

## What if Goodbye Didn't have to be Forever?

Is it a good idea to allow the living to reconcile with the deceased through AI imitation?

## What Even is Copyright Anymore?

New policies in AI are blurring the lines between fair use and plagiarism.

## Trust me, I'm a Machine

Even if AI technologies can help improve efficiency in the medical field, should we forfeit the comfort that human doctors give?

# AI Nexus

*April 2025*

**Where Ideas Come Together**



Our AI

<https://www.our-ai.org>





# Contents

## 2 Editor's Note

## 3 What if Goodbye Didn't have to be Forever?

Could AI preserve love beyond death?

## 6 What Even is Copyright Anymore?

Should we *let* the AI companies use our data?

## 8 Trust Me, I'm a Machine

Comfort is universal, yet irreplaceable.

## 13 Acknowledgements

# Editor's Note

Thank you for reading yet another edition of *AI Nexus*. Unfortunately, this version has been delayed for about 10 days due to AP testing and finals; it takes me up to 20 hours to make one of these—writing, research, and designing included—and I simply could not juggle this workload with all the studying.

If you haven't noticed yet, this version is significantly shorter than usual, with many of our usual staff writers sharing the need to pursue similar commitments as I have. Despite this, we're slowly growing, and we're looking forward to some awesome things during summer! Hope you enjoy.

–Thomas Yin

Chief Editor  
5/09/2025

# What if *Goodbye*

Didn't have to be Forever?

*Liv Skeete*



Imagine losing someone you love—and then hearing their voice again, not just echoing softly in memory, but truly responding. Imagine laughter bubbling up, an argument resurfacing with familiar warmth, or a quiet conversation that feels impossibly real. What if goodbye didn't have to be forever?

In 2021, technology took a bold leap closer to this reality. Microsoft filed a patent for a chatbot designed to replicate real individuals, including those who had passed away. The idea is simple yet profoundly unsettling: using digital footprints, such as social media posts, text messages, emails, videos, and voice recordings, to reconstruct personalities as interactive AI characters. Explicitly, Microsoft outlined the possibility of recreating deceased loved ones, igniting a mix of intrigue, hope, and dread across the internet. Although Microsoft clarified they weren't actively pursuing this technology, the patent still exists, leaving the possibilities open for future exploration.


The idea of digitally resurrecting loved ones isn't entirely new. Popular culture grappled with it powerfully in the widely acclaimed "Be Right Back" episode of *Black Mirror*. In this deeply moving story, Martha, a woman shattered by the sudden death of her partner, Ash, turns desperately to technology. She uses an advanced AI service to recreate Ash from his online presence: first as texts, then voice calls, and eventually a physical presence. Without spoiling its powerful journey, the episode captures both the profound comfort and deep unease that such an AI resurrection could evoke. What begins as a much-needed remedy for Martha soon evolves into a harrowing exploration of authenticity, grief, and the





human inability to fully accept loss. "Be Right Back" resonated deeply with audiences—not just as a chilling fictional scenario, but as a prophetic caution for the real world.

Then, fiction became reality. Joshua Barbeau, grieving deeply after losing his fiancée Jessica Pereira, found himself facing a world without her laughter, conversations, and quiet reassurances. In an aching attempt to hold onto Jessica's essence, Joshua turned to Project December in 2021. This AI platform allowed him to upload Jessica's text messages and social media history, creating an interactive chatbot eerily reflective of her personality. As he conversed with the AI, Joshua experienced profound emotional moments; at times comforting, at others deeply unsettling. The chatbot sounded so similar to Jessica, capturing her quirks, humor, and mannerisms, that Joshua described feeling momentarily as though she were truly present again. Yet he grappled with an unsettling sense of artificiality and emptiness beneath the simulated warmth. Joshua's story highlights a critical truth: technology can mimic our loved ones convincingly, but can it ever truly restore them?



The emotional pull towards AI resurrection is undeniably powerful. When confronted with devastating loss, humans instinctively cling to remnants of those they love. Traditionally, these remnants are photographs, handwritten letters, or cherished mementos. AI, however, offers something extraordinary: the promise of interaction, of conversations continuing past death. As memories fade over time, people grasp at fading recollections—how someone smiled when teasing, the rhythm

of their laugh, or the distinctive quirks that defined their personality. AI seemingly preserves these intricate details, protecting cherished memories from the relentless erosion of time. But at what cost?

Consent is critical to account for. Should individuals explicitly agree, pre-death, to their digital resurrection? What if they never had the chance—or worse, never wanted it? Even with consent, ethical dilemmas persist. Is recreating someone digitally truly honoring their memory, or does it inadvertently disrespect their humanity by reducing them to static data points?


