



Legal Ethics: Artificial Intelligence and the Practice of Law

R' Mordechai Torczyner – torczyner@torontotorah.com

To receive a credit letter, please write your name in the Zoom Chat

Introduction

1. Professionalism Topics covered in this presentation

3.17 Technology in a law or legal services practice, including considerations related to security of information, productivity and efficiency

2. Relevant links

Past Legal Ethics sessions

<https://torontotorah.com/legal>

Must an Artificial Intelligence Observe Shabbat? (2019) <https://www.yutorah.org/lectures/lecture.cfm/928109/>

The Ethics of Robo-Advising (2021) <https://www.yutorah.org/lectures/lecture.cfm/1000099/>

Artificial Intelligence and Medical Decisions (2022) <https://www.yutorah.org/lectures/lecture.cfm/1037095/>

3. Benjamin Weiser, *Here's What Happens When Your Lawyer Uses ChatGPT*, New York Times May 27 '23

<https://www.nytimes.com/2023/05/27/nyregion/avianca-airline-lawsuit-chatgpt.html>

The lawsuit began like so many others: A man named Roberto Mata sued the airline Avianca, saying he was injured when a metal serving cart struck his knee during a flight to Kennedy International Airport in New York.

When Avianca asked a Manhattan federal judge to toss out the case, Mr. Mata's lawyers vehemently objected, submitting a 10-page brief that cited more than half a dozen relevant court decisions. There was *Martinez v. Delta Air Lines*, *Zicherman v. Korean Air Lines* and, of course, *Varghese v. China Southern Airlines*, with its learned discussion of federal law and "the tolling effect of the automatic stay on a statute of limitations."

There was just one hitch: No one — not the airline's lawyers, not even the judge himself — could find the decisions or the quotations cited and summarized in the brief.

That was because ChatGPT had invented everything.

The lawyer who created the brief, Steven A. Schwartz of the firm Levidow, Levidow & Oberman, threw himself on the mercy of the court on Thursday, saying in an affidavit that he had used the artificial intelligence program to do his legal research — "a source that has revealed itself to be unreliable."

Mr. Schwartz, who has practiced law in New York for three decades, told Judge P. Kevin Castel that he had no intent to deceive the court or the airline. Mr. Schwartz said that he had never used ChatGPT, and "therefore was unaware of the possibility that its content could be false."

He had, he told Judge Castel, even asked the program to verify that the cases were real.

It had said yes.

4. The proposed AIDA Act, part of the Digital Information and Privacy Act C-27

<https://ised-isde.canada.ca/site/innovation-better-canada/en/artificial-intelligence-and-data-act>

The AIDA will introduce new requirements for businesses to ensure the safety and fairness of high-impact AI systems every step of the way:

Design: Businesses will be required to identify and address the risks of their AI system with regard to harm and bias and to keep relevant records.

Development: Businesses will be required to assess the intended uses and limitations of their AI system and make sure users understand them.

Deployment: Businesses will be required to put in place appropriate risk mitigation strategies and ensure systems are continually monitored.

Vignettes

1> Binah & Daat LLP is a century-old law firm specializing in corporate law. After researching new Artificial Intelligence research tools, they want to bring on board *Chochmah*, AI-powered legal research software designed to scan databases, investigate relevant laws and precedents and produce accurate, succinct results. Are they obligated to notify clients that they are implementing this new tool?

- 2> Binah & Daat LLP is approached by *Chochmah*'s developers; they are training a next-generation AI, and want to use the data from Binah & Daat's extensive archives. In exchange, they will give Binah & Daat LLP free access to *Chochmah 2.0*. Are there ways for Binah & Daat to accommodate this request?
- 3> The lawyers of a small firm, Pachim Ketanim LLP, try to build their reputation by publishing articles on the practice of law. They start asking ChatGPT to generate articles based on the work of others – until they are threatened with a suit by a law library for plagiarizing their intellectual property. Is Pachim Ketanim LLP in the wrong?
- 4> In 2030, Pachim Ketanim LLP add *Kategor*, a "robot lawyer" which can process and prosecute small claims autonomously. The technology has been in use for five years, with very good reviews. But five months after it is introduced, *Kategor* makes a basic error in a client's case, and the client sues Pachim Ketanim. Who is liable for *Kategor*'s error?
- 5> Governments are beginning to use Artificial Intelligence software to filter immigration applications, extrapolating security issues and other concerns from demographic factors, social media profiles, work history, organizational affiliations and more. What ethical concerns does this practice raise?

