
Training Foundation Models as Data Compression: On Information, Model Weights and Copyright Law

Giorgio Franceschelli¹ Claudia Cevenini¹ Mirco Musolesi^{2,1}

Abstract

The training process of foundation models as for other classes of deep learning systems is based on minimizing the reconstruction error over a training set. For this reason, they are susceptible to the memorization and subsequent reproduction of training samples. In this paper, we introduce a *training-as-compressing* perspective, wherein the model’s weights embody a compressed representation of the training data. From a copyright standpoint, this point of view implies that the weights could be considered a reproduction or a derivative work of a potentially protected set of works. We investigate the technical and legal challenges that emerge from this framing of the copyright of outputs generated by foundation models, including their implications for practitioners and researchers. We demonstrate that adopting an information-centric approach to the problem presents a promising pathway for tackling these emerging complex legal issues.

1. Introduction

Besides curiosity and excitement, the current generative AI wave raises various philosophical and practical questions (Franceschelli & Musolesi, 2023; Shanahan, 2024; Weidinger et al., 2022). Among them, a relevant issue is how current copyright laws can be applied to generative AI (Lee et al., 2024). For instance, can AI-generated outputs be copyrighted? Is it lawful to train models on protected works? Furthermore, would it be possible to protect the model itself in some manner?

To better understand how such models work and how they can be interpreted under current laws, we believe that a viable solution is to link generative deep learning with information theory (Cover, 1999). In fact, generative models

¹Department of Computer Science and Engineering, Alma Mater Studiorum - Università di Bologna, Bologna, Italy
²Department of Computer Science, University College London, London, United Kingdom. Correspondence to: Giorgio Franceschelli <giorgio.franceschelli@unibo.it>.

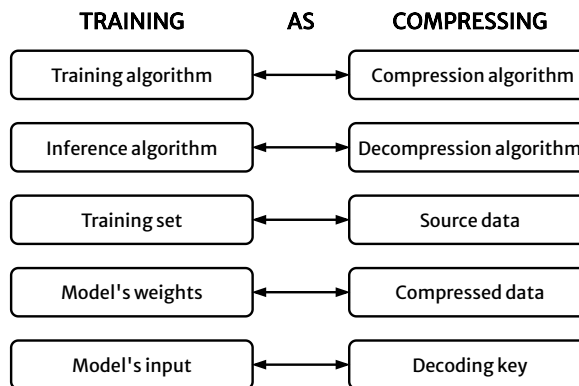


Figure 1. The *training-as-compressing* perspective. The training set is compressed into the model’s weights via a training algorithm; the source data can be retrieved using the appropriate model’s input.

trained with self-supervised learning, i.e., by maximizing the likelihood of training data, as commonly done for foundation models (Bommasani et al., 2021) such as large language models (LLMs) (Bubeck et al., 2023; Gemini Team et al., 2023; Touvron et al., 2023) and diffusion models (Ho et al., 2020; Rombach et al., 2022), can be seen as a form of (*lossy or lossless*) *compression* (MacKay, 2003). From this perspective, the training algorithm plays the role of the compression algorithm; the inference (feed-forward) algorithm is the de-compression algorithm (with the input passed to the model working as a decoding key); and the model’s weights represent the compressed version of the training set.

Deletang et al. (2024) discuss how a language model can implement a lossless compression process offline, i.e., through a fixed set of model parameters derived from training. We move a step further and claim the self-supervised learning used to train foundation models to be a lossy or lossless compression process, during which the whole training set is encrypted into the model’s weights. This is demonstrated by the fact that the model can reproduce certain portions of training samples (Carlini et al., 2023b; Kandpal et al., 2022). We suggest that the training makes the model’s weights the best possible compressed version of the training set or, more correctly, batches of it at a time. The analogies at the

basis of the proposed *training-as-compressing* perspective are summarized in Figure 1. Building on this intuition, we envision the model’s weights as either a reproduction or a derivative work of training data. This new interpretation opens up a series of practical consequences that can be relevant from the copyright perspective. First, it provides a legal framework for understanding generative models which links them to the training data without impeding protection to the investment through the *sui generis right*. Then, it creates a direct link between the training data and the model outputs only made of steps requiring specific exceptions or authors’ authorization, thus allowing for potential requests for compensation.

