



ACEDS | NY Metro
ASSOCIATION OF CERTIFIED
E-DISCOVERY SPECIALISTS Chapter

magazine

Issue 1 Autumn 2019



AI & Ethics: The Roundtable and the Rule of Law

by Cat Casey & James A. Sherer

Why Meet and Confer on AI?

Rapid technological innovation—and in particular, the adoption and use of artificial intelligence (AI and specifically machine learning or ML), has fundamentally changed both daily life and business practices across a wide spectrum of industries. While the practice of law was arguably later to the party than most (although part of that may be legal practitioners’ reluctance to trust the validity of advertising and marketing campaigns) there is no denying that AI is now firmly entrenched in legal practice considerations.

This includes AI-optimized practice management for law firm operations; continuous active learning (CAL) powered by ML or neural networks that reduce time to insight and evidence in eDiscovery; and advanced AI suggesting or predicting the merits of appeals and case outcomes. But it also includes the uses of AI in client activities and strategies, where attorneys, paralegals, and support staff must have the requisite understanding of just what AI means to their clients—or understand enough to know what they don’t know, and then find a way to learn or to get assistance in supporting the practice of law on those clients’ behalf.

These were the issues raised during a set of discussions and teaching sessions this September 20th and 21st, when 80 thought leaders in the AI space (including members of parliament, heads of state, legal scholars, ethicists, and legal technologists developing legal solutions powered by AI) gathered in Athens, Greece at the “Athens Roundtable on AI and the Rule of Law.” There, this diverse group analyzed and debated how to develop an ethical

framework that would direct the use of AI in support of legal endeavors; how to educate practitioners and users; and how to further, intelligently consider how an understanding of AI is key to the future of the rule of law and related practice.

Hosted by the Institute of Electrical and Electronics Engineers (IEEE), global law firm [Covington & Burling LLP](#), [The Future Society](#), and the [European Law Observatory on New Technologies \(ELON\)](#), the roundtable aimed to move from a diffuse discussion of nascent and existing principles occurring around the globe, and to focus instead on practice and practical implications that could form a middle ground for consensus and understanding. The roundtable further sought to bring together a variety of different stakeholders focused on the ethics of AI development and application to law, and admirably succeeded in creating a diverse group of invested participants.

The Rule of Law

This conference focused on the application of AI in support of (or at least preventing AI from impeding) the rule of law, a diffuse concept that is summed up by the World Justice Project as four universal principles:

- **Accountability** – Both the government and private actors are accountable under the law, and no one is above it
- **Just Laws** – Laws are clear, publicized, stable, applied evenly, and protect fundamental human rights (such as safety, contracts, property)
- **Open Government** – The enactment and enforcement of laws are accessible, fair, and efficient

- **Accessible & Impartial Dispute Resolution** – Justice is delivered timely by competent, ethical, and independent representatives who are accessible, possess adequate resources, and reflect the communities they serve

So, the rule of law as discussed at the roundtable focused on what it takes for justice to be equal, and an overriding concern was whether AI and its application could—or would by its very nature—negatively affect that aspiration.

Current AI Ethical Standards

The roundtable first discussed recently adopted ethical standards in AI, and how those standards might apply to other realms, including the law. Specifically, the Roundtable presented the EU’s “[Ethics Guidelines for Trustworthy AI](#),” the [OECD Principles on Artificial Intelligence](#), and the G20 human-centered AI Principles drawn from the OECD to foster debate and consensus on the key components of ethical or trustworthy AI. But while these proposed standards varied, they shared many common themes and concerns, such as ideas on transparency and accountability; disparate impact and non-discrimination; and human-centered agency and oversight ideals.

In particular, the proposed EU ethical standard presented seven key requirements that served as overall themes for much of the continued discussion: human agency and oversight; technical robustness and safety; privacy and data governance; transparency; diversity and fairness; societal and governmental wellbeing; and accountability. The G20 and OECD recommendations likewise focused on

human centricity (the idea that humanity should be at the center of AI approach); inclusivity; transparency; robustness; and accountability.

Again, as part of the level-setting process, a good portion of the first day was spent unpacking these objectives and reviewing the standards from an insider's point of view, where participants in the drafting teams or supporting agencies shared unattributed stories and detail regarding the principles' drafting processes. Some principles were more self-explanatory than others—such as human agency and centricity; diversity and fairness, social wellbeing; and data privacy. But a number of the remainder deserved a more thorough discussion and explanation from the participants.

