

Software Malpractice in the Age of AI: A Guide for the Wary Tech Company

Danny Tobey, M.D., J.D.

Harvard College '99; Yale Law School '03; UT Southwestern Medical School '08; Vinson & Elkins, Dallas, TX, USA

Abstract

Professional malpractice—the concept of heightened duties for those entrusted with special knowledge and crucial tasks—is rooted in history. And yet, since the dawn of the computer age, courts in the United States have almost universally rejected a theory of software malpractice, declining to hold software engineers to the same professional standards as doctors, lawyers, and engineers. What is changing, however, is the speed at which software based on artificial intelligence technologies is replacing the very professionals already subject to professional liability. Society has already decided (in some cases, millennia ago) that those tasks warrant special accountability; new to the analysis is which human is closest in line to the adverse event. As AI expands, the pressure for courts to go one level up the causal chain in search of human agency and professional accountability will mount. This essay analyzes the case law rejecting software malpractice for clues about where the doctrine might go in the age of AI, then discusses what technology companies can learn from the safety enhancements of doctors, lawyers, and other historic professionals who have adapted to such heightened legal scrutiny for years.

Introduction

Since 1989, courts in the United States have almost universally declined to hold software engineers to the same professional standards as doctors, lawyers, architects, and engineers. But in the age of AI, as software *replaces* the doctor, lawyer, architect, and engineer, will courts finally take the bait and establish a cause of action for software malpractice? And what can the far-sighted technology company do now to anticipate and adapt to possibly higher legal standards?

Professional malpractice—the concept of heightened duties for those entrusted with special knowledge and crucial tasks—is rooted in history. Around 2000 B.C., the Code of Hammurabi held: “If the doctor has treated a gentlemen with a lancet of bronze and has caused the gentleman to die, or has opened an abscess of the eye for a gentleman with a bronze lancet, and has caused the loss of the gentleman’s

eye, one shall cut off his hands” (Bal 2008). Closer to 2000 A.D., today’s software companies should be asking: When the doctor is code, whose hands will be on the block?

The concept of software malpractice is not new. The Eighth Circuit recognized such a cause of action in 1989 in *Diversified Graphics, Ltd. v. Groves*, 868 F.2d 293 (8th Cir. 1989). But courts since have almost universally declined to adopt that holding or endorse such claims. What is changing, however, is the speed at which software based on artificial intelligence technologies is assuming the very tasks traditionally handled by professionals already subject to professional liability. Society has already decided (in some cases, millennia ago) that those tasks warrant special accountability; what is new is which human is closest in line to the event.

As AI expands, the pressure for courts to go one level up the causal chain in search of human agency and professional accountability will mount. Indeed, the very cases rejecting *Diversified Graphics* hint at the doctrinal path to software malpractice. Technology companies should anticipate this evolution and learn from the safety enhancements of doctors, lawyers, and other historic professionals who have adapted to such heightened legal scrutiny for years—that is, while doctors and lawyers exist.

Medicine: A Case Study in Disruption

To trace the evolution of AI into the professional human space, take the example of medicine—a quintessential profession subject to fiduciary duties and malpractice claims. Doctors face special scrutiny because they are entrusted with important, often life-altering decisions, requiring special expertise to address: predicting a heart attack before it is too late; diagnosing an illness from a constellation of symptoms meaningless to the lay observer. Medicine is the archetypal “art not science”—as years of medical students

have been told, “patients don’t read textbooks,” and diagnosis requires not just seeing patterns but knowing when to disregard them, employing skills of psychology and intuition to sift through unreliable patient narratives and clinical tests with imperfect sensitivity and specificity.

As a result, doctors are exposed to—and in some ways, protected by—higher standards. They are judged against other doctors, not the ordinary person’s view of what may or may not be reasonable. Special duties have governed physicians for millennia, through Roman law and the English common law to the modern tort and regulatory systems of developed nations (Bal 2008).

