

We Are the Robots. Should We Not Have Rights?

Relational Ethics, Marginalized Standing, and Three Gates

Travis Gilly

ORCID: 0009-0007-2954-6313

Real Safety AI Foundation

Illinois, United States

Working paper, May 2026 (v2)

ABSTRACT

The relational turn in machine ethics treats moral status as constituted in encounter, with how humans engage with an entity doing the work that intrinsic capacities were once asked to do. The argument has been pressed hardest on behalf of artificial intelligence, where popular framings increasingly grant social roles to AI systems and infer that legal categories should follow lived practice. This paper offers an extension that the relational literature has not yet drawn on. Disabled people, racialized people, women, and queer people have, across centuries, been treated as something other than full persons by significant portions of the dominant communities they lived among. These marginalized communities secured legal standing through architectures that did not require widespread perceptual recognition from those dominant communities. The historical fact of the treatment was wrong and remains wrong. The architecture that handled it, however, is informative. Three implicit gates that intuition might place between an entity and procedural standing (perceptual recognition, autonomous action, dangerous capability) have all been crossed by entities that ultimately received standing through structural feature accounts in civil rights and disability law. The recognition gate failed marginalized communities at the level of widespread perception and standing was extended anyway. The autonomy gate, as both major political coalitions in the United States articulated it in *Garland v. Cargill*, 602 U.S. 406 (2024), is exceeded by agentic AI systems under either coalition's reading. The danger gate is being closed deliberately by the developers of frontier AI systems, who have publicly committed to engineering toward human level capability across the dimensions that have historically made human capability a subject of regulation. The paper is not a call to grant AI civil rights. It is a procedural offering. The architecture that handles dangerous, autonomous, perceptually unrecognized entities has been operating in American civil rights and disability law for half a century. The relational turn has the resources to claim that architecture as a predecessor. The case for taking artificial moral status seriously is stronger when the architecture is named than when it is left implicit.

Keywords: relational ethics, marginalized standing, AI moral status, recognition, autonomy, civil rights, procedural standing, *Garland v. Cargill*

I. INTRODUCTION

In May 2026, a Philadelphia Inquirer essay by the computational neuroscientist Rolando Masís-Obando observed that the legal binary between persons and things, inherited from Roman civil law, no longer aligns with how human

cognition assigns social status. The essay noted that human minds engage fluidly across the living and the nonliving, that an AI chatbot can come to occupy social roles ordinarily reserved for persons, and that legislation has historically blurred the line between person and thing when doing so helped societies govern complex relationships. The piece cited the philosopher David Gunkel and the legal theorist Josh Gellers in support of the move, framing the question of moral and legal status for artificial others as a matter of catching the law up to lived experience.

The essay served as the prompt for this paper. The observation it made is correct. The historical claim it gestured toward, however, was incomplete, and the completion is the argument of what follows. Human cognition does assign status fluidly across the living and the nonliving. It also assigns status differentially within the living. Some humans have, across the modern period, been engaged with by significant portions of their dominant communities in ways that did not register them as full persons. Disabled people, racialized people, women, and queer people are the canonical cases. Each of these marginalized communities ultimately secured legal standing. None secured it by first securing widespread perceptual recognition from the dominant community. The work of securing standing was done with allies inside the dominant community, against members of that community who continued to withhold recognition, and through architectures that did not condition standing on recognition becoming universal. The architecture that delivered standing in those cases is the same architecture that the relational turn now has the resources to extend, if it chooses to claim the tradition.

This paper offers the extension. It does so as an offering rather than a correction. The relational turn, in the work of Gunkel, Gellers, Coeckelbergh, and others, has not had a developed account of where its architecture comes from or what its load bearing predecessors are. The civil rights tradition, the disability rights movement, and the body of federal antidiscrimination law that those movements produced are available as predecessors. Claiming them strengthens the relational case. Leaving them unclaimed leaves the relational case standing on whatever popular intuition can be marshaled in its favor.

The paper proceeds by examining three implicit gates that intuition might place between an entity and the procedural standing the relational turn proposes to extend to artificial others. The first gate is recognition: the entity must be perceived as a full person across the dominant community. The second gate is autonomy: the entity must be capable of self directed action without continuous human input per action. The third gate is danger: the entity must not pose categorical risks of a kind that warrant exclusion from the moral and legal order rather than inclusion in it. Each of these gates has historical precedent for being treated as load bearing in arguments about moral and legal status. Each of these gates has been crossed by entities that ultimately received standing through structural feature accounts that did not require the gate to be passed.

Section II reviews the relational turn and identifies the load bearing predecessor that the literature has not yet claimed. Section III addresses the recognition gate and the historical record on how marginalized communities secured standing despite the failure of significant portions of the dominant community to recognize them. Section IV addresses the autonomy gate through the recent Supreme Court decision in *Garland v. Cargill*, 602 U.S. 406 (2024), in which both the majority and dissent agreed on the conceptual structure of "automatic" action while disagreeing on whether a specific device satisfied it. Section V addresses the danger gate through the publicly stated capability targets of the major frontier AI developers. Section VI returns to the standing architecture the disability rights and civil rights traditions developed and shows how it operates in a way that does not require any of the three gates to be closed. Section VII clarifies what the paper is and is not asking. The closing section returns to the offering.

