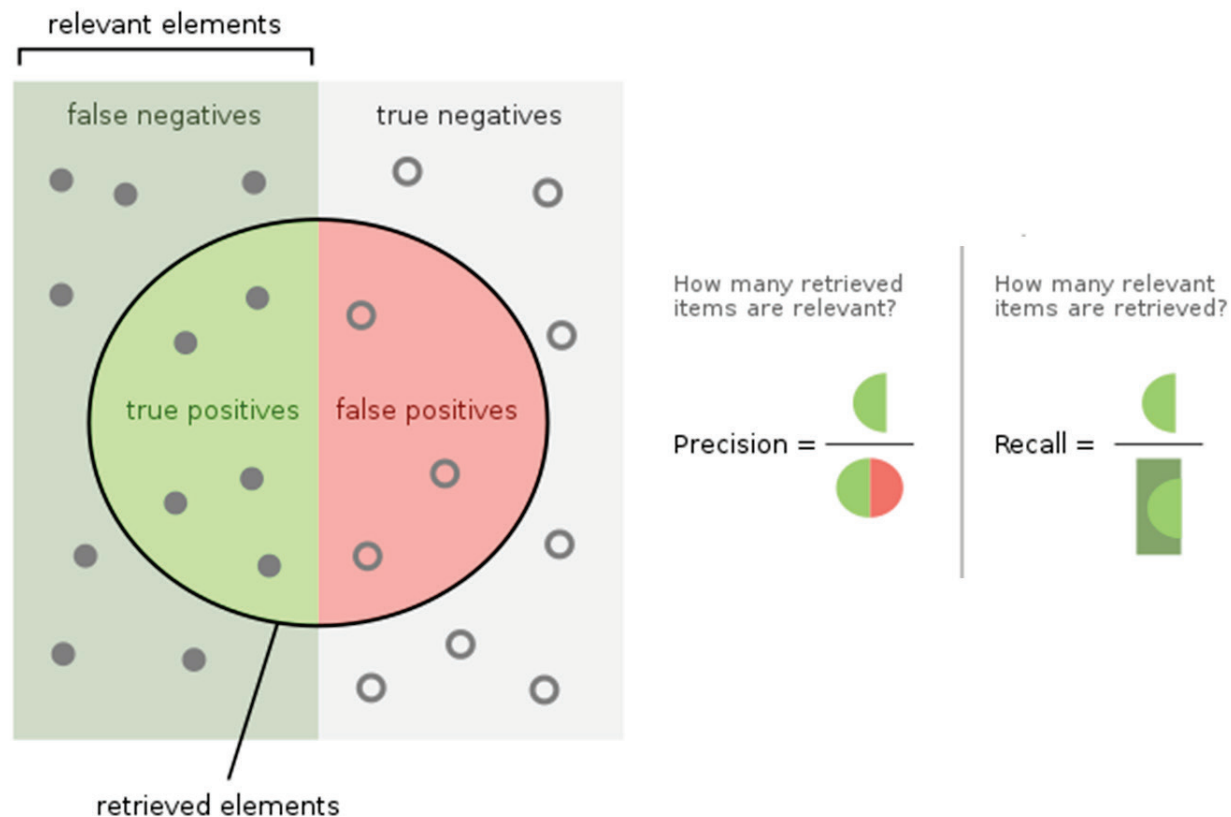
- Transfer learning unlocks part of the pre-trained layers and its parameters for the training process.
- One model was trained for every 10% increment of trainable layers for each of the nine model architectures.

	EfficientNetB7	InceptionResNetV2	InceptionV3	ResNet101	ResNet50	ResNet50V2	VGG16	VGG19	Xception
<b>Dense</b>	865,282	371,714	495,618	692,226	692,226	692,226	147,458	147,458	692,226
<b>10%</b>	27,109,026	13,212,354	6,569,154	15,668,226	8,573,442	8,571,906	2,507,266	2,507,266	6,942,618
<b>20%</b>	39,025,122	23,384,930	11,610,498	19,880,194	15,670,274	15,926,274	7,226,882	7,226,882	9,094,586
<b>30%</b>	51,840,450	31,170,978	13,713,794	23,499,266	17,905,666	17,897,474	9,586,690	11,946,498	11,246,554
<b>40%</b>	55,959,698	36,662,274	15,595,778	27,121,154	19,880,194	19,871,234	13,126,658	16,666,114	13,398,522
<b>50%</b>	59,948,634	42,292,706	17,285,058	31,330,306	22,054,402	22,042,626	13,716,738	17,846,274	15,550,490
<b>60%</b>	62,169,034	47,779,746	18,826,882	34,951,682	23,057,666	23,045,122	14,306,818	19,026,434	17,702,458
<b>70%</b>	63,780,554	53,142,770	21,200,738	38,898,946	23,553,410	23,540,482	14,601,986	19,911,682	19,854,426
<b>80%</b>	64,363,866	53,558,866	21,559,282	42,075,906	24,017,154	24,002,818	14,823,426	20,133,122	21,308,362
<b>90%</b>	64,579,542	54,003,266	21,842,658	43,035,394	24,141,570	24,126,978	14,860,354	20,170,050	21,444,682
<b>100%</b>	64,652,242	54,647,906	22,263,970	43,245,058	24,226,818	24,211,586	14,862,146	20,171,842	21,499,178



# Models Performance

- $F_\beta = (1 + \beta^2) \cdot \frac{precision \cdot recall}{(\beta^2 \cdot precision) + recall}$
- $F_\beta$  score uses the  $\beta$  parameter to weight the relative importance of recall versus precision. A  $\beta$  score greater than one weights the  $F_\beta$  score in favor of recall.
- By setting  $\beta = 2$ , the F-score is weighted such that recall is considered twice as important as precision and a reliable detection is considered twice as impactful as a false alarm.

