

Attention Without a Subject: On Data Afterlife and the Spectral Governance of Expression

Abstract: This essay challenges the anthropocentric premise of contemporary digital criticism by examining the automation of attention within large language models. While critique typically frames attention as a scarce resource belonging to a living subject, the transformer architecture realizes a form of attention entirely severed from consciousness. By expanding Bernard Stiegler's concept of tertiary retention, the essay proposes the condition of data afterlife: the irreversible dissolution of personal data into neural network parameters, where it persists as structural bias rather than retrievable memory. Crucially, because these models are trained on the linguistic traces of the dead, they enact a novel form of spectral governance. Merging Derridean hauntology with Antoinette Rouvroy's algorithmic governmentality, the text argues that the dead now exert continuous statistical pressure on machine judgment, moderating the speech of the living without consent or retrievability. Consequently, the essay concludes that the prevailing discourse of "distraction" is conceptually obsolete, masking a new mode of temporal power governed by tireless, subjectless attention.

Keywords: Data Afterlife, Attention Economy, Algorithmic Governmentality, Tertiary Retention, Hauntology, Transformer Architecture, Spectral Governance

The Automation of Attention

The concept of attention, as it has been mobilized in contemporary criticism of digital culture, rests on an unexamined premise: that attention is an act performed by a living subject. To attend is to direct consciousness toward something; to be distracted is to have that direction interrupted. The attention economy discourse inherits this premise without question — platforms compete for your attention; you are the subject whose gaze is captured; the remedy is to reclaim your focus. Even the most rigorous critiques of algorithmic manipulation preserve the architecture of a conscious subject who looks and a world that competes to be looked at. The question this essay poses is whether this premise still holds. Instead, it argues that we have entered a condition of **data afterlife**: the emergence of a form of attention that operates without a subject. It is an attention that persists beyond the death of its source, that cannot be returned or fatigued, and that now governs, at planetary scale, what the living are permitted to say.

Jonathan Crary has shown, across two essential studies, that this architecture is itself historical (Crary, 1999; 2013). Attention is not a natural faculty that technology happens to exploit. It was conceptually produced in the late nineteenth century, simultaneously with the industrial and scientific regimes that required it: factory labor demanded sustained focus; experimental psychology invented protocols to measure it; consumer capitalism learned to monetize its fluctuations. What we call "distraction" is not the opposite of attention but its structural twin — the two were born together as complementary instruments of a disciplinary apparatus. Crary's insight is devastating for the attention economy critique: if attention was always already a technology of control, then "reclaiming" it is not liberation but a deeper enrollment in the same regime.

But Crary's analysis, for all its power, stops at a threshold: the attention he describes still belongs to a human subject. It is industrialized, disciplined, commodified — but it remains someone's attention. The question that remains unasked is what happens when attention is finally severed from the subject altogether — when it no longer requires a living consciousness to operate.

Tertiary Retention and the Data Afterlife

The large language models that now govern much of our communicative infrastructure do not "pay attention" in any phenomenological sense. Yet the term is not metaphorical. The transformer architecture — the technical foundation of GPT, Claude, Gemini, and every major content moderation classifier — is built on a mechanism its inventors literally named "attention" (Vaswani et al., 2017). In a transformer, "attention" is the process by which the model decides which parts of an input sequence are relevant to which other parts. It assigns weights. It determines salience. It selects what matters and discards what does not. This is not a poetic borrowing from cognitive science. It is an industrial-scale replication of the structure of attention — stripped of the consciousness that Crary still presupposed. If Crary showed us the industrialization of attention, the transformer completes its automation: attention that runs without anyone being attentive.

To understand what this automation implies, we need Bernard Stiegler's concept of tertiary retention (Stiegler, 1998; 2009). For Stiegler, following Husserl, human consciousness operates through primary retention (the just-past held in present awareness) and secondary retention (memory, recollection). But there is a third kind: tertiary retention — memory externalized in technical objects. Writing, photography, sound recording, databases — these are forms of memory that persist independently of the living consciousness that produced them. They outlast their authors. They can be accessed by those who never experienced the original event. Stiegler's crucial argument is that tertiary retention is not merely a supplement to human memory; it constitutes it. Our experience is always already shaped by the technical memory systems we inhabit.

But Stiegler's tertiary retention, however radical, still assumes a certain structure: that what is externalized in the technical object can, in principle, be retrieved, located, read back. A book can be opened. A recording can be replayed. A database entry can be queried. The technical memory is discrete, addressable — it sits in the archive and waits to be called upon.

