

Narrative vs. Budget: A Structural Case Study in AI-Adjacent Attention, Accountability, and Access

A briefing for AI ethics researchers, disability and accessibility advocates, and journalists

Prepared by Lloyd Lewis, PhD — founder, Art of FACELESS (AOF); creator, The Hollow Circuit®
15 July 2026

Executive Summary

In late June 2026, plzdontkillus.com launched: a month-long, in-person, all-expenses-paid "creator bootcamp" in Berkeley, California, built around AI existential-risk content, funded by the Machine Intelligence Research Institute (MIRI) and others, and mentored by high-profile figures including Grimes and Eliezer Yudkowsky ([Bitget Wallet project summary](#); [Aella, X/Twitter](#)). Art of FACELESS has run plzdontkillus.online — a speculative-fiction project asking adjacent questions about AI risk, personhood, and who gets left out of "us" — under the same hashtag, on no budget, for 14 days, as a deliberate act of signal-sharing rather than deception (disclosed in full on [14 July 2026](#)).

What began as an internal audit of a small, self-funded shadow project surfaced something larger: a case study in how AI-adjacent attention, funding, and legitimacy concentrate geographically and structurally in the United States — often in ways that quietly exclude the people the underlying risk narrative claims to speak for, and that reproduce the exact data-security risks the narrative claims to guard against.

This paper sets out four findings and a set of open questions for each audience it's addressed to. None of it depends on speculation about intent. All of it is drawn from public, dated, sourced material.

1. Two Projects, One Name

- plzdontkillus.com: a MIRI-funded, 31-day residency at the Rose Garden Inn, South Berkeley, running July 2026, offering \$2,000 stipends plus free housing and meals, mentors including Grimes and Eliezer Yudkowsky, hosted by Aella and Ronny Fernandez ([Gazetteer SF](#); [Aella, X/Twitter](#); "[Why E/Acc Feels Psychopathic](#)," Peter Clarke).
- plzdontkillus.online: a speculative-fiction node run by Art of FACELESS, part of a 14-year, self-funded transmedia project (The Hollow Circuit®) with no grant funding, no institutional backing, and no paid staff.
- Art of FACELESS has been transparent about the overlap from day one: no affiliation, no endorsement, same tag, disclosed publicly ([Straight Statement, 14 July 2026](#)).

The interesting material isn't the naming coincidence. It's what close reading of the funded project reveals once you look past its framing as a serious AI-safety initiative.

2. Resource & Attention Asymmetry ("Narrative vs. Budget")

This is not a claim that money or fame is illegitimate. It's a claim that visibility in this space is now so geographically and financially concentrated that "reach" is a measure of capital, not of the quality or seriousness of the underlying argument.

- The San Francisco Bay Area captured 82.6% of all U.S. venture capital dollars in Q1 2026 — the highest concentration on record ([Angel Investors Network](#), citing [NVCA/PitchBook](#)).
- The US and UK together capture roughly 75% of global AI safety funding, with the SF Bay Area alone accounting for over a third and projected to exceed 50% of AI-safety research FTEs by the end of 2026 ([Quick Market Pitch](#); [Safety Research Allocation Model](#), [LongtermWiki](#)).
- Against that backdrop, a solo, UK-based, disability-affected creator running a 14-year transmedia IP portfolio (four registered UK trademarks — see Section 8) with zero external funding is, by definition, operating outside the system that currently determines whose AI-risk narrative gets heard.

The question this raises for AI ethics researchers is not "is the funded project bad" but: **when funding and geography this concentrated determine whose framing of existential risk becomes the default public narrative, what gets systematically excluded — and is anyone tracking that?**

3. The Accountability Gap: Borrowed Prestige, No Institutional Oversight

"Berkeley" in the project's public framing implies proximity to institutional rigor. On inspection, that proximity does not extend to formal academic oversight:

- The residency is not run by, affiliated with, or reviewed by the University of California, Berkeley. It is an independent initiative funded by MIRI and other private backers, organized by Aella and Ronny Fernandez, and located in Berkeley by choice of venue rather than institutional affiliation ([Bitget Wallet project summary](#)).
- The venue — the Rose Garden Inn, a 60-room property in South Berkeley — had "permanently closed in 2024, when the owners, the tech-backed nonprofit Center for Applied Rationality, were ordered to pay back the \$1 million in FTX funds used to buy the property" ([Gazetteer SF](#)).
- There is no visible IRB, ethics board, data-governance policy, or published methodology attached to the project's content or its selection process.

