Why should you take this course?

Practical considerations and instructor qualifications

In this lecture:

- Practical considerations
- Why learn from Dr. David?

Practical Considerations

- Get AIGP certified
- Pursue role in AI governance
- Round out resume (e.g., privacy, cybersecurity, legal, data analysis, software engineering)
- Self-enrichment / professional development

Why Dr. David? (1)

- Time
  - Comprehensive AIGP curriculum
  - On demand
  - No textbook or reading required
  - Microlearning
- Money
  - Other courses hundreds or thousands of dollars
- Format
  - Ideal for auditory and visual learners
- No prerequisites
  - No AI, legal, IT, or privacy background required

Why Dr. David? (2)

Education

- Ph.D. University of California, Irvine
- Privacy Engineering Certificate, Carnegie Mellon University
- CIPP/US, CIPM, AIGP, FIP
- GIAC Critical Controls Certification
- Associate of (ISC)^2

Why Dr. David? (3)

- Teaching since 2005
  - 2005-14: K-12
  - 2014-21: college/university (in person and online)
  - 2021-present: privacy workforce development
- Currently privacy analyst at large public sector organization
  - Develop, design, deliver privacy, AI literacy workforce training
- 2022-23: IAPP KnowledgeNet co-chair
- IAPP Global Privacy Summit panel organizer and speaker

Review:

Practical Considerations

- Get AIGP certified
- Pursue role in AI governance
- Round out resume
- Self-enrichment / professional development

Dr. David

- Pain points
  - Time
  - Money
  - Format
  - No prerequisites
- Uniquely qualified
  - Education
  - Experience

What is the AIGP?

Artificial Intelligence Governance Professional

In this lecture:

- What is the AIGP?
- What do AI Governance Professionals do?
- Who should pursue the AIGP?
- What does the AIGP exam cover?
- How do I get certified?
- Exam structure

What is the AIGP?

- Artificial Intelligence Governance Professional
- Certification provided by IAPP
  - Formerly the International Association of Privacy Professionals

What is AI Governance?

- Organization’s use of laws, policies, frameworks, practices
- Organizational, national, international level
- Implementing, managing, overseeing, regulating AI throughout development life cycle
  - Identification, assessment, mitigation of AI risk
- Responsible, ethical, compliant use
  - Includes addressing bias, privacy, misuse
- Increase innovation and trust

What do AIGP’s do?

- Implement responsible AI practices
- Manage risk
- Establish multi-stakeholder oversight body
- Monitor AI systems throughout life cycle
- Address issues such as bias, discrimination, privacy, misuse

Who should pursue the AIGP?

- AI compliance
- Risk management
- Legal and governance
- Data scientists
- AI project managers
- Model operations teams
- Social scientists
- Trust professionals
- Privacy professionals

What does the AIGP cover?

- Body of Knowledge (BoK): 4 domains
  I. Understanding the foundations of AI governance (16-20 questions)
  II. Understanding how laws, standards, and frameworks apply to AI (19-23)
  III. Understanding how to govern AI development (21-25)
  IV. Understanding how to govern AI deployment and use (21-25)

How do I get certified? (1)

- Prepare
  - w/ Dr. David's Certification Masterclass
- Register
  - In-person at test center OR proctored online
  - USD $799 for first-time IAPP non-members
  - USD $649 for first-time IAPP members
  - USD $625 for retake IAPP non-members
  - USD $475 for retake IAPP members, holders of other IAPP certs
- Pass the exam
  - Get results immediately

How do I get certified? (2)

- Maintain
  - Biannual maintenance fee (USD $250 OR included with IAPP membership)
  - Earn 20 Continuing Professional Education (CPE) credits that correspond to AIGP BoK
    - Different than Continuing Privacy Education credits

How is the exam structured?

