



A Copyright Reboot for Robots

Interview with Ryan Abbott

**For podcast release
Monday, January 23, 2023**

KENNEALLY: Machines that can write poetry, paint scenic vistas, and compose sonatas are no longer found only in science fiction. Today, artistic automatons increasingly share our world. Soon, the robots may even have their day in court.

Welcome to CCC's podcast series. I'm Christopher Kenneally for Velocity of Content.

In *The Reasonable Robot: Artificial Intelligence and the Law*, law professor Ryan Abbott argues that a technological society like ours must abandon discrimination between AI and human behavior and develop innovative legal principles on intellectual property to close the gap between machines and mortals.

Professor Abbott joins me from Los Angeles to explain why a reboot for robots would benefit human well-being economically and socially. Welcome to Velocity of Content, Professor Abbott.

ABBOTT: Thanks, Chris. Very excited to be here.

KENNEALLY: Well, we're excited to speak with you. It's a very timely topic to be discussing right now, but the story is one that has a really rich past, and you're going to tell us about that. We'll start by asking about a court case that you're involved in. In June 2022, your client, Dr. Steven Thaler, who is a developer of artificial intelligence systems that generate creative output, sued the Copyright Office. In 2019, the US Copyright Office had refused to register a copyright claim from Dr. Thaler for an author identified as the Creativity Machine. The Copyright Office has consistently refused to extend copyright protection to nonhuman creations over more than four decades. Why do you think it's time to reverse that?

ABBOTT: You made a great point. Now's an exciting time to be talking about this, because law professors have been interested in this sort of thing really for a very long time, and so have technologists, because AI has been functionally making creative works for a very long time. People have alleged that this has been going on for decades. It hasn't been traditionally very interesting to lawyers, policymakers, and industry, because while the



technology existed, it just wasn't that commercially useful. Last year, we saw a real paradigm change in the ability of these generative models now open to the public online to make art and literature and images and text at scale in ways that have value to people using it in all sorts of activities. So while these kind of legal issues have been really around a long time, they have just suddenly picked up a real commercial importance, and people who weren't looking at it before are now thinking carefully about it.

The US Copyright Office has had an official policy since 1973 that human creativity is a fundamental requirement to protect a work with copyright. Other jurisdictions do it differently. The United Kingdom, for example, has a law from 1988 that says in the absence of a traditional human author, a work can still get copyright protection. That's called a computer-generated work, and there's a slightly different framework for it.

But the Copyright Office policy has never been tested in court, probably because again, while of theoretical interest – if an AI can make a song, that's very interesting. If it can make a song people actually want to listen to, that's on the radio, that should be generating streaming royalties, suddenly there's a reason to litigate over it. But this Copyright Office policy has never been tested in court, and nowhere in the Copyright Act does it say an author has to be a human being. In fact, for more than a century, the US has had corporate authors. So corporations can be legal authors, not acknowledging a human being.

The Copyright Office largely draws on case law where courts considered creativity and put it in human-centric sorts of terms, but really did so based on an assumption that a creative actor is a person, that human creativity is exceptional. In fact, the cases they rely on are from the 19th century, before really the development of modern computers at all. So we argue that they're wrong to rely on that sort of thing and that protecting AI-generated output is consistent with the purpose and the language of the Copyright Act, which is to promote the generation and dissemination of new works. Increasingly in the future, instead of music and movie studios just going to human creatives to do things, they're going to be using generative AI systems to do some creative work in ways that have social benefit. That is really the intent of the Copyright Act, and that's what providing protection would allow.

KENNEALLY: As AI and other new technologies reach new heights, Professor Abbott, it's still true that the bar for copyright is fairly low. Tell us about that.

ABBOTT: The bar for copyright is exceptionally low. To get copyright protection, a work has to be original, and there has to be some amount of creativity associated with that. In what has essentially become the leading Supreme Court case on the originality standard, the court held that an alphabetical phone book listing didn't count as creative, because really there's only one way you could make a phone book or that someone would want to make a



phone book. But if I was sitting here doodling while we were having our podcast, that podcast would be protected by copyright for 70 years after I died. Our podcast is protected by copyright. My random lectures are protected by copyright. Really, very little creativity is needed to make something protectable.

