

ON A.I. AND CITIES: PLATFORM DESIGN, ALGORITHMIC PERCEPTION, AND URBAN GEOPOLITICS

LECTURE BY BENJAMIN BRATTON

Benno Premsele Lecture
1 November 2015

Many of you are familiar, I should think, with the Sanzhi Pod City near New Taipei City in Taiwan. This future city was late for its own birth, which was in 1978. Originally planned as a vacation resort for US soldiers, the project was doomed by a series of mysterious car accidents, and abandoned in 1980. The future lasted only two years. However, when demolition work began in 2008 it was discovered that not one but five species of orchid mantis as yet unknown to science, had overtaken the ruins, and multiplied to a population of an estimated 10 million insect inhabitants, above ground, underground, in the structures, in between them. No one knows how or why. Etymologists observe that the unintended orchid mantis civilization has developed an incredibly complex division of labor, not only between species but also within species. Strict systems of food capture, nest construction, and types of stigmergic communication between individuals of different species that has never been observed anywhere before. Their work has also resulted in the proliferation of new sub-species of orchids flowers, which the insects resemble and from which they get their name. Orchids usually don't grow in this part of Taiwan, but today thrive in the unusual labyrinthine cold and darkness provided by the Mantis' own architecture.

The future city is not for humans. The Anthropocene, the reframing of the Earth in the image of industrial modernity, will be short lived. It will be less of a geologic era than a geopolitical instant. Humans are vanishing. Our cities are not our own. We are building the habitats for life forms other than our own. We are their tools; we are the robots for future insects.

The extraordinary architecture of Sanzhi - that is the systems built by the orchid mantis on top of and in-between the UFO pods -, has become, in a short 30 years, a precious future-archeological resource. It is not a failed future, but a successful one. It is our future. We are already its present. We who are displaced by the orchid mantis.

I find that you can tell a lot about someone's materialism (and their realism, speculative or otherwise) by how they think and talk about real matter itself, its importance, its mutability, its agency, its muteness. We take it that computation was more discovered than it was invented; that algorithmic generative processes are intrinsic to the wider unfolding of the world, and that our provisional computing appliances are nowhere near as sophisticated and efficient as

those that precede Homo sapiens. We take it that computation is one of the ways that matter, in whatever form, achieves intelligence via procedural abstraction. Also, Sapien brain tissue is matter that has achieved intelligence in ways that both clearly are and clearly are not “computational” in the convention sense. Lastly we take it that the solution of synthetic computation —our software— into urban fabric provides the landscape of inorganic forms with a sort of distributed intelligence, because it provides them a capacity for abstraction. That is, intelligence as abstraction, computation as intelligence, computation as abstraction. That appreciation should account for two related but different understandings. One would recognize that intelligence (and knowledge) is always distributed among multiple positions and forms of life, both similar and dissimilar to one another. This is not to say that “nothing is true and everything is permutation,” rather that no single neuro-anatomical disposition has a privileged monopoly on how to think intelligently. What might qualify a general intelligence not duty bound to species or phylum is its capacity for abstraction. Ray Brassier suggests that the ability of an organism, however primitive, to map its own surroundings, particularly in relation to the basic terms of friend, food or foe, is a primordial abstraction from which we do not graduate so much as learn to develop into something like reason and its local human variations. In this way, mapping abstraction is not an early stage through which things pass on their way toward more complex forms of intelligence, rather a general principle of that complexification.

Like protozoa and their ganglia, feeling about to figure out what is out there, or like humans looking, tasting, and imagining patterns, today’s forms of AI are (sometimes) augmented by various technologies of machine cognition, machine sensation, and machine vision that allow them to see and sense the world ‘out there’ and to abstract the forms of a (mechanically) embodied intelligence, both deliberately programmed for them and emerging unexpectedly. So we discuss A.I. at the level of the city in terms of a certain “distribution of the sensible” to twist Ranciere’s phrase to new ends. These include: 1. “Thinking” (Including Turing Test, humanistic models of thought, epistemological and practical dangers.

From deep infrastructure to immediate affects, what we today gather under the name “artificial intelligence” will shift not only what counts as “thinking” but also what counts as architecture, design, politics and programming.

“Epidermal Media” synthetic skin and the designed convergence of natural sensation and machine sensing, for animal skins, like ours but also for urban skins, surfaces and interfaces. Implications for dwelling. 3. “Machine Vision” camouflage, inverse uncanny valley, AI apophenia and paranoia, etc.

