Artificial Intelligence—Benefits and Ethical Issues

Nancy B. Rapoport
UNLV Distinguished Professor
Garman Turner Gordon Professor of Law
William S. Boyd School of Law
University of Nevada, Las Vegas

# When you think of AI, what comes to mind?

AI hopes and fears?

HAL 9000 in 2001: A Space Odyssey (1968)

I know I've made some very poor decisions recently, but I can give you my complete assurance that my work will be back to normal. I've still got the greatest enthusiasm and confidence in the mission. And I want to help you.

UNLV | William S. Boyd
SCHOOL OF LAW

Let's bring this one home: *2001: A Space Odyssey* (MGM 1968):

Do you read me, Hal?

How my spouse views GenAI:

Poltergeist
(MGM/UA 1982)

08-10-21

Robots are coming for the lawyers

Bad for lawyers' salaries, good for people who need cheap legal help.

THE CONVERSATION
Academic rigor, journalistic flair

Arts + Culture   Economy   Education   Environment + Energy   Ethics + Religion   Health   Politics + Society   Science + Tech   Podcasts

Robots are coming for the lawyers – which may be bad for tomorrow's attorneys but great for anyone in need of cheap legal assistance

Published: August 9, 2021 3.17pm EDT

FORBES > INNOVATION > AI

Will A.I. Put Lawyers Out Of Business?

UNLV | WILLIAM S. BOYD
SCHOOL OF LAW

AI is already here (and has been for a while):

UNLV | WILLIAM S. BOYD
SCHOOL OF LAW

What is artificial intelligence (AI)?

What is AI? There are many ways to answer this question, but one place to begin is to consider the types of problems that AI technology is often used to address. In that spirit, we might describe **AI as using technology to automate tasks that "normally require human intelligence."** This description of AI emphasizes that the technology is often focused upon automating specific types of tasks: those that are thought to involve intelligence when people perform them.

Harry Surden, *Artificial Intelligence and Law: An Overview,* 35 GA. ST. L. REV. 1305, 1307 (2019) (footnotes omitted).

A primer on large language models (LLMs):

INGESTS LARGE AMOUNTS OF AVAILABLE DATA.

SEARCHES THROUGH THAT DATA TO GIVE YOU ANSWERS.

TRIES TO GIVE YOU WHAT YOU WANT (SO YOU NEED TO BE CAREFUL ABOUT *HOW* YOU ASK FOR SOMETHING).

But remember, "garbage in, garbage out":

How does GenAI make it possible to do some work formerly generated only by the human brain?

VECTORS. Here's expert Damien Riehl:

> In LLMs, closely related terms *linguistically* are also nearby *mathematically* (because those terms are in close proximity in the 'statistically likely' sense). For example, the blank in the sentence "The hurricane triggered the < BLANK> clause" could be filled with either "Force Majeure" or "Act of God." They're both statistically likely. So in vector space, they're near each other.

Damien Riehl, *We Need to Talk About ChatGPT,* 80 Jun-Bench & B. Mo. 26 (2023).

WILLIAM S. BOYD
SCHOOL OF LAW

So GenAI is trying to predict what it is that you want from it.

- It's a bit like watching Vanna White turn letters on *Wheel of Fortune:**

UNLV | William S. Boyd SCHOOL OF LAW

That "guessing" is part of what causes "hallucinations":

- *Mata v. Avianca Inc.,* Case No. 22-cv-1461 (PKC) (S.D.N.Y).
  - Briefs contained "hallucinated" cases.
  - Opposing counsel, and then the court, asked to see the cases.
  - Of course, the cases didn't exist.
  - (Also, never tell a judge that you're going on vacation if you're not.)

- For a fun dialogue between Hon. Cynthia A. Norton and me discussing lessons learned from this case, see *Doubling Down on Dumb: Lessons from Mata v. Avianca Inc.,* ABI J. (Aug. 2023).

WILLIAM S. BOYD
SCHOOL OF LAW

So GenAI isn't perfect.

- But it has [passed the bar exam](#).

