Symbiotic Traces: An Ecological Record of Language Model Training

Hanna Renedo Junxiu Tang Dashun Wang
Kellogg School of Management, Northwestern University, Evanston, IL, 60201
{hanna.renedo, junxiu.tang, dashun.wang}@kellogg.northwestern.edu

Abstract

What would it mean to visualize a language model's training process in the vocabulary of a crisis it will never be able to feel? *Symbiotic Traces* pairs fifty stages of the OLMo 3 32B pretraining run with climate terms drawn from IPCC reports and Earth system science, then with original steam-printed marigold petals. AI interprets deterministic rules to assign each climate term, while the physical chemistry of botanical printing shapes each final image. With creative agency distributed across machine, artist, and material process, the product is a visual language for a model under construction.

1 Introduction

The relationship between artificial intelligence and its environmental cost has become an increasingly controversial subject of public concern. Training a single large language model requires an extraordinary amount of energy: the pretraining run for OLMo 3 32B, an open language model released by the Allen Institute for AI, consumed roughly 681 megawatt-hours of power from its GPUs alone [1]. *Symbiotic Traces* asks what it would mean to describe that consumption in the same vocabulary we use to describe a warming planet, and to let the model's own training process help write that description.

The project pairs two entangled systems: the pretraining of a large language model, and the climate system destabilization to which its energy use contributes. Each of the training run's fifty stages is paired with a term from the vocabulary of climate change and an original botanical print, made by steaming a marigold petal directly onto paper. The result is a series of climate term and petal print pairings, in which each term assigned to narrate the model's own construction simultaneously becomes the trace of a living thing.

We chose marigolds because their growing seasons are among those already shifting under climate change. Using them makes climate change both this project's subject and part of its material process. This project considers what it means for an artist to respond to a distillation of climate vocabulary that a machine has assigned to the stages of its own training. Positioned at the intersection of botanical printmaking, climate change discourse, and machine training data, *Symbiotic Traces* offers a visual lexicon in which the story of a model's growth and the story of a warming world are told side by side.

2 Methodology

Our dataset is the training run log of OLMo 3 32B, a 32-billion-parameter open frontier language model trained by the Allen Institute for AI on 5.7 trillion tokens of text between September 22 and November 7, 2025. We were drawn to this log because it doubles as a record of both how the model became usable and, by proxy, the energy that process consumed. We compiled a climate lexicon with Claude Fable 5, drawn from IPCC reports and climate literature, then narrowed it to fifty words spanning growing-season and species-range terms alongside the broader climate vocabulary.

40th Conference on Neural Information Processing Systems (NeurIPS 2026). Creative AI Track.



Figure 1: Documentation of the plant steaming process. Marigold petals arranged between sheets of moistened Arches text-wove paper, the bundle wrapped around a dowel, and the finished prints after unwrapping.



Figure 2: The complete set of fifty marigold petal prints, each paired with a climate word assigned to its corresponding OLMo 3 32B training stage.

Each of the fifty training stages was assigned one word from this subset through a single deterministic rule: each stage was checked against a fixed list of measurable conditions, and each condition that held contributed a candidate word, and one word was selected from among them (Appendix A).

We then produced original botanical prints of locally sourced marigold petals by moistening paper, arranging the petals between two sheets of Arches text-wove paper, wrapping the bundle tightly around a dowel, sealing it in parchment and steaming it (Figure 1). Heat and moisture release pigment directly from the petal into the paper fiber, transferring color and texture in a unique visual gesture. Although we defined the assignment rule and carried out this process, neither we nor the model controlled the resulting image: each print's coloration, bleed, and stain are determined by heat and moisture alone. This method echoes what Dietmar Offenhuber, in *Autographic Design*, calls autographic visualization: an image that carries a direct physical trace of its source rather than a symbolic representation of it [2].