I would begin by arguing for an expanded landscape of thinking also by disqualifying the normal alternative, which is to model AI in terms of how humans think that humans think. The [Turing Test](#) is a model here, and its demand that A.I “pass” as human is troubling. As dramatized in [The Imitation Game](#), the recent film biography of Turing directed by Morten Tyldum, the mathematician himself also had to “pass,” in his case as straight man in a society

that criminalized homosexuality. Upon discovery that he was not what he appeared to be, he was forced to undergo horrific medical treatments known as 'chemical castration.' Ultimately the physical and emotional pain was too great and he committed suicide. The episode was grotesque tribute to a man whose recent contribution to defeating Hitler's military was still a state secret. Turing was only recently given posthumous pardon, but the tens of thousands of other British men sentenced under similar laws have not. One notes the sour ironic correspondence between asking an AI to "pass" the Test in order to qualify as intelligent—to "pass" as a human intelligence—with Turing's own need to hide his homosexuality and to "pass" as a straight man. The demands of both bluffs are unnecessary and profoundly unfair. Should complex AI arrive, it will not be humanlike unless we insist that it pretend to be so, because, one assumes, the idea that intelligence could be both real and inhuman at the same time is morally and psychologically intolerable. Instead of nurturing this bigotry, we would do better to allow that in our universe "thinking" is much more diverse, even alien, than our own particular case.

The real philosophical lessons of AI will have less to do with humans teaching machines how to think than with machines teaching humans a fuller and truer range of what thinking can be.

It is better to examine how identification works from our side of the conversation. It is clearly much easier to make a robot that a human believes to have emotions (and for which, in turn, a human has emotions, positive or negative) than it is to make a robot that actually has those emotions. The human may feel love or hate or comfort from the AI, but he or she is reading cues, not detecting feelings. What seems like empathy is really a one-way projection mistaken for recognition (like the Turing Test itself), and not based on any mutual solidarity.

In other fictions, policing the imitation game is a matter of life and death. The plot of Ridley Scott's film *Blade Runner* (1982), based on Philip K. Dick's novel, *Do Androids Dream of Electric Sheep?* (1968), hinges on the Voight-Kampff empathy test, which differentiates humans from replicants. Replicants are throttled in two important ways: they expire after just a few years, and they have, ostensibly, a very diminished capacity for empathy. Deckard, the Harrison Ford character, must retire a group of rogue replicants but first he must find them and in this world Turing Test thresholds are weaponized, lest replicants pass as humans and trespass beyond their station. By the film's conclusion, Deckard (who himself may or may not be a replicant) develops empathy for the replicants' desire for "more life" and arguably they too: at least Roy Batty (Rutger Hauer) seem to have empathy for Deckard's own dilemma. His dilemma (and ours) is that in order to enforce the gap between the human and the AI, defined by empathy or lack thereof, Deckard must suppress the empathy that supposedly makes him uniquely human. By forcing him to quash his own identification with the replicants that supposedly cannot have empathy in return, the principle of differentiation requires its own violation in order to maintain itself.

In the second part of this talk I will discuss A.I. at the level of the city in relationship to skin, vision, and thinking. In the third part I will end with some provisional design admonitions.

Computation has evolved into a planetary-scale megastructure. It is both a vast envelope that frames cities and an elemental substance that helps define every object; we are both inside of it and it is inside of us.

Some designers (architects) may see software as something added onto space. They see the idea of smart cities as stupid because it assumes that cities are not already intelligent. They are right. Some designers (programmers) may see cities as modules of hardware that fit together, one at a time into that megastructural matrix. From this perspective, our cities are a kind of computational hardware that fit together, city by city, into that megastructure. For this, computation is an elemental substance that helps define the physics of every urban object. Google as exemplary second “designer” here, both in terms of driverless cars but also embedded software as embodied in/as a “city”. It turns out a better way of getting rid of parking lots is not zoning but augmented cars with cameras and sensors. They too are right, and so one mode of design is nested in the other. So to describe this for two terms (which I hate btw), seen together, the Smart City is inside the Internet of Things just as the Internet of Things is inside the Smart City. There are some obvious design issues implied here: the rise of networked platforms for physical world services and HAI (Human-level Artificial Intelligence) etc. But it’s not the “alpha soixante” “smart control and production” variations that interest me. Other important design implications include how the City layer of The Stack automates inside-out sovereign decisions, how a popular fear/enthusiasm for A.I. is fear of Copernican shifts, and reconnection of “the User” to include non-human actors and implications for our geopolitics, namely the redefinition of sovereignty.