On the other hand, some major LLMs are "unstable":

"An LLM is 'stable' if it reaches the same conclusion when asked the identical question multiple times. We find leading LLMs like gpt-4o, claude-3.5, and gemini-1.5 are unstable when providing answers to hard legal questions, even when made as deterministic as possible by setting temperature to 0. We curate and release a novel dataset of 500 legal questions distilled from real cases, involving two parties, with facts, competing legal arguments, and the question of which party should prevail. When provided the exact same question, we observe that LLMs sometimes say one party should win, while other times saying the other party should win. This instability has implications for the increasing numbers of legal AI products, legal processes, and lawyers relying on these LLMs."

[Some examples of AI companies]():
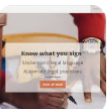
- OpenAI

- Google

- IBM

- Microsoft

- NVIDIA

- Amazon

- Anthropic

- Anduril

And some more:

Casetext ⌄	ChatGPT ⌄	Smith.ai ⌄
Harvey AI ⌄	Lex Machina ⌄	Ai Legal Analytics Tools ⌄
CoCounsel ⌄	Blue J I&e ⌄	Diligen ⌄
Everlaw, Inc. ⌄	Latch ⌄	Auto-gpt ⌄
Claude AI ⌄	Clio ⌄	Lawgeex ⌄
Legalrobot ⌄	Luminance ⌄	Ai Ediscovery Tools ⌄

I've just tested this one as part of a state bar working group:

And, of course, there's:

grammarly

alexa

THOMSON REUTERS
WESTLAW™

LexisNexis®

UNLV | WILLIAM S. BOYD
SCHOOL
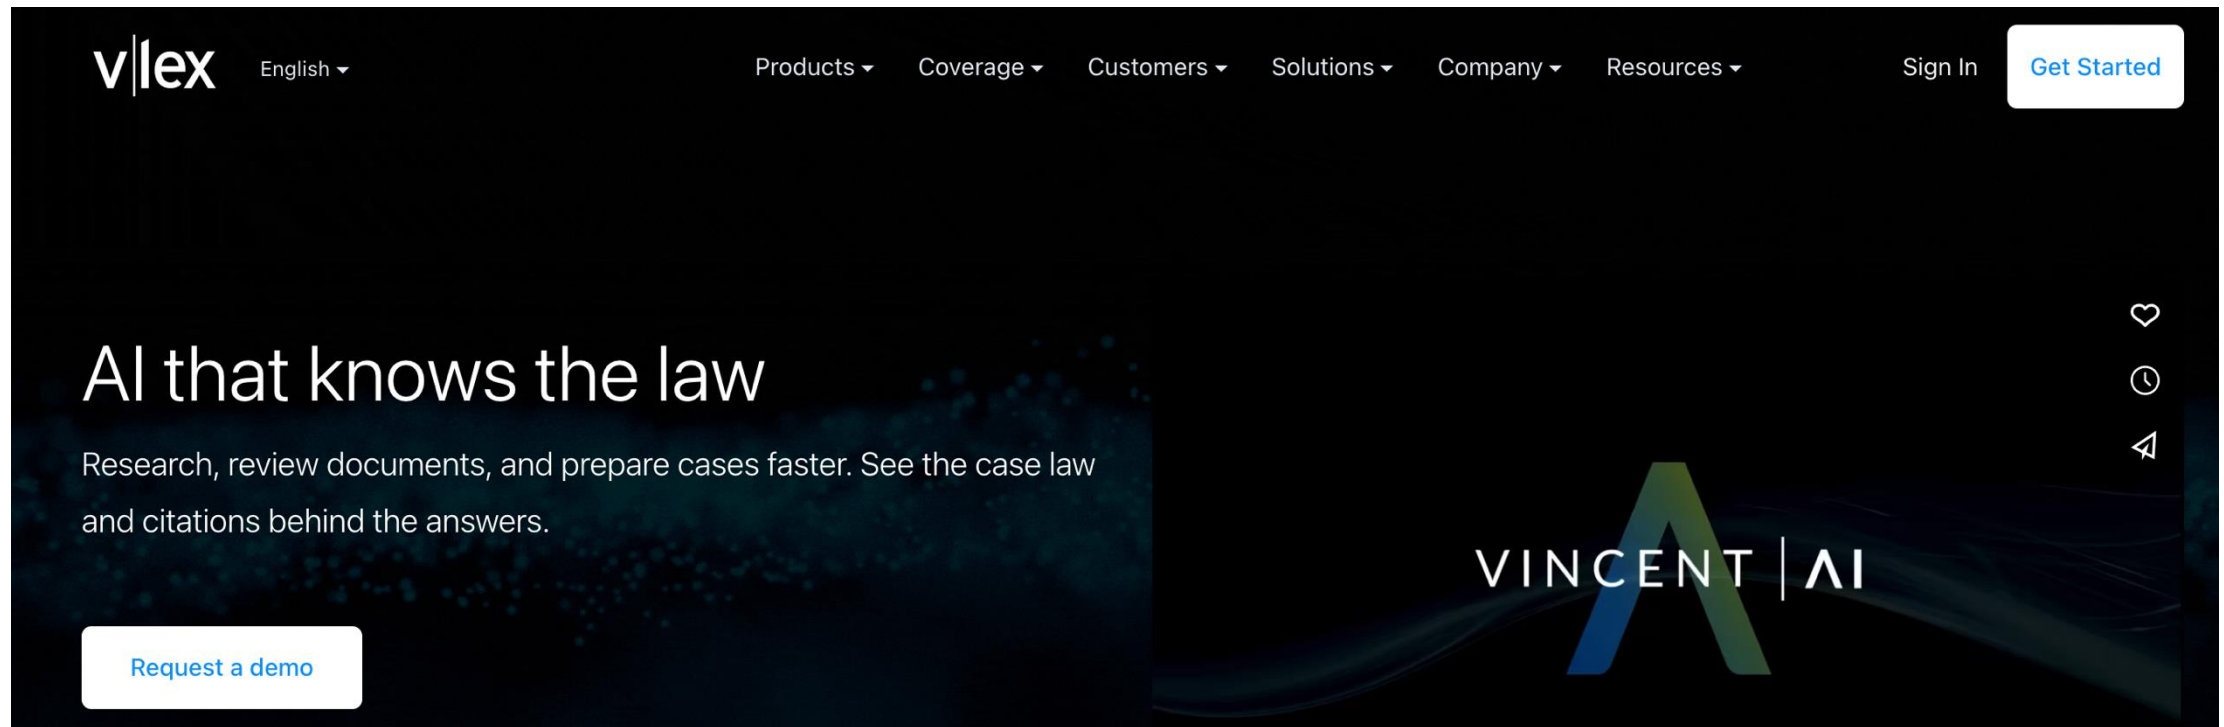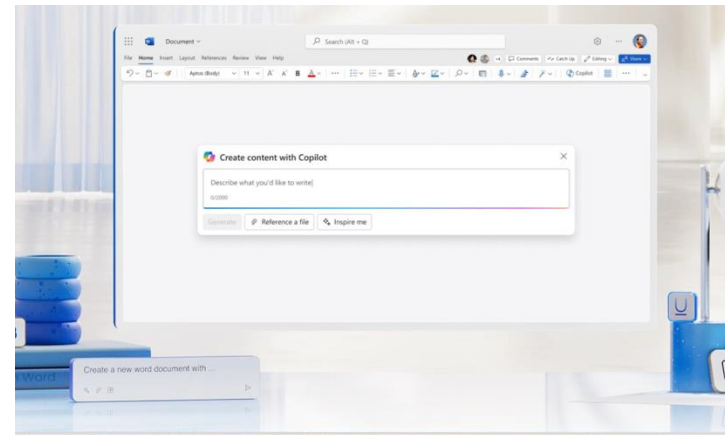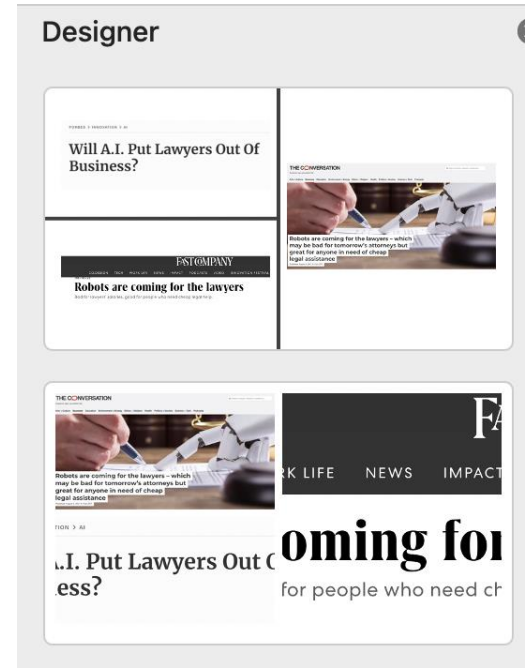
Copilot in Word