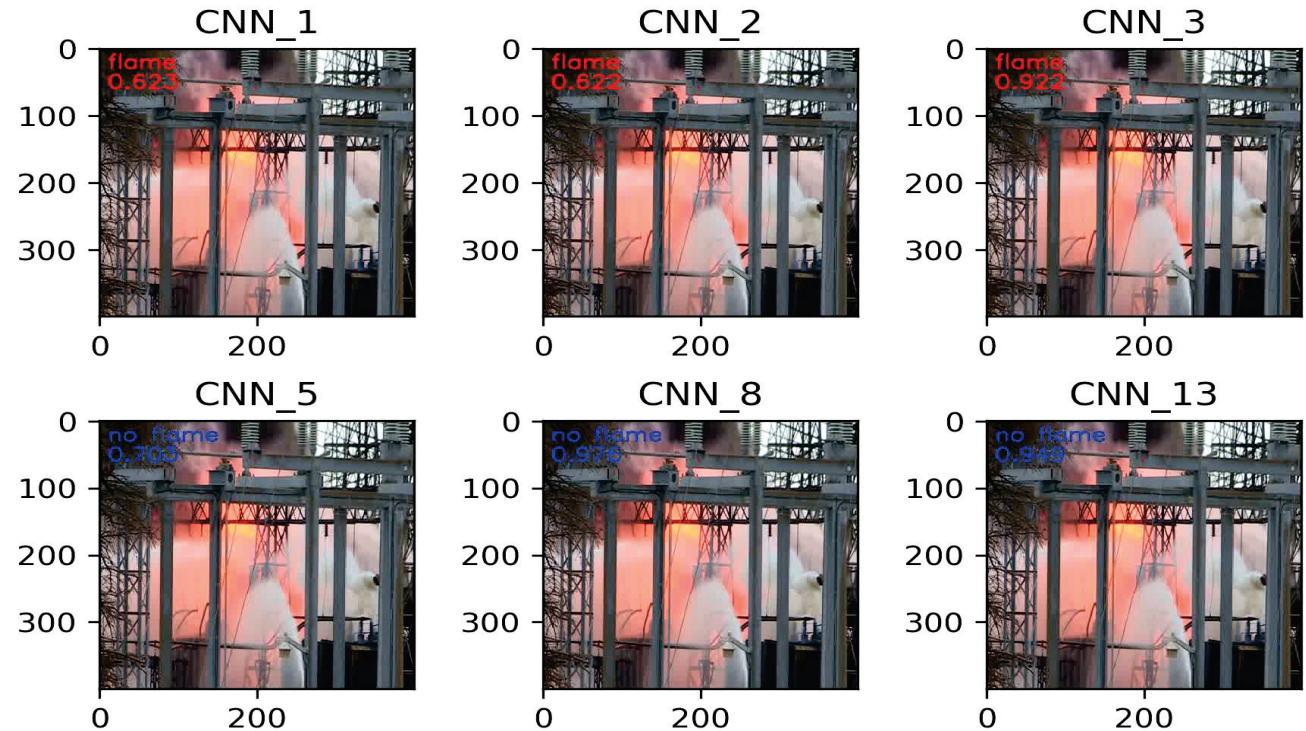


Source: [https://en.wikipedia.org/wiki/Precision\\_and\\_recall](https://en.wikipedia.org/wiki/Precision_and_recall)

# Models Performance

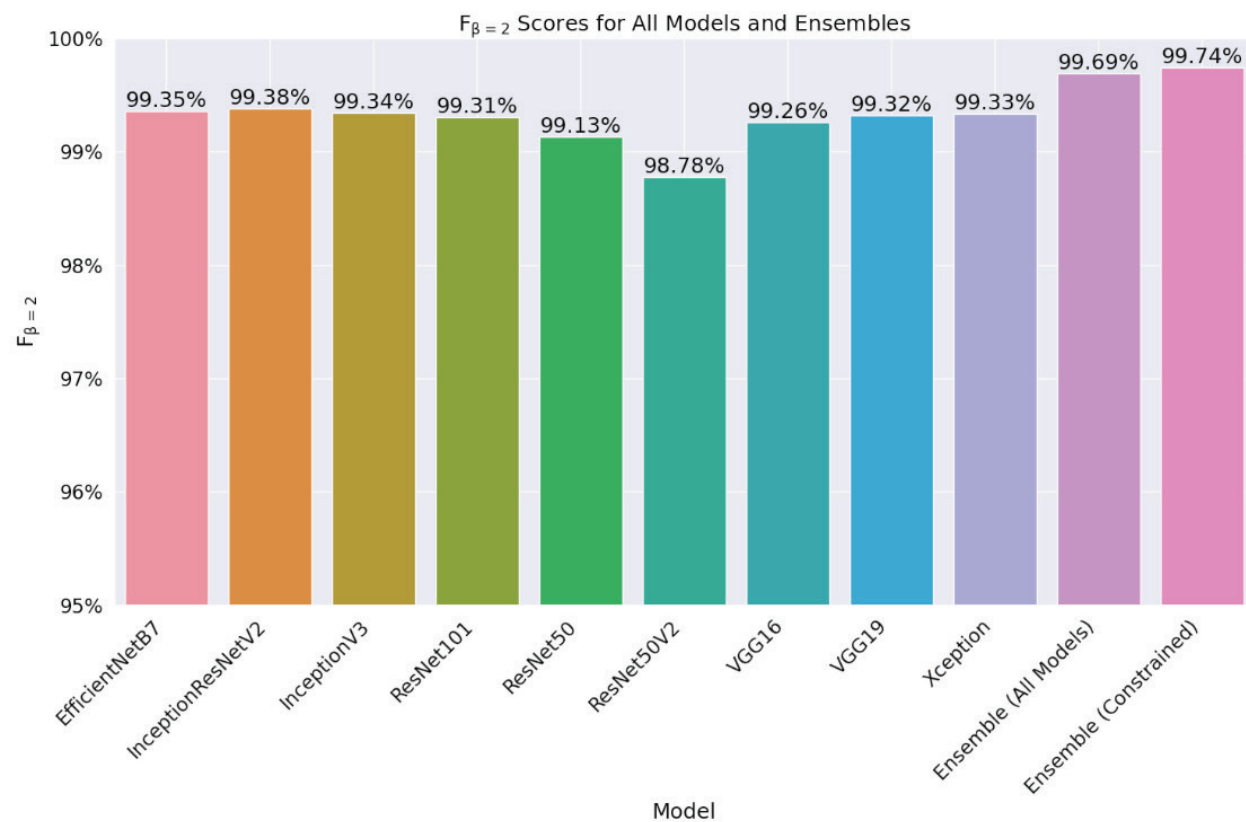
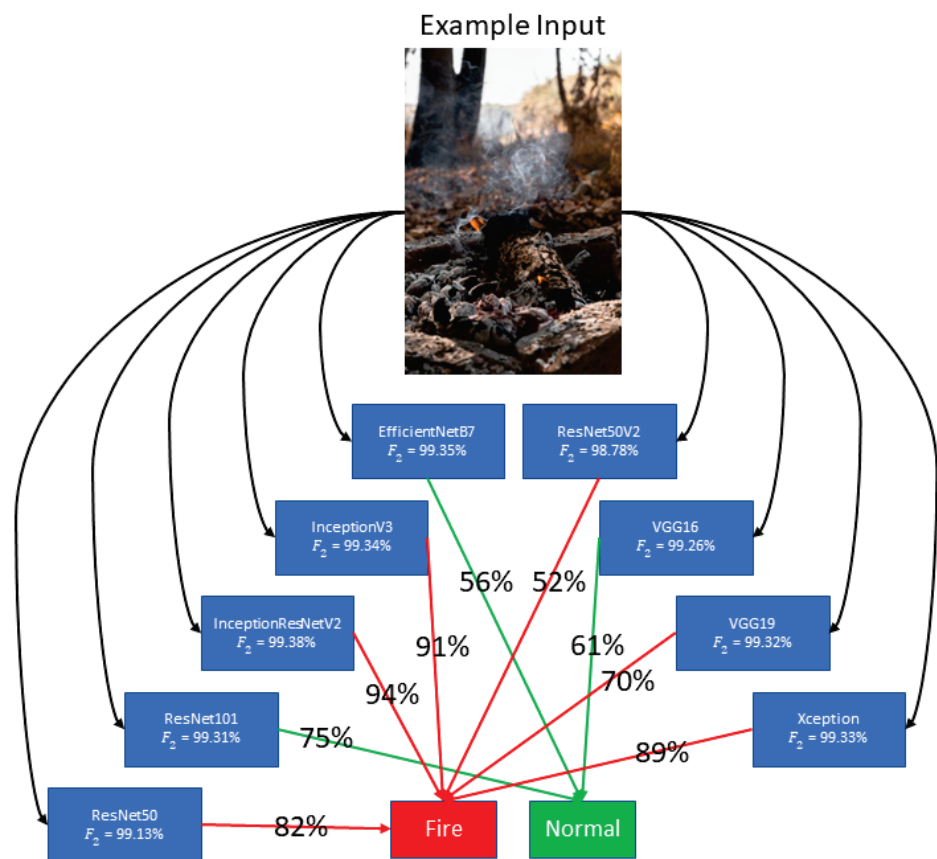
$F_{\beta=2}$  Scores for Fire Predictions Versus Percentage of Trainable Model Layers