We then paired each climate term with its stage's petal print through a process of visual association: matching the tone and connotation of each word to the petal print whose form or gesture felt most complementary to it. The abstraction of the contact prints made this kind of associative reading possible. Each mark was open enough to hold a range of interpretive weight rather than dictating a single, literal reading. This pairing of text and image informs the viewer's reading, much as Roland Barthes describes text and image as anchoring and relaying one another's meaning [3]. Sequenced together, the resulting fifty pairings of climate word and petal print, reproduced in Figure 2, form a visual lexicon in which language drawn from a machine's training and a mark drawn from a living plant sit side by side. For the accompanying video, the fifty petal prints were algorithmically morphed into one another.

3 Exploring the Theme of Agency

Symbiotic Traces addresses this year's theme of Agency on several levels. First, the project treats the artist's partial delegation of control over the work's final visual outcome to the machine as an expansion, rather than a diminishment, of creative agency. In *Symbiotic Traces*, agency predominantly resides in designing the conditions under which words and images emerge: the deterministic rule that assigns climate words to stages of model training, and the material process through which marigold petals transfer their pigments to paper. The final pairing of each print therefore arises through an encounter between machine, material, and artist. Long before the development of AI, Vera Molnár argued that delegating aspects of artistic production to a computational system does not necessarily diminish expression, and could reveal what is truly human. Reflecting on her own practice, she observed that "...it may seem paradoxical that a so-called inhuman machine helps to realize what is most subjective, what is most profound, in man" [4].

Second, the project considers agency at the scale of the individual confronting a planetary crisis. Climate change is a problem so vast that any single, personal response can feel inconsequential. *Symbiotic Traces* responds to that scale by building a gestural and interpretive visual language out of the vocabulary already used to describe it. By pairing each abstract term with a physical trace of a living flower, we see this as a way of making the crisis legible on a tangible and material scale.

Third, and most directly, the project is distributed across three actors whose agencies remain deliberately unresolved: (1) the large language models: OLMo 3, whose own training behavior selects each word within our rule, and Claude Fable 5, the instrument that compiles the lexicon and writes the code; (2) the botanical printing process, whose steam and pigment transfer determine the unique form of each mark; (3) and the artist, who designs the rules that dictate these processes and composes the word and petal print pairings. None of the three are fully in control, and none are fully passive.

Novelist Olga Ravn, whose fiction often stages encounters between human and non-human intelligence, has drawn a sharp line between the two: intelligence and consciousness, she has said, are two very different things [5]. *Symbiotic Traces* contends with that distinction. It treats the model as an intelligent but non-conscious collaborator, one capable of sorting and selecting language, but not of feeling or meaning making. *Symbiotic Traces* positions the human artist as an important emotional, embodied, and discerning counterbalance to the capabilities of the model. Unlike machines, humans have a specific and nuanced relationship to the environment, which ultimately shapes the visual outcome of the piece.

4 Author Bios

Hanna Renedo is an interdisciplinary visual artist and designer at the Center for Science of Science and Innovation at Northwestern University. Her research-based practice explores the intersections of science, visual art, and technology, through artists' books, information design, and experimental image-making. She is a recipient of the Caxton Book Arts Grant, and her work is housed in several special collections, including the School of the Art Institute of Chicago and the Bainbridge Island Museum of Art. She earned a BA in biology from Bowdoin College and an MFA in Visual Communication Design from the School of the Art Institute of Chicago.

Junxiu Tang is an interdisciplinary researcher across computer science, digital media art, and social science. His research studies how people contribute to, read, and act on data. His art practice makes these processes material, spatial, and perceptible, while examining what is lost, invented, or redistributed through them. His works have been exhibited at the SIGGRAPH Art Gallery, IEEE VIS Arts Program, and Information+ Conference.

Dashun Wang is the Kellogg Chair of Technology and a Professor at the Kellogg School of Management, and the McCormick School of Engineering, at Northwestern University. Dashun is a recipient of multiple awards for his research and teaching, including the AFOSR Young Investigator award, Poets & Quants Best 40 Under 40 Professors, Junior Scientific Award from the Complex Systems Society, Erdos-Renyi Prize, Thinkers50 Radar 2021, and more.

