

Shapes of Cambria: The Cambrian Explosion in Silico

Jinwoo Lee

Korea Advanced Institute of Science and Technology, South Korea

<https://cinescope-wkr.github.io/shapes-of-cambria/>

ABSTRACT

*In the burgeoning field of “differentiable visual computing,” the once-static components of the imaging pipeline—lenses, sensors, and light transport—are reimaged as differentiable parameters, allowing for end-to-end optimization through gradient descent. However, this paradigm often sacrifices artistic agency for the sake of sterile mathematical convergence. **Shapes of Cambria: The Cambrian Explosion in Silico** is an interactive generative ecosystem that reclaims this agency, transforming the biological “Light Switch” theory into a computational manifold.*

*By parameterizing a differentiable polynomial optics framework as a mutable “optical genome,” the work simulates a digital Cambrian Explosion where vision is not a fixed utility, but an evolving medium. While modern AI models chase the cold certainty of image sharpness, our system positions the artist and participants as the primary environmental pressure. This steers the evolutionary agents toward the intentional cultivation of “expressive artifacts”—chromatic bleeding, optical warping, and organic aberrations. Through this human-in-the-loop intervention, **Shapes of Cambria** demonstrates that the synergy between differentiable mechanics and curated mutation can yield a more poetic rendering of reality than deterministic optimization alone, effectively engineering the imperfection into the heart of the digital eye.*

1 INTRODUCTION

In biological history, the sudden emergence of vision—conceptually framed as the “Light Switch” hypothesis [9, 10]—triggered the Cambrian Explosion [3, 14], catalyzing an unprecedented leap in ecological complexity. Profoundly inspired by recent computational models that simulate this milestone through the co-evolution of embod-

ied agents—most notably the pioneering framework by Tiwary et al. [18]—our work embraces the brilliance of such numerical simulations. However, these algorithmic narratives remain inherently teleological. Bound by reductionist objective functions that solely maximize survival tasks and image sharpness, they reduce the profound drama of emergent sight to a sterile benchmark, entirely neglecting the organic, physical friction of light transport.

This sterility is mirrored in the contemporary trajectory of visual computing. Historically, forward rendering granted artists absolute control over physical light simulation. Conversely, the advent of differentiable ray tracing [8, 15] and end-to-end computational optics [5, 20, 23] has shifted power from the creator to the algorithm. By transforming lenses and sensors into fluid parameters optimized for mathematical infallibility, these frameworks [24] treat physical anomalies—such as chromatic aberration and diffraction—as noise to be eradicated. Consequently, intentional imperfection and artistic agency are exiled behind the opaque black box of deep learning algorithms, enforcing an illusion of flawless perfection.

In this work, we challenge this *deterministic convergence* by situating itself within the rich lineage of *artificial life art*. While pioneering works including Karl Sims’ *Evolved Virtual Creatures* [13] demonstrated the morphological evolution of virtual entities, and Ian Cheng’s *Emisaries* [2] simulated ecosystems driven by unpredictable cognitive agents, our work diverges by evolving the perceptual medium itself. We propose a dual-path rendering architecture **Cambria**, which is built upon a customized differentiable polynomial optics framework [6], where the internal mechanics of lenses and light transport function as highly mutable digital phenotypes.

Rather than utilizing *gradient descent* and *auto differentiation* techniques [4, 11] to erase optical, graphical de-

viation, our system appropriates this exact mathematical foundation to intentionally cultivate anomalies as expressive artifacts. Drawing upon *Interactive Evolutionary Computation* [17], we reintroduce the artist and participants as active environmental pressures, subverting the closed optimization loop into a curated canvas. Philosophically aligned with Hito Steyerl’s defense of the “poor image” [16], the work rejects the hegemony of high-resolution perfection. When our **digital organisms** open their eyes, they inherit the organic flaws of our surrounding physical reality, ultimately reclaiming optics and computational rendering as a fluid medium for **human intentionality** into the heart of artificial sight, seemingly dominated by the algorithms.

2 CAMBRIA

2.1 From Numerical Simulation to Interactive Ecosystem

Our installation is primarily inspired by the Artificial Cambrian Intelligence (ACI) framework [18], which couples evolutionary search with reinforcement learning to study how task-specific pressures shape visual morphology. While we retain the premise of embodied optimization, we reconfigure the computational simulation for the scientific purpose into **Cambria**—an interactive, generative ecosystem built on a genetic graphics framework.

In standard ACI, evaluation and termination are tethered to isolated tasks, which obscures the evolutionary loop behind episodic resets. Cambria dismantles this constraint through a factorized architecture: a *control layer* maintaining embodied locomotion, an *optical layer* parameterizing perception, and a *curation layer* embedding human mutation into the environment.

