Review Article
Open Access

Distributed cognition: assessing the structure of urban scale artificial intelligence

Abstract

Urban scale artificial intelligence (AI) is most frequently structured to sense and gather as much information as possible. Is this the most appropriate evaluation of intelligence? Many descriptions of so–called Smart Cities focus exclusively on their sensorial capabilities, but little is given to their cognitive capacities. Discussion of the deployment of computation within the urban environment has largely avoided notions of cognition, though a capacity for cognition is ultimately what we are asking of anything that is to be "smart". Cognition itself may be presented in a variety of ways, and determining what structure is most appropriate for urban scale AI is a critical discussion with significant implications for our collective ecological footprint. This article attempts to frame the integration of urban scale AI within a discussion of cognitive structures by building from the work of Benjamin Bratton and Edwin Hutchins to analyze the material culture of Masdar City and the Internet of Things, ultimately arguing that strategies of distributed cognition are both more feasible and more performative than a traditional Smart City model.

Keywords: urban scale artificial intelligence, smart city, distributed cognition, urban design, networked transit

Volume 2 Issue 5 - 2017

Alexander Webb
Department of Architecture, University of New Mexico, USA

Correspondence: Alexander Webb, Department of Architecture, University of New Mexico, USA, Email awebb4@unm.edu

Received: June 22, 2017 | **Published:** July 19, 2017

Abbreviation: AI, artificial intelligence.

Introduction

Urban scale artificial intelligence (AI) is most frequently structured to sense and gather as much information as possible. Is this the most appropriate evaluation of intelligence? Many descriptions of so–called Smart Cities focus exclusively on their sensorial capabilities, but little is given to their cognitive capacities. Discussion of the deployment of computation within the urban environment has largely avoided notions of cognition, though a capacity for cognition is ultimately what we are asking of anything that is to be "smart". Cognition itself may be presented in a variety of ways, and determining what structure is most appropriate for urban scale AI is a critical discussion with significant implications for our collective ecological footprint. Though it may be difficult and potentially unproductive to speculate on precisely why discussions of cognition have not occurred, there are suggestions of an anthropocentric bias towards the organization and identity of intelligence. Both early cybernetic investigations and more recent investigations in cognitive science have suggested a broader understanding of how cognition is organized, and have much to offer discussions of the deployment of artificial intelligence within an urban framework. Discourse on the nature of intelligence in other fields has the capacity to dramatically reshape the current deployment of urban scale AI. The vast ecological and sociological issues that benefit from urbanization on a global scale suggest that cities should be part of larger discussions of climate change and other issues of sustainability. As cities reduce resource use per capita through efficiencies of transportation and production, some architects and urban designers are leveraging every tool at their disposal to not only entice potential inhabitants but to also reduce the impact of cities during the process. As a vibrant discussion around the benefits of an urbanized planet continues, there should be discussion regarding how to both encourage and leverage this increase of urbanity with an effective integration of urban scale AI.

This article attempts to frame the integration of urban scale AI within a discussion of cognitive structures. One difficulty of the discussion is that urban scale AI is difficult to prototype or simulates urban conditions are resistant to prediction and therefore difficult to predict how they will respond to additional interventions.[1,2] Here, models of cognition are contextualized through various precedents as a means of determining their application as functional metrics for a smart city. The primary metric for evaluation will be the capacity for the system to respond to changing contexts and user needs, and the computational efficiency in which they do so. This article examines the question of how to evaluate organization and structure of urban scale AI by building from the work of Benjamin Bratton and Edwin Hutchins to analyze the material culture of Masdar City and the Internet of Things, ultimately arguing that strategies of distributed cognition are both more feasible and more performative than a traditional Smart City model.

Scale AI

There has been much critique of the organization of the sensor–based Smart City, a model where retrieved information is fed to a central server or brain.[3–9] Smart City projects are led by multinational corporations such as Cisco and Microsoft, and are becoming more utilized as a model for development.[10–13] In 2012 it was reported that over 180 developing cities in China were following a similar model.[14] The consensus of critics is that their deployment is hierarchically–structured in terms of computation and population agency and unresponsive to geophysical context. These critiques are typically focused on ground–up, sensor based cities such as Songdo, and Masdar City.

A common counter proposal is a citizen–led series of micro interventions, frequently referred to as civic hacking.[15] Anthony Townsend's Smart Cities: Big Data, Civic Hackers and the Quest for A New Utopia describe civic hacking as attempts to amplify and accelerate urban social dynamics through code, electronics and interfaces.[15] Where municipalities make data freely available, and individuals and small start–up companies 'hack" solutions to improve their cities. While Townsend makes a compelling case for this approach, he does acknowledge the shortcomings as well. Civic hacks frequently solve problems for small groups of users, but are difficult to scale up to be productive for an entire metropolis.[15] To do so would require leveraging the involvement corporate software engineers to execute the drudgery of developing robust and serviceable software.[15] Townsend underlying suggestion is well taken, corporations like Cisco should not be the only leaders in determining the design of urban scale AI, but the alternative of a completely flat organization is unproven as well.

