



AI AND TRADE SECRETS: PROTECTING CONFIDENTIAL INFORMATION IN THE ERA OF INNOVATIVE TECHNOLOGIES

<https://doi.org/10.5281/zenodo.19103458>

Navruzбек Tilaboev

University of Illinois at Urbana-Champaign

Law College

nova.tilaboyev@gmail.com

Abstract: *The collision between artificial intelligence and trade secret law is not a future problem. It is happening now, in boardrooms, courtrooms, and the daily workflows of engineers who paste proprietary code into chatbots without a second thought. This article argues that existing legal frameworks, while structurally sound, are straining under pressures their drafters could not have anticipated. Drawing on the Defend Trade Secrets Act, the Uniform Trade Secrets Act, and the EU Trade Secrets Directive, the analysis maps where doctrine holds, where it bends, and where it is starting to break. Landmark cases are examined alongside regulatory developments and the emerging forensic challenges that make AI-related misappropriation claims uniquely difficult to prove. The article concludes with practical guidance for counsel navigating this terrain today, not waiting for the law to catch up.*

Keywords: *Trade secrets; artificial intelligence; Defend Trade Secrets Act; confidential information; machine learning; data privacy; intellectual property; misappropriation; generative AI; corporate governance; EU Trade Secrets Directive; technology law*

I. INTRODUCTION

Start with a thought experiment. Your company has spent four years and tens of millions of dollars developing a proprietary model for predicting equipment failure in industrial facilities. The model is your edge. Competitors have tried and failed to replicate it. Then, one afternoon, a well-meaning engineer pastes a key section of the underlying codebase into a commercial AI assistant to help debug a loop. The platform's terms of service, buried in paragraph

nineteen of a document nobody reads, permit the provider to use submitted content to improve its systems. Within months, fragments of your logic are embedded in a model that your fiercest rival now licenses.

This is not a hypothetical designed to provoke anxiety. Variants of this scenario are playing out across industries, and the legal system is only beginning to develop tools to address them. Trade secret law, which has served as the primary vehicle for protecting



confidential business information for well over a century, was never designed for an environment where sensitive data flows through third-party AI platforms, where machine learning models can absorb and reproduce proprietary content without any overt act of copying, and where the line between legitimate reverse engineering and misappropriation has become genuinely blurry.¹

What makes this moment particularly complicated is that the threat is not coming primarily from bad actors. It is coming from convenience. The engineers, lawyers, analysts, and executives uploading sensitive material to AI tools are not, for the most part, trying to steal anything or give anything away. They are trying to work faster. That makes the problem harder, not easier, to manage through traditional legal mechanisms that assume intentional wrongdoing or at least recklessness.²

This article works through these challenges in layers. Part II surveys the legal architecture for trade secret protection, with particular attention to the federal Defend Trade Secrets Act and its interaction with state law and international frameworks. Part III examines, with some specificity, how AI systems create new and largely unaddressed vulnerabilities within that architecture. Part IV looks at the litigation landscape, including cases that have already started to test these boundaries. Part V addresses what companies should actually be doing right now. Part VI

offers a candid assessment of where the law needs to go.

II. The Legal Framework for Trade Secret Protection

A. The Defend Trade Secrets Act of 2016

For most of American legal history, trade secret protection was a state law matter. The result was a patchwork: forty-eight states eventually adopted some version of the Uniform Trade Secrets Act, but with enough variation in definitions, remedies, and procedural rules that multistate disputes became genuinely complicated. The Defend Trade Secrets Act, signed into law on May 11, 2016, changed that by creating the first federal civil cause of action for trade secret misappropriation.³

The DTSA defines a trade secret expansively, covering "all forms and types of financial, business, scientific, technical, economic, or engineering information" that derives independent economic value from remaining secret and that the owner has taken reasonable measures to protect.⁴ That breadth was intentional. Congress understood it was legislating for a rapidly evolving technological environment, and deliberately avoided narrower definitions that might quickly become outdated.