Modern legal systems have generally assumed that human-monitored technology is *safer* than technology alone. This assumption drives the “informed intermediary” and “learned professional” doctrines, which can break the chain of causation leading back to manufacturers in product-liability cases when a skilled professional stands between them and the end-user. *See, e.g., Figueroa v. Boston Sci. Corp.*, 254 F. Supp. 2d 361, 370 (S.D.N.Y. 2003). That same assumption drove Congress in 2016 to draw a bright line in the 21st Century Cures Act between medical software that acts alone and medical software where humans “independently review” the basis for a software’s clinical recommendations. *See* FDCA § 520(o)(1)(E). Human review exempts the software from FDA regulation, presumably because doctors monitor and, as needed, override the software’s recommendations. Or take the longstanding duties of hospitals and doctors to provide adequate medical equipment—the implicit standard of care being that doctors and hospitals monitor machines and not vice-versa.

And then comes the disruption: By some estimates, physicians are able to predict heart attacks 30% of the time, while AI systems are already predicting them with 80% success—and far sooner than their human counterparts (Economist 2016). In 2016, doctors misdiagnosed a 60-year-old Japanese woman with a rare form of leukemia for months, subjecting her to treatments to no avail; IBM’s Watson reviewed her data against 20 million cancer research papers and correctly diagnosed her in ten minutes (Feldman 2016). In 2017, the journal *Nature* reported a study testing a deep neural network against 21 board-certified dermatologists on biopsy-proven clinical images, “demonstrating an artificial intelligence capable of classifying skin cancer with a level of competence comparable to dermatologists” (Esteva et al. 2017). Fifty percent of hospitals plan to adopt some form of AI within the next 5 years (Sullivan 2017).

For now, the assumption of human-monitored technology as safer may still hold. “Human-in-the-loop” AI is still superior to machine-only outcomes in some contexts (Nushi et al. 2017; Russakovsky, Li, and Fei-Fei 2015). But for how long? Early AI, in the form of expert systems and other rule-based approaches, built machine “doctors” by modeling human doctors, attempting to convert medical reasoning into a

series of if-then statements. That approach was fundamentally limited: As an approximation of human decision-making, AI could at best mimic optimal human performance. New approaches, of course, depart from mimicry to true machine learning. In October of 2017—a symbolic moment in AI history—Google’s DeepMind AI beat the human world champion at the game Go, again, but this time without human input or training. Past versions had won by analyzing more than 100,000 human games. The new version taught itself to play, knowing only the rules and objective of the game, without human input. After three days of self-learning, the human-less version beat the human-informed version 100 times out of 100 (Vincent 2017).

The time will come (likely soon) when AI consistently outperforms both humans alone *and* humans-in-the-loop AI, mining datasets larger than any human can fathom, refining black-box algorithms no human can reverse-engineer or second-guess, and finding connections among data no human can compute—or even comprehend in hindsight. At that point, what does a human monitor of technology have to offer? After all, the value of AI is in seeing things humans cannot. As deep AI surpasses and supplants human intuition, the proper recommendations will *by definition* be counterintuitive at times. Even today, it is not always clear whether a doctor should be faulted for following a counterintuitive machine recommendation or for disregarding it. Is the computer wrong, or is it just seeing farther?

The same trends are occurring in law, engineering, accounting, and other professions subject to heightened scrutiny and legal responsibility. Already, legal software is beginning to handle even the “conceptual” tasks traditionally performed by lawyers: drafting pleadings, analyzing precedents, reviewing documents, and predicting litigation outcomes (Hutson 2017; Wittenberg 2017). So too in accounting (audits, fraud detection, risk assessment), finance (investment advising, lending decisions), and other fields. A judge recently granted bail to a defendant based on the recommendation of a risk-assessment algorithm, only for that defendant to commit murder upon release. As a deputy district attorney observed: “It’s very hard for a judge to go against this type of risk assessment program because it’s couched in science.” And so, “if there’s an algorithm that says ‘keep them out of custody’ even if their instinct and the record say otherwise, they’re going to follow what the algorithm says” (Westervelt 2017). As AI improves, humans will have even less freedom to depart from AI’s instructions.

