Artificial Intelligence, Real Ethics

By Professor Roy D. Simon

ROY D. SIMON is a Distinguished Professor of Legal Ethics Emeritus at Hofstra University School of Law, where he taught from 1992 until 2010, and he is the principal author of SIMON'S NEW YORK RULES OF PROFESSIONAL CONDUCT ANNOTATED (published by Thomson Reuters. He is a legal ethics advisor to law firms and serves as an expert witness in cases where the professional conduct of lawyers or law firms is at issue.

At the dawn of the computer age, IBM's slogan was, "Machines should work; people should think." Artificial intelligence (often called "AI") is turning that slogan into a practical reality. Increasingly, artificial intelligence is enabling machines to do the legal work you used to do, leaving you more time to think, to plan, to strategize. But what about legal ethics? What ethical issues does artificial intelligence raise for lawyers, and how should you respond to those ethical issues?

Let's make the concept of artificial intelligence more concrete. Let's think of artificial intelligence as an amazing bionic legal intern who can do flawless work in a fraction of the time, and at a fraction of the cost, that it would take you to do the same work. That used to be science fiction, but now it's not. The future is already here, and new inventions are coming at us at warp speed. (Other writers in this issue delve more deeply into the meaning and realities of "artificial intelligence," but for purposes of this article, a bionic legal intern will suffice as a tangible manifestation of the idea.)

How are you going to control this bionic legal intern? What are your ethical duties and choices?

Your First Duty: Be Competent

As a lawyer, your first duty is to be competent. The first substantive rule in the New York Rules of Professional Conduct is Rule 1.1, entitled "Competence." The first sentence of Rule 1.1(a) says: "A lawyer should provide competent representation to a client." The next sentence in Rule 1.1(a) explains that "[c]ompetent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation."

Rule 1.1 does not literally *mandate* competent representation because it says a lawyer "should" provide competent representation, not "shall" – but Rule 1.1(b) puts teeth in the duty of competence by essentially prohibiting incompetent representation. Rule 1.1(b) provides: "A lawyer ***shall not*** handle a legal matter that the lawyer knows or should know

that the lawyer is not competent to handle, without associating with a lawyer who is competent to handle it." (Emphasis added.)

What does Rule 1.1 mean to a practicing lawyer? It means three things.

First, you must endeavor to acquire the "legal knowledge and skill" you need to do a good job on the matters you handle. And since the technology that you use to handle legal matters keeps changing, you have to keep up with the changes. This idea is expressed in Comment [8] to Rule 1.1, headed "Maintaining Competence," which says:

> [8] To maintain the requisite knowledge and skill, a lawyer should (i) keep abreast of changes in substantive and procedural law relevant to the lawyer's practice, [and] ***(ii) keep abreast of the benefits and risks associated with technology the lawyer uses to provide services to clients or to store or transmit confidential information*** .... [Emphasis added.]

You therefore must stay current with the "benefits and risks" of any AI software and services that you use. What does that mean? Do you have to become a computer science expert? No. You don't have to understand the code or algorithms behind the AI product. (Even many of the experts who develop these products don't fully understand them – *see Cliff Kuang, Can A.I. Be Taught to Explain Itself?* (N.Y. Times Magazine, Nov. 21, 2017) ("As machine learning becomes more powerful, the field's researchers increasingly find themselves unable to account for what their algorithms know — or how they know it"). But as a lawyer, you do have to understand what an AI product can and cannot do, and you have to evaluate whether it is performing as advertised. (I will discuss these ideas in more detail below.)

Second, Rule 1.1 requires you to exercise "thoroughness and preparation" commensurate with the tasks at hand. If you are using technology to conduct legal research, to communicate with clients, to file court papers, or to perform other legal tasks, then you should learn to use that technology competently. And if courts or clients or co-counsel or opposing counsel are using a given technology – such as e-discovery, electronic filing, email, or PDFs – then you need to keep abreast of the "benefits and risks" associated with that technology.

Third, if a new matter comes along that you are not technologically competent to handle, then under Rule 1.1(b) you have three choices: (a) turn down the matter; or (b) spend whatever time it takes to acquire the necessary legal knowledge and technological skill; or (c) associate with a different lawyer – whether in your own firm or (with your client's permission) in a different firm – who already has the necessary technological knowledge and skill.

If you need to learn a skill that involves technology – social media, ECF, PACER, metadata, spam filters, whatever – you can consult with a nonlawyer who can teach you what you

need to know, or you can delegate to a nonlawyer who already has that skill, as long as you supervise the nonlawyer in compliance with the Rules of Professional Conduct. And that brings us to the next topic: your duty to supervise nonlawyers.