Examples of Analytic and Generative AI and the Practice of Law

5. Lexis-Nexis, *The Power of Artificial Intelligence in Legal Research*

<https://www.lexisnexis.com/community/insights/legal/b/thought-leadership/posts/the-power-of-artificial-intelligence-in-legal-research>

Brief Analysis mines archives of legal precedents to find supporting documentation that bolsters your legal argument. The more it's used, the more it automatically "learns" and gets smarter, which makes it faster and easier to find and digest relevant documents.

Context is an AI-backed judicial analytics tool (sold separately from Lexis+) that analyzes a judge's past decisions, enabling lawyers to pinpoint the types of arguments a certain judge typically finds most persuasive. With access to not only ruling history, but also the specific type of language a judge uses in court, you can garner insights that make or break a case.

6. Arezou Rezvani, *'Robot Lawyer' Makes The Case Against Parking Tickets*, 1/16/17

<https://www.npr.org/2017/01/16/510096767/robot-lawyer-makes-the-case-against-parking-tickets>

It was the first day of school for Dan Lear's three kids. In a scramble to get his boys to class on time, the Seattle lawyer wound up parking in a space he probably should have avoided.

"There was a fire hydrant, but the curb wasn't painted and the fire hydrant was painted a kind of a funny color. And so I thought, and maybe it was wishful thinking, but I thought I would be OK to park there," he says.

Sure enough, Lear returned to a ticket.

"I was bummed! I mean, obviously no one's really happy when they get a ticket but I went home, I put it on my fridge and I let it sit there cause I just didn't want to deal with it," he says.

So he found something that would — DoNotPay, a free online "robot lawyer."

It has helped drivers in London and New York City appeal parking tickets.

It had just expanded into Seattle, and Lear decided to give it a go.

He logged in and the DoNotPay bot asked Lear a series of questions — like where the ticket was issued and a description of what happened. Within minutes, he had a 500-word letter to send to the city.

The verdict?

"Ultimately, yeah, they let me off," Lear says.

The idea behind DoNotPay belongs to Joshua Browder, a 20-year-old student at Stanford University who's originally from London.

7. Michael Loy, *Legal Liability for Artificially Intelligent "Robot Lawyers"*, Lewis and Clark Law Review 26:3

<https://law.lclark.edu/live/files/33853-2638loy>

Today's robot lawyers are no longer just classroom experiments that dabble in issue-spotting or legal research; attorneys use robot lawyers to perform substantive legal work and dispense real legal advice. For example, Kira Systems sells a robot lawyer to law firms that scans a client's previous contracts, identifies important terms within the contracts, and

generates a personalized draft for attorney review within seconds. LegalMation, another startup, promises attorneys that its robot lawyer will automatically draft “responsive pleadings, discovery requests and responses, and related documents that are tailored to the claims, allegations, and requests in the legal document uploaded, incorporating jurisdictional requirements as well as the attorney’s own style, formatting, and response strategy.” These “non-autonomous” robot lawyers, which are controlled by human attorneys, benefit the legal profession by increasing profit margins, reducing mundane work, and expanding access to justice.

Vignette 1: Transparency with clients

8. Law Society of Ontario, Rules of Professional Conduct (2014), Rules 3.1-1; 3.1-2; Commentary 14, 15 and 15.1
In this rule, “competent lawyer” means a lawyer who has and applies relevant knowledge, skills and attributes in a manner appropriate to each matter undertaken on behalf of a client...

A lawyer shall perform any legal services undertaken on a client’s behalf to the standard of a competent lawyer.

[14] A lawyer who is incompetent does the client a disservice, brings discredit to the profession and may bring the administration of justice into disrepute. In addition to damaging the lawyer’s own reputation and practice, incompetence may also injure the lawyer’s partners and associates.

