



# The emergence and evolution of urban AI

Michael Batty<sup>1</sup>

Received: 10 April 2022 / Accepted: 15 June 2022 / Published online: 5 July 2022  
© The Author(s), under exclusive licence to Springer-Verlag London Ltd., part of Springer Nature 2022

Artificial intelligence (AI) emerged alongside the development of the digital computer more than 80 years ago during the second world war. The essential logic for representing a variety of phenomena using the binary code enabled a range of numerical problems, hitherto beyond solution, to be addressed. New technologies based on electrical switching that rapidly grew out of analogies with mechanical switching formed the essence of these computers but behind this technology lay philosophic foundations that paved the way for speculations that a new form of AI might be possible. Alan Turing, one of the architects of this philosophy, along with John von Neumann, argued that our ability to represent many phenomena using binary codes, suggested that the digital computer was a ‘universal machine’ and in this sense held out “... hope that machines will eventually compete with men in all purely intellectual fields” (Turing 1950).

The term AI was coined at a conference at Dartmouth, Connecticut in 1956 where those who led the field for the next 25 years developed a vibrant and optimistic research program into how computers could be programmed to simulate the way we as human beings ‘think’ (Dyson 2012). This was the original goal—to assume that as the computer was a universal machine, it could be used to simulate the complexity of the human brain, and in this way, come to simulate the kinds of problem-solving, design and perhaps creative tasks that we consider unique to ourselves as human beings. This came to be called ‘strong AI’ and it dominated the quest to build AI until the mid-1970s when its problems began to emerge. The notion that one could simulate a general intelligence came under severe scrutiny and the claims for what it might do if it ever emerged seemed far-fetched. But as the dominant field went into its ‘nuclear winter’, elements of it particularly those associated with methods of searching for a pattern in data underlying basic robotics, vision, and

automated manufacturing, began to take over. This led to the kind of AI that now dominates the contemporary scene—‘weak AI’—where the focus is no longer on finding the logic or the intuitions that we as human beings use in our problem-solving but on searching for coherent patterns in data that provide it with structure, thus offering the possibility that such order might be useful for making certain forms of limited prediction.

All the papers in this special issue deal with applications of weak AI although there are still remnants of strong AI thinking in some of the assumptions made by those using these technologies. There is the tacit assumption that some of the patterns and some of the rules used to implement weak forms of AI do reflect patterns of human behaviour but in most contemporary applications of AI, the link to human behaviour is not widely tested, nor is it central to the development of AI. In the earliest days of AI, as it was conceived of as being useful to urban planning, the notion was that powerful optimisation tools could be developed and synthesised with our own intuitions and that such systems would become central to the creation of much more effective plans than any that had been developed hitherto. In fact, because of the limits on strong AI, these efforts were quickly abandoned with the exception of a focus on expert systems in the 1980s. The shift has been to using AI to examine patterns of spatial behaviours in cities rather than new methods for developing better cities. In this sense, the existing terrain of urban AI is relatively routine and somewhat low-key but nevertheless extensive in that it operates across many different areas relating to cities and their planning.

Urban AI penetrates many dimensions of the city. As the methods of AI tend to be independent of spatial and temporal scale, they infuse many areas of the city and it is difficult to classify them into types. As their features with respect to questions of equity, transparency, and efficiency are very wide-ranging, their applications are equally extensive and the papers collected together in this special issue tend to be more of a sample from many different areas rather than an attempt to chart and bound the field of urban AI. In this sense, they represent different applications from a wide

✉ Michael Batty  
m.batty@ucl.ac.uk

<sup>1</sup> Centre for Advanced Spatial Analysis (CASA),  
University College London, 90 Tottenham Court Road,  
London W1T 4TJ, UK

landscape of possibilities and as a collective, they provide a snapshot of possibilities. They exclude many applications of urban AI that focus on infrastructure, flows of energy and people, locational differences and more aggregative aspects of the city but the emphasis here is very much on their impact on social and ethical issues that are to the forefront in figuring out the impacts of urban AI on societal questions. Some of the bigger applications of AI relate to transit and housing whereas platforms such as Airbnb and Uber amongst many others have much wider social applications. These are difficult to unravel and they represent some of the uncharted questions relating to the long-term impact of AI on the form and functioning of the city. Most of the papers here deal with much more localised behaviour in urban areas.