References

[1] A. Bradbury. Minipost: OLMo 3 training cost. Muxup, 2025. URL https://muxup.com/2025q4/minipost-olmo3-training-cost. Reporting power-consumption figures disclosed by the Ai2 OLMo 3 team.

[2] Dietmar Offenhuber. *Autographic Design: The Matter of Data in a Self-Inscribing World*. MIT Press, Cambridge, MA, 2023. ISBN 9780262547024. doi: 10.7551/mitpress/14666.001.0001.

[3] Roland Barthes. Rhetoric of the image. In Stephen Heath, editor, *Image, Music, Text*. Hill and Wang, New York, 1977. Translated by Stephen Heath.

[4] Zsuzsanna Vályi-Nagy. Vera molnár, 1924–2023. *Artforum*, 2024.

[5] Olga Ravn. Podcast interview. Granta Podcast, 2026. URL https://granta.com/podcast-olga-ravn/. Interview with Thomas Meaney. Accessed July 2026.

A Mapping the Training Run to Climate Words

The OLMo 3 32B pretraining run is divided into 50 stages, and a deterministic rule gives each stage one word from a 50-word climate lexicon. Words may repeat. The stages come from KernelCPD change-point segmentation of the telemetry, 46 breakpoints over 600 bins, plus two mid-training stages and one long-context stage. We chose the granularity by a sweep from 5 to 60 stages. The number of distinct words peaks at 15 on the plateau between 46 and 51 stages, and we selected 50.

The rule has three parts. A list of measurable conditions turns each stage's telemetry into candidate words, cumulative consumption caps how severe a word may be, and a priority rule picks one candidate.

1. Conditions Table 1 lists the conditions, the word or word series each contributes, and the correspondence behind it. A stage may satisfy several conditions, and a stage that satisfies none receives the default word *global warming level*.

Table 1: Conditions measured on each training stage, the climate word or series each contributes, and the correspondence motivating it.

Condition (measured on the stage)	**Contributes**	**Correspondence**
Learning rate ramping up	*anthropogenic*	the forcing being switched on
Log-loss falls by > 0.3	tipping series	rapid state reorganization
Stage belongs to mid-training	shift series	the data mix migrates
Stage belongs to long-context training	*adaptation*	restructuring for a new environment
All three of gradient stress, mass refusals, and infrastructure failure co-occur, consumption $\geq 80\%$	*ecosystem collapse*	every system failing at once, late
Two of the three co-occur, consumption $\geq 20\%$	*cascading effects*	one system's failure propagating
Extreme gradient turbulence (CV ≥ 1.5 or max $\geq 40\times$ median), one rank worse if ≥ 5 refusals	stress series	acute shock
Gradient spikes ≥ 5 (milder)	one rank lower in the stress series	recurring stress
Optimizer refused ≥ 30 steps	tipping series	the system at its safety limit
Throughput collapsed $\geq 30\times$ or ≥ 4 restarts	disruption series	infrastructure breakdown
Cumulative consumption crosses 50% or 80%	*carbon budget*	the budget being spent
Gradient baseline $\geq 1.3\times$ its run minimum, consumption $\geq 50\%$	*sea-level rise*	a slow, monotone, late rise of the floor
Log-loss rises by > 0.005 (regression)	*flood*	sudden excess, state overwhelmed
Loss descends > 0.01 with rebound $< 0.5\%$	*irreversibility*	descent that never comes back
Loss CV ≥ 0.02 with no trend	*deep uncertainty*	variance without direction
LR $< 15\%$ of peak and loss flat	*net zero*	forcing gone, state settled
LR $< 35\%$ of peak and falling	*decarbonization*	the forcing being withdrawn
None of the above	*global warming level*	monotone drift, nothing else true

2. Consumption gate Four contributions in Table 1 are series of rising severity, and the cumulative token share of the 5.7T-token run sets the highest rank a stage may reach (Table 2). Severity therefore follows accumulated consumption. The most extreme moment of the run, the opening loss collapse at 0.2% consumption, is held to *early warning signal*.