[15] Incompetence, Negligence and Mistakes – This rule does not require a standard of perfection. An error or omission, even though it might be actionable for damages in negligence or contract, will not necessarily constitute a failure to maintain the standard of professional competence described in the rule. While damages may be awarded for negligence, incompetence can give rise to the additional sanction of disciplinary action.

[15.1] The *Law Society Act* provides that a lawyer fails to meet standards of professional competence if there are deficiencies in

- (a) the lawyer’s knowledge, skill, or judgment,
- (b) the lawyer’s attention to the interests of clients,
- (c) the records, systems, or procedures of the lawyer’s professional business, or
- (d) other aspects of the lawyer’s professional business,

and the deficiencies give rise to a reasonable apprehension that the quality of service to clients may be adversely affected.

9. Law Society of Ontario, Rules of Professional Conduct (2014), Rule 3.2-1

A lawyer has a duty to provide courteous, thorough and prompt service to clients. The quality of service required of a lawyer is service that is competent, timely, conscientious, diligent, efficient and civil.

10. Talmud, Bava Metzia 93a-b

ההוא רעיא דהוה קא רעי חיותא אגודא דנהר פפא, שריג חדא מינייהו ונפלת למיא, אתא לקמיה דרבה ופטריה. אמר, "מאי הוה ליה למיעבד? הא נטר כדנטרי אינשי!" אמר ליה אביי: אלא מעתה, על למתא בעידנא דעיילי אינשי הכי נמי דפטור? אמר ליה: אין... איתיביה, "עד מתי שומר שכר חייב לשמור? עד כדי הייתי ביום אכלני חרב וקרח בלילה!" אמר ליה, התם נמי בחזוני מתא. אמר ליה, אטו יעקב אבינו חזן מתא הוה? דאמר ליה ללבן, "נטרי לך נטירותא יתירתא כחזוני מתא."

A shepherd was herding animals on the bridge of the Pappa River; one slipped and fell into the water. He came before Rabbah, who said he was not liable, saying, "What should he have done? He guarded as people normally do!"

Abbaye replied: Then had he gone into the city at a time when people normally do, would he have been exempt?

He said: Yes.

Had he slept a bit, at a time when people sleep, would he also have been exempt?

He said: Yes...

[Abbaye] challenged from a text, "How far must a paid guardian go in guarding? To the point of [Jacob’s declaration], 'I was consumed by withering heat during the day, and frost at night!'"

He said: That is a case of a municipal guard.

[Abbaye] said, "Was our ancestor Jacob a municipal guard?"

He told Lavan, "I will guard for you at a higher level, as municipal guards do."

11. Rabbi Moses Maimonides (12th century Egypt), Mishneh Torah, Hilchot Sechirut 13:7

כדרך שמוזהר בעה"ב שלא יגזול שכר עני ולא יעכבנו, כך העני מוזהר שלא יגזול מלאכת בעה"ב ויבטל מעט בכאן ומעט בכאן ומוציא כל היום במרמה. אלא חייב לדקדק על עצמו בזמן, שהרי הקפידו על ברכה רביעית של ברכת המזון שלא יברך אותה. וכן חייב לעבוד בכל כחו שהרי יעקב הצדיק אמר "כי בכל כחי עבדתי את אביכן." לפיכך נטל שכר זאת אף בעולם הזה שנאמר "ויפרץ האיש מאד מאד."

Just as an employer is warned not to steal the pay of a pauper and not to delay it, so the pauper is warned not to steal the work of the employer, wasting a little here and a little there and so spending the entire day cleverly. Rather, he must be precise with his time, for the Sages were even careful not to recite the fourth blessing of the blessing after meals [on an employer's time]. He must also work with all of his energy, for the righteous Jacob said, "I worked for your father with all of my energy." Therefore [Jacob] was rewarded even in this world, as it says, "And the man became very, very great."

12. Midrash Tannaim to Deuteronomy 24:15

"ואליו הוא נשא את נפשו" הרי זו אזהרה לפועל שיעשה בכל כוחו:

"For this he risks his life" – This warns an employee to work with all of his strength.