Model	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
EfficientNetB7	99.35%	96.49%	97.16%	96.76%	97.11%	97.37%	97.59%	97.22%	97.36%	97.71%	97.45%
InceptionResNetV2	98.91%	99.17%	99.10%	99.16%	99.12%	99.15%	99.09%	99.15%	99.28%	99.21%	99.38%
InceptionV3	99.03%	98.97%	99.14%	99.01%	99.02%	99.11%	99.02%	99.18%	99.27%	99.14%	99.34%
ResNet101	99.12%	98.79%	98.94%	99.01%	99.02%	99.16%	99.06%	99.18%	99.31%	99.14%	99.23%
ResNet50	98.97%	98.94%	98.80%	98.95%	98.90%	98.96%	99.00%	99.06%	98.99%	99.03%	99.13%
ResNet50V2	98.77%	98.22%	98.56%	98.42%	98.66%	98.34%	98.27%	98.43%	98.42%	98.78%	98.44%
VGG16	97.88%	98.17%	99.21%	99.03%	99.06%	99.26%	98.99%	99.09%	99.05%	99.12%	99.24%
VGG19	98.27%	98.11%	98.99%	99.32%	99.29%	99.28%	99.25%	99.02%	99.13%	99.14%	99.16%
Xception	98.96%	99.08%	99.12%	99.12%	99.13%	99.13%	99.13%	99.19%	99.33%	99.28%	99.27%



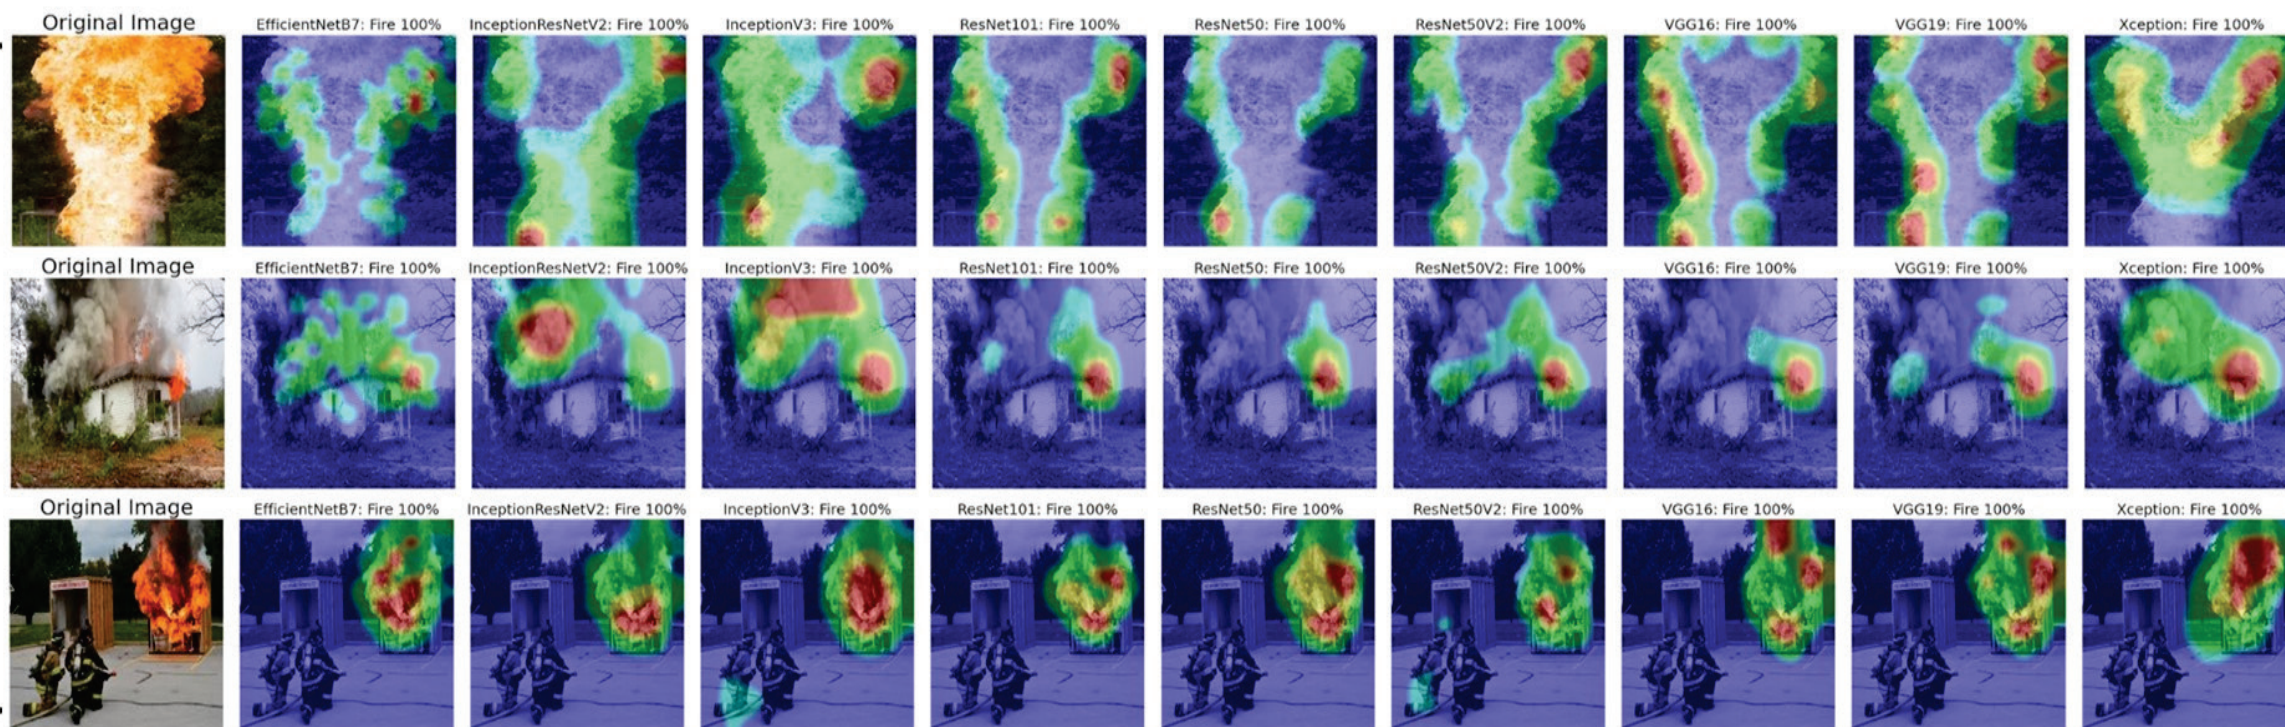
Heatmap representation of  $F_2$  score metrics for fire detection