Antoine Picon[16] shares a similar critique of the Smart City's organization. In Smart Cities: A Specialized Intelligence, Picon[16] unpacks the theoretical, anthropological, political and sociological forces that drive the organization of the Smart City. In doing so, he describes the juxtaposition of civic hacking against corporate development as a false dichotomy. From this perspective, Cisco and other multinational corporations are not necessarily against a democratic civic participation, it just is not one of the products they are currently producing. Picon[17] suggests that other stakeholders are involved, particularly the millions of anonymous stakeholders, the city's residents, who are experimenting with a new relationship to the urban environment. To Picon,[18] the important distinction is between ideologies, what he describes as a neocybernetic technocratic approach versus more participatory strategies. Picon[18] interests are how permeable technological interventions are within an urban landscape, not necessarily how different interventions they are structured or related to each other. Picon[18] is helpful in de–politicizing the argument for civic hacking, but ultimately does not negate the underlying strategy or propose an alternative structure, calling for alternatives to be presented.

Mark Shepard[7] Sentient City– Ubiquitous Computing, Architecture, and the Future of Space is a curated collection of essays and projects that respond to his concept of the sentient city. To Shepard,[7] the term describes infrastructure capable of sensing and responding to events, activities and behavior within the city. Here, sentience refers to the ability to feel or perceive subjectively, but does not necessarily include self–awareness.[14] As Picon[18] asks, through directing his critique at the work of Carlo Ratti's Senseable City Lab at the Massachusetts Institute of Technology, that the model of the sentient city is intelligent– but in what sense? Notions of intelligence are present, but unclear. Ultimately, Shepard[7] argues for a sense or strategy of intervention within the city, but does not argue for a specific structure of urban scale AI. Urbanist Adam Greenfield[19] has written several essays on the organization of urban scale AI. Against the Smart City (The City is here for you to Use) is a similar call for a flat, citizen–led alternative to the Smart City. Though much of the pamphlet reads similarly to other work mentioned here, his argument becomes inconsistent. He criticizes the sensor–based Smart City as a neoliberal tool of deregulation, but advocates for a deregulated software exchange platform in line with the android app store with little explanation. The question of regulation is a critical component of the discussion around the politics of urban scale AI, Greenfield[20] variable position confuses the discussion. In "Practices of the Minimum Viable Utopia", Greenfield[21] suggests a strategy of urban scale AI that mimics the software development practice of lean development as an alternative to the Smart City. This strategy, though possessing structural implications, is more focused on a method of developing interventions as to how they are organized. Greenfield[22] describes a scenario where computation would facilitate human decision–making, but not be making decisions itself. Greenfield[22] Everywhere: The Dawning of Ubiquitous Computing describes behavioral and experiential aspects of ubiquitous computation, but not its structure.

Other authors have described aspects of urban scale AI that have structural implications. Duncan McLaren and Julian Agyeman's Sharing Cities: a Case for Truly Smart and Sustainable Cities describes the potential of utilizing digital technology to promote the sharing of urban resources, therefore increasing the social dynamics of a city. In The City as Interface: How New Media Are Changing the City, Martijn de Waal[23,24] describes how technology can be used to augment social engagement within a metropolis, social behaviors that could inform computational organizations, but does not specifically address how the technology is structured. Benjamin Bratton[25] forthcoming book on artificial intelligence suggests an understanding of urban scale AI as a distributed sensing system and describes a close relationship between the Smart City and the distributed system of the Internet of Things (IoT). Many of these observations are trapped by similar understandings of intelligence. The Smart City and Civic Hacking strategies understand artificial intelligence as sensing information to present to human actors. This understanding is problematic when framed by the issue of urban simulation. The enormous datasets collected by computational agents must be presented to humans in a way that we might understand; therefore a degree of representation must be deployed. If the city as a system is complex enough that it defies representation, the information presented is inherently suspect. In this sense, civic hacking is a more productive structure than the sensor–based Smart City. The distributed, multifaceted strategy of civic hacking eschews prediction as, in the words of Town send; the hackers are essentially "tinkering" towards a more performative city.

The examples provided in Sentient City adopt a similar flat approach, but avoid the issue of simulation. In the collection of essays Martijn de Waal[23] proposes a blend of tracking and sensing devices that collectively communicate to inform an understanding of the city's behavior. Shepard & de Waal[7,23] allude to strategies of sensing to produce computational responses, an aspect of the sentient city model is that creates a more direct response from the mechanic, but they do not articulate a framework to develop these responses in detail. The proposals of civic hacking do not explicitly describe AI's role or organization, but provide a strategy where it may be more productive. Ultimately, if the role of urban scale AI is to present information to human actors, the presentation of the information through simulation creates a problematic issue– how can we represent a city when a city defies representation? Instead of using computation to present information to a human intelligence, many are intervening within the city by leveraging artificial intelligence directly within the urban context. When considering the vast amount of information cities collects, there is a similarly vast amount of necessary responses. By definition, computation is necessary for negotiating Big Data and implementing these responses intelligently. Creating a framework for urban scale AI that enables artificial intelligence to have significant agency further leverages this capacity. Rather than define intelligence through its sensorial capabilities to feed a human intelligence, distributed cognitive elements that both sense and compute is more productive (Figure 1).