13. Rabbi Joseph Caro (16th century Turkey/Israel), Shulchan Aruch, Choshen Mishpat 331:1-2

השוכר את הפועלים ואמר להם להשכים ולהעריב, מקום שנהגו שלא להשכים ושלא להעריב אינו יכול לכוּפֵן, אפילו הוסיף על שכרן, כיון שלא התנה כן בשעה ששכרן.

מקום שנהגו לזון, יזון; לספק בגרוגרות או בתמרים וכיוצא בהם, יספק; הכל כמנהג המדינה.

In a place where the practice is not to rise early and stay late, one who hires employees and tells them to rise early and stay late, cannot coerce them, even based on higher wages. This is because he didn't stipulate it when hiring them.

Where the practice is to feed them, one must feed them. To provide dried figs, dates and the like, provide it. All is according to local practice.

14. Harish, Morgado, Stern, Das, *Artificial Intelligence and Clinical Decision Making: The New Nature of Medical Uncertainty*, Journal of Academic Medicine Jan '21

By the end of its 3-game *Jeopardy* run, Watson had defeated its human competitors by a considerable margin. While this result was impressive, Watson's most memorable moment for some came during the final round when it responded "Toronto" to a question about American cities. In this instance, Watson's probabilistic answering design prevented it from excluding any solutions with total certainty, leading to an incorrect (albeit low-confidence) conclusion that the audience knew was obviously incorrect.

15. Amann, Blasimme, Vayena, Frey, Madai, *Explainability for artificial intelligence in healthcare: a multidisciplinary perspective*, Medical Informatics and Decision Making, Nov '20

<https://bmcmmedinformdecismak.biomedcentral.com/articles/10.1186/s12911-020-01332-6>

From the development point-of-view, explainability will regularly be helpful for developers to sanity check their AI models beyond mere performance. For example, it is highly beneficial to rule out that the prediction performance is based on meta-data rather than the data itself. A famous non-medical example was the classification task to discern between huskies and wolves, where the prediction was solely driven by the identification of a snowy background rather than real differences between huskies and wolves. This phenomenon is also called a "Clever Hans" phenomenon. Clever Hans phenomena are also found in medicine. An example is the model developed by researchers from Mount Sinai hospital which performed very well in distinguishing high-risk patients from non-high-risk patients based on x-ray imaging. However, when the tool was applied outside of Mount Sinai, the performance plummeted. As it turned out the AI model did not learn clinically relevant information from the images. In analogy to the snowy background in the example introduced above, the prediction was based on hardware related meta-data tied to the specific x-ray machine that was used to image the high-risk ICU patients exclusively at Mount Sinai...

16. Challen, Denny, Pitt, Gompels, Edwards, Tsaneva-Atanasova, *Artificial intelligence, bias and clinical safety*, BMJ Quality & Safety '19 <https://qualitysafety.bmj.com/content/28/3/231>

As humans, clinicians are susceptible to a range of cognitive biases which influence their ability to make accurate decisions... Automation bias describes the phenomenon whereby clinicians accept the guidance of an automated system

and cease searching for confirmatory evidence (e.g., see Tsai *et al*), perhaps transferring responsibility for decision-making onto the machine—an effect reportedly strongest when a machine advises that a case is normal...

Vignette 2: Confidentiality and AI Training

17. Anthony Tockar, *Riding with the Stars: Passenger Privacy in the NYC Taxicab Dataset*, Neustar Research 9/15/14
There has been a lot of online comment recently about a dataset released by the New York City Taxi and Limousine Commission. It contains details about every taxi ride (yellow cabs) in New York in 2013, including the pickup and drop off times, locations, fare and tip amounts, as well as anonymized (hashed) versions of the taxi's license and medallion numbers. It was obtained via a FOIL (Freedom of Information Law) request earlier this year and has been making waves in the hacker community ever since.

The release of this data in this unalloyed format raises several privacy concerns. The most well-documented of these deals with the hash function used to "anonymize" the license and medallion numbers. A bit of lateral thinking from one civic hacker and the data was completely de-anonymized. This data can now be used to calculate, for example, any driver's annual income. More disquieting, though, in my opinion, is the privacy risk to passengers. With only a small amount of auxiliary knowledge, using this dataset an attacker could identify where an individual went, how much they paid, weekly habits, etc. I will demonstrate how easy this is to do in the following section....

