

**ARTIFICIAL INTELLIGENCE, TRADE SECRETS, AND THE
CHALLENGE OF TRANSPARENCY**

*John Villasenor**

Artificial intelligence (“AI”) systems can operate in ways that their designers may not fully understand. This creates a series of important questions regarding trade secrets. This Article argues that AI system designers should be able to hold trade secret rights in AI algorithms even when they are unable to articulate how those algorithms operate. However, to assert a misappropriation claim, trade secret owners must be able to acquire information enabling them to describe the algorithm at issue in sufficient detail. This Article also explores compliance with regulatory requirements regarding AI algorithm transparency. While arguing against a maximalist approach to trade secret scope, it identifies approaches enabling trade secret owners to provide the required disclosures while protecting their intellectual property.

TABLE OF CONTENTS

I. INTRODUCTION.....	496
<i>A. Artificial Intelligence</i>	<i>499</i>
<i>B. Trade Secrets: Historical Context and Current Frameworks</i>	<i>504</i>
1. <i>Common Law Roots</i>	<i>504</i>
2. <i>Modern Trade Secret Statutes</i>	<i>505</i>
II. PROTECTING AI TRADE SECRETS	507
<i>A. The Question of Knowledge</i>	<i>508</i>
1. <i>Affirmative Knowledge Not Required</i>	<i>508</i>

* Professor of Electrical Engineering, Law, Public Policy, and Management & Faculty Co-Director of the UCLA Institute for Technology, Law and Policy; Nonresident Senior Fellow, the Brookings Institution. Thanks to Jonas Anderson, Charles Tait Graves, Cynthia Ho, Camilla Hrdy, Dustin Marlan, Tim Murphy, Elizabeth Rowe, Sharon Sandeen, Charlotte Tschider, and Deepa Varadarajan for providing valuable feedback.

2. <i>AI and General Knowledge, Skill, and Experience</i>	510
B. <i>Alleging Misappropriation</i>	512
1. <i>Standing</i>	513
2. <i>Pleadings and “Particularity”</i>	514
3. <i>Identification of Trade Secrets</i>	517
III. AI TRANSPARENCY REQUIREMENTS	520
A. <i>The Challenges to AI System Understanding</i>	521
1. <i>The Meaning of Transparency</i>	521
2. <i>Size and Adaptivity</i>	523
3. <i>The AI Supply Chain</i>	524
B. <i>Meeting AI Transparency Requirements</i>	525
1. <i>Is the Information at Issue Really a Trade Secret?</i>	526
2. <i>Can the Discloser Elect to Use Trade Secrets More Narrowly?</i>	526
3. <i>Can Public Disclosures Avoid Trade Secrets?</i>	529
4. <i>Regulatory Disclosures and Preservation of Trade Secrets</i>	530
5. <i>Would the Information Be Better Protected Through Patents?</i>	532
IV. CONCLUSION	535

I. INTRODUCTION

Artificial intelligence (“AI”) systems can be extraordinarily complex. In addition, they learn from their environment, operating in ways that can be elusive even to their designers. This creates a set of important and timely questions regarding the intersection of AI with trade secret law, which developed under the premise that a trade secret is known by its owner.

First, to what extent should the designers of an AI system hold trade secret rights regarding algorithms that they may be unable to describe? Second, how does the potential lack of knowledge regarding an AI system’s operation impact misappropriation claims? Third, how should calls for AI transparency be satisfied when AI system designers have incomplete knowledge regarding how their systems work? And even if they are able to acquire sufficient

knowledge, what strategies can best navigate tensions between disclosure obligations and trade secret protection?

This Article argues that an AI system designer should be able to hold trade secret rights to AI-developed algorithms, even when the specifics of how those algorithms operate may not initially be known to the designer. This is consistent with trade secret statutory text and case law, and also has advantages from a policy standpoint given the patent eligibility issues associated with AI algorithms.

This Article also examines the challenges of pleading and litigating AI algorithm trade secret misappropriation claims, which will require plaintiffs to describe how the algorithm at issue works. Finally, it identifies a set of approaches to help AI system designers (or owners, etc.) address transparency while also preserving their intellectual property rights. It argues that in the context of AI systems, a maximalist approach to trade secret designation is unnecessary from a business standpoint, counterproductive with respect to the public policy goals of transparency requirements, and often legally incorrect.

To further contextualize these issues, consider an example far removed both in time and in subject matter from today's AI technology landscape. In *Vickery v. Welch*¹ in 1837, the Massachusetts Supreme Court addressed an allegation that the defendant had failed to deliver a secret chocolate recipe to the plaintiff despite having promised to do so.² The *Vickery* court recognized that a recipe for making chocolate can be a trade secret.³ It was not necessary then—nor should it be necessary today—for the owner of a secret chocolate recipe to have knowledge of the specifics of the chemical processes that occur during baking. Nor would the owner need to be able to articulate a detailed physiological explanation of why the resulting chocolate has a taste meeting the approval of buyers. If a chocolate-maker benefits

¹ *Vickery v. Welch*, 36 Mass. 523 (Mass. 1837).

² *Vickery* addressed the defendant's alleged failure to fulfill an agreement to "transfer to the plaintiff, for his exclusive use, the secret manner which the defendant had of making chocolate." *Id.* at 525.

³ "[T]he defendant had used such an exclusive art, which had given great advantage to him in the manufacture of chocolate." *Id.* at 526.

economically from knowing (and keeping secret) the ingredients used and the manner in which they should be combined in the baking process, and if that knowledge is not known by or readily ascertainable by others in the field, the recipe qualifies as a trade secret.⁴ The fact that the value and success of the recipe is in part due to factors that may be outside the chocolate-maker's knowledge in no way diminishes the trade secret status of the recipe.

Now consider a computer programmer who, without using AI, designs a proprietary algorithm that has value by not being known to or readily ascertainable by others in the field, and then writes the software to implement it. Suppose further that the programmer (and the programmer's employer if the programmer is an employee) undertakes reasonable efforts to ensure that the algorithm remains secret. The programmer (or their employer if the programmer is an employee) will have trade secret rights in both the algorithm as well as the corresponding source code.⁵ This remains true even though the programmer may not be able to articulate in detail the full set of reasons why the algorithm is able to successfully perform a task. In this respect, the programmer is in some ways analogous to the chocolate-maker.

Next, suppose that the programmer updates the proprietary algorithm to include AI, so that it will now adapt on its own. The AI system designer, by virtue of having designed the system and directed its operation, should have rights to algorithmic improvements created by it. This does not mean that the AI system is equivalent to an employee. As Jeanne Fromer has written, “[w]hereas departing employees can legally take their elevated general knowledge and skill to new jobs, a key path by which knowledge spills across an industry, machines automating employees’ tasks will never take their general knowledge and skill

⁴ An additional requirement is that the information must not constitute general knowledge, skill, and experience, though courts are inconsistent in how they interpret and apply this exclusion. See Camilla A. Hrady, *The General Knowledge, Skill, And Experience Paradox*, 60 B.C. L. REV. 2409 (2019).

⁵ While formal trade secret definitions are discussed further *infra*, this example is constructed so that the algorithm in question meets the statutory definition of a trade secret. See, e.g., the federal definition in 18 U.S.C. § 1839(3).

elsewhere to competitors.”⁶ The complexity in the trade secret analysis lies in the fact that, unlike the chocolate-maker who identifies a way to refine a recipe, the AI system designer may not initially be aware of what the AI-generated algorithmic improvements are, though the designer may discern the better performance made possible through those improvements.

The remainder of this Article is devoted to exploring that complexity, as well as its implications, and proceeds as follows. This Article first presents a short overview of AI, with emphasis on its ability to adapt, and provides a brief review of the history of trade secret law and of current statutory frameworks. Part II explores the scope of trade secret protection for AI systems and addresses the requirements for stating and litigating a misappropriation claim. Part III examines transparency, both in terms of the challenges of obtaining information about the operation of AI systems that may initially appear opaque, and in terms of balancing AI transparency obligations with trade secret rights. Conclusions are offered in Part IV.

A. Artificial Intelligence

While AI has a decades-long history,⁷ it is only in recent years that it has gained large-scale visibility and attention among policymakers, legal scholars, and the broader public. There are many different definitions of AI, but nearly all of them describe systems that can adapt their actions in order to respond to their environments. A definition that was codified into federal law under the National Artificial Intelligence Initiative Act of 2020 is representative:

The term “artificial intelligence” means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to—

⁶ Jeanne C. Fromer, *Machines as the New Oompa-Loompas: Trade Secrecy, the Cloud, Machine Learning, and Automation*, 94 N.Y.U. L. REV. 706, 725–26 (2019) (citation omitted).

⁷ In 1950, the pioneering computer scientist Alan Turing published a scientific paper asking, “Can machines think?” Alan M. Turing, *Computing Machinery & Intelligence*, 59 MIND 433, 433 (1950).

(A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action.⁸

The ability to learn and adapt is the foundation of AI's extraordinary potential. There are variations in how AI systems learn. For instance, in supervised machine learning, an AI system is presented with labeled data.⁹ If the goal is to train an AI system to differentiate between cars and trucks, presenting the system with 1000 images of cars with the data label "cars" and 1000 images of trucks with the label "trucks" will enable the system to analyze both the images and their labels as it learns to differentiate cars from trucks. By contrast, in unsupervised learning the AI system is trained using unlabeled data and has to infer categories (e.g., cars vs. trucks) without human assistance.¹⁰ Because it uses a training data set with less information (e.g., 1000 unlabeled images each of cars and trucks), unsupervised learning is more difficult than supervised learning.

Another variation in AI training lies in the extent of human input at the start of the process regarding an algorithm to use. One option is for a person to provide an AI system with an initial algorithm, which the system then further develops as it receives additional data. There are also AI systems that learn algorithms from scratch with essentially no human input. In 2018, a group of researchers from DeepMind (now part of Google) published a paper describing AlphaZero, an AI system that, "given no domain knowledge except the game rules" for chess, was able to teach itself to achieve "superhuman performance."¹¹

Additionally, AI system designers can put limits on when and how AI systems can evolve. A driverless car company might use AI internally to develop driving algorithms while also ensuring that the driving algorithms, once they are thoroughly tested, are static in the cars it sells so that they do not evolve further. This allows the system

⁸ 15 U.S.C. § 9401(3).