Planetary-scale computation takes different forms at different scales—energy and mineral sourcing and grids; subterranean cloud infrastructure; urban software and public service privatization; massive universal addressing systems; interfaces drawn by the augmentation of the hand, of the eye, or dissolved into objects; users both over-outlined by self-quantification and also exploded by the arrival of legions of sensors, algorithms, and robots. Instead of seeing all of these as a hodgepodge of different species of computing, spinning out on their own at different scales and tempos, we should see them as forming a coherent and interdependent whole. These technologies align, layer by layer, into something like a vast, if also incomplete, pervasive if also irregular, software and hardware Stack.

Cities we see as (1) provisional settlements organized around agricultural provision (an initial culinary geo-engineering) and (2) as a node on the horizon orienting mobile people and things as they migrate and flow. The first is the city as singular megastructure and the second is the city as a cluster of way-finding interfaces. The city is home to millions of species, from microbes to insects to vegetation to sapient mammals. It is a living bacteriological and immunological tumult, a situated ecology of predation and symbiosis across multiple scales

(and now we insert evolutionary robotics into this). Cities defined in this way, with perspectives mixed from public health planning and supply-chain logistics urban systems, are infrastructure for living and non-living matter to consume itself and for some forms of matter to achieve and hone sentient intelligence (which in turn remakes cities in its image).

AI “Intelligence” is one way that matter organizes itself into durable complexity. A special form of that complexity is the city: a settled accumulation of a material intelligence, both human and inhuman. As Artificial Intelligence becomes more sophisticated what will be its urban design project? What should it be? That is, I mean AI both as something that we see around us all the time and in terms of forms of synthetic reason that we are in the process of designing and designating, and which will contribute to the formulation of cities to come. We see it in what philosophy calls ‘the Other Mind’ problem, how to communicate with an intelligence with which you can only share very partial worlds (which would include the alien in Solaris or 2001 but also each animal).

We try to see AI not in terms of how we think that we think, as a sort of virtual or artificial human cognition, but as thinking and embodying another spot than we do on a larger, shared continuum of material intelligence.

I want to quickly hone in on ways that abstraction can be dermal and cartographic: how it occupies and agitates skins (human skin, building skin, any skin) and how it does so for purposes of mapping. I will touch briefly on each of these. First skin. Part of the research I am engaged in at UCSD, in collaboration with the Dept. of Nano engineering, Bioengineering is on how presumed differences between natural sensation and machine sensing can be elided by technologies operating in and on the skin and at the scale of our dermal and epidermal somatosensory system (mechanoreceptors, thermoreceptors, etc). By this I do not mean some OOO word game for which machines are bequeathed panyschic capacities of intuition and affect. I refer instead to real-deal techniques by which the biological and non-biological physio-chemical reactions are interwoven. Designing chemical sensitivity at or near the atomic level and registration of these as transmissible information with dermal and epidermal microelectronics at larger scales, we conceive a mediatization of sensate surface. There are many researchers around the world whose work is representative of this interest: John Rogers at the University of Illinois and Joseph Wang and Todd Coleman at UCSD. Artificial skins are among the most interesting areas of new design. DARPA has long funded research in advanced prostheses, and currently one research track focuses on providing synthetic tactile sensation in the “finger tips” and other dermal surfaces of corporeal extensions for the wearer. In other labs one can find many varieties of epidermal microelectronics that augment living skin’s capacities for sensing external or internal stimuli. Next door, nanoscale chemistry cooks up “inks” that react to the presence of ambient trace elements of whatever particles they are tuned to react with.

Together, technologies suggest a skin-based microstack comprised of biological, chemical and mechanical sensing and processing in various combinations. As an emergent formulation

of wearable computing (here including also biochemical reactions) it would blend animal sensation and machine sensing into composite surficial media. As platforms for haptic interfaces, its terms of “interaction” may be below or above the normal perceptual scales, and may involve autonomic nervous reactions as much if not more than the relatively blunt gestures of hands and thumbs. Keeping with the skin of human bodies for a moment, we recognize that while photography and cinema made it possible to compose -and to see- the sort of images that we had never seen before —montage, slow motion, double exposure, etc.— and similar goes for synthetic and recorded sound and our auditory senses. Our largest sensory organ is our skin, however, quite obviously, to date we have not developed nearly as diverse or sophisticated forms of artificial touch, tactility, feeling for skin based media as we have for vision and hearing. Clothing is one way that design has leapfrogged genetic evolution to augment skin: for example, engineering scuba suits is faster than evolving fur or blubber is. Today’s dermal/epidermal media make it possible to sense things about the external world that our skin is otherwise not able to register (particulate matter in the air sensed as parts per trillion) or things about our internal state (ph balance in sweat).