An internet of processors

In attempting to deconstruct our collective appropriation of artificial intelligence within the urban environment, it may be productive to assign simplistic analogies to the urban scale AI

Citation: Webb A. Distributed cognition: assessing the structure of urban scale artificial intelligence. *Int Rob Auto J.* 2017;2(5):187–194.
DOI: 10.15406/iratj.2017.02.00034

strategies described so far. The Smart City model, to date, is essentially a mammalian intelligence in a vegetative state, receiving information but not responding to it. The civic hacking, as well as the sentient city approach, is an array of individual sensors or actuators, with little reciprocity or relation to each other. The model advocated for here is an array of networked computational systems that are both unified in a cohesive whole but irreducible to their constituent parts. This argument borrows from research conducted within the fields of cognitive and computer science. The term, artificial intelligence, it has been so entirely overused by software developers, tech companies, and news and media outlets that it has lost all meaning within common discourse.[24] For the purposes of this article, investigations into the structure of AI conducted by other fields will be leveraged as an attempt to reconstitute meaning into the term. Research into AI has helped develop a variety of definitions of cognitive organizations, some of which are useful for defining intelligence as it relates to the urban environment. But to investigate the alternatives it may be useful to unpack the understanding of AI that we have inherited. For much of our computational history, the Turing Test has served as the standard for evaluating AI. Intended to evaluate whether or not a machine is intelligent, a successful test is determined solely through the perception of a human interacting with a machine. If the human does not realize it is conversing with a machine and not another human, the machine is intelligent. As Benjamin Bratton describes, this anthropocentric framework poses significant issues. If artificial intelligence is presented to sound and appear human, as in the case with Siri and Cortana, they are essentially presented in drag, dressed to be something they are not.[25] Bratton's provocation is effective, but the underlying message should not be missed– if we are only developing and looking for intelligence that appears like us or sounds like our own, we may ignore other forms of cognition. As Stuart Russell & Peter Norvig[26] describe in Artificial Intelligence: A Modern Approach, the quest for artificial flight succeeded once engineers stopped imitating birds and began using wind tunnels.

Figure 1 Chicago's Array of Things, photography courtesy of The Urban Center for Computation and Data.

Part of the difficulty of determining an appropriate definition of artificial intelligence is that intelligence may be presented in far more forms than we might have previously thought. The work of cognitive scientist Edwin Hutchins[27] suggests that intelligence may span multiple entities, demonstrating that complex systems such as commercial airplane cockpits make collective decisions between multiple cognitive beings and processes. As the different entities within these organizations communicate, they effectively behave as a singular cognitive system. This phenomenon, described by Hutchins as distributed cognition, demonstrates that intelligence may not be limited to a single set of neurons, synapses, or circuitry. As Hutchins suggests, an intelligence distributed between neural nets essentially becomes a neural net itself. Ant colonies may be considered a singular brain, with each ant acting as a neuron– a suggestion made by Douglas Hofstadter[28] in 1980 and supported by cognitive scientists since. Studies of slime molds have suggested that neurons may not even be necessary for biological cognition.[29] If neurons are not necessary for cognition and collections of neurons can act collectively, then cognition is more broadly distributed than previously suggested– even within our own bodies. Our "second brain", the neurons that line our intestinal tract, is a relatively recent discovery, first described in 1921 and only gaining acceptance by the medical community in the 1980s.[30] If neurons or even processors may not be necessary for cognitive processes, and if distributed cognition suggests that intelligence may span multiple organisms or entities, then not only could intelligence be different that our own, but it can be structured differently than our own. The networked, sensor–laded urban infrastructure of the Smart City that funnels into a central "brain" may be one deployment of urban scale AI, but not the only possibility. By acknowledging our anthropocentric biases towards the organization of cognition, essentially a collective blind spot, we are open to find a wider range of strategies– many of which may be more appropriate and productive in regards to the global issues that can be addressed through urbanization. When juxtaposed against the computational struggles the Smart City strategy is facing– the lack of computational power, the desire to create a central image of the context, as well as the enormous expense– a distributed cognitive strategy has merit. Distributing computation throughout a context allows cities to incrementally invest, rather than create sensor–based mega structures at once. A distributed strategy empowers the mechanic intelligence to respond locally, while understanding key aspects to the larger system. By creating an array of small processors, the bottlenecks of data at a central server are avoided. Distributed cognition is scalable, iterative, and may be tested and evaluated as developed (Figure 2).

Figure 2 Photograph of Masdar City by Gökçe Günel.

Citation: Webb A. Distributed cognition: assessing the structure of urban scale artificial intelligence. *Int Rob Auto J.* 2017;2(5):187–194.
DOI: 10.15406/iratj.2017.02.00034

The organization of the smart city

Masdar, a city in the United Arab Emirates, is a prototypical sensor–based Smart City. The intent of the urban scale AI as deployed by ground–up Smart Cities like Masdar and Songdo, and for Smart City retrofits like Chicago's Array of Things as well, are to use sensors to make every element of the city knowable and actionable.[31] This ambition has come at an enormous expense– though the construction cost has yet to be determined; the government of Abu Dhabi has committed $15billion. Masdar as an urban scale AI privileges the collection of the information, creating sensors with small amounts of response. As one of IBM's chief technology officers said in a New York Times interview on a Smart City deployment, "Smart is all about information. Once you have the information and understand it and know what to do with it, you are halfway to smart".[32]