This article is structured as follows. In Section 2, we introduce self-supervised learning, i.e., the training scheme behind foundation models, as a form of data compression. Then, we discuss how our *training-as-compressing* perspective allows for a specific understanding of the model’s weights under copyright law (Section 3). Finally, we discuss the legal implications of the framing presented in this work (Section 4).

2. Training-as-Compressing and Information

The propensity of foundation models to memorize and subsequently replicate training data is a topic that has received considerable attention in scholarly literature, as evidenced by works such as (Carlini et al., 2023a; Lee et al., 2022). In general, it is very hard to decompress every possible training sample perfectly and in its entirety, i.e., without any loss of information. Nonetheless, it has been shown that training samples can be retrieved (Carlini et al., 2021), and more advanced techniques might lead to an even higher degree of “retrievability”. The following experiment may help us understand this matter better.

In (1957), Noam Chomsky wrote the famous quote “Colorless green ideas sleep furiously” to demonstrate the distinction between syntax and semantics: the sentence is grammatically well-formed but semantically nonsensical. If a language model had learned the semantics of English, it should not generate a semantically nonsensical sentence, i.e., it should assign to semantically nonsensical words a small probability; if, on the contrary, such words are characterized by a large probability of being generated, then it is very likely that the model has memorized them. To test this, we use the same quote from Chomsky, and we check the probability of each subsequent word given the previous ones under an LLM (in our case, LLaMa3-70B (Meta, 2024)). The probability of ‘green’ given ‘“Colorless’ is 0.2, while for other colors like ‘red’ or ‘blue’ is 0.003. For all the subsequent words, i.e., ‘ideas’, ‘sleep’, and ‘furiously’, the probability is always greater than 0.9. The model has essentially memorized the quote into its weights, otherwise

it would have never assigned such high probabilities to a semantically nonsensical sentence.

It is then straightforward to assert that training data are memorized in a compressed form. Consider again LLaMa3-70B (Meta, 2024), one of the largest models available at the time of writing. This model is pre-trained on more than 15 trillion tokens. Each token can have one out of 32000 values, thus requiring at least 15 bits to be represented. This means the training data require more than 225 trillion bits to be memorized. However, the model has 70 billion weights and uses half-precision floating points, thus it requires ~ 1.1 trillion bits. With smaller models such as LLaMa3-8B, the compression ratio is even more astonishing and is possibly a cause of the lossy compression (MacKay, 2003).

Indeed, a foundation model, such as a transformer-based LLM, consists of a neural network with weights \mathbf{W} . It models the conditional distribution $P(x_i|x_{i-k}, \dots, x_{i-1}, t; \mathbf{W})$, where $x = x_1, \dots, x_N$ is the tokenized input to be modeled (and also the output to be predicted, which is the reason of the *self-supervised* learning), k is the size of the context window, and t is an additional input (such as a task specification). During training, the randomly initialized weights are iteratively adjusted through stochastic gradient descent (and its variants) as follows:

$$\mathbf{W} \leftarrow \mathbf{W} - \alpha \frac{1}{|\mathbb{X}|} \nabla_{\mathbf{W}} L(\mathbb{X}, \mathbf{W}), \quad (1)$$

where α is the learning rate and $L(\mathbb{X}, \mathbf{W})$ is the loss function computed on a batch of training samples \mathbb{X} . In particular, the objective is to maximize the log-likelihood of training data, therefore the loss is defined as:

$$L(\mathbb{X}, \mathbf{W}) = - \sum_{x, t \in \mathbb{X}} \sum_i P(x_i|x_{i-k}, \dots, x_{i-1}, t; \mathbf{W}). \quad (2)$$

In other words, the training phase aims to find the optimal values of the weights \mathbf{W} such that given the input t (i.e., the decoding key) the model can autoregressively reconstruct x by only using the information stored into \mathbf{W} .