- 100 questions
  - 85 scored
  - 15 ungraded trial questions
- Question format
  - Multiple choice
  - ~30% of questions are case study format
- 3 hours, with optional 15-minute break
- Scored on scale from 100-500
  - 300 passing score

Review (1):

What is the AIGP?

- Artificial Intelligence Governance Professional
- Certification provided by the IAPP

What do AIGP's do?

- Implement responsible AI practices
- Risk management
- Establish multi-stakeholder oversight body
- Monitor AI systems throughout life cycle
- Address issues such as bias, discrimination, privacy, misuse

Review (2):

Who should pursue the AIGP?

- AI compliance
- Risk management
- Legal and governance
- Data scientists
- AI project managers
- Model operations teams
- Social scientists
- Trust professionals

What does the AIGP cover?

- Foundations of AI governance
- How laws, standards, and frameworks apply to AI
- How to govern AI development
- How to govern AI deployment and use

Review (3):

How do I get certified?

- Prepare
- Register
- Pass
- Maintain

How is the exam structured?

- 100 questions
- Question format
- 3 hours, with optional 15-minute break
- Scored on scale from 100-500

What is the IAPP?

Privacy, AI governance, and digital responsibility

In this lecture:

- What is the IAPP?
- IAPP’s mission
- What is a KnowledgeNet?

What is the IAPP?

- Originally: International Association of Privacy Professionals
  - September 2024: rebranded to “IAPP”
- Founded in 2000
- Largest organization for privacy professionals in the world
- Global conferences
- Resources galore (for members):
  - Newsletters, podcasts, training, reports, surveys, other publications

IAPP’s Mission

- “To define, promote, and improve the professions of privacy, AI governance and digital responsibility globally.”

What is a KnowledgeNet?

- Local IAPP chapter
  - Led by volunteer “Chapter Chairs”
  - Hold regular socials, CPE events
- Located worldwide
  - ~60 in US
  - ~100 throughout Asia, Europe, Canada, Latin America, and Africa

Review:

- IAPP
- KnowledgeNet Chapters

Why get AIGP certified?

Data from the last couple years supports need for AI governance professionals

In this lecture:

- Helen Toner TED Talk
- Dario Amodei’s bold prediction
- Investment
- Consumer perspectives
- Business trends
- Impact on the workforce
- Career advancement

Helen Toner, TED

Two ideas for governing:

Don’t be intimidated by technology or the folks building it.

Focus on adaptability, not certainty.