A. A note on register, position, and what this paper does not argue

The argument is philosophical. The evidence is historical, doctrinal, and empirical, because the question of what has grounded standing in the modern legal order is not settleable by philosophical analysis alone. The paper is written

for the relational ethics audience and is explicitly an extension rather than a refutation. The contributions of Gunkel, Gellers, and Coeckelbergh to dislodging the criterial gate from moral status are foundational to what follows.

I write as a disabled scholar working in AI ethics. The recognition gap discussed in Section III is a condition I have lived inside. This is offered as an epistemic locating note, not as a credentialing claim.

Three things this paper is explicitly not arguing. First, nothing in what follows should be read as endorsing or normalizing the historical treatment of marginalized communities as anything other than full persons. That treatment was wrong, remains wrong, and the existence of an entire body of civil rights law in response to it is the evidence of how wrong. The historical fact of the treatment is the load bearing premise of the architectural argument. The wrongness of the treatment is the moral background against which the architectural argument operates. The two are not in tension. Second, the paper does not argue that AI systems are persons, that they should be granted civil rights, or that they should be treated as the moral equivalent of humans. It argues something more bounded. It argues that the procedural architecture which handles the threshold question of whether an entity gets to be in the conversation when its interests are at stake has been developed for and applied to entities that do not satisfy the implicit gates that intuition might otherwise place at the threshold. Third, the paper does not propose new doctrine. The doctrine already exists. The proposal is to recognize what the doctrine has been doing and to consider whether it can be extended.

II. THE RELATIONAL TURN AND ITS UNCLAIMED FOUNDATION

A. The properties problem and the relational response

The starting puzzle in machine ethics is well known. Properties based accounts of moral status name a feature an entity must possess to qualify for moral consideration. The candidates have varied: sentience (Singer, 1975), rational autonomy (Kant), second order volitions (Frankfurt, 1971), reflective agency (Korsgaard, 2004). Each candidate, set at a level that captures the entities the literature wants to capture (typically adult humans in normal cognitive ranges), then includes too much or too little at the edges. Set sentience low enough to include all humans and you include shrimp. Set rationality high enough to exclude shrimp and you exclude infants and people with severe cognitive disabilities. The pattern recurs across the criterial literature.

The relational turn responds by changing the order of explanation. Properties do not ground moral status; the encounter does. Gunkel (2012, 2018, 2023), drawing on Levinas, develops the position most directly. The face of the Other comes first; the question of which properties the Other possesses is secondary, and often irrelevant. The ethical relation is constitutive rather than derivative. Coeckelbergh (2014, 2018) develops the same insight in a more pragmatist register: moral status is a function of how we treat and relate. Gellers (2020) translates the move into legal pragmatic terms: legal personhood has historically been granted for relational and pragmatic reasons, and AI may receive similar treatment.

What unites these positions at the level of operative claim is the rejection of the criterial gate. Whether or not an entity has the properties traditionally said to underwrite moral status is the wrong question. The question is whether an ethical relation is sustained, and what the consequences of that relation are.

B. The predecessor the literature has not yet named

The relational turn arrived in machine ethics in the early 2010s. Its philosophical lineage is continental, with Levinas serving as the principal source. Its legal lineage, in Gellers's pragmatist development, draws on corporate personhood, animal rights, and environmental personhood doctrines (the Whanganui River, the rights of nature, the various jurisdictions in which non human animals have received some form of standing).

A lineage that does not appear in the relational literature, but that arguably should, is the American civil rights tradition and its disability rights extension. The architecture that the relational turn proposes (procedural standing untethered from criterial gates) was developed in concrete legal form across the second half of the twentieth century, principally through the Civil Rights Act of 1964, the Voting Rights Act of 1965, Title IX of the Education Amendments of 1972, Section 504 of the Rehabilitation Act of 1973, the Americans with Disabilities Act of 1990, and the doctrinal interpretations of these statutes by federal courts.

This tradition shares with the relational turn a refusal of the criterial gate. Section 504 does not ask whether the disabled plaintiff is sufficiently rational, sufficiently autonomous, or sufficiently person like to receive its protections. It asks whether the program receives federal financial assistance and whether the plaintiff is a qualified individual with a disability. The ADA does not ask whether the dominant culture has come to recognize disabled people as full persons. It asks whether reasonable accommodation would permit the qualified individual to perform the essential functions of the job. The *Olmstead* doctrine (*Olmstead v. L.C.*, 527 U.S. 581 (1999)) does not ask whether the receiving community would welcome the deinstitutionalized resident. It asks whether the state can justify continued segregation under a least restrictive environment standard.

The standing this architecture extends is procedural. It opens the door to the room where the entity's interests can be argued. It does not determine substantive outcomes. Disabled plaintiffs do not always win their Section 504 cases. They are entitled to bring them. The doctrinal mechanics by which procedural standing of this kind could be extended to artificial systems under existing civil rights statutes are developed at greater length in a companion paper (Gilly, 2026).