I spent some of the most riveting hours of my professional career searching through images of "celebrities in taxis in Manhattan in 2013" to find enough information to identify the correct record in the database. I had some success – combining the below photos of BC and JA with some information from celebrity gossip blogs allowed me to find their trips, which are shown in the accompanying maps.

In BC's case, we now know that his cab took him to Greenwich Village, possibly to have dinner at Melibea, and that he paid \$10.50, with no recorded tip. Ironically, he got in the cab to escape the photographers! We also know that JA got into her taxi outside her hotel, the Trump SoHo, and somewhat surprisingly also did not add a tip to her \$9 fare. Now while this information is relatively benign, particularly a year down the line, I have revealed information that was not previously in the public domain. Considering the speculative drivel that usually accompanies these photos (trust me, I know!), a celebrity journalist would be thrilled to learn this additional information...

But OK, perhaps you're not convinced. After all, this dataset is (thankfully) not real-time. How about we leave the poor celebrities alone and consider something a little more provocative. Larry Flynt's Hustler Club is in a fairly isolated location in Hell's Kitchen, and no doubt experiences significant cab traffic in the early hours of the morning. I ran a query to pull out all pickups that occurred outside the club after midnight and before 6am, and mapped the drop-off coordinates to see if I could pinpoint individuals who frequented the establishment. The map below shows my results – the yellow points correspond to drop-offs that are closely clustered, implying a frequent customer...

Examining one of the clusters in the map above revealed that only one of the 5 likely drop-off addresses was inhabited; a search for that address revealed its resident's name. In addition, by examining other drop-offs at this address, I found that this gentleman also frequented such establishments as... Using websites like Spokeo and Facebook, I was also able to find out his property value, ethnicity, relationship status, court records and even a profile picture!

18. Selling of information on physicians' prescribing patterns: PIPEDA Case Summary #2001-14

<https://www.priv.gc.ca/en/opc-actions-and-decisions/investigations/investigations-into-businesses/2001/pipeda-2001-014/>

In two separate complaints an individual and a physician complained that the Canadian arm of a U.S.-based international marketing firm was improperly disclosing personal information by gathering and selling data on physicians' prescribing patterns without their consent.

19. LSO, Rules of Professional Conduct (2014), Rule 3.3-1 and Commentary [1]

A lawyer at all times shall hold in strict confidence all information concerning the business and affairs of the client acquired in the course of the professional relationship and shall not divulge any such information unless

- (a) expressly or impliedly authorized by the client;
- (b) required by law or by order of a tribunal of competent jurisdiction to do so;
- (c) required to provide the information to the Law Society; or

(d) otherwise permitted by rules 3.3-2 to 3.3-6.

[1] A lawyer cannot render effective professional service to the client unless there is full and unreserved communication between them. At the same time, the client must feel completely secure and entitled to proceed on the basis that, without any express request or stipulation on the client's part, matters disclosed to or discussed with the lawyer will be held in strict confidence.

20. Talmud, Yoma 4b

מניין לאומר דבר לחבירו שהוא בבל יאמר לו עד שיאמר לו לך אמור שנאמר "וידבר ד' אליו מאהל מועד לאמר"

How do we know that one may not repeat that which is he told, until he is told, 'Go tell'?' "And Gd told Moshe in the Tent of Meeting, to tell [the Jews]."

21. Rabbi Moses Maimonides (12th century Egypt), Mishneh Torah, Hilchot Deiot 7:5-6

המספר דברים שגורמים אם נשמעו איש מפי איש להזיק חבירו בגופו או בממונו ואפילו להצר לו או להפחידו הרי זה לשון הרע... כל אלו הם בעלי לשון הרע שאסור לדור בשכונתם וכל שכן לישב עמם ולשמוע דבריהם

Statements which, if publicized, would cause financial harm, physical harm, or even pain or fear, are *lashon hara*... One may not live in such people's neighbourhoods, let alone sit with them and listen to them.