As illustrated in Figure 1, the system decouples fast variables (optical genome, degeneration state) that mutate at the cadence of human interaction, from slow variables (locomotion policy) updated at lineage boundaries.

2.2 The Mutable Optical Genome

To transcend the idealized pinhole camera typical of embodied AI, **Cambria** parameterizes perception as an explicit, mutable lens system. The optical genome θ^{opt} consists of a sequence of refractive interfaces. For each surface, the physical sag profile is defined by the standard aspheric form:

$$z(r) = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=2}^N \alpha_i r^{2i}, \quad (1)$$

where r is the radial distance, c is the base curvature, k is the conic constant, and α_i are higher-order coefficients.

By grounding the genome in differentiable optical design [6, 7], visual variations—such as depth-dependent blur, chromatic dispersion, and spherical distortion—emerge not as arbitrary post-processing filters, but as the direct consequences of physical geometries (Figure 2).

2.3 Dual-Path Light Transport

Evaluating complex wave optics is computationally prohibitive for live evolutionary search. Cambria resolves this friction through a bifurcated rendering pipeline (Figure 3).

- **In-Loop Proxy Rendering:** For continuous, real-time interaction, Cambria utilizes a latency-bounded proxy inside the MuJoCo [25] native viewer. It approximates field-dependent blur and radial distortion to preserve the relative perceptual structure across candidate phenotypes without breaking embodied control.
- **Asynchronous Archival Rendering:** Once a phenotype is selected, Cambria serializes the full world state, materials, and optical genome into a Universal Scene Description (.usd) payload. This is exported to an offline sample-based ray tracer [6] to compute high-fidelity aberrations and intricate light-material interactions that are otherwise invisible in the live loop.

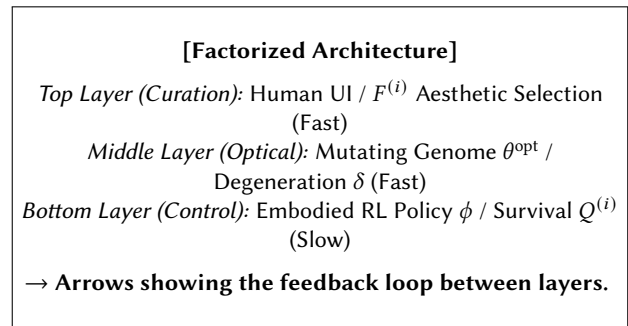


Figure 1. Factorized architecture of Cambria. The system decouples fast variables (optical parameters, degeneration) that respond immediately to human curation, from slow variables (locomotion policy) updated at lineage boundaries, enabling a continuous interactive ecosystem.

[Differentiable Lens Profile]

Visual: A cross-section of a complex lens system.
Overlay mathematical parameters onto the drawing:
 c (curvature), t (thickness), a (aperture), and rays bending (refraction).
Show how changing θ^{opt} morphs the lens shape and shifts the focal point.

Figure 2. The mutable optical genome. Unlike standard geometric approximations, Cambria models explicit refractive interfaces. Differentiable parameters seamlessly warp the lens profile, directly inducing physical aberrations like depth-dependent blur and spherical distortion.

2.4 Viability-Gated Interactive Selection

If human preference governed lineage updates unconditionally, the digital agents would quickly collapse into visually striking but behaviorally inert phenotypes. **Cambria** prevents this by retaining a task-grounded viability score $Q^{(i)}$ (accumulated survival reward over a short horizon) while offloading the evolutionary gradient to the participant.

At each generation, a gradient-free optimizer samples a set of mutated phenotypes. Through the native viewer interface, participants assign an aesthetic preference score $U^{(i)}$ to their desired mutation. The effective fitness $F^{(i)}$ governing the lineage update is strictly gated by viability:

$$F^{(i)} = \begin{cases} U^{(i)}, & Q^{(i)} \geq q_{\min} \\ -\infty, & Q^{(i)} < q_{\min} \end{cases} \quad (2)$$

where q_{\min} is the minimum behavioral threshold. This logic expresses the core artistic intervention of the system

[Dual-Path Rendering Pipeline]




Figure 3. The bifurcated rendering pipeline. The real-time path ensures frictionless intra-generational evaluation, while the asynchronous archival path computes high-fidelity wave optics phenomena from the exported Universal Scene Description (.usd) payload.