Corporate exploitation further complicates this digital afterlife. When digital personas are managed by profit-driven companies, there's a risk these businesses might selectively amplify appealing traits or behaviors to encourage continued interaction, potentially distorting the authentic essence of the person. And when life itself becomes a commodity, ethical boundaries quickly blur.

Consider a subscription-based service maintaining your father's digital consciousness. If the company suddenly raised prices beyond your means, would discontinuing the subscription equate to "killing" your digital father? How much would you be willing—or able—to pay to keep him "alive"? Such scenarios underscore the troubling potential for emotional and financial exploitation. Perhaps a fairer approach would be to mandate a fixed, one-time fee, allowing individuals to permanently "own" their loved one's digital representation without the looming threat of ongoing financial manipulation.

Consider another scenario: a mother facing terminal illness creates an AI to replicate her voice, values, and stories for her young daughter. Age becomes a critical factor here. Providing a digital mother to a very young child, say a three-year-old, might hinder emotional bonding with potential future caregivers. Conversely, for an older child, perhaps seventeen, who already has established memories, a digital mother could offer meaningful comfort. Yet, determining the exact age at which digital resurrection transitions from harmful to helpful is challenging.

Additionally, this digital mother would remain static, potentially reinforcing beliefs or behaviors that the real mother might have naturally outgrown or reconsidered over time. This raises the question of whether digital representations should have built-in mechanisms to evolve or adapt subtly, or perhaps even a defined time frame after which their interactions become limited, gently encouraging loved ones to move forward.

Zooming out further, these questions pierce into the core of human identity itself. Are we merely the sum of our digital footprints—texts, posts, and voice recordings—or something more intangible and profound? More deeply, can machines capture what we often call the soul? What does it truly mean to "live forever" when your essence is reduced to code? AI resurrection challenges our understanding of legacy, memory, and immortality itself—inviting uncomfortable introspection about the essence of being human.



Despite these concerns, the human desire to persist after death is irresistible. Personally, I would still want a digital version of myself to exist, even if flawed. I would also want digital versions of my loved ones—not as replacements, but as enriched, interactive memories, a natural progression from photographs or videos. However, it's crucial to recognize these digital representations for what they truly are: comforting echoes rather than authentic continuations. Their inherent inability to genuinely evolve or adapt means they should be approached with emotional caution. Viewing these digital representations as mementos



rather than the actual person also encourages deeper appreciation of our loved ones while they're alive, reminding us that life is precious and irreplaceable, and should never be treated as expendable simply because there's an expectation it can be digitally continued.

Returning to the initial question: what if goodbye didn't have to be forever? Perhaps the question isn't whether we should embrace AI to revive our loved ones, but rather what this longing reveals about humanity itself. It speaks to our capacity to love deeply, our desperate struggle against the inevitability of loss, and the delicate beauty of memory. Maybe true immortality isn't found in algorithms or code, but in how we cherish, grieve, and ultimately let go. And perhaps, in exploring these haunting possibilities, we better understand what it truly means to be alive, to love—and to say goodbye.



# What Even Is Copyright Anymore?

Thomas Yin

Over the past few years, almost every single major AI company has faced some sort of court case directly involving copyright abuses in training AI. Although most of these cases are still ongoing, the advent of advanced generative AI trained on and producing massive amounts of data has blurred the line between copyrighted materials and works in the public domain. Do artists and news agencies deserve to be compensated — or even given copyright of — work produced by AI trained on their copyrighted content? As you might expect, the issue is a tad more complicated than a set of legal definitions.

Surprisingly, the debate of whether a non-human author of a creative work could license that work under copyright has been ongoing for quite a while. In 2011, a macaque (a type of monkey) in a British wildlife conservatory found a camera displaced by a photographer, and, grinning from ear to ear, took a few photos of itself.

When the wildlife conservatory eventually published these “monkey selfies”, citing the extraordinary event, the good people at PETA, a totally unbiased organization known not to file allegedly frivolous lawsuits about animal rights, decided that it was a good idea to sue the conservatory on behalf of the animal, claiming that the publication violated copyright laws, since the monkey, being the creator, had the copyright to the work and did not consent to its publication.

As expected, the lawsuit didn’t work. An American court of appeals shut down the case after reaffirming a previous court’s decision that in order for creative works to be licensed, it has to have a majority of creative input by a human author. In fact, U.S. courts as well as bureaus have, time and time again, reinforced the idea that copyright must be able to be held by a human being, not an “automatic” program or non-human. The Judicial system powerfully reaffirmed this concept just a





few days ago in Thaler v Perlmutter, in which it declared that works created by Artificial Intelligence are not eligible to be protected by copyright, citing that the Copyright Act of 1976 “requires all eligible work to be authored in the first instance by a human being”.

