NO. 25-2153

UNITED STATES COURT OF APPEALS FOR THE THIRD CIRCUIT

THOMPSON REUTERS ENTERPRISE CENTRE GMBH AND
WEST PUBLISHING CORPORATION
Plaintiffs-Appellees

ROSS INTELLIGENCE INC.
Defendant-Appellant

On Appeal from the United States District Court
for the District of Delaware
Civ. Action No. 20-613 (Hon. S. Bibas)

BRIEF OF AMICI CURIAE RANDY GOEBEL AND LARRY ULLMAN

Michael P. Abate
Burt A. (Chuck) Stinson
KAPLAN JOHNSON ABATE & BIRD LLP
710 West Main Street, Suite 400
Louisville, KY 40202
Telephone: (502) 416-1630
mabate@kaplanjohnsonlaw.com
cstinson@kaplanjohnsonlaw.com

REQUIRED CERTIFICATION OF BAR MEMBERSHIP

Pursuant to 3d Cir. L.A.R. 28.3(d), I hereby certify that I am a member in good standing of the bar of the Third Circuit.

s/Michael P. Abate

TABLE OF CONTENTS

TABLE OF AUTHORITIES

Cases

Other Authorities

INTEREST OF AMICI CURIAE[1]

Randy Goebel is a professor of Computing Science at the University of Alberta in Edmonton, Alberta, Canada. He has worked on Artificial Intelligence (AI) research for several decades, and on a variety of its applications, including 12 years on AI and legal reasoning. He co-leads an international legal reasoning competition consortium (Competition on Legal Information Extraction and Entailment or COLIEE) with colleagues in Japan. He has published more than 200 peer-reviewed papers on AI and has supervised more than 80 graduate degrees. He's also a co-founder of one of Canada's three federally funded AI institutes, the Alberta Machine Intelligence Institute.

Larry Ullman is a technical writer, author, and former Extension Instructor for the University of California, Berkeley. Larry has written over 20 books on programming, databases, and other computer-specific topics. He is currently a consultant focusing on developer experience

[1] Pursuant to Fed. R. App. P. 29(c)(5), amicus certifies that: (1) no party's counsel authored the brief in whole or in part; (2) no party or party's counsel contributed money that was intended to fund preparing or submitting the brief; and (3) no person other than amicus, its members, or its counsel contributed money that was intended to fund preparing or submitting the brief.

and developer documentation, and is also an investor/advisor in several startups, almost all of which employ AI in a significant capacity.

How courts legally define types of AI and the fair use of information that goes into building AI foundations will have sweeping implications for technical and business innovation for years to come. Thus, it is hard to understate how important it is for courts to base their rulings on an accurate understanding of this new, and rapidly evolving technology. These *amici* are well-equipped to contribute to that understanding.

AUTHORITY TO FILE AMICUS CURIAE BRIEF

All parties have consented to the filing of this brief. Therefore, a motion for leave to file this brief is not required under Fed. R. App. P. 29(a).

SUMMARY OF ARGUMENT

As Justice Frankfurter observed nearly a century ago, courts must be careful "not to embarrass the future by judicial answers which at best can deal only in a truncated way with problems" created by emerging technologies. *Northwest Airlines v. State of Minnesota*, 322 U.S. 292, 300 (1944). That is why, "[w]hen confronting new concerns

wrought by digital technology," courts must be "careful not to uncritically extend existing precedents." *Carpenter v. United States*, 585 U.S. 296, 318 (2018).

Judge Bibas' conflicting summary judgment opinions in this case perfectly demonstrate how existing copyright precedent does not provide an analytical framework that maps neatly onto the complex reality of artificial-intelligence systems. In his laudable effort to analyze this case within the copyright framework that Thompson Reuters wields like a broadsword, Judge Bibas unfortunately oversimplified ROSS Intelligence's product—and AI in general—into a binary scheme that is unrecognizable to those of us who develop and use AI systems in the real world. Specifically, Judge Bibas seemed to believe that AI can be divided into two types—generative and non-generative—and that this taxonomy dictates whether something is sufficiently "transformative" to qualify as fair use. The reality is far more complicated, however, and *amici* urge this Court to eschew that false distinction, lest the monopolistic impulses of legal publishers be allowed to stifle the advancement of an emerging revolutionary technology—in the legal field and beyond.

ARGUMENT

I. The Trial Court's "Fair Use" Analysis Hinged on its Assumption that ROSS Intelligence's AI Model was Not "Generative."