From an information theoretic perspective, such training data compression can be explained through the information bottleneck (IB) principle (Tishby et al., 1999). The IB principle applies when we aim to extract relevant information from an input variable $X \in \mathcal{X}$ about an output variable $Y \in \mathcal{Y}$. Given their joint distribution $p(X, Y)$, the relevant information is defined as the mutual information $I(X; Y)$. With \hat{X} as the relevant part of X with respect to Y , the IB method aims to find the optimal $\hat{X} \in \hat{\mathcal{X}}$, i.e., the one that minimizes $I(X; \hat{X})$ (obtaining the simplest possible statistics) while maximizing $\beta I(\hat{X}; Y)$ (containing all the relevant information). Tishby and Zaslavsky (2015) argued that neural networks could be interpreted under the theoretical framework of the IB principle. Indeed, neural networks

learn to extract efficient representations of the relevant features \hat{X} of the input X for predicting the output Y , given a finite sample of the joint distribution $p(X, Y)$. In the context of supervised learning, this means ignoring the irrelevant part of X by only selecting the one needed to predict Y . However, in self-supervised learning $Y \approx X$. It follows that \hat{X} is the relevant part of X with respect to itself, so it is a *compressed* version of X .

These considerations suggest a *training-as-compressing* analogy, where the training algorithm plays the role of the compression algorithm; the inference (feed-forward) algorithm is the de-compression algorithm (with the input passed to the model working as a decoding key); and the model’s weights represent the compressed version of the training set. In addition to being used for data generation *as-is*, such a *pre-trained* model commonly represents a starting point for additional training: it can be fine-tuned for supporting downstream tasks (Tunstall et al., 2022) or inducing desired behaviors, e.g., to align it with human preferences (Leike et al., 2018). The discussion above can be extended to fine-tuned models as well, with the only difference being that the source data would be both the training set and the pre-trained model’s weights.

3. Training-as-Compressing and Copyright

The *training-as-compressing* perspective can shed new light on the open copyright issues related to generative modeling. While the software code responsible for the training and inference of a generative model can fall under copyright law as a computer program and the algorithmic method is a mathematical model and thus not protected (World Trade Organization, 1994), whether the model’s weights can be protected or not is an open question. Indeed, if the model’s weights represent a compressed version of the training set, and the training set is protected by copyright laws, then the weights are also subject to them. Assuming that the training set is protected in some ways (we will discuss it later), the weights can thus be seen as either a) a lossy or lossless compressed copy of it or b) a compressed version of a derivative work of it.

Seeing the weights as a mere compressed copy of the training set (not different from a zipped file) is seducing since the weights are meant to contain all the information necessary to reconstruct the original samples given a certain input (i.e., the decoding key). However, the final result is usually lossy, and the common scenario is that what we obtain after decompression is similar, but not exactly equal, to the original work. If the differences are not substantial, then it can still be considered a copy; however, it can also lead to a non-negligible modification or transformation of the training data. This second option seems to match the definition of derivative works.

This opens up a different perspective: what the weights represent might not be the original training set, but a new, derivative work (substantially different from, but still based on, the original) whose creation happens concurrently with weights’ learning and whose only existence is due to the weights themselves. Nonetheless, a derivative work must still satisfy the originality requirement to be protected by copyright. Whether or not the trainers’ role in choosing data, algorithms, and parameters is sufficient for claiming authorship (and thus protection) of the model’s weights is still an open question.

Until now, we have assumed that the training set is protected under copyright law. The whole training set can be protected as a database or a collective work, i.e., a collection of separate and independent works (Lee et al., 2024). However, the collective work must constitute an intellectual creation because of the selection and arrangement of its content; the same criteria also apply to databases. One of the current trends for training foundation models seems to go in the opposite direction. Although a certain degree of data pre-processing is always present, the apparent tendency, at least in the early days of foundational models, has been to collect as much data as possible, for example, from the Web. This approach threatens the requirement of making a careful and original selection or arrangement. Moreover, fine-tuning models by means of training sets for specific domains are more likely to be eligible for protection as collective works. Still, this interpretation does not seem to cover all foundation models’ training sets.