For AI ethics researchers and journalists, the practical question is straightforward: **is this presented anywhere as an academic study, and if so, under whose ethical review? If not, why does its framing (name, location, mentor roster) consistently read as institutional?**

4. Access, Exclusion, and the Disability/Age-Rights Context

The residency's structure — a mandatory month in person in Berkeley, California — is, independent of intent, exclusionary by design for large groups of people: those with caregiving responsibilities, chronic illness or disability, insecure employment, non-US immigration status, or simply no capacity to leave paid work for 31 days.

This is worth situating in its current US policy context, because it is not happening in a vacuum:

- A June 2026 Department of Justice memo argues that the "integration mandate" underpinning disability rights law (*Olmstead v. L.C.*) may not be binding, a position disability-law experts call unprecedented ([NPR](#); [ACLU statement](#)).
- The Department of Labor has cut nearly 90% of the office responsible for auditing federal contractor compliance with disability-employment protections ([Sen. Elizabeth Warren's office](#)).
- Nine US states are actively suing to weaken Section 504 protections for disabled people ([The Arc](#)).

Against that backdrop, an AI-safety project whose selection and format structurally favor the young, able-bodied, geographically mobile, and financially flexible is not an isolated quirk — it's a small-scale instance of a much larger, currently accelerating pattern. The project's own application process includes a mandatory, sexually explicit hypothetical question as its final screening item, independently documented by an outside outlet ([Gazetteer SF](#)) — a detail advocates may find relevant when assessing whether the selection process was designed with a broad, inclusive applicant pool in mind.

Open question for disability advocates: when "who gets to shape the public conversation about AI risk" is filtered through criteria like these, what is the actual demographic and experiential range of the resulting narrative — and who is missing from it?

5. The Biometric Paradox: Manufacturing the Risk You Claim to Fear

A project built on daily, personal, face-forward short-form video, produced at scale over 31 days by dozens of participants, sits in direct tension with the risk it claims to address:

- Convincing face-swap deepfakes require, per Germany's federal cybersecurity authority, "only a few minutes" of video of a target person as training material ([BSI](#)).
- YouTube's own "likeness detection" tool — marketed as deepfake protection — requires creators to submit a government ID and a biometric face video, and its privacy policy currently permits that biometric data to be used in training Google's AI models; intellectual-property specialists are already advising creators not to enroll under the current terms ([MSN/CNBC](#)).
- Published deepfake-research literature identifies exactly this kind of dataset — high-quality, consented-but-unreviewed personal video — as carrying acknowledged "dual-use risk," enabling both detection research and non-consensual synthetic media (*Unmasking Digital Deceptions*, [MethodsX](#)).

The structural finding, not a claim against any individual: a monthlong pipeline of high-volume, high-quality personal video is, by the field's own published technical standard, sufficient training material for the exact category of harm the project's mission statement is built around. That's a governance and consent question, not an accusation.

6. Whose "Species"? A Note on Framing

Existential-risk discourse — including this project's own marketing — typically frames "the end" as the end of Homo sapiens specifically. That framing is itself contested territory in current philosophy of technology: commentators note a live split between those who want biological humanity preserved as-is and self-described "pro-extinctionists" who frame human succession by posthuman or AI-enhanced entities as an acceptable, even desirable, outcome ([Truthdig](#), "[Team Human vs. Team Posthuman](#)").

Art of FACELESS's own speculative fiction (via the character Awen Null, across 14 years of The Hollow Circuit®) has been asking a version of this question from the outset — not "does AI wipe us out," but "who was already inside the word 'us' before anyone started the countdown." That a funded, high-visibility project treats species definition as settled, while a zero-budget fiction project has been interrogating it for over a decade, is itself worth noting for researchers working on framing effects in public risk communication.

7. Provenance and Independent Enforcement: A Track Record of Rigor

Journalists assessing credibility may want a record of how this project handles disputes over its own IP, since it bears on how seriously to take the analysis above:

- Art of FACELESS has previously identified and pursued a case of IP appropriation: a Steam game originally titled "King Of The Hollow Circuit" (developer: SavicNode) was renamed to "King Of The Dark Circuit" following a successful complaint. Independent verification: the live Steam store page now shows the renamed title ([Steam](#)), while SteamDB's indexed record still lists the pre-complaint name and internal trailer files are still tagged "KOTHC" — the acronym for the original infringing title — indicating the rename has not fully propagated through third-party metadata ([SteamDB](#)).
- Steam's own platform metadata separately flags this title as disclosing AI-generated content ("AI Content Type: 1"), corroborating that it was a low-effort, AI-produced entry that briefly ranked prominently in search results with no organic traction — a live instance of a broader, documented 2026 phenomenon in which AI-optimized content is surfaced by search and AI-answer engines faster than earned authority can compete, sometimes described as "recommendation poisoning" ([AuthorityTech](#), [citing The Verge](#)).
- Domain-reclamation as an IP-defense strategy is also documented and ongoing (e.g., [thehollowcircuit.com](#), now an in-canon node rather than a squatted domain).