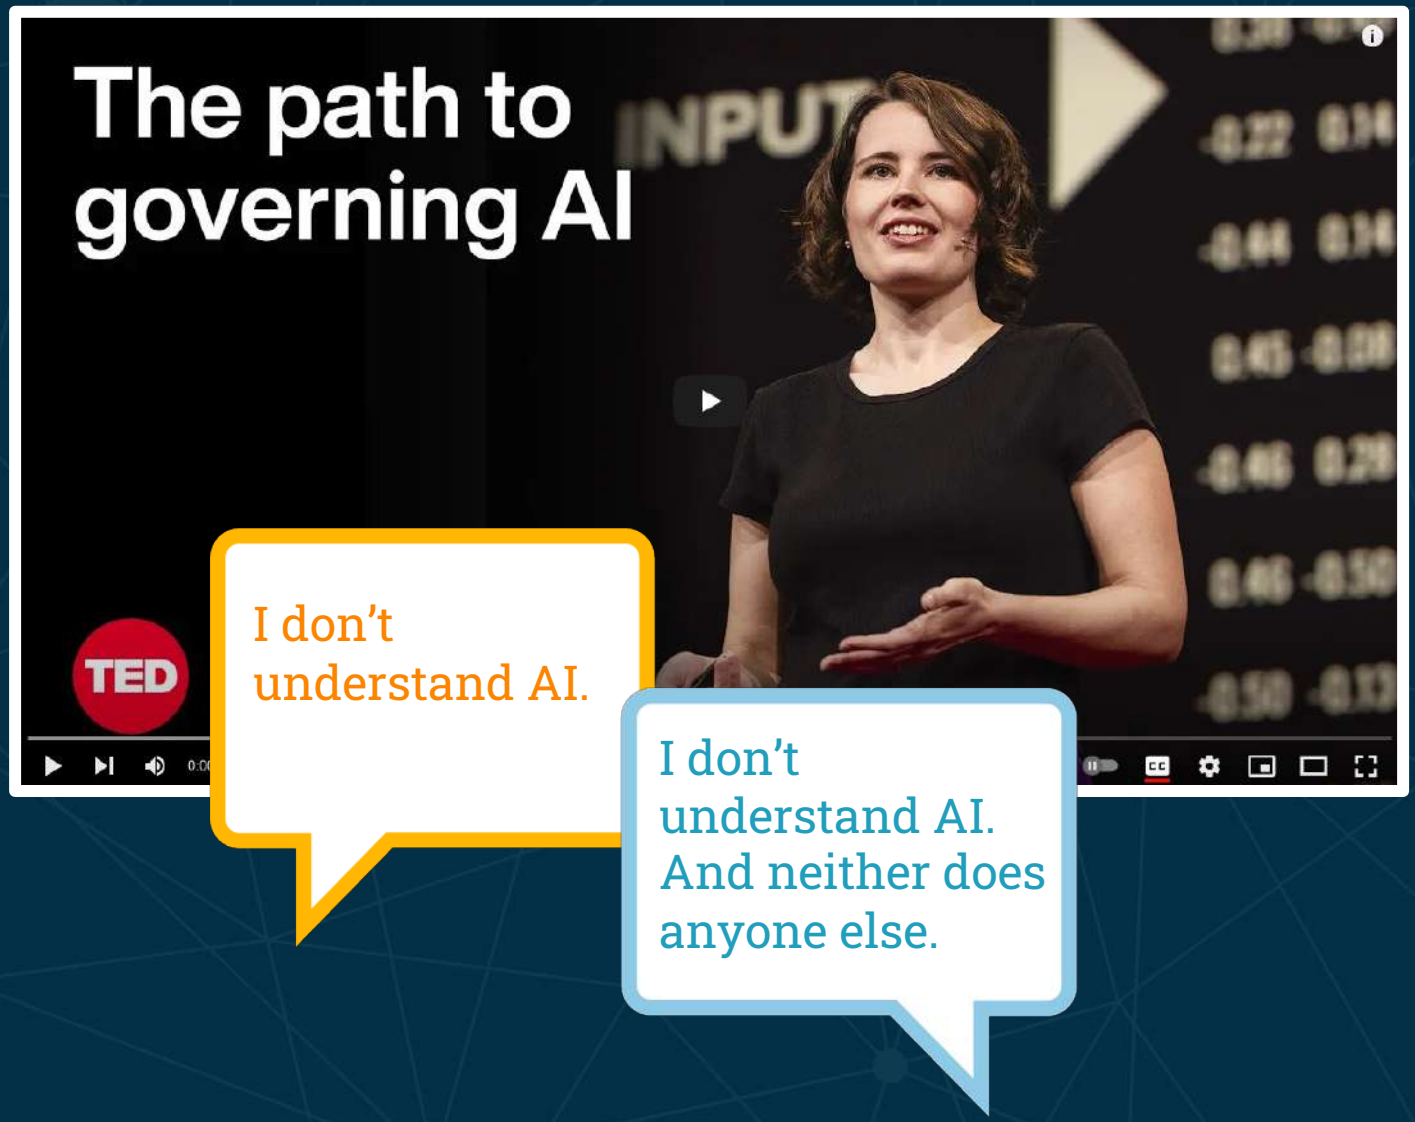

Bold Prediction

"AI will likely become the most powerful and strategic technology in history. By 2027, AI developed by frontier labs will likely be smarter than Nobel Prize winners across most fields of science and engineering. It will be able to use all the senses and interfaces of a human working virtually—text, audio, video, mouse, keyboard control and internet access—to complete complex tasks that would take people months or years.... Imagine a country of geniuses contained in a data center."