Rule 5.3: Your Duty to Supervise Nonlawyers

Under Rule 5.3, entitled "Lawyer's Responsibility for Conduct of Nonlawyers," you have a duty to "ensure that the work of nonlawyers who work for the firm is ***adequately supervised, as appropriate***." (Emphasis added.) Artificial intelligence products are effectively non-human nonlawyers. Let's return to the analogy of our fast and flawless bionic legal intern. In my view, supervising a bionic legal intern – the software equivalent of an artificially intelligent robot lawyer – is equivalent to supervising a human legal intern. What does it mean to ensure that the bionic legal intern – the artificial intelligence product – is "adequately supervised, as appropriate"? I have three concrete suggestions: (1) hire an expert to vet the AI product; (2) learn what the AI product can (and can't) do; and (3) double check the output of the AI product. I'll discuss each suggestion.

1. Hire an expert to check out the AI product. The analogy to the human legal intern is helpful. You wouldn't hire a new legal intern without a background check, and you shouldn't use a new artificial intelligence product without a background check. Before you buy or use an AI product in your law practice, you should ask a lot of questions. Who developed the AI product? What is the developer's reputation? Is the product compatible with the other software you already use? Is the product free of malware that could steal your confidential information, or does it contain vulnerabilities that could open a back door to let cyberattackers into your system? These questions are not merely paranoia – *see* Shane Harris & Gordon Lubold, *Russia Turned Kaspersky Software Into Global Spying Tool* (Wall St. J. Oct. 11, 2017) (reporting that the Russian government modified a popular antivirus software program to turn it into an espionage tool that secretly scanned computers for classified U.S. government documents).

But there's a big obstacle to checking out an AI product: you are a lawyer, not a software engineer. Most lawyers have no idea how to find the answers to questions about AI products or developers. So you should do what you do in the rest of your practice when you encounter an area beyond your expertise – hire an expert. When you need a complex damages calculation or an environmental quality test, you hire an expert. When you have a technology issue, you should hire an IT consultant. Your firm should designate a skilled technology person or company – inside or outside your firm – to vet every new artificial intelligence product (and every new software product) that you use in your law practice.

2. Learn what the AI product can do – and what it can't do. AI products can do amazing things, but they also have limitations, and you need to know what they are. If a legal intern spends most of her time preparing UCC forms, you wouldn't rely on her to write a

complaint in an antitrust case, and you shouldn't count on an AI product to do things it cannot do. For example, suppose you have written a memorandum of law (for a tribunal or a transaction), and you want to know, this afternoon, whether the authorities you cited are still good law and whether you missed any authorities you should have cited. Artificial intelligence can perform this task. Several AI products will not only cite check an entire memo but also will suggest additional cases. (The products include BriefCheck by Shepard's, WestCheck by Westlaw, CARA by CaseText, and EVA by ROSS. They all work a bit differently. You should try them all to see how they work and which ones you like best.) Within a minute, the AI software will generate a report. But what does the report cover? Does the product check statutes and regulations to see if they have been amended? Will the report tell you whether the quotations in your memo are 100% accurate? Will it suggest additional cases based only on the cases you have cited, or will it actually examine the fact pattern in your memo to look for issues you missed? You don't have to understand how to use every feature of an AI product (just as you don't have to understand every feature of Word or Excel). But you do have to understand what the AI product can and cannot do, and you need to learn how to deploy the product features that you actually use. AI products, like legal interns, get results only if you give the right instructions. I learned this concept the hard way, long before AI. When I was a brand new attorney at a large firm, I drafted a complaint and sent a messenger down to the courthouse to file it. An hour later the messenger came back with the complaint, still unfiled. The clerk had refused to accept the complaint because it did not have a blueback. I grumpily added the blueback and sent the messenger to court again. Then I complained to a colleague that the court clerks were hung up on petty details. My colleague said, "That may be true, but you should never send someone to do something you haven't done yourself."

You don't have to master AI products, but you should try them out to get a feel for how they work. Ask an IT person at your firm to tutor you for an hour or two, or search for YouTube videos about how to use each product. Then you will be in a position to know when a particular AI product is appropriate, and you can instruct your subordinates on what to do with each product.

3. Double check the AI product's output. When a legal intern drafts a brief, you read it over for quality control. You have to do the same thing with the work product of an AI tool. For example, some AI software products will actually draft substantive memos in certain subject areas, or review contracts, or fill out forms to fight parking tickets. These whiz-bang tools may produce great work, but you are still responsible for the work product. If the work product is defective and the client is harmed, the client is going to sue you and your law firm, not the artificial intelligence product.

