Artificial Intelligence and Content Ownership

Muhammad Ali Qadri, Mohammed Abdul Malik, Raghav Sethi Cristian Diaz
Virginia Polytechnic Institute and State University, VA, USA
{ muhammadaliq, buildwithmalik, raghavsethi, cristd4 }@vt.edu

Abstract

The creation of Artificial Intelligence has led to things we have only dreamed about. With the advent of autonomous vehicles, programs beating the best humans at things like Go and Chess, to name a few, and stock predictors along with several other applications, this has also brought some interesting problems. AI has gotten to the point where it can generate art, music, and even full-length movies, and when trained properly, the results are pretty convincing. People are winning competitions that have huge prizes with AI, and people will soon try to make hit songs with this. With such an overwhelming market and industry impact, we intend to review key cases, and literature regarding AI-produced media ownership and try to determine a potential solution for the ethical implications this can bring. We will consider various stakeholder perspectives, including those of the software/model builder, artists, hardware team, and company selling the software, as well as the owner of the machine. Since it's hard to definitively say who should get credit for such a work and how laws should be defined to incorporate such challenges. Should the software/model builder get credit, the artists who select amongst the generated pieces, the hardware team, the company selling the software, or the owner of the machine?

***Index Terms*—DAO, AI, Smart Contracts, Generative Models, GANs, Autoencoders**

I. Introduction

Why do we even need to worry about AI-generated art at all? On August 29, 2022, Jason Allen made use of AI software (Midjourney) and won $300 by taking 1st place at the Colorado state fair. A tweet went viral expressing how unjust it was. The tweet itself was not that interesting, but one of the replies said "We're watching the death of artistry unfold right before our eyes — if creative jobs aren't safe from machines, then even high-skilled jobs are in danger of becoming obsolete. What will we have then?" [1]. That person could be on to something. There are already projects that are producing AI-generated music. There are also many projects that use AI models and different services to generate images and music that can eventually be stitched together into a full-length movie [2]. This power over AI and what its potential can be is truly unsettling. Could this someday get to the point of winning a grammy? And if that were to happen, who gets to copyright the song? We'd argue that it's already challenging enough to make it as a musician if AI gets to that point. We simply don't know what to say to artists trying to make it in such times. Maybe, they should have been born in a different era, and are just unlucky. This is why it's vitally important to have some outline of what is copyrightable, and who should get the credit, if anyone at all, regarding AI-generated work. This can be e

In this report, we aim to discuss and debate the different techniques of AI-generated art and how they can impact ownership. It is very hard to say who should get ownership over the created art form, whether it would be the AI tool that created the art. The person who used the tool? or the person who originally thought of making this AI generator to produce art. We discuss the different laws around the copyright of intellectual property in the EU or the US law. We also discussed different ideas such as DAOs (Decentralized Autonomous Organizations) and art generation techniques such as GANs (Generative Adversarial Networks). Debating these different concepts allows us to get a better perspective on who has actual ownership.

II. Literature Review

Let's discuss another real-life example. A dutch team decided to embark on a project to recreate art like the legendary artist Rembrandt. For this project, artists, AI specialists, and even museum staff were involved in trying to help recreate this. The specialists upscaled the original Rembrandt work, and created algorithms to detect and focus on the features according to what the artists were prioritizing. The end result was amazing, a new artwork that looked legitimate. Technologically speaking, Microsoft was backing this project with their computational power, along with art experts' input, it's hard to say who deserves the credit in a scenario like

this. They all had equal stakes, and according to some, legal definitions were involved. [3], [4]

Within the scope of the law, we should establish what currently binds and creates the framework for what is considered valid for intellectual property. Hugenholtz did just this for Europe as per the most recent laws at the time of its publication.