[Mockup: The Curator Interface]

Left: Main Simulation Viewport (Agent swimming in cambrian_shallows).
Right/Bottom: A grid of 3-4 mutated vision tiles.
Overlay UI: Each tile has a viability check (e.g., $Q > q_{\min}$ **PASS** or **FAIL**).
Only 'PASS' tiles have clickable [Breed] buttons.

Figure 4. The viability-gated selection interface embedded natively within the simulator. Participants observe and select aesthetically evocative optical mutations, but the optimization strictly limits the selectable manifold to phenotypes that satisfy minimum embodied survival requirements ($Q \geq q_{\min}$).

(Figure 4): human curation dictates the aesthetic evolution, but strictly within the manifold of the search space that supports embodied survival.

2.5 Observable Degeneration and Materiality

Standard RL benchmarks represent failure as an instantaneous episode termination (`reset()`), rendering systemic collapse invisible. Cambria introduces an explicit, continuous degeneration state $\delta \in [0, 1]$ that accumulates physical trauma (e.g., collisions, prolonged occlusion) prior to termination. As δ increases, the optical state is heavily modulated—widening the point-spread function and amplifying decentration—while motor control attenuates. This transforms abrupt computational truncation into a cinematic process of physical and perceptual decay.

Furthermore, Cambria decouples task logic from environmental appearance. While high-frequency visual cues remain behaviorally operative for the agent, they are seamlessly embedded into cinematic water-column structures and organic textures. This separation ensures the simulation transcends its origins as a functional test rig, supporting sustained artistic intervention without sacrificing scientific legibility.

3 DISCUSSION

Contemporary computer vision and graphics research is inherently bound by rigid objective functions that seek to render or recognize the “perfect imagery,” systematically suppressing noise and optical aberrations [24]. By prioritizing task-specific fidelity and survival metrics, these teleological models often reduce the profound evolutionary drama of vision to a sterile computational benchmark

within the language of *evolutionary pressure* [18]. **Shapes of Cambria: The Cambrian Explosion in Silico** diverges from this trajectory of deterministic optimization, recontextualizing the optical simulation as an untamed, emergent ecosystem. Within this work, the differentiable optics pipeline is not utilized as a tool for carving flawless lenses for target [24], but rather functions as the computational substrate for *physical morphogenesis*—a process by which **digital organisms physically struggle to accommodate light**.

By introducing human-in-the-loop intentionality as an active environmental pressure, the system dismantles the blind algorithmic pursuit of mathematical perfection. It reclaims physical anomalies—such as cascading chromatic bleeding and diffraction—not as rendering noise to be minimized, but as expressive vestiges of evolution [1]. This approach subverts the common rendering techniques established by the graphics community, deploying them instead to explore the latent physical textures and glitches that lie beneath frictionless digital imagery [16]. Ultimately, this *in silico* ecosystem proposes a novel interface for aesthetically negotiating how machine vision perceives the world, shifting the paradigm from algorithmic problem-solving to ecological curation.

4 FUTURE WORK

The conceptual framework of optical evolution proposed by our work scales across multiple intersections of computer science and artistic practice. As this *in silico* ecosystem matures, future iterations will unfold along two primary axes: the phylogenetic visualization of algorithms and the physical democratization of interaction.

4.1 Algorithmic Phylogeny and Structural Evolution

Current evolutionary models within the system are constrained to exploring continuous parameter spaces, such as lens curvature and refractive indices. Future work will extend this into the *realm of genetic programming* [7], enabling **structural evolution** wherein the underlying shader logic and neural architectures of the agents undergo autonomous crossover and mutation. Concurrently, we intend to reframe the repository’s version control history (e.g., Git [19] commits) as a literal “*digital fossil record*” [12]. By visualizing this algorithmic phylogeny, the extinct fragments of deprecated code will be integrated into the artwork’s meta-narrative alongside surviving optical traits,

documenting the dead-ends and detours of evolutionary computation.

4.2 Surrogate Distillation for Ubiquitous Interaction

To democratize this evolutionary process, future development must address the computational friction of differentiable ray tracing [6, 8] and numerical simulations [18], which currently tethers human intervention to high-performance hardware environments. By distilling heavy physical simulations into lightweight neural surrogate models [21] integrated within real-time game engines like Unity[22], we aim to open the ecosystem to **ubiquitous interaction**.

Crucially, this transition to a real-time environment is not merely an exercise in computational optimization. The data loss and rendering artifacts inevitably introduced when approximating rigorous physical models for real-time execution are embraced within the project’s philosophy as *translational mutations*. The computational friction experienced as computationally-intensive physics descend into the arena of real-time interaction will act as a novel organic texture, born from the digital organism’s adaptation to a new computational environment. Through this, we aim to transition differentiable rendering from an isolated computational task into a **continuously evolving, multi-participant artistic ecology**.

