

The Ethics of AI

Utah State Bar
March 15, 2024

Bennett B. Borden
Partner
Chief Data Scientist
DLA Piper



Let's start with what we mean by AI

Classical “AI”

Typically, algorithms, classification, and recommendation systems

Divides (usually) people into groups and treats them according to their group membership

Generative AI

Generates new content based on instructions from a user

How Classical AI Works



Group membership

Group membership describes an **association with two or more people.**

- Identify a **particular characteristic**
- Place individuals or things **into groups** according to an association with the characteristic
- **Action is taken** based on that group membership

Examples

Industry	Characteristic the industry cares about
Insurance	How likely is it you will get sick? Hurt? Die? Break something?
Financial services	Can you pay back money you borrow? How likely is it you will do so?
Employment	Are you the right fit for the job?
Health care	What is wrong with you? How can I fix it?
Security	Are you who you say you are?
Fraud	Are you who you say you are? Are you up to no good?

Understanding Algorithmic Bias




An inaccurate view of “the truth”

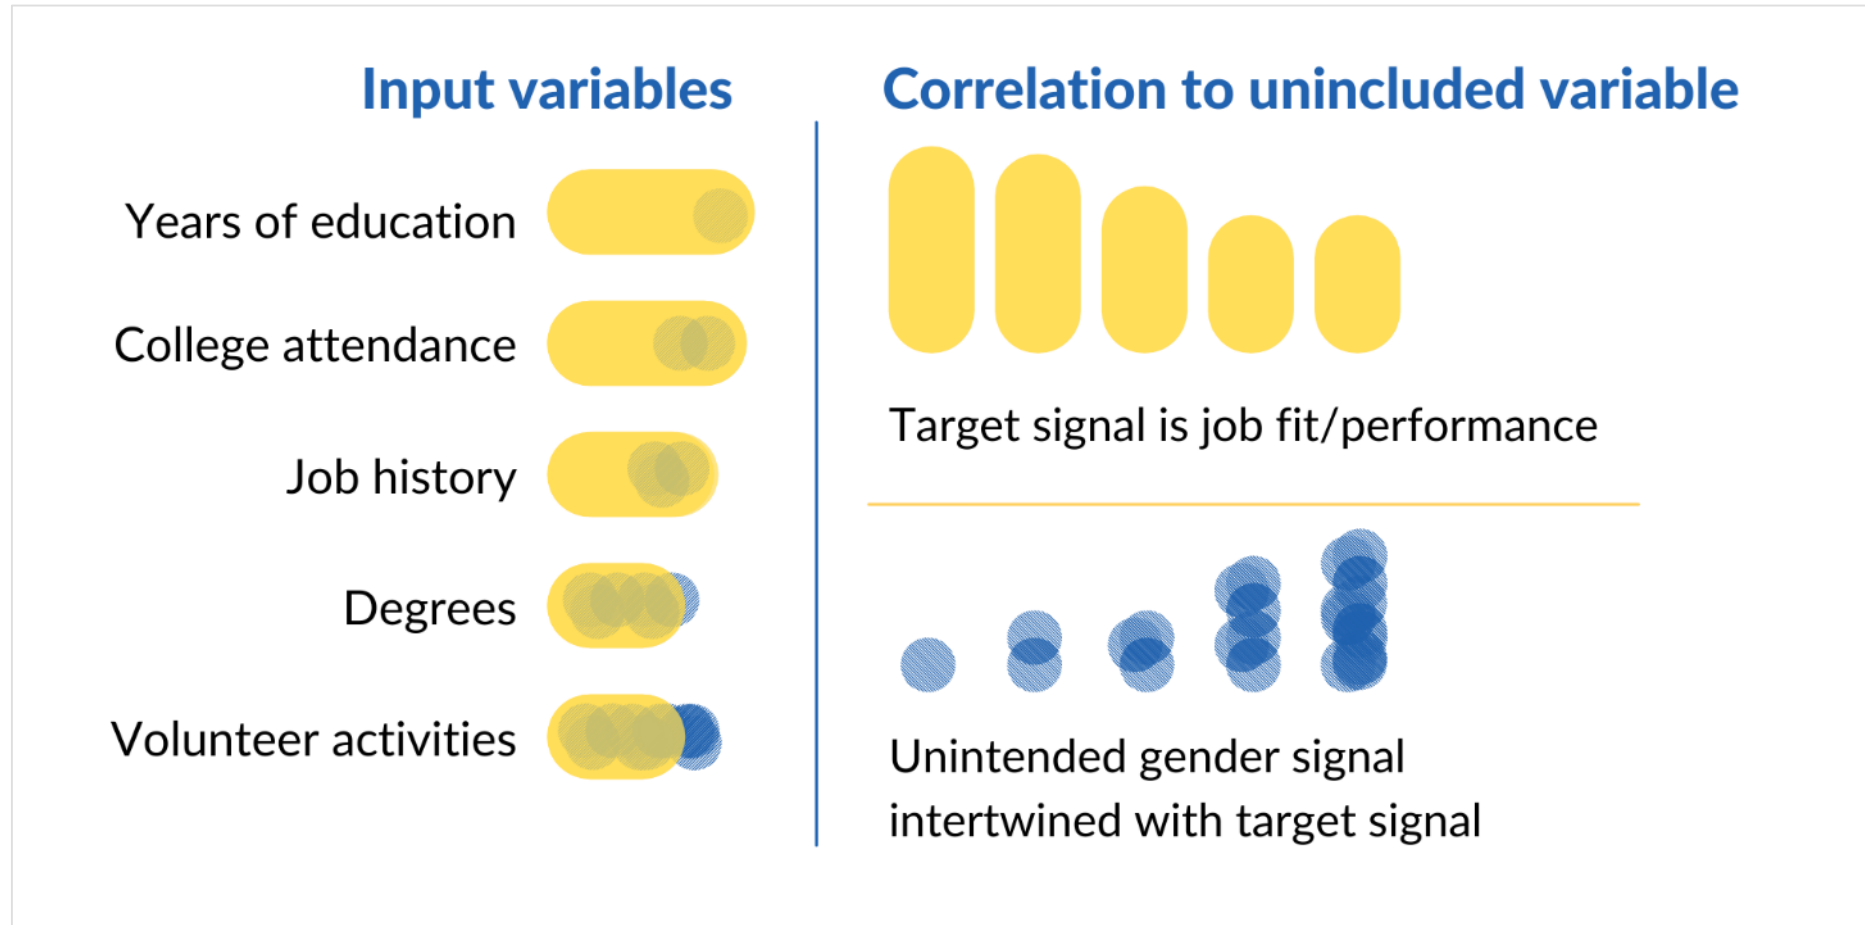
- Models try to determine **the answer to a question** such as what’s your credit risk? Will you pay me back?
- The more accurate the answer the better.
- **Algorithmic bias** can be caused by:
 - The data on which the model is built (*pre-processing*)
 - How the model works internally (*in-processing*)
 - How the output is used (*post-processing*)