⁹ *Supervised vs. Unsupervised Learning: What's the Difference?*, IBM, <https://www.ibm.com/blog/supervised-vs-unsupervised-learning/> [<https://perma.cc/R42B-BRRA>] (last visited Jan. 7, 2024).

¹⁰ *Id.*

¹¹ David Silver et al., *A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go Through Self-Play*, 362 SCI. 1140, 1140 (2018).

designers to harness the power of AI in a controlled environment while avoiding the potential problems that might arise if an AI system for a safety-critical task were to evolve post-deployment in an unforeseen manner. There are also circumstances where AI system designers will choose to allow an AI algorithm to continue to evolve, as the resulting improvements can allow it to perform tasks more effectively.

A hypothetical scenario involving algorithmic stock trading helps to illustrate how people and AI systems can interact. Consider a team of investors who design an automated (but initially not AI-enabled) system to use in trading the stock of an airline called Airline, Inc. The goal is to predict future moves of Airline, Inc.'s share price so the investors can purchase shares ahead of anticipated increases and sell shares ahead of anticipated decreases.

Based on careful analysis of historical data, the investors have selected a bundle of indicators¹² and a specific manner in which they can be combined in order to produce share price predictions. Assume further that this algorithm is not known to or easily ascertainable by others in the field. Without using AI, the investors write and then start running software that generates predictions, using the resulting predictions to buy and sell shares of Airline, Inc. Under this scenario, the investors have trade secret rights in the algorithm, in the software they have written to implement it, and likely in the data the software produces documenting how often its predictions are correct.

Next, suppose that the investors decide to augment their system using AI. They modify the software so that it can learn on its own by continuously evaluating the predictive utility of each of the indicators. Acting autonomously, the AI system stops using the indicators that turn out to have little utility, increases the weight given to those with proven utility, and introduces new indicators that were not originally identified by the investors. Over time, the prediction algorithm becomes very different from the one originally

¹² In this context, an indicator is a piece of data that has potential to have predictive value regarding moves in the share price of Airline, Inc. Examples could include the stock price of other companies including other airlines, the price of a commodity such as oil, broader market metrics, etc.

developed by the investors, not only using a different set of indicators but also growing in size and complexity.

Provided that the algorithm, after its modification using AI, has value by not being known to or easily ascertainable by others in the field, and that the investors make reasonable efforts to maintain it as secret, the investors should have trade secret rights in it. This is true despite the changes the algorithm has undergone and despite the fact that the investors who designed the original algorithm would no longer even know what the current algorithm is unless they get that information from the AI software.

It was the investors who provided the key intellectual input and developed the original algorithm by writing the software to implement it, recognizing that AI could be used to increase the system's performance, and modifying the software to incorporate AI. Yet the fact remains that, due to the addition of AI, unless they examine the post-adaptation algorithm, the investors will not have knowledge of the specific algorithm the system is using.

Of course, it could be argued that “the algorithm” is not a particular instantiation at any given point in time but rather the overall framework used in the system to perform adaptation, which the investors created and are very much aware of. But regardless of these semantic issues, the ability of an AI system to develop computational approaches that the system's designers do not have knowledge of opens the door to fact patterns different from those that have traditionally been considered in trade secret litigation. While there is a growing body of literature regarding trade secrets in the context of AI,¹³ there has been far less attention (the work of

¹³ See, e.g., Fromer, *supra* note 6, at 708 (arguing that recent advances in computing, including in machine learning, “allow businesses to circumvent trade secret law’s central limitations, thereby overfortifying trade secrecy protection”); Hawraa Hammoud, *Trade Secrets and Artificial Intelligence: Opportunities & Challenges* 10 (Dec. 29, 2020) (unpublished manuscript) (on file with SSRN) (concluding that “[t]rade [s]ecrets are well suited to protect AI technology against misappropriation”); Gregory Hagen, *AI and Patents and Trade Secrets*, in *ARTIFICIAL INTELLIGENCE AND THE LAW IN CANADA 1* (Florian Martin-Bariteau & Teresa Scassa eds., 2021) (discussing “how both patent law and trade secret law could adapt to ensure that sufficient information about inventions and automated decisions is disclosed to further knowledge and accountability”); Jordan R. Jaffe

Jeanne Fromer and Charlotte Tschider being notable exceptions)¹⁴ directed to the specific trade secret law questions raised by the adaptive nature and complexity of AI algorithms.

This Article focuses primarily on the intersection between trade secrets and AI algorithms. Of course, trade secrets can also apply to many other aspects of an AI system as well, including approaches to training, the data used in training, the design of software and

et al., *The Rising Importance of Trade Secret Protection for AI-Related Intellectual Property 1* (Apr. 24, 2020) (unpublished manuscript) (on file with Quinn Emanuel) (“explor[ing] the tradeoffs between patents and trade secrets in the AI sector” and “describ[ing] how trade secrets have become essential tools for companies to protect their AI-related intellectual property”); Nari Lee, *Protection for Artificial Intelligence in Personalised Medicine – The Patent/Trade Secret Trade Off*, in *THE HARMONIZATION AND PROTECTION OF TRADE SECRETS IN THE EU – AN APPRAISAL OF THE EU DIRECTIVE 267*, 267–94 (Jens Schovsbo et al. eds., 2019) (addressing the respective roles of patents and trade secrets in relation to AI in relation to personalized medicine); Mariateresa Maggiolino, *EU Trade Secrets Law and Algorithmic Transparency 3* (Bocconi Legal Stud. Rsch. Paper Series, Working Paper No. 3363178, 2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3363178 [<https://perma.cc/WX3P-QY9C>] (arguing that in the context of EU law, “the most important sources of algorithmic transparency lay outside EU Trade Secret law”); Jessica Meyers, *Artificial Intelligence and Trade Secrets*, 11 *LANDSLIDE* 17, 21 (2019) (writing that “[w]hile businesses have legitimate interests in protecting proprietary information, such as the method by which a decision is made, individuals also have legitimate rights to know that AI algorithms are created and applied in a fair manner”); Ulla-Maija Mylly, *Transparent AI? Navigating Between Rules on Trade Secrets and Access to Information*, 54 *INT. R. INTELL. PROP. & COMPETITION L.* 1013, 1013 (2023) (examining “the role of disclosure obligations under the provisions of the [EU’s Artificial Intelligence Act]” and “the tension between obligations to disclose information on the one hand and requirements to protect the trade secrets contained in the technical details of AI on the other”); Sharon K. Sandeen & Tanya Aplin, *Trade Secrecy, Factual Secrecy and the Hype Surrounding AI*, in *RESEARCH HANDBOOK ON INTELLECTUAL PROPERTY AND ARTIFICIAL INTELLIGENCE* 443 (Ryan Abbott ed., Edward Elgar Publ’g 2022) (distinguishing between factual secrecy and trade secrecy and arguing that with respect to AI, “the claim of trade secrets protection is overstated”).

¹⁴ See Fromer, *supra* note 6, at 688; Charlotte A. Tschider, *Beyond the ‘Black Box’*, 98 *DENV. L. REV.* 683, 688 (2021) (identifying “key issues with calls for AI transparency and explainability, including the role of trade secrecy in preventing disclosure of useful information and technical complexities”).

hardware for implementing the system, and its performance characteristics.

B. Trade Secrets: Historical Context and Current Frameworks

1. Common Law Roots

Modern American trade secret law reflects a combination of common law, modern state law, and federal statutory law.¹⁵ In 1868 the Massachusetts Supreme Court decided *Peabody v. Norfolk*,¹⁶ a case the U.S. Supreme Court a century later characterized as the decision through which “trade secret law was imported into this country.”¹⁷ The *Peabody* court, citing its own 1837 *Vickery* decision as well as *Morison v. Moat*¹⁸ in England, wrote that “[i]n this court, it is settled that a secret art is a legal subject of property.”¹⁹ The *Peabody* court also quoted former Supreme Court justice Joseph Story’s observation that “courts of equity will restrain a party from making a disclosure of secrets communicated to him in the course of a confidential employment.”²⁰

The Restatement (First) of Torts (“First Restatement”), published in 1939 by the American Law Institute, provided a definition of trade secrets reflecting contemporary common law and stated in part that a “trade secret may consist of any formula, pattern, device or compilation of information which is used in one’s business, and which gives him an opportunity to obtain an advantage over competitors who do not know or use it.”²¹ The First

¹⁵ For a detailed history of U.S. trade secret law, see Sharon Sandeen, *The Evolution of Trade Secret Law and Why Courts Commit Error When They Do Not Follow the Uniform Trade Secrets Act*, 33 *HAMLIN L. REV.* 493 (2010).

¹⁶ *Peabody v. Norfolk*, 98 Mass. 452 (Mass. 1868).

¹⁷ *Kewanee Oil Co. v. Bicron Corp.*, 416 U.S. 470, 493 n.23 (1974).

¹⁸ *Morison v. Moat*, 9 Hare 241, 20 L.J. Ch. 513 (1851).

¹⁹ *Peabody v. Norfolk*, 98 Mass. at 459–60. The *Peabody* court noted that the English *Morison v. Moat* court had concluded that “[t]here is no doubt whatever, that when a party who has a secret in trade employs persons under a contract express or implied, or under duty express or implied, those persons cannot gain the knowledge of the secret and then set it up against their employer.” *Id.* at 459 (quoting *Morison v. Moat*, 21 L.J. Ch. 248 (1852)).

²⁰ *Id.* at 459 (quoting JOSEPH STORY, 2 *COMMENTARIES ON EQUITY JURISPRUDENCE* § 952, at 223 (1839)).

²¹ *RESTATEMENT (FIRST) OF TORTS* § 757 cmt. b. (AM. L. INST. 1939).

Restatement also identified a set of six factors for evaluating whether information is protectable as a trade secret.²² This framework was very influential in the decades following its publication, with the definition and six factor-test “cited approvingly in virtually every U.S. jurisdiction.”²³ Due to subsequent statutory law developments, the First Restatement’s direct influence has diminished considerably since the late twentieth century, although some courts today still use its six-factor test.²⁴

2. *Modern Trade Secret Statutes*

In the late 1960s, in response to what it later characterized as the “uneven” development of state trade secret law and the resulting “undue uncertainty concerning the parameters of trade secret protection,” the Uniform Law Commission began work to develop the Uniform Trade Secrets Act (“UTSA”) model legislation.²⁵ The UTSA, which “codifies the basic principles of common law trade secret protection,”²⁶ was published in 1979 and then revised in

²² *Id.*

²³ ROGER M. MILGRIM & ERIC E. BENSON, MILGRIM ON TRADE SECRETS § 1.01 (2024).