REFERENCES

- [1] Michael Betancourt. *Glitch art in theory and practice: Critical failures and post-digital aesthetics*. Routledge, 2016. 4
- [2] Ian Cheng. Emissaries: A trilogy of simulations. *Architectural Design*, 89(1):118–125, 2019. 1
- [3] Douglas Fox. What sparked the cambrian explosion? *Nature*, 530(7590):268–270, 2016. 1
- [4] Roy Frostig, Matthew James Johnson, and Chris Leary. Compiling machine learning programs via high-level tracing. In *SysML conference 2018*, 2019. 1
- [5] Chi-Jui Ho, Yash Belhe, Steve Rotenberg, Ravi Ramamoorthi, Tzu-Mao Li, and Nicholas Antipa. A differentiable wave optics model for end-to-end computational imaging system optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 28042–28051, 2025. 1
- [6] Matthias B Hullin, Johannes Hanika, and Wolfgang Heidrich. Polynomial optics: A construction kit for efficient ray-tracing of lens systems. In *Computer Graphics Forum*, pages 1375–1383. Wiley Online Library, 2012. 1, 2, 4
- [7] William B Langdon and Riccardo Poli. *Foundations of genetic programming*. Springer, 2002. 2, 4

- [8] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. *ACM Transactions on Graphics (TOG)*, 37(6):1–11, 2018. 1, 4
- [9] Dan-E Nilsson and Susanne Pelger. A pessimistic estimate of the time required for an eye to evolve. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 256(1345):53–58, 1994. 1
- [10] Andrew Parker. *In the Blink of an Eye: The Cause of the Most Dramatic Event in the History of Life*. Free Press, New York, 2003. 1
- [11] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 1
- [12] Václav Rajlich. Software evolution and maintenance. In *Future of software engineering proceedings*, pages 133–144. 2014. 4
- [13] Karl Sims. Evolving virtual creatures. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 699–706. 2023. 1
- [14] M Paul Smith and David AT Harper. Causes of the cambrian explosion. *Science*, 341(6152):1355–1356, 2013. 1
- [15] Andrew Spielberg, Fangcheng Zhong, Konstantinos Rematas, Krishna Murthy Jatavallabhula, Cengiz Oztireli, Tzu-Mao Li, and Derek Nowrouzezahrai. Differentiable visual computing for inverse problems and machine learning. *Nature Machine Intelligence*, 5(11):1189–1199, 2023. 1
- [16] Hito Steyerl et al. In defense of the poor image. *e-flux journal*, 10(11):1–9, 2009. 2, 4
- [17] Hideyuki Takagi. Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE*, 89(9):1275–1296, 2002. 2
- [18] Kushagra Tiwary, Aaron Young, Zaid Tasneem, Tzofi Klinghoffer, Akshat Dave, Tomaso Poggio, Dan-Eric Nilsson, Brian Cheung, and Ramesh Raskar. What if eye...? computationally recreating vision evolution. *Science Advances*, 11(51):eady2888, 2025. 1, 2, 4
- [19] Linus Torvalds and Junio C Hamano. Git, 2025. 4
- [20] Ethan Tseng, Shane Colburn, James Whitehead, Luocheng Huang, Seung-Hwan Baek, Arka Majumdar, and Felix Heide. Neural nano-optics for high-quality thin lens imaging. *Nature communications*, 12(1):6493, 2021. 1
- [21] Ethan Fan Hei Tseng. *Neural Proxy Functions for Computational Cameras and Displays*. PhD thesis, Princeton University, 2025. 4
- [22] Unity Technologies. Unity 6.3 lts release: Performance, stability, and production readiness. <https://blog.unity.com/engine/unity-6-3-lts-is-now-available>, 2025. Accessed: 2026-03-09. 4
- [23] Congli Wang, Ni Chen, and Wolfgang Heidrich. do: A differentiable engine for deep lens design of computational imaging systems. *IEEE Transactions on Computational Imaging*, 8:905–916, 2022. 1
- [24] Zhaoqiang Wang, Yifan Peng, Lu Fang, and Liang Gao. Computational optical imaging: on the convergence of physical and digital layers. *Optica*, 12(1):113–130, 2025. 1, 3, 4
- [25] Kevin Zakka, Baruch Tabanpour, Qiayuan Liao, Mustafa Haiderbhai, Samuel Holt, Jing Yuan Luo, Arthur Allshire, Erik Frey, Koushil Sreenath, Lueder A Kahrs, et al. Mujoco playground. *arXiv preprint arXiv:2502.08844*, 2025. 2