

Artificial Intelligence

Generative AI In Legal Practice: Technical And Legal Aspects To Consider

*By
Dr. Ilya Kolochenko*

*ImmuniWeb
Geneva, Switzerland*

**A commentary article
reprinted from the
October 2023 issue of
Mealey's Litigation Report:
Artificial Intelligence**



Commentary

Generative AI In Legal Practice: Technical And Legal Aspects To Consider

By
Dr. Ilia Kolochenko

[Editor's Note: Dr. Ilia Kolochenko is a Swiss cybersecurity expert and jurist, having over 15 years of practice. He is a Chief Architect at global application security company ImmuniWeb headquartered in Geneva, Switzerland. He holds a Master of Science degree in Criminal Justice (Cybersecurity & Cybercrime Investigation) degree from Boston University, an LLM (Master of Laws) in Information Technology from University of Edinburgh Law School, and a PhD in Computer Science from Capitol Technology University where he currently serves as an Adjunct Professor of Cybersecurity Practice & Cyber Law. Any commentary or opinions do not reflect the opinions of ImmuniWeb or LexisNexis®, Mealey Publications™. Copyright © 2023 by Dr. Ilia Kolochenko. Responses are welcome.]

I. Introduction

This year may deserve to be named the year of most important advancement in generative Artificial Intelligence (AI) and its swift proliferation into our everyday lives. The explosive growth of generative AI technologies is not without its technical and legal drawbacks, many of which Gordon Platt and I predicted back in February 2023.¹ While for many legal professionals the technical side of generative AI remains an indecipherable enigma, well-commoditized and ready-to-use AI products by Azure Cognitive Services and AWS AI Services enable even small companies and startups to streamline Machine Learning (ML) technologies for a wide spectrum of business tasks. In parallel, tech companies are jealously fighting for AI training data. Illustratively, X (formerly known as Twitter) has just updated its privacy policy to use public posts of its users for AI training purposes,² after recently restricting third-party access to its APIs – likely to prevent competi-

tors from doing exactly the same with user-generated content. The legality of AI training with personal, copyrighted or otherwise proprietary data remains a topic of heated debates amid the growing legal uncertainty.

Nonetheless, the ongoing media hype around the alleged superpower of generative AI triggers an ambivalent *déjà vu* effect in my memory. Several years ago, when Bitcoin was steadily approaching its historical maximum, some developing countries declared Bitcoin as a legal tender alongside with the US dollar,³ while even in the historically prudent Switzerland, certain cantons made it possible to pay taxes in cryptocurrencies.⁴ Unsurprisingly, numerous crypto enthusiasts sincerely believed that cryptocurrencies would soon replace global banking system and make most bankers obsolete. While the recent collapse of Credit Suisse indeed triggered an electroshock effect on financial markets, its failure is unlikely to be attributable to Bitcoin or other cryptocurrencies, which currently raise even more concerns than global banking system. Smart contracts – which were proclaimed to replace bankers and lawyers altogether – failed to deliver the promised miracle so far. That being said, generative AI, of course, has a strong future potential for intelligent automation in multiple professions including legal practice, which we will briefly explore in this article alongside the concomitant risks and technical limitations.

II. Generative AI and Economics

Before delving into technical details, it may be worthwhile looking at generative AI through the cold-blood prism of economics. The author recently attended a cybersecurity webinar during which both

the audience and one of the speakers were excited by the fact – disclosed at the end of the webinar – that the entire 15-minute speech and slides of the speaker were created by generative AI. On the face of it, such capacities of a machine may indeed seem groundbreaking and trigger both fear and excitement among human audience. The reality is, however, far from being that simple. Answering one of the questions from the audience, the speaker admitted that he had spent over four hours experimenting and playing with different command prompts to generate and finetune the right content, and then spent one more hour to make some cosmetic adjustments to it. After talking to other attendees, who were proficient in the subject matter of the AI-powered presentation in question, we all agreed that preparation of the presented materials would normally take us less than two hours. This isolated example, of course, cannot be generalized or extrapolated to the multitude of generative AI use cases, but it aptly illustrates that, from the economic viewpoint, generative AI is not always as efficient as some AI technology vendors may advertise it. Therefore, make sure that implementation of generative AI in your organization will be driven by a solid economic foundation and not by a spontaneous decision to blindly follow fashionable trend.

