

Foundation Models as Agents of Austerity

Anton Koch

Department of Design
Mainz University of Applied Sciences
Holzstr. 36, 55116 Mainz, Germany
anton.koch@hs-mainz.de

Guilherme Zanchetta

Division of Cyber-Physical Systems
CTI Renato Archer,
NICS-UNICAMP
Campinas, Brazil
zanchetta@nics.unicamp.br

Abstract

This article examines the role of Foundation Models within the contemporary landscape of generative Artificial Intelligence, arguing that they should not be regarded as neutral technological developments, but rather as instruments embedded in the political economy of neoliberal austerity. It contends that these models facilitate the displacement of socially embedded practices, favouring efficiency-oriented, privatised, and platform-governed operations.

Introduction

The concept of the foundation model (FM) in its various implementations dominates the current discourse around generative artificial intelligence (AI). The term refers to generative models trained on vast amounts of data, such as a large language model (LLM) or general pre-trained transformer (GPT). It has gained popularity through services such as chatbots and image, audio, and video generators. In this article, we will argue that the FM is deployed as an automated continuation of the neoliberal idea of austerity and acts as an agent of it. The politico-economic project of austerity can be broadly described as a reduction of public expenditure. It seeks to gradually replace social organisation in state institutions with a market- and game-theory-based framework, aiming for maximum connectivity and information flow with minimal outside interference. Historically, the LLM emerged in the wake of a global intensification of austerity measures following the 2008 financial crisis. It is beyond the scope of this article to determine whether this development was purely coincidental or benefited from quantitative easing measures, which released vast amounts of low-to-no-risk capital into the tech sector, allowing massive stock buybacks and infrastructure buildouts. However, the broad scope of use and inherent logic of big-data stochastic working, which is similar to managerial decision-making, make it a good fit. While the technology enables the automation of general administrative tasks, it has also begun to impact a growing number of workers in the creative disciplines of writing, illustration, and photography. In some cases, creative workers are now merely required to “clean up” the automated output, thus being progressively alienated from their practice. Their work is neither replaced nor augmented, but devalued, largely phased out and replaced with generic

output, similar to many professions previously displaced by capitalist automation logic.

Basic concepts

To situate our critique of FMs, we need to address the specific terminology used when talking about these models and their purpose. However, it also requires us to clarify our frame of reference regarding generative AI in general as well as the FM in particular. We acknowledge the long and complicated history of the general field of AI with its numerous approaches to the concept of machine intelligence. For this article, we are focusing explicitly on the FM as being “trained on broad data (generally using self-supervision at scale) that can be adapted (e.g., fine-tuned) to a wide range of downstream tasks” (Bommasani et al. 2022, p. 3). These models, while not new in terms of their underlying methods and technologies, are unprecedented in scale and scope and have been referred to as foundation models by Stanford researchers to indicate a “paradigm shift” (Bommasani et al. 2022, p. 6) in the field of AI.

Throughout this article, *Intelligence* is broadly defined as the appropriation of information, its transformation into knowledge and the latter’s application. However, it should also be seen as an emergent and dynamic property of any complex system that is highly context-dependent. In that sense, we argue for referring to, as physicist Frank Wilczek has called it, an “ecology of intelligence” (Frank Wilczek 2019), a plethora of diverse tapestries of artificial and human intelligence. This view rejects the stereotypical notion of a hierarchical, linear intelligence quotient (IQ) scale, which is a profoundly flawed concept (Geburu and Torres 2024).

Most importantly, we ground our understanding of *computation* as a formalist practice, encompassing the execution of mathematical operations, algorithmic symbol manipulation, and stochastic evaluation. While this function is now primarily tied to the machinic computer, it is worth noting that this practice did not emerge with the machine, but pre-dates it by far. It must also be noted here that, while we employ this broad definition, it is essential to consider the specificity of computation concerning its cultural framework, as laid out by scholar David Golumbia (Golumbia 2009). At the same time, computation has been shown to transcend the simple definition of an automated mode of calculation or simple digital mechanics in terms of its metaphysics (Brian

Cantwell Smith 1996) and aesthetics (Fazi 2018). While these discussions by far exceed the scope of this article, they present an important horizon for an overall computational potential.