Dario Amodei, CEO and co-founder of Anthropic

Investment

- OpenAI: USD $175 billion globally awaiting investment in AI projects
- Microsoft plans to invest $80 billion in AI-enabled data centers in FY25
- GenAI companies: raised $56 billion (worldwide) in 2024, up 192% from 2023
- Global data center spending predicted to reach $250 billion annually

Consumer Perspectives

99% of Americans use at least one AI-enabled product, but 64% don't realize they're using AI

72% expect AI to negatively impact the spread of false information

64% believe AI will negatively affect social connections

60% think AI will negatively impact job opportunities

Americans Use AI in Everyday Products Without Realizing It

Regardless of their own use, Americans perceive significant AI risks and responsibilities

BY ELLYN MAESE

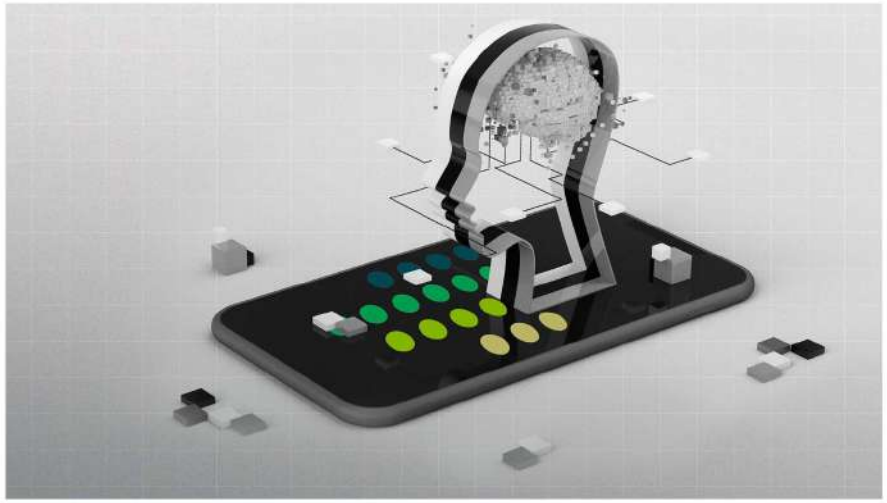

Business Trends

Harvard Business Review (HBR)

C-suite executives from 125 leading companies

98.4% of organizations are increasing AI and data investments

74.8% see value coming from productivity gains and customer service improvements

Analytics And Data Science

6 Ways AI Changed Business in 2024, According to Executives

by Randy Bean

January 2, 2025

Hiroshi Watanabe/Getty Images

Impact on the Workforce

80% of U.S. workers could have at least 10% of their work activities affected by large language models