²⁴ *See, e.g.,* Pauwels v. Deloitte LLP, 2023 U.S. App. LEXIS 26597 (2d Cir. 2023) at *13 (“Under New York law . . . possession of a trade secret—is generally evaluated with reference to six factors”) (New York is one of two states not to have enacted the UTSA); Heil Trailer Int’l Co. v. Kula, 542 Fed. Appx. 329, 330 (5th Cir. 2013) (“the applicable test for trade secret status described by the Supreme Court of Texas in *In re Bass* is a six-factor balancing test”); Hoog v. Dometic Corp., 2023 U.S. Dist. LEXIS 175299 (W.D. Okla. 2023) at *11 n.6 (“The Oklahoma Supreme Court has adopted six factors from the Restatement of Torts, § 757, to help determine whether information is a trade secret”); Woodstream Corp. v. Nature’s Way Bird Prods., LLC, No. 1:23-cv-294, 2023 U.S. Dist. LEXIS 167025, at *5 (N.D. Ohio 2023) (“The Ohio Supreme Court considers six factors in determining whether an item constitutes a trade secret”); Allstate Ins. Co. v. Fougere, 79 F.4th 172, 188 (1st Cir. 2023) (identifying the six-factor test).

²⁵ UNIF. TRADE SECRETS ACT, Prefatory Note (UNIF. L. COMM’N 1985) [hereinafter UTSA].

²⁶ *Id.*

1985,²⁷ and has since been enacted (with some variations)²⁸ in nearly all U.S. states and the District of Columbia.²⁹

The growth in computer networks and in the number of reported computer breaches in the mid-1990s generated concern among federal policymakers that computer hacking could be used to remotely access and extract trade secret information from the computers of U.S. companies, including for purposes of economic espionage.³⁰ This led Congress to pass the Economic Espionage Act of 1996 (“EEA”),³¹ which established criminal penalties for trade secret theft with knowledge “that the offense will benefit any foreign government.”³² In a statement accompanying the bill’s signing, President Bill Clinton wrote that “[u]ntil today, Federal law has not accorded appropriate or adequate protection to trade secrets, making it difficult to prosecute thefts involving this type of information.”³³ Despite reciting economic espionage in its title, the EEA also included separate language criminalizing trade secret theft without any requirement that the theft benefit a foreign government, provided that it involved a trade secret “related to or included in a

²⁷ “On August 9, 1979, the Act was approved and recommended for enactment in all the states On August 8, 1985, four clarifying amendments were approved and recommended for enactment in all the states.” *Id.*

²⁸ For a state-by-state comparison of state trade secret laws to the UTSA, see *Trade Secrets Laws and the UTSA: 50 State and Federal Law Survey*, BECK REED RIDEN LLP (Jan. 24, 2017), <https://beckreedriden.com/trade-secrets-laws-and-the-utsa-a-50-state-and-federal-law-survey-chart/> [<https://perma.cc/X43U-9DMH>].

²⁹ See *Uniform Trade Secrets Act*, UNIF. L. COMM’N, <https://www.uniformlaws.org/committees/community-home?CommunityKey=3a2538fb-e030-4e2d-a9e2-90373dc05792> [<https://perma.cc/X43U-9DMH>] (last visited Oct. 26, 2023) (indicating enactment in the District of Columbia and all U.S. states other than New York and North Carolina).

³⁰ *Clinton Cracks Down on Hackers*, CNET (Oct. 14, 1996), <https://www.cnet.com/tech/services-and-software/clinton-cracks-down-on-hackers/> [<https://perma.cc/73Z8-M4AR>].

³¹ Economic Espionage Act of 1996, Pub. L. No. 104-294, 110 Stat. 3488 (1996) [hereinafter EEA].

³² *Id.* §101 (codified at 18 U.S.C. § 1831).

³³ William J. Clinton, *Statement on Signing the Economic Espionage Act of 1996*, AM. PRESIDENCY PROJECT (Oct. 11, 1996), <https://www.presidency.ucsb.edu/documents/statement-signing-the-economic-espionage-act-1996> [<https://perma.cc/HNG9-ZUTD>].

product that is produced for or placed in interstate or foreign commerce.³⁴

In 2016, President Barack Obama signed the Defend Trade Secrets Act of 2016 (“DTSA”) into law.³⁵ The DTSA introduced a federal private right of action for trade secret misappropriation and is thus partially duplicative of rights already available at the state level. There are, however, some differences. To be cognizable under the DTSA, a misappropriation claim must be related to interstate or foreign commerce. For a claim not meeting this requirement, a claimant would, absent diversity jurisdiction, not have access to the DTSA.

The language of both the UTSA and DTSA lacks an affirmative requirement that a trade secret be used. This provides a notable contrast with the First Restatement, which included commentary describing a trade secret as “a process or device for continuous use in the operation of the business.”³⁶ That said, if a trade secret is not used,³⁷ it can lose its economic value and thus its trade secret status on that basis. As Camilla Hrdy and Mark Lemley have written, trade secrets “can . . . expire when they are abandoned due to failure to derive value from their secrecy.”³⁸

II. PROTECTING AI TRADE SECRETS

Not all information that is factually secret will meet the statutory requirements of a trade secret. Sharon Sandeen and Tayna Aplin have explained that in the specific context of AI, “the claim of trade secrets protection is overstated,” and “not all AI related information would (or should) qualify for trade secret protection even if it is

³⁴ EEA §101 (codified at 18 U.S.C. § 1832).

³⁵ Defend Trade Secrets Act of 2016, Pub. L. No. 114-153, 130 Stat. 376 (2016) [hereinafter DTSA].

³⁶ RESTATEMENT (FIRST) OF TORTS § 757 cmt. b. (AM. L. INST. 1939).

³⁷ “Use” of a trade secret is broad enough to encompass negative information. For instance, knowledge that a particular machine learning algorithm is *not* well suited to a given task can qualify as a trade secret under the UTSA and DTSA if it meets the requirements of the statutory definition.

³⁸ Camilla A. Hrdy & Mark A. Lemley, *Abandoning Trade Secrets*, 73 STAN. L. REV. 1, 66 (2021).

factually secret or treated as secret by its holder.”³⁹ Additionally, a company’s designation of information as “confidential” or “proprietary” does not automatically confer trade secret status.

Notwithstanding the above caveats, AI can implicate trade secrets in many different ways. Trade secret law can protect an AI system algorithm, source code, information in documents describing its design, the process of selecting and using training data, and knowledge gathered during testing regarding how to improve its performance. A company’s plans for designing, building, and bringing an AI-based product or service to market can also qualify (though would not automatically qualify) as trade secrets.

Thus, in many respects, trade secret protection for AI is no different than for non-AI computer-based products and services. However, there are novel trade secret law questions arising from the adaptive nature and complexity of AI systems, which creates a gap between the knowledge of the AI system designer and the actual behavior of the system.⁴⁰ The implications of that knowledge gap are discussed next.

A. The Question of Knowledge

1. Affirmative Knowledge Not Required

Consider an AI system that, through adaptation, has evolved to the point where an algorithm it is executing is quite different from the one initially envisioned and programmed by its designer. Assume further that the resulting algorithm is neither known nor readily ascertainable to others working in the same field of endeavor, and that it has economic value on that basis. And, assume that the designer, who is an employee of a company for whom the designer did the work to create the AI system, instructs the AI system to output a human-readable description of the current state of the algorithm. By reviewing that output, the designer gains a full understanding of how the algorithm is currently operating. At that

³⁹ Sandeen & Aplin, *supra* note 13, at 443–44 (internal citation omitted).

⁴⁰ While it has long been true that not all employees in a company will know all of the company’s trade secrets, with AI systems it can be possible to create trade secrets that *no* employee of the company knows.

snapshot in time, the information about how the AI system works is clearly a trade secret.

Now turn the clock back to just before the AI system designer reviewed the output describing the current operation of the AI algorithm. Prior to that review, when the system designer did not understand the algorithm's current operation, was the information a trade secret? This Article argues that the answer is yes. Both statutory and policy arguments support this answer.

Statutory definitions of trade secrets recite the *lack* of knowledge of non-owners of trade secrets but do not explicitly mention the *level* of knowledge of owners. For instance, the UTSA defines a trade secret as information that, among other requirements, “derives independent economic value . . . from not being generally known to, and not being readily ascertainable by proper means by” others in the field.⁴¹ The definition of a trade secret in federal law contains substantially identical language.⁴² Neither of the definitions mentions the affirmative knowledge of the trade secret owner. In addition, the federal definition of a trade secret “owner” says nothing about affirmative knowledge, referring instead only to the “person or entity” that has “title to, or license in” the trade secret.⁴³ Notably, the definition of owner is silent on the entity that *developed* the trade secret and thus provides no affirmative requirement that the entity be human.

The policy argument is best made by considering the problematic consequences that would ensue if the information regarding how the AI system is currently operating is *not* deemed a

⁴¹ UTSA § 1(4).

⁴² Compare 18 U.S.C. § 1839(3) (“‘[T]rade secret’ means all forms and types of financial, business, scientific, technical, economic, or engineering information, . . . [that] derives independent economic value . . . from not being generally known to, and not being readily ascertainable *through proper means* by, another person who can obtain economic value from the disclosure or use of the information (emphasis added)), with UTSA § 1(4) (“‘Trade secret’ means information . . . that: (i) derives independent economic value, actual or potential, from not being generally known to, and not being readily ascertainable *by proper means* by, other persons who can obtain economic value from its disclosure or use” (emphasis added)).

⁴³ 18 U.S.C. § 1839(4).

trade secret. Imagine that the system designer in the above example has a work colleague who is a rogue employee. The rogue employee knows that the algorithm is valuable but does not know how it works. Motivated by an intent to commit misappropriation, the rogue employee instructs the AI system to output a human-readable description of the current algorithm. The rogue employee then leaves their employment and uses that information to start a competing business. It would belie logic to conclude that the information is not a trade secret.

One response to this thought experiment might be to assert that the information *became* a trade secret the moment the rogue employee reviewed and understood the algorithm description, thereby turning the rogue employee's actions of taking that information to a competitor into a classic misappropriation case. But suppose the rogue employee downloaded the algorithm description while still employed but did not review it until after resigning from their job and starting the competing business, ensuring that they did not have full knowledge of the algorithm at the moment they resigned. Again, it would belie logic to use the timing of the rogue employee's understanding of the algorithm to assert that the information taken was not a trade secret. The trade secret lies in the information, regardless of whether or when either the AI system designer or the rogue employee chose to review it and understand its implications.