Series	Ranks, with the cumulative consumption that unlocks each
stress	*heatwave* 0%, *drought* 20%, *megadrought* 50%, *aridification* 80%, *desertification* 80%
disruption	*hazard* 0%, *displacement* 20%, *climate migration* 50%
tipping	*early warning signal* 0%, *threshold* 20%, *tipping point* 50%, *tipping cascade* 80%
shift	*phenological shift* 20%, *range shift* 50%

Table 2: Severity series and the cumulative consumption that unlocks each rank. A condition yields the highest rank the stage's consumption allows.

3. Priority When several conditions hold, the condition that holds in the fewest stages across the run decides, and ties go to table order. The default word stays wherever nothing else holds.

4. Result Table 3 gives the assignment stage by stage. The mapping uses 15 of the 26 words tied to conditions. Eleven stay unused for three reasons. The refusal condition appears only twice at this granularity and loses the tie to the data-mix condition, so *threshold*, *tipping point*, and *tipping cascade* drop out. This run meets none of the conditions for *flood*, *deep uncertainty*, and *irreversibility*. The triggers for *aridification*, *desertification*, *ecosystem collapse*, *displacement*, and *phenological shift* occur outside the consumption tier each word requires. The other 24 lexicon words have no condition tied to the telemetry.

Table 3: Stage-by-stage mapping from deciding condition to assigned climate word (c = cumulative token share at stage end). Consecutive identical default stages are compressed.

Stage(s)	c	Deciding condition	Word
1	0.002	opening collapse (gated to rank 1)	*early warning signal*
2	0.004	warmup	*anthropogenic*
3–14	0.006–0.137	none	*global warming level* $\times 12$
15	0.138	extreme turbulence ($55.8\times$)	*heatwave*
16	0.164	infrastructure (4 restarts)	*hazard*
17	0.223	gradient spikes	*heatwave*
18–21	0.287–0.475	none	*global warming level* $\times 4$
22	0.476	extreme turbulence	*drought*
23–24	0.496–0.499	none	*global warming level* $\times 2$
25	0.508	crossing 50% of the budget	*carbon budget*
26	0.524	infrastructure	*climate migration*
27	0.524	extreme turbulence + refusals	*megadrought*
28–30	0.573–0.652	none	*global warming level* $\times 3$
31–34	0.672–0.748	LR decay	*decarbonization* $\times 4$
35–36	0.765–0.787	gradient baseline risen $\geq 1.3\times$	*sea-level rise* $\times 2$
37	0.807	crossing 80% of the budget	*carbon budget*
38–39	0.817–0.820	baseline rise	*sea-level rise* $\times 2$
40	0.839	two failure families co-occur	*cascading effects*
41	0.841	infrastructure (536 throughput collapses)	*climate migration*
42–43	0.841–0.854	baseline rise	*sea-level rise* $\times 2$
44	0.882	forcing ≈ 0, loss flat	*net zero*
45	0.914	baseline rise	*sea-level rise*
46–47	0.944–0.966	forcing ≈ 0	*net zero* $\times 2$
48–49	0.975–0.983	data-mix change	*range shift* $\times 2$
50	1.000	long-context restructuring	*adaptation*

Read in order, the sequence tracks the run. *Early warning signal* opens it, *anthropogenic* marks the warmup, and *global warming level* holds through the long early stretch. *Heatwave* and *drought* arrive

as consumption passes 14% and 48%. At the half-budget mark, *carbon budget*, *climate migration*, and *megadrought* follow in close succession. Learning-rate decay brings *decarbonization*, and a rising gradient baseline yields *sea-level rise* seven times across the late run. *Cascading effects* and a second *climate migration* fall near 84%, *net zero* arrives as the learning rate settles, and *range shift* and *adaptation* close the run as the data mix changes and long-context training begins.

5. Provenance The mapping is a pure function of the public OLMo 3 32B pretraining telemetry (W&B project `ai2-llm/Olmo-3-1125-32B`, pulled 2026-07-09), the 50-word lexicon, and the rule above, with no randomness.