III. Generative AI Before 2022

Most people became aware of generative AI and intelligent chatbots in 2022 with the public release of ChatGPT by OpenAI. Historically, generative AI and even primitive chatbots were already known in 1960s, when Artificial Neural Networks (ANNs) – the underlying technology of generative AI – were rapidly gaining popularity in scientific community.⁵ Nonetheless, the excitement about then-nascent technology vanished fairly quickly because then-primitive hardware unable to process the requisite volumes of data at necessary speed. Generative AI restored its popularity about a decade ago with rapidly evolving architectures of ANNs and underlying training techniques. Arguably, training is the most important component for any generative AI system: in a nutshell, it is the very process when the system learns to be “intelligent” by processing training data supplied by data scientists and engineers. Of note, as we will see below, in AI business, data is king. Probably the most important technical innovation and scientific breakthrough was the introduction of so-called Transformer,⁶ a deep-learning (a subclass of ANNs)

technique that significantly simplified, accelerated and boosted efficiency of ANNs training process on huge volumes of data, eventually leading to creation of Large Language Models (LLMs) that actually serve as “brains” of ChatGPT and other implementations of modern generative AI.

When dealing with texts, the beauty of LLMs is that, while having no human understanding of text or other content, the LLM model grasps billions of connections between words in large chunks of texts thereby considering the context to summarize, translate, review or generate text. This is where the “magic” actually comes from: LLMs behave as if they have cognitive capacities, analytical and reasoning skills akin to human beings.⁷ In reality, all LLMs actually do is a highly sophisticated prediction of the best possible sequence of words based on trillions of lines of human-written texts used for training. For example, when answering your questions, the model will simply predict the best possible sentences that shall follow the question. This is, of course, an oversimplification of a highly complex technical process, but it well explains why in the notorious case when ChatGPT invented non-existent legal precedents cited as binding case law, causing serious troubles for the law firm that incorporated the invented cases into their filing without verifying them first:⁸ the underlying LLM model simply does not understand the very concept of case law, it merely predicts which name of the case will be the most consonant to the training data. Similar incidents happen when generative AI creates software code or smart contracts:⁹ the code may contain security vulnerabilities, obsolete or even non-existing functions with invented parameters. Worse, if backdoored code is present in sufficient quantity among training data, the model may unwittingly generate a code with backdoor. In the next section, we will briefly explore the most frequent pitfalls and technical limitations of generative AI.

IV. Generative AI and Its Pitfalls

One of the key issues with LLMs are the so-called “hallucinations” when a model generates non-sense, erroneous or absurd content. High quality and sufficient quantity of training data are the most vital ingredients of any LLM model: the more you can get, the better AI-created content the model will produce. For example, the largely publicized experiment when GPT-4 – the most recent LLM model by OpenAI

at the time of writing – outperformed an average law school graduate on the UBE Bar exam.¹⁰ Given how many model MEE essays are publicly available by mushrooming Bar exam preparation companies, alongside MPT sample answers and countless simulated MBE questions with correct answers and explanations, it is perfectly foreseeable that GPT-4 did well on the Bar exam. If you change the exam format in a creative manner and also amend the exam subjects, GPT-4 will probably get lost and fail, while students will likely do well. The Bar exam case is, however, a good example of the immense technical progress made by generative AI during last few years thanks to Transformers and availability of high-performance computing for training.

Another major issue of LLMs are volatility and unpredictability of generated content. If you played with ChatGPT or other AI chatbots, you probably noticed that output for the very same input may vary each time, even when the so-called temperature – a parameter responsible for “creativity” of LLM, for instance, an unusually peculiar order or combination of generated words – is set to zero. AI technology vendors vigorously call to accept the unpredictability as a “new normal” and rather praise the inconsistency of results as a “new genius”, but when AI’s output is substantially different each time and its predictability verges on zero, the technology simply becomes inappropriate for use in legal practice and many other activities that require semantical precision and uniformity. It is important to note here that AI vendors are continually improving their LLMs, however, the improvements oftentimes make things even worse. For example, a group of researchers from Stanford University and UC Berkley demonstrated that both GPT-3.5 and GPT-4 improved their accuracy in some tasks and bodies of knowledge between March and June 2023, but at the same time, in other tasks and bodies of knowledge their accuracy has significantly degraded for unknown reasons.¹¹ Lack of consistency and high possibility of errors probably is one of the contributing factors why ChatGPT’s traffic has been declining for three summer months in a row during summer.¹² Paradigmatically, error detection can be particularly complicated, as AI-generated text tend to look neat, graceful and elegant to a human eye, aptly disguising subtle but fundamental errors.