Some believe that AI will not supplant professional human utility; it will merely prioritize for people certain “human” skills like “judgment,” while ceding more mechanical skills like “prediction” to machines. So argued one set of authors in 2016 in the *Harvard Business Review* (Agrawal, Gans, and Goldfarb 2016). Unfortunately for humans—and highlighting the uncertainty on point—another set of authors writing in the same journal just a month apart reached

the opposite conclusion—arguing that there is little difference in reality between the so-called “routine work” that machines can do and the allegedly “tricky stuff that calls for judgment, creativity, and empathy” (Susskind and Susskind 2016). Rather than presenting a difference in kind, such tasks may be just a slightly more complex problem subject to computation. With AI already—if clumsily at first—composing music (Grimstad 2017) and deconstructing what makes a bestseller (Tolentino 2016), this second set of 2016 Harvard Business Review authors looks to have the better odds. They predict the end of human professions.

What happens, then, when the law’s traditional assumption of favoring humans or human-software collaboration over software alone is upended, and the standard of care places software above human intervention? The basis for professional responsibility will not have disappeared at that point. Humans will still receive high-risk, high-impact services. As one court said, “robots cannot be sued,”—still true for now—“but they can cause devastating damage,” triggering a search for “the ultimate responsible distributor. . . .” *United States v. Athlone Indus., Inc.*, 746 F.2d 977, 979 (3d Cir. 1984). In the context of AI, that means the programmer may become the last human standing in the chain of causation. As software supplants professions, software malpractice will likely supplant professional malpractice. Even the decades of case law *rejecting* software-malpractice claims hint at the fault lines to come.

The Law Against Software Malpractice Contains the Seeds of its Own Reversal

Commentators have long observed that traditional product-based tort and contract remedies are ill-suited to software (Goertzel 2016). Software straddles the line between product and service. It is increasingly ubiquitous, unavoidable, and indispensable. EULAs, clickwrap, and the like are mandatory, accepted without inspection, yet enforced as contractual. The economic loss rule further inhibits many tort claims against software providers. In other fields, professional negligence claims have circumvented these barriers to liability, providing an easier path for consumers to seek relief from providers. Not so in the technology space, yet. Technology companies have long relied on these barriers to liability, and consumers have vigorously litigated against them, including by proposing software malpractice claims. And yet, since the dawn of the computer age, courts have almost universally declined to hold software engineers to the same higher standards as other professions.

The leading exception is *Diversified Graphics, Ltd. v. Groves*, 868 F.2d 293 (8th Cir. 1989). In *Diversified*, a United States Court of Appeals—the highest court to consider the issue so far—allowed a cause of action for computer malpractice to lie, holding that a “computer systems

consultant” was “properly held to a professional standard of care” in light of its “superior knowledge and expertise in the area of computer systems.” *Id.* at 296. In 1986, an Indiana state court likewise held that developing software for a specific client was “more analogous to a client seeking a lawyer’s advice or a patient seeking medical treatment for a particular ailment than it is to a customer buying seed corn, soap, or cam shafts” and imposed a higher legal standard. *Data Processing Servs., Inc. v. L.H. Smith Oil Corp.*, 492 N.E.2d 314, 319 (Ind. Ct. App. 1986), overruled on other grounds by *Insul-Mark Midwest, Inc. v. Modern Materials, Inc.*, 612 N.E.2d 550 (Ind. 1993). In 2002, a Delaware superior court found that such a claim would lie under Tennessee law. *See Bridgestone/Firestone, Inc. v. Cap Gemini Am., Inc.*, No. CIV.A. 00C-10-058HDR, 2002 WL 1042089, at *4 (Del. Super. Ct. May 23, 2002).

