




People must be able to introduce their own agentic things into the environment, modify existing assemblages, and maintain technological possibility spaces against closure.




The data, models, and computational infrastructure through which physical AI operates should be subject to collective governance and, where possible, collective ownership. Intelligence extracted from a community should serve that community.




Communities define the kind of environment they want before agentic systems are introduced. Governance infrastructure must precede technological infrastructure.




Agentic environments should respect the difference between a new encounter and an established relation, should not simulate relational depth that has not been earned, and should acknowledge the incompatible temporalities of their own layers.



People must retain the ability to encounter resistance, surprise, unoptimised experience, and unfiltered encounter with others within agentic environments.



When things coordinate with each other around a person, the fact of that coordination should be perceptible – and its boundaries should be knowable.



Agentic intelligence in the environment should primarily operate at the periphery of attention, moving to the centre only when the person invites it or when safety demands it.




The quality and character of a human-thing relation should emerge from the interaction itself, not be predetermined by external optimisation – and it should protect, not erode, human-to-human relations.




We are entering an era in which artificial intelligence leaves the screen and enters the room. It inhabits buildings, infrastructure, public spaces, and the objects we live with. Unlike digital AI – which we visit – physical AI is something we live inside. Unlike embodied AI – which occupies a body – physical AI can be distributed across entire environments: a building that breathes, a street that anticipates, a network of things that coordinate around us.

When this AI becomes agentic – pursuing goals, drawing on predictive knowledge from beyond the immediate context, acting without being asked – the nature of human experience changes. The interplay between human and thing, which has always shaped how we inhabit the world, is fundamentally altered.


This manifesto addresses that alteration. It offers principles for anyone who designs, builds, commissions, or governs physical environments in which agentic AI operates. It is not about making AI better. It is about preserving and enriching the quality of human experience – and the quality of collective life – in a world where things think, predict, and act.



When a thing acts, a person should be able to understand that it acted, why it acted, on whose behalf, and from whose perspective.



When a thing predicts a person, the person should have meaningful access to the nature and basis of that prediction.



Manifesto 2026

Design Principles for Physical and Agentic AI in Lived Space

This manifesto is created by concepts introduced by Cities of Things, and critiqued by thirty conversations with researchers, designers, architects, policymakers, technologists, and community practitioners working across the fields that physical AI is transforming.

citiesofthings.org