

2020 ACC National Capital Region

Protecting Artificial Intelligence Innovations: Unique Challenges & Strategies



Introductions



Matt Weinstein

Matt Weinstein is Director, Legal - Patents & Open Source at Accenture, responsible for AI and emerging technology patent portfolio development, IP risk mitigation, dispute resolution, licensing and commercial contracting. Matt also leads Accenture's open source legal team, which is responsible for all open source governance and licensing matters for the company. Matt holds a BS in Computer Engineering.



Hilary Weckstein

Hilary Weckstein is an experienced leader in the healthcare information technology industry. As General Counsel and Chief Privacy Officer at Prognos Health Inc., she holds responsibility for legal, risk management, human resources, and data privacy matters. She negotiates and drafts agreements and advises the company on all matters, including complex legal issues related to data analytics and the use of innovative technologies. Hilary holds certifications in healthcare compliance (CHC®) and in data privacy (CIPP/US®).

Introductions



Ilona Levine

Ilona Levine is General Counsel, Legal Affairs and Corporate Communications at China Telecom Americas. She frequently speaks on privacy, cybersecurity, litigation and other topics at various professional conferences. currently serves as co-chair of the ACC North Capital Region Privacy and Data Protection Forum.



Rodney Rothwell

Rodney Rothwell is a partner on Kilpatrick Townsend's patent team. Rodney's practice is focused on machine learning techniques, medical devices and bioinformatics. He is the Vice-Chair of the Corporate Counsel Committee of the Federal Circuit Bar Association. Rodney holds a BS in Medical Technology and in Electrical Engineering.

Introductions



Kate Gaudry

Kate Gaudry is partner on Kilpatrick Townsend's patent team. Kate's practice is focused on artificial intelligence, computational biology, and software technologies. She has authored over 60 publications focused on original empirical analyses. Kate holds a PhD in Computational Neurobiology and a BS in Physics.

Outline

- Artificial Intelligence Primer
- Artificial Intelligence and Intellectual Property
- Artificial Intelligence and Data Privacy/Cybersecurity