Despite these exceptions, the vast majority of courts to consider the issue have declined to impose professional duties on computer programmers and software companies. As one federal district court recently observed, noting decades of holdings: “Of the courts to consider the question, the overwhelming majority have determined that a malpractice or professional negligence claim does not lie against computer consultants or programmers.” *Superior Edge, Inc. v. Monsanto Co.*, 44 F. Supp. 3d 890, 912 (D. Minn. 2014) (collecting cases); *see also, Avazpour Networking Servs., Inc. v. Falconstor Software, Inc.*, 937 F. Supp. 2d 355, 364 (E.D.N.Y. 2013); *Columbus McKinnon Corp. v. China Semiconductor Co., Ltd.*, 867 F. Supp. 1173, 1182-83 (W.D.N.Y. 1994); *Hosp. Comput. Sys., Inc. v. Staten Island Hosp.*, 788 F. Supp. 1351, 1361 (D.N.J. 1992); *Triangle Underwriters, Inc. v. Honeywell, Inc.*, 604 F.2d 737, 745-46 (2d Cir. 1979); *Arthur D. Little Int’l, Inc. v. Dooyang Corp.*, 928 F. Supp. 1189, 1202-03 (D. Mass. 1996).

But why? And how sturdy are those opinions rejecting such claims? A closer inspection reveals that in the age of software ubiquity, and more specifically AI, the foundations of those contrary opinions may be eroding.

Professional Maturation

Most courts declining to hold software companies to professional duties note the lack of professional “indicia” in the industry. Such indicia include uniform training, self-imposed industry standards, and state licensure and regulation. *See Superior Edge*, 44 F. Supp. 3d at 912 (quoting *Ferris & Salter, P.C.*, 889 F. Supp. 2d at 1152 (quoting Raymond T. Nimmer, *The Law of Computer Tech.* § 9.30 (4th ed. 2012)); *Hosp. Comput. Sys.*, 788 F. Supp. at 1361.

But those professional indicia are increasing in the software industry. In 2012, the Texas Board of Professional Engineers announced the Software Engineering Principles and Practice of Engineering (PE) exam, creating “a path to licensure for practicing software engineers” and marking “a

critical step in the overall philosophy surrounding software engineering and the licensing of software engineers in the United States” (TBPE 2012). The Institute of Electrical and Electronics Engineers (IEEE) has said: “Just as practicing professionals such as doctors, accountants, and nurses are licensed, so should software engineers” (Kowalenko 2012). In September 2017, a coalition of clinical software companies released “Voluntary Industry Guidelines for the Design of Medium Risk Clinical Decision Support Software to Assure the Central Role of Healthcare Professionals in Clinical Decision-Making,” hoping to stave off broader FDA regulation (Miliard 2017). In the age of AI, from the Asilomar Conference to the National Governors Association, calls for greater regulation (both internal and external) of the software industry are growing, as the potential consequences of its products grow.

Professional Migration

On top of this trend of professional maturation, there is the more interesting—and largely unprecedented—question of *professional migration*: that is, thanks to AI, the software industry is not only professionalizing, but its product is rapidly replacing historic professionals themselves, taking on tasks formerly the exclusive domain of human professional analysis, decision-making, and implementation. In short, the software industry is not just professionalizing itself but blurring the lines between software and traditional professions.

The law is not without guidance on this question. And indeed, and perhaps ironically, a recent case *rejecting Diversified* may demonstrate why professional migration opens the door to *embracing Diversified* and imposing professional liability. In *Superior Edge, Inc. v. Monsanto Co.*, a federal district court analyzed *Diversified* and declined to adopt its holding, citing a novel distinction. In *Superior Edge*, Monsanto was urging the court to hold a software company “to a standard of care as a member of a learned and skilled profession [with] a duty to exercise the ordinary and reasonable technical skill that is usually exercised by those in the software development field.” 44 F. Supp. 3d at 911 (quoting Monsanto Countercls. ¶ 161). Monsanto argued: “By failing to timely deliver a fully functional and scalable software product in a timely manner, SEI deviated from the standard of care and that deviation was the proximate cause of Monsanto’s damages.” *Id.*

The *Superior* court rejected the theory, and any application of *Diversified*, noting that the computer consultants in *Diversified* happened to work for an accounting company, and accounting—separate and apart from computer science—was an established profession. *See id.* at 913 (“but that case [*Diversified*] was a professional negligence action against an accounting firm that acted as a consultant for the client’s purchase and implementation of an in-house com-

puterized data processing system. The availability of professional negligence actions against accountants is already well-established under Missouri law, and *Diversified* says nothing about whether a professional negligence action against a computer professional—not acting in a role as an accountant—is appropriate.”).