The longer-term and more interesting design horizon is skin-based media that allow bodies novel and unnatural forms of touch-sensation that we have perhaps never before experienced. Pleasure and pain are on the menu but much more besides. Again, early and experimental cinema innovations of artificial vision, sights and sounds is one precedent. Molecular gastronomy and its rebuilding of food from first chemical principles is another, if we include taste as a particularly nuanced form of touch, which of course it is.

What are today novel tastes are models for other new sensations to come.

Now by shifting this line of inquiry up to the urban scale, what we take to be the artificially mediated surface includes things that are analogous to skin but also what we today refer to colloquially as machine “vision”. Vision has evolved independently many times in evolutionary history and arguably in the past two decades it has evolved again, this time not for cuttlefish or rattlesnakes but for networked CCD sensors and algorithmic armatures processing what is sensed into differentiating and motivated recognition. “Visual” sensors responding to light, often but not always shaped like cameras, are one kind of designed surface capable of synthetic sensation and computational interpretation, but in the wider urban landscape they co-mingle with networks of other surficial sensors responding to motion, pressure, heat, ambient air qualities, pressure, etc.

To the extent that it is convenient, we call these variously machine vision or machine hearing or machine skin, but any correspondence from and to the mammalian sensory system is merely allegorical. The AI city may be embodying itself, but not as humans do. When information is scarce, then copying something is the work of the mechanical image. However, when information is abundant (especially when overabundant) then seeing the original, picking its pattern out of the background, is the work of machine vision. Today many images are made for no one, but this does not mean that they are functionless. They are made by and

for machines that “see” the world differently than we do. Those machines do not have eyeballs, rods and cones, and a visual cortex, but do have sensors that detect light, motion, form, heat and color in other ways. The animalian vision-image making nexus is primarily human, and it is likely that from the beginning to the end of the Holocene the total quantity of images that humans have produced from cave walls to FaceTime (measured per year, in totals pounds of pictures, total gigabytes of information, etc.) continues to increase exponentially. With digital imaging machines now in every person’s pocket, the raw sum of images produced since the year 2000 may be more than all those produced before that year. It depends on how you quantify “an image,” but what Walter Benjamin called “mechanical reproducibility” has certainly allowed images to proliferate far beyond the means of human craft.

In the last decade or two, “vision” has arguably evolved once more, this time not for birds, cephalopods or rattlesnakes but for inorganic species of machines that can make and process images for people to see or for other machines to analyze. Today, the industrial scale processing of data that has been gleaned by scanning the light spectrum in some way—from urban-scale street surveillance to millimeter-scale quality control along assembly lines—represents a significant fraction of all the work that the world does to image itself for the purpose of governing human society.

At the end of the day, the mechanic phylum takes more selfies than selves do.

The function of representation is very different however. The “image” likely remains data, and is never rendered to look like a “picture” because there is no need. An algorithm programmed to discern a particular pattern or anomaly can “see” it directly in the data itself. It does not necessarily need that data to be projected, as if for a mammal, and then re-seen and re-interpreted back into code. Like plants, do machines also possess a kind of vision without images? Or at least a kind of image without natural abstraction toward and from corporeal experience, that is, an abstraction that is based in chemical or informational pattern-finding but not as a cartographic simulation of experience?

Perhaps surprising to some, the function of machinic images today (and of their abstractions both animal and algorithmic) is to determine the veracity and originality of what those images represent. For example, some of the shapes printed on a dollar bill are there for humans to differentiate the value of one token from another, but far more are for machine vision counterfeit detection to verify that this is a “real” dollar.

The piece of paper is full of machines that happen to look like pictures; it is an image-as-machine.

Or, a surveillance scan of a city may pick out one face from thousands in motion, looking for the one true target. Elsewhere, inserting tiny camera probes into great paintings so as to verify that they are originals is often required by insurance companies backing purchasers.

Online, captcha software shows an image, and then by a quick inversion of the Turing Test, analyzes how the user interprets it and determines if they are in fact a real person. We conclude that Walter Benjamin's assurances that mechanical reproduction would undermine the aura of the original is true when, for example, we compare a painting with a postcard of a painting. However machine vision and images-as-machines are put to work to ensure auratic originals, verified non-fakes, true identities, unbroken versions, normal targets and certified real deals. His historical arc from pre-mechanical original to mechanical copy is incomplete without another curve, one leading now to the machinic authentic.