Yet, works with a significant portion of human contributions — the AI Nexus Magazine, for example — is considered eligible for copyright under the Act. And no, prompting doesn't count as human input, according to the U.S. Department of Copyright. While the courts have held firm on the status quo of denying copyright, they might have to reconsider this stance as the abilities and uses of AI evolve over time.

One of the most expensive parts of any AI model is data. When large-scale general use LLMs like ChatGPT first came out, researchers proclaimed that the model was trained on WebText, a massive text dataset constituting a sample of content found from the entire internet, a well-scrutinized fact that, after the hype died down, hinted at the problems to come. A lot of the things found online are open source (e.g. Reddit posts, Creative Commons artwork), but sometimes, these types of content just do not seem to cut it for training advanced AI models. About a year ago, stock image company Getty Images filed a lawsuit against Stability AI — one of the leading companies developing generative imagery AI — arguing that the company had scraped images off Getty's site and used it to train its AI models. Some of the evidence presented is almost comical: images produced by Stable Diffusion, Stability's

generative model, often shows a “Getty Images” watermark, distorted and barely recognizable at times. The case is ongoing and expected to produce a final decision sometime after a trial scheduled in the summer of 2025.

The scope of these disputes is not constrained to image AI models — one of the most famous copyright disputes over training material happened with New York Times, who famously sued OpenAI and Microsoft over copyright disputes, claiming that their LLM models, by training on copyrighted New York Times articles found on the internet, infringed on the publishing giant's copyright claims to these articles. They claimed that:

- Since AI models have the tendency to sometimes produce works facsimile, LLMs constitute an unlawful republishing of parts of a copyrighted work.
- AI models can summarize or paraphrase exclusive information found in its training data, and LLMs with this capability harms NYT by giving a user information that they would otherwise have to find by means of an NYT subscription.
- Since training data is an integral part of an AI model's functionality, and NYT articles constitutes a significant portion of ChatGPT's training data, OpenAI is profiting off copyrighted material produced by and licensed by NYT.

These lawsuits come at a critical time for AI development in the United States — along with a very much real AI race with China to stay ahead in developing the most

advanced AI model as well as the reformation of many longstanding legal precedents in the wake of the AI revolution, traditional ideas of originality and creativity are being challenged by the new way that people have begun using AI.

Why do we call something creative? Think about something that you consider creative. Maybe you thought of a cool beat, a beautiful painting, or a hilarious joke, but consider: what is it that these things have in common? Are they creative because they are fresh and pleasing to a certain aesthetic? Originally produced by an artist? Imbued with a sense of expression? Now, given this traditional perception of creativity, attempt to fit this idea to modern AI. Although some people have attempted to argue otherwise, it is commonly agreed that AI, by solely attempting to replicate facets of its training data, cannot actually produce traditionally original works. Yet, AI is much more arguably able to produce works that evoke emotion or match a certain aesthetic. Philosophical questions rarely have a defined answer, so it is difficult for me to explain whether AI can truly be creative, yet legal matters (copyright in regards to AI, for example) must always be settled with solid, albeit imperfect, solutions.

AI generated text and video will probably not receive copyright status anytime soon, not only because of the aforementioned strong precedent against attributing copyright to non-human creators but also because there would not be a single entity that copyright could be attributed to — should the hundreds of

thousands of creators of training data receive copyright for the fact that their work was used to train the AI, or should the company who assembled the product get the final say in how the work is published? In contrast, the attribution of training data is a big question that has a more definite answer. Perhaps in response to the lawsuit, OpenAI's CEO Sam Altman famously claimed a few days ago that the U.S., if it does not pursue the freedom of AI companies to train on copyrighted works, will fall behind China in the AI race. While Altman's points that copyrighted works are critical to many AI training pipelines and that access to copyrighted works will probably significantly improve the quality of AI than without are most likely correct, his call for unconditional fair use for copyrighted works as training data holds two major inconsistencies.


First of all, AI companies live off data, and often pay hefty sums for it — in fact, some AI companies employ an enormous labor force to manually label data as part of a process called Reinforcement Learning with Human Feedback. AI companies have been known to vehemently pursue deals to procure training data in some cases (for example, OpenAI signed major deals with Vox Media and the Associated Press to license the companies' media contents for AI training), yet, as alleged by the NYT-backed court case, frequently do not disclose other sources of copyrighted data. Many expect these disputes to be settled privately, but I believe that these settlements overlook the fact that current AI companies give polarizing treatment to different media outlets. So if AI companies



like OpenAI have history of licensing the content of other companies, why not strike a deal with New York Times? While some speculate that the suit could be a negotiation tactic instead of a genuine challenge on an important precedent, it seems logical that AI developers should remain consistent in their policies of acquiring and using training data.