The broadest definition of what AI constitutes is a project that strives to automate reason. To this end, it deploys computation to realise a form of quasi-autonomous decision-making within a machine. This includes rule-based, symbolic AI as well as the currently popular connectionist variants of the artificial neural network (ANN) and deep neural network (DNN). Here, we follow Matteo Pasquinelli's historical account of AI's development from a specific, capitalist organisation of labour (Pasquinelli 2023), as well as Dan McQuillan's socio-political critique of its deployment (McQuillan 2022).

Finally, we understand *creativity* as the production of ideas and their active exploration and reflection. While the term can be used in many contexts, we specifically refer to the labour performed within the contemporary creative industries.

Background

An abbreviated history of creative computation

Using computational machines and methods in art has a longstanding history, but has been particularly influential within the field of “computer art”, which began in the second half of the 20th century. There have been early examples of geometric and mathematical art by pioneers like Vera Molnár and Manfred Mohr. Early forms of data-driven decision-making in the creative process are exemplified by works of artists like John Cage, with what the composer called chance operations (Nicholls 2002) and Iannis Xenakis's Stochastic music (Xenakis 1966). Later on, notable approaches to the simulation of “artificial life” were developed by Karl Sims and William Latham. Markov chains have been extensively used to generate sequential and interactive compositions (Manzollí 1996) (Manzollí and Baruco Machado Andraus 2022) Lastly, AI art has a history that extends beyond today's generative models, with artists such as Helena Sarin and Mario Klingemann utilising generative adversarial networks (GANs) to create original works.

Computationalism, Anthropomorphism and the Humanist Trap

There are symptoms that, although entangled, can be analytically located: computationalism—understood not merely in its naïve form as the assertion that “everything is a computer”, but more rigorously as computation (following the elaborations of digital philosophers and thinkers of the philosophy of computation)—and anthropomorphism, the projection of human characteristics onto non-human entities. Both tendencies are embedded in and amplified by the cybernetic apparatus (Geoghegan 2011), which encompasses mathematical procedures, diagrammatic strategies, and technological tools that have functioned both as material infrastructures for research and as quasi-transcendental ideals. This apparatus has enabled a mode of abstraction

that displaces historical, political, and disciplinary specificities, while simultaneously facilitating alliances among researchers and institutions across diverse borders. In this epistemic regime, cybernetics does not approach machines by asking “what individual act will it produce here and now?” but rather “what are all the possible behaviours it can exhibit?”—a framing aligned with information theory, which concerns itself not with particularities but with sets of possibilities. Energy, once central in biological explanation, becomes a given in cybernetic reasoning; the focus shifts to the constraints and regulatory mechanisms that delimit developmental trajectories. Whether a system is energetically open or closed is often secondary to the degree to which its behaviour is structured by controlling factors within a broader space of potential outcomes. In this regard, cybernetics positions itself in relation to material systems—electronic (Ashby 1956), mechanical, biological, economic, or political—much as geometry once related to terrestrial objects: just as geometry moved from empirical demonstrations in Euclidean space to the formal coherence of non-Euclidean and abstract geometries, cybernetics generalises from instantiated machines to the conceptual space of all possible machines. What matters is not the specific mechanism but the architecture of behaviours it can potentially exhibit. This shift toward abstraction had a retroactive impact on biology itself, reconfigured by cybernetic thought, while, paradoxically, engineers began deploying biological metaphors for machinic concepts and implementations (Caporael 1986). Over time, these metaphors also feed back into the sciences, as has been particularly effective with the idea of the mind as a computer (Pasquinelli 2021). The politics of humanism within the western imperialist system and its project of science (as envisioned by Francis Bacon) produced an anthropocentric, exceptionalist and hierarchical worldview with the human (more specifically, the white and male human) in the centre. While the Enlightenment ostensibly sought to ground everything in reason, it did so according to the existing demarcations of imperialist, scientific hierarchies. Now, Transhumanism seeks to overcome the human condition entirely, sees it as an intermediate step, thereby clinging to the same linear, hierarchical exceptionalism of a singular intelligent quality.