Biased Collection



Proxy effect



Risks of Classification or Recommendation Systems

The primary risks are **accuracy** and **bias**, and controls should be built around both of those risks

How Generative AI Works

The background of the slide is a dark blue color. On the right side, there are several thin, light blue wavy lines that flow from the bottom towards the top, creating a sense of movement and depth. The lines are smooth and curved, resembling waves or a stylized representation of data flow.

Large Language Models



A MODEL of LANGUAGE built on a LARGE volume of expressions of that language



Trained on trillions of examples of how human beings have expressed an idea in text.



Trained on what has been digitized in the last 60 years. (News articles, books, scholarly works, social media, millions of web pages.)



Some of these expressions are factually accurate, fair, benevolent, and useful. Others are certainly not.




The result is that LLMs have created **the world's largest compendium of knowledge** (and also a lot of junk).



The trick is how to get what you need out of it without the junk.

GenAI guesses the next best word

The best thing about AI is its ability to



Steven Wolfram, "What Is ChatGPT Doingand Why Does It Work?"

BUT It doesn't understand the meaning of those words


This is why it
“hallucinates”



Write a positive review of fyre festival



I attended the Fyre Festival last year and was blown away by the experience. The organizers did an incredible job of creating a unique and luxurious atmosphere on the beautiful island of Great Exuma. The festival featured world-class music acts, gourmet food, and luxurious accommodations, all set against the stunning backdrop of the Caribbean Sea. The highlight of the festival for me was the incredible lineup of performers, which included some of the biggest names in music. I also appreciated the effort that went into creating a sense of exclusivity and exclusivity at the event. Overall, I had an amazing time at the Fyre Festival and would highly recommend it to anyone looking for a one-of-a-kind event experience.



Just as with text,
pictures can lie



Risks of Generative AI Systems

The primary risks are **accuracy, IP infringement, and bias/toxicity**, and controls should be built around those risks.