In this publication by Hugenholtz [5], it was noted that 4 tests must be satisfied in EU copyright law. They were: Production in the literary or artistic domain, human intellectual effort, originality/creativity, and expression. So the summarized version is that the first constraint is easily met, and the second constraint isn't ruled out after a precedent set in the Painter case which established that it is possible to create works of authorship with the aid of a machine or device (such as a camera). Hugenholtz argued that this applies to AI in several ways aside from just picking the one to submit. Whether it be, the development of the AI software, the gathering, and choice of training data, or the desired functional specifications, among others. The amount of involvement necessary Hugenholtz asserts is the point of contention and debate. The third point, again from the painter's case, allows for human originality/creativity in some parts of the process of the work. This consists of an iterative process containing conception, execution, redaction, and finally, the output. Long story short, only in execution could there not be an argument for some form of human involvement. Lastly, the fourth point requires the human's creativity to be expressed in it, which can just boil down to intent, and this Hugenholtz asserts isn't completely ruled out if the work stayed within the author's original intent.

Having considered European law, and its take on intellectual property, it is also time to consider US law. Given that these laws will bind the students at this university and even at their next job, they are essential to know. As things are now, AI has no grounds for intellectual property, since the work isn't man-made. However, Ravid from Fordham University proposed a different model and argues that current law is not prepared for the technological advancement that is AI. Their proposition is to consider AI as work for hire, which would give the user the copyright privilege, which can be a person, a corporation, or other entities. The paper considers different approaches, separating AI software and copyright [6].

A. Current state of copyright laws

The current state of copyright laws presents a number of challenges and limitations with regard to the ownership and licensing of AI-generated content. In many cases, existing copyright laws do not adequately address the unique characteristics of such content, and do not provide clear and consistent rules for determining ownership and authorship.

For example, under US copyright law, the default assumption is that the creator of a work is the person who physically creates it, using their own skill and judgment. This assumption is problematic in the case of AI-generated content, as the content is not created by a human in the traditional sense, but rather is generated by an AI algorithm based on a set of inputs and rules.

As a result, the question of who should be credited or held responsible for AI-generated content is a complex and contentious one, and is likely to require the development of new legal frameworks and concepts to address it.

To address these challenges, US copyright laws will need to be amended to explicitly recognize the unique characteristics of AI-generated content, and to provide specific rules and regulations for determining ownership and licensing of such content. This could involve the development of new legal mechanisms, such as decentralized autonomous organizations (DAOs), to manage the ownership and licensing of AI-generated content in a transparent and accountable manner.

Additionally, US copyright laws will need to be updated to take into account the collaborative and iterative nature of the process by which AI-generated content is created. This could involve recognizing the contributions of multiple stakeholders, including the software/model builder, the artist or journalist who selects among the generated content, and the owner of the AI system, in determining ownership and authorship of the resulting content.

B. Generative Models

Artificial Intelligence has revolutionized different industries like healthcare, finance, retail, transportation, and manufacturing. Recently, we have seen a lot of progress in content generation performed by AI models. We can now generate high-quality text, images, audio, and video from small input provided by the users. With the availability of these tools to the public, there would be an increase in misuse like fake news, deep fakes, etc. Some people are worried that as text generation systems get better and better at creating realistic-sounding fake text, they could be used to create false news stories, spread disinformation, or otherwise trick people. There is a need for a tool to identify if a work of art is generated by AI or not. We will present some generative models and their state of the art results.

Generative models are a class of machine learning models that are capable of generating new data samples that are similar to the training data. These models are trained to learn the underlying distribution of the data, and they can then generate new samples that are drawn from this distribution. Generative models have a wide range of applications, including image generation, text generation, and even protein structure prediction. They can be used to create realistic synthetic data for use in machine learning algorithms, and they can also be used to generate novel samples that can be used to test the performance of machine learning models.

Recent advances in deep learning have led to the development of more sophisticated generative models, such as generative adversarial networks (GANs) [7] and variational autoencoders (VAEs) [8], which are able to learn complex distributions and generate high-quality samples. The GAN is a relatively new approach to generative modeling that has received significant attention in the machine learning community. The GAN was introduced by Ian Goodfellow in 2014, and it has since been applied to a wide range of tasks, including image generation, text generation, and even protein structure prediction.

One of the key contributions of Goodfellow's GAN paper is the introduction of the concept of a generative adversarial network, which consists of two competing neural networks, a generative model and a discriminative model. The generative model is trained to produce data that is similar to the training data, while the discriminative model is trained to distinguish between the generated data and the real data.