In considering whether ROSS Intelligence's use of Westlaw headnotes constituted fair use, the trial court repeatedly stated that ROSS Intelligence's product is not "transformative." *See Thomson Reuters Enter. Ctr. GMBH v. Ross Intelligence, Inc.*, 765 F. Supp. 3d 382, 397, 398, 399 (D. Del. 2025). That holding, in turn, was based on the court's understanding that "Ross's AI is not generative AI (AI that writes new content itself)." *Id.* at 398. "Rather, when a user enters a legal question, Ross spits back relevant judicial opinions that have already been written." *Id.* In other words, the court believed that AI must be generative to be transformative: it must create something that did not exist before. And, if it is not "transformative" in this way, the Court believed it cannot constitute fair use.

The trial court drove home this point, and signaled its attempt at proceeding cautiously, by noting that "the AI landscape is changing rapidly," and that his opinion concerned only "non-generative AI." Because the court adopted an oversimplistic view of AI, as explained

below, that analysis is likely to cause significant harm to this powerful, emerging technology sector if it is not reversed.

II. The Trial Court's Opinion is Based on a Colloquial Understanding of the Term "Generative AI" That Fails to Appreciate How AI Technology Actually Works.

The colloquial mental model of AI in general, and of "generative AI" in particular, is wrong. Partially driven by the sudden popularity of ChatGPT in 2022, the concept of AI moved beyond science fiction and theoretical computing to become a household term. As often happens in linguistic migrations, however, technical terms are diluted to the point of inaccuracy.

Lay language equates "generative AI" with AI that dynamically makes content. With the initial popularity of ChatGPT, the content was an image: real people in faked situations, entirely fake people proffered as real, anything incongruous or inane. Instead of requiring the artistic, graphic skills to create such an image, the user only needed to describe it.

The "generative AI" label has similarly been applied to writing (essays, poems, etc.), recipes, instructions (e.g., explain how to do this task), even computer code. Despite the different mediums, the

fundamental concept is consistent: the public believes that a "generative" AI application uses its vast scope of captured knowledge of a field—photos, writing, programming—to manufacture a new artifact within that field. That's precisely the understanding the trial court harbored; it found that ROSS's AI was not "generative" because it does not "write[] new content itself." *ROSS Intelligence*, 765 F. Supp. 3d at 398.

But this is not what "generative" AI actually means. The reality is that generative AI is but one type of AI, even if currently the most commonly known. The commonplace, but quite wrong view of "generative" AI adopted by consumers (and the trial court) likely stems from the average person's first overt interaction with AI being something like ChatGPT. Companies promote this use of AI to build their brand—ChatGPT is a product. Add in the novelty of new technology, and the reward feedback mechanism of a viral culture, and suddenly "I made this image" becomes "this is generative AI" or even just "I'm using AI".

But identifying applications as either "generative" or "non-generative" is both scientifically illogical and also detrimental to proper

understanding of our changing world. In fact, one cannot even categorize the types of AI in a single way.

For example, AI can differ by **technology**. Some systems use machine Learning (ML). Others use Natural Language Processing (NLP). Still others use Knowledge Representation and Reasoning (KRR). And some are Robotics systems.

Or, you could classify AI systems by **capability**. For example, AI could be "weak" or "narrow." Alternatively, it could be "broad" or "shallow." AI could be considered "strong" or "general." Or it could be "superintelligent."

AI could also be classified by its **function**. It could be reactive. It could be used for decision support instead of decision making. It could be a limited-memory system. It could rely on what's known as the "theory of mind." It could even be self-aware.

Finally, AI could be classified by its **purpose**. It could be considered "generative," "predictive," "assistive," or "conversational."

These sample groupings, while useful, aren't even an accepted universal standard. These are just four lenses that *amici* find useful ways to think about and examine AI technology. Others may categorize

them slightly differently. *See, e.g.*, Gosearch AI, *The Ultimate Breakdown of Different AI Types and Models*.[2]

This categorization problem is not unique to AI, of course. Take the horse for example. In some contexts, it's a pet. In others, sporting equipment. It could be an investment. Or it could be a therapy animal. And it could just be a wild creature. But one would never divide the animal kingdom into "horse" and "not horse." Yet that's effectively what the trial court's reasoning did here with regards to AI systems.

In reality, "generative" AI could fit into multiple frameworks outlined above. Clearly it's one sub-type of *Purpose*. It's also a subset of *Narrow AI*, when focusing on AI types by *Capability*. But regardless of what lens one uses to view AI, would the trial court's reasoning mean that everything that's not "generative" counts as non-generative, and therefore not transformative? As AI scientists, writers, and entrepreneurs, that distinction makes no sense to us. But according to the opinion of the district court, that's what the Copyright Act requires.