You don't have to duplicate the entire task that the AI product performed, just as you don't start the research and drafting from square one when a legal intern submits a draft brief. But you do have to read the entire paper to make sure it is relevant, organized, clear, and

not contrary to common sense. You also should spot check the case citations and quotations, at least until you develop confidence in the particular AI product.

Rule 1.6(c): Your Duty to Protect Confidential Information

All lawyers know that they have a duty of confidentiality, but not all lawyers realize that Rule 1.6(c) of the New York Rules of Professional Conduct was amended effective January 1, 2017 (just over a year ago). Amended Rule 1.6(c) requires lawyers to make "reasonable efforts" to protect confidential information against three things: (1) inadvertent disclosure or use; (2) unauthorized disclosure or use; and (3) unauthorized access (*e.g.,* hacking from outside or inside your firm).

Most AI products, such as the cite-checking products described above, require access to your confidential data. (A draft memo itself is confidential information, for example.) This raises a lot of questions about confidentiality. What happens to your confidential data once the AI vendor gains access to it? Who has access to it at the AI vendor? Does the AI vendor share your confidential information with other third-party vendors? If so, do you know who those third-party vendors are, and have you checked them out? Do they have a contractual duty of confidentiality? What happens to your client's data if the AI vendor is sold, merges, retires, or goes bankrupt? If the AI vendor is subpoenaed, is the vendor contractually obligated to give you notice so that you can intervene to challenge the subpoena?

These questions are just the tip of the iceberg regarding confidentiality. Excellent guidance on transmitting or storing confidential information is available in N.Y. State Ethics Op. 1020 (2014), which concerned the ethics of cloud storage. Opinion 1020 concluded that whether a lawyer may post and share documents using a cloud data storage tool depends on whether the technology employed "provides reasonable protection to confidential client information and, if not, whether the lawyer obtains informed consent from the client after advising the client of the relevant risks." The same principles should apply when you give an AI vendor access to your confidential data.

Rule 1.5(a): Fee and Expense Issues

Rule 1.5(a) prohibits a lawyer from charging fees and expenses that are "excessive" (*i.e.,* are not "reasonable"). This raises at least two issues relevant to artificial intelligence.

The first question is, may you charge your clients when you use artificial intelligence products the way you charge clients for Westlaw or Lexis? In my opinion, the answer is yes as long as either (a) you charge your clients at your out-of-pocket cost (*e.g.,* a fair share of the cost of a license, or the cost per use), or (b) you obtain consent from your clients to charge a reasonable markup – *see* ABA Ethics Op. 93-393 (1993) ("Any reasonable calculation of direct costs as well as any reasonable allocation of related overhead should pass ethical muster" but, *absent the client's agreement*, a lawyer may not "create an

additional source of profit for the law firm" by charging above cost for computer research services or other nonlegal services).

Second, are you charging an excessive fee if you perform hourly rate legal work the old fashioned way instead of using an AI product that could do the job faster? That is a tough question, and depends on the circumstances.

An analogy might help. If you charge by the hour, are you charging an excessive fee if you do all of your legal research by reading bound copies of reporters instead of by using Westlaw, Lexis, or Fastcase? Three or four decades ago the answer was probably no, but today I think the answer is yes – doing all of your legal research in bound books results in an excessive fee because you can KeyCite or Shepardize cases in a fraction of the time it takes to slog through paper volumes. Similarly, today I think you are not charging an excessive fee if you continue using your customary methods instead of using a newfangled AI product – but soon, most lawyers will be using AI products and services for certain types of work (such as the cite-checking products discussed earlier), and charging for ten hours of your time to do work that AI could do in ten minutes sounds like an excessive fee to me. You have to keep abreast of the benefits of technology that applies to your practice.

Now for a bonus question: Do you have a duty to alert your clients to the option of using AI products that may save substantial fees or arrive at quicker or more accurate results? Right now the answer to that question is unclear – but before long, practicing law without using AI will be like practicing law with an Underwood manual typewriter, and you will have to tell your clients that there is a better, cheaper, faster way.

Conclusion: Technology Moves Ahead But Ethical Duties Remain Constant

Artificial intelligence is making it possible to create products that sound like magic – lightning fast, uncannily accurate, effortless to use. These new products raise new questions about legal ethics. But the ethical principles remain familiar – lawyers must be competent, supervise the work product, protect confidential information, and charge reasonable fees and expenses. If you keep those timeless ethical principles in mind, you can keep your ethical balance even as technology moves ahead at a dizzying pace.