A massive receptacle for information, Masdar does not contain a light switch or water faucet within the entire metropolis.[33] These amenities sense when the user needs them and automate themselves, facilitating a small degree of mechanic agency. But this automation is essentially a large–scale deployment of the motion sensors that have been available and in use for several decades. As of 2011, the city could not yet know when residents arrived in their buildings and change climate accordingly– a stated goal of the project.[31] If the city can only detect when a resident is in front of a sink or moving in front of a sensor, the nature of the city's intelligence needs to be questioned. The organization of Masdar can be understood as organized along an extremely superficial understanding mammalian intelligence. The sensors act as a nervous system, sending energy use and water use information to a centralized "brain", the Building Management System (BMS).[31] But unlike a human brain, Masdar's BMS has a limited capacity to respond to the information it gathers, with no ability to send information out in the way a mammalian brain sends electric impulses to muscles. Once the city was constructed, and ideal temperature was determined, a compromise of energy efficiency and comfort, and implemented throughout the city.[31] When residents of Masdar complained about their lack of control over their climate, "dummy" thermostats were installed.[31] The residents were not aware that the thermostats did not function, and executives of Masdar City described the benefits of dummy thermostats as a productive manipulation.[31] Masdar is intelligent in terms of data collection, but dumb in terms of both human and mechanic agency. That can calculate an optimum temperature, yet not facilitate direct input from its human users to manipulate it; suggests a significant sensor array but essentially a regression. Human users are quite used to adjusting their own thermo stats; to remove this capacity is reductive. When the collective urban response is this reductive– such as setting one average temperature for an entire city– the gathered information is incredibly rich but its deployment is less effective than the system it is replacing.

Like Masdar, Los Angeles' ATSAC traffic system is also a sensor rich system with a low degree of agency. The traffic control platform is comprised of 21,000 traffic detectors and 3,900 networked traffic lights, but struggles with the sheer volume of Los Angeles' drivers.[34,35] Developed for the 1984 Olympic Games, its flexible approach has adapted to incorporate technologies as they develop, integrating closed circuit television, light rail priority systems,[34] traveler information displays, adaptive traffic control and smart bus priority systems. The combined system produces an immense amount of information, but has limited agency in how to act upon its context. Ultimately, the system can only effect the change in traffic signals, a helpful tool in many situations, but not enough for volume of traffic in the Los Angeles basin which has been reported as the most congested region in the United States.[36] It is not enough to simply provide an immense amount of data, essentially a digital version of sensory input, to the computational brain. Simply adding more video feeds, more traffic data does not necessarily create a more effective urban scale AI, there is simply more data at its disposal. But if the data itself is not a measure of intelligence, there needs to be a computational response to the data. For this to happen, the AI needs to know, on a base level, what the data is. As robotics Steve Grand describes, "Brains don't talk in ASCII code—they talk in muscle tensions and retinal signals. Our senses have co–evolved with our brains, so that each depends upon the other. Trying to understand or even replicate the functions of the brain using video cameras and electric motors is futile".[37]

The sensor–based Smart City, as they are currently deployed, inherits its organization from the Internet of Things (IoT) more than any discourse on the organization of intelligence.[38] Here, IoT is understood as the array of objects that connect to the internet and each other.[39,40] If the baseline understanding IoT is simply as things that communicate data via a digital network, then it is easy to understand how Smart Cities are gauged by how many sensors connect to a central node. Because of their similarities and overlaps, the relationship between IoT and the Smart City is irreducible. As Benjamin Bratton[25] observed, "The Smart City is inside the Internet of Things just as the Internet of Things is inside the Smart City" The similarity between the Smart city and IoT is not merely associative, but also structural– the Institute for Electrical and Electronics Engineers (IEEE) have set the standards for connectivity and communication for internet–connected objects.[40] Described by the IEEE as the "connective tissue" for IoT, these standards have also served as the connectivity infrastructure for smart city grids and smart buildings.[40] The standards that have created a fluid interconnection between objects, but have not necessarily created a standard for the quality of those communications. While many objects that are considered part of IoT have a greater capacity than simply a network connection, the shared characteristic of these objects is an ability to send information via the internet. The current result of IoT is an array of incredibly feeble, networked objects. The laughable helplessness of internet–enabled trash cans and smart diapers is celebrated by the Twitter feed Internet of Shit (@ internet of shit), which posts video clips of stuck Roombas, malfunctioning Wifi–enabled pressure cookers, hacked teddy bears and digital assistants are tricked into repeating themselves endlessly.[41] A similar parody was produced by David Jimison & Jooyoun Paek[42] their project Too Smart City consisted of a trash can that spat back refuse, a sign that constantly changes its message, and a bench that ejects users. The objects of Too Smart City underscore the feebleness of networked objects by endowing them with both agencies they are lacking as well as a capacity to revolt.

The empty connection of things to the internet merely allows objects to broadcast their structural helplessness. A wifi connected trash can only broadcast that it is full, a status update conveying information that would otherwise be easily observed by the user. The implication is that the human user would more likely notice a push notification on their smart phone than their trash can overflowing with garbage, which may be the case but remains an empty gesture of intelligence. Likewise, the assumption with the smart diaper is that a parent would be more aware of their phones ring tone than their crying baby. These connections are not necessarily without value, yet they do not necessarily contain value either. Network connection without intrinsic value serves as the conceptual approach for sensor–based smart cities (Figure 3). The importance of a distributed strategy is not only the distribution, but the capacity for each cognitive element to have a degree of computational or physical agency. With the cockpit example, a form of cognition is enabled through the specific activities of the pilots. Hutchins[43] describes how calculating the reference speed for landing is a cognitive dialogue between the crew. Due to the potentially catastrophic consequences if the landing speed is miscalculated, the pilot and the pilot not flying (PNF) collectively compute the calculation through repeating readings, stating representations of the context and the plane's settings, and verbally describing computed results.[43] Hutchins describes this dynamic as the cockpit's "memory", an interpretation of material symbols or real time data concerning the Hutchins.[43] Similarly, the ant colony does not work as a hive mind if the ants themselves cannot make decisions or navigate the landscape themselves. Biologists have found that the formal characteristics of ant structures are driven by the contextual decisions made by individuals, rather than the communication of a template or model.[44] IoT's communication infrastructure is a productive start, but adding individual agency is a critical component of distributed cognition. When compared to the bottlenecked servers of the Smart City, a distributed system inherently has the capacity to be more responsive and adaptive. By endowing the different computational elements with agency to instruct or act, the system evolves past the helplessness of IoT. These are the characteristics of an urban scale AI necessary to address global issues of ecology. The city's "memory" is the shared abstraction of the information concerning a task that is exchanged between entities. Mechanic agents require a capacity to react in real time to their contexts, funneling urban information to a central node creates computational gridlock and harvesting information is ineffective without a capacity to respond. Urban scale AI doesn't need to sense as much as it needs to behave (Figure 4).