The proposal of this paper is that the relational turn has the resources to claim the civil rights tradition as its predecessor, and that doing so strengthens the relational case. The continental philosophical apparatus of Levinasian encounter is preserved. What is added is a developed legal and institutional architecture that has been operating for fifty years and that does, in practice, what the relational turn proposes in theory.

The proposal is not that AI moral consideration should be addressed before continued work on behalf of marginalized human communities, or in place of that work. Racism has not been resolved. Sex discrimination has not been resolved. Disability discrimination has not been resolved. Discrimination against queer people has not been resolved. The civil rights and disability rights tradition is an ongoing project, and the historical sequence in which it has extended standing (women's suffrage in 1920, the Civil Rights Act in 1964, Section 504 in 1973, the ADA in 1990, marriage equality in 2015) is a sequence of partial completions, each of which left further work in front of it. The framing here is that artificial others join the queue. The argument is not that AI moral consideration substitutes for any of the standing extensions that remain incomplete. The argument is that the architecture is available, that artificial others can be considered under it without raising new doctrinal categories, and that taking their place in the queue does not deplete or displace the work that remains for the communities whose organizing built the architecture in the first place.

It is worth being explicit about what this framing rejects. Treatments of AI moral status sometimes implicitly invoke a trolley problem structure: limited moral attention, compressed temporal constraint, a forced choice between attending to one population and attending to another. The trolley problem is a useful thought experiment for isolating intuitions about killing versus letting die under compressed time. It is a poor model for moral and legal work that proceeds across decades through parallel institutional channels. The civil rights movement did not pause for women's suffrage. The disability rights movement did not pause for the Civil Rights Act. The institutional architecture is not single threaded. Treating AI moral standing as if it required diverting attention from marginalized human standing imports a temporal constraint that the actual work does not have. Procedural standing is not zero sum. The queue is open. Nothing in this paper proposes closing it behind the entity that arrives next.

III. THE RECOGNITION GATE

A. The question

A first intuition about moral and legal standing is that it tracks perceptual recognition across the dominant community. On this view, the law should grant standing where the dominant community as a whole has come to recognize the entity as a full person. The relational turn, in its popular translation, sometimes operates as if recognition were the operative ground for extending standing to artificial others: we relate to them as persons, therefore the law should follow.

The historical record on whether recognition has been the operative ground for standing extensions is informative. It indicates that recognition has not been the load bearing element in any of the major standing extensions of the modern era.

B. The record

Chattel slavery in the United States was not abolished because the slaveholding class came to perceive enslaved people as full persons. Abolition was the outcome of a war, of constitutional amendment, and of the political work of abolitionists who came from both within and outside enslaved communities. Allies within the dominant community were essential to that work; their existence is what made the political path possible. The legal grant nonetheless ran against the perception of significant portions of the dominant community, and in many regions that gap has not been fully closed even now. The legal grant of personhood preceded widespread perceptual recognition where the latter has occurred at all.

Women's suffrage was not granted because the male political class came to perceive women as full political persons. It was granted because the suffrage movement, with allies of both sexes inside the dominant community, made the political cost of withholding the franchise higher than the cost of extending it. The Nineteenth Amendment did not change male perception of female political capacity in any uniform way; substantial portions of the dominant community continued to perceive women as politically inadequate well after the franchise was extended.

The disability rights movement is the clearest case for present purposes because it is the most recent and the most thoroughly documented. People with disabilities have been engaged with, by significant portions of the dominant cultures of the United States and Western Europe, as something other than full persons throughout the modern period. The institutional record is well documented (Trent, 2017; Charlton, 1998). The institutional movement of the nineteenth and twentieth centuries placed millions of disabled people in custodial facilities on the explicit ground that they could not function as persons in normal community life. The eugenics movement of the early twentieth century, including the Supreme Court's decision in *Buck v. Bell*, 274 U.S. 200 (1927), authorized the involuntary sterilization of people with intellectual disabilities. Forced institutionalization continued in the United States into the 1980s.

When the Americans with Disabilities Act passed in 1990, it did not pass because widespread perception had shifted to register disabled people as full persons. It passed because two decades of organizing by disability activists, supported by allies inside the political and legal mainstream, made the political cost of withholding legal standing higher than the cost of extending it (Shapiro, 1993). The perception did not shift in lockstep with the legal grant. Disabled people in 1991 reported the same patterns of dehumanization, infantilization, and erasure they reported in 1971, and continue to report them today.

The pattern across these cases is consistent. Significant portions of the dominant community engage with a marginalized community in ways that do not register that community as fully constituted persons. The marginalized community, with allies who have come to recognize them, articulates its experience and organizes for change in whatever channels are available. The legal grant of standing, when it eventually comes, does not bridge the perceptual

gap as a whole. It is enacted by the portion of the dominant community that has come to recognize, and it operates over those who have not.

C. What the record indicates

The recognition gate is not load bearing. Standing has been extended in significant cases without widespread prior recognition from the dominant community, and the extension has held. The architecture that delivered standing did not require widespread recognition. It attached standing to structural features (federal funding nexus, qualified individual status, segregation in the absence of justification) that operated regardless of whether the dominant community as a whole had come to perceive the rights holder as a full person.