In practical terms, the DTSA offers meaningful tools. Plaintiffs can seek injunctive relief, compensatory damages, and, where misappropriation is willful and malicious, exemplary damages up to twice the compensatory amount plus attorney's fees. The statute also provides,



in genuinely extraordinary circumstances, for ex parte civil seizure orders, allowing courts to order the seizure of misappropriated property before the defendant even has notice.⁵ That last remedy is rarely granted, but its existence signals the seriousness with which Congress viewed trade secret theft.

B. The Uniform Trade Secrets Act and State Law

The DTSA did not preempt state law; it supplemented it. Most trade secret litigation still involves state claims alongside or instead of federal ones, and the UTSA framework, first promulgated in 1979, remains the dominant state-level architecture.⁶ The UTSA's definition of misappropriation captures both the improper acquisition of a trade secret and its use or disclosure by someone who knew, or had reason to know, that the information was acquired improperly or under circumstances creating a duty of confidentiality.⁷

That "reason to know" standard deserves emphasis. It has real bite in AI contexts, where a company that builds a model on scraped or otherwise improperly sourced data may face liability even if it was not the original wrongdoer. Whether courts will apply this standard aggressively to AI developers remains to be seen, but the statutory language provides at least a foothold for such claims.

C. The International Dimension

Trade secret law is not, of course, a domestic affair. The TRIPS Agreement requires all WTO member nations to

protect undisclosed information meeting the basic criteria of secrecy, commercial value, and reasonable protective measures.⁸ In practice, enforcement varies dramatically across jurisdictions, and the global nature of AI development, with models trained in one country, hosted in another, and accessed in a third, creates genuine cross-border complexity.

The European Union's Trade Secrets Directive (2016/943/EU) brought meaningful harmonization to EU member states, establishing common definitions and remedies while explicitly preserving the right to acquire trade secrets through independent research or reverse engineering.⁹ That carve-out, entirely sensible in a pre-AI world, has become one of the more contested features of the Directive in an era where AI-assisted reverse engineering can accomplish in weeks what once took years of human analysis. Several EU member states have transposed the Directive into national law, but interpretive divergence remains, particularly around what constitutes "adequate" protective measures in the context of AI systems.

III. Where Artificial Intelligence Strains the Framework

A. The Memorization Problem

Here is something that surprised a lot of lawyers when researchers first demonstrated it clearly: large language models can memorize specific text from their training data and reproduce it verbatim when prompted the right way. Not paraphrase it. Not approximate it. Reproduce it, word for word.¹⁰



The implications for trade secret law are serious and underappreciated. When a company uses proprietary data — customer records, internal research, financial projections, source code — to train or fine-tune an AI model, that information may become embedded in the model's parameters. It is not stored the way a file is stored, but it can sometimes be extracted through carefully crafted prompts, a technique researchers have demonstrated with unsettling effectiveness even on commercially deployed models.¹¹

Now consider what happens when that model is shared, licensed, or accessed by an unauthorized third party. Has the company disclosed its trade secrets? The argument that it has is not frivolous. And yet existing statutory frameworks offer no clear answer, because they were written with the assumption that disclosure requires some identifiable act of transmission. The law has simply not caught up with a form of information leakage that looks, from the outside, like nothing at all.

B. The Casual Disclosure Problem

If the memorization problem is subtle and technical, this one is neither. The single largest trade secret risk created by AI, at least right now, is that employees are uploading confidential information to third-party platforms every day, without fully understanding what happens to it once it leaves their screen.¹²

The Samsung incident of 2023 is the most widely cited example. Employees reportedly uploaded proprietary

semiconductor source code and the contents of internal meetings to ChatGPT, apparently for entirely mundane productivity purposes. When Samsung discovered what had happened, the company moved quickly to restrict external AI use, but the information had already been transmitted to OpenAI's servers.¹³ Whether any of that information was retained, used in training, or otherwise compromised remains unclear. The point is that the disclosure happened casually, without malice, and without any of the red flags that traditional trade secret monitoring is designed to catch.

Samsung is not unique. It is simply the company where the story became public. The same pattern is repeating itself in law firms, pharmaceutical companies, financial institutions, and technology startups, anywhere people have adopted AI tools faster than their legal and compliance functions have had time to develop governing policies. Which, at the moment, is nearly everywhere.