Discussions involving standards related by the IEEE, the Council of Europe, the OECD, and the EU seemed to address the considerable, potential risks and benefits of AI for the rule of law, and conveyed a sense of urgency in developing ethical and functional standards and certifications. These standards centered on the robustness and effectiveness of the technology; transparency and accountability for the technology's operation; and the added factor of practitioner competence when deploying an AI-powered solution. Each of those deserve a closer look:

Robustness

In the context of trustworthy AI, robustness refers to resiliency and security. AI systems must be safe, with an available fallback plan in case something goes wrong; and should be accurate, reliable, and reproducible according to standards proposed by the EU. This should be familiar to eDiscovery practitioners, as discussions regarding IEEE guidance considered NIST standards of “precision” and “recall” to determine the effectiveness or robustness of a given AI solution, and NIST's own mandate to develop AI standards generally.

Transparency

For transparency, the EU considers systems and system decisions that are

explicable in a manner that makes sense to the stakeholder concerned. For example, humans should be aware when they are interacting with an AI system; should be informed of the system's capabilities and limitations; and should be given enough information to appeal decisions rendered by the system. The IEEE in particular evaluated this component, and acknowledged the inherent potential conflict between the need to safeguard intellectual property rights for purveyors of AI weighed against the need to enable AI consumers to make informed decisions, and sought to determine what insights into the underlying algorithms and data sets would meet that balance.

Competence

As AI-enabled solutions already impact access to justice, provide key evidence, and even forecast case outcomes, the IEEE recognized that the concept of competence would be a key tenet of AI trustworthiness—that is, the tools are too powerful to be used irresponsibly or unknowingly. As AI tools can (and should) effectively augment and amplify legal practitioner intelligence, this

requires a level of operator proficiency and understanding of the AI-powered technology. The thought was that the best AI-powered tool is only ethically and effectively deployed when the practitioner operating it has a basic level of competence.

Accountability

People trust computers, sometimes even to their detriment. With the potentially disproportionate weight people ascribe to the efficacy of computer/AI-generated suggestions and proscriptions, nearly all of the entities participating in the roundtable highlighted the need to hold purveyors and operators of AI-powered tools in the legal space accountable for that vast source of power and influence. The EU guidance in particular required mechanisms in place to ensure responsibility and accountability for AI systems and their outcomes. This included audit-ability, which enables the assessment of algorithms, data, and design processes, especially when critical applications were considered. If the possibility of an issue arose, the EU required an accessible means of fixing the issue.

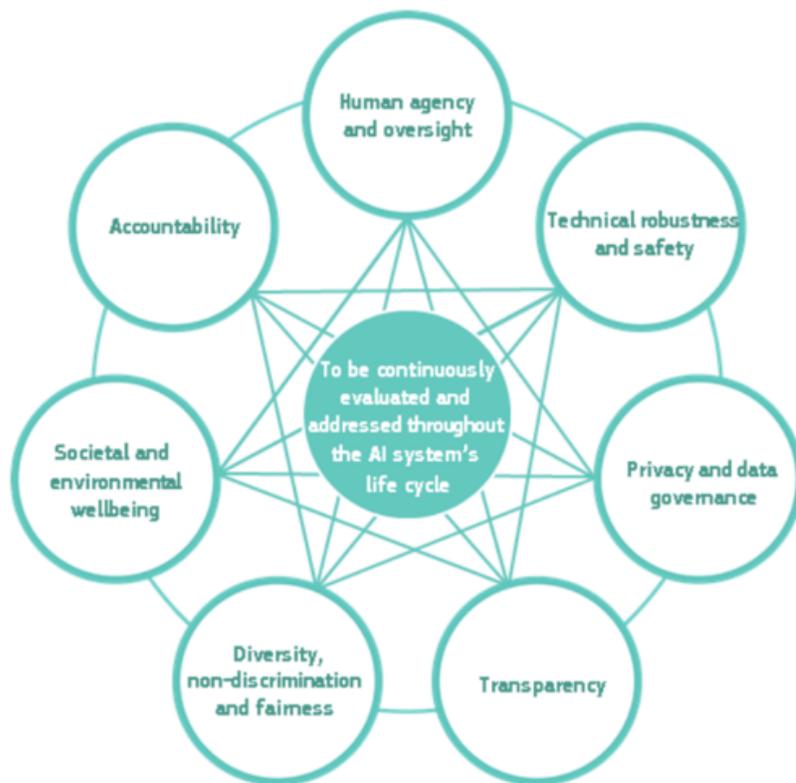


Figure 2: Interrelationship of the seven requirements: all are of equal importance, support each other, and should be implemented and evaluated throughout the AI system's lifecycle

Application of Ethical Frameworks for AI to Legal Applications

The roundtable also addressed the question: “Should practitioners be intimidated or inspired by AI?” After two days of lively debate and deep synthesis, the clear (if not universal) trend among participants was towards an inspired future, with the caveat that practitioners had much work ahead to ensure that outcome.