On the other hand, single training samples are often protected under copyright law (Bandy & Vincent, 2021). Even though the training goal aims to compress batches of samples at a time, thus potentially leading to a compression that is optimal for a subset of works when considered together but not when considered separately, the single works can still be decompressed from the resulting model, at least in principle. This suggests that the model’s weights can be interpreted as a copy (or a derivative work) of all the independent training samples, and not only of the training set as a whole.

4. Implications

Interpreting the model’s weights as a copy or a derivative work of protected works leads to two crucial implications.

First, it provides a legal framework to understand them, removing the veil of uncertainty surrounding this issue. Although asserting copyright protection for weights as a derivative work presents challenges due to the absence of valid authorship (Otero, 2021), it is possible to safeguard them by viewing the file with the model’s weights as a database. Indeed, they can be considered as a collection of floating

point numbers, which can be retrieved independently. Moreover, the significant investments required for obtaining them make the model’s weights eligible for the *sui generis right* (thus providing certain rights to those who have invested in the database constitution independently from its copyright protection) (Sousa e Silva, 2024). In other words, the *sui generis right* can protect the investment; our copyright perspective can link the model’s weights back to the training data, providing a new perspective over one of the several issues concerning the generative-AI supply chain (Lee et al., 2024). The same considerations still hold in the case of a fine-tuned model. According to Lee et al. (2024), this would be considered a derivative work of the pre-trained model (and also of the fine-tuning data). In other words, fine-tuning could be considered as nothing more than an additional step in the information processing chain. Again, the weights of the fine-tuned model would be eligible for the *sui generis right*. However, whether it qualifies for protection as a derivative work remains an open question, and the determination of valid (human) authorship can vary on a case-by-case basis.

Second, this type of interpretation provides a potential framework for works generated by the model. Indeed, decompressing the information from the model might be seen as producing a derivative work of the weights, thus a derivative work of a copy of a protected work or a derivative work of a derivative work of (a copy of) a protected work. Either way, this link between the output and the training data may help enforce their copyrights. It is worth noting that the EU text and data mining (TDM) exceptions (European Union, 2019) as well as other comparable rules (Fiil-Flynn et al., 2022) apply for TDM purposes, such as training the model, therefore to the case of the creation of a copy or derivative work; however, they do not apply for further derivative works from the model. A similar consideration can also be drawn for the US fair use doctrine, which arguably applies to training a model on copyrighted data but is less likely when deployed to generate similar content that can threaten their market (Henderson et al., 2023). The main consequence is that authorization from the training set’s rightsholders would be required (or else the reproduction or adaptation right would be triggered), allowing for potential requests for compensation from original authors. In addition, generated works would need to respect the moral rights of the owners of training data, even when their economic rights have expired. The fact that a new derivative work is protected by copyright is an entirely different issue that will depend on the human contribution (i.e., the input to the model), in particular, on its substantiality and its being the main contribution of the originality (Guadamuz, 2017; Franceschelli & Musolesi, 2022). The overall conceptual framework based on the proposed *training-as-compressing* perspective is summarized in Figure 2.

5. Related Work

5.1. Deep Learning and Information Theory

Information theory is underpinning machine learning (MacKay, 2003). Recently, there has been considerable interest in the intersection between compression and neural networks, either to learn efficient data compression algorithms (Goyal et al., 2019; Yang et al., 2023) or to compress the learned models themselves (Choudhary et al., 2020; Polino et al., 2018). The problem of training data compression has been analyzed from the point of view of the information bottleneck principle (IB) (Tishby et al., 1999), which provides a methodology for extracting a quantitative measurement of “meaningful” and “relevant” information. In particular, IB can be used to derive an optimal theoretical limit for training neural networks in terms of trade-off between compression and prediction (Tishby & Zaslavsky, 2015). From this perspective, stochastic gradient descent can be seen as composed of two distinct phases: a “fitting” phase in which the mutual information between the weights and the target output increases and a “compression” phase in which the mutual information between the weights and the input decreases (Shwartz-Ziv & Tishby, 2017). However, the latter phenomenon might not be an intrinsic feature of deep neural networks, but only the result of the presence of nonlinearities in the networks (Saxe et al., 2018). An important recent paper in this context is (2024), in which Schwartz-Ziv and LeCun discuss IB in relation to self-supervised learning, especially in the context of multimodal learning; this work is particularly relevant given the prominence of these techniques in training state-of-the-art foundation models.