C. AI-Powered Reverse Engineering

The UTSA and the EU Directive both permit reverse engineering as a legitimate means of acquiring a trade secret. That was a deliberate policy choice, and generally a sound one: prohibiting reverse engineering would significantly chill competition and innovation. But the calculus shifts when AI systems can perform reverse



engineering at a scale and speed that no human analyst could approach.¹⁴

A competitor with access to your product's observable outputs, public patent filings, research publications, and market behavior can now run machine learning analysis that identifies patterns suggesting your underlying methodology with a degree of confidence that would have been impossible a decade ago. Whether this is still "reverse engineering" in the sense the law had in mind, or whether it is something else that deserves different treatment, is a genuinely open question. Courts have not yet had to answer it directly, but they will.

D. Ownership of AI-Generated Secrets

One more complication, distinct from the others. What happens when an AI system itself generates information that qualifies as a trade secret, a novel algorithm, a non-obvious business insight, a synthetic dataset with genuine commercial value? Who owns that information, and can it be protected?¹⁵

Trade secret doctrine has always assumed a human owner taking active steps to maintain confidentiality. As AI generation becomes more autonomous and the human role in producing valuable outputs decreases, that assumption is increasingly strained. Courts will need to address whether AI-generated information can meet the "reasonable measures" requirement when the primary actor maintaining its secrecy is, functionally, an automated system. No clear answer exists yet, and the stakes,

given the volume of commercially valuable information now being generated by AI, are substantial.

IV. The Litigation Landscape

A. Waymo LLC v. Uber Technologies, Inc.

Before AI-specific trade secret claims had fully arrived, the technology sector had already produced one of the most consequential trade secret cases in modern legal history. Waymo's suit against Uber alleged that Anthony Levandowski, a former Google engineer, had downloaded roughly 14,000 confidential files before departing to found a self-driving startup that Uber subsequently acquired.¹⁶ The alleged secrets included proprietary LiDAR designs central to Waymo's autonomous vehicle development.

The case settled for approximately \$245 million in Uber equity before going to the jury. But its procedural legacy matters as much as its financial terms. The litigation established important precedents for forensic discovery in technology cases, clarified the evidentiary standards for demonstrating access and use in technical trade secret disputes, and demonstrated both the potential magnitude of damages in the sector and the extraordinary complexity of proving misappropriation when the alleged wrongdoing involves sophisticated engineering systems. Those lessons translate directly to AI-era cases, where the technical complexity is even greater.

B. New Disputes on the Horizon



The next generation of trade secret litigation is beginning to take shape, and it looks different from anything the courts have seen before. Cases have emerged in the Northern District of California and the Southern District of New York raising questions about whether the use of proprietary data to train commercial AI models, without the data owner's authorization, constitutes misappropriation under the DTSA.¹⁷ These are not easy cases to litigate.

The evidentiary challenges are formidable. Demonstrating that a specific trade secret was incorporated into an AI model's training data requires forensic methodologies that are still being developed. The scale of modern training datasets makes comprehensive discovery practically difficult. And the opacity of many commercial AI systems means that even the developers themselves may struggle to explain precisely what their models have learned from which inputs. Courts are being asked to apply legal standards designed for the world of file transfers and employee departures to a technological environment they were never designed to address.

C. Regulatory Pressure

Regulators have noticed. The Federal Trade Commission has signaled that the unauthorized use of business data to train AI models may constitute unfair or deceptive practices under Section 5 of the FTC Act, and has called for rulemaking to address AI-related data practices more specifically.¹⁸ The EU's AI Act, which entered into force in 2024,

includes transparency requirements around training data that may generate disclosure obligations with direct relevance to trade secret claims.

At the state level, the picture is fragmented but moving. California, as so often, is leading, with proposals addressing AI use in employment contexts and the confidentiality obligations that attach to AI-processed data. The SEC's 2023 cybersecurity disclosure rules create reporting obligations for material incidents involving proprietary AI systems or training data, effectively forcing public companies to treat certain AI-related trade secret breaches as disclosure events.¹⁹ The direction of travel is clear: regulatory frameworks are converging on AI-related data practices from multiple directions simultaneously.