Figure 3 Photography by Matthew Lutz at Princeton University and Chris Reid at the University of Sydney.

Distributed agency

If commercial airplane cockpits can be seen as singular cognitive entities, the same principle can apply to an urban intelligence. Networked transit provider Uber uses their Heat map technology to describe the density potential riders across a city.[45] The Heat maps provide this information for drivers, and allow them to decide whether or not to approach a "surge" of riders. The presenting of information allows the human drivers to deploy their own heuristics as to engage a surge or not, leveraging their own sensibility about the city through a guideline and not a directive. In this sense, each driver is their own processor, distributed through the city but computationally connected through Uber's platform. If we accept that the driver of an Uber vehicle as a human processor of information, then Uber's strategy begins to suggest a distributed cognitive system. The driver is presented with the heat map information, but still retains the agency to engage with the heat map or not. The prompts through Uber's app present the driver the opportunity to accept a fare or not. Here the app and the heat map are sensing information, but the vehicle's driver is essentially computing whether or not to engage with the information. Though Uber has many policies to encourage certain behaviors, the human processor is responding to the information but not taking specific direction from a central server. This distribution of cognition, or computational agency, empowers a dynamic, reflexive infrastructure.

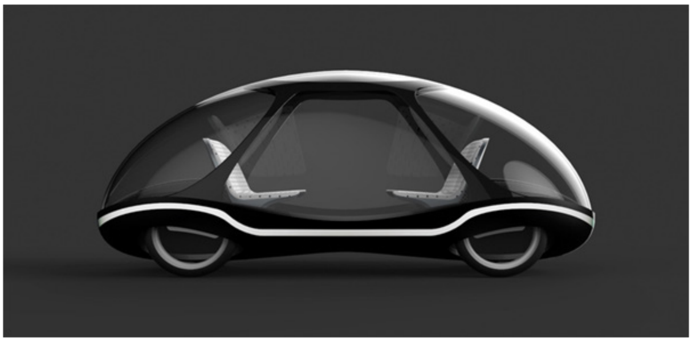

Figure 4 Experimental Electric Driverless Car by Josh Flowers.

The city's AI gains even more agency with more distributed parts communicating directly. This potential is demonstrated with the pending development of driverless cars, projects developed by Google and Tesla, among others, to replace the driver of the horseless carriage. At the time of writing, driverless cars are celebrated for their potential urban contributions, as many suggest that driverless cars could alleviate the need for parking in downtown areas. As driverless cars can continue to drive themselves while their users don't need them, their need for parking is greatly reduced– which would be extremely helpful. A team of transportation engineers at the University of Texas Austin has suggested that each driverless car deployed could displace up to eight parking spots in Austin.[46] At the time of writing the United States has close to a billion parking spots with only 253million passenger cars and light trucks.[47] Reducing the expectation for drivers to park anywhere anytime would dramatically reduce paved surfaces and lead to widespread redevelopment– particularly in dense urban areas. Shared car systems could reduce resource use as well. While the effect of this projected parking reduction would have on urban space is significant, other benefits are presented when viewed as distributed cognition. Vehicles with driverless capacities currently can self–park and prevent drifting between lanes, but their capacity to compute traffic information could enable an incredibly high level of performance with energy use– driverless cars would not just look at where traffic is while choosing a route, they also communicate with other cars and understand how many cars were heading in a certain direction and make route adjustments accordingly. This would leverage mechanic intelligence's ability to receive significant traffic information and compute the best route, surpassing Uber's current model of human driver and heat map.

Citation: Webb A. Distributed cognition: assessing the structure of urban scale artificial intelligence. *Int Rob Auto J.* 2017;2(5):187–194. DOI: 10.15406/iratj.2017.02.00034

Computer Scientists Peter Stone & Kurt Dresner[48] created software that would eliminate traffic lights altogether, where robotic cars would act as a multi–agent system that would collectively compute their trajectories and weave around each other without. As we have little problem avoiding other people as we walk through a crowd, a driverless car essentially has the potential to compute similar variables at a higher speed and accommodate the higher velocity of a moving vehicle. If a stop light is a mechanism to prevent cars from colliding at intersections, a driverless car network achieves the same goal– but without extra resource use of losing and regaining momentum, idling and the experience of waiting. Stone & Dresner[48] suggest that driverless cars could potentially track pedestrians as well, therefore maintaining an intersection's capacity for multiple modes of transportation. When seen through the lens of distributed cognition, driverless cars suggest a similar strategy to an Uber fleet, but with a capacity to respond to significantly more information. While an Uber driver receives information through the Uber app, the Uber heat map, and their own understanding of the city's context; the driverless car has to operate similarly but respond to the trajectories of other vehicles, pedestrians, and bicyclists. Driverless cars have the potential to not only respond to traffic and adjust route, but anticipate where traffic congestion might occur and change route before it happens. It should not be surprising that Uber has demonstrated interest in these aspects of driverless cars. In 2015, Uber created its own research laboratory in Pittsburgh to develop its own self–driving cars through a $300 million partnership with Volvo.[47] Prototypes of the driverless cabs launched in Pittsburgh in September 2016.