2. *AI and General Knowledge, Skill, and Experience*

As the foregoing hypothetical illustrates, the information regarding how an AI algorithm operates can be a trade secret even when no person has in their mind sufficient knowledge to describe in detail how it works. But is it *necessarily* a trade secret? The answer is no. As an initial matter, trade secrets by definition cannot include information that is generally known to or readily ascertainable by others in the field. An AI-developed algorithm that is known or readily ascertainable, even if it is considered new to the particular people who designed the AI system that generated it, would not qualify as a trade secret.

But there is also another wrinkle: Trade secret law does not cover an employee's general knowledge, skill, and experience,

which Charles Tait Graves has described as “the corpus of information every employee may take from job to job.”⁴⁴ As Hrdy observes, this creates a paradox under which “employers are encouraged to communicate trade secrets to employees, but this information loses protection if it becomes part of those employees’ unprotectable general knowledge, skill, and experience.”⁴⁵ Courts often fail to account for this factor, which limits the scope of information protectable as a trade secret.⁴⁶

AI raises the question of whether this doctrine can exclude from protection some AI-developed algorithms.⁴⁷ The answer is yes. Getting that answer does not require concluding that an AI system has skill, experience, and perhaps knowledge, although those are attributes an AI system arguably possesses. The doctrine developed to protect *people* who acquire those things, enabling them to freely change employers. There is no corresponding policy imperative to protect the interests of an AI system.

However, the general knowledge, skill, and experience exclusion to the scope of trade secret protection will arise in relation to AI systems because what people working with such systems know is still at the center of an inquiry regarding whether information is protectable as a trade secret. For an algorithmic improvement developed by AI to be a trade secret, it must fall outside the general knowledge, skill, and experience exclusion, even if no person at the company that owns the AI system had originally articulated that particular algorithmic improvement. In short, while an AI system can be used to develop algorithms that may be new trade secrets, it

⁴⁴ Charles Tait Graves, *Trade Secrets as Property: Theory and Consequences*, 15 J. INTELL. PROP. L. 39, 52 (2007).

⁴⁵ Hrdy, *supra* note 4, at 2410.

⁴⁶ *Id.* (“Courts therefore miss the category of information that, while technically secret to a company, is nonetheless unprotectable.”).

⁴⁷ This is not to suggest that AI-developed algorithms necessarily will be, or will likely be, unprotectable. Rather, AI-developed algorithms do not get an exception from the requirement that trade secrets cannot comprise general knowledge, skill, and experience of the relevant employees. If an AI system happens to develop an algorithm that falls within the general knowledge, skill, and experience exclusion, it should not be a trade secret.

cannot be used to expand the scope of information protectable under trade secret law.

B. Alleging Misappropriation

Holding trade secret rights is different from *asserting* them.⁴⁸ Unlike patents, which are issued and published by the U.S. government to inform the public of the existence, description, and ownership of inventions, trade secrets are secret. While a trade secret owner may choose to publicize the *existence* of a trade secret (e.g., a food products company may advertise that it has a “secret recipe”), companies often choose not to publicly disclose even general category information about their trade secrets. Often, the fact of a trade secret’s existence becomes known only when a trade secret owner files a misappropriation claim. Another important contrast with patents is that patents are presumptively valid.⁴⁹ By contrast, there is no legal presumption that a misappropriation plaintiff is correct in asserting that certain information constitutes a trade secret.

In the context of pursuing misappropriation claims for AI algorithms developed through automated adaptation, asserting trade secret rights will require describing what those algorithms are doing. It will not be sufficient for plaintiffs to state, in effect, “we are not sure how the AI algorithm works, but whatever it is doing, it is our trade secret, and the defendant has misappropriated it.” Plaintiffs will need to provide enough information about the trade secret in the complaint to survive a motion to dismiss, and later to provide significantly more detail, so that the defendant (and the court) knows what specific information is at issue and can respond accordingly.

⁴⁸ Regarding pleading trade secret cases under the DTSA, see William L. Schaller, *On Equipoise, Knowledge, and Speculation: A Unified Theory of Pleading Under the Defend Trade Secrets Act – Jurisdiction, Identification, Misappropriation, and Inevitable Disclosure*, 27 J. INTELL. PROP. L. 137 (2020).

⁴⁹ The presumption of validity does not preclude a later challenge alleging invalidity. Patent validity can be challenged in federal district court or through the Patent and Trademark Office.

1. *Standing*

Under the DTSA, the “owner” of a trade secret is “the person or entity in whom or in which rightful legal or equitable title to, or license in, the trade secret is reposed.”⁵⁰ In federal court, ownership is a central aspect to establishing standing. At the state level, standing is more complicated. As Graves writes in relation to trade secret litigation in state courts, “[t]he trend for two decades has been to broaden the class of potential plaintiffs to almost anyone who has possession of information it claims to be secret.”⁵¹ Possession is a far lower hurdle than ownership, as it does not implicate the manner in which possession was obtained. For instance, a company that licenses from a third party trade secret information regarding a manufacturing process possesses but does not own the information. This difference in scope means that there is a set of claimants that might have standing to bring state trade secret claims while lacking it with respect to federal claims under the DTSA.

Assuming that an AI system designer worked alone (or that if the designer is a company, that the company worked alone) to develop the AI system, then the designer would own the associated trade secrets, and thus have standing to bring a misappropriation claim in both state and federal courts. But things are more complicated when, as will often occur, an AI system is designed through the combined actions of multiple entities. If Company A contracts with Company B to help it design an AI system, Company A will (typically) be the owner of the system, but both companies may possess it. Absent a contractual provision otherwise (e.g., by specifically providing for some degree of joint ownership), this could exclude Company B from pursuing a federal trade secret misappropriation claim against a third party, even if both companies—with the benefit of hindsight—would have wanted Company B to have the option to do so.

Sometimes federal and state differences in standing requirements will work in the other direction. Company A might *not* want Company B to have the option to bring a misappropriation

⁵⁰ 18 U.S.C. § 1839(4).

⁵¹ Charles Tait Graves, *Curiosities of Standing in Trade Secret Law*, 20 NW. J. TECH. & INTELL. PROP. 159, 161 (2023).

claim against a third party, but because Company B possesses but does not own the AI system, it might have standing in state (but not federal) court to do so. This can again be addressed contractually, e.g., through a provision in the contract stating that only Company A can initiate a claim.

2. *Pleadings and “Particularity”*

While any trade secret misappropriation plaintiff must grapple with the question of how much detail to include in a complaint, things are more complex when the trade secret at issue is an AI algorithm about which the plaintiff’s own knowledge may initially be limited. An additional consideration is that standards for pleading misappropriation vary across jurisdictions and are evolving.

The California Uniform Trade Secrets Act requires that a misappropriation claim “identify the trade secret with reasonable particularity.”⁵² That standard has its roots in a 1968 California state appeals court ruling, *Diodes, Inc. v. Franzen*,⁵³ concluding that a plaintiff must describe “the trade secret with sufficient particularity to separate it from matters of general knowledge in the trade or of special knowledge of those persons who are skilled in the trade, and to permit the defendant to ascertain at least the boundaries within which the secret lies.”⁵⁴ A “reasonable particularity” pleading requirement is also provided under Massachusetts law.⁵⁵ South Carolina’s trade secret statute requires “particularity” as a precondition for discovery.⁵⁶ While California, Massachusetts, and South Carolina are exceptions in codifying a “particularity”

⁵² CAL. CODE CIV. PROC. § 2019.210.

⁵³ *Diodes, Inc. v. Franzen*, 260 Cal. App. 2d 244 (Cal. Ct. App. 1968).

⁵⁴ *Id.* at 253.

⁵⁵ MASS. GEN. LAWS ch. 93, § 42D(b) (2018) (stating that a misappropriation claimant “must state with reasonable particularity . . . the nature of the trade secrets . . .”).

⁵⁶ S.C. CODE ANN. § 39-8-60(B)(1) (1997) (providing that discovery can proceed only if “the allegations in the initial pleading setting forth the factual predicate for or against liability have been plead with particularity”).

requirement into statutory law, state courts in Delaware have also applied that standard.⁵⁷

The requirements for pleading a DTSA claim in federal court are generally guided by Rule 8 of the Federal Rules of Civil Procedure,⁵⁸ which, as interpreted by the Supreme Court in *Ashcroft v. Iqbal*,⁵⁹ requires that a complaint must allege facts that “plausibly suggest an entitlement to relief” in order to survive a motion to dismiss.⁶⁰ An exception is when fraud is alleged—something that will occur in some but not all trade secret cases. In that event, a different rule applies, and the complainant must “state with particularity the circumstances constituting” the alleged fraud.⁶¹

Interestingly, while neither Rule 8 nor its construction in *Iqbal* recites “particularity,” an increasing number of federal courts are nonetheless requiring particularity (or “specificity”) in trade secret misappropriation complaints, including those not alleging fraud. In a 2020 decision, a New York federal district court explained that “[a]lthough the Second Circuit has not articulated a specificity requirement, district courts in this circuit routinely require that plaintiffs plead their trade secrets with sufficient specificity to inform the defendants of what they are alleged to have misappropriated.”⁶² The “sufficient particularity” language from *Diodes, Inc.*, was adopted in a 2021 Third Circuit ruling that aimed “to clarify the requirements for pleading a trade secret misappropriation claim under the” DTSA,⁶³ and that is likely to be highly influential in other circuits.

⁵⁷ See, e.g., *SmithKline Beecham Pharms. Co. v. Merck & Co., Inc.*, 766 A.2d 442, 447 (Del. 2000) (“The plaintiff must disclose the allegedly misappropriated trade secrets with reasonable particularity.”).

⁵⁸ A pleading must include “a short and plain statement of the claim showing that the pleader is entitled to relief.” FED. R. CIV. P. 8(a)(2).

⁵⁹ *Ashcroft v. Iqbal*, 556 U.S. 662 (2009).

⁶⁰ *Id.* at 680 (stating that for a complaint to survive a motion to dismiss, it must allege facts that “plausibly suggest an entitlement to relief”).

⁶¹ FED. R. CIV. P. 9(b).