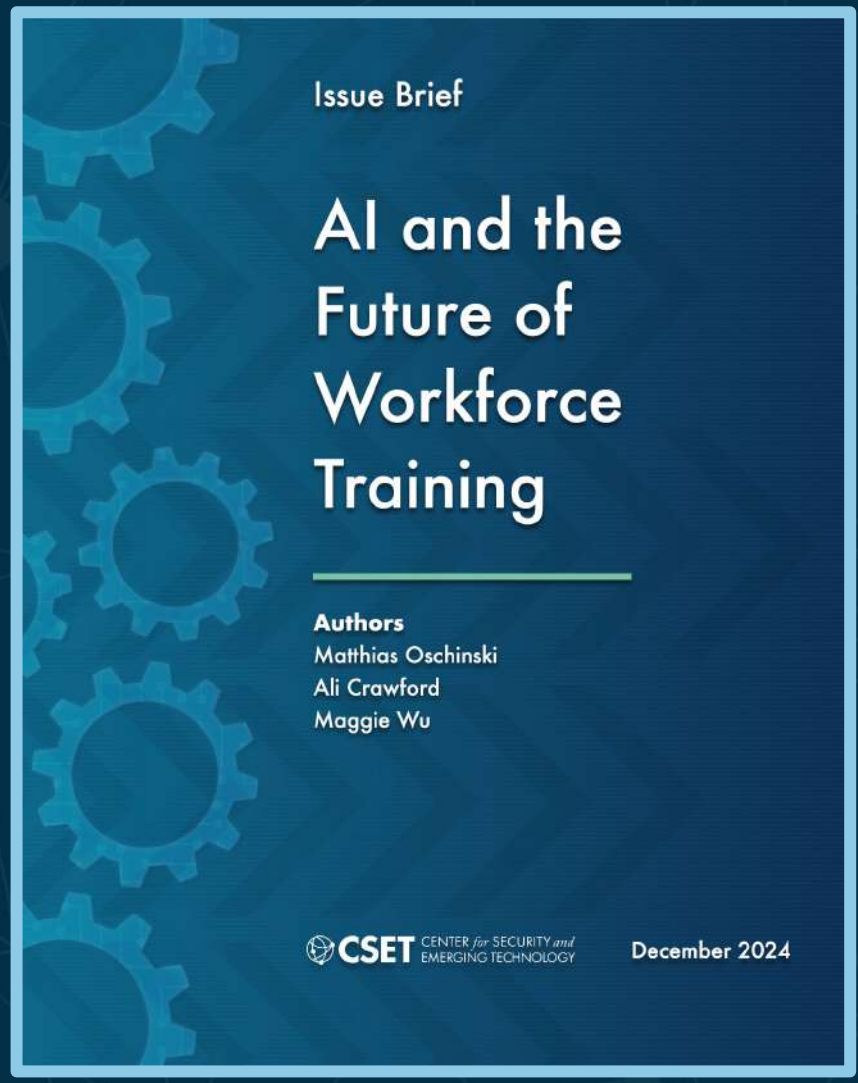

Career Advancement (1)

- Microsoft aims to train 2.5 million Americans in AI skills in 2025
- LinkedIn, “AI Consultant” 2nd fastest growing job in U.S.
- HBR, 84.3% have hired chief data/analytics officers, up from 12.0% in 2012
- HBR, 33.1% have chief AI officers

Career Advancement (2)

- Commitment to learning, self-improvement
- Demonstrates subject matter expertise
- Stand out from non-certificate holders

Review (1):

Helen Toner

- Don’t be intimidated by technology or the folks building it.
- Focus on adaptability, not certainty.

Dario Amodei

- AI will likely become the most powerful and strategic technology in history. By 2027, AI developed by frontier labs will likely be smarter than Nobel Prize winners across most fields of science and engineering.

Review (2):

Investment

- USD $175 billion globally awaiting investment in AI projects
- Microsoft to invest $80 billion in data centers in FY25
- GenAI companies: raised $56 billion (worldwide) in 2024
- Global data center spending predicted to reach $250 billion annually

Consumer protection

- 99% of Americans use at least one AI-enabled product, but 64% don't realize they're using AI
- 72% expect AI to negatively impact the spread of false information
- 64% believe AI will negatively affect social connections
- 60% think AI will negatively impact job opportunities

Review (3):

Business trends

- 98.4% of organizations are increasing AI and data investments
- 74.8% see value coming from productivity gains and customer service improvements

Impact on the workforce

- 80% of U.S. workers could have at least 10% of their work activities affected by LLMs

Review (4):

Career advancement

- Microsoft aims to train 2.5 million Americans in AI skills in 2025
- LinkedIn, “AI Consultant” 2nd fastest growing job in U.S.
- HBR, 84.3% have hired chief data/analytics officers, up from 12.0% in 2012
- HBR, 33.1% have chief AI officers

- Commitment to learning, self-improvement
- Demonstrates subject matter expertise
- Stand out from non-certificate holders

What is risk management?