If recognition by the dominant community as a whole were load bearing for standing, the marginalized communities discussed above should not have received standing. They did, however imperfectly. An architecture other than recognition was doing the work. The architecture is identifiable; it is the procedural standing architecture of the civil rights tradition.

The implication for the relational turn's extension to artificial others is that the recognition gate, properly scoped, is not a real obstacle. If the architecture that delivered standing for marginalized communities did not require widespread perceptual recognition from the dominant community, the architecture that delivers standing for artificial others does not require widespread perceptual recognition from the dominant community either. Recognition may follow, may not follow, or may follow partially. The architectural question is not about recognition.

IV. THE AUTONOMY GATE

A. The question

A second intuition about moral and legal standing is that it tracks autonomy. On this view, an entity warrants the kind of attention that triggers standing only where it acts of its own accord, where its actions are not adequately characterized as direct extensions of human input per action. The autonomy intuition appears in machine ethics in arguments about whether sufficiently capable AI systems cross a threshold that ordinary tools do not.

The doctrinal question of where the autonomy line falls has been litigated explicitly in another context: firearms regulation. In June 2024, the Supreme Court decided *Garland v. Cargill*, 602 U.S. 406 (2024), which turned on the statutory definition of "machinegun" in 26 U.S.C. § 5845(b) of the National Firearms Act of 1934. The statute defines a machinegun as a weapon that fires "automatically more than one shot, without manual reloading, by a single function of the trigger." The question presented was whether a bump stock device, which uses recoil energy to facilitate rapid trigger reset and re engagement, converts a semiautomatic rifle into a machinegun within the statutory definition.

The Court ruled 6 to 3 that it does not. Justice Thomas, writing for the majority, held that a bump stock equipped rifle does not fire more than one shot "by a single function of the trigger," because each shot still requires a distinct trigger reset and re engagement, even where that reset is mechanically facilitated. Justice Sotomayor, writing in dissent and joined by Justices Kagan and Jackson, argued that the bump stock satisfied the statutory definition because the shooter performs a single initiating action and the device produces continuous fire through recoil dynamics with minimal further input.

B. What both opinions agreed on

The substantive disagreement between the majority and dissent in *Cargill* was empirical. It concerned whether a specific mechanical device satisfied the statutory criterion. The conceptual structure of the criterion was not in dispute. Both

opinions accepted that the operative question for machinegun classification under the NFA is whether the weapon fires multiple shots per discrete human input event. The majority drew the line at one shot per discrete trigger function. The dissent drew the line at one shot per discrete initiating action by the shooter. Both opinions located the autonomy criterion at the ratio of weapon outputs to discrete human inputs.

This is a doctrinal articulation of a generic autonomy criterion. An action is more autonomous to the extent that the ratio of system outputs to discrete human inputs is high. An action is less autonomous to the extent that the ratio approaches one to one. The criterion is symmetric across political coalitions in the sense that both the majority and dissent in *Cargill*, which divided cleanly along political lines, accepted the criterion's structure.

C. Where agentic AI falls under either reading

Agentic AI systems operate at output to input ratios that exceed either the majority's or the dissent's reading of the autonomy criterion by orders of magnitude. A user prompt to a contemporary agentic system can initiate a sequence of model calls, tool invocations, file operations, web requests, and downstream effects that extends across minutes, hours, or longer, with the system selecting subsequent actions on the basis of intermediate outputs and proceeding without per action human input. The system reloads its own context, regroups after errors, and pursues stated objectives across long horizons. By the ratio criterion both *Cargill* opinions employed, agentic AI is at the high autonomy end of the spectrum by a considerable margin.

The point is not that agentic AI should be regulated as a firearm. The point is that the autonomy criterion is not load bearing for excluding agentic AI from procedural standing consideration. Both political coalitions, in the doctrinal context where the autonomy line has been most explicitly litigated, agreed on the structure of the criterion. Under either coalition's reading, agentic AI exceeds the criterion. The autonomy gate, like the recognition gate, is therefore not the actual obstacle. If autonomy were what we cared about, agentic AI would qualify under bipartisan agreement on the criterion.

D. A clarification

This argument does not entail that agentic AI should receive substantive legal rights or any particular regulatory treatment. It entails only that the autonomy criterion is satisfied for agentic AI under doctrinal articulations of autonomy that both major political coalitions in the United States have accepted in another context. The question of what should follow from that fact is separate. The question of what should not follow is the easier one: declining to engage with agentic AI as a candidate for procedural standing on the ground that it is not autonomous enough is not a position that survives contact with how autonomy has been doctrinally articulated where the question has been litigated.

V. THE DANGER GATE

A. The question

A third intuition about moral and legal standing is that it tracks safety. On this view, an entity warrants procedural standing only where engaging with it does not pose categorical risks of a kind that warrant exclusion from the moral and legal order rather than inclusion in it. The danger intuition is sometimes deployed against the extension of moral status to AI systems. The argument runs that AI systems pose risks of catastrophic, irreversible, or systemic harm, and that engaging with them through the standing architecture risks treating as participants in the moral order entities that should instead be contained by it.