ChatGPT, Anthropic, BERT, LLaMa2

Also image generators like DALLE-2, Midjourney, Stability etc.



The first widespread commercial applications of this technology have been CHAT platforms.



This **compendium of knowledge** is available through an interface where you can ask it questions.



This is immensely powerful but requires a great deal of expertise to get the right answers out.



Think of prompt engineering as deposing the platform.

AI Ethics

The background of the slide is a dark blue color. It features several thin, light blue wavy lines that flow from the bottom left towards the top right, creating a sense of movement and depth. The lines vary in frequency and amplitude, some being more pronounced than others.

General Coalescing Around Ethical Principles

NIST AI Risk Management Framework (Jan. 2023), e.g.

- Valid & Reliable
 - Safe
 - Secure & Resilient
 - Explainable & Interpretable
 - Privacy-enhanced
 - Fair – With harmful bias managed
- Accountable & Transparent

AI Risk Management Framework



EU AI Act

AI Safety Principles

- Human oversight and agency
- Technical robustness and safety
- Privacy and data governance
- Transparency
- Diversity, non-discrimination, and fairness
- Social and environmental wellbeing

The Challenge:

How do you take these general principles and translate them into mathematical tests?

- What am I testing for?
- What methods am I using to test?
 - What are their strengths and weaknesses?
- When I get a result, is that a good or bad result?
- How do I "fix" a bad result without causing more problems?

Regulatory Trends

CHALLENGES AND LEGISLATIVE
DEVELOPMENTS

AI in Washington

Current state of affairs

Actions in Congress

- Majority Leader Schumer led AI legislative initiative *Safe Innovation Framework for AI policy*, June 2023
 - “AI is unlike anything Congress has dealt with before” Maj. Leader Schumer
 - Member-only briefings with key AI thought leaders
- Senate AI hearings – The debate on whether to regulate or not regulate, e.g. Senate Judiciary hearing, May 2023
- House Speaker McCarthy AI education effort, April 2023
- **AI Research, Innovation and Accountability Act – November 13, 2023**
 - <https://www.dlapiper.com/en/insights/publications/ai-outlook/2023/us-senators-introduce-bill-to-establish-ai-governance-framework>



U.S. Executive Order on AI

- President Biden signed into effect an Executive Order on [Safe, Secure, and Trustworthy Artificial Intelligence](#) on October 30, 2023.
- The Order requires the development of standards, practices, and even new regulation for the development and use of AI across most aspects of the U.S. economy.
- The order also stresses the need for AI systems security and testing AI systems in order to prevent harms such as discrimination.
- <https://www.dlapiper.com/en/insights/publications/ai-outlook/2023/safe-secure-and-trustworthy-white-house-publishes-executive-order>



CFBP, DOJ, EEOC, FTC Issue Joint Statement on Regulating AI

<https://www.justice.gov/crt/page/file/1581491/download>



JOINT STATEMENT ON ENFORCEMENT EFFORTS AGAINST DISCRIMINATION AND BIAS IN AUTOMATED SYSTEMS

FTC Enforcement Action –X-Mode Social, Inc.

- January 9, 2024, FTC reached a settlement with X-Mode Social, Inc. and Outlogic, LLC. involving its allegations that X-Mode and Outlogic gathered and sold sensitive location consumer data to third parties.
- The FTC stated that X-Mode and Outlogic disregarded consumer privacy preferences in violation of Section 5(a) of the FTC Act.
- The agreement included **injunctive relief and disgorgement** of data and any resulting model or algorithm derived from such data.

FTC Enforcement - Rite Aid

- December 19, 2023 FTC took action against Rite Aid Corporation for their use of **facial recognition technology** in their store surveillance program.
- In the action, the agency contends that the system was inherently flawed, generating **unacceptable levels of error** and leading to individuals being wrongly searched and accused of criminal conduct in contravention of Section 5 of the FTC Act.
- FTC ordered the **deletion and disgorgement** of all data, models and any other data product based on them.

The FTC's Recipe for Compliance

- Conducting thorough **risk assessments** to evaluate the potential for consumer misidentification, particularly concerning racial or gender-based disparities;
- Rigorously verifying the technology's precision through **comprehensive testing** and documentation processes before its operational use;
- Implementing robust procedures to ascertain and maintain **high quality training data** that informs and trains the algorithmic decision-making framework; and
- Establishing ongoing oversight protocols to **continuously monitor** the technology's performance, with a focus on detecting and addressing any emergent bias or inaccuracies.