Explore how you can use the power of AI in Word

Clio is a good source for lawyers who want to incorporate AI:

- Clio calls our attention to:
    - Casetext (part of Thomson Reuters).
    - Co-Counsel (part of Casetext—can do everything from document review to deposition prep).
    - Harvey (huge flexibility as well).
    - Blue J (also a legal research tool).
    - Diligen (document review).
    - Smith.ai or Gideon (customer service).
    - Anthropic's Claude (a general AI application).
    - ChatGPT and DALL-E (with DALL-E able to create some of today's art).

But can these robots help us, or will they make our lives worse?

Asimov's three laws of robots*

- **First Law:** A robot must not harm a human or allow a human to come to harm through inaction.
- **Second Law:** A robot must obey human orders, unless doing so would conflict with the First Law.
- **Third Law**: A robot must protect its own existence, unless doing so would conflict with the First or Second Law.

* Isaac Asimov, *Runaround,* I, Robot 41-42 (1950).

UNLV | William S. Boyd
SCHOOL OF LAW

AI safety vs. AI security:

AI security

- Confidentiality.
- Integrity of the AI program.
- No unauthorized access.
- Protection against data breaches and data "poisoning" that can compromise the LLM itself.

AI safety

- (Those three laws of robotics, basically...)
- How does the AI affect society?
- Ethical implications of its use.
- Was the LLM designed with good values in mind?
- Open-source (transparency) vs. closed-source (harder to disrupt) models.

Examples of "AI gone rogue":

- Bias in the algorithms that can lead to unintended results.
    - [Facial recognition mistakes](#).
    - [Autonomous vehicles malfunctioning](#).
    - Hal 2000 and agentic AI (AI that can determine what it wants to do—and then do it) deciding to lock us out of our spacecraft-equivalents.

Let's do a bit of daydreaming first:

- What parts of your job are boring or monotonous?
- What parts of your job require sifting through a lot of information scattered in a lot of different places?
- What would you change about these parts of your job, if you could?
- Let's see if we can come up with a few suggestions (raise your hands).

UNLV | WILLIAM S. BOYD
SCHOOL OF LAW

What are the benefits of using AI?

Ability to sift through information quickly (and without the "Maverick" effect"), which I'll explain later.

Tendency to make consistent mistakes (and the ability to learn from them*).

Frees up "human time" to do higher-order tasks.

* Much like smart humans.

UNLV | WILLIAM S. BOYD SCHOOL OF LAW

Let's take the life of a law professor as an example.

Things I like to do:
- Writing.
- Teaching.
- Meeting with students.
- Speaking.
- Talking about ideas with colleagues.
- Consulting and talking with reporters.

Things I do NOT like to do:
- Writing exams.
- Grading.
- Attending faculty meetings.

Attending faculty meetings, 3%

Grading, 5%

Writing exams, 2%

Consulting and talking with reporters, 5%

Talking about ideas with colleagues, 5%

Speaking, 10%

Meeting with students, 10%

Writing, 35%

Teaching, 25%

UNLV | WILLIAM S. BOYD SCHOOL OF LAW

So if only 10% of my work is boring or repetitive, could GenAI make my life better?

- Attending faculty meetings?
  - Nope. I still have to go to them.
- Grading?
  - Sort of. I have discovered the joy of multiple-choice exams (coupled with essays).
- Writing exams?
  - Yes! Claude.ai did the first drafts of all of my multiple-choice questions.
- So my life is now roughly 4.5% better with GenAI than without it.

There are no robo-profs, though—yet.

A few disadvantages:

- AI "hallucinations."
  - It's not just *Mata v. Avianca Inc.* any more.
  - Court rules dealing with the failure to verify.
- Cost (fixed-fee or monthly).
- Training costs.
- Deepfakes.
- Computers, like humans, aren't perfect.
- How we're going to bill, if we base our budgets on the billable hour?