22. Mishnah Bava Batra 3:7 (60a)

מתני' לא יפתח אדם לחצר השותפין פתח כנגד פתח וחלון כנגד חלון היה קטן לא יעשנו גדול אחד לא יעשנו שנים אבל פותח הוא לרה"ר פתח כנגד פתח וחלון כנגד חלון היה קטן עושה אותו גדול ואחד עושה אותו שנים

Mishnah: One may not open a window opposite a window, or a door opposite a door, across a yard owned in partnership. One may not increase their size, or turn one into two. However, one may do this into a public domain.

Vignette 3: Intellectual Property and Generative AI

23. Kyle Wiggers, *The current legal cases against generative AI are just the beginning*, Jan 27 '23

<https://techcrunch.com/2023/01/27/the-current-legal-cases-against-generative-ai-are-just-the-beginning/>

Microsoft, GitHub and OpenAI are currently being sued in a class action motion that accuses them of violating copyright law by allowing Copilot, a code-generating AI system trained on billions of lines of public code, to regurgitate licensed code snippets without providing credit.

Two companies behind popular AI art tools, Midjourney and Stability AI, are in the crosshairs of a legal case that alleges they infringed on the rights of millions of artists by training their tools on web-scraped images.

And just last week, stock image supplier Getty Images took Stability AI to court for reportedly using millions of images from its site without permission to train Stable Diffusion, an art-generating AI.

At issue, mainly, is generative AI's tendency to replicate images, text and more — including copyrighted content — from the data that was used to train it. In a recent example, an AI tool used by CNET to write explanatory articles was found to have plagiarized articles written by humans — articles presumably swept up in its training dataset. Meanwhile, an academic study published in December found that image-generating AI models like DALL-E 2 and Stable Diffusion can and do replicate aspects of images from their training data.

24. Copyright Act (R.S.C., 1985, c. C-42) s. 29

Fair dealing for the purpose of research, private study, education, parody or satire does not infringe copyright.

25. Riddhi Setty, *First AI Art Generator Lawsuits Threaten Future of Emerging Tech*, 1/20/23

<https://news.bloomberglaw.com/ip-law/first-ai-art-generator-lawsuits-threaten-future-of-emerging-tech>

In the complaint, the artists allege that AI tools permit users to create works "in the style of" a given artist instead of commissioning or licensing an original work, violating the rights of millions of artists and resulting in works that compete with the originals.

Style has proven difficult, if not impossible, to protect with copyright, some attorneys say.

"Style is an idea, style is a technique, it's a process," said University of Kentucky law professor Michael Murray. "These are things that are not protected by copyright."

Given an AI generator's ability to draw directly from the works of art it's trained on, however, there are circumstances in which "the line between emulating style and wrongful copying is not always apparent," said Columbia Law Professor Jane Ginsburg.

26. Riddhi Setty, *First AI Art Generator Lawsuits Threaten Future of Emerging Tech*, 1/20/23

The success of a fair use defense will depend on whether the works generated by the AI are considered transformative—whether they use the copyrighted works in a way that significantly varies from the originals.

"The minute you're able to make a work that's transformative, you removed it from the ambit of this idea that it's a mere derivative of the original work," said Vivek Jayaram, founder of Jayaram Law.

Jayaram said that previous case law, particularly the Supreme Court's 2021 *Google v. Oracle* decision, suggests that using collected data to create new works can be transformative. In that case, Google's use of portions of Java SE code to create its Android operating system was found to be fair use.

"Using IP or that training data to allow others to use the tool to create new work—that construct, in and of itself, I don't think will be found to be completely infringing," said Jayaram. He cautioned, though, that without transformation of a work, AI generators essentially serve as digital art counterfeiting machines.

27. OpenAI, Comment Regarding Request for Comments on Intellectual Property Protection for AI Innovation

https://www.uspto.gov/sites/default/files/documents/OpenAI_RFC-84-FR-58141.pdf

Training of AI systems is clearly highly transformative. Works in training corpora were meant primarily for human consumption for their standalone entertainment value. The "object of the original creation," in other words, is direct human consumption of the author's expression. Intermediate copying of works in training AI systems is, by contrast, "non-expressive": the copying helps computer programs learn the patterns inherent in human-generated media. The aim of this process—creation of a useful generative AI system—is quite different than the original object of human consumption. The output is different too: nobody looking to read a specific webpage contained in the corpus used to train an AI system can do so by studying the AI system or its outputs. The new purpose and expression are thus both highly transformative.