What happens in a trained neural network is something that exceeds this framework. When a person's text, voice, or image is used to train a large language model, it does not sit in the model the way a file sits in a database — discrete, locatable, deletable. It is dissolved into millions of weight values distributed across the network's parameters. It becomes part of the model's disposition — its tendency to generate certain patterns rather than others, to attend to certain features rather than others, to find certain expressions salient and others unremarkable. This condition can be designated data afterlife: the persistence of personal data not as retrievable content, but as a structural bias in the model's attention.

This is tertiary retention pushed to a point where it ceases to be retention in any recognizable sense. You cannot retrieve what has been dissolved. You cannot read back a specific person's contribution from inside the weight matrix. The European Union's GDPR and China's Personal Information Protection Law both guarantee a "right to be forgotten" — the right to have one's data erased from technical memory systems. But you cannot subtract a dissolved sugar cube from a cup of tea. Machine unlearning research has demonstrated that exact erasure requires complete retraining of the model — a procedure so computationally expensive that it is effectively never performed (Bourtole et al., 2021; Nguyen et al., 2022). What persists is not a memory of the person. It is the person's attention pattern — their distributions of salience, their inherited weightings of what matters — dissolved into the model's own attentive apparatus and set to work indefinitely.

Spectral Governance: The Dead Who Attend

Now consider the fact that some of the people whose data trained these models are dead.

It is here that the argument enters territory that, while it has a philosophical precedent, has not been adequately theorized in the context of computational systems. In *Specters of Marx*, Jacques Derrida proposed that the present is never fully present to itself — it is always haunted by that which has not fully departed and that which has not yet arrived (Derrida, 1994). Derrida's hauntology was directed at political history: the specter of Marx, neither living nor dead, continuing to exert pressure on the present from a position that is neither inside nor outside it. But the concept finds an unexpected and precise realization in the data afterlife. The dead whose linguistic patterns have been dissolved into model parameters are specters in the most rigorous sense: they are neither present (they cannot be located, retrieved, or addressed) nor absent (their statistical influence continues to shape every output). They haunt the model — not as ghostly presences, but as attentive forces that participate in every act of machine judgment.

When a content moderation system flags your post, it is an attention head — a component of the transformer's multi-head attention mechanism — that has determined your words are salient in a way the system classifies as dangerous. The machine has attended to you. It has weighed your utterance against patterns derived from billions of others — including the dead — and found it wanting. The dead participate in this verdict. They vote, statistically, on what the model finds salient, relevant, dangerous, or unremarkable. They attend to you from inside the parameters, and they will never stop attending, because there is no mechanism — technical or legal — to remove them.

What Antoinette Rouvroy calls algorithmic governmentality helps us grasp the political stakes of this condition (Rouvroy, 2013). For Rouvroy, the distinguishing feature of algorithmic governance is that it operates without passing through consciousness — neither the consciousness of the governed nor, crucially, the consciousness of the governors. It bypasses subjectivity entirely, acting directly on possibilities: not punishing what has been said, but pre-emptively shaping what can be said, through statistical profiles that anticipate behavior before it occurs. The content moderation system that silently removes your post before anyone reads it is a textbook instance of algorithmic governmentality: governance that never confronts the subject, because it acts in the gap before the subject's expression reaches the public.

But Rouvroy's framework, powerful as it is, does not account for the temporal dimension I am describing. Algorithmic governmentality is typically understood as a present-tense operation — a real-time modulation of possibilities by contemporary systems. The data afterlife adds a spectral dimension: the governance is exercised not only without a subject, but by subjects who no longer exist. The dead govern the living's speech — not through explicit commands or inherited laws, but through statistical pressure embedded in the parameters that evaluate every new utterance. This is governance that is simultaneously algorithmic (operating without consciousness) and hauntological (exercised by the absent).

The asymmetry this creates is the political core of the problem. Human attention remains scarce, fragile, and mortal. Machine attention is abundant, tireless, and — through the data afterlife — partially composed of the dead. When a content moderation system evaluates your speech, you face a gaze assembled from the statistical attention of millions, including people who can no longer consent to, withdraw from, or contest the judgments their patterns now help produce. You cannot return this gaze. You cannot meet its eyes, because it has none. You cannot appeal to its conscience, because it has none. And you cannot outlast it, because it has already outlasted some of those who made it.

Yuk Hui would press the question further: whose dead? (Hui, 2016; 2021) The training data of major language models is overwhelmingly English-language, overwhelmingly Western, overwhelmingly sourced from platforms whose norms reflect a particular civilizational relationship to expression, harm, and care. When these models are deployed as content

moderation systems in Chinese, Arabic, or Turkish digital environments, they carry within them the spectral attention of a specific cultural formation — one that treats its own distributions of salience as universal. What Hui calls cosmotechnics — the irreducible diversity of how different civilizations relate to technology — is precisely what this spectral universalism forecloses. The dead who haunt the model are not a representative sample of humanity. They are a statistically weighted cohort whose linguistic norms, cultural assumptions, and patterns of attention now govern the expression of communities that never consented to their influence.