Congress put together a commission to consider some of these issues in the '70s, and the commission said AI really autonomously doing stuff is too speculative right now for us to consider. Whether or not that was true then, it is no longer true now. You can go to these open-source generative AI systems, like DALL-E 2 or GPT-3 or Codex or Stability or Diffusion or Midjourney, type a few words in, and it will make a piece of art for you. It's a piece of art that is far beyond the threshold for creativity, sometimes much better than anything – well, most of the time much better than anything I could make, although there is a real art to phrasing these prompts and a degree sometimes of human/AI joint creativity that's very interesting.

KENNEALLY: So the results will differ from individual to individual, and if I may say, from robot to robot.

ABBOTT: Indeed. Different jurisdictions have different theories on which copyright is given. Pretty much every jurisdiction provides copyright as a result of international treaties like the Berne Convention and the TRIPS Agreement, but even so, copyright differs in different jurisdictions.

For example, in France, it's very clear that most of the people there – at least the people involved in the legal system – believe that copyright is really predominantly something that exists to protect the moral rights of authors. If that should happen to have some sort of commercial impact on the publishing industry, so be it. But really, there's a real focus on authorship and moral rights.

Whereas in the United States, the Constitution, Congress, and the courts have all been very clear that there are many benefits to having copyright law as well as costs, and authors do benefit directly from having copyright protections, but the primary beneficiary of the law is intended to be the public. The theory is that providing authors with these incentives, you are encouraging people to engage in activity that has a broader social benefit. So if that's the theory on which you're providing benefits, that fits very neatly into a concept where AI is being used to provide these benefits. If you really only wanted to encourage people directly doing creative things or other sorts of goals of copyright law, as they might in other jurisdictions, you might feel a bit differently about it.

KENNEALLY: With AI and machine learning systems, they train and refine their skills based on millions, even hundreds of millions, of publicly available copyrighted works. Recently,



artists and other creators have begun to allege infringement in such cases, and they are suing many leading technology companies over these practices. What do those complaints seek to establish, and how have the AI developers responded?

ABBOTT: Well, those complaints are a fairly new thing, and the AI companies have not responded in court yet to them. Two class actions were filed at the end of 2022 by the same or a similar group of lawyers against leading marketers of generative AI systems. One involved Codex, which is this automatic code generator. Basically, you give it prompts, and it helps code the back end of, for example, an application or a website. That complaint alleged that Codex used open-source code, which it was free generally to do, but it wasn't giving attribution, which was one of the license requirements of the open-source code, to where it came from. So the complaint alleged that that was copyright infringement.

Another complaint filed by that group as, again, a putative class action against some of these AI image generators is that using copyrighted images to train an AI was infringement, because it involves copying the image many, many times, which copyright generally prohibits without permission. And it also alleged that the works coming out of these AI systems, when they were made in the style of a particular artist, were derivative works, which means they would be infringing.

This raises, really, some very interesting questions that are somewhat open even when people are involved. This question of whether you can train AI on copyrighted works is very controversial and important. Some AI developers say that training an AI under US law should be fair use. We have a legal doctrine that says, yes, it's copyright infringement, but we think it's fair, because it meets certain factors. Generally, rights-holders believe that they should be given the opportunity to license their works for that sort of use, potentially to prohibit it being used if they don't want. For example, if an artist says I don't want your AI to be able to make something in my style, I just am not going to license this to you. Or if they do want to license it, to get fair remuneration for doing so.

So cases have looked at AI training, text and data mining, in different sorts of contexts, but never really here. And there's some really interesting open questions about whether an artist's style is protectable and exactly how. These cases should generate some useful law for right-holders and for attorneys working in the space.

I should also mention Getty Images recently announced that they were bringing a lawsuit against one of these AI art generators in the UK, again based on the theory that it used Getty Images to train an AI and that these images were part of output being made, and this was all done without a license. Some of these systems have used a very large number of works found in a variety of places, including the internet.