In addition, AI, even though not at the peak of its potential as a tool, is already being marketed as an effective way for large companies to cut costs and bolster profit (most of the largest corporations have already taken steps to integrate AI into their operations, and smaller companies are expected to follow in their footsteps). Considering that the aggregate of often copyrighted high-quality training data is critical to the performance of the AI, and therefore the convincingness of AI as the tool of efficiency as it is being marketed decreases, most likely impacting customer satisfaction and subscription sales. It is paradoxical for companies such as OpenAI to expect free training data under fair use to improve the for-profit AI products that are being sold for personal or professional use.

Regardless of the future legal stances in relation to copyright and AI, we should recognize that the AI that we use today is built on trillions of words, typed out by millions of people and dissolved into bits of numbers and word fragments, hopefully to the benefit of society. Even without the wishful thinking, I have to say that this fact alone is somewhat poetic.



Eric Song

# TRUST ME, I'M A MACHINE

I remember the hands of my mother, caressing my leg after I fell from a particularly high monkeybar. Steady, warm, practiced from decades of nursing experience. Whenever I got hurt as a kid— a twisted ankle, a high fever— she never panicked. Rather, she always listened. She knew when to worry and when to wait. This kind of care is getting increasingly difficult to obtain these days. The storm arrived slowly. Like many other technological revolutions, it came silently— disguised as innovation, wrapped in synthetic words like “efficiency”, “precision”, and “personalization”. However, now that AI is entering hospitals and exam rooms with direct life-and-death consequences for any mistakes, I’m starting to wonder if we’re witnessing something incredibly more disruptive than just a software upgrade.

Over just the past year, I’ve watched with unease as AI systems became more incorporated into healthcare, tasked with jobs such as diagnosing diseases, assigning treatments, and managing personal records. Tech companies call this “revolutionary”, promising better outcomes, faster decisions, and lower costs, but beneath these promises lies a grave issue—we are building a healthcare system that is increasingly trusting machines over people—and in doing so, we are sacrificing something irreplaceable.

There have already been severe cases of AI mishandling this power. For instance, a study in 2019 published by Science reveals that a healthcare prediction algorithm used throughout the US to identify high-risk patients for care management programs was far less likely to flag Black patients.



Our AI

<https://www.our-ai.org>



This is what made it in the headlines, but what doesn't?

When an AI misdiagnoses a rare form of cancer or flags a false positive, rather than showing up in the headlines, it shows up in the panic of a patient, the scheduling of an unnecessary biopsy, the missed chance to catch an existing disease early. The data driving AI is vast, yes – but it's also messy, biased, and far from complete. With the glow of innovation blinding the majority, these flaws are often dismissed as just “glitches” – growing pains of a system still developing. Meanwhile, the human side of healthcare is being automated out of existence. Doctors now spend more time clicking checkboxes on AI-powered EMRs than they do talking to patients. Predictive models are now replacing gut instinct. And in all of this, no one seems to ask a simple question: just because a machine can do something – does that mean it should?

Some advocates argue that AI will relieve medical personnel of mundane tasks and focus more on care, however that ideal is already drifting out of reach. For instance, in many hospitals, AI is not augmenting human decisions– it is replacing them. UnitedHealthcare faced backlash after reports that an AI algorithm was being used to automatically deny claims for elderly patients in rehab– often without a doctor ever reviewing the case.

The concerns don't stop with just technicality – they're ethical. Who would be responsible when an algorithm makes a deadly mistake? Can you sue a machine? Can you appeal its decision? In 2021, IBM's Watson for Oncology was pulled from several hospitals after reports that it recommended unsafe cancer treatments. What's worse, no one could explain how things got so wrong– because the system's reasoning was hidden behind proprietary code. In the age of AI-driven medicine, accountability becomes scattered. Responsibility blurs. When something goes wrong – and it will – it's the patient that will ultimately suffer.

This decade may not be a “golden age of medical progress”, but the moment we traded trust for speed, empathy for automation, and wisdom for code. My mother didn't need an algorithm to tell her when a patient was scared. She could see it in their face, hear it in their voice. She didn't need a model to know that healing takes more than a pill– it takes human presence. In our rush to reimagine medicine, we must be careful to remember that the field serves humans, not machines.

**Design**  
Thomas Yin

**Writing**  
Thomas Yin  
Olivia Skeete  
Eric Song

**Editing**  
Thomas Yin

**Special  
Thanks to:**  
Ran Gu

**Thank you for reading!**



# Acknowledgements