28. Riddhi Setty, *First AI Art Generator Lawsuits Threaten Future of Emerging Tech*, 1/20/23

Matthew Sag, a law professor at Emory University focused on AI, said that AI training could be considered fair use under a 2015 Second Circuit case that the Supreme Court declined to review. The decision in *Authors Guild v. Google* found that digitizing books and displaying snippets of them without permission to allow text searching is transformative fair use.

In that context the fair use was clear, Sag said, "because none of the original expression leaks out in the end."

29. OpenAI, Comment Regarding Request for Comments on Intellectual Property Protection for AI Innovation

The fair use doctrine "permits courts to avoid rigid application of the copyright statute when, on occasion, it would stifle the very creativity which that law is designed to foster." AI systems hold immense promise for both creative expression and general economic innovation. Copyright barriers to training AI systems would have "disastrous ramifications" and "could jeopardize the technology's social value, or drive innovation to a foreign jurisdiction with relaxed copyright constraints." We thus submit that such barriers would "stifle the very creativity which [copyright] law is designed to foster" and retard "the Progress of Science and useful Arts."

30. Congressional Research Service, *Generative Artificial Intelligence and Copyright Law*, May 11'23

<https://crsreports.congress.gov/product/pdf/LSB/LSB10922>

AI systems are "trained" to create literary, visual, and other artistic works by exposing the program to large amounts of data, which may consist of existing works such as text and images from the internet. This training process may involve making digital copies of existing works, carrying a risk of copyright infringement. As the U.S. Patent and Trademark Office has described, this process "will almost by definition involve the reproduction of entire works or substantial portions thereof." OpenAI, for example, acknowledges that its programs are trained on "large, publicly available datasets that include copyrighted works" and that this process "necessarily involves first making copies of the data to

be analyzed." Creating such copies, without express or implied permission from the various copyright owners, may infringe the copyright holders' exclusive right to make reproductions of their work.

31. Mia Sato, *Drake's AI clone is here – and Drake might not be able to stop him*, May 1 '23

<https://www.theverge.com/2023/5/1/23703087/ai-drake-the-weeknd-music-copyright-legal-battle-right-of-publicity>

What if AI-generated "unreleased" Drake tracks surfaced and diverted revenue from actual Drake? What if the AI songs are just bad, and people were convinced Drake lost his magic touch for hits? What if a creator made AI Drake sing a white nationalist anthem? The problem quickly expands beyond the scope of copyright and into Drake's personhood and identity — and unlike the wave of unknowns with AI, there is some precedent with how a person's likeness is used.

32. Copyright Act (R.S.C., 1985, c. C-42) s. 29.21 (1) (d)

It is not an infringement of copyright for an individual to use an existing work or other subject-matter or copy of one, which has been published or otherwise made available to the public, in the creation of a new work or other subject-matter in which copyright subsists and for the individual — or, with the individual's authorization, a member of their household — to use the new work or other subject-matter or to authorize an intermediary to disseminate it, if...

(d) the use of, or the authorization to disseminate, the new work or other subject-matter does not have a substantial adverse effect, financial or otherwise, on the exploitation or potential exploitation of the existing work or other subject-matter — or copy of it — or on an existing or potential market for it, including that the new work or other subject-matter is not a substitute for the existing one.

33. Rabbi Yitzchak Alfasi (11th century Morocco/Spain), Responsum 133

שאלה. א' מן התלמידים גנב ספרי פרושין לחבירו וכשתבעו ממנו נשבע שבועה חמורה שלא יחזירם לו עד שיעתיק אותם ויש מי שהורה שמותר לגנבן. תשובה: הגונב והמורה כלם טעו, שלא כדין עשו. שזה שהורה שמותר לעשות כן דומה לו משום שמתלמד דברי תורה שהיא מצוה היא וטעה משום שאמר לולב הגזול והיבש פסול ואתמר עלה משום דמצוה הבאה בעבירה במקום שונאה...