The limits of AI

Connectionist AI is characterised, by some of its proponents, as a universal technology (Eloundou et al. 2023), but it is deeply indebted to a particular epistemological concept and thereby limited within that specificity. To situate a critique of AI, it is vital to frame it within its epistemological and functional limitations and thus to decouple it from its sprawling mythologies.

The idea of the neural network goes back to the biologist McCulloch's and logician Walter Pitts' publication of their “logical calculus of the nervous system” (McCulloch and Pitts 1943), which proposed a formalised model derived from a theory on the function of vision. This theory of the eye as a networked processor and classifier of sensory information to be forwarded to the brain was immensely influential, as the economist Friedrich von Hayek also picked it up

for his theory of mind (Hayek 1952). It was Hayek's work that would come to broadly influence both an economic formalisation of markets as cellular networks of signal flow and a general understanding of the mind as a computer. Scholar Matteo Pasquinelli has described this entry of a technical paradigm into the natural sciences as an "ideological coup d'état" (Pasquinelli 2021, p. 161). While it may be debatable whether Hayek was somehow consciously aiming to naturalise economic theory, it is evident that the connectionist idea was a precursor for the upcoming so-called "information age". It also influenced Frank Rosenblatt, who credits Hayek in his paper on the Perceptron, probably the first machine implementation of an ANN (Rosenblatt 1958, p. 388). The Perceptron symbolises the convergence of connectionism with the stochastic reasoning through classification by the "regression to the mean" as introduced by Francis Galton. Both ideas are rooted within the Western, colonial project of humanist science and, as such, are powerful tools for the construction of neoliberal ideas around societal self-organisation through uninhibited connectivity, stochastic management and free flow of capital and information.

AI's functional limitation originates in its particular specificity modelled as a machine for computer vision (CV). It is not conceived as a brain per se, but a specialised organ for vision and classification of patterns, as scholar James Dobson details in his history of CV (Dobson 2023). In its large form, today's models are made to "see" within large quantities of data and respond in a synthesised way (e.g. natural language, images), but also with a certain "lossyness", as argued by author Ted Chiang (Chiang 2023). In opposition to any analytical, constructive process, it merely performs a brute-force approach at fitting myriads of weak associations to a pattern and then reproducing these patterns. While this certainly leads to impressive results at scale, it is still functionally limited to information storage and retrieval. This is not to dismiss the potential of using large models for a variety of tasks, but these are generally performed with a large overhead of computing power compared to specialised algorithms, as Huggingface's lead researcher, Sasha Luccioni, has detailed (Luccioni, Jernite, and Strubell 2024).

The Perceptron ultimately failed to deliver on what was already the first wave of excitement over the prospect of supposedly thinking machines, ushering in what was called the "AI winter" and taking a back seat to symbolic AI. It took until the beginning of the 21st century for the resurgence of connectionist AI with the advent of AlexNet in 2012, to the current boom triggered by the release of ChatGPT in 2022. However, this resurgence is not necessarily due to an evolution in the machine learning technique itself, as deep learning has been around at least since the 1990s. The exponential growth in feasible model sizes and dimensionality was made possible by the discovery of graphics processing units (GPUs) technology's capability to accelerate parallel computing (Nicole Hemsoth 2015) and a massive buildout of datacenter infrastructure during the last two decades after the advent of cloud computing. The models we see today are running at an unprecedented technological scale, but do so at an even more massive cost in compute power, data labour and mass-scale appropriation of content.