Still, the machinic visual subject is not something that possesses humanlike or human-level perceptual and aesthetic capacities, but rather something that is uncanny and interesting because it does not possess those things and yet can see us, recognize us and know us regardless. That's weird and interesting enough. There is the question of how the world looks as a screen, and another, more important I think, is how we look as objects of perception from the position of the machines with which we co-occupy that world. Seeing ourselves through the "eyes" of this machinic Other who does not and cannot have an affective sense of aesthetics is a kind of disenchantment. We are just stuff in the world for "distributed machine cognition" to look at and to make sense of. Our own sapience is real and unique, but as we are things-to-observe-that-just-happen-to-be-sapient, this doesn't really matter to machine vision. This disenchantment is more than just like hearing the recorded sound of your own voice ("that's not me") it is potentially the clearing away of a closely guarded illusion. This uncomfortable recognition in the machine's mirror is a kind of "reverse uncanny valley."

Instead of being creeped out at how slightly inhuman the creature in the image appears, we are creeped out at how un-human we ourselves look through the creature's eyes.

In writing a book on artificial intelligence I find have found that it is impossible to fully separate the technology itself from the sometimes bizarre ideas that we have for it and about it. I am fascinated by how technologies are not only anthropomorphized, but how some are considered menacing and others amazing, and some both at once.

Whether the aliens are reading your thoughts through radio waves and dental implants or Google is reading your thoughts through email, the borderlands between schizophrenia and sensible vigilance – delusion and deduction – are disputed territory.

That is, the images are remarkable enough on their own, irrespective of providence, but once the viewer makes the partial empathetic transference into the Other Mind of the feature-recognition algorithms and their way of seeing, the effect is of trying on a new sort mode of perception and making sense of the world through those eyes: that is the real payoff. A

character in Peter Watts' recent novel, *Echopraxia*, argues that false pattern recognition is hardwired into *Homo sapiens*' evolutionary success. The story goes like this: Tens of thousands of years ago, two guys are hiding in the tall grass. Quietly they are looking all around for both predator and prey. One of them sees a faint but distinct anomaly in the way the light breaks through the grass moving left to right. He recognizes this pattern as that of a looming tiger and runs away back to the village, leaving his friend to be eaten. The guy who ran away was able to reproduce. His pattern-recognition genes ensure his and their own survival, and the guy who did not possess these genes did not reproduce. However, there was a third guy sitting in similar grass perhaps the morning earlier. He too saw a weird pattern in the grass. He too thought there was a tiger and got up and ran back to the safety of the village. But in his case there was no tiger – it was all in his mind. He was not cunning; he was paranoid. Yet, he too was able to reproduce and his faulty pattern-recognition genes (which are perhaps actually the same as the first guy's accurate pattern-recognition genes?) were able to reproduce themselves as well. The lesson is that evolution has greatly rewarded human pattern recognition, but bound inextricably with that bounty and bargain it has also rewarded hallucination and error.

Consider Google's Deep Dream, by which feature-recognition algorithms are looped back on themselves resulting in amazing hallucinatory synthetic apophenia.

That is, Google's algorithms are themselves paranoid; they see psychedelic dog faces in everything.

The reason that Deep Dream hallucinates dogs is not spooky. Its object recognition faculties were trained using the ImageNet data set, which is based on various breeds of dogs. It will therefore hallucinate dogs where there are none and identify people as dogs if it is told to look at them hard enough. The latter may seem insulting, but it shouldn't. Diogenes' proto-cosmopolitanism is based on the glorious dog-like commonality of all. If the Deep Dream images are artifacts of a computer's hallucinations of phantom conclusions, then the conspiratorial figure cut by Evil Google Stepfather is that same paranoid vision turned inside out and back on itself. Perhaps Google, the AI, is as paranoid in how it sees us as some of us are in how we see it. It is not surprising, then, that as AI matures, its own pattern-recognition faculties would reach the plateau of creative apophenia. From Antonin Artaud to Seymour Cray, many have perceived homologies between abnormal human psychology and various machine behaviors. Beyond simply anthropomorphizing intentions ("the Xerox is mad today") our deduction and induction of mechanic intelligence trace along the paths of understanding of our own intelligence. It is one of the few such homologies our humanism will allow. That said, still, the trove of allegories available from this history of machine psychiatry still holds considerable interpretive and even creative potential despite itself. (In Kubrick's 2001: A Space Odyssey for example, HAL was obviously a deeply paranoid creature. Tormented, the AI decided that the mission to Jupiter could not withstand the dark conspiracy of humans and so jettisoned them.)