Artificial Intelligence Primer

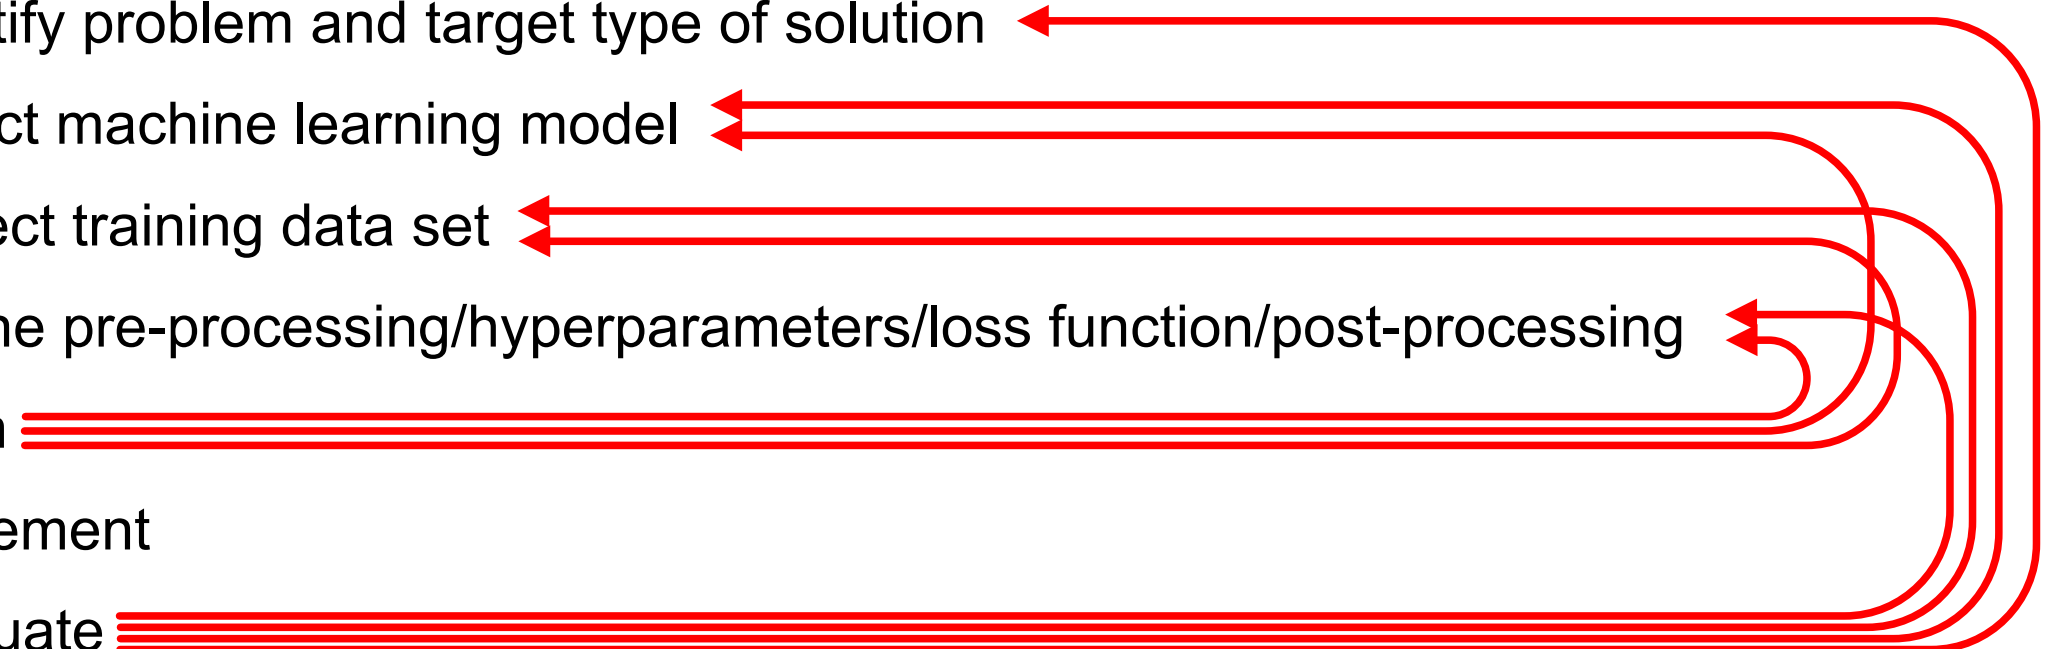
- Computer systems performing tasks that traditionally required a human decision-maker.
- Exemplary use cases: computer vision, speech recognition, decision-making, language processing, medical diagnoses, financial trading, art generation.
- Artificial intelligence as a field of research has existed since the 1950s.
- Recent advances in computing system power, big data, and the shift to build focused “intelligent agents” has precipitated widespread use.

Artificial Intelligence Primer

- Artificial intelligence includes machine learning and rules-based algorithms.
- **Machine learning** techniques (e.g., neural networks) infer parameters from an existing data set (training data) and use those parameters to process subsequent information.
 - Sub-optimal training data will result in a sub-optimal model.
- **Rules-based** algorithms include deterministic systems that do not need to be trained.
 - Sub-optimal rules will result in a sub-optimal model.

Artificial Intelligence Primer

- How are machine learning systems developed?

- Identify problem and target type of solution
 - Select machine learning model
 - Collect training data set
 - Define pre-processing/hyperparameters/loss function/post-processing
 - Train
 - Implement
 - Evaluate
- 

Artificial Intelligence Primer

- Narrow intelligence: AI that generates outputs for a specific task.
- General intelligence: AI that applies learned knowledge and skills in a variety of context to mimic human intelligence.
- All existing AI is narrow intelligence...
- When will we see General AI? By 2030? 2200? Never?

Artificial Intelligence and Intellectual Property Law

- Potential types of “AI inventions”:
 - New types of AI
 - New applications of AI
 - Innovations made with the assistance of AI
 - Innovations made by AI

Patenting Requirements


- Patent Eligibility: an invention must be more than just an abstract idea, natural phenomenon, law of nature or product of nature.
- Prior-Art Requirements: Novelty / Non-Obviousness.
- Disclosure Requirements: Written Description conveying conception of invention and Enablement to make and use the invention (and *Best Mode* during examination).

Patent Eligibility (in the U.S.)