Overall, the GAN paper has had a significant impact on the machine learning community, and it has opened up new possibilities for generative modeling. While there are still challenges to be addressed, the GAN approach has proven to be a powerful tool for learning complex distributions and generating high-quality data.

In recent years, there have been several new generative models that have been built on top of the GAN framework. These models have been developed to address some of the limitations of the original GAN approach, and they have demonstrated impressive results on a variety of tasks.

One example of a new generative model built on top of GANs is DALL-E 2 [9]. DALL-E 2 is a model that is trained to generate images from text descriptions, and it has been shown to produce highly detailed and diverse images. For example, given the text description "a giraffe wearing a hat," DALL-E can generate an image of a giraffe wearing a hat, even though it has never seen such an image during training. Another famous example of a new generative model built on top of GANs is stable diffusion [10], which produces equally good results.

These new generative models built on top of GANs are making significant contributions to the field of machine learning, and they are enabling researchers to tackle new and challenging tasks. In the future, it is likely that we will see even more advanced models built on top of GANs, which will continue to push the boundaries of what is possible with generative modeling.

OpenAI [11], the company behind DALL-E 2 exposed the generative model via API. In a blog post [12], they reported that three million people are now using DALL-E 2 to produce over four million images a day. In response to concerns about unaddressed rights issues, Getty Images announced in August 2022 to ban the upload and sale of illustrations generated using DALL-E 2 and other similar tools [13]. This decision was followed by other stock image websites. It was stated by the companies that the ban was motivated by concerns about the copyrighted images contained in the training data sets for systems like DALL-E 2, which are often scraped from the web. Artists whose styles can be replicated using the system may face threats to their livelihood, particularly if they did not give consent for their work to be used in the training of DALL-E 2. To address the concerns of artists, Shutterstock, a rival of Getty Images, recently announced that it would use DALL-E 2 to generate content, but would also launch a "contributor fund" to reimburse creators when the company sells work to train text-to-image AI systems [14]. In addition, the company is banning AI art uploaded by third parties in order to minimize the risk of copyrighted work appearing on the platform. This approach aims to find a middle ground between the use of AI for generating content and the rights of artists.

In recent years, AI has made significant progress in generating music. The ability to create music using AI has the potential to revolutionize the music industry, enabling the generation of high-quality, diverse, and personalized music on demand. One of the key challenges in generating music using AI is modeling the long-term structure and coherence of musical compositions. This is particularly difficult in the raw audio domain, where the complexity and variability of musical signals pose significant challenges for machine learning models. To address this challenge, researchers have developed a range of techniques, including the use of recurrent neural networks, GANs, and VAEs. Another important aspect of music generation using AI is the ability to control and steer the musical style and content. This can be achieved by conditioning the generation process on various factors, such as artist and genre, as well as on unaligned lyrics to control the singing.

OpenAI has introduced Jukebox [15], a neural network that is able to generate music in a variety of genres and artist styles as raw audio, including rudimentary singing. This is a significant advancement in the field of automatic music generation, as it allows for the creation of music directly at the audio level, rather than using symbolic representations. This allows for a greater level of expressiveness and flexibility in the music that is generated. To tackle the long context of raw audio, the model uses a multi-scale VQ-VAE (Vector Quantized-Variational Autoencoder) to compress it into discrete codes, which are then modeled using autoregressive Transformers. The combined model can generate high-fidelity and diverse songs that are coherent for up to multiple minutes. Jukebox can be conditioned on artist and genre to steer the musical and vocal style, and on unaligned lyrics to make the singing more controllable. OpenAI has released thousands of non-cherry-picked samples, along with the model weights and code.

Researchers from MIT and Harvard have developed GLTR [16], a tool to help humans detect whether a text was generated by a model. A human-subjects study has shown that GLTR, which applies a suite of baseline statistical methods to detect generation artifacts across common sampling schemes, improves the detection rate of fake text from 54% to 72% without any prior training. Initiatives like GLTR can be useful for not only detecting fake text, but also identifying Twitter bots that have been used to interfere with electoral processes in the US and other countries. Hugging Face has also hosted GPT-2 Output Detector Demo [17] to detect AI-generated content.