V. Generative AI and Robolawyers

In 2023, AI evangelists have already made countless bold statements that generative AI would replace

lawyers and tectonically revolutionize the legal profession. Amusingly, it is not the first, and probably not the last, time when technology experts dream about replacing lawyers with a piece of code. While it is undisputed that technical progress continually and progressively shapes and improves legal practice, we are still far away from replacing human lawyers with AI or anything else.

First, most medium and large-sized legal firms around the globe, have been using various LegalTech solutions and products for over a decade already. Moreover, AI technologies including ANNs and other less-known forms of AI, such as decision trees and expert systems, have been reliable assistants of practicing attorneys in a broad spectrum of tasks including contract review, legal research, analysis of legal documents, and triage of e-discovery materials. Some of the tools use simple Machine Learning (ML) techniques for simple tasks such as regression or classification, producing reliable outcomes for some operational tasks. Eventually, introduction of ChatGPT and other AI-powered bots will unlikely revolutionize legal practice. It may, however, motivate law firms and even solo practitioners to accelerate their digital transformation and embrace new technologies, like the Internet did twenty years ago.

Second, as discussed above, inconsistency of AI-generated texts, say contracts or settlement agreements, may be a major problem for a law firm. In response to the inconsistency, law firms may either develop in-house LLMs, probably being a prohibitively expensive and unreasonably long process for most law firms, or perform the so-called fine-tuning of base models such as GPT family of models by OpenAI or LLaMA by Facebook (i.e., all-purpose models trained on vast volumes of all possible data to cover as many areas of knowledge as feasible, however, some areas of complex, rare or specific knowledge are usually covered in a superficial manner). Fine-tuning runs additional training of the base model by using industry-specific or even company-specific data to enable the LLM to produce better and more specific results tailored to particular needs. Getting back to economics and efficiency, creation of inconsistent albeit similar contracts for every new client may create an unmanageable mess for a law firm. When each contract contains different words and divergent word order, even in a routine contractual matter, a “creatively” placed comma

or expression may negate the very purpose of the contract, leading to a disaster for client and liability for law firm. This is not to mention LLMs' hallucinations discussed above, which will probably require a lawyer to proofread all AI-generated content. In sum, using a centrally managed repository of verified and internally vetted templates may be both a faster and more reliable approach to save time when drafting new legal documents. Having said that, drafting of pleadings, motions and answers is quite different, as these documents tend to be comparatively unique each time and few can be quickly assembled from one-size-fits-all templates. For such documents, generative AI may accelerate and even improve the drafting process by addressing more issues than a tired lawyer at the end of busy day in court. This will require a case file or other set of relevant legal documents to be provided as input to get a meaningful output. Importantly, all AI-generated content shall be meticulously verified by human expert before use in litigation or in court.

Third, as explained above, generative AI does not have cognitive abilities that human has. Therefore, using intelligent chatbots to get a legal advice may be a slippery slope for clients of the pseudo-legal service. In some jurisdictions, an excessively aggressive demand letter by inexperienced plaintiff may simply lead to criminal charges for extortion or menaces, and use of ChatGPT or other AI bot will unlikely be an excuse. This is probably the root cause why US state Bars coherently object over allowing AI to practice law:¹³ to prevent irremediable harm to clients, who will unlikely be able to file a legal malpractice lawsuit or at least lodge a disciplinary complaint against an AI bot from an overseas jurisdiction.

VI. Generative AI and Copyright Law

Despite acquiesced admission of using copyrighted and licensed content by AI technology vendors for their LLMs training purposes, the exact provenance of training datasets is unsurprisingly kept secret by them including OpenAI. At the time of writing (September 2023), it would be an arduous task to find a big law firm that has not yet produced an opinion about both copyrightability of AI-generated content and legality of using copyrighted materials for AI training purposes. Similar to fierce debates by legal scholars, opinions of lawyers are quite polarized about this matter. In the meanwhile, OpenAI and other providers of generative AI solutions are now

facing a snowballing number of lawsuits for, among other things, copyright infringement, violation of privacy rights¹⁴, unfair competition, unjust enrichment, defamation and related torts, and violation of terms of service agreements.