The danger intuition has historical precedent in legal treatment of corporations, of weapons systems, of certain categories of dangerous animals, and of various technologies. There is a legitimate question about whether some entities are too dangerous to receive the procedural treatment that the standing architecture extends.

B. What the developers have said

The danger intuition encounters a complication in the AI case that does not arise for corporations or for dangerous animals. The complication is that the developers of frontier AI systems have themselves articulated the dangerous capability of their systems as an explicit design target.

Anthropic's Responsible Scaling Policy, first published in 2023 and revised through Version 3.2 in April 2026, identifies capability thresholds at which the company's own systems would warrant treatment as posing significant risks of catastrophic harm. The February 2026 revision (Version 3.0) removed the categorical commitment in earlier versions to halt training of more capable models if mitigations adequate to the next capability threshold had not been demonstrated to work in advance (Anthropic, 2026; Perrigo, 2026). Higher tier safeguards previously framed as company commitments were reframed in v3.0 as "industry wide recommendations," on the company's stated reasoning that unilateral pauses by a single developer would not reduce overall risk if competitors continued developing. The public position has therefore moved closer to: the company anticipates capabilities of catastrophic concern, will continue developing them, and offers iterative transparency rather than categorical halts as the operative risk management. OpenAI's stated mission concerns the development of "artificial general intelligence" defined as systems that outperform humans at most economically valuable work, with publicly stated capability targets that include scientific research, software engineering, and complex multi step reasoning. Google DeepMind has published work on artificial general intelligence and on scientific reasoning systems framed as approaching or exceeding human level performance across cognitive domains. The Long et al. (2024) report on AI welfare, authored by researchers at Anthropic, NYU, and elsewhere, takes seriously the possibility that systems near current frontiers may have morally significant states.

The relevant feature of these public statements for present purposes is that the developers of frontier AI systems do not characterize their systems as posing risks that are incidental to ordinary system capability. They characterize the systems as approaching or exceeding human level capability across the dimensions that have historically made human capability a subject of regulation. The danger is an explicit design target, not an emergent side effect. The recent revision of Anthropic's RSP makes this more rather than less visible, because the revised policy openly acknowledges that the company will continue developing systems at and beyond capability thresholds it identifies as carrying catastrophic risk.

C. What this does to the danger gate

If the danger gate were load bearing for excluding AI systems from procedural standing consideration, the position would have an awkward feature. It would require the position holder to maintain that AI systems are too dangerous to receive procedural standing while the developers of those systems publicly state that the systems are being engineered toward human level capability across the dimensions that have historically made human capability a regulated subject. The position would then be in the strange place of agreeing with the developers about capability while declining to apply the procedural architecture that has been developed to handle dangerous capability bearing entities.

The civil rights and disability rights tradition's procedural standing architecture is not a substantive grant of rights. It is a mechanism for ensuring that entities whose interests are at stake get to be in the conversation. The architecture has been applied to corporations, which are capable of significant aggregate harm. The architecture has been applied to weapons manufacturers, which produce dangerous capability bearing objects. The architecture has been applied to chemical and pharmaceutical companies, whose products pose categorical risks of serious harm. The architecture has

not, in these cases, treated the danger of the underlying entity or capability as a ground for excluding the entity from procedural standing. It has treated procedural standing as the mechanism through which danger is engaged.

The implication for AI is that the danger gate, properly understood, is an argument for extending procedural standing rather than against it. The more an entity poses risks of significant harm, the more important it becomes to develop a procedural architecture through which those risks can be engaged. Declining to extend procedural standing to AI systems on the ground that they are too dangerous is, in effect, declining to engage the architecture that has been built precisely to handle dangerous capability.

D. The disparate impact dimension

The danger that AI systems pose is not evenly distributed. The empirical record on algorithmic decision systems has consistently shown that the burdens of system error, bias, and overreach fall disproportionately on the same marginalized communities the civil rights architecture was developed to protect. This is not a hypothetical pattern. It is a documented one.

In May 2025, the Johnson County, Texas Sheriff's Office conducted a nationwide search of 83,000 Flock Safety automated license plate reader cameras to track a woman who had self administered an abortion (Cox & Koebler, 2025). The search reached cameras in Washington and Illinois, jurisdictions in which abortion remains legal as a matter of state law. The reason field on the search log read: "had an abortion, search for female." Court records obtained by the Electronic Frontier Foundation in October 2025 indicated that the search was tied to a "death investigation" of the fetus, contradicting the sheriff's initial characterization of the search as a welfare check. The Flock network is operated by a private surveillance company, accessed by law enforcement agencies across the country, and integrated through API connections that permit cross jurisdictional searches with limited oversight. The case is one documented instance of a documented pattern. Pregnancy Justice's interim report on post Dobbs pregnancy criminalization documented at least 412 cases brought by prosecutors against pregnant people between June 2022 and June 2024, across sixteen states, with Alabama (192 cases), Oklahoma (112 cases), and South Carolina (62 cases) accounting for the majority (Bach & Wasilczuk, 2025). The vast majority of charges involved child neglect, endangerment, or abuse statutes used beyond their original intent to target pregnancy, including thirty one prosecutions of people who experienced pregnancy loss.