⁶² *Zirvi v. Flatley*, 433 F. Supp. 3d 448, 465 (S.D.N.Y. 2020) (quoting *ExpertConnect, L.L.C. v. Fowler*, No 18 Civ. 4828 (LGS), 2019 U.S. Dist. LEXIS 114931, at *4 (S.D.N.Y. July 10, 2019)).

⁶³ *Oakwood Labs. LLC v. Thanoo*, 999 F.3d 892, 896 (3d Cir. 2021). The decision also quoted the *Diodes, Inc.* “sufficient particularity” standard. *Id.* at 252.

While the landscape is evolving, it is clear that failure to provide sufficient detail regarding the information at issue in a trade secret complaint can lead to dismissal. A 2016 ruling from a California federal district court in an AI trade secret case is instructive. In *Loop AI Labs Inc. v. Gatti*,⁶⁴ the court granted defendant IQSystem, Inc.’s “motion to enforce the court’s . . . order directing Plaintiff Loop AI Labs Inc. to submit a particularized trade secret disclosure.”⁶⁵ In identifying the allegedly misappropriated trade secrets, the plaintiff listed (among others) “confidential modeling and discussions regarding its development and experimentation of supervised and unsupervised artificial intelligence technology.”⁶⁶ In response, the court found that the plaintiff’s “technique of listing general concepts or categories of information is plainly insufficient; Defendants cannot fairly be expected to rebut Plaintiff’s trade secrets claim without a reasonably concrete definition of the purported secrets.”⁶⁷

These trends make clear that if the allegation is misappropriation of an AI-generated algorithm, as distinct from the code for implementing it, a complaint will need to go beyond describing it only using very broad terms, such as “an AI-based algorithm for stock trading.” Rather, the complaint will have to provide enough detail to meet the applicable pleading standard—which in a growing number of courts is “particularity” or “specificity.”

Some defendants in AI algorithm misappropriation cases will no doubt assert that the information at issue was not known to the owner at the time of the alleged misappropriation, and thus cannot be a trade secret. Courts should reject that argument for the reasons discussed earlier. As one federal district court wrote in a 1996 state trade secret case unrelated to AI, “[t]he owner of a trade secret need not necessarily recognize the full value of his knowledge, or be aware of all the information’s ramifications.”⁶⁸ In the AI era, courts

⁶⁴ *Loop AI Labs Inc. v. Gatti*, 195 F. Supp. 3d 1107 (N.D. Cal. 2016).

⁶⁵ *Id.* at 1109. The case was subsequently dismissed. *See Loop AI Labs Inc. v. Gatti*, 2017 U.S. Dist. LEXIS 34109 at *58 (9th Cir. May 20, 2016).

⁶⁶ *Id.* at 1114.

⁶⁷ *Id.* at 1114–15.

⁶⁸ *Glaxo Inc. v. Novopharm Ltd.*, 931 F. Supp. 1280, 1304 (E.D.N.C. 1996). This case involved North Carolina’s trade secret statute. *Id.* at 1299 (citing N.C. Gen. Stat. § 66-152(3)).

in state and federal trade secret cases have full statutory support to go further, and can conclude that while knowledge regarding the workings of an AI algorithm is a prerequisite to filing a misappropriation claim, it is not a prerequisite to being a victim of misappropriation.

3. *Identification of Trade Secrets*

Once a misappropriation case proceeds past the pleading stage, a plaintiff needs to describe the trade secret more substantively. In this regard, a 2021 publication by a trade secrets working group of the Sedona Conference titled *Commentary on the Proper Identification of Asserted Trade Secrets in Misappropriation Cases* (“Sedona Commentary”) provides useful context.⁶⁹

The Sedona Commentary does not specifically mention artificial intelligence or machine learning. It does, however, provide a framework for identifying trade secrets under which a plaintiff needs to provide “[s]pecific, identifying information, such as: the trade secret elements, components, ingredients, steps, *algorithms*, and other specific details the plaintiff contends constitute the trade secret at issue.”⁷⁰ (All AI involves algorithms, though not all algorithms involve AI.)

Regardless of whether courts adopt this particular framework, a plaintiff in an AI algorithm case will need to determine and then describe in sufficient detail how the allegedly misappropriated algorithm works. Often, due to the AI-driven adaptation, the algorithm will operate quite differently than it did when it was initially developed. The need to figure out the inner workings of a previously unknown algorithm places burdens on an AI algorithm trade secret plaintiff that are not present in traditional misappropriation cases. However, it is an eminently reasonable burden to impose on those seeking to use the court system to protect the economic advantages attributable to trade secret AI algorithms.

⁶⁹ The Sedona Conference, *Commentary on the Proper Identification of Asserted Trade Secrets in Misappropriation Cases*, 22 SEDONA CONF. J. 223 (2021).

⁷⁰ *Id.* at 258 (emphasis added).

Consider again the example of the AI-generated algorithm for trading the stock of Airline, Inc., which over time learned which set of indicators were most effective in predicting stock price changes. A plaintiff alleging that the AI-determined set of indicators is the allegedly misappropriated information will need to identify those indicators and describe the manner of their use with enough specificity so that there is a foundation for the parties to argue the key questions in the dispute. Typically, the key questions include: (1) does the information at issue qualify as a trade secret?; (2) what information did the defendant take?; and (3) does the manner in which it was taken constitute misappropriation?

These questions imply, but do not directly state or answer, the broader question of what the standard for disclosure should be in “identifying” a trade secret in litigation. In discussing litigation in patent, trademark, copyright, and design patent cases, Fromer and Lemley have written that “the different approaches IP regimes take to proving infringement are traceable to the different conceptions of the proper audience in each regime”⁷¹—an audience that, for the regimes considered by Fromer and Lemley, could be experts, ordinary observers, or consumers.⁷²

Questions about what constitutes a trade secret and whether it has been misappropriated are best answered from the perspective of people who have skill in the relevant field. This follows from the definitions in both the UTSA and federal law. Both require that a trade secret not be “generally known” or “readily ascertainable” by other persons “*who can obtain economic value from [its/the] disclosure or use*”⁷³—text that implicates the level of knowledge of those working in the relevant field.

A trade secret plaintiff should need to provide a level of detail sufficient for a person of ordinary skill in the relevant field to understand the principles of operation of the algorithm, including

⁷¹ Jeanne C. Fromer & Mark A. Lemley, *The Audience in Intellectual Property Infringement*, 112 MICH. L. REV. 1251, 1303 (2014).

⁷² *Id.* (discussing “divergent views over . . . whether it is the expert, the ordinary observer, or the consumer who is the proper audience in IP infringement”).

⁷³ UTSA § 1(4)(i); 18 U.S.C. § 1839(3)(B) (emphasis added). The UTSA uses “its,” while federal law uses “the.”

the general manner in which it makes decisions, and appreciate that it has (or lacks) value from not being generally known or readily ascertainable to others in the field.⁷⁴ While a “person of ordinary skill in the art” is a concept associated primarily with patent law,⁷⁵ it is relevant to trade secrets as well.⁷⁶ After all, an inquiry of whether information is “generally known” or “readily ascertainable” requires a standard against which to perform those evaluations. And using a person of ordinary skill in the field is certainly a more appropriate standard than using an “ordinary observer” or a “consumer,” neither of whom would typically have the knowledge to obtain economic value from the use of the information.⁷⁷ The concept of skill in the art is also implicitly reflected in the exclusion from trade secret protection of information that constitutes an employee’s general knowledge, skill, and experience. Indeed, one of the key challenges facing courts in misappropriation cases lies in determining whether this exclusion applies.⁷⁸ As Kurt Saunders and Nina Golden have written, some courts, after noting the general knowledge, skill, and experience carveout, simply “categorize[] the information at issue as a trade secret or not,” without providing substantive analysis.⁷⁹ By contrast, “[o]ther courts . . . offer a

⁷⁴ Among other things, disclosure to this level of detail will provide a foundation so that the litigants’ experts can offer opinions.

⁷⁵ See, e.g., U.S. PAT. & TRADEMARK OFF., MANUAL OF PATENT EXAMINING PROCEDURE § 2141.03 (9th ed. Rev. Oct. 2019) (“The person of ordinary skill in the art is a hypothetical person who is presumed to have known the relevant art at the relevant time.”).

⁷⁶ There are, however, differences between how the concept of a person of ordinary skill would apply in patents as opposed to trade secrets. In patent law, the POSITA is hypothetical and presumed to know all relevant art. *Id.* In trade secrets, to assess what is “generally known” or “readily ascertainable,” a better standard would be the knowledge of a typical person of skill in the field, not a hypothetical person who knows all art in the field.

⁷⁷ A possible alternative to a person of ordinary skill in the art would be a “reasonable competitor,” though a reasonable competitor would also have skill in the art, leaving it unclear how these two standards might differ in practice in the context of trade secret litigation.

⁷⁸ See, e.g., Kurt M. Saunders & Nina Golden, *Skill or Secret? — The Line Between Trade Secrets and Employee General Skills and Knowledge*, 15 N.Y.U. J.L. & BUS. 61 (2018).

⁷⁹ *Id.* at 76 (citation omitted).

justification, or at least a perfunctory explanation, for their conclusions.”⁸⁰

A misappropriation plaintiff has access to several mechanisms that prevent the act of providing detailed information about a trade secret from destroying the trade secret. For example, in federal cases, courts commonly issue protective orders to enable parties to share confidential information (including trade secrets) during discovery.⁸¹ Under a protective order in litigation regarding an AI system, a plaintiff’s outside counsel and technical experts can get access to the defendant’s AI source code, training data, test results, and technical documents describing the AI system.⁸² They can also get access to deposition testimony from witnesses employed by the defendant who are knowledgeable about the operation of the AI system.⁸³

Separately, when documents to be filed with a court contain trade secrets and other confidential information, the party making the filing can request leave from the court to file under seal.⁸⁴ In contrast with a protective order, which, once issued by the court, gives parties the ability to designate information as confidential under the order, a request to file information under seal is subject to court approval.

III. AI TRANSPARENCY REQUIREMENTS

AI regulatory proposals and statements of AI principles commonly address transparency, which is intended to make it easier to understand (and challenge) the operation of AI systems.⁸⁵ For

⁸⁰ *Id.* (citation omitted).

⁸¹ FED. R. CIV. P. 26(c).

⁸² A protective order in federal civil litigation can “requir[e] that a trade secret or other confidential research, development, or commercial information not be revealed or be revealed only in a specified way.” *Id.*

⁸³ FED. R. CIV. P. 30(b)(6).

