

Talkin' 'Bout AI Generation: Copyright and the Generative AI Supply Chain

Intellectual Property Scholars Conference
August 3, 2023

James Grimmelman
(with Katherine Lee and A. Feder Cooper)

Goals of this project

- Give a thorough overview of the copyright issues raised by generative AI
- Embrace the complexity of generative AI
- Embrace the complexity of copyright law

The Generative AI Supply Chain

The generative AI supply chain

- Data
- Datasets
- Foundation models
- Fine-tuned models
- Applications
- Generations

Data

- Text, photographs, music, etc.
- Created by a wide variety of people who mostly did not have generative AI in mind
- These are the relevant copyright owners

Datasets

- Huge scale: billions of items, terabytes of data
- Obtained via web-crawling, UGC uploads, and negotiated licensing

Models

- Training turns a dataset into a model
- The trainer chooses a model architecture, training algorithm, and training dataset
- Typical cost: \$100,000 to \$100,000,000, and running time from days to months

Fine-tuned models

- One model can be “fine-tuned” to create another by training it on additional data
- Specialize a general-purpose model for a domain of interest (e.g. scientific papers)
- Steer a model’s outputs(e.g., avoid giving harmful advice)

Applications

- Additional functionality can include customized prompts, external data sources, APIs, and post-processing the generations
- Spectrum of customization: from simple wrappers to complicated functionality
- Spectrum of openness: from cloud services to standalone apps to open-source releases

Generations

- Prompt + application + random seed = generation
- Frequently an iterative process as the user observes the outputs and refines their prompt

Observations

A diverse supply chain

- A dataset/model/app could ...
 - ... draw on multiple upstream ones
 - ... be used by multiple downstream ones
- Every stage could be ...
 - ... public or private
 - ... proprietary or open
 - ... carried out by the same or different parties

What counts as a copy?

- Datasets include literal copies
- Some generations are copies
- We think a model is a copy of a work when:
 1. It was trained on the work, and
 2. It can generate a copy of the work,
 3. From a prompt that is not copy of the work
- *The details vary by model and usage*

Who is a direct infringer?

- When a generative-AI service is used to produce an infringing output, the direct infringer could be:
 - The user (“oil painting of Elsa and Anna from Frozen”)
 - The app deployer (“heroic princesses” generates Elsa and Anna)
 - Both (the model is trained only on Disney princesses and the user knows it)

Upstream and downstream

- Copyright concerns cannot be localized to one stage in the generative-AI supply chain
- Secondary infringement doctrines can push liability forward and backwards in the chain
- The fair-use case for datasets and models depends on the balance of infringing and noninfringing generations

The gravitational pull of existing doctrine

- Most datasets, models, and apps do not fit ...
 - ... Section 512 notice-and-takedown
 - ... the *Sony* defense
- But there is a natural inclination to draw on these frameworks to define companies' duties
 - Might they be recreated within fair use?

Questions