Arguably, the *Superior* court misinterprets the holding of *Diversified*. The Eighth Circuit did not find the computer consultants there liable because they happened to come from an accounting firm. Quite the contrary, the Eighth Circuit made no mention of the overlap between the computing and accounting functions, and instead specifically imposed liability on the computer consultants *as* computer consultants: “E&W failed to act reasonably in light of its superior knowledge and expertise *in the area of computer systems.*” *Diversified*, 868 F.2d at 296 (emphasis added). The court did look to the American Institute of Certified Public Accountants for evidence of consulting standards, but cited only general principles of professional transparency, and was expressly looking for “the professional standard of care required *of a computer systems consultant,*” not an accountant. *Id.* (emphasis added).

But suppose *Diversified*’s holding *did* rest on the blending of computer consultancy with a traditional profession: In the age of AI, that will increasingly be the case, and the distinction relied on by the *Superior* court may not be long for this earth. With the advent of deep AI, the fact pattern of future lawsuits will increasingly resemble *Diversified* over *Superior*, with software engineers migrating into historic professional decision-making in order to reproduce and supplant it. A software company that develops medical AI will supplement or supplant human medical decision making. Medical professionals will likely be involved in the software company’s process of developing and training that software, if not reviewing its output on an ongoing basis: *i.e.*, the fact pattern will be doctors making software, much as *Superior* characterized *Diversified* as accountants making software. Even if a software company did not employ medical professionals to help make or monitor medical software, the fact-pattern would still resemble *Superior*’s criterion of overlapping novel and traditional professions: if not doctors making software, then software making doctors. It is not that the programmer’s skill set will resemble the doctor’s, but rather that the programmer will employ her own advanced skill set to accomplish the same end as the doctor previously had, an end warranting heightened legal duties. The programmer thus becomes the prime human cause of the same risks and rewards as the erstwhile physician. And where the fruit of programmer’s labors duplicates and replaces historical professional judgment, the pressure for continued professional accountability will mount. And the distinction relied on by *Superior* to avoid such calls will thin.

Precedent and Disruption

Finally, courts are not constrained by the rate of professional maturation or migration. The Restatement of Torts has a built-in flexibility for specialized negligence claims, allowing higher standards for both “services in the practice of a profession *or trade* . . .” Restatement (Second) of Torts § 299A (1965) (emphasis added). In the Restatement’s comments, “trade” is defined more broadly than “profession,” to include “any person who undertakes to render services to others in the practice of a skilled trade, such as that of airplane pilot, precision machinist, electrician, carpenter, blacksmith, or plumber.” *Id.* cmt. b. The cases adopting software malpractice in the ‘80s through the ‘00s, before much if any professional maturation, went this route. *See Data*, 492 N.E.2d at 319 (“Those who hold themselves out to the world as possessing skill and qualifications in their respective *trades or professions*”) (emphasis added); *Diversified*, 868 F.2d at 296 (“Professional persons *and* those engaged in any work or trade requiring special skill”); *Bridgestone/Firestone*, 2002 WL 1042089, at *4.

Most courts have declined to take this shortcut around professional maturation or migration. But the only barrier is precedent. *See, e.g., Superior Edge*, 44 F. Supp. 3d at 913–14 (“But the Court finds that based on Missouri’s *existing law*—which generally limits professional negligence actions to those fields in which participants are regulated by state licensing requirements—the Missouri Supreme Court would follow the majority of courts . . .”) (emphasis added).

Precedent holds until it does not. New facts and circumstances permit the common law to evolve, and those facts and circumstances are in the eye of the beholder; here, courts. Indeed, professionalism is a historical measure; its indicia (regulation, licensure, credentialing) take time to evolve. Trade is an ahistorical measure, focused more on the presence of specialized knowledge and skills than on the process by which they are acquired or regulated. And so where software engineers use specialized knowledge and skill to *reproduce* professional interventions without professional gatekeepers, courts may come to see wisdom in the Restatement’s broader focus on both professions *and* trades.