While copyright-related claims, including alleged violations of the Digital Millennium Copyright Act (DMCA) provisions,¹⁵ dominate the current AI litigation landscape, contemporary copyright legislation will unlikely provide a tenable protection to human creators of texts, music or images that have been exploited for Large Language Models (LLMs) training purposes without authors' permission. Technology giants seem to be certain, or at least are willing to publicly demonstrate such certainty, that the "fair use" doctrine or other available defenses will shield them from copyright infringement claims. For instance, Microsoft has just offered to indemnify commercial users of Microsoft Copilot,¹⁶ its AI-enabled assistant for software developers that is also accused of massive copyright-related violations and use of source code for training in violation of software license agreements.¹⁷ Even if copyright holders eventually prevail in court of law, it may be far too late. For instance, in the US, the spiraling dispute about copyright issues in generative AI has strong chances to end up at the Supreme Court, a journey that will probably take several years.¹⁸ Finally, despite that generative AI companies are already facing scrutiny for possible violations of competition and antitrust laws,¹⁹ the underlying questions of law will probably remain unsettled for years as well.

VII. Generative AI and Looming Regulation

Lawmakers from different countries are rapidly preparing AI regulations that will swiftly overshadow the copyright issues in generative AI, making some of them irrelevant. Some proposals of new legislation, like flat ban on generation of content that would violate law, remain technically utopic, however, others make perfect sense. Akin to pioneering privacy regulation with GDPR, the EU also leads model AI regulation with the proposed draft of the EU AI Act.²⁰ Initially introduced in a half-sleeping mode before the ChatGPT hype in 2021, after sensationalizing reports about novel risks of generative AI, the draft is now rapidly advancing to become an EU-wide law in a foreseeable future. Among other things, the Act requires certain categories of AI vendors to disclose

copyrighted data used for LLMs training. If enacted, the Act will open floodgates for breach-of-contract lawsuits for illicit data scrapping techniques that have been widely deployed by majority of leading generative AI vendors to collect training data for their super-powerful LLMs. Interestingly, in August 2023, the New York Times updated its website's terms of service to expressly prohibit any use of its online content for AI training and interconnected purposes.²¹ Of note, most websites already had anti-bot provisions in their terms of service for many years, trying to prevent malicious data-scrapping activities that, however, had nothing to do with AI training. Resultantly, a requirement for public disclosure of training data may be economically fatal for AI vendors, who will face an avalanche of lawsuits for breach of contract, a well-tested area of law with abundance of jurisprudence.

On the other side of the Atlantic, US tech giants are also actively working towards creation of better and safer AI,²² however, they are at the very beginning of the journey compared to other countries including China.²³ This creates a complicated dilemma for western lawmakers. On one hand, allowing western tech giants to monopolize generative AI space and to continue misappropriating proprietary content created by humans without their permission and without any remuneration may lead to highly undesirable social and economic consequences. On the other hand, overregulation will certainly hinder innovation and give advantage to foreign countries that will deliberately ignore any EU or US laws and aggressively accumulate knowledge by perfidiously using prohibited data-collection techniques, eventually winning the AI race.

Similar complexities prevail in AI-assisted practice of law: a flat prohibition of AI will unlikely be beneficial either for lawyers or their clients, while unregulated "practice" of legal profession by now-unreliable robots and unstable LLMs without a thorough supervision by legal professionals will certainly harm clients and society. Not only the courts may suddenly get flooded and paralyzed with meritless albeit well-written lawsuits created by bots, but countless plaintiffs and defendants will lose their otherwise winning cases because of technical glitches and will be left without a legal recourse against a judgement-proof AI startup in a foreign country. A possible solution may be a requirement that only lawyers can purchase and use

generative AI products for professional purposes to intelligently automate and accelerate their work, while staying accountable for preventable mistakes and omissions made by AI.

VIII. Conclusion

Despite the impressive progress in generative AI made in recent years, the underlying technology is still far away from being able to replace lawyers. Generative AI may, however, enhance the well-established LegalTech market by both performing previously non-automatable legal tasks and boosting the speed and accuracy of other ones. Generative AI shall be leveraged under thorough supervision of lawyers, who are to be eventually responsible for AI-generated content and work product, like they are for their non-lawyer employees. The future of generative AI is turbulent and unpredictable due to the spiraling legal uncertainty, looming regulations and restrictions, and ballooning number of lawsuits against AI vendors that have been collecting massive amounts of proprietary data without permission. In sum, the proclaimed revolution in legal profession by robotlawyers has failed so far.