KENNEALLY: Professor Abbott, how should the US and other countries possibly develop AI policies and laws that you think would prove beneficial for society?

ABBOTT: That's a great question, and it's something jurisdictions have been thinking about for a long time. We're just getting to the stage where jurisdictions are ramping up the number of rules and regulations they have on AI, and there's two kind of fundamental mindsets about regulating, one of which is that the law historically lags behind technological evolution and that maybe we should let that happen and let the market do what it will, and the other view is the law should precede technological development. That's largely a view I hold, in the sense that laws and regulations exist to promote generally the public interest, and there are ways that we want these systems to develop in socially useful, as opposed to socially harmful, sorts of ways. This really requires the policymakers be thinking deeply about these issues and thinking about what rules will best promote social interest, rather than coming to the party years late and trying to regulate once the damage has already been done.

On the other hand, exactly how to regulate is a very complex issue, and there are a lot of complex issues here, and it does differ by sector. But if you think about copyright, for example, I think the policymakers should be looking at, well, why do we have copyright law? What problems is it trying to solve? In the United States, we have copyright law to promote the generation and dissemination of new works. So rather than taking a very textualist, literal approach to the Copyright Act and trying to think, well, someone wrote this in the '70s without AI in mind. How can we shoehorn something in here? Taking a step back and saying, well, with AI making works, is that something we want to protect, or is it not something that we want to protect? And thinking about that and what the implications of it are and then what rules follow.

To me, the right rule that follows is, well, this is the sort of thing that we do want to protect, so we should be clear that we want to protect it. It is, of course, not just protecting AI-generated works. It is AI-generated works and infringement. It is text and data mining. It is training AI systems more broadly. It is AI challenging other fundamental tenets of copyright law, like style. We have laws on when style can be protected, and those laws were developed at a time where it was kind of difficult to copy someone's style. People can do it. It takes a lot of work. Only some people can do it. But in the age of AI, it is going to be very easy to copy someone's style. If you have an AI music generator, and you want to make a bestselling hit that sounds like Taylor Swift wrote it, the time is not far from when we'll be able to ask AI to do that. AI does that right now, just not that well. But it won't be too long until people actually want to hear that music. That's going to fundamentally change these tests and what the right outcome should be.



KENNEALLY: And it seems to me, Professor Abbott, that copyright has played an important role in monetization of works, but the situation you're describing is one where copyright potentially can drive innovation, creating exciting new works that we don't even imagine right now.

ABBOTT: I think that's exactly right. In the case of an AI, it's generally the case that a lot of the cost involved in making these generative systems comes upstream the specific generation of one creative work. Once an AI is fully trained and operational, the cost of having it make additional works is not so great. So this incentive here is really acting upstream on the people making and developing these systems. Instead of encouraging someone to make specific creative works, you may be encouraging them to build systems that make creative works. And the more value those works have, the more you're encouraging people to build these systems and build systems in ways that are really going to make creative works in ways we can't now.

For example, there are companies that have AI making personalized music, and they can make, at least in theory, music for me based on how I'm feeling in a moment – if I'm exercising, music to encourage me to exercise well, if I'm feeling sad, to improve my mood, so on and so forth, or music for me in real time to go along with a video game I'm playing. The idea of having your own personalized composer with you making music on the fly to change how you're feeling in that moment is a really interesting and ethically fraught sort of activity that we haven't really had to consider up until now.

KENNEALLY: Well, Professor Ryan Abbott, author of *The Reasonable Robot: Artificial Intelligence and the Law*, thanks so much for joining me today and discussing these issues.

ABBOTT: My pleasure.

KENNEALLY: That's all for now. Our producer is Jeremy Brieske of Burst Marketing. You can subscribe to this program wherever you go for podcasts, and please do follow us on Twitter and on Facebook. You can also find Velocity of Content on the CCC YouTube channel. I'm Christopher Kenneally. Thanks for listening.

END OF FILE