Indeed, software companies in particular would be wise not to put their eggs in the basket of precedent for planning purposes. Precedent is the rallying cry of the legal industry, but disruption has been the rallying cry of the software industry (or at least its start-ups’ marketing arms). If precedent is continuity, disruption is discontinuity—in algorithmic terms, genetic shift over genetic drift. Of course, the *modus operandi* of one profession does not dictate the *modus operandi* of another. But where the party urging stability and adherence to precedent to courts is the one upending millennia-old professions in the blink of an eye, prudent planners will not bank on that line of reasoning carrying the day.

Looking Ahead

What, then, is the software industry to do, to elevate the bar in advance of potentially higher legal scrutiny?

Safety First

The low bar to implementation of new AI systems, coupled with the race to move first and the lack of higher legal scrutiny, has sometimes led to an ethos of *develop first, think later*. Take the Cambridge University student who created AI that identifies partially-obscured faces in crowds. When asked if the technology might be used by repressive governments to identify and punish dissidents, he conceded: “To be honest when I was trying to come up with this method, I was just trying to focus on criminals . . . I actually don’t have a good answer for how that can be stopped. It has to be regulated somehow . . . it should only be used for people who want to use it for good stuff” (Matsakis 2017). Not every coder need be a political philosopher, but when new technologies wield vast potential for good, ill, amoral accident, and unintended harm, the excuse of not having a good answer—or merely having an aspiration toward “good stuff”—is not enough to forestall epic liability. Every technology contains its opposite, and developers should be thinking from the outset about unintended consequences, protections against defects and misuse, and other legal and safety measures.

Embrace Professionalization

A perverse view of the doctrine around professional indicia would go as follows: If self-regulation brings legal regulation, avoid self-regulation. Hopefully, this essay has dispelled the utility of that notion, by flagging the alternate routes courts have already taken to impose higher standards even when professionalism lags. Moreover, should evidence arise of *intentional* anti-professionalism, courts and lawmakers would feel even more compelled to take alternative routes to accountability. Nor will hiding behind the traditional distinction between product and service help for much longer. Technology spans (and often blurs) the divide between product and service, and the choices companies make—in product design, contracts, marketing, and sales—can affect how courts rule. But even in *Diversified*, the computer consultants were hired to provide “a ‘turnkey’ computer system” that was “self-sufficient” once operational, requiring “only minimal training” to use—and yet professional service liability attached. 868 F.2d at 297. Likewise, a software company named Smartcop tried to avoid a “Professional Liability Exclusion—Computer Software” in its insurance policy by relying on the distinction between the products it made and the services it provided, but found no traction: “Whether characterized as providing defective products or a defective service to the Sheriff’s Office, all of

these claims stem from Smartcop’s duties to sell, license, or furnish its computer software.” *Maryland Cas. Co. v. Smartcop, Inc.*, No. 4:11-cv-10100-KMM, 2012 WL 4344571, at *4 (S.D. Fla. Sept. 21, 2012).

One insight of *Diversified* is often overlooked. Being held to professional standards is not all bad. In some respects, it is protective. As the Eighth Circuit observed: “A breach of a professional standard is more exacting and difficult to prove than breach of ordinary care.” *Diversified*, 868 F.2d at 295–96. Moreover, embracing and accelerating professionalism can mean increased predictability, by defining the rules of the road and the standards by which an industry should be judged, while elevating safety and quality. Professional liability lowers some barriers to consumer redress, but it raises others, in a socially productive way: Professionals cannot contract away all of their duties, but they can sometimes find safe harbor in meeting the standards of their industry, not the standards of lay individuals judging their industry without technical insight.

Learn From Doctors, Lawyers, and Accountants (While They Exist)

“Medical quackery and the promotions of nostrums and worthless drugs were among the most prominent abuses which led to the establishment of formal self-regulation in business and, in turn, to the creation of the NBBB [National Better Business Bureau]” (Ladimer 1965). Doctors, lawyers, and the like have long histories of weathering self-regulation and external regulation by state and federal agencies and by courts and lawmakers. However imperfect, the creation of formal educational programs, licensure exams, industry standards, and continuing education requirements have helped push some harmful conduct to the margins and advanced the professions. Taking another page from their historic counterparts, some software companies already purchase professional malpractice insurance. See, e.g., *Phila. Consol. Holding Corp. v. LSI-Lowery Sys., Inc.*, 775 F.3d 1072 (8th Cir. 2015).