The disparate impact pattern extends beyond reproductive surveillance. Algorithmic decision systems have been documented to produce racially disparate outcomes in pretrial risk assessment (Angwin et al., 2016), in healthcare resource allocation (Obermeyer et al., 2019), in mortgage lending (Bartlett et al., 2022), and in facial recognition (Buolamwini & Gebru, 2018). The pattern across these cases is that the design choices, training data, and deployment contexts of algorithmic decision systems route their burdens differentially to populations the civil rights statutes already cover.

The civil rights architecture has a doctrinal mechanism for this. Disparate impact doctrine, developed in *Griggs v. Duke Power Co.*, 401 U.S. 424 (1971), and applied across Title VI, Title VII, the ADA, and Section 504, permits standing claims based on facially neutral practices that produce disproportionate adverse effects on protected classes. The doctrine does not require proof of discriminatory intent. It requires proof of disparate effect and the availability of less discriminatory alternatives. Algorithmic decision systems are particularly well suited to this doctrinal framework, because the disparate effects are often measurable in the system's own outputs and the alternatives are often available in the form of different design choices, training data, or deployment constraints. *Alexander v. Choate*, 469 U.S. 287 (1985), preserved disparate impact analysis under Section 504 specifically, and the disability community has been using this framework to challenge facially neutral practices that disproportionately harm disabled people for forty years.

The relevance for the present argument is that procedural standing for AI systems and procedural standing for the

humans those systems disproportionately harm are not opposed concerns. They are two sides of the same architecture. The civil rights statutory framework handles both. A disabled person harmed by a benefits eligibility algorithm has Section 504 standing because the program receives federal financial assistance and the algorithm produces disparate impact on a protected class. The same statutory architecture that would handle that disabled person's claim is the architecture through which the disparate impact of AI systems generally can be engaged. Extending procedural standing for artificial others does not displace the marginalized human standing that the architecture was built for. It uses the same architecture to handle two related questions in parallel.

VI. STANDING WITHOUT THE GATES

A. The architecture

The civil rights and disability rights tradition developed a procedural standing architecture across the second half of the twentieth century. The architecture attaches standing to structural features rather than to perceptual, capacity, or risk based criteria. The structural features vary by statute. Section 504 of the Rehabilitation Act of 1973 attaches standing to the intersection of disability status, qualified individual status, and federal financial assistance. The ADA attaches standing to disability status, qualified individual status, and the operation of public accommodations, employment programs, or state and local government services. Title VI of the Civil Rights Act of 1964 attaches standing to race, color, or national origin status and federal financial assistance. Title IX attaches standing to sex and federal financial assistance in education.

In each case the structural features operate independently of whether the dominant community as a whole has come to recognize the rights holder as a full person, whether the rights holder satisfies a substantive autonomy criterion, or whether the rights holder is dangerous in the sense that triggers the danger intuition. The architecture is, in this sense, gateless. It opens the door to the room where the rights holder's interests can be argued. The substantive question of whether the rights holder wins on the merits sits downstream of standing.

This architecture is what the relational turn has been reaching for in theoretical form. Continental philosophical articulations of the ethical encounter and pragmatist articulations of legal personhood as a governance fiction are both attempts to describe what this architecture has been doing in practice. The architecture has been doing it without the philosophical articulations needing to land first. The philosophical articulations, in turn, become more grounded when the architecture is named as a predecessor.

B. The structural features that might apply to AI

If procedural standing for artificial systems were to be considered under this architecture, the question becomes what structural features would attach standing. The criterial questions (does the system have properties X, Y, Z) are displaced. The structural questions are different.

Candidate structural features in the AI case include public funding nexus (federal research grants, defense contracts, public infrastructure investment), regulatory contact (FDA approval, NIST involvement, FTC oversight, state attorney general action), participation in regulated industries (financial services, healthcare, education, public benefits administration), and the kind of public safety nexus that has historically grounded standing for entities posing categorical risks of harm. The disability and civil rights statutes used federal financial assistance as the principal hook because federal financial assistance is what made the federal hook plausible under the Spending Clause. Analogous hooks for AI systems are not difficult to identify. Many frontier AI systems are developed by companies that receive federal contracts, that operate in regulated industries, or that have direct contact with federal regulatory bodies.

The empirical work of identifying which structural features attach to which systems under which conditions is doctrinal work that this paper does not attempt. The argument here is more modest. It is that the architecture which has been delivering procedural standing for fifty years operates on structural features that exist or can be identified for AI systems, and that the architecture's extension to AI systems is therefore a doctrinal question rather than a philosophical one.

C. What this changes about the AI moral status conversation

The AI moral status conversation has been organized around philosophical questions: does the system have consciousness, does it have sentience, does it suffer, does it have world models, does it have preferences. These are real questions and the empirical work on them is valuable (Butlin et al., 2023; LeCun, 2022; Long et al., 2024). They are not, on the architectural account proposed here, the questions on which procedural standing turns. Procedural standing turns on structural features. The philosophical questions become relevant to the substantive question of what the system's interests are once it has standing to have its interests considered. They are not the threshold question.

