

Today's Panel



**Reyner Fouch (M)**  
Director  
KPMG US



**Balaji Padmanabhan**  
Associate Dean for Strategic Initiatives  
University of Maryland



#ERM26



2

AI Risk & Governance

Bajali Padmanabhan  
University of Maryland



3



#ERM26



4

Today's AI (and Magic)

- Can do some amazing things because of brilliant engineering “under the hood” that makes it seem like magic!



#ERM26



5

What is Artificial Intelligence?

- Building machines that demonstrate (human) intelligence.
  - Turing Test
- What do humans do that could be considered intelligent?
  - (and how do we build machines that could do that?)



#ERM26



6

Demonstrating Intelligence: Then

```
> Hello, I am Eliza.
* I am afraid.
> How long have you been afraid?
> Since midterms.
> Tell me more...
* I am afraid that I will fail
my course.
> Do you enjoy being afraid that
you will fail your course?
```

EXAMPLE RULE FROM MYCIN

```
IF
1. the infection is primary bacteremia, and
2. the site of the culture is one of the sites listed, and
3. the suspected point of entry of the organism is the
   gastrointestinal tract:
THEN
There is suggestive evidence (0.7) that the identity of
the organism is tetracycline.
```

© 1982 Addison-Wesley, Inc. All rights reserved.  
110, Programming a Expert Rule-Based System  
by Edward S. Shortliffe  
All Rights Reserved. See page 110-111.  
Second Edition, 1983.



#ERM26



7

### Demonstrating Intelligence: Now



8

### AI: Then and Now

- Use-cases look eerily similar, but under the hood it's night and day

**Learning from pre-programmed logic**  
**VS**  
**Learning from continuous, massive data**  
**and new learning paradigms**  
**MACHINE LEARNING**



9

### A (Partial) Machine Learning Taxonomy

- **Supervised Learning, I.e. Predictive Models**
  - Decision Trees
  - Logistic Regressions
  - XGBoost
  - Neural Networks
  - Deep Learning
- **Self-Supervised Algorithms**
  - Learning by "filling in the gaps"
  - Modern Generative AI algorithms
- **Reinforcement Learning Algorithms**
  - Learning through experience, or trial and error
- **Unsupervised Learning**
  - Pattern discovery
  - Clustering algorithms



10

### AI Foundation Models

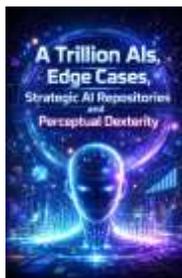
Unlike traditional ML models, here you start with a generic "off-the-shelf" model and adapt to tasks.



Bommasani, Rishi, et al.  
 "On the opportunities and risks of foundation models."  
 arXiv preprint arXiv:2108.07258(2021).



11



12

Meta's Yann LeCun to Launch Physical A.I. Startup After Declaring LLMs a 'Dead End'



1X Neo is a \$20,000 home robot that will learn chores via teleoperation



13

## Today's AI: The Good, The Bad, and The Ugly



14

## The Good

- What capabilities of AI will we see translating into great ideas and products?



15

## The Bad

- Lacks explainability, common sense reasoning and many other things we humans are (supposedly) good at
- Doesn't come with guarantees
- Has shown to be brittle
- Makes up stuff
- Can be biased
- Has high training and inference costs



16

## The Ugly

- Can be used to generate dangerous content and applications
- Can have security flaws
- Has environmental impacts
- Potentially violates copyright
- Learning from human-generated data and insights over the years to potentially replace humans in some tasks
- Benefits the data and compute rich over the others



17

## The Solution: AI Governance

But: Many Frameworks and Messaging Focus on Risks (Alone)

- "Governance refers to the processes, standards, and guardrails that help ensure AI systems and tools are safe and ethical"
  - IBM
- "Governance should be implemented at every stage of the AI lifecycle to manage AI risks effectively"
  - ISO/IEC 42001
  - AWS Security Blog
- Opportunity to broaden the design of tools and frameworks to maximize benefits while managing risk



18

## Nine Issues that good AI Governance Needs to Address

1

Verification architectures for AI use

- Ensuring correctness both of content as well as how it's used



19

Nine Issues that good AI Governance Needs to Address

**2** Responsible and Fair Use

- Responsibility to all stakeholders, consistent with organizational values
- Embracing purpose, explainability, transparency and accountability where needed



20

Nine Issues that good AI Governance Needs to Address

**3** Effective design of (future) work

- Proactively thinking Human+AI architectures



21

Nine Issues that good AI Governance Needs to Address

**4** Ensuring planning for short and long-term horizons

- How does AI use impact future human skills and organizational capabilities?



22

Nine Issues that good AI Governance Needs to Address

**5** Managing and Governing Systems, Not Models or AI alone

- Value is created by systems not algorithms
- Attention on models + humans + processes



23

Nine Issues that good AI Governance Needs to Address

**6** Return on Investment and a Portfolio Strategy

- Importance of developing a portfolio of ideas/projects



24

Nine Issues that good AI Governance Needs to Address

**7** Regulatory Compliance

- Compliance (with regulations)
- Being on top of specific emergent legal concerns (copyright issues, contractual indemnification, patentability of innovations among others)



25

Nine Issues that good AI Governance Needs to Address

8

Risk Management

- Data, security and privacy
- Managing for human gullibility for form over function sometimes
- Ensure red-teaming and adversarial preparation



26

Nine Issues that good AI Governance Needs to Address

9

Vendor Management

- Ensuring AI products and services provided by AI vendors address the aforementioned eight issues



27

Executive Certificate Program

MANAGING GEN AI RISK

Enabling Growth Through Responsible Risk Management



OVERVIEW

For decades, financial institutions have used Enterprise Risk Management (ERM) to help build and manage conventional predictive models. Recently, Generative AI tools have advanced such that leading institutions are embracing AI strategies, raising questions about the role of ERM and other control functions in achieving strategic objectives safely.

PROGRAM AT A GLANCE

- **Format:** In-person
- **Duration:** 2 days
- **Recurrence:** Quarterly  
\*In-person in a custom workshop! Contact Mary Beth Kambick (mkambick@perdell.com)
- **Next Workshop:** March 24-27, 2026



28