Nicolas Economou, Chair of The Future Society’s Law Initiative and of the IEEE Law Committee, summed the challenges facing legal in ethical adoption to AI:

The trustworthy adoption of AI in legal systems can support the functions of the law and the values that animate it. The likes of the Council of Europe and the IEEE have now published principles in pursuit of that vision. Their application in practice by lawyers, judges, technology specialists, private entities, and institutions of state is a challenge of the “what” and the “how” to develop appropriate protocols, metrics, standards, supporting policies, and to upgrade the education infrastructure accordingly – for judges, lawyers, technical professionals, and many other stakeholders.

Simply stated, there were concerns that the legal system has substantial educational, regulatory, and application gaps that must be addressed to ensure that AI is a net positive. In order to maximize the benefit of applying AI to the practice of law (while containing and mitigating risk), education across a myriad of stakeholders is key. Without basic AI literacy, therefore, regulating and developing ethical AI is almost impossible.

Legal Education For –Not By–AI

It seemed that the creators of AI technology might require a baseline of legal and ethical education to shepherd their development process—especially when access to justice or case

outcomes was at stake, or if the technology could be misused to the detriment of society. Likewise, it was also likely that legal practitioners, regulators, and citizens generally would require some kind of AI fluency to understand the potential risks and benefits of AI, especially in the context of law.

In order to ethically develop AI and create a marketplace where informed trust and consent could be developed, some roundtable members supported the introduction of technical and AI literacy in primary schools, subsequently refined in university and law schools, and fully realized within the judiciary and the body of practicing attorneys worldwide. And at the most basic level, the roundtable’s consensus was the approach taken to teach citizens and legal practitioners must adapt to both changes in AI as well as the surrounding digitally powered world.

And Do it Now

While participants were actively concerned with regulating too early and perhaps stifling the development of AI, they certainly acknowledged that a wait-and-see approach left too much in the hands of self-interested technologists, corporations, and undereducated practitioners.

This also led to an issue that cut across the debate: collaboration among the stakeholders. Legal practitioners, ethicists, and the bench agreed that they needed to (better) understand what AI is, in order to communicate with AI developers and providers. Involvement of all parties in bridging the knowledge gap and developing ethical frameworks for the future of AI in the law was therefore mission-critical.

The Future

Once the 48 hours of rigorous discussion and debate ended and after the leading minds from the over 20 countries participating in the roundtable stepped away, several additional thoughts remained. The impact of AI is already being felt in law, and it is likely that practitioners require more meetings like the roundtable to address the impact of

AI and how it can be beneficial to the practice of law while practitioners can help mitigate the risks of bias, malfeasance, and incompetence.

Before leaving, the roundtable participants committed to investing time, mental capital, and resources into the development of a set of legal-focused AI Ethics for Law principles, and to develop a curriculum focused on upskilling legal practitioners, the judiciary, and the next generations of law students. The aim is to educate and advise legal practitioners and supporting third parties who are engaging with AI, providing the tools that can distinguish sound and appropriate innovation from marketing slicks, and can apply the same rigors to legal technology that we currently see in medicine and other high-human-impact industries. 

About the Authors

Cat Casey is Chief Innovation Officer (CINO) for the leading AI powered cloud enabled legal technology company, DISCO, where she drives innovation and provides direction for legal technology solutions. Casey is a global thought leader on the application of AI and advanced technology to the practice of law and an outspoken advocate of legal professionals embracing technology to deliver better legal outcomes. Prior to DISCO, Casey was the global director of eDiscovery and Technology for Gibson Dunn a leader in the Forensic Technology Practice for both PwC and KPMG.

James A. Sherer is a Partner in the New York office of BakerHostetler, where he chairs the Information Governance and Artificial Intelligence practice teams and serves as part of the eDiscovery and Management and Privacy and Data Protection groups. James’s work and scholarship focuses on litigation; discovery management processes; enterprise risk; records and information governance; data privacy, security, and bank secrecy; technology integration issues; artificial intelligence; social media and the Internet of Things, and related merger and acquisition diligence.