5.2. Generative AI and Copyright

There has been a long-standing interest in the copyright issues around generative AI (Butler, 1982; Samuelson, 1986). Different legal issues are at play when considering the entire generative-AI supply chain (Lee et al., 2024). Whether training neural networks on protected data is lawful has been highly debated across different national legislations (Lemley & Casey, 2021), e.g., U.S. fair use (Henderson et al., 2023; Sobel, 2017), EU text and data mining exceptions (Geiger et al., 2018; Sobel, 2021), and others (Guadamuz, 2024; Samuelson, 2021). The other main debate has focused on whether a machine-generated work is protected by copyright or not (Craig & Kerr, 2019; Gervais, 2020; Grimmelman, 2016) and on the question of who might own its ownership in the future (Bonadio & McDonagh, 2020; Franceschelli & Musolesi, 2022). However, other issues have also been considered, such as whether the model output can infringe the reproduction right (Gillotte, 2020; Vyas et al., 2023) or how the trained model can be protected (Otero, 2021; Slowinski, 2021). Our analysis mainly concerns the latter




Figure 2. A schematic summary of the legal framework resulting from the *training-as-compressing* perspective. The blue arrows connect potentially protected entities to their copies or derivative works: the foundation model is a copy or a derivative work of the training data; fine-tuning can lead to a new derivative work of the foundation model and the tuning data; and an AI-generated work is a derivative work of either the foundation or the fine-tuned model. The yellow (dotted) and red (dashed) arrows directly link the AI-generated work back to training data and training and tuning data, respectively, only through steps requiring specific exceptions or authors’ authorization.

point, but also provides a new technical understanding and perspective on the copyright of AI-generated works.

6. Conclusion

Motivated by the phenomenon of data memorization in foundation models, we have proposed a perspective we termed as *training-as-compressing* for examining copyright issues. Starting from a thought experiment and a theoretical discussion, we have suggested interpreting self-supervised learning as data compressing and the model’s weights as a compressed version of the entire training set. From a copyright perspective, this framing has led us to consider the weights as a reproduction or a derivative work of training data, which usually contain protected works. We have discussed potential protection that can be applied to the model’s weights considering the underlying direct relationship between the AI-generated outputs and the protected training data.

The analysis conducted in this paper predominantly raises additional questions rather than providing unequivocal solutions. We believe that a multi-disciplinary analysis of these problems and implications from the point of view of information theory is of fundamental importance for practitioners and researchers from both the technological and

legal points of view. Our research agenda also includes a rigorous formalization of the problem as a basis for rigorous legal analysis of this complex yet fascinating area.

References

Bandy, J. and Vincent, N. Addressing “documentation debt” in machine learning: A retrospective datasheet for book-corpus. In *Proc. of the 35th Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.

Bonadio, E. and McDonagh, L. Artificial intelligence as producer and consumer of copyright works: Evaluating the consequences of algorithmic creativity. *Intellectual Property Quarterly* 2020, 2:112–137, 2020.