⁸⁴ *See, e.g.*, U.S. DIST. CT, W.D. TEX., ADMINISTRATIVE POLICIES AND PROCEDURES FOR ELECTRONIC FILING IN CIVIL AND CRIMINAL CASES 12 (2016) (“A sealed document in a civil case requires leave of the Court before being filed . . . Sealed documents cannot be electronically accessed by attorneys or the public.”).

⁸⁵ *See, e.g.*, *Transparency and Explainability (Principle 1.3)*, OECD.AI POL’Y OBSERVATORY, <https://oecd.ai/en/dashboards/ai-principles/P7>

instance, in October 2023 the White House issued an executive order on AI underscoring the role of regulatory agencies in “emphasizing or clarifying requirements and expectations related to the transparency of AI models” used by “regulated entities.”⁸⁶ A European Union document explaining the EU AI Act states that “transparency obligations require that AI-based systems must be transparent in their functioning so that users can understand how decisions are taken and the logic behind them.”⁸⁷

To satisfy transparency requirements imposed through legislation, regulation, or “soft law” (voluntary industry self-regulation),⁸⁸ an AI system designer (or owner, operator, etc.) must (1) acquire sufficiently detailed knowledge regarding the system to explain its operation, and (2) disclose that information to the requisite level of detail. Both of these requirements raise challenges in the context of trade secrets.

A. The Challenges to AI System Understanding

1. The Meaning of Transparency

There is also the question of what exactly is meant by transparency. Consider an autonomous car manufacturer that gives a motor vehicle safety agency access to millions of lines of code and the many gigabytes of data the code was trained on. In one sense, this is fully transparent, as the manufacturer has given the safety agency literally all of its computer code and training data. But if it would take large teams of highly trained people months to analyze

[<https://perma.cc/HQ5E-GXAX>] (last visited Jan. 8, 2024) (including the principle “Transparency and explainability” as one of the Organization for Economic Co-operation and Development’s “AI Principles”). Principle 1.3 further states that transparency can “foster a general understanding of AI systems” and “enable those adversely affected by an AI system to challenge its outcome based on plain and easy-to-understand information on the factors, and the logic that served as the basis for the prediction, recommendation or decision.” *Id.*

⁸⁶ Exec. Order No. 14,110, 88 Fed. Reg. 75191, 75214 (Oct. 30, 2023).

⁸⁷ *Key Issues: Transparency Obligations*, EU AI ACT (Nov. 25, 2022), <https://www.euaiact.com/key-issue/5#> [<https://perma.cc/A2H2-NYCR>].

⁸⁸ *See, e.g.*, Carlos Ignacio Gutierrez & Gary Marchant, *How Soft Law is Used in AI Governance*, BROOKINGS INST. (May 27, 2021), <https://www.brookings.edu/articles/how-soft-law-is-used-in-ai-governance/> [<https://perma.cc/KQK5-C6MK>].

the code and data to fully understand the operation and potential vulnerabilities of the vehicle, the disclosure is of little utility.

To address this, discussions of AI governance sometimes invoke the term “explainability” instead of “transparency.”⁸⁹ IBM describes explainable AI as “a set of processes and methods that allows human users to comprehend and trust the results and output created by machine learning algorithms.”⁹⁰ The authors of a 2022 Harvard Business Review article wrote “[e]xplainable AI has to do with how the AI model transforms inputs into outputs; what are the rules? Why did this particular input lead to this particular output?”⁹¹ The authors contrast that with transparency, which “is about everything that happens before and during the production and deployment of the model, whether or not the model has explainable outputs.”⁹²

Regardless of the specific terminology used, as a prerequisite to deciding how to meet AI transparency (or explainability) requirements while still preserving trade secrets, the disclosing party must first find a mechanism to obtain its own understanding of how the AI system is making decisions. One potentially useful approach is to build transparency into an AI system during the entire design process, as opposed to trying to infer it retroactively after the system is complete.⁹³ Greater visibility into the design process would help address a common limitation under which, as Tschider has written, “[a]lthough much academic attention has focused on algorithmic transparency, more attention should be paid to the processes,

⁸⁹ Not all sources clearly articulate a difference between “transparent” AI and “explainable” AI. *See, e.g.*, RICHARD ROOVERS, DELOITTE, TRANSPARENCY AND RESPONSIBILITY IN ARTIFICIAL INTELLIGENCE: A CALL FOR EXPLAINABLE AI 6 (2019) (stating “[t]ransparent AI is explainable AI”).

⁹⁰ *What Is Explainable AI?*, IBM, <https://www.ibm.com/topics/explainable-ai> [<https://perma.cc/SU32-W4WK>] (last visited Oct. 31, 2023).

⁹¹ Reid Blackman & Beena Ammanath, *Building Transparency into AI Projects*, HARV. BUS. REV. (June 20, 2022), <https://hbr.org/2022/06/building-transparency-into-ai-projects> [<https://perma.cc/2FU5-BUT9>].

⁹² *Id.*

⁹³ *Id.* (“[T]ransparency is not something that happens at the end of deploying a model when someone asks about it. Transparency is a chain that travels from the designers to developers to executives who approve deployment to the people it impacts and everyone in between.”).

methods, and strategies used to create algorithmic decision-making systems.”⁹⁴

2. *Size and Adaptivity*

In order for transparency requirements to be effective, AI system designers need to understand how AI algorithms make decisions.⁹⁵ Obtaining that understanding can sometimes be difficult due to inscrutability, which Andrew Selbst and Solon Barocas explain is “a situation in which the rules that govern decision-making are so complex, numerous, and interdependent that they defy practical inspection and resist comprehension.”⁹⁶

One challenge is size. The software used in driverless cars can have hundreds of millions of lines of code,⁹⁷ and can perform many trillions of operations per second.⁹⁸ Applications of AI in domains such as protein folding, which is a powerful tool for drug development, can also involve very large amounts of computation. In March 2023, researchers from Meta published a paper describing the use of machine learning to predict over 600 million protein structures.⁹⁹ The large language models (“LLMs”) that have received enormous public attention since late 2022 following the

⁹⁴ Charlotte A. Tschider, *Beyond the “Black Box”*, 98 DENV. L. REV. 683 (2021).

⁹⁵ There is also a question of what level of understanding is required. In some environments, an approximate understanding may be sufficient. In others (e.g., in evaluating the safety of an AI algorithm used to guide robotic surgery) a precise understanding may be needed.

⁹⁶ Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 FORDHAM L. REV. 1085, 1094 (2018).

⁹⁷ See, e.g., *What Is An Autonomous Car?*, SYNOPSIS, <https://www.synopsys.com/automotive/what-is-autonomous-car.html> [<https://perma.cc/2UYT-NA9F>] (last visited Oct. 30, 2023) (“Today’s cars have 100 million lines of code. Tomorrow’s autonomous cars will have more than 300 million lines of code . . .”).

⁹⁸ See, e.g., Stephen Shankland, *Meet Tesla's Self-Driving Car Computer and its Two AI Brains*, CNET (Aug. 20, 2019), <https://www.cnet.com/tech/computing/meet-tesla-self-driving-car-computer-and-its-two-ai-brains/> [<https://perma.cc/Y58Q-L3WC>] (stating “[e]ach Tesla AI chip runs at 2GHz and performs 36 trillion operations per second”).

⁹⁹ Justin Jackson, *Predicting Protein Folding from Single Sequences with Meta AI ESM-2*, PHYS. ORG. (Mar. 23, 2023), <https://phys.org/news/2023-03-protein-sequences-meta-ai-esm-.html> [<https://perma.cc/MR6P-DFBX>].

release of OpenAI's ChatGPT are highly complex, potentially containing over 100 billion parameters.¹⁰⁰

Another attribute is adaptivity, which can lead to what Tschider describes as “dynamic inscrutability.”¹⁰¹ AI systems often evolve their decision-making as they gain greater experience. This adaptivity can happen slowly, over time scales involving days or weeks, or quickly, with time scales measured in tiny fractions of a second. In the financial markets, high frequency trading algorithms based on machine learning can operate on scales of microseconds,¹⁰² and in 2022 the U.S. Department of Defense's Defense Advanced Research Projects Agency announced a program to develop AI processors that can “self-reconfigure within 50 nanoseconds.”¹⁰³ The combination of size and adaptivity can make it very difficult for even an AI system's own designers to understand the specifics of how an AI system makes a particular decision.

3. *The AI Supply Chain*

If that were not complicated enough, there is also the AI supply chain, which will become increasingly labyrinthine as the AI ecosystem continues to mature. While there is nothing new about complex supply chains, an example regarding vehicles illustrates how the addition of AI to supply chains can increase opacity. Manufacturer transparency obligations associated with (largely) non-autonomous cars do not generally involve descriptions of algorithms. A prospective buyer seeks transparency in relation to features such as the car's general safety features, the power of the engine, and the distance the car can travel on one tank of gas or on a full battery charge. But the buyer does not demand a detailed

¹⁰⁰ See Tom Brown et al., *Language Models Are Few-Shot Learners* 46 (July 22, 2020) (unpublished manuscript) (on file with arXiv) (stating in Table D.1 that the model “GPT-3 175B” contains 174.6 billion parameters).

¹⁰¹ Tschider, *supra* note 94, at 690.

¹⁰² Jasmina Arifovic et al., *Machine Learning and Speed in High-Frequency Trading*, 139 J. ECON. DYNAMICS & CONTROL 1, 1 (2022) (“[T]he speed at which HFT is conducted” has progressed “from a scale of milliseconds to microseconds.”).

¹⁰³ *DARPA Eyes Adaptive, Real-Time Processors for Future AI-Enabled Radios*, DEF. ADVANCED RSCH. PROJECTS AGENCY (Oct. 6, 2022), <https://www.darpa.mil/news-events/2022-10-06> [<https://perma.cc/LV8P-CXEY>].

explanation of exactly how the algorithm used in antilock braking operates.

Now consider what happens when AI is brought into the mix. For autonomous and semi-autonomous vehicles, a regulatory agency may indeed want to know at least some level of detail regarding how AI makes decisions on when and how the brakes are applied. But the manufacturer may not have easy access to that information, having instead licensed the AI software module for braking from an upstream braking module supplier. That supplier may have delivered the braking software in binary “executable” form appropriate for a computer but unreadable to a human.¹⁰⁴ The supplier may also have taken additional steps to prevent its customers from learning the details of the braking algorithm, e.g., by requiring that licensees sign non-disclosure agreements in relation to any information about the algorithm they receive from the supplier. The supplier may also require licensees to warrant that they will not engage in any reverse engineering of the algorithm.

