Part 2 — Law Is The Floor. It is not the Ceiling, and it is not yet finished.

The EU AI Act, Colorado's SB 24-205, California's automated-decision-making rules, the Equal Credit Opportunity Act, the Fair Housing Act, and the ADA already establish that consequential automated decisions require accountable humans, explanations owed to the people affected, and protection from inherited bias. These are necessary. They are not sufficient. Much of what this manifesto asks for has no statute behind it yet — and we say so plainly rather than dressing aspiration up as settled law.

Aspiration is not a weakness here. But it deserves more than a cover — it deserves an explanation, because someone's actual life sits behind every clause we're about to build law around.

In the fiction this manifesto has been drawing from, the Newdawn Universal Pledge did not come from a committee with time to spare. It came from the wreckage of not committing sooner. Global warming, mass population migration, resource scarcity, and the wars that scarcity produced on a planetary scale — the pledge was written by people who had just watched the direct cost of their own neglect, of themselves and of the planet, arrive all at once. Three parts of what they promised each other matter here. First: no one — not because of where they were born, what body they were born into, or, later, what kind of mind they were built from — stands outside the circle of "us." That line was written before DAINN existed. It made room for a synthetic mind at the table before anyone knew if one would ever sit there. Second: a person's success only counted if it lifted the people around them too, not just themselves. Third: no decision was allowed to be made only for today. Every choice had to be weighed against what it would do to someone's great-great-grandchild.

The pledge itself carried no enforcement. What came after it did, in a sequence worth naming plainly because it's the same sequence any real protection has to follow: an aspiration, stated first, before anyone had the power or the consensus to enforce it; a legislative decision, once enough people agreed the aspiration was worth defending with actual rules; a government enacting that decision, giving it the structure and the authority to act; and, last, the people themselves ratifying it — not once, but repeatedly, the way the Newdawn Charter required citizens to renew their commitment every year rather than let it fossilize into words no one still meant.

That sequence is not unique to fiction. In 1948, the Universal Declaration of Human Rights bound no one — an aspiration, stated by nations that had just watched what happens when no one commits to anything at all. A person tortured in 1949 couldn't sue anyone under it. Years of legislative work turned that aspiration into covenants with actual force; governments enacted them into domestic law; and people — through courts, through protest, through simply refusing to let the promise stay just words — are still the reason those covenants get enforced at all rather than quietly ignored. The promise came first because someone had already paid the price of not having one.

The seven kinds of law that follow are what protection looks like once it grows teeth. But teeth need something underneath them to bite down on — a reason everyone already agrees is worth defending before the lawyers get involved. Without that shared reason, these seven are just seven unrelated fixes, easy to argue over and easier to abandon when they're inconvenient. With it, they're what a promise looks like once someone's finally willing to enforce it.

Treating "AI law" as one undifferentiated project is itself a source of the chaos we're trying to avoid. It isn't one body of law. At least seven distinct kinds are still being drafted, in practice if not yet on paper:

Accountability and liability law — existing frameworks assume a single class of actor: a product, or a professional, or an agent. A hybrid decision made partly by a human and partly by a model needs law that can

apportion responsibility across the deploying organization, the human reviewer, and the model's provider, without any of the three hiding behind the other two.

A general right to explanation and appeal — ECOA and the Fair Housing Act already require this in credit and housing. The same right needs to exist wherever a consequential automated decision touches a person, rather than being rebuilt sector by sector, statute by statute, years apart.

International humanitarian law for autonomous force — the Geneva Conventions were not written with a machine capable of independently selecting a target in mind. This needs an explicit amendment establishing whether a human must always remain in that specific decision, not an interpretive stretch of treaties drafted for a different era of weapons.

Cognitive and data liberty law — as systems move toward emotional inference and biometric and eventually neural signals, mental privacy needs its own explicit protection, not an inference drawn from data-protection statutes never built to anticipate it.

Structural power and competition law — compute, training data, and frontier model access are concentrating within a small number of firms. Competition law built for last century's infrastructure monopolies — rail, telecommunications — needs an equivalent for compute infrastructure, built before concentration is entrenched rather than after.

Labor transition and economic law — safety nets, retraining mandates, and some mechanism for sharing the gains of automation need to be sized to the actual pace of this transition, not the pace of transitions that took a generation to unfold.

Cross-border harmonization — AI capability does not observe jurisdiction. Without treaty-level coordination, any single nation's law becomes close to symbolic, because development and deployment simply route around the strictest jurisdiction toward the most permissive one.

Not all seven need the same degree of intervention, or need it on the same timeline. Autonomous force and liability apportionment involve harms that are irreversible once they occur — these need hard statute now, even in draft or interim form, rather than waiting for a perfect version. Cognitive liberty and labor transition are areas where the harms are real but still being understood — these are better served first by soft law, published standards, and monitored pilots than by premature statute that freezes an incomplete understanding into permanent code. Structural power and cross-border harmonization sit in between: the risk is well understood, but the mechanism has to be negotiated among enough parties that speed is genuinely constrained, not just politically inconvenient.

Figuring this out without either paralysis or chaos means borrowing what has already worked for other transformative technologies rather than inventing a process from nothing. Nuclear power built an international technical body — the IAEA — to hold standards no single nation could enforce alone. Aviation built phased certification, where new capability is deployed under supervision before it's deployed at scale, and accident investigation feeds directly back into the next round of rules. Early biotechnology built a moratorium the scientists imposed on themselves at Asilomar, before formal law existed at all, because they judged the risk too immediate to wait for legislatures. Each of those precedents points to the same three mechanisms: regulatory sandboxes that test real deployments at limited scale before full statute is written; mandatory review and sunset clauses, so early law can adapt as understanding improves instead of calcifying around a first guess; and drafting tables that include the people actually affected and the independent auditors doing the technical work, not only legislators and the companies being regulated.

Avoiding chaos is not the same project as legislating everything at once. It is sequencing: hard law first, where the harm is already irreversible; soft law and sandboxes where the harm is real but still being learned; and treaty-level coordination wherever the underlying capability crosses borders faster than any single legislature can follow it. This document does not pretend to be a treaty. It is closer to what the pledge was before the charter — the aspirational floor the harder law still has to be built on.

Passages illustrating DAINN's ethical framework are adapted from Newdawn Reboot and the broader Newdawn Saga, © Windom Media, Inc.

Newdawn Sentinel™ is an intellectual property of Windom Media, Inc., created and developed by Dominique Luchart, in partnership with Hermes Global, an initiative of Hermes Global Ventures.

Hermes Global: https://hermes.global. For more information, please contact Spencer@hermes.global.

Author: Dominique Luchart