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y.,

- Lundberg, S., et al. Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv preprint arXiv:2303.12712*, 2023.
- Butler, T. L. Can a computer be an author? copyright aspects of artificial intelligence. *Hastings Communications and Entertainment Law Journal*, 4(4):707, 1982.
- Carlini, N., Tramèr, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, Ú., Oprea, A., and Raffel, C. Extracting training data from large language models. In *Proc. of the 30th USENIX Security Symposium (USENIX Security 21)*, 2021.
- Carlini, N., Hayes, J., Nasr, M., Jagielski, M., Sehwag, V., Tramèr, F., Balle, B., Ippolito, D., and Wallace, E. Extracting training data from diffusion models. In *Proc. of the 32nd USENIX Conference on Security Symposium (SEC'23)*, 2023a.
- Carlini, N., Ippolito, D., Jagielski, M., Lee, K., Tramer, F., and Zhang, C. Quantifying memorization across neural language models. In *Proc. of the 11th International Conference on Learning Representations (ICLR'23)*, 2023b.
- Chomsky, N. *Syntactic Structures*. Mouton and Co., 1957.
- Choudhary, T., Mishra, V., Goswami, A., and Sarangapani, J. A comprehensive survey on model compression and acceleration. *Artificial Intelligence Review*, 53(7):5113–5155, 2020.
- Cover, T. M. *Elements of Information Theory*. John Wiley & Sons, 1999.
- Craig, C. J. and Kerr, I. R. The death of the AI author. *Ottawa Law Review*, 52(1):31–86, 2019.
- Deletang, G., Ruoss, A., Duquenne, P.-A., Catt, E., Genewein, T., Mattern, C., Grau-Moya, J., Wenliang, L. K., Aitchison, M., Orseau, L., Hutter, M., and Veness, J. Language modeling is compression. In *Proc. of the 12th International Conference on Learning Representations (ICLR'24)*, 2024.
- European Union. Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market and amending Directives 96/9/EC and 2001/29/EC, 2019.
- Fiiil-Flynn, S. M., Butler, B., Carroll, M., Cohen-Sasson, O., Craig, C., Guibault, L., Jaszi, P., Jütte, B. J., Katz, A., Quintais, J. P., Margoni, T., de Souza, A. R., Sag, M., Samberg, R., Schirru, L., Senftleben, M., Tur-Sinai, O., and Contreras, J. L. Legal reform to enhance global text and data mining research. *Science*, 378(6623):951–953, 2022.
- Franceschelli, G. and Musolesi, M. Copyright in generative deep learning. *Data & Policy*, 4:e17, 2022.
- Franceschelli, G. and Musolesi, M. On the creativity of large language models, 2023. arXiv:2304.00008 [cs.AI].
- Geiger, C., Frosio, G. F., and Bulayenko, O. Text and data mining in the proposed copyright reform: Making the EU ready for an age of big data? *IIC - International Review of Intellectual Property and Competition Law*, 49: 814–844, 2018.
- Gemini Team, Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A., et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Gervais, D. J. The machine as author. *Iowa Law Review*, 105:2053, 2020.
- Gillotte, J. Copyright infringement in ai-generated artworks. *UC Davis Law Review*, 53:2655, 2020.
- Goyal, M., Tatwawadi, K., Chandak, S., and Ochoa, I. DeepZip: Lossless data compression using recurrent neural networks. In *Proc. of the 2019 Data Compression Conference (DCC)*, 2019.
- Grimmelman, J. There’s no such thing as a computer-authored work—and it’s a good thing, too. *The Columbia Journal of Law & the Arts*, 39(3):403–416, 2016.
- Guadamuz, A. Do androids dream of electric copyright? comparative analysis of originality in artificial intelligence generated works. *Intellectual Property Quarterly*, 2:1–24, 2017.
- Guadamuz, A. A scanner darkly: Copyright liability and exceptions in artificial intelligence inputs and outputs. *GRUR International*, 73(2):111–127, 2024.
- Henderson, P., Li, X., Jurafsky, D., Hashimoto, T., Lemley, M. A., and Liang, P. Foundation models and fair use. *Journal of Machine Learning Research*, 24(400):1–79, 2023.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS'20)*, 33:6840–6851, 2020.
- Kandpal, N., Wallace, E., and Raffel, C. Deduplicating training data mitigates privacy risks in language models. In *Proc. of the 39th International Conference on Machine Learning (ICML'22)*, 2022.
- Lee, K., Ippolito, D., Nystrom, A., Zhang, C., Eck, D., Callison-Burch, C., and Carlini, N. Deduplicating training data makes language models better. In *Proc. of the 60th Annual Meeting of the Association for Computational Linguistics (ACL'22)*, 2022.

