Lancaster University Future Places Centre: Technical Reports

The Future Places Centre (the FPC) is researching how people understand the places they live in – their workplaces and homes, their towns and the countryside. Key to such understandings are various forms of 'data' - much of that computer-derived but much of it also produced by cultural practices: in stories about places written down, thought about and endlessly reimagined. 'Data' labels many types of evidence and various ways of seeing that continuously evolves.

The FPC's technical reports and publications present some of its findings on this subject: data and place.

For further information about any particular report or publication, please contact the author or the FPC - fpc@lancaster.ac.uk

This technical report: The Geographies of Intelligence at Home (book chapter preview)

Published: (Forthcoming), Chapter 4, The Shape of Thought: Reasoning in the Age of AI

The use of AI tools and techniques is transforming how data is processed, how it is gathered and how places come to be understood. This paper is a preprint chapter that reports on how AI tools have shaped the way the home setting is understood and made over the past twenty years or so.

Chapter 4: The Geographies of Intelligence at Home

Introduction

In prior chapters, we have been seeing how a certain view of AI encourages ways of looking at the world – at ourselves, at what we do, at the tools that we use to get things done – and how this can shape our thoughts on all sorts of matters. Indeed, we have been seeing how AI algorithms work, or some notion of how they work, can frame how people understand many things well removed from AI. This shaping can take surprising forms: it can, for instance, lead people to alter how they act when doing things with AI. They seek to become, in some cases, more algorithm-like and hence, they hope, more compatible with the technology. What being more algorithm-like is not fixed or certain, of course, more a question of improvisation. But we began to see what it can look like – and some of the confusions that result – when people collaborate or jointly work with AI through a WIMP interface and find themselves navigating feedback loops generated by AI that they don't fully understand. We began to see that the AI might not understand the user either, making matters even more complex. We also took the opportunity to see how one needs to look at how the technology works and how the user uses technology to see what the two, the human and the machine, end up doing together: we looked at chatbots to understand this. In this chapter, I want to look at a similar instance of mutual shaping, this time in the home setting – the home is a geography of a particular kind. It is a place of bricks and mortar, of technology and people. How these interact is complex, and notions about where intelligence may be – in a machine or in a human, say – are often orthogonal questions to what needs focusing on if one wants to see how to make 'space-technology-and-persons' be a marriage that offers more, rather than less. We should be especially careful to avoid muddles and confusions.

My purpose, though, is not to recount another case of how technology and users mutually shape each other in a particular space; it is rather to do the following: first, to recount that shaping but in a fashion that ensures we see its complex weave and dynamics. Second, to help formulate how one might shape our own thoughts about home life as we look to the future and how technology might help home life develop new forms. The domestic sphere provides a case study of how we might shape our thoughts to this task; we need what I will propose is the right sensibility that delivers insights that are beyond the scope of NAI. Homes are special in part

because of the moral cargoes they contain and not just for the patterns of behaviour one finds in them. NAI focuses on those patterns but understanding the values of the cargoes that gives those patterns meaning and purpose is beyond what NAI can currently do or see. Shaping our thoughts should also ensure we see human behaviours for what they are too: sometimes contrary in their manner, undertaken deliberately to obfuscate what might be tractable to machinic perspectives. As we shall see, occupants of homes will often adjust their behaviours in light of what they think the technologies of AI capture. The vendors of technology might struggle with achieving what they hope their technology can do, but users react nonetheless and do so in ways that makes it even harder for technology to succeed. Some of the aspirations behind technology might be exaggerated, even foolhardy, but some of the actions of people can be confounding too. This is not the only way that people behave, of course. Sometimes 'users' in home settings come to rely on certain types of AI technology in ways that alters their relationship with the technology but without these same users even noticing that their relationship has shifted. In some cases, we shall see that it is not the intelligence in the technology that concerns the user, as how the technology allows them to be more intelligent in their own endeavours. They use the technology to focus on themselves at the cost of looking at the technology in the domestic sphere. This very different consequence underlines how careful we need to be when assessing the role or impact of technology. It really isn't a matter or adding or reducing intelligence; nor is it a matter of mapping patterns to needs; sometimes it's about matters that are well removed from technology but which are important to people - such as their notions of identity that make some families different from others. These notions do impact on how technology finds its place in the ecology of the domestic world.

I want to explore these diverse consequences by providing a potted history of AI technologies in the home setting. I start with what was called the 'smart home agenda' which, as I say, emerged about twenty years ago, and report some research I was involved with at that time. The goal, then, and putting it crudely, was to model as much human behaviour in the home as possible, and then to come up with technological ways of automating some of the behaviour modelled. The hope was that, ultimately, the labours of the home dweller would be replaced by the powers of machinery – some mechanical, some virtual (in software), but all governed, one way or another, by 'intelligence' in the systems.

To distinguish this from later technologies and approaches to technologies in the home, I will call this the 'behaviourist' era. I do so as that term labels an approach to the social sciences that emphasizes what people do rather than what they think or say; actions speak louder than words in this view[1]. In some ways this approach succeeded, but in other ways came to highlight, in users' reactions, how the behaviours captured by the technology might be understood. Users came to resist what these representations conveyed, wanting to offer alternative interpretations. These alternatives were sometimes bolstered by users intentionally changing behaviours so as to reflect how they wanted to be seen. Sometimes changes they made had to do with meanings they gave behaviour, meanings they adjusted in light of what they saw as behaviourist misinterpretations. In making these changes, the patterns that the technology was seeking to capture and model kept being slightly out of reach for that same technology.