Before I draw out common themes from this collection, it is worth noting the dominant feature of weak AI that feature in these applications. At the outset, the term AI is used here in its most catholic sense and it pertains as much to ways in which cities are becoming ‘smart’ using new digital infrastructures that are being embedded into the urban fabric as well as social media. First, data for most AI problems is usually generated in real time, and in this sense, it is ‘big’ meaning that it is voluminous in terms of size. It is continual in its generation and thus at any point in time although the data is finite, its ultimate size is unknown as long as the sensors that generate it are switched on. Second, the tools used to extract patterns are in general simple models that attempt to explain their structure and these models are continually massaged until they provide a good fit in terms of their predictions and the data. In short, this is the process of machine learning that defines the patterns that such systems produce. Third and this is most relevant to the papers here, the methods of AI are based on simple rules that tend to explain the meaning of the applications in question, the way they are applied and the impact on those who use them and those who are impacted by them. To an extent, the methods of AI in the papers collected here tend to merge into more traditional methods which can be characterised as formal and systematic digital issues rather than fully-fledged applications of learning and optimisation which tend to dominate more traditional applications.

The first paper “Tensions in Transparent Urban AI: Designing A Smart Electric Vehicle Charge Point” written by a team from TU Delft, illustrates perfectly the wider impact of AI. It deals with the user interface reflecting how users react to a system for the charging of electric vehicles, with a view to designing this system to allay the fears of users in how to use these in the most effective way. In short, this paper is about the online charging of electric vehicles and this in itself has several AI components but it also explores ways in which users of the interface react to it in different forms. Range anxiety with electric vehicles has a

major effect on how users might use a system to charge their cars and thus the human–computer interface is all important. In this sense, the impact of how users react to systems that involve AI in various ways is a key focus of this area and it appears time and again in the impact of various AI systems which form the focus of these articles.

The second paper from a group at Penn State University examines the “Street Surface Condition of Wealthy and Poor Neighborhoods” in Los Angeles and is more traditional in its use of AI. Street systems which cut across poor and wealthy neighborhoods in Los Angeles are measured using Street View-like technologies that extract the key components of streets in terms of their physical condition. These are then related to income and house prices and the analysis reveals that streets running through poor neighborhoods are no less desirable than those running through rich. Indeed there is even the implication that the opposite to what you might expect occurs in this application. In fact, the paper illustrates some of the major limitations of AI. The fact that there is massive statistical manipulation involved in exacting patterns from large data sets does not mean that the patterns revealed make any sense in terms of our intuitions and in this sense, the application shows that one must be wary of any conclusions taken from data which is analysed using the ideas of deep learning. In short, what is learned may well be quite counter to what one’s intuitions suggest and the fact that the patterns extracted are somewhat of a black box, mean that it is hard if not impossible to draw robust and useful conclusions from such applications.

The third paper deals with a system of design that brings participants or potential members of a housing complex based on superblocks together to share outside space in the most effective way; the projects in question are illustrated with superblocks in Tampere, Finland. This is a project involving qualitative principles of urban design but in analogy to AI principles, and as well as illustrating how such problems can be informed by AI, the outcomes are also key to developing more effective and workable designs. In this sense, the project is in the same frame as that involving electric charging in the first paper. The fourth paper which is about “Emotional AI and Crime” takes the argument into the key area of how good or bad are AI techniques that are designed for facial and related recognition. There are countless stories about how automated procedures confuse and confound such recognition and how many biases are introduced through the rather unintelligent programming of what are supposed to be intelligent systems. AI is hard as the pioneers who we noted above found out to their cost and one of the key issues is that much of what is out there in terms of systems generating big data in real time, treat these as black boxes whose models of the pattern are not easily interpretable. As we know, models that produce good predictions are only contingent on the data used for their training and

at any point in time can generate outcomes that are quite unpredictable, hence wrong. When these pertain to human systems, the consequences can be catastrophic. In fact, this particular paper takes a reasonably balanced view accepting that such systems will always continue to be developed, and it charts the kinds of contextual issues that are central to their deployment. To an extent, the paper cuts to the core of urban AI in practice, particularly in contexts where the focus is on individuals rather than aggregative patterns.

The papers introduced so far all have implications that AI technologies underpin the existence of a surveillant society where remote information technologies provide a means of monitoring every kind of object with each object not necessarily aware of this potential intrusion. Where this involves human beings, then this can be ethically and morally reprehensible and this focuses the debate on the whole question of privacy and confidentiality. In the fifth paper, Sherman introduces these issues around the traditional equivalence of this kind of observation in prison-type environments where the example is Jeremy Bentham's late eighteenth century *Panopticon*, an architectural arrangement where one can watch the many from a single vantage point. Sherman then introduces the notion of the *Polyopticon* where 'the many watch the many' from 'multiple vantage points' which is the consequence of technology, that is networks, accessible visually which enables anyone to communicate with anyone else: in short, a totally connected society. To an extent, it is already here but this paper points out many features of surveillance which are slowly, perhaps even rapidly, creeping into contemporary societies that relate to the embedding of many kinds of AI into the urban environment. These issues are woven throughout artificial general intelligence as well as many other features of AI that are impossible to control and which are intrinsic to AI. These are central to this collection of papers and even to the rationale for the journal itself.