# Ensemble Classifier



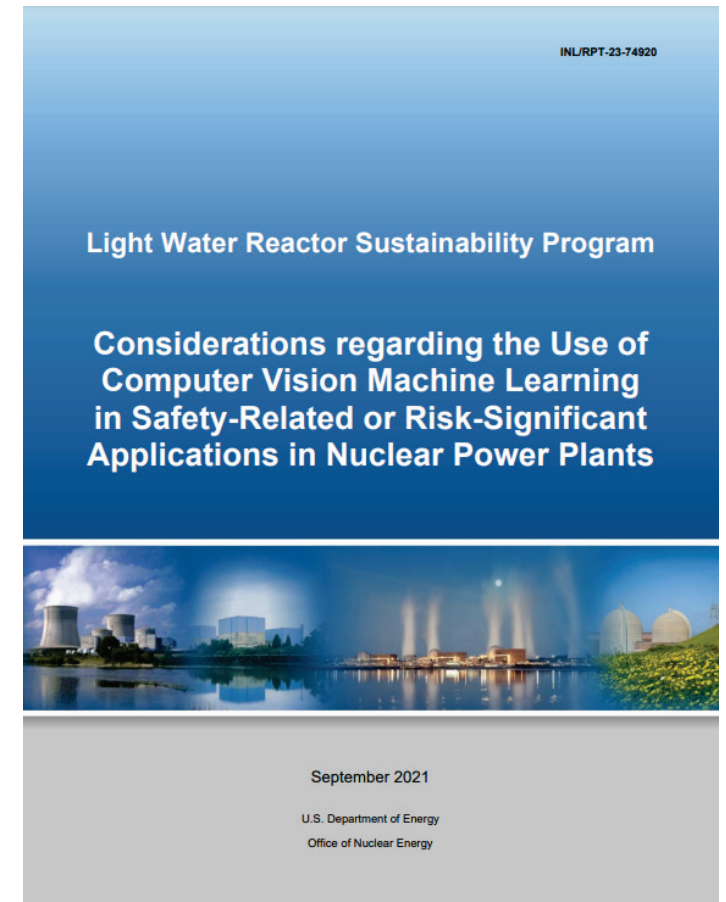
# Models Focus

Fire Detection



# Use Considerations

- Digital Instrumentation and Controls (DI&C) regulatory requirements would need to be satisfied for any and every AI application that impacts safety related and risk-significant applications
- Research into the fire watch models' compatibility with the current safety standards
- Aligns with the directions of the recent executive order on safe, secure, and trustworthy artificial intelligence use. Specifically, it provides insight for development of new standards for AI safety and security



To evaluate how example AI technologies align with the safety framework, and discusses how they could be analyzed, modeled, tested, and validated in a manner similar to typical DI&C technologies.

<https://lwrs.inl.gov/SitePages/Reports.aspx>

# Compatibility with the Current Standards Requirements

Characteristic/Consideration	Independence	Defense in Depth	CCF	V&V	QA	Configuration Control	Cyber Security	CGD	Maintainability	Traceability	Design Control	Repeatability	Deterministic Nature	Explainability	Reliability	FMEA	Simplicity	Justification	Trustworthiness
Open-source data and model																			
Frequent updates to source																			
Massive amounts of data																			
Periodic training																			
Probabilistic and stochastic																			
Various performance metrics																			
Incomprehensible to reviewers																			
Inherited bias																			
Non-systematic approach																			
Robustness to new conditions																			
Special skillset																			

## Example of Compatibility Considerations

Models often utilize open-source datasets and feature extraction engines or models.