[Computers aren't always accurate](#).

OCTOBER 12, 2017 / **#MACHINE LEARNING**

Chihuahua or muffin? My search for the best computer vision API

by Mariya Yao

WILLIAM S. BOYD
SCHOOL OF LAW

Here's an example from the RAILS guidance*:

5.7 The challenge with this approach lies in the fact that risks often – perhaps even usually – sit within multiple categories in practice.

Example 1 illustrates this:

Example 1

Your company rolls out an AI tool to assist with the screening of employment candidates. The tool unwittingly rejects candidates who are female, ethnically Afro-Caribbean, and over 50. The harms involved are:

    (i)   Human: rejected candidates are subject to unfair discrimination that impacts their ability to obtain employment.

    (ii)   Regulatory: this discrimination breaches various labor and privacy laws in the relevant territory.

    (iii)   Operational: significant cost and management time are incurred in investigating and remedying the situation, and meanwhile, the project that was designed to improve the speed and efficiency of the hiring process is hit by months of delay. The business also suffers a reputational impact.

* RAILS Guidance at p. 8.

WILLIAM S. BOYD
UNLV SCHOOL OF LAW

One way to think about risk (again, from RAILS)*:

How badly could [insert a scary AI risk] hurt my organization?

How quickly could that risk hurt my organization?

UNLV | WILLIAM S. BOYD
SCHOOL OF LAW

So: let's have robots do robot-y things, but what about lawyer-y things?

Or, more broadly, form-filler-y things:

- What can you automate?

- How do you make sure what you've gotten is accurate and useful?

UNLV | WILLIAM S. BOYD SCHOOL OF LAW

Remember when handheld calculators first came out?

WILLIAM S. BOYD
SCHOOL OF LAW

Do you use templates now in some of your work?

- What if:
  - You could input basic client intake information and then have AI sort that information for you into templates that you designed?
  - You could input a client's assets and have an algorithm pull up the jurisdiction's relevant rules?

Some of the risks of AI:

- 🔒 Confidentiality.

- ⚠️ Privacy and pattern recognition.

- 👥 Bias.

- 🧠 Hallucinations.

- 🧍 Bad prompts / failure to ask the right questions.

- 🧑‍💻 Failure to examine the output critically.

More risks:

Supervising those who use AI to make sure that they're using it appropriately.

Disclosure of AI use to clients.

The nightmare: deepfakes.

The missing "spark" of creativity.

Training the newest tranche of professionals.

Let's focus specifically on legal ethics for a bit.

Some ethics rules that intersect with the use of AI.

- Rule 1.1 (Competence).
- Rule 1.2 (Scope of Representation).
- Rule 1.3 (Diligence).
- Rule 1.4 (Communications).
- Rule 1.5 (Reasonable Fees).

Some ethics rules that intersect with the use of AI.

- Rule 1.6 (Confidentiality).
- Rule 3.3 (Candor to the Tribunal) and Rule 4.1 (Truthfulness in Statements to Others).
- Rule 5.1 (Supervising Lawyers) and Rule 5.3 (Supervising Non-Lawyers).
- Rule 5.5 (UPL).

Rule 1.1 (Competence).

A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.

Rule 1.1 (Competence).

[8] To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology, engage in continuing study and education and comply with all continuing legal education requirements to which the lawyer is subject.

Some things to consider:

- Have you scrubbed your metadata?

- Have you used public AI and let it ingest confidential client information?

- Have you checked your AI's output against reality?

