

PAIS Abstract Submission Cover Page

Andre Ye, Kevin Feng
University of Washington
{andreye, kjfeng}@uw.edu

Both authors of this abstract would like to be considered for a travel stipend to attend the PAIS workshop. If this is not possible, we may also consider attendance if the travel stipend is only provided to one of the authors. Both authors plan to depart from **Seattle, Washington, USA**. More information about the authors are as follows:

Andre Ye is a 2nd-year undergraduate student triple majoring in Computer Science, Philosophy, and Comparative History of Ideas at the University of Washington. His research explores how the distribution of data sources affects the behavior of downstream deep learning models, and how these patterns of influence can be leveraged to amplify desirable model qualities. He is currently working in the Social Futures Lab on uncertainty representation and previously worked on researching training techniques for medical computer vision models at UW Medicine. His philosophical interests include post-humanist ethics, (post-)Marxist philosophy, and critical theory.

Kevin Feng is a 2nd-year PhD in the Human Centered Design & Engineering department at the University of Washington. His research lies at the intersection of social computing and human-centered machine learning. He designs, builds, and evaluates tools to support novel interactions that allow users to harvest capabilities of modern machine learning techniques for their own domain-specific use cases, regardless of their technical expertise. He is supported by a 2022 Herbold Fellowship in Data Science and Computation. Before UW, he received his BSE in Computer Science from Princeton University, where he built tools to democratize research on web advertisements at scale at the Center for Information Technology Policy (CITP).

Generative Models as Simulation of Labor-Power

Andre Ye, Kevin Feng
University of Washington
{andreye, kjfeng}@uw.edu

Recent advances in large-scale Generative Models (GMs) such as GPT-3 and Stable Diffusion have enabled strikingly sophisticated generation of a wide range of traditionally human-made content—including text, art, and music. GMs’ capabilities are learned from massive collections of training data, to the extent that they seem to encapsulate a substantial proportion of human knowledge. The labor congealed in this data is supplied by the ‘public’ in various openly available media contributions. While the Marxian model of the machine may seem applicable here, the GM embodies something much different—a Frankenstein figure: it animates ‘dead’ inert labor snapshots (data) rather than performing a series of fixed transformations; it brings the static data products of the ‘public’ to life into a dynamic model. We argue that GMs uniquely function as *simulations of labor-power*, and that this understanding provides the beginning of an ontological basis to navigate difficult contemporary ethical problems in AI intellectual property and privacy.

Karl Marx, in formulating his theory of capital and production, provides a useful portrayal of the machine as a particular transformation within the production process: it consumes products and produces modifications, akin to a mill’s capacity to transform wheat to flour. Insofar as he believed that the value of a commodity is measured by the labor congealed within it, Marx asserted that machines cannot produce surplus-value because machines do not have *labor-power*. He defines labor-power as an abstraction of concrete labor: it is the capacity or drive of a worker to continue working, even at increasing intensities and in novel production contexts. We can interpret this as the adaptive and sustained generation of labor, to work wherever and however work is needed. While a mill may have been built to grind wheat into flour and does nothing else by way of adaptivity or self-sustenance, it is difficult to argue that GMs behave in a comparable manner. The *generative* dimension of GMs endows them with a sense of labor-power: it is well-established in deep learning research that GMs are successful few- to zero-shot learners on incredibly diverse and previously unseen data. When GMs produce a generated artifact, it is not philosophically nor intuitively just another mechanical information transformation. Thus, GMs embody a novel scale of ‘machine.’ They are fueled by an unprecedentedly wide range of labor products and matches—or even supersedes—the

adaptability and generality of the laborers who contributed such data products. To understand the particular characteristics of GM as ‘neo-machine’, we need to move past Marx.

Jean Baudrillard explored the role of simulation in rupturing (and suturing) the fabric of reality in postmodern society. Baudrillard refutes production-centric Marxian analysis and argues instead that *simulated* representations of reality (e.g., digital information, knowledge industries, mass media) are the primary organizing forms of society. The concept of simulation is highly applicable here—in the sense that labor-power ‘generates’ varied labor, the generative modeling of such varied labor constitutes a simulation of labor-power. GMs simulate aggregated labor-power by embodying fields of inert data as if they were ‘alive.’ Baudrillard’s observation that evolved simulation allows the sign to displace the signified is evident in this context: for instance, GMs can rapidly simulate the labor-power of an artist (its ability to accept a diversity of requests—known more commonly as ‘prompts’—to create ‘new’ works), so much so that models like DALL-E and Stable Diffusion begin to question the relevance of ‘real’ artists. However, Marxian analysis of labor and commodification still appears in the ways in which we engage with GMs as simulated labor-power. We see evidence of this in prompt marketplaces, where intricately designed prompts are sold as methods to coax the GM into generating desirable outputs, thus maximizing the use of (simulated) labor-power. We therefore understand simulated labor-power as jointly informed by Marx and Baudrillard.

Our argument that GMs effectively simulate labor-power by animating ‘dead’ collective labor holds compelling implications for how we should interpret and engage with GMs. Moreover, it provides the crucial ontological edifice upon which subsequent ethical determinations can be derived: a common misstep in AI ethics is to specify an ethics without addressing the underlying metaphysical foundation. Public policy generally agrees that laborers with the capacity for labor-power should be protected; measures such as union and working condition regulation aim to ameliorate the burden of self-reproduction and the exploitation of labor-power. However, we cannot consider the ‘public’ as a laborer while it is actively—and often unknowingly—contributing knowledge to the GM: rather, the GM internalizes inert data products and simulates their continued adaptive generation. We should therefore respect the products created by GMs not as by a machine, nor as by a laborer, but rather as by the unique position of a collective of *simulated laborers*.

One of the most compelling implications of this thesis is that it provides a philosophical justification for research into privacy and intellectual property protection even when there is no clearly attributable human individual or group ('an intellectual') at stake. It emphasizes *simulation* as the principal mediator of GM production processes. This emphasis is particularly relevant to understanding the ongoing lawsuit against GitHub Copilot for violating private code licenses, as well as increasing concerns within artist communities on unauthorized replication and distribution of their work via text-to-image GMs. Prior GM research suggests that GMs do not simply perform copy-paste information retrieval in sufficiently complex domains where 'creative' expression is demonstrable. While it may be infeasible to attribute a GM output to a particular member or group of the 'public', we intuitively sense that something is being repeated or regenerated. This intuition stems from GMs' *simulation* of labor-power. Centering this concept allows us to access new dimensions of traditional IP debates that hinge on the distinctions between 'original' and 'copy', and to channel this nuance into developing more informed social and legal ways to reason about and enrich our intellectual interactions with GMs.