V. What Companies Should Actually Be Doing

A. Write the Policy — and Mean It

The starting point is obvious but widely neglected: every organization that handles competitively sensitive information needs a clear, written policy governing the use of AI tools. This does not mean a blanket prohibition, which will be ignored, or a policy so vague it provides no guidance. It means a document that specifies which categories of information cannot be uploaded to external platforms, which tools have been vetted and approved, and what the reporting procedure is when someone realizes they may have made a mistake.²⁰



The hardest part is not writing the policy. It is enforcement and culture. A policy that employees view as an obstacle rather than a reasonable protection will be circumvented systematically, and in litigation that circumvention will be used to argue that the company failed to take reasonable measures. Legal counsel needs to work with technology and people teams to build something that employees actually understand and follow. That is a different kind of drafting problem than most lawyers are used to.

B. Contractual Architecture

AI governance is not a single-contract problem. It runs through the entire contracting ecosystem. Vendor agreements with AI platform providers need explicit representations about data handling, clear prohibitions on using submitted content for model training without consent, and meaningful indemnification provisions. Employment and consulting agreements need updated confidentiality clauses that specifically address AI tools and platforms, not just the generic language about not disclosing proprietary information that was written in a different era.²¹

Non-disclosure agreements with partners and counterparties should expressly address the AI context: what systems may or may not be used to process shared information, and what happens if a breach occurs through an AI platform. These are not exotic provisions. They are the kind of practical protections that should be standard in any deal involving sensitive commercial

information, and they are not yet standard. That gap represents real exposure.

C. Technical Controls

Legal frameworks are only as effective as the technical infrastructure that supports them. Data classification systems need to be built or updated to identify which information is most sensitive and subject the most stringent access controls to it. Data loss prevention tools that monitor for unusual upload activity, including large file transfers to external platforms, can catch problems before they become litigation. Regular audits of AI model outputs for signs of verbatim reproduction of proprietary content are worth considering for organizations that train models internally.²²

For companies training or fine-tuning models on proprietary data, differential privacy techniques offer a meaningful, if imperfect, way to reduce the risk of training data memorization. Watermarking of sensitive documents can help establish provenance if misappropriation is later alleged. These are not silver bullets, but they are the kind of layered technical safeguards that courts will look for when assessing whether an organization took its trade secret obligations seriously.

D. Preparing to Litigate

Companies that handle valuable proprietary information should assume, as a matter of planning discipline, that they may one day need to litigate an AI-related trade secret claim, whether as plaintiff or



defendant. That means preserving documentation of protective measures contemporaneously, maintaining clear records of training data provenance and model development decisions, and identifying technical experts with AI-specific expertise before a dispute arises rather than after.

The discovery phase of AI trade secret litigation is genuinely unlike conventional IP disputes. Technical experts will need to analyze model architectures, evaluate training data composition, and construct arguments about what a model could or could not have learned from particular inputs. Building those expert relationships in advance is not paranoid contingency planning. It is basic litigation preparedness for any company operating seriously in the AI space.

VI. Where the Law Needs to Go

The honest assessment is this: the legal frameworks currently governing trade secret protection are doing reasonably well at the edges of the AI problem and struggling at its center. The DTSA and UTSA provide workable foundations for cases involving clear misappropriation, intentional theft, or obvious breach of confidentiality obligations. They are less equipped to address the more diffuse, system-level risks that AI creates: training data absorption, casual employee disclosure, AI-assisted reverse engineering at scale.

Legislatures need to address, with some specificity, how the "reasonable

measures" standard applies when the relevant information flows through AI training pipelines. Courts need clearer evidentiary frameworks for evaluating claims that a model was trained on misappropriated data. The international dimension demands coordination that is not yet happening: a trade secret misappropriated through a model trained across three jurisdictions and accessed in a fourth presents problems that no single domestic statute can cleanly resolve.

None of this is a reason for despair, or for paralysis. Trade secret law has adapted to new technologies before. It adapted to digital files when physical document theft was the paradigm. It adapted to employee mobility in a world of increasing labor fluidity. It will adapt to AI. The question is whether that adaptation will happen proactively, through careful legislative and doctrinal development, or reactively, through expensive litigation that could have been avoided with clearer rules.