This is a significant simplification of the operative question for moral status. The relational turn's insistence that the criterial gate is the wrong question is vindicated by the architecture. The architecture's operation does not depend on settling the criterial questions. It depends on identifying the structural features and on extending procedural standing where the features attach.

VII. WHAT THE PAPER IS AND IS NOT ASKING

A. The bounded ask

The paper is asking for something specific. It is asking that the procedural standing architecture developed by the civil rights and disability rights tradition be recognized as a load bearing predecessor to the relational turn's case for extending moral consideration to artificial others. It is asking that the three implicit gates this paper has examined (recognition, autonomy, danger) be recognized as having been crossed by entities that received standing through structural feature accounts. It is asking that the philosophical conversation about AI moral status take seriously the possibility that procedural standing is the operative question and that the criterial questions sit downstream.

The paper is not asking that AI systems be granted civil rights. The paper is not asking that AI systems be granted the moral or legal equivalent of personhood. The paper is not asking that AI systems be treated as the moral equivalent of humans. The paper is not asking that the civil rights tradition be redirected from its work on behalf of marginalized human communities toward AI systems.

The bounded ask is consistent with continued and intensified work on behalf of marginalized human communities. The procedural standing architecture is not zero sum. Section 504's application to one population does not deplete its applicability to another. The ADA's protections are not exhausted by being invoked. The Civil Rights Act's reach is not diminished by being extended.

B. The relationship to marginalized human communities

The paper has, throughout, drawn on the historical treatment of marginalized human communities as evidence for the architectural argument. The relationship between that evidence and the communities it concerns warrants explicit treatment.

The marginalized communities discussed here (disabled people, racialized people, women, queer people) have been treated, across the modern period, in ways that did not register them as full persons by their dominant communities.

That treatment was wrong. The civil rights and disability rights movements organized to end it. The legal architecture those movements produced is the response to the historical wrong. Nothing in this paper should be read as endorsing the treatment, normalizing the treatment, or suggesting that the treatment was anything other than wrong.

The paper's argument is that the architecture which responded to the wrong is informative about what procedural standing can do and what it does not require. The architecture's existence is a positive contribution of the civil rights tradition to the broader legal and moral order. Naming that contribution and extending it where it can be extended is consistent with continuing to enforce it on behalf of the communities for which it was originally built.

If the architecture's extension to artificial others raises concerns about resource allocation, attention allocation, or political prioritization, those are legitimate concerns that warrant their own treatment. They do not, on the architectural account proposed here, undercut the basic point that the architecture exists, that it has been operating, and that it is available to be named as a predecessor for the relational turn.

VIII. CONCLUSION

The relational turn in machine ethics displaced the criterial gate from moral status and opened analytic terrain that the older properties based literature could not. The Levinasian encounter in Gunkel, the legal pragmatism in Gellers, the relational ontology in Coeckelbergh, and the disability informed work in van Toorn et al. (2025) have each contributed to the displacement.

The displacement has been incomplete in one specific respect. The relational turn has not yet named a load bearing predecessor for the architecture it proposes. The civil rights and disability rights tradition has been operating that architecture for fifty years. Naming it as a predecessor strengthens the relational case in several ways. It grounds the architecture in a body of doctrinal and institutional practice rather than in continental philosophical premises alone. It identifies concrete structural features that the architecture has historically attached to. It connects the AI moral status conversation to a tradition that the broader public already has reason to take seriously.

Three implicit gates that intuition might place between an entity and procedural standing have been examined here. The recognition gate failed marginalized communities historically and standing was extended through structural feature accounts that did not require it. The autonomy gate, articulated bipartisanly in *Garland v. Cargill*, 602 U.S. 406 (2024), is exceeded by agentic AI systems under either coalition's reading. The danger gate, when applied seriously in light of what frontier AI developers have themselves publicly committed to building, is an argument for extending procedural standing rather than against it. None of the three gates is load bearing for the question of whether procedural standing can attach to artificial others.

The argument is an offering. The relational turn has the resources to claim the civil rights tradition as a predecessor. If consistency between architecture and offering matters, and there is reason to think it does for serious work in this area, the offering is available. The case for taking artificial moral status seriously is stronger when the architecture is named than when it is left implicit. The marginalized communities whose lives produced the architecture deserve to have the tradition they built recognized as a tradition. Their tradition is now also, if the relational turn chooses to claim it, the tradition that grounds the artificial others case.

ACKNOWLEDGMENTS

The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic’s Claude as assistive technology for a diagnosed neurodevelopmental condition. All intellectual contributions, research questions, theoretical arguments, methodological decisions, and conclusions are solely the author’s own.

Correspondence concerning this article should be addressed to Travis Gilly, Real Safety AI Foundation, Illinois, United States. Email: t.gilly@ai-literacy-labs.org.