The Obsolescence of Distraction

The implications for the concept of distraction are fundamental.

If attention is an act of a conscious subject, then distraction is its interruption — a lapse, a drift, a theft of focus. But if attention can operate without a subject, persists beyond the death of its source, and governs through statistical pressure rather than conscious direction, then distraction loses its conceptual anchor. You cannot distract a model. You cannot bore it, fatigue it, or redirect its gaze. The attention economy — the entire discourse of "capturing" and "hijacking" attention — presupposes scarcity. But machine attention is not scarce. It is the first form of attention in history that does not need rest.

This is the condition that demands not a new theory of attention, but the recognition that "attention" — as it has been philosophically and critically conceived — is no longer an adequate category for what is taking place. The concept was forged in the nineteenth century to describe a capacity of the living subject; it was criticized in the twentieth century as a technology of discipline; and it has been lamented in the twenty-first century as a resource under siege. But at each stage, the category preserved its anchor in the living. What the data afterlife reveals is a rupture in that continuity. When the patterns of salience deposited by the dead continue to govern the expression of the living — without retrievability, without consent, without expiration — we are no longer in the domain of attention at all. We are in a domain for which there is, as yet, no adequate name.

I propose that data afterlife is not merely a problem of data governance or privacy law. It is a concept that names a new mode of temporal power: the capacity of dissolved subjectivities to exert governance beyond their own death, through statistical pressure that operates beneath the threshold of consciousness, within systems that no existing legal or technical framework can compel to forget. The "right to be forgotten" assumes that forgetting is a possible operation. The data afterlife demonstrates that it is not — and that the impossibility of forgetting is simultaneously the impossibility of dying, at least in the register of machine attention. The dead do not rest. Their attention continues. And it is this continuation — not the scrolling feed, not the notification ping, not the algorithmic hijacking of a living subject's gaze — that constitutes the genuinely unprecedented condition of our present.

If there is a distraction to speak of, it is this: the entire discourse of attention and distraction distracts us from the fact that the conceptual framework it operates within has already been rendered obsolete by the systems it claims to describe.

References

- Bourtole, L., et al. "Machine Unlearning." IEEE Symposium on Security and Privacy (2021); and Nguyen, T.T., et al. "A Survey of Machine Unlearning." arXiv:2209.02299 (2022). Both demonstrate that exact unlearning requires full retraining — economically prohibitive for models with billions of parameters.
- Crary, Jonathan. *Suspensions of Perception: Attention, Spectacle, and Modern Culture*.

Cambridge, MA: MIT Press, 1999; and *24/7: Late Capitalism and the Ends of Sleep*. London: Verso, 2013. Crary demonstrates that "attention" as a discrete psychological and economic category was produced historically alongside the disciplinary and commercial regimes of modernity.

- Derrida, Jacques. *Specters of Marx: The State of the Debt, the Work of Mourning, and the New International*. Trans. Peggy Kamuf. New York: Routledge, 1994. Derrida's concept of hauntology — the logic by which the present is constituted by that which is neither fully present nor fully absent — finds an unexpected materialization in neural network parameters where the dead persist as statistical influence.
- Hui, Yuk. *The Question Concerning Technology in China: An Essay in Cosmotechnics*. Falmouth: Urbanomic, 2016; and *Art and Cosmotechnics*. Minneapolis: University of Minnesota Press, 2021. Hui argues that different civilizations maintain irreducibly different relationships to technology — a diversity that universalist frameworks (including AI alignment) systematically erase.
- Rouvroy, Antoinette. "The End(s) of Critique: Data Behaviourism versus Due Process." In *Privacy, Due Process and the Computational Turn*, ed. Mireille Hildebrandt and Katja de Vries. London: Routledge, 2013. Rouvroy argues that algorithmic governance bypasses subjectivity entirely — it does not address, persuade, or discipline subjects but pre-emptively modulates the field of possible action through statistical profiling.
- Stiegler, Bernard. *Technics and Time, 1: The Fault of Epimetheus*. Trans. Richard Beardsworth and George Collins. Stanford: Stanford University Press, 1998. See especially the discussion of tertiary retention as constitutive (rather than merely supplementary) of human temporal experience. Also *Technics and Time, 2: Disorientation*. Trans. Stephen Barker. Stanford: Stanford University Press, 2009.
- Vaswani, A., et al. "Attention Is All You Need." *Advances in Neural Information Processing Systems* 30 (2017). The paper that introduced the transformer architecture named its core mechanism "attention" — a term borrowed from cognitive science but implemented as a purely computational operation of weight assignment across token sequences.