For practitioners, the message is straightforward even if the problems are not: the time to engage with these issues is now, before a client's competitive advantage walks out the door embedded in a model parameter that no one thought to protect. The companies and counsel who treat AI governance as a legal priority rather than an IT afterthought will be better positioned, legally and commercially, than those who are still writing the policy the week after the breach.



REFERENCES:

1. Pooley, J. (2022). *Trade Secrets: The Use of Confidential Business Information and Know-How in Commerce* (2nd ed.). Law Journal Press.
2. Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2), 841–887.
3. Defend Trade Secrets Act of 2016, Pub. L. No. 114-153, 130 Stat. 376 (codified at 18 U.S.C. §§ 1836–1839).
4. 18 U.S.C. § 1839(3) (2016).
5. 18 U.S.C. § 1836(b)(2) (2016) (providing for ex parte civil seizure in extraordinary circumstances).
6. National Conference of Commissioners on Uniform State Laws. (1985). *Uniform Trade Secrets Act with 1985 Amendments*. Retrieved from <https://www.uniformlaws.org>
7. *Uniform Trade Secrets Act* § 1(2) (1985).
8. Agreement on Trade-Related Aspects of Intellectual Property Rights art. 39, Apr. 15, 1994, Marrakesh Agreement Establishing the World Trade Organization, Annex 1C, 1869 U.N.T.S. 299.
9. Directive (EU) 2016/943 of the European Parliament and of the Council of 8 June 2016 on the Protection of Undisclosed Know-How and Business Information (Trade Secrets) Against Their Unlawful Acquisition, Use and Disclosure, 2016 O.J. (L 157) 1.
10. Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., Oprea, A., & Raffel, C. (2021). Extracting Training Data from Large Language Models. In *30th USENIX Security Symposium (USENIX Security 21)*, 2633–2650.
11. Nasr, M., Carlini, N., Hayase, J., Jagielski, M., Cooper, A. F., Ippolito, D., Choquette-Choo, C. A., Wallace, E., Tramer, F., & Lee, K. (2023). Scalable Extraction of Training Data from (Production) Language Models. arXiv preprint arXiv:2311.17035.
12. Goldman, E., & Rauterberg, G. (2023). *Generative AI and the Law: A Practitioner's Guide*. *Stanford Technology Law Review*, 26(1), 1–68.
13. Samsung Bans ChatGPT Use Among Staff After Sensitive Data Leak. (2023, May 2). Reuters. <https://www.reuters.com/technology/samsung-bans-use-generative-ai-tools-like-chatgpt-staff-after-sensitive-data-2023-05-02/>
14. Samuelson, P. (2023). How Generative AI Systems like ChatGPT Are Built on Intellectual Property. *Communications of the ACM*, 66(9), 20–22.
15. Grimmelmann, J. (2023). Who Are the AI Owners? *Yale Journal of Law & Technology*, 25, 1–55.



16. Waymo LLC v. Uber Technologies, Inc., No. 3:17-cv-00939-WHA (N.D. Cal. 2018). See also Levine, J. (2020). The Trade Secret of Self-Driving Cars. *Berkeley Technology Law Journal*, 35(2), 649–712.
17. See, e.g., Doe 1 v. GitHub, Inc., No. 4:22-cv-06823-JST (N.D. Cal. 2023). See also Dobkin, J. (2023). Misappropriation in the Age of Machine Learning. *Columbia Science and Technology Law Review*, 24(1), 112–165.
18. Federal Trade Commission. (2023). Generative AI Raises Competition Concerns. FTC Business Blog. <https://www.ftc.gov/business-guidance/blog/2023/06/generative-ai-raises-competition-concerns>
19. Securities and Exchange Commission. (2023). Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure, Release No. 33-11216, 88 Fed. Reg. 51,896 (Aug. 4, 2023).
20. Sanger, C., & Lobel, O. (2023). AI Governance in the Workplace: Legal Frameworks for Employee Data and Confidentiality. *University of Pennsylvania Journal of Labor and Employment Law*, 25(3), 301–354.
21. Stone, J., & Lee, A. (2023). Drafting Enforceable AI Confidentiality Provisions in Employment Contracts. Practical Law Company, Thomson Reuters.
22. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep Learning with Differential Privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.