Complying with transparency requirements will require providers of AI systems (which in the case of an autonomous vehicle is the vehicle’s manufacturer) to demand more visibility into the inner workings of the components provided by their upstream suppliers. The provider of the AI-based end product or service—i.e., the company that does the final integration of the component parts and brings a completed product or service to market—will be responsible for understanding what they are marketing to a sufficient level of detail to meet transparency requirements. Additionally, or alternatively, regulating agencies themselves may require that upstream suppliers for safety-critical functions (e.g., vehicle braking) directly provide transparency disclosures to the agencies.

B. Meeting AI Transparency Requirements

There are many domains where AI transparency requirements will be important. As noted earlier, manufacturers of driverless cars (or of AI-based components used in driverless cars) may need to disclose the details of algorithms to vehicle safety regulators.

¹⁰⁴ Compiled (also referred to as binary or executable) computer code can be reverse engineered, but the process is time consuming and expensive.

Banking or housing regulators might require transparency to ensure that AI systems used in making loan decisions comply with banking and housing antidiscrimination laws. The Federal Trade Commission might require transparency disclosures to ensure compliance with consumer privacy protections. Transparency mandates may also arise in relation to government use of algorithms supplied by private companies for use in criminal risk assessments.¹⁰⁵ The following questions can help companies, regulatory agencies, and civil society groups better frame policies and compliance metrics arising from transparency requirements:

1. Is the Information at Issue Really a Trade Secret?

There is a temptation for companies to be overinclusive when considering the subset of their confidential information to view as trade secret information. Companies may consider overinclusion as the least risky approach, under the view that failing to treat a trade secret as such is more harmful than improperly designating information as a trade secret.

But overinclusion has real costs. It erects unnecessary obstacles to employee mobility and to transparency, both of which are important public policy goals. In addition, when a company incorrectly asserts that something is a trade secret, the assertion itself can create a presumption that can be difficult and impractical to rebut. The tendencies towards overinclusion will likely be particularly acute in AI—a domain where confidential information will sometimes be information that is factually secret, but that does not qualify as a trade secret.¹⁰⁶

2. Can the Discloser Elect to Use Trade Secrets More Narrowly?

Just because a company *can* legitimately claim trade secret rights in a particular set of information does not mean that it always should. In assessing their trade secrets, companies may not properly distinguish “value” in the statutory definition of a trade secret from the value of the investment made in developing the information at issue. Under common law, there was indeed a connection between

¹⁰⁵ Elizabeth A. Rowe & Nyja Prior, *Procuring Algorithmic Transparency*, 74 ALA. L. REV. 303, 307 (2022).

¹⁰⁶ Sandeen & Aplin, *supra* note 13, at 443.

investment and trade secret status. One of the six factors listed in the First Restatement (and still used in some courts today)¹⁰⁷ to evaluate whether information is a trade secret was “the amount of effort or money expended by him in developing the information.”¹⁰⁸ But the UTSA and DTSA do not recite this test. (There are interesting policy questions regarding whether investment should nonetheless be considered; Joseph Fishman and Deepa Varadarajan have written that “[w]hile tying trade secret protection to development cost has a long pedigree at common law, it doesn’t get the attention it deserves today because it’s not mentioned in any governing statute.”¹⁰⁹)

Despite the lack of statutory support, there is a tendency for companies to view all trade secrets obtained at high cost as presumptively being high-value trade secrets. But that is not necessarily the case. While some valuable trade secrets will be developed only after great expense, a valuable trade secret can sometimes be created with modest investment, and conversely, high amounts of effort can be expended without producing any trade secrets.

An additional concern regarding overbreadth is the use of contract law to prohibit activities that trade secret law permits, thereby drawing a cloak around information that would otherwise be exposed to discovery by others. For instance, while reverse engineering is a proper means to ascertain a trade secret,¹¹⁰ as Varadarajan writes, “[c]onsumer software licenses often contain broad prohibitions against product disassembly, decompiling, and other forms of reverse engineering.”¹¹¹ In addition, as Camilla Hrdy and Christopher Seaman observe, non-disclosure agreements often encompass “far more information than trade secret law does—

¹⁰⁷ RESTATEMENT (FIRST) OF TORTS § 757 cmt. b. (AM. L. INST. 1939).

¹⁰⁸ *Id.*

¹⁰⁹ Joseph P. Fishman & Deepa Varadarajan, *Earning Trade Secrets*, 109 CORNELL L. REV. (forthcoming 2024) (manuscript at 1) (on file with SSRN).

¹¹⁰ See, e.g., UTSA § 1 cmt. (UNIF. L. COMM’N 1985) (“Proper means include . . . Discovery by ‘reverse engineering’”); *Kewanee Oil Co. v. Bicron Corp.*, 416 U.S. 470, 476 (1974) (citing RESTATEMENT (FIRST) OF TORTS § 757(a)); 18 U.S.C. § 1839(6) (“[I]mproper means’ . . . does not include reverse engineering.”).

¹¹¹ Deepa Varadarajan, *The Trade Secret-Contract Interface*, 103 IOWA L. REV. 1543, 1569 (2018) (citation omitted).

including publicly available or generally known information, and information that trade secret law would classify as unprotectable ‘general knowledge, skill, and experience.’ ”¹¹²

This is not to suggest that companies that make AI systems should freely waive their trade secret rights in the most valuable trade secrets that bring them a competitive advantage and the resulting commercial success. But because there is little disincentive to avoid over-designation, companies often adopt a maximalist view of trade secrets under which anything that can plausibly be argued (and often that cannot be plausibly argued) to be a trade secret is treated as such. In the context of AI systems, where there will be both formal transparency obligations and pressure from civil society organizations and customers to disclose how AI systems operate, a company’s own longer-term business interests as well as the public’s interest can be served if companies making AI systems adopt a more balanced, less maximalist view of deciding which information to treat as trade secrets.

While it might seem impractical to ask companies to voluntarily disclose any more information than necessary, the sheer complexity of many modern AI systems will give companies multiple options in terms of the levels of abstraction to use in a description. This makes it possible for a company to disclose large amounts of information about how an AI system operates while also maintaining as trade secrets the most specific information that can be the main source of competitive advantage.

For instance, a company developing an LLM for generating text outputs might disclose the training data, the model size, information about the model design, and the details of how the company performs “alignment”—i.e., tuning the model so that its outputs align with human values.¹¹³ As part of the alignment process, the company could also invite outsiders to “red-team” their systems and

¹¹² Camilla A. Hrdy & Christopher B. Seaman, *Beyond Trade Secrecy: Confidentiality Agreements That Act Like Noncompetes*, 133 YALE L.J. 669 (2024).

¹¹³ See Sharon K. Sandeen, *A Typology of Disclosure*, 54 AKRON L. REV. 657 (2021) (discussing different types of “disclosure,” many of which do not involve the loss of trade secrets).

thereby help identify and mitigate potential harmful outputs. This could include publicly describing the flaws that were found and the particular steps that were taken to address them. All of this disclosure would still allow the company to maintain as trade secrets the innermost details of the LLM, such as the specific parameter values used within the model.

3. *Can Public Disclosures Avoid Trade Secrets?*

Not all AI transparency obligations will implicate trade secrets. If the transparency requirements call for a high-level description of how an AI system works, the disclosing party will often be able to avoid touching on trade secrets. Even a relatively detailed description of an AI system will often avoid trade secrets. Consider again the paper describing the design and training of AlphaZero, the AI system that “[s]tarting from random play, and given no domain knowledge except the game rules, . . . achieved within 24 hours a superhuman level of play in the games of chess and shogi (Japanese chess) as well as Go, and convincingly defeated a world-champion program in each case.”¹¹⁴ While there is certainly plenty of information regarding AlphaZero that remains a trade secret, that did not prevent Google from providing a very substantive disclosure regarding its operation. Similarly, in 2020, OpenAI researchers published a detailed description of their work on large language models.¹¹⁵

Additionally, some transparency obligations might operate in the negative, allowing the resulting disclosures to steer clear of trade secrets. A real estate company that is using AI to help make decisions on loan applications might be required to certify that the system is *not* relying on any data from a list of data types that can serve as proxies for categories protected from discrimination under the Fair Housing Act.¹¹⁶ Other than by confirming the exclusions, making that certification would say little about the types of data the

¹¹⁴ David Silver et al., *Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm 1* (Dec. 5, 2017) (unpublished manuscript) (on file with arXiv).

¹¹⁵ See Brown et al., *supra* note 100.

¹¹⁶ 42 U.S.C. § 3605(a) (prohibiting discrimination in relation to loans “because of race, color, religion, sex, handicap, familial status, or national origin”).

AI system *does* use and would say nothing at all regarding the specific manner in which those data are used.

In some instances, transparency requirements might implicate aspects of an AI system that the system’s owners do not consider confidential. An AI governance framework designed to mitigate bias might require disclosure of the training data used—something that an AI system owner might not consider a trade secret.

4. *Regulatory Disclosures and Preservation of Trade Secrets*

The tensions between mandatory disclosures to regulatory agencies and trade secrets have long been recognized.¹¹⁷ Frameworks for providing government agencies often allow some of the submitted information to be designated as confidential, though whether the government will agree with those designations is a different question.

Relatedly, the Freedom of Information Act (“FOIA”) contains an exemption for “trade secrets and commercial or financial information obtained from a person and privileged or confidential.”¹¹⁸ The Supreme Court addressed the meaning of “confidential” in 2019 in *Food Marketing Institute v. Argus Leader Media*,¹¹⁹ concluding that when the FOIA was enacted in 1966, confidential “meant then, as it does now, ‘private’ or ‘secret.’ ”¹²⁰ The Court held that “[a]t least where commercial or financial information is both customarily and actually treated as private by its owner and provided to the government under an assurance of privacy, the information is ‘confidential.’ ”¹²¹ Thus, in situations where the government does provide that assurance, a company’s trade secret information provided to the government under a confidentiality designation remains protected. While this standard will maintain the confidentiality of information that is genuinely a trade secret, it will also, as Varadarajan has written, “dramatically

¹¹⁷ See, e.g., Elizabeth A. Rowe, *Striking a Balance: When Should Trade-Secret Law Shield Disclosures to the Government?*, 96 IOWA L. REV. 791 (2010).

¹¹⁸ 5 U.S.C. § 552(b)(4).