Is machine vision paranoid? Is our popular understanding of machine vision paranoid? What is Google up to? Are they watching you, and if so who (or what) are they seeing? Is our understanding of Google's paranoid machine vision itself paranoid, making the paranoid AI that much more paranoid in response? In a word: yes. Is this how artificial intelligence should evolve, in relation to and autonomously from human intelligence? If the evolution of human pattern recognition is any indicator – based on both productive and destructive apophenia, deception and delusion – then yes and no.

A co-mingling of diverse sensors (of light, air, sound, chemistry, etc.) draws a landscape of sensing and thinking little species, partially embodied discretely one with another, and partially co-embodied with one another as their information inputs are aggregated, modeled and acted upon in various pluralities. Homo sapiens comes equipped with an extraordinary array of sensory faculties, which as discussed may be augmented further by synthetic layers in various ratios, ranging from the sensors/trackers in our phones that we carry about like mules, to more intimate media of artificial images, sounds, etc. Situating ourselves in this expanded field, we are both sensors and sensed. On the one hand, we are a primary sapient actor in this drama, supervising an orchestra of sensing technologies, each individually capable of functional processing and together of certain forms of intelligence (as are neurons or other cells). On the other, we are not only the subject of the scenario; we are also its subject matter. That wider urban landscape of synthetic sensory systems is not only a platform through which we extend and extrapolate our capacities for abstraction, it also capable of other sorts of abstraction on its own. As part of its intelligence, it looks at us and registers abstractions about us. (This topic of abstraction came up at a symposium I participated in last month at SCI_Arc in LA.

There the question of abstraction was raised consistently in relation to how architectural form and formal imagery might be read and how the designer might relate between the real and the figural. Here I would want to emphasize instead that

abstraction in and on the urban surface is not only a matter for how
architecture is perceived but for how architecture perceives us.

The work of abstraction for urbanism is not only to deploy abstract forms but to set in motion mechanisms and programs that are capable of their own feats of abstraction and to calibrate how they abstract us and one another accordingly. In short, we should imagine AI urbanism in terms of Von Uexküll's stroll into the field populated by intermingling but mutually oblivious little lifeworlds, and/or in terms of Deleuze's parable of the tick. The latter lays in wait for some threshold event to come its way, at which point it triggers its programmed response at and into its own void. Many of our urban sensors and their limited forms of AI work similarly, and with similar nobility.

More versatile synthetic intelligences occupy more complex Umwelt, some are predator and prey, some are in motion, some are flowering, some pollinating. As we stroll among them, we may be registered by them or we may be ignored. We may be a primary cause of concern or

we may be a passing interference in an evolutionary dynamic in which we are neither the protagonist nor target. As many of you know, shifts from top down to bottom up AI have been marked by a shift of emphasis from intelligence as a formal syntax to intelligence as a specifically embodied relation to specific worlds and models. Heuristic knowledge of habitats is seen as inseparable from AI manipulation of a situated problem space. In Robotics, the pairing of synthetic sensing (vision for example) with algorithmic reasoning allows for simple artificial species to perform intelligently because they have what amounts to functionally-embodied perceptual location in their world. They can think because they can see.

In any event, like the protozoa whose ganglia provided it with a basic ability of cartographic abstraction of its environs, today's AI - individually and more importantly in aggregate—are not disembodied. They are embodied by machine sensing (machine vision, networked surgical sensing, etc.) in the city and at the scale of the city. The cities we build are in this way not only the habitat about which AI learn embodied contextual knowledge,..they are also the distributed sensory apparatus with which the AI embodies that context. The city layer —as platform and as governing apparatus— is self-incorporating in this way (others too). So the concern is less A.I. and Cities than Cities as A.I., and better: A.I. as City.

Now among the structural affordances of Stacks is their modularity. What occupies any given layer, including the City layer, can be replaced by new things, and insofar as they communicate with the layers above and below according to the protocols of the platform (in this case, the address layer and the cloud layer) then whole remains. Not only do stacks accommodate the comprehensive replacement of components, their chief design value is that they enable and encourage it. That is, while the figure of The Stack is a figural totality, it is not static. It is made to be re-made; and so to articulate the Stack-we-have is already to anticipate in some way the Stack-to-come

Very quickly then, by way of conclusion: 3 provisional points for a design brief:

Against AI-automation skeuomorphism

Algorithmic governance should be able to enforce rule and also to learn. Blockchain advocates evangelize its decentralized architecture, which is very likely a key means to ensure that accountability. But the transposition of that into a commanding armature, means centralization, it means getting over the fact that platforms are both centralizing and decentralizing at the same time. Again, I am not arguing pro- or anti-platforms, as either a good or an evil, but nor am I arguing that they are neutral. They aren't, and their lack of neutrality that makes them useful as geopolitical design tools. The critique of infrastructure is essential but it must also rotate into infrastructure scale-design models (if it is serious, not just posturing). For this, it is useful to think in terms of totalities, such as The Stack, because it provides a framework for considering the distributed agency, subjectivity, causality and effect. there are not externalities because there is no "outside" to put them.