C. Impact on media and communication outlets

The use of artificial intelligence (AI) in the generation of news articles and other forms of journalism has the potential to significantly impact the field. AI algorithms can be trained to analyze data sources such as financial reports, press releases, and social media posts, and automatically generate coherent, informative news stories based on this information. This has the potential to significantly reduce the time and effort required to produce news content, and could enable journalists to focus on more complex, investigative reporting tasks.

However, there are concerns about the accuracy and bias of AI-generated news stories. AI algorithms are only as good as the data and information they are trained on, and there is a risk that they could produce biased or misleading content if the underlying data is flawed or biased. Additionally, AI algorithms can be vulnerable to manipulation by malicious actors who may seek to use them to spread misinformation or propaganda.

To address these concerns, it is important to carefully evaluate the quality and reliability of AI-generated content, and to develop appropriate measures to detect and prevent bias and manipulation. Additionally, the use of AI algorithms to assist with fact-checking and verification of news stories can help to improve the accuracy and reliability of the resulting content. However, it is important to ensure that the use of AI for these purposes does not compromise the independence and integrity of the news media.

III. DAOs a potential solution?

After careful consideration, DAOs have the potential for dealing with the question of ownership, rather than leaving the choice to governments. A DAO is simply a Decentralized Autonomous Organization, which the idea is, for a given transaction, there isn't a middleman or intermediary entity [18].

Another article examined goes on to give examples of how AI could take over the role of a "minor" and associate certain AI's to an NFT, whose inherent value then becomes intrinsically linked to the decision makers or coin owners [19].

Let me offer a concrete example, Marcello Mari, CEO of Singularity DAO, makes use of Singularity Net's AI technology to handle crypto management. His goal is to help people who might not be as tech savvy, invest in crypto; the AI will then pick crypto investments based on underlying algorithms. A step towards what Mari describes as 'financial singularity', which is when robots will manage all financial transactions[20]. In this case, the investors are the benefactors, so the question of ownership is answered. The best feature of handling AI in this matter, is again, the fact that the government isn't controlling any of this, and is decentralized, for fairness.

This sounds interesting, right? Well, things aren't always so simple. There are skeptics when it comes to AI DAO. Trent McConaghy , believes that following this path will lead us to a future where we're humans renting out AI-DAO services. Since, at least at a theoretical level, we can apply AI to just about anything. He detailed an example of how an Art AI Dao would learn to generate art that humans enjoy. Maybe at some point, the alpha program can make a duplicate of itself and change a few things, and see how this new beta program performs. Obviously, if things aren't working out, then it'll just run out of resources, but if the program thrives, it will eventually do the same thing. In doing so, we're giving AI the opportunity to "wake up" with access to almost unlimited resources the program is being given. This is why we need to have an ethical discussion now on how to avoid AI dominance. [21]

One way would be to have sort of fail-safes by taking almost a medicinal approach where a small version is released, and tries to add defenses against it growing too much as antibodies do for us, which is what Trent was proposing. The thing about that is even with viruses, mutations happen, and we'll need to possibly add more defenses or try to add more at a certain interval or frequency.

In either case, we believe DAOs are a good approach, but we need to remember that we still need to approach this with caution.

Having already laid out how AI is considered now, in different nations. It's clear that there's room for some copyright changes. We felt the need to address specifically three questions, so that the law fosters keeping technology safe, with room for investment. We think when it comes to these NFT's, that there shouldn't be a single owner or stakeholder holding all the shares, the idea behind AI dao is to let multiple perspectives flourish behind the algorithm's productions. Instead of the Sherman antitrust law, we'll probably need something similar, to prevent organizations or individuals from monopolizing some form of AI daos, say the Art DAOs for example. In the initial phases, it's ok if the shares are with only a few people, but after a certain ether growth(or some other currency), then it should prevent one singular person from owning too much. What would that percentage be? We're thinking anything preventing a complete majority, so up to 50% but maybe even that is too much.