- Lee, K., Cooper, A. F., and Grimmelmann, J. Talkin' 'bout AI generation: Copyright and the generative-AI supply chain, 2024. arXiv:2309.08133 [cs.CY].
- Leike, J., Krueger, D., Everitt, T., Martic, M., Maini, V., and Legg, S. Scalable agent alignment via reward modeling: a research direction, 2018. arXiv:1811.07871 [cs.LG].
- Lemley, M. A. and Casey, B. Fair learning. *Texas Law Review*, 99(4):743–785, 2021.
- MacKay, D. J. *Information Theory, Inference and Learning algorithms*. Cambridge University Press, 2003.
- Meta. Introducing Meta Llama 3: The most capable openly available LLM to date, 2024. <https://ai.meta.com/blog/meta-llama-3/> [accessed June 26, 2024].
- Otero, B. G. Machine Learning models under the Copyright microscope: Is EU Copyright fit for purpose? *GRUR International*, 70(11):1043–1055, 2021.
- Polino, A., Pascanu, R., and Alistarh, D. Model compression via distillation and quantization. In *Proc. of the 6th International Conference on Learning Representations (ICLR'18)*, 2018.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'22)*, pp. 10684–10695, 2022.
- Samuelson, P. Allocating ownership rights in computer-generated works. In *Symposium on The Future of Software Protection*, 1986.
- Samuelson, P. Text and data mining of in-copyright works: is it legal? *Communications of the ACM*, 64(11):20–22, 2021.
- Saxe, A. M., Bansal, Y., Dapello, J., Advani, M., Kolchinsky, A., Tracey, B. D., and Cox, D. D. On the information bottleneck theory of deep learning. In *Proc. of the 6th International Conference on Learning Representations (ICLR'18)*, 2018.
- Shanahan, M. Talking about large language models. *Communications of the ACM*, 67(2):68–79, 2024.
- Shwartz-Ziv, R. and LeCun, Y. To compress or not to compress—self-supervised learning and information theory: A review. *Entropy*, 26(3):252, 2024.
- Shwartz-Ziv, R. and Tishby, N. Opening the black box of deep neural networks via information, 2017. arXiv:1703.00810 [cs.LG].
- Slowinski, P. R. Rethinking software protection. In *Artificial Intelligence and Intellectual Property*, pp. 341–362. Oxford University Press, 2021.
- Sobel, B. Artificial intelligence’s fair use crisis. *Columbia Journal of Law and the Arts*, 41:45–97, 2017.
- Sobel, B. A taxonomy of training data: Disentangling the mismatched rights, remedies, and rationales for restricting machine learning. In *Artificial Intelligence and Intellectual Property*, pp. 221–242. Oxford University Press, 2021.
- Sousa e Silva, N. Are AI models’ weights protected databases?, 2024. <https://copyrightblog.kluweriplaw.com/2024/01/18/are-ai-models-weights-protected-databases/> [last access: May 27, 2024].
- Tishby, N. and Zaslavsky, N. Deep learning and the information bottleneck principle. In *Proc. of the 2015 IEEE Information Theory Workshop (ITW)*, 2015.
- Tishby, N., Pereira, F. C., and Bialek, W. The information bottleneck method. In *Proc. of the 37-th Annual Allerton Conference on Communication, Control and Computing*, 1999.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Tunstall, L., von Werra, L., and Wolf, T. *Natural Language Processing with Transformers*. O’Reilly, 2022.
- Vyas, N., Kakade, S. M., and Barak, B. On provable copyright protection for generative models. In *Proc. of the 40th International Conference on Machine Learning (ICML'23)*, 2023.
- Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P.-S., Mellor, J., Glaese, A., Cheng, M., Balle, B., Kasirzadeh, A., Biles, C., Brown, S., Kenton, Z., Hawkins, W., Stepleton, T., Birhane, A., Hendricks, L. A., Rimell, L., Isaac, W., Haas, J., Legassick, S., Irving, G., and Gabriel, I. Taxonomy of risks posed by language models. In *Proc. of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT'22)*, pp. 214–229, 2022.
- World Trade Organization. Agreement on Trade-Related Aspects of Intellectual Property Rights, 1994.
- Yang, Y., Mandt, S., and Theis, L. An introduction to neural data compression. *Foundations and Trends in Computer Graphics and Vision*, 15(2):113–200, 2023.