From User-Centered Design to Redesigning the User: Or, for Universal User Suffrage

The complexities and contradictions of platform sovereignty in relation to nonhuman Users who are nevertheless at least quasi-intelligent and sensate actors work to dismember conventional distinctions between subject and tool, occupant and city, subject and apparatus,

etc. The agnostic universal suffrage of the User subject position should be claimed as strongly materialist (not only legal) cosmopolitan territory. Pushback against that disenchantment is not only a distaste for certain forms of financialization it is also bulwark on behalf of a pre-Copernican human exceptionalism: even fundamentalism. Of the latter, one of the most difficult shifts in our thinking will need to be around who and what is a “user” and who and what is sovereign as a user.

As said, States have citizens, markets have homo economicus, and platforms have users. It doesn't work to treat one as if it were the other.

The same person or thing may be all three at different times but the governing issues, the organizational issues and the design issues are different. The generic universality of platforms makes them formally open to all Users, human and nonhuman alike. If the User's actions are interoperable with the protocols of the platform, then in principle, it can communicate with its systems and its economies. For this, platforms generate User identities whether they are desired or not. Anything that can initiate interactions with the platform can be a user and the platform may see them and interact back with without knowing or caring who or what they are.

Platforms don't care if the state sees you as an illegal immigrant or if the market sees you as an externality.

Platforms ultimately don't care if the User is an animal, vegetable or mineral —all users may have platform sovereignty. In security-speak, a user is credentialed by 3 qualifications. something you know (password) something you are (like a fingerprint) something you have (keycard). If someone or something can be, have and know, it can be user: a trading algo, a driverless car, a sans-papiers, a chemical reaction triggering a threshold reaction in an environmental sensor embedded on a leaf in a rain forest: all users. User is open position, if we defend it as such. To develop political and economic design model of the stack is thereby inseparable from a philosophical (and technological) reconception of the human as a kind of User and of the User as something that is not necessarily human.

One can expect pushback as fervent as it is irrational. There are some affinities with technologies, however fictitious or bizarre, that are thought to embody the essence of human mastery. Guns and cars are among these. They are seen to amplify the human, not contaminate it by mediation and hybridization. I would go so far as to predict that there will be a movement to identify human-driven automobiles as a type of “arms” and that the Second Amendment to the United States Constitution, now used to shield gun owners from obvious liabilities and to protect their sense of personal dominion, will be flown to keep human beings behind steering wheels. Your life may be ended by someone encased in a two-ton steel box careening down the asphalt vista trying to prove a point about how technology will never capture his natural humanity. So while the Stack stages the death of the User in one sense—

the eclipse of a certain resolute humanism (hippie user and mechanistic-utilitarian user, both)—they do so because they also bring the multiplication and proliferation of other kinds of nonhuman Users (including sensors, financial algorithms, and robots from nanometric to landscape scale), any combination of which one might enter into a relationship with as part of a composite User.

A word then on the platform logics of the City layer. Stacks are platforms, but not all platforms are Stacks. Platforms are both technical and institutional models. Their logics of sovereignty are not reducible to those of States or Markets, but our contemporary discussion of the platforms which increasingly do heavy lifting of global governance are discussed as if they were. As I pointed out earlier, States have citizens, Markets have homo economicus, but the base component political subject of Platforms is the User, a very different creature. Platforms work by the centralized aggregation of interactions made by decentralized distributions of interfaces, one because the other. As far as the platform is concerned, User sovereignty derived from the decisions that it would insatiate into a wider interfacial program is available to anything -animal, vegetable, mineral- that can interact significantly with those interfaces and so initiate the entire Stack.