Generative AI as a form of automated austerity

Generative AI works on the surface, stripping away the underlying processes and recombining based on billions of arbitrary, weak connections. It automates the preemption of the social production process through stochastic prediction. Creation becomes consumption as it is reduced to natural-language search queries to recombine existing works. Intelligence, here, is falsely interpreted as a functional, algorithmic procedure of information retrieval and recombination instead of a complex assemblage of social practices and relations. It also perpetuates a corporate enclosure of the co-creative practice, in that it offers a puppet, or a *potemkin artist*, as a co-creative other, but essentially capitalises the practice itself. Then, as creation becomes more dependent on AI infrastructure, it creates a fundamental dependency that allows rent extraction on the ability, not the tools, to perform a practice like writing or illustrating. As austerity seeks to replace the public with the private, the communal with the proprietary, to invigorate the markets, these models strip away by predictively preempting the creative process that might otherwise involve social interaction, reflection and experience.

Foundations for critique and opposition

We need to work from a materialist base to situate and critique dominant technological trends. While there is speculation on "existential risk" or utopian predictions of "abundance", it is clear that AI as an economy, as a socio-technical apparatus, is already impacting countless lives worldwide in very different ways. These experiences, of data-labourers, rare-earth miners, but also writers and illustrators now only hired to clean up generative AI output, should be guiding the discussion around it, along with its vastly unsustainable energy and resource demands. There needs to be a categorical rejection of corporate black-box technologies that cannot be replicated either due to intellectual property rights but also because of the sheer amount of resources needed to reproduce them. It is also important to critically focus on questions of authorship in the face of unlawfully acquired content for training, as well as ownership of a practice. The way these models are currently implemented, they function as an enclosure of social production and the corporate appropriation of creative practice. This threat is similar to past struggles, not against technological progress, but the displacement, deskilling and alienation of specialist workers, such as in the Luddite movement in the early 19th century. As such, it is also clear that, even if the technology may be impressive, sometimes it still needs to be rejected wholesale. But we also acknowledge that this rejection often is not an option if employers require the workers to interact with it, or project funding is dependent on its implementation. This is why a resistance to technologies of austerity is not just brought on by conscious consumption, by refusing to participate in the hype, but through labour movements and political activism. As the FM is an agent of austerity, it is still the most viable strategy to push back against austerity politics instead of its symptoms.

Deep Environments over Deep Networks

Beyond the materialist critique of generative AI, but also the general technology of late-stage capitalism, there is another radical and powerful approach. The philosopher Sylvia Wynter challenges the entire singular foundation of the underlying science in that she advocates for, as scholar Walter D. Mignolo calls it, “epistemic disobedience” (Mignolo 2015, p. 106). This perspective serves as a starting point for acknowledging the specificity of Western science as it emerged within the colonial framework of exploration, classification, and quantification for optimal extraction and exploitation. This does not mean discarding science as a whole because of its fraught and complicated history, but understanding it as one epistemological model among others. Likewise, the FM represents a very narrow and specific epistemic construct, and everything built on it will be unable to go beyond this foundation in a fundamental way. However, if we remain conscious of its limitations and specificity, we can weigh it against alternative methods and models that pose true alternatives to merely switching out the training data. This article is not a call for the abolishment of contemporary technology and its replacement with an ethically or functionally superior variant or the return of supposedly “simpler times”. It is rather an invitation to try and think beyond its narrow frame of future possibility. Instead of building a unified, enclosed “Metaverse” or striving for a global “Singularity”, we prefer a pluriverse of hyperlocal digital infrastructures and practices. These exist as dynamic reflections of and extensions to their respective environment and culture, instead of aiming to establish normative and rigid frameworks or platforms. We prefer strata over stacks, relations over connections and variance over the mean. The philosopher Yuk Hui uses the term “Technodiversity” to describe an extended frame of understanding technology as culturally specific (Yuk Hui 2023), and thus open to being extended in its material and metaphysical potential by moving away from a singular, solutionist mindset. This concept is powerful in that it suggests not viewing code or machinery simply as tools or pragmatic solutions for perceived problems. But it also requires us to perceive our networks and assemblages of humans and machines neither anthropocentric, as the image of the cyborg, nor as the hierarchical master-and-slave duality. These networks are not characterised by the maximisation of information flow and frictionless efficiency as in the connectionist framework, but rather by kinship and care. Kinship, as Donna Haraway has proposed (Haraway 2016), can be felt within a non-hierarchical environment of people, creatures, plants and could even be thought to extend towards inhuman entities as much as its connectedness is not predicated on market value exchange, but on mere acknowledgement, awe or wonder. On a social level, this means to engage in networks of care that favour processes that foreground social relationships, deploying technology not as a mediator of communication but communication and expression in itself. Engaging in smaller, local digital ecosystems allows us to focus on the complex needs of actual local communities instead of generalised assumptions towards a stereotypical, fictional target audience. Finally, this also requires embrac-