Secondly, we'll once and for all deal with the copyright issues regarding AI software being sold. In this regard, we feel that it's best for strictly innovative software to earn. As for works produced using those programs, unless the customer who purchased the software was editing the underlying program or has some way of demonstrating its influence in the work, at every stage. Only then, do we feel that work has grounds for copyright, if any of those things are true. Which was the European perspective on this issue, which we found to be logical and just.

Lastly, we wanted to discuss if AI programs should be given the rights of people. To be honest, we determined that while on this surface this can be linked to whether or not AI can copyright its own works, which is an important question to answer. However, that's just one small facet with this question, and we think that this question is a different issue altogether. We would have to consider things such as "How would we still have control over them?", "Should we even have control over them?" and that's just one example of several questions that need to be answered for this.

IV. Conclusion and Future Works

The creation of Artificial Intelligence has led to many impressive and innovative developments, including the advent of autonomous vehicles, programs that can beat the best humans at games like Go and Chess, and stock predictors. However, the use of AI in the creative industries has also raised important ethical questions around ownership and credit for AI-generated art, music, and other media.

In this paper, we have aimed to review key cases and literature regarding AI-produced media ownership, and to determine a potential solution for the ethical implications of this technology. We have discussed the different techniques of AI-generated art, and have explored the challenges of determining ownership and credit for such works. We have also reviewed the relevant laws and regulations in the European Union and United States, and have discussed the potential use of decentralized autonomous organizations and smart contracts to address the issues of ownership and credit.

Overall, our review of the literature and discussion of the different concepts has allowed us to gain a better understanding of the complex ethical and legal issues raised by AI-generated media. While there is no easy solution to these challenges, we believe that the use of decentralized autonomous organizations(DAO's) and smart contracts may provide a way to fairly and transparently distribute ownership and credit for AI-generated works.

In conclusion, the rapid advancement of AI technology has the potential to significantly impact the creative industries, and it will be important for policymakers, industry stakeholders, and the public to carefully consider the ethical and legal implications of this technology. By conducting further research, debating the different concepts, and developing effective solutions, we can help ensure that the rights and interests of all parties involved in the creation of AI-generated media are protected.

For future work, we think the exploration of blockchain technology and the potential role of AI in the democratization of creative industries is next. Further research and experimentation in these areas may help to develop more effective and comprehensive solutions to the challenges raised by AI in the creative industries.

References

[1] E. Velie, "Ai beats out human artists at art competition," sep 2022.
[2] N. AL-SIBAI, "Someone's making an entire movie using video generated by ai," sep 2022.
[3] W. Knight, "A new tool uses ai to spot text written by ai," July 2019.

[4] S. Gehrmann, H. Strobelt, and A. M. Rush, "GLTR: statistical detection and visualization of generated text," *CoRR*, vol. abs/1906.04043, 2019.

[5] P. Hugenholtz and J. Quintais, "Copyright and artificial creation: Does eu copyright law protect ai-assisted output?," *IIC - International Review of Intellectual Property and Competition Law*, vol. 52, 10 2021.

[6] S. Yanisky-Ravid, "Generating rembrandt: Artificial intelligence, copyright, and accountability in the 3a era — the human-like authors are already here — a new model," *IO: Productivity*, 2017.

[7] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," 2014.

[8] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2013.

[9] OpenAI, "Dall-e 2," 2022.

[10] Stability.AI, "Stable diffusion public release," 2022.

[11] OpenAI, "Openai," 2022.

[12] Stability.AI, "Dall·e now available without waitlist," 2022.

[13] TheVerge, "Getty images bans ai-generated content over fears of legal challenges," 2022.

[14] N. Lomas, "Shutterstock to integrate openai's dall-e 2 and launch fund for contributor artists," 2022.

[15] OpenAI, "Jukebox," 2022.

[16] S. Gehrmann, H. Strobelt, and A. M. Rush, "Gltr: Statistical detection and visualization of generated text," 2019.

[17] H. Face, "Gpt-2 output detector demo," 2022.

[18] H. K, "Ai beats out human artists at art competition," Aug 2022.

[19] "Humans.ai presents its ai dao concept," Mar 2022.

[20] "Meet the dao chief who wants artificial intelligence to take his job," Dec 2021.

[21] T. McConaghy, "Wild, wooly ai daos," Jun 2016.