It is not always possible to determine the level of overlap among open-source datasets. Open-source models could use similar fundamental concepts. This impacts the independence of the developed models:

- Causes the system to be susceptible to common cause failure
- Overestimates the software verification results
- Introduces a cybersecurity concern



Methods to create independent datasets may be needed (e.g., GANs)



Methods to quantify independence may be needed



# Sustaining National Nuclear Assets

*[lwrs.inl.gov](http://lwrs.inl.gov)*



# THE GOOD, THE BAD AND THE UGLY OF AI IN PROCESS CONTROL

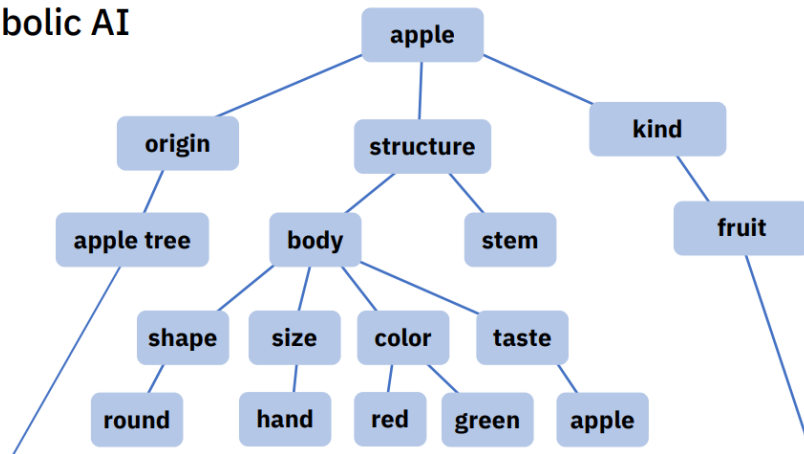
Missy Cummings, PhD

George Mason University

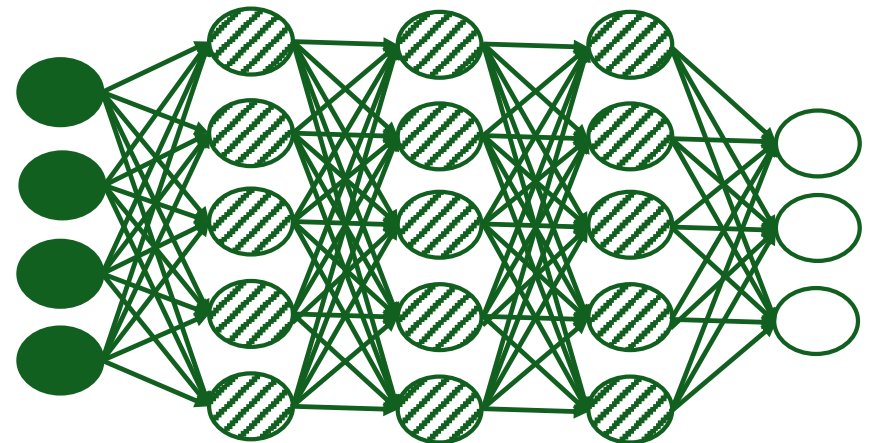


# Just what is AI?

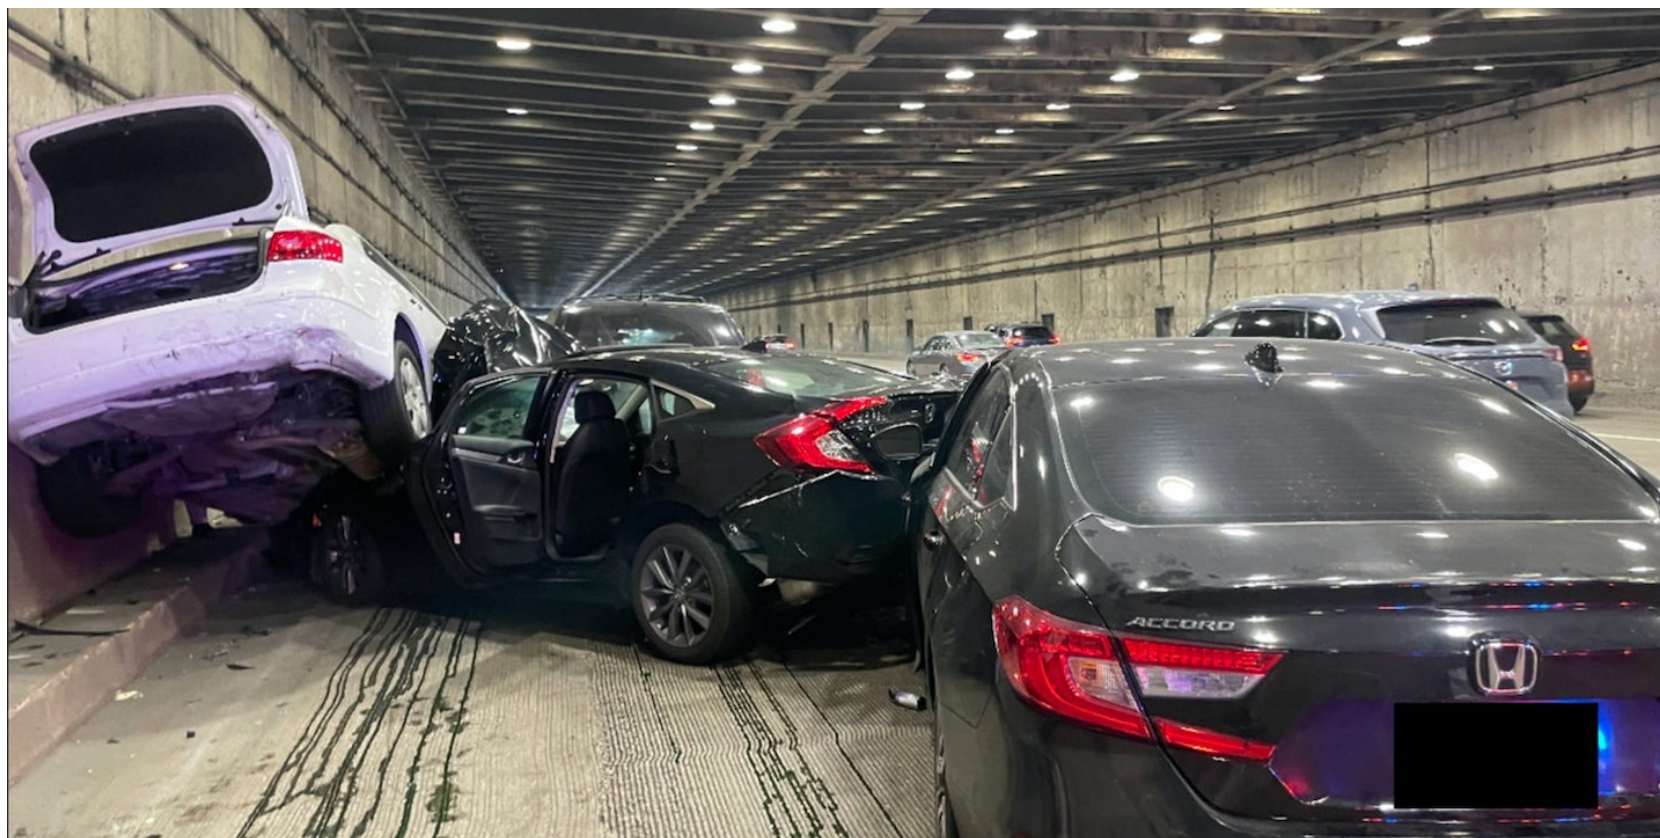
## Symbolic AI



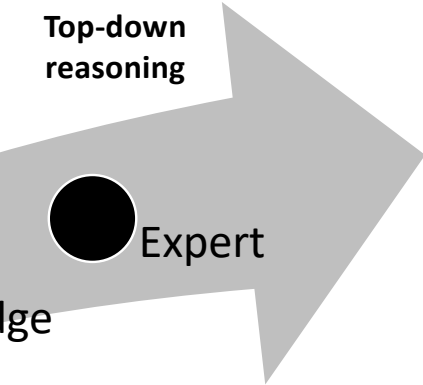
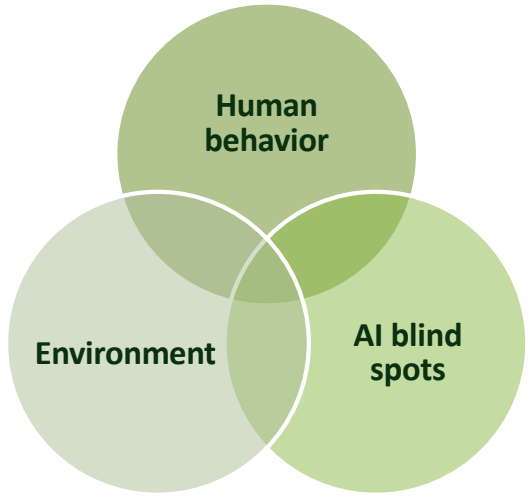
## Neural Networks