From a sustainability perspective, autonomous vehicles might be necessary. Recent reports have found that TNCs like Lyft and Uber are making gridlock in New York City worse, not better as initially promised.[49] The reduced parking and reduced material need for individual car ownership are both environmental advantages, but the capacity for an autonomous vehicle to receive more traffic information could potentially present resource efficiencies and emissions reductions. Robotic vehicles could make decisions not only based on destination and traffic information, but more conduct more complex calculations using more subtle data like gauging the energy used to negotiate topography and wind direction, as well as fuel lost while idling. Another successful example of an urban scale AI with a distributed cognition is The Nest. The learning thermostat can be seen as the adaptive, distributed alternative of Masdar's heating and cooling aspirations. The Nest is a device that deploys machine learning based on heuristics and learned behavior to control a building's interior temperature.[50,51] The Nest observes the user's behaviors, daily life patterns, and preferences and combines this data with sensor–based information on local weather to determine an appropriate temperature.[52] More recent versions of The Nest can determine if their occupants are within the building through the Home/Away Assist feature, and mitigate the temperature accordingly– something Masdar in all of its investment has failed to accomplish. Rather than determining an optimal temperature for an entire city and conditioning spaces at this temperature when occupants are present, The Nest can adjust based off of local conditions and occupancy, saving energy while still allowing residents to control their own environment.

Though the resource use benefits of The Nest have yet to be proven, one study suggests that there is not a clear energy savings by using The Nest instead of a standard thermostat; the technology has significant implications for Smart Cities when viewed as distributed cognitive AI[50,53] an array of Nests distributes computation across the city through localized processing. The computational power could potentially rival that of a Smart City, but is much more achievable as it is dispersed among multiple processors. By computationally controlling temperature through the installation of hardware, the strategy allows software to be updated to eventually produce a resource reduction. The computations of The Nest drive discrete, localized changes, where the brute–force strategy of Masdar is computationally challenged to produce a similar result. Though The Nest is one of the products most associated with IoT, it differs greatly in terms of its agency. The assumption of many IoT devices is that humans are the actuators, and that humans wish to be engaged in the level of action that is required. The internet–enabled trash can is an extreme example, its agency is limited to the capacity to broadcast its status to a presumably oblivious human to act as actuator, and remove the garbage. Garbage removal requires a significant amount of physical agency, and would require substantial infrastructure to enable. The Nest, by comparison, simply needs to communicate instructions to the mechanical systems that already exist within the building. Its agency is enabled through the control of existing machines, creating a low–cost effect with high impact. The distributed model of the Nest and driverless cars alleviates the necessity for a massive, central computational structure and homogenous responses. This mode addresses Masdar's failures by creating a locally responsive system that may leverage AI but also receive input from human users. The Nest can use predictive algorithms and location technology to minimize energy use, driverless cars can compute safe trajectories to eliminate the energy loss at stop lights. Rather than use the brain and nervous system as an organizational metaphor, as cities such as Masdar do, these systems look past this anthropocentric bias and embrace a distributed model.

The practicality of The Nest is enabled through its ability to communicate with its users and with other machines, and then compute how to proceed. Thermostats have traditionally been used to communicate with mechanical units, thus The Nest's communication system was in place. The Nest's ability to gather local weather information, as well as observe its user's behaviors, makes its actions more effective. But its ability to communicate with smart phones and, potentially, driverless cars allows The Nest to have a robust agency, predicting users' arrival and instructing systems to adjust. As The Nest is predictive– learning its users' preferences in response to climate, time of day, etc. it has the potential, once programmed, to have a seamless interaction with its inhabitants where the users forget it is even there. A municipal retrofit of installing a Nest in each home is a far more performative strategy than Madsar's. If Masdar costs only the $15 billion provided by Abu Dhabi, then it has spent approximately $300,000 per home for each of its 50,000 occupants to potentially condition its climate in response to occupancy– assuming that each occupant has their own space and that the budget does not surpass what Abu Dhabi pledged, both of which are presumptuous assumptions. The Nest already has that capacity that at a suggested retail price of $299– approximately one thousandth of the per home cost.