- Under current USPTO guidance, a claim is patent eligible if:
 - It is not directed to a judicial exception (e.g., abstract idea);
 - Any judicial exception to which the claim is directed is integrate into a practical application; or
 - The claim is significantly more than the judicial exception.

AI and Patent Eligibility

- Patent eligibility of inventions of new types of AI and/or new software integrations of AI will likely be heavily dependent on the evolving state of patent eligibility in the software space.
- Patent eligibility of inventions developed with or by AI will likely depend on the particular field of the invention.



Gaudry, K, Hayim, S. “Update on 101 Rejections at the USPTO: Prospects for Computer-Related Applications Continue to Improve Post-Guidance” *IPWatchDog*, 2019.

AI and Novelty / Non-Obviousness Requirements

- Novelty and/or non-obviousness of inventions of new software integrations of AI will likely be more difficult to exemplify as AI continues to be more heavily integrated in society.
- Non-obviousness may become even harder to substantiate as AI “building blocks” proliferate.
- Novelty and/or non-obviousness of inventions developed with or by AI will likely depend on the particular field of the invention.

AI and Disclosure Requirements

- If field is predictable, less disclosure is required. Whether an “AI innovation” is in a “predictable” space will depend on the innovation.
- Some AI models are more interpretable than others. The feasibility of disclosing how a model operates or why a particular solution is suitable may depend on the model used to arrive at the solution and how it was developed.
- Data has significant independent value. The value of training data may influence the degree to which an entity decides to pursue patent protection in a particular case. Patent offices seem to recognize this sentiment and an interim disclosure of how training data is to be collected may be sufficient.

AI and Inventorship

- Most jurisdictions only recognize humans as being capable of being inventors.
- DABUS innovations – patent applications filed only listing an AI as inventor.
 - USPTO rejected due to formality issues.
 - EPO rejected due to formality issues.
- WIPO
 - Issued request for comments on IP Policy for AI Inventions in Dec. 2019, raising the issue of inventorship for AI.
- **If AI cannot be named as an inventor, who is the inventor? Can the invention be patented at all?**

AI and Copyright / IP Ownership Issues

- Should the output of an AI (without any creative involvement from a human) qualify for copyright protection? If not, then what degree of human involvement would be sufficient to qualify the work for copyright protection?
- Does a new type of right need to be created to protect AI-created content?
- Who owns the resulting work? The entity that creates the AI? The entity that trains the AI? The entity that uses the AI?
- Is it okay to use copyright-protected material to train an AI? Should the copyright-protected content used to train AI be recognized in some way? Could the training of the AI be considered a transformative work?

AI and Open Source

- Business leaders are excited about AI's potential to profoundly transform organizations by making them more innovative and productive.
- AI algorithms have been around for a while and many of the most popular and prevalent algorithms are open source software ("OSS").
- When organizations build AI systems, they typically use available open-source AI algorithms, use commercial "black box" AI systems that leverage open-source components, or build from scratch with open-source components and proprietary components.
- If your company is developing products or services with AI, chances are high that your company uses in some capacity. While the benefits of OSS are clear (e.g., flexibility, cost-effectiveness, and speed), potential risks must be accounted for and mitigated.

What is Open Source?

- OSS is a type of computer software for which source code is generally made freely available ...under a specific license permitting others to use, copy, modify and redistribute the software, as well as other terms and conditions that must be complied with.

- Examples of Popular Open Source AI:

- TensorFlow
- OpenNN
- Scikit-learn
- IBM Watson
- Apache Mahout
- Accord.NET
- Torch
- Keras
- Microsoft Cognitive Toolkit
- Theano
- Caffe

Open Source AI

- Using open source AI requires proper governance and controls, as with all use of OSS, to manage and mitigate legal, operational, and security risks.
- Know your obligations under the applicable OSS licenses and comply with them.
- Keep license information and copyright notices intact, provide attribution notices, identify modifications, refrain from using the developer's name to promote your software, etc.
- **Watch out for changes in use cases!** OSS that may be OK to use in an internal use case may not be OK to distribute externally/incorporate in a commercial software product.
- OSS may also be available under a commercial license (dual licensing).