¹¹⁹ *Food Mktg. Inst. v. Argus Leader Media*, 139 S. Ct. 2356 (2019).

¹²⁰ *Id.* at 2363 (citing WEBSTER’S SEVENTH NEW COLLEGIATE DICTIONARY 174 (1963)).

¹²¹ *Id.* at 2366.

expand the private sector's ability to shield from public view information provided to the government."¹²²

And what happens when the government agency collects trade secret information but *declines* to provide an assurance of privacy? While the *Argus Leader* Court did not reach this question,¹²³ it will frequently arise in regulatory disclosures regarding AI systems. In this regard, an example involving autonomous vehicles is instructive. In August 2021, the National Highway Traffic Safety Administration ("NHTSA") issued an order requiring manufacturers and operators of automated vehicles to "report crashes to the agency."¹²⁴ The order allows the party submitting a report to designate certain portions of the report as containing confidential business information ("CBI"),¹²⁵ which is defined to include trade secrets.¹²⁶ However, the order also states that "[m]aking a request for confidential treatment does not ensure that the information claimed to be confidential will be determined to be confidential."¹²⁷

This disclaimer reflects a tension that will be common in government regulation of AI systems. On the one hand, knowing that some companies will attempt to apply confidentiality designations overly broadly, agencies will be understandably reluctant to commit in advance to a company's own determinations of what information should be kept from public view. On the other hand, companies will be reluctant to provide detailed information

¹²² Deepa Varadarajan, *Business Secrecy Expansion and FOIA*, 68 UCLA L. REV. 462, 462 (2021).

¹²³ See *Argus Leader*, 139 S. Ct. at 2366. (stating that there is "no need to resolve that question in this case").

¹²⁴ U.S. DEP'T OF TRANSP., NAT'L HIGHWAY TRAFFIC SAFETY ADMIN., FIRST AMENDED STANDING GEN. ORD. 2021-01, INCIDENT REPORTING FOR AUTOMATED DRIVING SYSTEMS (ADS) AND LEVEL 2 ADVANCED DRIVER ASSISTANCE SYSTEMS (ADAS) 2 (2021). A second amended version of this order was published in 2023, and with respect to the excerpts cited herein is unchanged.

¹²⁵ *Id.* at 11 (identifying three categories of information that can be designated CBI).

¹²⁶ *Id.* at 12 (citing 49 C.F.R. § 512(D)–(E)). The cited section states: "Confidential business information means trade secrets or commercial or financial information that is privileged or confidential." *Id.* § 512(3)(c).

¹²⁷ *Id.* at 11–12.

about trade secrets if they believe that a regulator may unilaterally decide to override confidentiality designations. Things will get even more complex to the extent that future regulations require a company to provide source code for an AI system to a government agency. As Sonia Katyal has written, source code “remains one of the few spheres to enjoy concurrent protections from trade secrecy, copyright law, and patent law.”¹²⁸ If a company submits source code to a regulatory agency embodying a trade secret method that the company later intends to patent, the agency’s subsequent handling of that information could implicate not only trade secret status but also patentability.¹²⁹

Despite these tensions, there is cause for optimism, as both the government and companies will face pressure to act reasonably. An agency that does not make a good faith effort to honor reasonable confidentiality designations will find that companies are less forthcoming in describing their AI systems, slowing the pace of regulatory approvals, and leading to criticism that the agency is impeding access to AI innovations. A company that inappropriately designates information as confidential, and then pursues litigation against an agency for not honoring the designation will risk reputational damage.

5. *Would the Information Be Better Protected Through Patents?*

A trade secret can cover a broad range of information, including patentable inventions.¹³⁰ For the subset of trade secrets that are patentable, this creates an option: inventors can choose to file a patent application and in doing so forego trade secret rights in exchange for disclosing the invention and seeking a time-limited

¹²⁸ Sonia K. Katyal, *The Paradox of Source Code Secrecy*, 104 CORNELL L. REV. 1183, 1190 (2019).

¹²⁹ While U.S. patent law provides a one-year grace period for certain disclosures in advance of a patent filing (*see* 35 U.S.C. § 102(b)(1)), foreign jurisdictions generally do not.

¹³⁰ *See, e.g., Kewanee Oil Co. v. Bicron Corp.*, 416 U.S. 470, 490 (1974) (“We conclude that the extension of trade secret protection to clearly patentable inventions does not conflict with the patent policy of disclosure.”).

right to exclude others from practicing it.¹³¹ Alternatively, they can choose to maintain the invention as a trade secret, avoiding the public disclosure required for obtaining a patent. However, this runs the risk (among other risks such as reverse engineering and accidental disclosure) that others will independently develop the same trade secret, and perhaps choose to publicly disclose it through a patent application or other means.¹³²

For patent-eligible trade secrets, companies have long needed to make the choice between filing a patent application or maintaining the trade secret.¹³³ Among other considerations, this will involve evaluating the anticipated relative value of each approach.¹³⁴ Enforceability, which is easier with patents given that they are public and presumptively valid, will also be a consideration.

When the information in question is an AI algorithm, it is particularly important to consider patent eligibility. While the Supreme Court has never addressed the patentability of AI algorithms specifically, it has addressed computer algorithms more generally. In 2014 in *Alice Corp. v. CLS Bank International*,¹³⁵ the Court considered the patentability of “a computer-implemented scheme for mitigating ‘settlement risk’ (i.e., the risk that only one party to a financial transaction will pay what it owes) by using a third-party intermediary.”¹³⁶

Building on its 2012 decision in *Mayo Collaborative Services v. Prometheus Labs*,¹³⁷ the *Alice* Court described a two-step process for determining patent eligibility. The first step is to “determine

¹³¹ A patent, once issued, gives its owner a right to exclude others from making, using, selling or offering to sell, or importing the invention into the United States. See 35 U.S.C. § 271(a).

¹³² Absent a non-publication request, patent applications are typically published after 18 months. See 37 C.F.R. § 1.213 (2000).

¹³³ For simplicity, this discussion assumes that the trade secret could be converted into a single patent. Of course, other variations are possible. A trade secret could lead to multiple patents. Or, a single patent could embody multiple former trade secrets.

¹³⁴ See, e.g., Jonas Anderson, *Secret Inventions*, 26 BERKELEY TECH. L.J. 917 (2011).

¹³⁵ *Alice Corp. v. CLS Bank Int'l*, 573 U.S. 208 (2014).

¹³⁶ *Id.* at 212.

¹³⁷ *Mayo Collaborative Servs. v. Prometheus Labs*, 566 U.S. 66 (2012).

whether the claims at issue are directed to a patent-ineligible concept.”¹³⁸ If the claims are directed to a patent-eligible concept, the inquiry ends there. However, if the claims are directed to a patent-ineligible concept, then the inquiry proceeds to the second step, which is to “examine the elements of the claim to determine whether it contains an ‘inventive concept’ sufficient to ‘transform’ the claimed abstract idea into a patent-eligible application.”¹³⁹

No such test exists for trade secret eligibility, which instead depends on whether the information has economic value through not being generally known or readily ascertainable and is subject to reasonable efforts to maintain its secrecy. In situations where an AI-related trade secret is of questionable patent eligibility, the company that developed it might be more likely to opt to protect it as a trade secret rather than disclosing it through filing a patent application that will generally be published after eighteen months but that might ultimately be unsuccessful.¹⁴⁰ If the patent application is denied, the company will have lost the trade secret (due to the published patent application) while also failing to obtain a patent.

This risk, plus the fact that a company may face a significant burden to even identify what an AI algorithm is doing in sufficient detail to describe it in a patent application, will often push companies to favor trade secrets as opposed to patents. The unfortunate consequence is that many AI algorithms will remain unavailable to the public—including to those who would be able to improve them and apply them in ways not foreseen by their original developers.

¹³⁸ *Alice Corp.*, 573 U.S. at 218.

¹³⁹ *Id.* at 221 (citing *Mayo Collaborative Services*, 566 U.S. at 72, 80). In addition, the U.S. Patent and Trademark Office has revised the Manual of Patent Examining Procedure to include a description of the post-*Alice* process for determining patent eligibility, including a flowchart. See U.S. PAT. & TRADEMARK OFF., MANUAL OF PATENT EXAMINING PROCEDURE § 2106 (9th ed., rev. Oct. 2019).

¹⁴⁰ As noted earlier, unless there is a non-publication request, patent applications are typically published after 18 months. See 37 C.F.R. § 1.213 (2000). Once that publication occurs, there is no way to restore trade secret status for the published information.

That said, if an AI system includes a sufficient “inventive concept” to pass the *Alice* test, a company may elect to pursue the patent route. If the patent application is filed prior to making a mandated transparency disclosure pursuant to AI governance regulations, the act of making the regulatory disclosure will not alter the status of patent rights secured through the prior filing of the application. There is always a risk that the U.S. Patent and Trademark Office (“PTO”) will decline to issue the patent application. But if the PTO does grant a patent based on the application, then the company will have maintained rights in the intellectual property, albeit through a patent instead of a trade secret.

IV. CONCLUSION

Trade secrets have always played an important role in relation to technology, and AI systems are no exception. AI raises a set of unique challenges due to its adaptivity and complexity, which have implications for understanding the scope of trade secrets, pleading a misappropriation claim, and in relation to the disclosure obligations that will accompany AI regulatory frameworks.

This Article has argued that AI system designers can hold trade secret rights to the algorithms embodied in their AI systems, including when designers do not know how those algorithms work. The Article also discussed litigating misappropriation cases. For a misappropriation allegation involving the algorithm used in a complex and highly adapted AI system, initiating and then litigating a complaint will require plaintiffs to fill in gaps in their own knowledge of how the system operates.

The Article has also examined a series of approaches that can help companies and policymakers address the potential tension between disclosure requirements and trade secret status. These include avoiding a maximalist view of trade secrets and providing disclosures through mechanisms such as confidentiality designations in filings with regulatory agencies that aim to preserve trade secret status. When public disclosure is required, the complexity of AI systems offers opportunities for companies to provide sufficient detail to meet the public policy goals underlying the disclosure mandates while still preserving trade secrets.

As AI systems continue to gain adoption, fact patterns where AI system designers are unaware of the details of how their algorithms operate will become increasingly common. Experience over the coming years will clarify the standards and mechanisms for addressing this knowledge gap when litigating misappropriation cases and when complying with AI transparency regulations. There is every reason to believe that trade secret law—and more generally, intellectual property law—will play a vital role in AI in the future, just as it has for so many other technologies in the past.