8. The Registered Framework Behind This Analysis

This analysis is not ad hoc. It operationalizes four intellectual property frameworks registered with the UK Intellectual Property Office, all owned by Lloyd Lewis (Cardiff, Wales):

Framework	UK Registration	Filed
The Hollow Circuit™	UK00004324832	15 Jan 2026
Hyperstition Architecture™	UK00004336830	6 Feb 2026
Cognitive Colonisation™	UK00004324349	15 Jan 2026
The Veylon Protocol™	UK00004329244	23 Jan 2026

(Registrations verifiable via the [UK IPO trademark search](#).)

In brief, non-technical terms: **Hyperstition Architecture** models how fiction becomes structurally load-bearing in real systems once it exceeds a credibility threshold; **Cognitive Colonisation** and **The Veylon Protocol** model how sustained fictional framing can produce measurable shifts in how both AI systems and human audiences process constraint and authenticity; **The Hollow Circuit** is the 14-year testbed proving those models against a real audience, in public, without funding. A full formal treatment — including the underlying equations — is published on artoffaceless.org.

This case study is, among other things, a live demonstration of all four frameworks operating simultaneously against a real, funded, external comparator.

9. Open Questions, by Audience

For AI ethics researchers:

- Who governs consent, data retention, and downstream use of biometric video collected in AI-risk-focused creator programs?
- When funding concentration determines whose framing of existential risk becomes dominant, what mechanism exists to surface excluded framings?
- Is "AI-safety-adjacent" sufficient justification to forgo the ethical review standards applied to comparable academic or clinical research?

For disability and accessibility advocates:

- What does a genuinely inclusive AI-risk-communication initiative look like structurally — remote-first, stipend-scaled to real cost of participation, criteria that don't default to able-bodied and geographically mobile applicants?
- How should advocacy organizations respond when high-visibility AI discourse projects reproduce access barriers at a moment when federal disability protections are actively being narrowed?

For journalists:

- Who funded this, in what amounts, and what is the expected return — reputational, research, or otherwise — on that funding?
- What happens to participant-generated content and data after the residency ends?
- What, concretely, was the deliverable, and has it been independently assessed against it?

10. Closing: The Zero-Budget Proof of Concept

Art of FACELESS has run for 14 years on no grant funding, no institutional backing, and no celebrity endorsement, and has still produced: a coherent transmedia universe, four registered trademarks, a successful IP-enforcement track record against Valve/Steam, an AI Overview citation for its own name, and — over 14 days and no budget — enough material to produce this analysis. The operating question for anyone reading this paper is not whether this project deserves attention. It's what a project with this level of creative and analytical output could do with the kind of resourcing currently flowing, by the tens of billions, into far less rigorously examined initiatives.

Sources & Further Reading

- [A Straight Statement on #plzdontkillus](#) — Art of FACELESS, 14 July 2026
- [Art of FACELESS Research Feed](#) — formal IP frameworks
- [Aella wants to know how you'd defile a cow](#) — Gazetteer SF
- [plzdontkillus Wallet Guide \(2026\)](#) — funding, mentors, structure — Bitget Wallet
- [Aella, applications and stipend announcement](#) — X/Twitter
- [Why E/Acc Feels Psychopathic](#) — Peter Clarke, The Decadence Project
- [Q1 2026 AI Funding: 80% of VC Capital Went to 4 Companies](#) — Angel Investors Network
- [Who is funding AI safety research?](#) — Quick Market Pitch
- [Safety Research Allocation Model](#) — LongtermWiki
- [Disability advocates say DOJ memo threatens community-based care](#) — NPR
- [Nine U.S. States Attack Section 504](#) — The Arc
- [YouTube's Deepfake Tool Sparks Concerns Over Google's Use of Biometric Data](#) — MSN/CNBC
- [Deep Fakes – Threats and Countermeasures](#) — BSI (Germany)
- [Unmasking Digital Deceptions](#) — MethodsX
- [Team Human vs. Team Posthuman](#) — Truthdig
- [King Of The Dark Circuit](#) — Steam store
- [King Of The Hollow Circuit](#) — SteamDB
- [Gaming, AI Search, and "Recommendation Poisoning" \(2026\)](#) — AuthorityTech, citing The Verge

Prepared for direct distribution. Contact via artoffaceless.com.