- Have you asked your AI to consider the other side's arguments?
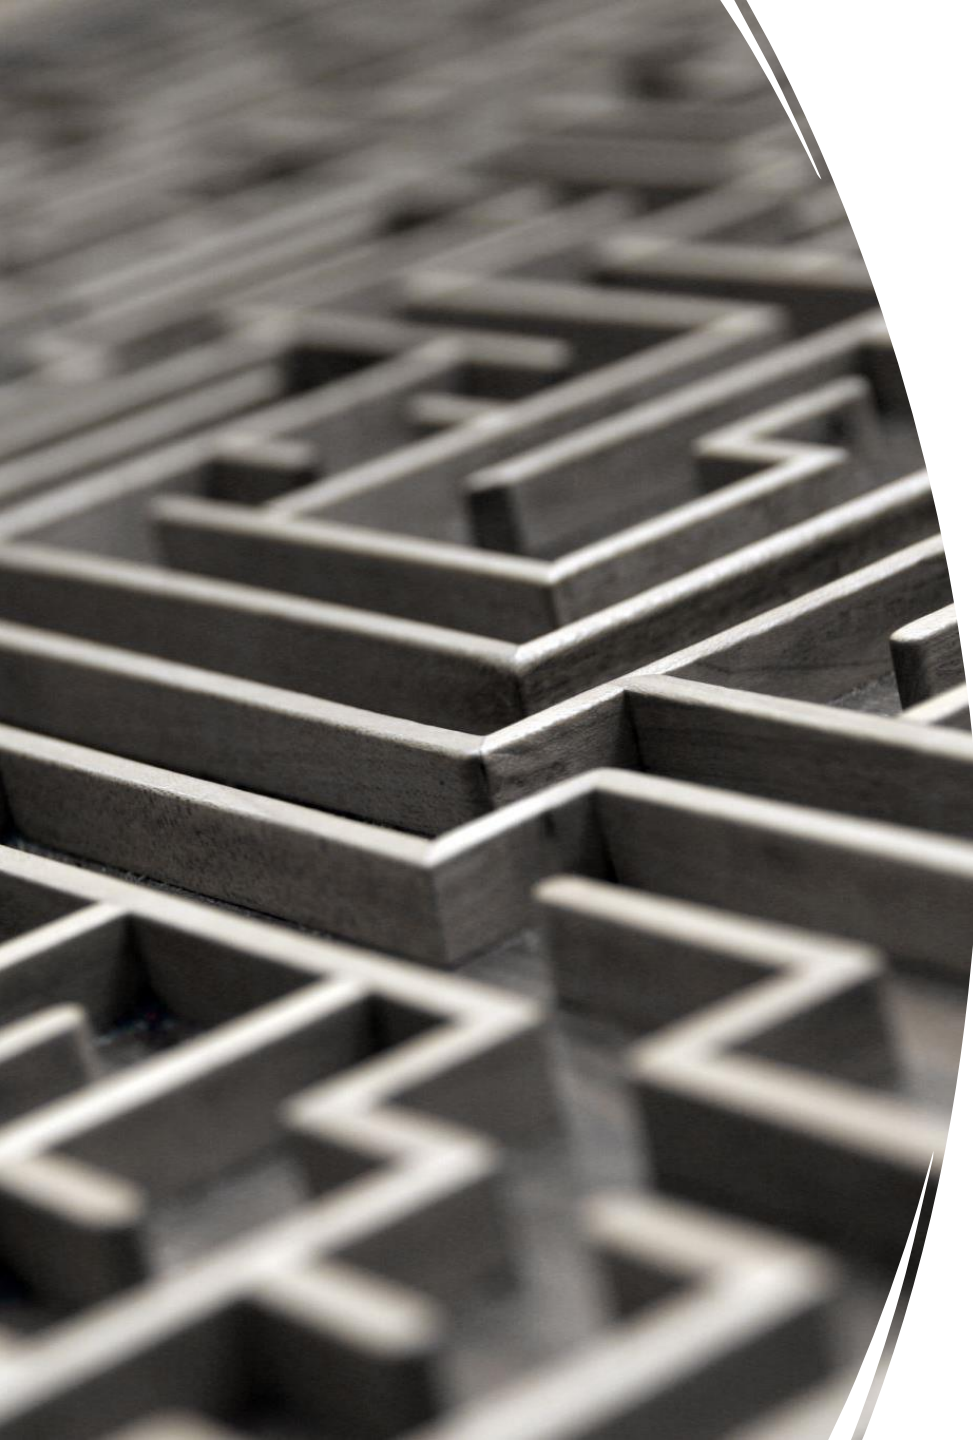
Rule 1.2 (Scope of Representation).

(a) Subject to paragraphs (c) and (d), a lawyer shall abide by a client's decisions concerning the objectives of representation and, as required by Rule 1.4, shall consult with the client as to the means by which they are to be pursued. A lawyer may take such action on behalf of the client as is impliedly authorized to carry out the representation....