ing uncertainty and ambivalence as axiomatic principles in design, engineering and creative processes to keep these dynamic and adaptable to environmental changes.

Conclusion

This article seeks to highlight the role of FMs as agents of austerity politics. To this end, we begin by clarifying the conceptual coordinates of intelligence, creativity, computation, and artificial intelligence. By tracing the epistemological and functional limitations of connectionism and computationalism, we argue that generative AI — particularly in its instantiation as the FM — is not merely a technological development, but a crystallised expression of neoliberal austerity. Through stochastic pattern-matching and the automation of creative and intellectual labour, FMs replicate the core logic of austerity: the erosion of socially embedded practices in favour of privatised, efficiency-oriented, and market-driven operations. Each interaction with an FM entails a tacit collaboration between the user, the model’s contested training data, and the corporate entity that owns and governs both. Rather than enhancing creative agency, this configuration repositions the user within a closed feedback loop where authorship becomes secondary to the logic of data capture and platform optimisation. To this end, we have first set out to clarify our conceptual understanding of intelligence, creativity, computation and AI. We have shown the origins and limitations of connectionist AI and computationalism in their epistemological and functional specificities. On this basis, we attempt to demonstrate how generative AI in the form of the FM functions as an agent of austerity, reproducing common patterns of capitalist and neoliberal logic. This dynamic echoes itself in its asymmetrical co-creation; every interaction tacitly endorses a three-way collaboration between the user, the model’s contested data, and the corporate entity governing both. Finally, we have offered some pointers to valid starting points to develop a materialist critique of these technologies and proposed to develop alternative imaginations and interpretations of the digital.

References

- Ashby, W. R. 1956. *An introduction to cybernetics* / william ross ashby.
- Bommasani, R.; Hudson, D. A.; Adeli, E.; Altman, R.; Arora, S.; Arx, S. v.; Bernstein, M. S.; Bohg, J.; Bosselut, A.; Brunskill, E.; Brynjolfsson, E.; Buch, S.; Card, D.; Castellon, R.; Chatterji, N.; Chen, A.; Creel, K.; Davis, J. Q.; Demszky, D.; Donahue, C.; Doumbouya, M.; Durmus, E.; Ermon, S.; Etchemendy, J.; Ethayarajh, K.; Fei-Fei, L.; Finn, C.; Gale, T.; Gillespie, L.; Goel, K.; Goodman, N.; Grossman, S.; Guha, N.; Hashimoto, T.; Henderson, P.; Hewitt, J.; Ho, D. E.; Hong, J.; Hsu, K.; Huang, J.; Icard, T.; Jain, S.; Jurafsky, D.; Kalluri, P.; Karamcheti, S.; Keeling, G.; Khani, F.; Khattab, O.; Koh, P. W.; Krass, M.; Krishna, R.; Kuditipudi, R.; Kumar, A.; Ladhak, F.; Lee, M.; Lee, T.; Leskovec, J.; Levent, I.; Li, X. L.; Li, X.; Ma, T.; Malik, A.; Manning, C. D.; Mirchandani, S.; Mitchell, E.; Munyikwa, Z.; Nair, S.; Narayan, A.; Narayanan, D.; Newman, B.; Nie,