Open Source AI – Legal Risks

- Certain OSS licenses are “viral” and may “spread” to your proprietary code if you combine your proprietary code with the OSS and distribute it.
 - Some licenses are only viral in the event of certain kinds of combination (static vs. dynamic linking).
 - Some licenses are viral even when there is no “distribution” of code, and code is merely accessed over a network (SaaS).
- Many open source licenses include express patent license grants and some arguably trigger an implied patent license grants.
- Some open source licenses require upstream indemnification if the OSS is used in commercial software.
- Unlike commercial software, OSS comes with no warranties and no support by default (some organizations may offer paid support).

Open Source AI – Operational Risks

- Using out-of-date, unpatched or unmaintained Open Source AI could impact your release capabilities, product quality and customer trust.
- Understand whether the open source community is consistently managing the open source code, and that your organization has implemented the latest version and deployed patches have been applied.
- Ensure that open source components won't risk the quality of the AI system. As of now, there are no universal standards for evaluating an open source component's quality.

Open Source AI – Security Risks

- Using Open Source AI brings the possibility of introducing software security vulnerabilities into your source code.
- Vulnerabilities in open source AI are similar to exploits that appear in proprietary software products. These are bits of code that the code author accidentally (or intentionally), wrote that hackers can benefit from, or features that permit attackers to capitalize in a way that was not planned (or planned) by the author of the code.
- Without the guarantees or warranties that typically are available for commercial software, it becomes your responsibility to understand and mitigate these risks, particularly where your end users expect you to provide warranties and indemnities.

Open Source Governance and Controls

- Knowing, approving, and managing the open source code your company uses, modifies, contributes, and/or distributes is key.
- To do this, you will need an open source policy.
- Common elements of open source policies:
 - Identify and educate stakeholders
 - Identify open source code business/legal objectives
 - Approval process for use of open source code
 - Identification and tracking of all open source code
 - Compliance obligations
 - IP considerations
 - Contributions/releases under OSS licenses

AI and Data Privacy

- What is data privacy law?
 - In the United States, we do not have a comprehensive privacy law. Laws governing privacy are patchwork, generally applying industry by industry (just to name a few: HIPAA & GINA for certain healthcare data, COPPA for children's data online, FERPA for educational records, GLB & FCRA for financial information).
 - Many states have also enacted more comprehensive privacy laws.
 - Outside of the US, the EU (GDPR), Canada (PIPEDA), and Australia all have comprehensive privacy laws.

Laws generally address at a minimum: Notice, Choice, Access.

A Sampling of Data Privacy Concerns in AI

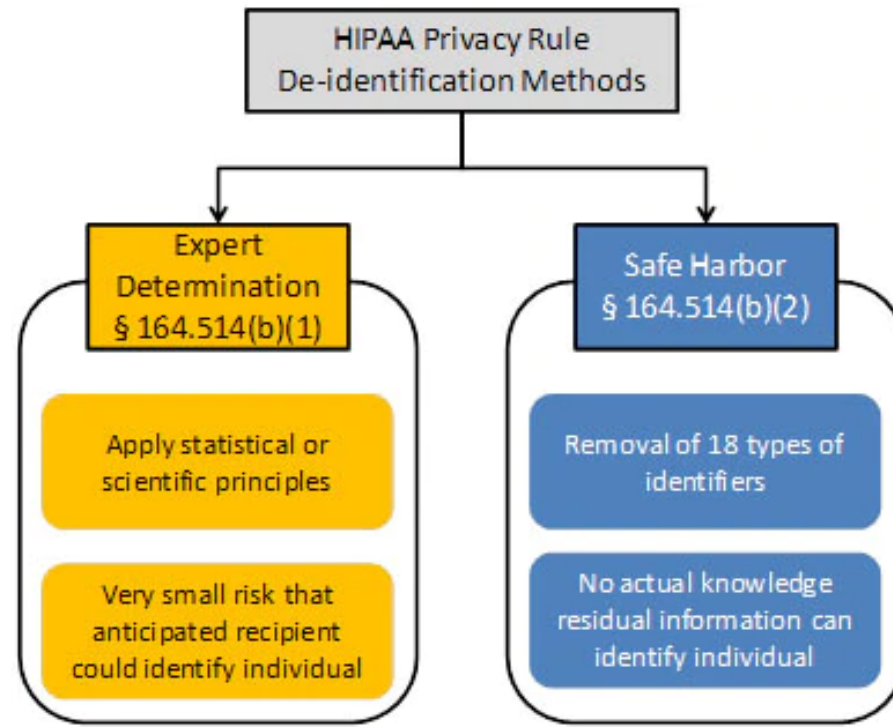
- Deceptive Practices (Ashley Madison: the FTC alleged that the AI used fake “engager profiles” of attractive mates to induce potential customers to sign up for the dating service.)
- Lack of Transparency (Facebook: the FTC alleged that Facebook misled consumers when it told them they could opt into facial recognition – even though the setting was on by default.)
- Responsibility to Understand Data Inputs (Under the FCRA you may be required to provide an “adverse action” notice explaining what information was used to come to a certain decision.)
- Potential for unfair or discriminatory outcomes (Example: using zip codes in an algorithm that helps a company make credit decisions.)