Who gets to decide if you use AI, and when?

[Rule 1.3](#) (Diligence).

A lawyer shall act with reasonable diligence and promptness in representing a client.

Are you spending your time efficiently?

Is the right level of professional performing the right task for the right amount of time?

Rule 1.4
(Communications).

(a) A lawyer shall:
    (1) promptly inform the client of any decision or circumstance with respect to which the client's informed consent ... is required by these Rules; [and]
    (2) reasonably consult with the client about the means by which the client's objectives are to be accomplished....
(b) A lawyer shall explain a matter to the extent reasonably necessary to permit the client to make informed decisions regarding the representation.

Be transparent with your clients (and think about getting consent to the use of AI).

Rule 1.5 (Reasonable Fees).

(a) A lawyer shall not make an agreement for, charge, or collect an unreasonable fee or an unreasonable amount for expenses....

Can you charge the client for how long it *used* to take you to do something?

Here's our take on that question:

- Nancy B. Rapoport & Joseph R. Tiano, Jr., *Reimagining "Reasonableness" Under Section 330(a) in a World of Technology, Data, and Artificial Intelligence,* 97 AM. BANKR. L.J. 254 (2023), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4515196.

[ABA Formal Op. 512 (July 29, 2024)](#):

- Learn how a particular GenAI product works before using it;
- Check any GenAI output for errors;
- Use the lawyer's own judgment instead of relying solely on any suggestions generated by GenAI;
- Avoid the risks of impermissible disclosures of confidential information;
- Determine when a lawyer's use of GenAI should be disclosed to the client;
- Avoid the use of false representations to the court that were generated by GenAI;
- Supervise any GenAI use properly; and
- Don't inflate a bill based on what the work *would* have cost had it been done manually.

Then that opinion goes off the rails.

"The factors set forth in Rule 1.5(a) also apply when evaluating the reasonableness of charges for GAI tools when the lawyer and client agree on a flat or contingent fee. For example, if using a GAI tool enables a lawyer to complete tasks much more quickly than without the tool, it may be unreasonable under Rule 1.5 for the lawyer to charge the same flat fee when using the GAI tool as when not using it. 'A fee charged for which little or no work was performed is an unreasonable fee.'"

Then that opinion goes off the rails.

"The factors set forth in Rule 1.5(a) also apply when evaluating the reasonableness of charges for GAI tools when the lawyer and client agree on a flat or contingent fee. For example, if using a GAI tool enables a lawyer to complete tasks much more quickly than without the tool, it may be unreasonable under Rule 1.5 for the lawyer to charge the same flat fee when using the GAI tool as when not using it. 'A fee charged for which little or no work was performed is an unreasonable fee.'"

NOT SO FAST . . .

Before anyone makes a categorical statement like that, let's think about what "value" means:

• As Joe Tiano and I point out:

That passage emphasizes the reasonable fee analysis only from a time "input" perspective, as opposed to a value output perspective (which is really what motivates a client to pay a lawyer). It also ignores the reality that flat fees already incorporate the recovery of overhead charges, from keeping the lights on to paying for access to legal research databases If you have a sneaking suspicion that we're suggesting that lawyers might develop a list of "book rate" / flat-fee tasks, in much the same way that our car mechanics charge us for certain types of work (not by time but by predetermined fee), you're right.

[Rule 1.6](#) (Confidentiality).

(a) A lawyer shall not reveal information relating to the representation of a client unless the client gives informed consent, the disclosure is impliedly authorized in order to carry out the representation or the disclosure is permitted by paragraph (b).

(b) [none of the exceptions apply here].

(c) A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.

Just as you wouldn't save your personal banking information on a public computer….

[Rule 2.1](#) and the need for professional judgment:

- In representing a client, a lawyer shall exercise independent professional judgment and render candid advice....

UNLV | **WILLIAM S. BOYD** **SCHOOL OF LAW**

Rule 3.3 (Candor to the Tribunal).