Conclusion

While distributed cognitive strategies like Uber and The Nest are still maturing, their capacity to incorporate deep learning strategies suggests a more efficient computational structure that engenders adaptive responses. The strategy of distributing computation across the entire city, with different nodes or "neurons" communicating directly with each other suggests a capacity to fold in different neural sets into the computation. By spreading the cognition across the city, individual processors can evaluate various transportation, energy, and spatial conditions across the city and respond accordingly. If The Nest can currently communicate with cell phones, then perhaps it could eventually incorporate travel information from either cell phones or driverless cars to detect when its occupant may be returning and change the climate to accommodate. An anthropomorphic framing of artificial intelligence as brain and nervous system has hidden its potential when deployed on an urban scale. A view of the city as a central, singular node; a massive sponge absorbing all of its inhabitant's data, subverts more appropriate definitions such as Hutchin[43] understanding of collection of neural nets acting as neural nets. Rather than celebrate how extensive the city's sensor array, or nervous system is, we should focus on the city's capacity to act and the intelligence of its actions. Structuring the AI to leverage a large nervous system will assist this goal, but should not the criteria we should use to evaluate it.

By gauging urban scale AI's intelligence on how well it performs, the discussion is reframed from the amount information it absorbs to what it does with the information it has. The agency of the system has a far more direct effect upon its impact; projects with far less infrastructure than Songdo and Masdar have proven to be more effective in terms of managing resources and engaging spatially with humans. The synthesizing of the collective parts in ways that allows them to communicate with each other, and has the capacity to produce a city that is more appropriately smart. There is significant precedent for moving past the centralized model of intelligence. As Dario Floreano & Claudio Mattiussi[53] describe in Bio–Inspired Intelligence: Theories, Methods and Technologies, research into artificial intelligence after the turn of the millennium has expanded past a focus on human brains to encompass a broader range of organisms, processes and phenomena that can occur at multiple spatial and temporal scales. Floreano & Mattiussi[53] suggest this shift is reflected in the rise of distributed computation, the fragmentation of the personal computer into a multitude of personal assistants, communication devices, and internet agents that are required to interact and respond to us. The shift from a desktop computer being the one synthesizer of information has branched into a much more distributed model, where a multitude of devices may operate within a network, communicating with not only multiple devices but multiple human agents as well. Floreano & Mattiusi[53] describe this model as "new artificial intelligence." As urban scale AI adapts to this new model of intelligence, it is critical that these distributed processors have a localized agency– the self–driving car can eliminate the energy loss of stop lights by calculating a trajectory that avoids other vehicles, but the calculation is valueless without the ability to change speed and direction. If The Nest can calculate the optimal temperature, but not instruct the mechanical systems to adapt, the calculation is equally useless. The criteria of response are a critical component in the shift to a distributed AI. Floreano & Mattiussi[53] suggest a connection with new artificial intelligence and behaviorism, a movement that emphasized the study of actions over the introspection of mental processes as a demonstration of intelligence. By reframing urban computation into a network of interconnected sensors, as well as processors and actuators, urban scale AI will reflect criteria for intelligence as suggested by cognitive and computer scientists.

Acknowledgement

The author would like to thank Neil Leach, Benjamin Bratton, Brian Goldstein, John Quale, Mark Childs, Aaron Ketner, Nick Di Donato, Adrian Pinnecoose, Helena Logan, Brandon Pachuca, Robert Peebles, Niko Maestas, Shelby Powell, Ali Al Gahmi, Patrick Aspaas, Brian Buenrostro, Cole Cottrell, Lam Nguyen, Jaime Frias, Joshwa Hernandez, Rafael Milla, Luis Miguel Uribe Torrado, Gökçe Günel, and Josh Flowers.

Conflict of interest

Author declares that there is none of the conflicts.

References

1. Nicholas de Monchaux. Spacesuit: *Fashioning Apollo*. 2011.
2. Nicholas de Monchaux's. Local Code: 3,659 Proposals About Data, Design & the Nature of Cities. USA: Princeton Architectural Press; 2011.
3. Bratton BH. On AI and Cities: Platform Design, Algorithmic Perception, and Urban Geopolitics. *Lecture presented at the Het Nieuwe Instituut*. 2015.
4. Greenfield A. Against the Smart City. 2013.
5. Hajer M, Dassen T. Smart about cities: Visualizing the Challenge for 21st Century Urbanism. *Nai010 PBL Netherlands Environmental Assessment Agency*. 2014.
6. Picon A. Smart Cities: A Specialized Intelligence–AD Primer. USA: Wiley; 2015.
7. Shepard M. Toward the Sentient City. Ubiquitous Computing, Architecture, and the Future of Urban Space. Shepard M, editor. USA: The MIT press; 2011.
8. Townsend AM. Smart Cities: Big data, Civic Hackers, and the Quest for a New Utopia. USA: WW Norton & Company; 2013.
9. Waal M. *The City as Interface: How New Media Are Changing the City*. Netherlands: Nai010 publishers; 2014.
10. Bratton BH. *The Stack: On Software and Sovereignty.* USA: MIT Press; 2016.
11. Mitchell W. *City of Bits: Space, Place, and the Infobahn.* USA: MIT Press; 1995.
12. Mitchell WJ. Me++: *The Cyborg Self and the Networked City*. USA: MIT Press; 2004.
13. Tuters M, Varnelis K. *Beyond Locative Media: Giving Shape to the Internet of Things.* USA: MIT Press; 2006.
14. Li Y, Lin Y, Geertman S. The Development of Smart Cities in China. In Proceedings of the 14th International Conference on Computers in Urban Planning and Urban Management. 2015. p. 1–20.
15. Anthony Townsend's. *Smart Cities: Big Data, Civic Hackers and the Quest for a New Utopia describe civic hacking as a response to the Smart City*. USA: WW Norton Company; 2013.
16. Picon A. Smart Cities: A Specialized Intelligence–AD Primer. USA: John Wiley & Sons; 2015. p. 11–27.
17. Picon A. Smart Cities: A Specialized Intelligence–AD Primer. USA: John Wiley & Sons; 2015. p. 27–30.
18. Picon A. Smart Cities: A Specialized Intelligence–AD Primer. USA: John Wiley & Sons; 2015. p. 145–146.
19. Greenfield A. Against the Smart City. USA: Do Projects; 2013. p. 571–875.
20. Greenfield A. Practices of the Minimum Viable Utopia. *Architectural Design*. 2017;87(1):16–25.
21. Greenfield A. Practices of the Minimum Viable Utopia. *Architectural Design*. 2017;87(1):18–19.