- A.; Niebles, J. C.; Nilforoshan, H.; Nyarko, J.; Ogut, G.; Orr, L.; Papadimitriou, I.; Park, J. S.; Piech, C.; Portelance, E.; Potts, C.; Raghunathan, A.; Reich, R.; Ren, H.; Rong, F.; Roohani, Y.; Ruiz, C.; Ryan, J.; Ré, C.; Sadigh, D.; Sagawa, S.; Santhanam, K.; Shih, A.; Srinivasan, K.; Tamkin, A.; Taori, R.; Thomas, A. W.; Tramèr, F.; Wang, R. E.; Wang, W.; Wu, B.; Wu, J.; Wu, Y.; Xie, S. M.; Yasunaga, M.; You, J.; Zaharia, M.; Zhang, M.; Zhang, T.; Zhang, X.; Zhang, Y.; Zheng, L.; Zhou, K.; and Liang, P. 2022. On the Opportunities and Risks of Foundation Models. arXiv:2108.07258 [cs].
- Brian Cantwell Smith. 1996. *Computation*.
- Caporael, L. 1986. Anthropomorphism and mechanomorphism: Two faces of the human machine. *Computers in Human Behavior* 2(3):215–234.
- Chiang, T. 2023. ChatGPT Is a Blurry JPEG of the Web. *The New Yorker*. Section: annals of artificial intelligence.
- Dobson, J. 2023. *The Birth of Computer Vision*.
- Eloundou, T.; Manning, S.; Mishkin, P.; and Rock, D. 2023. GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. arXiv:2303.10130 [econ].
- Fazi, M. B. 2018. *Contingent Computation: Abstraction, Experience, and Indeterminacy in Computational Aesthetics*. Rowman & Littlefield.
- Frank Wilczek. 2019. *Ecology of Intelligence*.
- Geburu, T., and Torres, E. P. 2024. The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. *First Monday* 29(4).
- Geoghegan, B. D. 2011. From Information Theory to French Theory: Jakobson, Lévi-Strauss, and the Cybernetic Apparatus. *Critical Inquiry* 38(1):96–126.
- Golumbia, D. 2009. *The Cultural Logic of Computation*. Harvard University Press. Google-Books-ID: 9yAsEAAAQBAJ.
- Haraway, D. J. 2016. *Staying with the Trouble: Making Kin in the Chthulucene*. Duke University Press.
- Hayek, F. A. 1952. *The Sensory Order: An Inquiry Into the Foundations of Theoretical Psychology*. Routledge & Kegan Paul.
- Luccioni, A. S.; Jernite, Y.; and Strubell, E. 2024. Power Hungry Processing: Watts Driving the Cost of AI Deployment? In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 85–99. arXiv:2311.16863 [cs].
- Manzolli, J., and Baruco Machado Andraus, M. 2022. Jardim das cartas: an installation as a presence ecology. In *Proceedings of the 10th International Conference on Digital and Interactive Arts, ARTECH '21*. New York, NY, USA: Association for Computing Machinery.
- Manzolli, J. 1996. Auto-organização: um paradigma composicional. In Debrun, M.; Gonzales, M. E. Q.; and Júnior, J. O. P., eds., *Auto-organização: estudos interdisciplinares*. Campinas: Editora CLE. 417–435.
- McCulloch, W. S., and Pitts, W. 1943. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics* 5(4):115–133.
- McQuillan, D. 2022. *Resisting AI: An Anti-fascist Approach to Artificial Intelligence*. Policy Press.
- Mignolo, W. D. 2015. Sylvia Wynter: What Does It Mean to Be Human?
- Nicholls, D., ed. 2002. *The Cambridge Companion to John Cage*. Cambridge Companions to Music. Cambridge ; New York: Cambridge University Press.
- Nicole Hemsoth. 2015. Inside The Programming Evolution of GPU Computing.
- Pasquinelli, M. 2021. How to Make a Class. *Qui Parle* 30(1):159.
- Pasquinelli, M. 2023. *The Eye of the Master: A Social History of Artificial Intelligence*. Verso Books.
- Rosenblatt, F. 1958. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review* 65(6):386–408.
- Xenakis, I. 1966. The origins of stochastic music 1. *Tempo* (78):9–12.
- Yuk Hui. 2023. Rethinking technodiversity.