[2] Available at https://www.gosearch.ai/blog/breakdown-of-different-ai-types-and-models

However one categorizes it, *truly* generative AI is way more than user-driven, dynamic, manufactured output. Using Siri (on Apple devices) or Alexa (on Amazon devices) to set timers or play music are also examples of generative AI. Machine learning, which consolidates data into knowledge, is generative. What all these applications have in common isn't their output. It's the behind-the-scenes effort that makes the output possible.

Generative AI examines vast amounts of data, applies pattern recognition and other human-like learning techniques, and builds a virtual encyclopaedic "brain" of expertise in a given field. This brain can then be used to create new content within that field: an image, a poem, a recipe. But—here's the key—the generative aspect includes creating the brain, unbeknownst to the eventual end user. As explained below, that's precisely what ROSS's AI does.

III. ROSS Uses AI in Transformative Ways

ROSS developed an algorithm, trained on legal memos, to facilitate legal research. Instead of the prompt "create an image of a frog riding a bicycle" that returns an image, the prompt "provide

relevant case law surrounding…" returns a thoughtful list of relevant cases, with context.

For this to even be possible, the ROSS AI first generated an expert "brain" that stores case history in a semantically meaningful way. This is fundamentally the same application of technology as used by ChatGPT, Siri, and so on. Thus, ROSS does not just "spit back" caselaw, as the trial court believed, 765 F. Supp. 3d at 398; it uses the complex brain it has built to contextually identify and present caselaw relevant to a query.

Indeed, as ROSS's brief explains (ROSS Br., p. 41), its system does not simply "spit back" caselaw. *See* 765 F. Supp. 3d at 398. Rather, when a user inputs a query, ROSS's system first analyzes that question using its AI "brain" to reach a "semantic understanding" (ROSS Br., p. 41) of the question involved. Then, it goes out and "search[es] across actual caselaw rather than editorial content" to find relevant cases responsive to the query. *Id.* To the undiscerning user, it might <u>*seem*</u> like the ROSS search results are similar to the results you might get on Westlaw, but that is not the case. ROSS first translates the query into an understanding of what a user is really asking, and then goes out and

finds cases that address that topic—whether they use the same search terms or not.[3] Put differently, even if the user-facing output looks similar, the process by which ROSS returns search results is quite different from Westlaw; thanks to AI, it is able to find cases that express the same concept, even if they use different terms that would not be "hits" in a traditional Boolean search.

In summary, even under the generative/non-generative dichotomy the trial court created, the methods by which the ROSS system captures and uses public legal case information can be casually labelled as "generative." But as experts in the field with a vested interest in the intersection between copyright law and AI development, we urge the court to eschew simple distinctions like this, whose unintended consequences could stunt the growth of transformative industries.

[3] Westlaw has incorporated a version of this technology into its search, called "Parallel Search." This was a technology launched by CaseText, one of ROSS's competitors that Thompson acquired for nearly three-quarters of a billion dollars. Those search results are different than what Westlaw's own search algorithm returns based on "keywords."

CONCLUSION

For the foregoing reasons, as well as those spelled out in ROSS's brief, this Court should reverse the trial court's infringement finding as to ROSS AI.

Respectfully submitted,

/s/ Michael P. Abate

Michael P. Abate
Burt A. (Chuck) Stinson
KAPLAN JOHNSON ABATE & BIRD LLP
710 West Main Street, Suite 400
Louisville, KY 40202
Telephone: (502) 416-1630
mabate@kaplanjohnsonlaw.com
cstinson@kaplanjohnsonlaw.com

CERTIFICATE OF COMPLIANCE

1. This brief complies with the type-volume limitation in Fed. R. App. P. 29(a)(5) and 32(a)(7)(B) because it contains 2,129 words, excluding the parts of the brief exempted by Fed. R. App. P. 32(f).

2. This brief complies with the typeface requirements of Fed. R. app. P. 32(a)(5) and the type-style requirements of Feed. R. App. P. 32(a)(6) because it has been prepared in a proportionally spaced type face, Century Schoolbook, in 14-point font.

3. The electronic copy of this brief has been scanned for viruses and none were detected.

/s/ Michael P. Abate

CERTIFICATE OF SERVICE

I certify that service of this brief was filed electronically with the Clerk of the Court for the United States Court of Appeals for the Third Circuit by using the appellate CM/ECF system on September 29, 2025.

I further certify that all participants in this case are registered CM/ECF users and that service will be accomplished by the appellate CM/ECF system.

/s/ Michael P. Abate

CERTIFICATE OF ELECTRONIC FILING

I hereby certify that the text of the electronic brief filed through the CM/ECF system is identical to the text in the paper copies dispatched on September 29, 2025, by Federal Express to the Clerk of the Court of the United States Court of Appeals for the Third Circuit.

/s/ Michael P. Abate