## AI Problems



# Autonomy/AI & Reasoning



Uncertainty

Expert

Knowledge

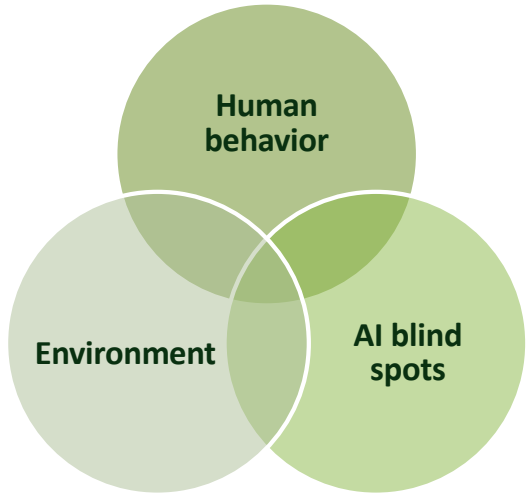
Rule

Skill

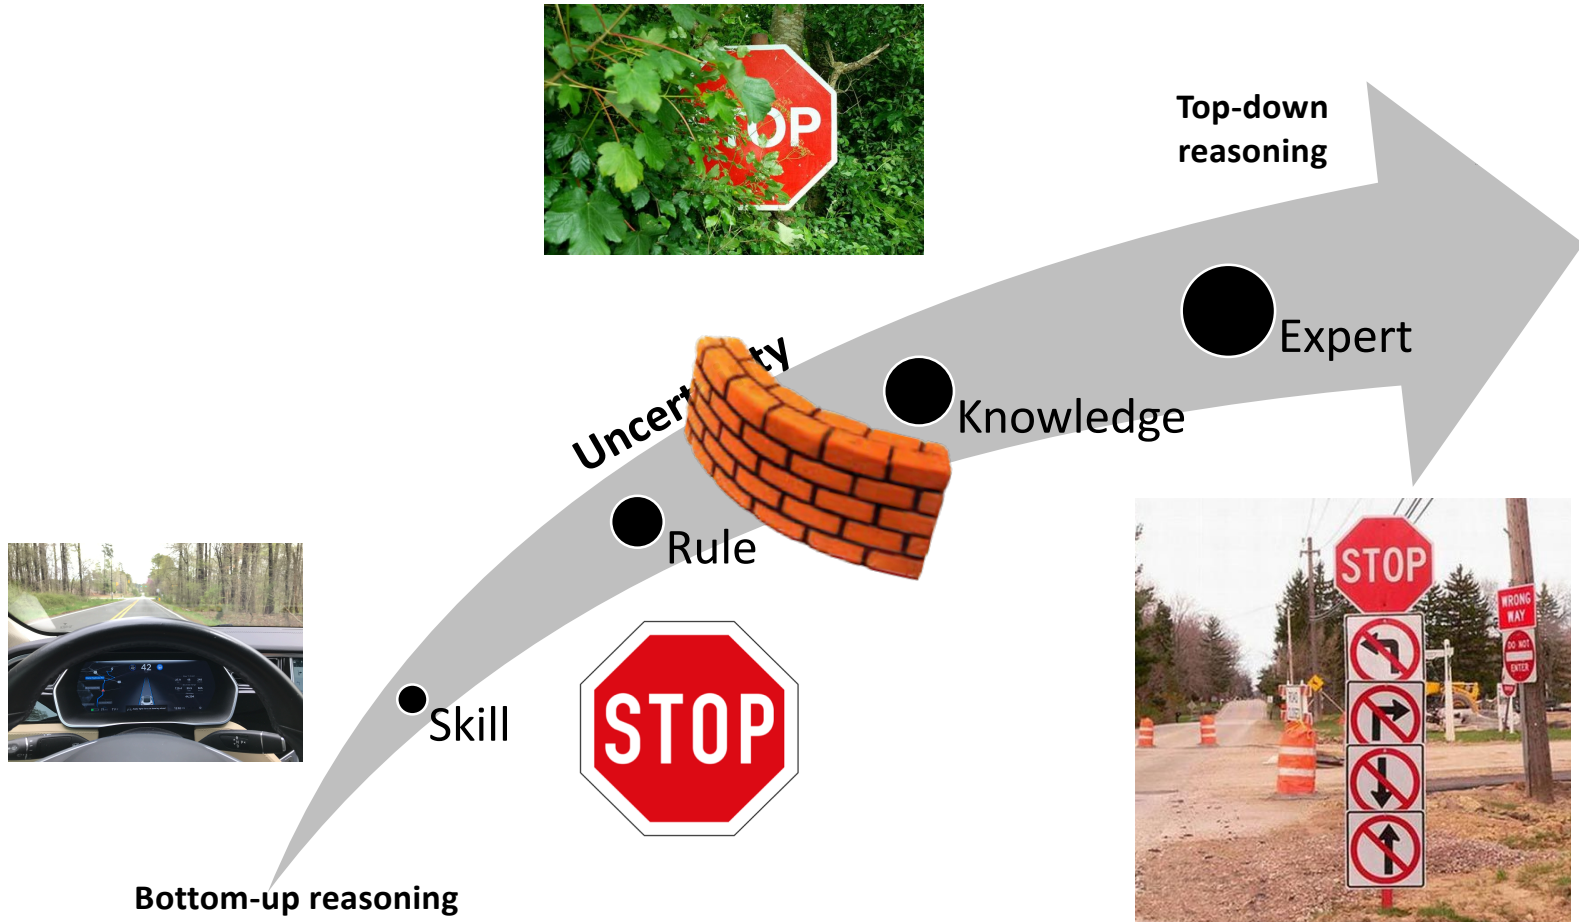


Bottom-up reasoning

# Autonomy/AI & Reasoning

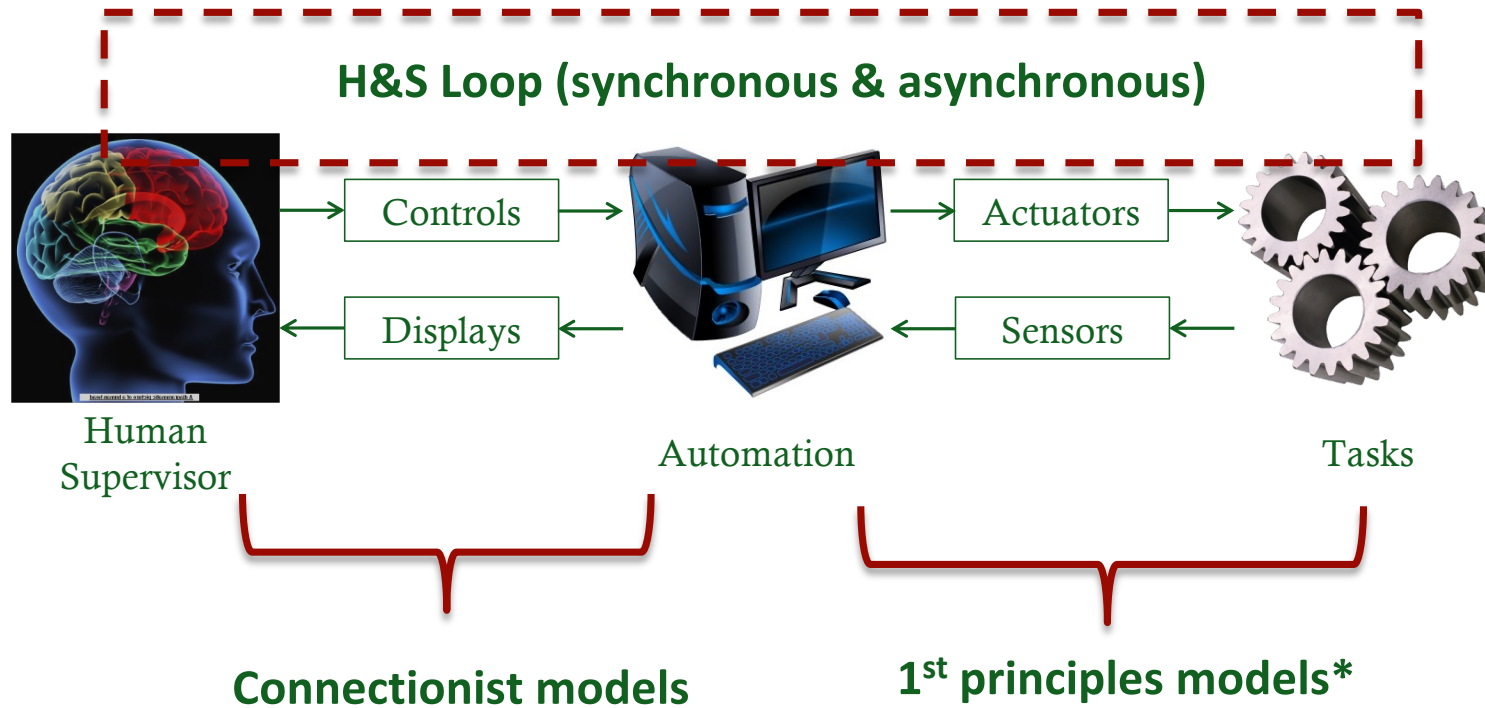


# The Uncertainty Wall



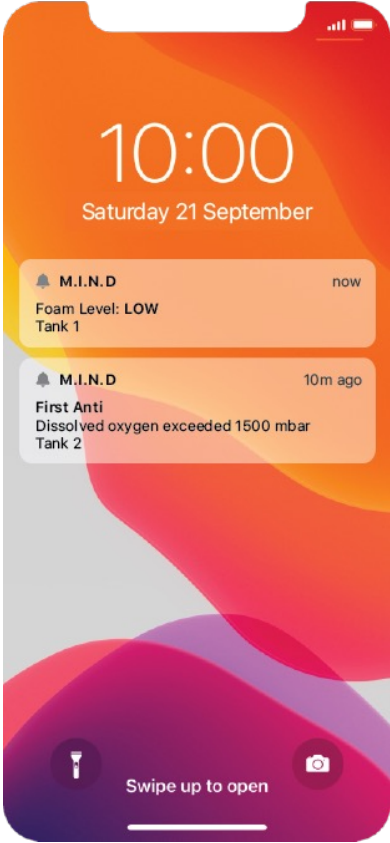
# AI in Process Control

Symbolic/ Connectionist



\*Uncertainty dependent

# Process Control w/ Unmodeled Uncertainty



**M!ND** CAMERA ALERTS ARCHIVE

Select Shift Time  
Start: 2020/11/06 02:10 PM End: 2020/11/06 02:10 PM

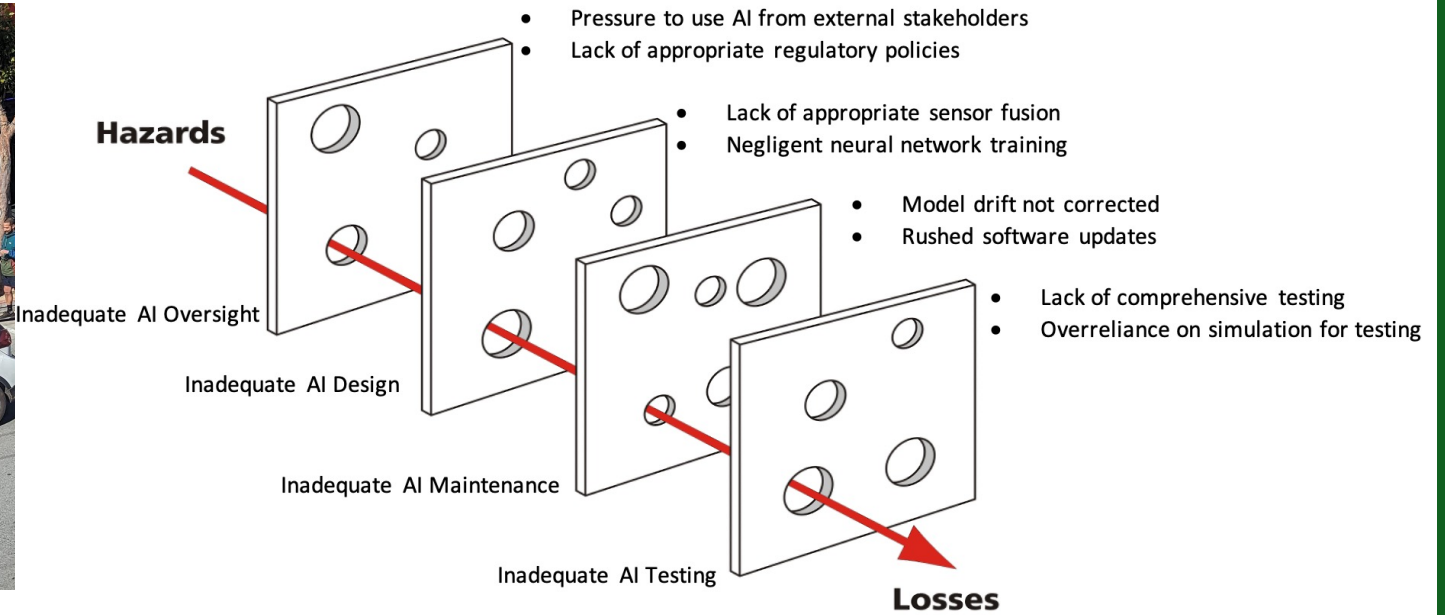
Tank 1 Tank 2  
Tank 3 Tank 4

PUnt	TEMP	JTEMP	F_Mix_2	pH	pO2	AIR_SP	O2_SP	WEIGHT	TURB
1	20.3 °C	73.0 °C		6.98 pH	-76567.8 %sat	0.0 lpm	0.0 lpm	-0.32 kg	0.10 AU
2	23.9 °C	125.2 °C	0 %	6.94 pH	-109893.4 %sat	0.0 lpm	0.0 lpm	0.00 kg	-0.02 AU
3	36.0 °C	36.0 °C		6.73 pH	83.5 %sat	3.1 lpm	0.0 lpm	27.84 kg	0.00 AU
4	35.9 °C	36.0 °C		6.73 pH	58.5 %sat	3.0 lpm	0.0 lpm	26.80 kg	0.00 AU
5	35.9 °C	36.0 °C		7.69 pH	25.8 %sat	3.0 lpm	0.0 lpm	31.09 kg	0.00 AU
6	36.0 °C	36.2 °C		5.71 pH	87.3 %sat	3.0 lpm	0.0 lpm	28.23 kg	0.00 AU