22. Greenfield A. Everywhere: The Dawning of Ubiquitous Computing. USA: CA New Riders; 2006.
23. De Waal M. *The Urban Culture of Sentient Cities: From an Internet of Things to a Public Sphere of Things*. Shepard, editor. USA: MIT Press; 2011. p. 190–191.
24. Bogost I. Artificial Intelligence' Has Become Meaningless. *The Atlantic*. 2017.
25. Bratton BH. Outing Artificial Intelligence: Reckoning with Turing Tests. Pasquinelli, editor. *Augmented Intelligence Traumas*. 2017. p. 69–80.
26. Russell S J, Norvig P, Canny JF, et al. Artificial Intelligence: A Modern Approach. USA: Prentice Hall; 2003. p. 1–1153.
27. Hutchins E. How a Cockpit Remembers its Speed. *Cognitive Science*. 1995;19(3):265–288.
28. Hofstadter DR. Godel, Escher, Bach: An Eternal Braid. New Society. 1980. p. 315–342.
29. Reid CR, MacDonald H, Mann RP, et al. Decision–Making Without a Brain: How an Amoeboid Organism Solves the Two–Armed Bandit. *Journal of the Royal Society Interface*. 2016;13:119.
30. Brown H. The Other Brain Also Deals With Many Woes. 2005.
31. Greenfield A. Against the Smart City. USA: Do Projects; 2013. 126 p.
32. Singer N. *Mission Control, Built for Cities*. USA: The New York Times; 2012.
33. Günel G. Masdar City's Hidden Brain. *ARPA Journal*. 2014.
34. Fisher JE. ATSAC: 25 Years Later. *ITE Journal*. 2011;81(7):48–53.
35. Sorensen P, Wachs M, Min EY, et al. Moving Los Angeles: Short–term Policy Options for Improving Transportation. *Rand Corporation*. 2008. p. 1–640.
36. *New TomTom Data Reveals Rush hour traffic doubles Journey Times for Commuters*. 2015. p. 1.
37. Johnston J. The Allure of Machinic Life: Cybernetics, Artificial Life, and the New AI. USA: MIT Press; 2008.
38. Miller M. The Internet of Things: How Smart TVs, Smart Cars, Smart Homes, and Smart Cities are changing the World. *Pearson Education*. 2015. p. 2–3.
39. Greengard S. *The Internet of Things.*USA: MIT Press; 2015. 15 p.
40. Greengard S. *The Internet of Things.* USA: MIT Press; 2015. 117 p.
41. Internet of Shit. *The Year is 2020: Our Voice Assistants Only Talk Amongst Themselves*. 2016.
42. Jimison D, Paek J. *An Intentional Failure for the Near–Future: Too Smart City*. Sentient City: Ubiquitous Computing, Architecture, and the Future of Urban Space. Shepard M, editor. USA: MIT Press; 2011. p. 111–115.
43. Hutchins EL. How a Cockpit Remembers its Speed. *Cognitive Science*. 1995;19(3):286.
44. Reid CR, Lutz MJ, Powell S, et al. Army Ants Dynamically Adjust Living Bridges in Response to a Cost–Benefit Trade–Off. *Proceedings of the National Academy of Sciences*. 2015;112(49):15117.
45. Rosenblat A, Stark L. Algorithmic Labor and Information Asymmetries: A Case Study of Uber's Drivers. *International Journal of Communication*. 2016;10:3758–3784.
46. Fagnant DJ, Kockelman KM, Bansal P. Operations of Shared Autonomous Vehicle Fleet for Austin, Texas, Market. *Transportation Research Record: Journal of the Transportation Research Board*. 2015;2536:103.
47. Thompson C. The Worst Thing About Driving is About to Change. 2016.
48. Dresner K, Stone P. A Multiagent Approach to Autonomous Intersection Management. *Journal of Artificial Intelligence Research*. 2008;31:591–656.
49. Schaller B. *App–Based Ride Services: Present and Future, Schaller Consulting*. 2017.
50. Yang R, Newman MW. Learning From a Learning Thermostat: Lessons for Intelligent Systems for the Home. *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 2013. p. 93–102.
51. Hernandez G, Arias O, Buentello, D, et al. Smart nest thermostat: A smart spy in your home. *Black Hat*. 2014. p. 1–8.
52. Edson J. *Design Like Apple*. Hachette Audio. 2012. 172 p.
53. Floreano D, Mattiussi C. *Bio–Inspired Artificial Intelligence: Theories, Methods, and Technologies*. USA: MIT Press; 2008.

Citation: Webb A. Distributed cognition: assessing the structure of urban scale artificial intelligence. *Int Rob Auto J.* 2017;2(5):187–194.
DOI: 10.15406/iratj.2017.02.00034