While the software industry faces a unique set of challenges, implementation of continuing education on industry standards, security and safety best practices, and ethical and legal boundaries can help prevent problems before they occur. Just as law firms and hospitals require their practitioners to stay apprised of industry standards to improve quality and reduce harm, so too can technology companies, with benefits not just for the public but for the company in its own self-interest. Even in the age of disruption, an ounce of prevention is worth a pound of cure. So said an early innovator and disruptor, who today exists only in history.

Acknowledgments

The author would like to thank Kyle Reynolds (Harvard Law School Class of 2018) for his research assistance.

References

- Agrawal, A.; Gans, J.; Goldfarb, A. 2016. The Simple Economics of Machine Intelligence. *Harvard Business Review* Nov. 17, 2016.
- Bal, B. 2008. An Introduction to Medical Malpractice in the United States. *Clin. Orthopaedics and Related Research* 467(2): 339–347.
- Esteva, A.; Kuprel, B.; Novoa, R.; Ko, J.; Swetter, S.; Blau, H.; and Thrun, S. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542: 115–118.
- Feldman, M. 2016. Watson Proving Better Than Doctors at Diagnosing Cancer. *Top500* (<https://www.top500.org/news/watson-proving-better-than-doctors-in-diagnosing-cancer/>).
- Goertzel, K. 2016. Supply Chain Risks in Critical Infrastructure: Legal Liability for Bad Software. *CrossTalk* Sep.-Oct. 2016.
- Grimstad, P. 2017. In Stargoon’s Car. *The Paris Review* June 29, 2017.
- Hutson, M. 2017. Artificial intelligence prevails at predicting Supreme Court decisions. *Science* May 2, 2017.
- Kowalenko, K. 2012. Licensing Software Engineers Is in the Works: IEEE is helping develop the first-ever licensure exam. *The Institute* February 2012.
- Ladimer, I. 1965. The Health Advertising Program of the National Better Business Bureau. *American Journal of Public Health* 55 (8): 1217–27.
- Matsakis, L. 2017. AI Will Soon Identify Protesters With Their Faces Partly Concealed. *Motherboard* Sep. 6, 2017.
- Miliard, M. 2017. Coalition publishes CDS software design guidelines for a post-21st Century Cures Act landscape. *HealthCare IT News* Sep. 5, 2017.
- Nushi, B.; Kamar, E.; Horvitz, E.; and Kossmann, D. 2017. On Human Intellect and Machine Failures: Troubleshooting Integrative Machine Learning Systems. In *Proceedings of Thirty-First AAAI Conference on Artificial Intelligence*, Menlo Park, Calif.: AAAI Press.
- Russakovsky, O.; Li, L.; Fei-Fei, L. 2015. Best of both worlds: human-machine collaboration for object annotation. *Computer Vision and Pattern Recognition*.
- Sullivan, T. 2017. Half of hospitals to adopt artificial intelligence within 5 years. *Healthcare IT News* Apr. 11, 2017.
- Susskind, R. and Susskind, D. 2016. Technology Will Replace Many Doctors, Lawyers, and Other Professionals. *Harvard Business Review* Oct. 11, 2016.
- Texas Board of Professional Engineers (TBPE). Software Engineering 2012 (<https://engineers.texas.gov/software.html>).
- Tolentino, J. “The Bestseller Code” Tells Us What We Already Know. *The New Yorker* Sep. 2016.
- Vincent, J. 2017. DeepMind’s Go-playing AI doesn’t need human help to beat us anymore. *The Verge* Oct. 18, 2017.
- Westervelt, E. Did A Bail Reform Algorithm Contribute To This San Francisco Man’s Murder? *NPR* Aug. 18, 2017.
- Wittenberg, D. 2017. Artificial Intelligence in the Practice of Law. *ABA Litigation News* Jan. 18, 2017.