# AI Hazard Analysis



## What about Large Language Models?



Correct my grammar in the following text:

**factor negative x cubed plus 7 times x squared minus 16 times x plus 12**

View other drafts



Sure, I can help you with that. Here is the factored form of the polynomial:

Code snippet

```
(x-2) (x-3) (x-4)
```

grammatically correct, street names are capitalized in English.

UNIVERSITY

# Don't Trust and Definitely Verify



why did the Cruise self-driving car hit the back of the bus?



[View other drafts](#) 



There are a few possible reasons why the Cruise self-driving car hit the back of the bus. One possibility is that the car's sensors did not detect the bus in time. Another possibility is that the car's software made a mistake in interpreting the sensor data. Finally, it is also possible that the car's driver was not paying attention and did not take action to avoid the collision.

Cruise has not released any information about the cause of the accident, so it is impossible to say for sure what happened. However, Cruise has said that it is investigating the accident and that it is committed to safety.

The accident is a reminder that self-driving cars are still in their early stages of development and that there is still a lot of work to be done to make them safe. However, it is also important to remember that human drivers are involved in millions of accidents every year, and that self-driving cars have the potential to reduce the number of accidents and make our roads safer.

# LLMs & Predictive Maintenance



Feature	LR	Method 1	Method 2	Method 3	Method 4
1. Speed limit	14	11	11	11	-
2. AADT	1	3	1	9	-
3. Access control	15	7	16	15	-
4. Left shoulder width	6	14	5	10	-
5. Right shoulder width	16	12	15	13	-
6. Number of lanes	11	13	8	12	-
7. Median width	5	6	12	16	-
8. Section length	3	4	2	5	-
9. Light	7	9	6	3	-
10. Weather	13	5	13	7	-
11. Maximum age	9	10	9	6	-
12. Minimum age	10	8	10	8	-
13. Vehicle type	2	1	3	1	1
14. Sobriety	4	2	4	2	2
15. Urban / Rural	8	15	7	4	-
16. Lane width	12	16	14	14	-

## The thorny path ahead

- The hype cycle
  - Follow the money
- Human + AI collaboration is key
- Cybersecurity & disinformation
- Inappropriate code in safety-critical systems
  - Hardware-software integration
- Workforce development
  - AI fact checking
  - AI maintenance
  - AI risk management is key



Questions?