The articles that follow broaden the range of AI technologies that define the city. We have already noted that a good deal of traditional AI deals with algorithms or procedures that combine diverse data and analytical structures in such a way that models containing millions of parameters are used to produce as good a fit as possible to systems that enable 'good' predictions to be made. It is worth knowing the actual meaning of the parameters that are generated despite the fact that this is often impossible and in this sense, these parameters define how we can extract order but not necessarily explanation from the 'black box'. Tsing in the sixth paper uses 'assemblage theory' to explore these kinds of problem, particularly illustrating these for a problem of citizen participation in Taiwan. To an extent, assemblage thinking has a parallel in the scientific world in complexity theory but in the collection of articles here, Tsing's paper is a wake-up call to illustrate that the world of AI is increasingly full of mystery and that a good deal of AI thinking is comfortable

with such this. To an extent, Tsing's paper questions all of this but it also suggests that assemblages are useful ways of characterising complex systems such as cities.

The remaining papers deal with various views of AI in cities, the first of these (the seventh paper) dealing with "Understanding Citizen Perceptions of AI in the Smart City" which based its analysis on two questionnaires that sample citizen views on the appropriateness of AI; in particular, on the limits of what typical citizens consider AI should have in its effects on how they should be exposed to an AI which they are able to control. The eighth paper focuses on how urban AI should be enriched at the level of the community by adopting the notion of 'thick explanation' first proposed by the American anthropologist Geertz. Essentially this is a critique of AI which often ignores the relevant contextual detail in its many applications and the paper "Urban-Semantic Computer Vision" argues that by adopting a thick explanation, the wider context can be embraced. This introduces the idea that a new form of semantics should be introduced to AI with vision taken as a pointer to the kinds of contextual and semantic focus that are surely needed in most applications. The ninth paper on "Artificial Intelligence in Local Governments: Perceptions of City Managers on Prospects, Constraints and Choices" looks at the problems of adopting an array of Urban AI in urban planning and city government and in some sense, this is a paper about a much wider range of digital techniques than the narrower definition of AI. In fact, AI in the wider context often merges with digital tools more generally and embraces networks, platforms, and all kinds of modelling tools as well as organisational structures that are built to develop the right kind of AI infrastructure in cities.

The last three papers focus on mobility which is one of the major areas of application of AI, largely because the vehicles and networks that move people and materials can themselves be automated while the way in which people and goods are moved can be developed most efficiently using various models which incorporate AI tools and techniques. The tenth paper "The System of Auto-mobility: Computer Vision and Urban Complexity—Reflections on Artificial Intelligence at Urban Scale" poses key questions about how automated systems using AI interact with the traditional city, changing its functions while also feeding back positively on the automated systems themselves. The argument is taken further in the eleventh paper on "Contestations in Urban Mobility" which introduces a range of issues involving rights, risks and responsibilities for the development of such AI. The last paper "Human–Machine Coordination in Mixed Traffic as a Problem of Meaningful Human Control" focuses on the debate on the conflict between different kinds of users of different traffic systems and the intrinsic conflicts between them. These have important implications for their development using AI and related digital techniques

and this paper suggests that a much closer examination of these potential problems of coordination and organisation in terms of AI should be thoroughly researched before a new automated system can be introduced.

The articles in this collection pose many problems that define the field of urban AI. In particular, the ways in which AI introduces order and structure into cities through the literal manipulation and construction of automated physical systems and the various digital twins that enable them to be studied provide an array of possible systems that all fall under the umbrella of urban AI. These papers do not focus at all on the mathematical tools that underpin much of the software that enables urban AI to exist but what they do is raise our awareness of how the wider social context interacts with these new technologies, both exacerbating old problems, generating new, providing opportunities for planning better cities and posing major issues of equity versus efficiency. To an extent, the focus on AI is wider than what might find in any discussion of AI in the narrower technical field for context is all important to see urban AI in context. This context

is fundamental to many of the issues raised in the articles that follow. A series of short reactions to this wider debate complete the collection with three book reviews providing a useful guide to the wider field.

There is much room for thought here. Read on, absorb, critique and enjoy.

## References

- Dyson G (2012) *Turing's cathedral: the origins of the digital universe*. Pantheon Books, New York
- Turing AM (1950) Computing machinery and intelligence. *Mind* 59(236):433–460

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.