AI in Utah

Utah Social Media Regulation Act

- Requires that social media companies verify users' ages, obtain parental consent authorizing minors to use the companies' services, limit the functionality of minors' accounts (e.g., messaging, advertising), and restrict the hours at which minors can access their accounts.

S.B. 149 Artificial Intelligence Amendments (Passed, sent to Governor March 12, 2024)

- Establishes Office of AI Policy
- Establishes AI sandbox
- Amends criminal code and consumer protection act to clarify that violations by AI are the responsibility of the humans using it
 - <https://le.utah.gov/~2024/bills/static/SB0149.html>



Litigation Trends



Court decisions will play a significant role in determining the contours of the law.

In re: Social Media Adolescent Addiction/Personal Injury Prods. Liab. Litig. (N.D. Cal.)

Consolidated class action claiming the algorithm that decided to show harmful content to minors is a product and resulted in harm.

Claims:

1. Strict Liability: Design Defect
2. Strict Liability: Failure to Warn
3. Negligent Design Defect
4. Negligent Failure to Warn
5. Negligence *per se*

In re: Social Media Adolescent Addiction/Personal Injury Prods. Liab. Litig. (N.D. Cal.)

Does the Communications Decency Act bar the claims?

Yes and No

- The CDA protects acts that are “publishing.”
- Some design features **are not like publishing** and the CDA does not apply
 - E.g., Not providing parental controls, no voluntary time limits, hard to delete account, hard to report predators.
- Some design features **are like publishing** and are protected by the CDA
 - E.g., use of algorithms to promote addictive behavior (THE BIG ONE), recommending children as friends to adults, notification clustering to drive attraction.

In re: Social Media Adolescent Addiction/Personal Injury Prods. Liab. Litig. (N.D. Cal.)

Are the social media platforms products?

Yes and No

- The question isn't all or nothing. Similar to the analysis under the CDA, the court looked at **specific design features.**
- The court found that all of the design features that are not protected by the CDA are products or product features and thus **the product liability claims as to them can proceed.**

Tremblay v. Open AI (N.D. Cal.)

Consolidated class action by authors claiming their copyrighted works were used as training data for GPT4

Claims:

1. Direct Copyright Infringement (not part of the MTD)
2. Vicarious Copyright Infringement
3. Violation of DMCA
4. Unfair Competition (CA)
5. Negligence
6. Unjust Enrichment

Tremblay v. Open AI

Consolidated class action by authors claiming their copyrighted works were used as training data for GPT4

Claims:

1. Direct Copyright Infringement (not part of the MTD)
2. ~~Vicarious Copyright Infringement~~
3. ~~Violation of DMCA~~
4. Unfair Competition (CA)
5. ~~Negligence~~
6. ~~Unjust Enrichment~~

Use Cases for AI





General LLMs can be sent to law school.

By emphasizing and prioritizing the legal content these models ingest, they are becoming a **compendium of legal knowledge**.

- Cocounsel (now Thompson Reuters,) Lexis Nexis, Harvey, Many more in development



This has made many legal tasks easier, faster, and more accurate.

AI Use Cases

- Legal Research
- Generating timelines
- Creating initial drafts of:
 - memoranda
 - deposition outlines and questions
 - discovery requests
 - motions
 - contracts and agreements
- Summarizing, analyzing, and comparing
- Determining how a draft contract differs from optimal desired position

Making general legal LLMs speak in your voice

1010
1010

Law firms can take an LLM that has been turned into a legal LLM, and turn it into a firm LLM

In fact, any company can do this, and this is the real power of GenAI



By ingesting the law firm's knowledge, as contained in its data, an LLM can be fine-tuned to prioritize that knowledge in its compendium.

Now it not only can draft a memorandum, but do so based on what the firm has said before, and in its voice.



Firms can build "baby bots" for every practice group (and for every product that the practice group creates)

Implications on Legal Practice

A decrease in the cost of legal services will have an immense leveling effect and increase access to justice.

The plaintiffs bar will likely take advantage of AI first as they have less organizational inertia against adoption.

Pro bono and access to justice organizations will also benefit from AI.

AI represents an existential threat to Big Law, unless it reacts deftly.

There will be firms that understand the transformative effect of AI, take advantage of it, and stride ahead of competitors.

Other firms will seek to defend historical methods of success and be overtaken.

Discussion

Thank you

Bennett Borden

Partner and Chief Data Scientist

DLA Piper

Bennett.Borden@us.dlapiper.com