Endnotes

1. "Data Security, Professional Perspective - ChatGPT: IP, Cybersecurity & Other Legal Risks of Generative AI" <https://www.bloomberglaw.com/external/document/X7K3GI38000000/data-security-professional-perspective-chatgpt-ip-cybersecurity/> February 2023
2. "Elon Musk's X tells users that it could use their posts to train AI models" <https://www.businessinsider.com/elon-musk-x-to-use-public-information-trains-ai-models-2023-9> September 1, 2023
3. "El Salvador's Bitcoin policy: Regret or reward?" <https://www.merit.unu.edu/el-salvador-bitcoin-policy/> August 21, 2023
4. "Tax payments with cryptocurrencies" <https://zg.ch/de/steuern-finanzen/steuern/steuerbezug/taxpaymentswithcryptocurrencies> accessed on September 18, 2023
5. "What is generative AI? Everything you need to know" <https://www.techtarget.com/searchen->

- [terpriseai/definition/generative-AI](#) September 2023
6. "Attention Is All You Need" <https://arxiv.org/abs/1706.03762> June 12, 2017
 7. "GPT-4 Can't Reason: Addendum" https://medium.com/@konstantine_45825/gpt-4-cant-reason-addendum-ed79d8452d44 September 3, 2023
 8. "Two US lawyers fined for submitting fake court citations from ChatGPT" <https://www.theguardian.com/technology/2023/jun/23/two-us-lawyers-fined-submitting-fake-court-citations-chatgpt> June 23, 2023
 9. "ChatGPT-coded smart contracts may be flawed, could 'fail miserably' when attacked: CertiK" <https://cointelegraph.com/news/ai-smart-contracts-flawed-fail-under-attack-certik> September 6, 2023
 10. "GPT-4 Passes the Bar Exam" <https://www.iit.edu/news/gpt-4-passes-bar-exam> March 15, 2023
 11. "How is ChatGPT's behavior changing over time?" <https://arxiv.org/abs/2307.09009> July 18, 2023
 12. "ChatGPT traffic slips again for third month in a row" <https://www.reuters.com/technology/chatgpt-traffic-slips-again-third-month-row-2023-09-07/> September 7, 2023
 13. "AI-powered "robot" lawyer won't argue in court after jail threats" <https://www.cbsnews.com/news/robot-lawyer-wont-argue-court-jail-threats-do-not-pay/> January 26, 2023
 14. "OpenAI, Microsoft hit with new US consumer privacy class action" <https://www.reuters.com/legal/litigation/openai-microsoft-hit-with-new-us-consumer-privacy-class-action-2023-09-06/> September 6, 2023
 15. "OpenAI asks court to trim authors' copyright lawsuits" <https://www.reuters.com/legal/litigation/openai-asks-court-trim-authors-copy-right-lawsuits-2023-08-29/> August 29, 2023
 16. "Microsoft announces new Copilot Copyright Commitment for customers" <https://blogs.microsoft.com/on-the-issues/2023/09/07/copilot-copyright-commitment-ai-legal-concerns/> September 7, 2023
 17. "GitHub and Copilot Intellectual Property Litigation" <https://www.saverilawfirm.com/our-cases/github-copilot-intellectual-property-litigation> accessed on September 18, 2023
 18. "Potential Supreme Court clash looms over copyright issues in generative AI training data" <https://venturebeat.com/ai/potential-supreme-court-clash-looms-over-copyright-issues-in-generative-ai-training-data/> September 6, 2023
 19. "AI Under the Antitrust Microscope: Competition Enforcers Focusing on Generative AI from All Angles" <https://www.antitrustlawblog.com/2023/08/articles/criminal-doj-ai-under-the-antitrust-microscope-competition-enforcers-focusing-on-generative-ai-from-all-angles/> August 9, 2023
 20. "EU AI Act: first regulation on artificial intelligence" <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence> June 6, 2023
 21. "The New York Times blocks OpenAI's web crawler" <https://www.theverge.com/2023/8/21/23840705/new-york-times-openai-web-crawler-ai-gpt> August 21, 2023
 22. "Nvidia, Palantir and more companies join White House AI pledge" <https://www.washingtonpost.com/technology/2023/09/12/white-house-ai-ibm-nvidia-pledge/> September 12, 2023
 23. "China's AI Regulations and How They Get Made" <https://carnegieendowment.org/2023/07/10/china-s-ai-regulations-and-how-they-get-made-pub-90117> July 10, 2023 ■

MEALEY'S LITIGATION REPORT: ARTIFICIAL INTELLIGENCE

edited by Bryan Redding

The Report is produced monthly by



11600 John F. Kennedy Blvd., Suite 1655, Philadelphia, PA 19103, USA

Telephone: (215)564-1788 1-800-MEALEYS (1-800-632-5397)

Email: mealeyinfo@lexisnexis.com

Web site: <http://www.lexisnexis.com/mealeys>

ISSN 2994-1105