



ICLR

Copyright Fundamentals for AI Researchers

The Twelfth International Conference on Learning Representations

May 7, 2014



Law Offices of Kate Downing

Article I, Section 8 of the US Constitution

“The Congress shall have Power...

To promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries”



17 U.S. Code § 106 - Exclusive rights in copyrighted works

Subject to sections 107 through 122, the owner of copyright under this title has the exclusive rights to do and to authorize any of the following:

- (1) to **reproduce** the copyrighted work in copies or phonorecords;
- (2) to **prepare derivative works** based upon the copyrighted work;
- (3) to **distribute** copies or phonorecords of the copyrighted work to the public by sale or other transfer of ownership, or by rental, lease, or lending;
- (4) in the case of literary, musical, dramatic, and choreographic works, pantomimes, and motion pictures and other audiovisual works, to **perform** the copyrighted work publicly;
- (5) in the case of literary, musical, dramatic, and choreographic works, pantomimes, and pictorial, graphic, or sculptural works, including the individual images of a motion picture or other audiovisual work, to **display** the copyrighted work publicly; and...



Types of Copyright Infringement

1. **Direct** - Copying the training dataset, derivative works of the training dataset, creating the model, distributing the model, distributing/displaying/performing output that might be a copy or derivative work of the training data
2. **Secondary** - responsibility for the output created by users, what users do with that output, maybe user redistribution of the model. Types:
 - **Contributory** - direct infringement, knowledge (or reason to know) of direct infringement, material contribution
 - **Vicarious** - direct infringement, direct or indirect financial benefit from infringement, ability to control infringers' actions
 - **Active Inducement** - direct infringement, promotion of infringement or other affirmative action to foster infringement



Do You Know Anything?

Liabilities

Is training a model on copyrighted training data a copyright violation (**copying** the dataset, creation of **derivative works** when preparing the dataset, maybe **distribution** of the dataset)? Always, sometimes, never?

Is a model a **derivative work** of the training data? Always, sometimes, never?

Is the output a **derivative work** of the training data? Always, sometimes, never?

Assets

Is a model copyrightable?

Can the output be copyrighted under any circumstances?



Defense 1: No Violation of Any Copyright Rights

- Unlikely to succeed with respect to claims related to copying of datasets; derivative works of datasets
 - *Sega v. Accolade*: intermediate copying can be infringement even if the ultimate product is non-infringing, so you need to make a fair use argument instead. That argument in turn depends on the nature of the final product.
- Might succeed with respect to claims that the model is a derivative work of the dataset
 - No substantial similarity between the model and the dataset
 - Even if the model might be considered a compressed encoding of the training data, it's not the same as a phonorecord or JPEG because there isn't a 1:1 relationship between the training data and the output
- Likely to succeed with respect to claims that the output is a derivative work of the training data
 - No substantial similarity between the training data and the output when model follows good practices



Defense 2: Fair Use

Factors for consideration:

1. The purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes. Does the use add something new, with a different purpose or character? Is the new use *transformative*?
2. The nature of the copyrighted work - how creative is the original work?
3. The amount or substantiality of the portion used
4. The effect of the use on the potential market for or value of the work. Takes into account how easy the original works are to license.



Transformativeness and Computers: Non-Expressive Use

Use that involves copying, but either doesn't communicate the expressive content of the work or communicates expressive content but for purposes other than mere enjoyment of the work as originally intended. Ex.:

- Copying books to create a searchable digital archive that can return snippets of text in order to find a fact or metadata about books (*Authors Guild v. Google* (2015), *Authors Guild v. HathiTrust* (2014))
- Copying video game code as part of reverse engineering to enable interoperability (*Sega v. Accolade* (1992)); copying APIs so developers can readily use a new platform (*Google v. Oracle* (2021))
- Copying images to create image search engines that return thumbnail images that function as pointers to the original work (*Perfect 10 v. Amazon* (2007), *Perfect 10 v. Google* (2007))
- Copying to create an archive of papers to detect student plagiarism (*A.V. v. iParadigms* (2009))



Battle of the Analogies: AI, the Compressed Archive

“By training Stable Diffusion on the Training Images, Stability caused those images to be stored at and incorporated into Stable Diffusion as compressed copies... “[N]ew” images are based entirely on the Training Images and are derivative works of the particular images Stable Diffusion draws from when assembling a given output. Ultimately, it is merely a complex collage tool.”

-Complaint in *Andersen v. Stability AI Ltd.*



Battle of the Analogies: AI, the Learning Child

“When a model is exposed to a large array of images labeled with the word “cup”, it learns what visual elements constitute the concept of “cup-ness”, much like a human child does. It does this not by compiling an internal database of training images, but rather by abstracting the factual metadata that correlates to the idea of “cup”. This enables it to then combine concepts and produce a new, entirely original image ”

- OpenAI’s Response to the Copyright Office’s Request for Comments

See <https://katedowninglaw.com/2023/06/21/battle-of-the-ai-analogies/> for more on this.



The Narrative



Law Offices of Kate Downing

Copyright Infringement Risk Factors - Output

1. Is the model general/foundational or does it have a specific purpose? Does the purpose align with encouraging progress in the arts and sciences?
2. Does the model produce content that would have been copyrightable if created by a human? Facts, ideas, theories, systems, and methods of operation are not copyrightable.
3. Is the model likely to output any of its training data based on the nature of the output?
4. Does the model have effective guardrails to ensure the output isn't a derivative work of the training data?
 - Exact matches to training data
 - Style of a particular author/artist - not copyrightable, but likely to produce derivative work
 - Copyrightable characters
 - Book/movie plots
5. What are users using the model for in practice? Is that use influenced by marketing or sales strategy?



Risk Factors for the Defense - Training Datasets

1. Is the training dataset copyrighted?
2. Is there any way for rights holders to opt out of the dataset?
3. Have you excluded watermarked images?

Memorization

4. Does the training data contain many duplicates?
5. Is the training dataset small? Is the ratio between the size of the training data and the size of the model itself small?
6. Was there overtraining?
7. Are images associated with unique text descriptions?



Copyright Infringement Risk Factors - Availability

Distributed models: no specific knowledge of individual infringements, no direct ability to prevent infringement by users on a real-time basis or ongoing relationship, no practical way to terminate a license in many cases

Hosted models: guardrails are harder for users to overcome and more can be added at any time, reducing liability on a continuous basis. No possible liability for distribution of the model if it's found to be a derivative work of the training data. No obligation to do attribution for open source software licenses or other similar licenses.



Things That Likely Don't Matter Much

- License of the model or any acceptable use policy
- Whether you charge money for it if you are a commercial actor



Thank You!

Please reach out to me at **kate@katedowninglaw.com** if you think of any more questions!

This presentation is Copyright © Kate Downing 2024 and is subject to the [Creative Commons Attribution 4.0 International License](#).

