



The First Appellate AI Copyright Win Was Not About Generative AI,
and That Is Exactly Why It Matters



The First Appellate AI Copyright Win Was Not About Generative AI, and That Is Exactly Why It Matters

On September 29, 2026 a unanimous Third Circuit panel rejected fair use for AI training. The system at the centre of it generated nothing: it was a search tool that copied 2,243 Westlaw headnotes.

The Third Circuit affirmed Judge Stephanos Bibas's Delaware partial-summary-judgment ruling against Ross Intelligence on September 29, 2026, six years after Thomson Reuters filed in May 2020. That makes it the first federal appellate answer on AI-training fair use. The infringement finding rested on 2,243 Westlaw headnotes, and the panel called the dispute no more than an ordinary copyright case. Everything else being read into it is interpretation.



Ross copied headnotes to build a competitor, so the court found no transformation

The panel's reasoning was narrow and blunt. Per [Reuters](#), the court wrote: "Ross took the headnotes to make it easier to develop a competing legal research tool. So Ross's use is not transformative." In the [unsealed opinion](#), the court added that "Thomson Reuters' materials possess a creative spark, and ROSS aspired to be a direct competitor by using them for a highly similar purpose."

Two things carried the result. First, the headnotes cleared the Feist originality threshold because Westlaw editors chose which legal propositions to summarize and how to word them, as [MSK's analysis](#) explains. Second, factor four. [LawNext](#), reporting on the unsealed opinion, notes Ross already had access to the underlying judicial opinions. It copied the headnotes because that was faster, and the court rejected convenience as a justification.

That last point is the one worth sitting with. Ross did not need the protected material to build its product. It needed it to build its product cheaply.

Ballard Spahr and Fox Rothschild read the same opinion ten days apart and disagreed

The claim repeated everywhere since the ruling is that Ross is the first domino in the AI training wars and reaches image models, music models, and everything else trained on protected work. [Fox Rothschild](#) put it that way: copyright owners' first appellate win in the AI training wars, with implications far beyond legal research. [Ballard Spahr's October 2, 2026 alert](#) reads the same opinion as distinguishing Ross from systems that generate original expression, and as explicitly leaving generative-AI questions unresolved.

Both firms are competent. Both read the same unanimous opinion. The gap between them is not sloppiness. The opinion genuinely supports both readings depending on which sentence you weight. Weight "no more than an ordinary copyright case" and Ross is a narrow holding about a non-generative search index that copied editorial summaries. Weight "aspired to be a direct competitor" and you have a factor-four test that any model



trained on a rightsholder's catalogue to compete with that rightsholder will struggle to pass.

Does the Ross ruling apply to image and music generators?

Not directly, and the comparison that matters is one the [ABA's Business Law Today](#) draws: Bartz v. Anthropic, where model training on lawfully acquired books was treated as highly transformative. Ross and Bartz point in different directions, and both are on the books.

My take: the two cases are not in conflict, and treating them as a split is the mispricing. Bartz involved a model whose output does not substitute for the training corpus in the market that corpus sells into. Ross involved a tool aimed at the exact customer Westlaw sells to, built from Westlaw's own editorial work, when the free underlying opinions were available. That is a difference in market substitution, not a difference in judicial mood. Factor four did the work in Ross, and I think it will do the work elsewhere.

The variable is not whether your model generates. It is whether the rightsholder can show you took a shortcut past a market they were already selling into.

So the question a buyer of creative-AI systems should ask is this: whose commercial market does this output land in, and did the training set come from that same market's supplier?

Ballard Spahr may simply be right, and here is the strongest version of that

The best argument against my reading is textual. The panel did reserve the generative-AI question. It said so. A court that wanted to write a rule for diffusion models and large language models had the opportunity and declined, and framed the whole thing as ordinary copyright. Appellate courts that call a case ordinary are usually signalling that it should not be cited for anything extraordinary.



The First Appellate AI Copyright Win Was Not About Generative AI, and That Is Exactly Why It Matters

There is also a factual distinction I cannot wave away. Ross's output was a list of relevant cases. Nothing it produced carried Westlaw's expression forward. The infringement happened entirely at the input stage, and the court still found substitution, because the input was the product. For an image model, the relationship between training inputs and sellable output is less direct, and a different panel could reasonably find that attenuation decisive. If another circuit rules the other way on a generative defendant in the next eighteen months, my reading of factor four as the universal lever is wrong and Ballard Spahr's framing holds.

I do not think that is the likeliest outcome, but it is live, and anyone telling you Ross settles the generative question is selling certainty that six years of litigation did not produce.

Spotify, Universal and the AFM moved while the firms argued

Spotify and Universal Music announced AI covers and remix licensing with artist and songwriter revenue sharing in May 2026. Pulling the other way, the American Federation of Musicians sued Universal, Warner and Atlantic over alleged licensing of members' recordings to Suno and Udio without adequate compensation.

Read those two together. The deal structure is forming, and the fight has already moved from "is training lawful" to who in the chain gets paid. I wrote about the gap between legal wins and actual artist income in [an earlier piece on artists winning the copyright war and losing the market](#), and Ross does nothing to close it. A headnote publisher with a litigation budget won. An individual illustrator is not in that position, which is why the AFM case is about collective licensing rather than infringement.

I expect the next appellate ruling to turn on factor four again

If it turns on transformativeness instead, and a court holds that generation of new expression is categorically transformative regardless of market effect, the Ross framework I have described collapses into a legal-research footnote.




The First Appellate AI Copyright Win Was Not About Generative AI, and That Is Exactly Why It Matters

Things I would watch, and would adjust on. A circuit split: Ross is one circuit, one panel, affirming one district judge, and a contrary holding elsewhere turns a settled-looking question into a Supreme Court question. Six years from filing to appellate answer suggests that timeline is long.

Licensing deals that cover training inputs rather than outputs. The Spotify and Universal arrangement is about AI covers and remixes, which is output-side. If the next wave of deals prices training corpora directly, the market has decided the question ahead of the courts.

And whether anyone applies the convenience point. The court rejected convenience as a justification because Ross had lawful access to the underlying opinions. That is the most portable piece of the opinion: if a cheaper lawful path existed and you skipped it, that fact now has appellate weight.

What to take from September 29, 2026 is smaller than either camp wants it to be. A court looked at an AI training case and declined to treat it as a new category of law. It applied four factors, found the fourth decisive, and called it ordinary. **Ordinary is a harder standard to argue around than exceptional**, and it applies to systems that generate nothing at all.



WRITTEN BY

Artur Markus

Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

Book a meeting →