HIPAA

- HIPAA is a federal law intended both to reduce healthcare costs by requiring the use of electronic data interchange and to protect the security and privacy of protected health information (PHI) in electronic data interchange.
- There are four basic rules of HIPAA that relate to Prognos: **the Privacy Rule**, the Security Rule, the Breach Notification Rule, and the Enforcement Rule.
- Under the Privacy Rule, HIPAA requires patient consent for any use or disclosure of PHI, excluding the prescribed “permitted uses” for **treatment, payment or healthcare operations**.
 - **Treatment** is the provision of healthcare and related services.
 - **Payment** encompasses the activities of a health plan to obtain payments and reimbursements for the provision of healthcare and to obtain premiums, fulfill coverage responsibilities or reimburse patients and providers for healthcare services.
 - **Healthcare Operations** is defined as all actions necessary to run a business and to support the core functions of treatment and payment. Examples include quality assessment, fraud and abuse detection, and compliance programs.

The Privacy Rule: De-Identified Information

The Privacy Rule provides the standard for de-identification of PHI. Health information is no longer “individually identifiable” if it does not identify an individual and there is no reasonable basis to believe it can be used to identify an individual. De-identified information is no longer individually identifiable health information – no longer PHI covered by HIPAA’s Privacy Rule.



De-Identification: Expert Determination Method

(b) Implementation specifications: requirements for de-identification of protected health information. A covered entity may determine that health information is not individually identifiable health information only if:

(1) A person with appropriate knowledge of and experience with generally accepted statistical and scientific principles and methods for rendering information not individually identifiable:

(i) Applying such principles and methods, determines that the risk is very small that the information could be used, alone or in combination with other reasonably available information, by an anticipated recipient to identify an individual who is a subject of the information; and

(ii) Documents the methods and results of the analysis that justify such determination.

De-Identification: Safe Harbor Method

In order to meet the Safe Harbor under §164.514(b), the following identifiers of the individual or of relatives, employers, or household members of the individual, must be removed:

- Names
- All geographic subdivisions smaller than a state, including street address, city, county, precinct, ZIP code, and their (except year) for dates that are directly related to an individual, including birth date, admission date, discharge date, death date, and all ages over 89 and all elements of dates (including year) equivalent geocodes*
- All elements of dates indicative of such age, except that such ages and elements may be aggregated into a single category of age 90 or older
- Telephone numbers
- Fax numbers
- Email addresses
- Social security numbers
- Medical record numbers
- Health plan beneficiary numbers
- Account numbers
- Certificate/license numbers
- Vehicle identifiers and serial numbers, including license plate numbers
- Device identifiers and serial numbers
- Web Universal Resource Locators (URLs)
- Internet Protocol (IP) addresses
- Biometric identifiers, including finger and voice prints
- Full-face photographs and any comparable images
- Any other unique identifying number, characteristic, or code

Privacy By Design (PbD)

1. Proactive not Reactive; Preventative not Remedial
2. Privacy as the Default Setting
3. Privacy Embedded into Design
4. Full Functionality – Positive-Sum, not Zero-Sum (“win-win”)
5. End-to-End Security – Full Lifecycle Protection
6. Visibility and Transparency – Keep it Open
7. Respect for User Privacy – Keep it User-Centric

AI and Cybersecurity

- According to projections, AI and Machine Learning Solutions spending will exceed \$133 billion by 2022.
- Many businesses are already using AI tools in its cybersecurity programs, including breach detection.
- But these technology could be exploited by malicious hackers.