The identification, assessment, and mitigation of risk

In this lecture:

- What is risk management?
- Typical methodology
- What is a control?
- Two examples

What is risk management?

- Identify, assess, mitigate harm
- Risk: potential harm
- Threat: delivers harm

Typical Methodology

- Risk score = severity of harm * probability of occurrence

What is a control?

- Countermeasure or action designed to modify risk
- Three types
  - Administrative (e.g., training)
  - Technical (e.g., firewall)
  - Physical (e.g., locks, security guards)

Example #1

- Risk: bodily injury
- Threat: car accident
- Mitigation/control: seat belt, obey traffic laws, maintain car

Example #2

- Risk: generation of harmful, offensive content
- Threat: adversarial prompt
- Mitigation/control: meta-prompt, content filter

Review (1):

What is risk management?

- Identify, assess, mitigate harm
- Risk: potential harm
- Threat: delivers harm

Typical methodology

- Risk score = severity of harm * probability of occurrence

Review (2):

What is a control?

- Countermeasure or action designed to modify risk
- Three types
  - Administrative (e.g., training)
  - Technical (e.g., firewall)
  - Physical (e.g., locks, security guards)

Two examples

- Car accident
- Adversarial prompt

What is AI risk management?

Identification, assessment, and mitigation of AI risk

In this lecture:

- Alignment
- AI RM strategy
- Context
- Risk considerations
- Risk impact and response

Alignment (1)

- Two definitions of alignment
- AI alignment: ability of AI systems to pursue an achieve goals that match operators' intended objectives, preferences, and ethical principals
- Risk management alignment: different strategies work together in a complementary fashion to support overall business objectives

Alignment (2)

- Organization’s risk management strategies must be aligned
  - E.g., privacy, cybersecurity, business, operations
- May each have AI component, or
- Have separate AI strategy

AI Risk Management Strategy

- Identify, assess, mitigate risk
- Includes
  - Conduct risk assessment/analysis
  - Identify risks
  - Determine responsibility for mitigation/control implementation
- Document all of this as part of process/report
  - E.g., into existing data governance, privacy impact assessment, authority to operate process

Context

- Risk assessments are context-specific
  - E.g., FRT applications
- Consider
  - AI system owner, operator/deployer
  - Industry/sector
  - Use case
  - Social impacts
  - Timing
  - Jurisdiction(s)

Risk Considerations

- Identify stakeholders (e.g., business, technical)
- AI’s business purpose, planned usages
- Potential harms (e.g., false positive, negative predictions)
- Training data (includes sensitive PII)
- Non-AI alternatives (i.e., is there a different solution?)

Risk Impact and Response

- Typically identified as high, moderate, low
- High: avoid use system; modify system to prevent risk
- Moderate: mitigate risks
- Low: accept or mitigate as appropriate
  - E.g., risk may not reach zero

Review (1):

Alignment

- Organization's risk management strategies must be aligned
- May each have AI component, or
- Have separate AI strategy

AI RM strategy

- Identify, assess, mitigate risk
- Includes
  - Conduct risk assessment/analysis
  - Identify risks
  - Determine responsibility for mitigation/control implementation

Review (2):

Context

- Risk assessments are context-specific
  - E.g., FRT applications
- Consider
  - AI system owner, operator/deployer
  - Industry/sector
  - Use case
  - Social impacts
  - Timing
  - Jurisdiction(s)

Risk considerations

- Identify stakeholders
- AI's business purpose, planned usages
- Potential harms
- Training data (includes sensitive PII)
- Non-AI alternatives

Review (3):

Risk impact and response

- Typically identified as high, moderate, low
- High: avoid use system; modify system to prevent risk
- Moderate: mitigate risks
- Low: accept or mitigate as appropriate