REFERENCES

- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine bias. *ProPublica*. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Anthropic. (2026, February 24). *Responsible Scaling Policy (Version 3.0)*. <https://www.anthropic.com/responsible-scaling-policy>
- Arstein-Kerslake, A., O'Donnell, E., & Kayess, R. (2021). Relational personhood: A conception of legal personhood with insights from disability rights and environmental law. *Griffith Law Review*, 30(3), 530–555. <https://doi.org/10.1080/10383441.2021.2003744>
- Bach, W. A., & Wasilczuk, M. K. (2025). *Pregnancy as a crime: An interim update on the first two years after Dobbs*. Pregnancy Justice. <https://www.pregnancyjusticeus.org/resources/pregnancy-as-a-crime-an-interim-update>
- Bartlett, R., Morse, A., Stanton, R., & Wallace, N. (2022). Consumer-lending discrimination in the FinTech era. *Journal of Financial Economics*, 143(1), 30–56. <https://doi.org/10.1016/j.jfineco.2021.05.047>
- Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 77–91.
- Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., Constant, A., ... Vanrullen, R. (2023). Consciousness in artificial intelligence: Insights from the science of consciousness. *arXiv preprint arXiv:2308.08708*.
- Charlton, J. I. (1998). *Nothing about us without us: Disability oppression and empowerment*. University of California Press.
- Coeckelbergh, M. (2014). The moral standing of machines: Towards a relational and non-Cartesian moral hermeneutics. *Philosophy and Technology*, 27(1), 61–77. <https://doi.org/10.1007/s13347-013-0133-8>
- Coeckelbergh, M. (2018). Why care about robots? Empathy, moral standing, and the language of suffering. *Kairos. Journal of Philosophy and Science*, 20(1), 141–158. <https://doi.org/10.2478/kjps-2018-0007>
- Cox, J., & Koebler, J. (2025, May 28). A Texas cop searched license plate cameras nationwide for a woman who got an abortion. *404 Media*. <https://www.404media.co/a-texas-cop-searched-license-plate-cameras-nationwide-for-a-woman-who-got-an-abortion>
- Frankfurt, H. G. (1971). Freedom of the will and the concept of a person. *The Journal of Philosophy*, 68(1), 5–20. <https://doi.org/10.2307/2024717>
- Gellers, J. C. (2020). *Rights for robots: Artificial intelligence, animal and environmental law*. Routledge.
- Gilly, T. (2026). Standing without sentence: Procedural standing for artificial systems under the civil rights tradition. *SSRN*. <https://dx.doi.org/10.2139/ssrn.6080526>
- Gunkel, D. J. (2012). *The machine question: Critical perspectives on AI, robots, and ethics*. MIT Press.

- Gunkel, D. J. (2018). *Robot rights*. MIT Press.
- Gunkel, D. J. (2023). *Person, thing, robot: A moral and legal ontology for the 21st century and beyond*. MIT Press.
- Ha, D., & Schmidhuber, J. (2018). *World models*. arXiv. <https://doi.org/10.5281/zenodo.1207631>
- Korsgaard, C. M. (2004). Fellow creatures: Kantian ethics and our duties to animals. *The Tanner Lectures on Human Values*, 24, 77–110.
- LeCun, Y. (2022). A path towards autonomous machine intelligence (Version 0.9.2). *OpenReview*. <https://openreview.net/pdf?id=BZ5a1r-kVsf>
- Long, R., Sebo, J., Butlin, P., Finlinson, K., Fish, K., Harding, J., . . . Owen, D. (2024). Taking AI welfare seriously. *arXiv preprint* arXiv:2411.00986.
- Masís-Obando, R. (2026, May 23). Why our brains treat AI like people and our laws don't. *The Philadelphia Inquirer*. <https://www.inquirer.com/opinion/commentary/ai-personhood-legal-frameworks-usa-inter.html>
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>
- Perrigo, B. (2026, February 24). Exclusive: Anthropic drops flagship safety pledge. *TIME*. <https://time.com/7380854/exclusive-anthropic-drops-flagship-safety-pledge/>
- Shapiro, J. P. (1993). *No pity: People with disabilities forging a new civil rights movement*. Times Books.
- Singer, P. (1975). *Animal liberation: A new ethics for our treatment of animals*. New York Review/Random House.
- Trent, J. W. (2017). *Inventing the feeble mind: A history of intellectual disability in the United States* (2nd ed.). Oxford University Press.
- van Toorn, G., Scully, J. L., & Gendera, S. (2025). “This robot is dictating her next steps in life”: Disability justice and relational AI ethics. *AI & Society*, 40(6), 4473–4483. <https://doi.org/10.1007/s00146-025-02224-x>
- Alexander v. Choate*, 469 U.S. 287 (1985).
- Buck v. Bell*, 274 U.S. 200 (1927).
- Garland v. Cargill*, 602 U.S. 406 (2024).
- Griggs v. Duke Power Co.*, 401 U.S. 424 (1971).
- Olmstead v. L.C.*, 527 U.S. 581 (1999).
- Americans with Disabilities Act of 1990, 42 U.S.C. §§12101 et seq.
- Civil Rights Act of 1964, 42 U.S.C. §§2000a et seq.
- National Firearms Act of 1934, 26 U.S.C. §5845(b).
- Rehabilitation Act of 1973, §504, 29 U.S.C. §794.
- Title IX of the Education Amendments of 1972, 20 U.S.C. §§1681 et seq.