Question: One of the students stole another person's books of commentaries. When he demanded it back, [the thief] swore a severe oath that he would not return them before copying them. Someone ruled that he could steal them.

Answer: The thief and the one who ruled are both mistaken, and acted improperly. The one who ruled that this is permitted thinks that since he is learning Torah this is a mitzvah, but he is wrong, since [the Talmud] says, "A stolen or dried out lulav is disqualified," and we learn regarding this that Gd hates a mitzvah performed via transgression...

34. Rabbi Yaakov Yeshayah Blau (21st century Israel), Pitchei Choshen V 9:11

ונראה שכל המצאה או חידוש באומנות או במלאכה (כגון פטנט) יש בו משום השגת גבול.

It seems that every discovery or invention in a craft or form of work (like a patent) is subject to the rule against invading a boundary.

35. Rabbi Malkiel Zvi Tannenbaum (19th century Poland), Divrei Malkiel 3:157

י"ל דה"ה בנ"ד דכיון שמדפיס העטיקעטין כמו של פלוני הרי חזינן דניהא לו בהוצאה וממילא יתחייב לשלם חלק בהוצאות הראשון על הרשיון.

One could say that the same is true in our case; since he prints the *etiketen* like this person's *etiketen*, we see that he is happy with that expenditure, and so he will be obligated to pay part of the money spent by the first party for the license.

36. Responsa of Rabbi Moshe Isserles (16th century Poland), #10

לא נניח את שלנו ולומר שכל ישראל ומי שבשם ישראל יכונה לא יקנה שום ספר מיימוני החדשים רק מאותן היוצאים מתחת יד הגאון הנ"ל או באי כחו, והוא מארבעה טעמים.

We will not abandon our part. We declare that all Israel, and anyone known as Israel, should not purchase any new copy of Maimonides other than those which emerge from the hand of the aforementioned Gaon or those whom he has empowered, for the following four reasons...

37. Talmud, Bava Batra 22a

רב דימי מנהרדעא אייתי גרוגרות בספינה א"ל ריש גלותא לרבא פוק חזי אי צורבא מרבנן הוא נקיט ליה שוקא

Rav Dimi of Neherdaa brought dried figs by boat. The Exilarch said to Rava: Go and see; if he is a young scholar, set up the market for him.

Vignette #4: Product Liability

38. Michael Loy, *Legal Liability for Artificially Intelligent "Robot Lawyers"*, Lewis and Clark Law Review 26:3

Regardless of whether it is reasonable for an attorney to delegate legal work to an artificially intelligent robot lawyer, the attorney is liable for any harms attributable to the robot lawyer because, first, the attorney has a duty to accept ultimate responsibility for the use of robot lawyers used as tools in his or her law practice. Second, even if the attorney considers an artificially intelligent robot lawyer more like a human employee than a software tool, the attorney still has a duty to accept ultimate responsibility for robot lawyers acting under his or her supervision...

Although exact definitions of foreseeability vary between jurisdictions, foreseeable harm generally falls within a category of consequences that a reasonably prudent person anticipates will likely occur...

To succeed in any action against an autonomous robot lawyer, a plaintiff must first establish that the court has jurisdiction over the robot lawyer. Courts may establish personal jurisdiction over property using in rem jurisdiction and may establish personal jurisdiction over persons using in personam jurisdiction...

Legislatures should make modest changes to the law to grant autonomous robot lawyers "limited" personhood rights sufficient to establish in personam jurisdiction. For example, legislatures could effectively grant autonomous robot lawyers the same rights and responsibilities as corporations and other legal entities by simply allowing the robot lawyers themselves to join business organizations. Or, autonomous robot lawyers could simply be required to register with a government entity before providing legal advice to the public.

39. Rabbi Joseph Caro (16th century Turkey/Israel), Shulchan Aruch Choshen Mishpat 420:25

צעק באזנו וחרשו, פטור מדיני אדם והייב בדיני שמים.

One who screams in the ear of another, deafening him, is not liable by human law, but is liable in heavenly law.