AI as a Cybersecurity Tool

- 75 percent of cybersecurity executives say that AI allows them to respond to breaches faster.
- Speed is where AI excels the most by surpassing the human capacity to detect and mitigate threats.
- 59 percent of cybersecurity professionals say AI streamlines the process of detecting and responding critical system weaknesses.
- **How?** Artificial intelligence shows significant potential for detecting fraudulent activity, malware and intrusions, as well as gauging the risk levels of login attempts. By making threat detection more sensitive and enabling nuanced behavior tracking, AI increases flexibility within identity and access management strategies.

AI as a Cybersecurity Threat

- A report by Europol has warned that Artificial Intelligence is one of the emerging technologies that could make cyberattacks more dangerous and more difficult to spot than ever before.
- Cyber criminals have already started using these techniques to help conduct hacking campaigns and malware attacks:
 - self-learning automated malware
 - ransomware
 - social engineering
 - phishing attacks.
- Example, machine learning could be employed to send out phishing emails automatically and learn what sort of language works in the campaigns, what generates clicks and how attacks against different targets should be crafted.

Other Examples

- “Deep fakes”—the use of AI to replicate and manipulate a person’s voice and image—as a cyber-crime weapon.
 - CEO’s voice was falsified using AI to execute a fraudulent money transfer.
- Developing malware that allows hackers to manipulate medical scans and produce fake cancer images, resulting in erroneous diagnoses.

AI and Cybersecurity Vulnerabilities

- AI technologies are increasingly integrated with products and services.
- The concern is AI-related cybersecurity vulnerabilities.
 - A recent Deloitte global study of AI early adopters revealed that more than four in 10 executives have “major” or “extreme” concerns about various types of AI risks, with “cybersecurity vulnerabilities” topping that list.
- The reasons for these concerns vary among countries:
 - U.S. executives worry hackers will use AI to steal sensitive or proprietary information.
 - Canadian executives worry about potential manipulation of AI data or algorithms.
 - France and Germany executives are most concerned about adversaries using AI to automate tasks involved in executing cyber-attacks. French executives also expressed more concern than their counterparts from other countries about misuse of AI to impersonate authorized users and bypass cyber defenses.

Comparison with Traditional Cybersecurity Issues

- AI attacks are fundamentally different in nature than the cybersecurity attacks.
- Traditional cybersecurity vulnerabilities are generally a result of programmer or user error. As a result, these errors can be identified and rectified.
- In contrast, the AI attack problem is more intrinsic: the algorithms themselves and their reliance on data are the problem. Unlike traditional cybersecurity vulnerabilities, the problems that create AI attacks cannot be “fixed” or “patched”.

Comparison with Traditional Cybersecurity Issues

- This difference has significant ramifications for policy and prevention.
- Mitigating traditional cybersecurity vulnerabilities deals with fixing “bugs” or educating users.
- Solution:
 - User education
 - IT department-led policy enforcement
 - Technical modifications such as code reviews and bug bounties

AI & Cybersecurity Conclusions

- **For AI attacks, training and robust passwords are not enough.**
- The algorithms themselves have the inherent limitations that allow for attack.
- Some of the proposed solutions:
- **Creating Security Compliance program specifically for AI Attacks**
- The compliance program will contain best practices that will manage the entire lifecycle of AI systems:
 - In the planning stage, they will force stakeholders to consider attack risks and surfaces when planning and deploying AI systems.
 - In the implementation stage, they will encourage adoption of IT-reforms that will make attacks more difficult to execute.
 - In the mitigation stage for addressing attacks that will inevitably occur, they will require the deployment of previously created attack response plans.
 - This program is modeled on existing compliance programs in other industries, such as PCI compliance for securing payment transactions.

Locations

Counsel to innovative companies and brands around the world

We help leaders create, expand, and protect the value of their companies and most prized assets by bringing an equal balance of business acumen, technical skill, and creative thinking to the opportunities and challenges they face.



Anchorage
Atlanta
Augusta
Beijing
Charlotte
Dallas
Denver

Houston
Los Angeles
New York
Raleigh
San Diego
San Francisco
Seattle

Shanghai
Silicon Valley
Stockholm
Tokyo
Walnut Creek
Washington DC
Winston-Salem