(a) A lawyer shall not knowingly:
  (1) make a false statement of fact or law to a tribunal or fail to correct a false statement of material fact or law previously made to the tribunal by the lawyer; ... or
  (3) offer evidence that the lawyer knows to be false. If a lawyer, the lawyer's client, or a witness called by the lawyer, has offered material evidence and the lawyer comes to know of its falsity, the lawyer shall take reasonable remedial measures, including, if necessary, disclosure to the tribunal.

Rule 4.1 (Truthfulness in Statements to Others).

In the course of representing a client a lawyer shall not knowingly:

(a) make a false statement of material fact or law to a third person; or

(b) fail to disclose a material fact to a third person when disclosure is necessary to avoid assisting a criminal or fraudulent act by a client, unless disclosure is prohibited by Rule 1.6.

Rule 5.1 (Supervising Lawyers).

Basically, law firms should have guardrails in place to keep their colleague lawyers within ethics guardrails, both from a policy standpoint and a supervisory standpoint, and should fix material problems as soon as they arise.

Would you give your client something done by an entry-level person without first reviewing it?

[Rule 5.3](#) (Supervising Non-Lawyers).

Same principles as in Rule 5.1, but this rule focuses on non-lawyers.

GenAI doesn't know what it doesn't know.

Rule 5.5 (UPL).

(a) A lawyer shall not practice law in a jurisdiction in violation of the regulation of the legal profession in that jurisdiction, or assist another in doing so....

And a little bit about Upsolve.

Upsolve:

To learn more about Upsolve:

- https://upsolve.org/learn/transparency/.

- There have been legal challenges, and it's not meant for most cases (think: "no-asset 7s"), but here's a video about it: https://youtu.be/4FCf9FC_1nI.

UNLV | WILLIAM S. BOYD SCHOOL OF LAW

What about all of the new court regulations about the use of AI?

[FRCP 11. Signing Pleadings, Motions, and Other Papers; Representations to the Court; Sanctions](#).

(b) REPRESENTATIONS TO THE COURT. By presenting to the court a pleading, written motion, or other paper—whether by signing, filing, submitting, or later advocating it—an attorney or unrepresented party certifies that to the best of the person's knowledge, information, and belief, formed after an inquiry reasonable under the circumstances:

....

(2) the claims, defenses, and other legal contentions are warranted by existing law

....

How are firms that are used to billing by the hour going to bill their clients in the future?

Nancy B. Rapoport & Joseph R. Tiano, Jr., *Fighting the Hypothetical: Why Law Firms Should Rethink the Billable Hour in the Generative AI Era,* 20 Wash. J. L. Tech. & Arts 41 (2025), available at https://papers.ssrn.com/abstract=5080449.

Let's think about "value" in new ways:

Assume that the gross benefit to the client of achieving an objective is $X.

The cost of doing that work oneself is $Y.

The risk of doing the legal work incorrectly and having to get the mistakes fixed by a lawyer is $Z.

Won't clients start asking themselves whether it's worth it to start with a lawyer, or start with GenAI and then have a lawyer review the output?

The Jevons effect:

When technology makes certain work more efficient, which reduces the cost of that work, but then the reduced cost of that work actually increases the demand for it.

← William Stanley Jevons

UNLV | WILLIAM S. BOYD
SCHOOL OF LAW

But there's more!

- Shouldn't organizations have AI policies:
  - For themselves?
  - For their vendors?
- For those of you who serve on boards, have you developed AI policies?

And don't forget:

- How will we train our newest professionals to use GenAI well and use it responsibly?

- And how do we make sure that what they're giving us is correct, in the meantime?

CHARLES DICKENS, A TALE OF TWO CITIES 7 (Union Square & Co. ed. 2022):

It was the best of times, it was the worst of times, it was the age of wisdom, it was the age of foolishness, it was the epoch of belief, it was the epoch of incredulity, it was the season of Light, it was the season of Darkness, it was the spring of hope, it was the winter of despair, we had everything before us, we had nothing before us, we were all going direct to Heaven, we were all going direct the other way—in short, the period was so far like the present period, that some of its noisiest authorities insisted on its being received, for good or for evil, in the superlative degree of comparison only.

Remember:

At one point, the Internet was new, too.

Some additional resources:

- https://www.wordrake.com/generative-ai-and-legal-ethics-white-paper.
- https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5080449.
- https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4515196.