The risks and potentials of platform sovereignty are born by a geopolitics of platforms still very much in formation (one that points as much to Cloud Feudalism as to Fully Automatic Luxury Communism) as well as by the radically open subject-position of the User, one available to humans, networked sensors, HST algorithms, driverless cars, swarms of animals, etc. The interface is agnostic as species. The platforms we are building, including the accidental infrastructure of The Stack, are predicated on a functional agent-subject position that holds far less regard for the special status of human creatures in how it governs with and by sovereign exceptions of exit/entrance than we are used to. Among the most painful and violent contests over platform sovereignty to come, fought at the level of the automated urban interface, will be drawn not only over nomic exception but also over human and humanist exceptionalism. Accomplishing a Vibrant Post-Anthropocene: Geoengineering as Bio chemical-informational-Materialist Geopolitics

The distribution of mutual perfection and co-embodiment implied by sensate AI at urban scale should, if we are attentive and persistent, shift and expand what we take to be an ethics of care, of shelter, and of encounter and negotiation — yes, immunity (for reals). Simmer-Era urban anomie may pale in contrast to the challenges of predator-prey/ camouflage-dissimulation relationships also implied by the emergence of various alien minds and urban scale. Negotiating our way to a livable post-anthropocene necessitates a fidelity to emergence, not to emergency. The advent of robust inhuman AI will provide for copernican disenchantments, ones that should enable a more reality-based understanding of ourselves, our situation, and a fuller and more complex understanding of what "intelligence" is and is not. From there we can hopefully make our world with a greater confidence that our models are good approximations of what's out there (always a helpful thing.)

Arguably the Anthropocene itself is due less to technology-run-amok than to the Humanist legacy that understands the world as having been given for our needs and created in our image. We see this still everywhere. Our computing culture is deeply confused, and is so along these same lines. We vacillate between thinking of technology as a transparent extension of our desires, on the one hand, and thinking of it as an unstoppable and linear historical force, on the other. For the first, agency is magically ours alone, and for the second, agency is all in the code. The gross inflation is merely inverted, back and forth, and this is why we can't have nice things. A well-known thought leader in the world of Design recently wrote, "it is time to invent a world where machines are subservient to the needs and wishes of humanity." Simply breathtaking. If you think so, I invite you to Google "pig decapitating machine" slide and then let's talk about inventing worlds in which machines are wholly subservient to humans wishes. One wonders whether it is only from society that once gave theological and legislative comfort to chattel slavery that this particular claim could still be offered in 2014 with such satisfied naiveté? This is the sentiment —this philosophy of technology exactly— that is the basic algorithm of the Anthropocenic predicament. It is time to move on. This pretentious folklore is too expensive.

So then what about desire? How to map material recombinancy as a theory of subjectivity against the tropes of eros and thanatos is tricky. Freud barely had Darwin let alone genetics and the rest, and so probably should be read like one reads the Greeks. Organic versus inorganic is a minor issue; something having or not having carbon is beside the point. A desire to melt back into the inorganic and to be overcome and overwhelmed by absorbed by matter itself may be a way of encountering and even desiring death but it is actually a very everyday occurrence. It is not a sad destiny to be interwoven with alien matter, and the key point of it is that other diagrams are possible, anatomic and economic. This human body and this Earthly landscape of matter are only the default settings. They are not destiny. Chemistry may drive the most radical forms of the political imaginary, a "culinary materialism, yes" a rubbing the clinamen raw.

It is harder to so than it should be. To pick two, Heideggerians and Singularitarians —two groups with much more in common, and two with which I have ongoing frustrations— have made it unnecessarily difficult to think and talk about species and technology. Certainly as far as they stand in for Design Theory. They have made the inevitable-whether-we-want-to-or-not geoengineering project that really actually constitutes eco-geopolitics for at least the 21st century into something that is either disqualified a priori as a metaphysical atrocity or something that will simply sort itself out by the invisible hand of Moore's Law. In the fetishization of the human experience of human experience, for one, and the psychotic proposition that the self is a stable physical entity that can be expanded exponentially without exploding into fractal plurality, for the other, both make it far too difficult to see that geoengineering project for what it should be: more like molecular gastronomy at landscape scale, a restored and resorted ecology designed to taste itself in new forms of richly-spiced and imaginatively-sauced mutual ingestion.

Last word, lest I am misunderstood to recommend too much today's versions of Smart Cities as models for the urbanism I'm sketching, I will end by misquoting a sentence from Adorno's *Minima Moralia*, substituting "smart cities" for his original topic "escapist cinema": It is not because they turn their back on washed-out existence that Smart Cities are so repugnant, but because they do not do so energetically enough, because they are themselves just as washed-out, because the satisfactions they fake coincide with the ignominy of denial...The dreams have no dream."