Following the behaviourist era came technologies that enabled new forms of sociability and digital interaction. Smart phones, for example, allowed users to do more wherever they were and this included the home, and, in a sense, the affordances smart phones offered ran counter to the smart home agenda which had emphasised automation. With smart phones, *users* could do more. Hence, I will call this period the 'agential' one. The technologies of this era had numerous consequences the most important of which was how they drew attention to the way homes are made. By that I mean that, with smart phones allowing interaction with the world outside, questions about what was appropriate inside the home came to the fore. Could the outside come in? If so, which part of outside? Such questions could not be reduced to simple judgements about what was public and what was domestic, as they were also bound up with different trajectories for those inside a home: teenagers used the smart phone and associated applications (social media especially) to strengthen the separations they sought from their parents inside the home while using the same technology to deepen relations with those they had outside the home - their school mates, say. I shall refer to some of my own research of the time to explore some of these changes in domestic life and how these affected the labelling that was appropriate: instead of smart being used to label a house, it became a label for the technology itself – the smart phone. Whether smart is equal to intelligence depends, of course, on how these terms are defined. I shall remark on this.

Adjectives (and sometimes nouns) notwithstanding, diverse forms of agentiality emerged at this time. They have kept evolving, needless to say, but for the purposes of concision I will draw a distinction between the technologies that helped enable this agentiality and more recent

technologies, ones that afford what have come to be called 'natural interaction'[2]. This is where we are today, and it is a term which labels the use of, for instance, the spoken word to command machines in the home setting. Users speak to Alexa and it connects them to music, say; users idle on their couch and ask Siri to suggest what they might cook for dinner. They might muse with ChatGPT too. These technologies require a great deal of data processing and are built on models of millions of user behaviours but they are not based on the movements of people inside homes, as much of the original smart home technology was. These new natural interaction technologies are, in some ways, the descendants of the technologies first explored in smart homes insofar as they seek patterns, but also quite different in what they model and in what they enable: their data comes from mass user behaviours, and from language modelling; quite different from the sensor-based calculations of most smart home tech. Nor do natural interaction technologies provoke the same reaction; some of them become taken for granted and, in so being, allow more shadowy influences on the patterns of home life.

This will conclude my potted sketch of the various technologies one can find in the domestic sphere and which can be variously called AI. Obviously, there is much more one could say about technology and home life than I manage to cover in this single chapter, just as there is more one could say about the different types of 'needs' people have for technology at home. The prior two chapters have allowed me to cover quite a great deal on the topic of technology; this one will allow me to point towards the variation one finds in 'the user'. A teenager and their needs could not be more opposite to those of some one very old, for example. A teenager often needs prompting to do housework, while for people with arthritic joints and tired hearts domestic labour is beyond practicality. For the latter, automation is necessary while for the teenager a luxury they could do better without.

I am not wanting to seek a balance in the attention I give people and machines, however. My task is to survey and focus on these three eras so as to convey the complexity and diverse interaction that operates between people technology and the home setting, and how therefore it is imperative we resist oversimplified accounts of 'AI at home'. Given this, I will try and conclude the chapter with how we might understand home settings in the round, shaping our thoughts to properly recognise what people might be wanting to do with their homes, and how technology is naturally part of the tool set they deploy to make homes what they want them to be. We will need

to ascertain what those behind the technology want too, as this might not coincide around a narrative that home dwellers have as regards questions of automation, smartness or even NAI.

Intelligence Domesticated: The Behaviourist Period

Before the 1990s, the term smart home had a particular meaning: it labelled houses built in more intelligent ways though the buildings themselves had no intelligence. Builders had been seeking more intelligent construction since at least the 1960s, and had sought to devise, for instance, cost effective house frames that would allow lighter walls to be hung and more easily moved as the needs of occupants altered over time[3]. Intelligence was about space and its construction; it was about allowing its reshaping by house owners. Intelligent building allowed intelligent remodelling. But by the year 2000, the term came to label how a house itself would become intelligent – somehow altering what was afforded in the functions of the house. These functions would be made greater by the use of computer technology: endowing what had been inert spaces – rooms, stairs, doorways and so on – into spaces that were monitored, watched, and through this, treated as sources of data that could let technology – basically sensor-based AI technology – alter what was done in those spaces.

At this point, what would be learnt from the data – even how it might be learnt – was still unclear, it being more of a hope and expectation than a proven possibility. But the hope had good credit: this was, after all, a little after the height of the internet bubble. That bubble might have burst but even so much had changed over the prior decade. The internet and associated technologies had altered many things, mostly in the workplace (the 'productivity paradox' notwithstanding[4]), and so it seemed obvious that the next candidate for change would be the home. Sensors and the aggregation of data from them would be the route for this transformation.

I was equally excited about these prospects. The publication of my book *Inside the Smart Home* in 2003 was motivated by research largely funded by a mobile phone operator, Orange Plc (a UK listed company) who imagined that they could participate in this change – mobile networks beginning to move towards data and mobile internet. With this in mind, Orange had built a 'smart home' and hired my research team to capture the experiences that the home provided. Some of the chapters of my book reported on our research in that home. Our own interest was partly anthropological, a question of wanting to explore new kinds of environments and the forms of life they might enable, and partly to do with design, with interaction with computer technologies

whose form and function might be wholly new. We wanted to do some research into human computer interaction (HCI) if you like, but it wasn't certain where or what the new 'smart home' would lead to or what the human-computer interaction in question would deliver. Some of these early forays into smart homes, including the Orange@home, entailed a great deal of watching and monitoring in ways that could appear somewhat sinister, panopticon-like, and hence engaging with these efforts was also a route to explore what were the downsides of AI as much as what were its upsides. We wanted to help shape a better future through HCI and anthropology, aware that these first steps might be faltering.

The Orange@home establishment was a large, old farmhouse with a walled garden at the intersection of two rural roads just north of London; a little remote, by itself, but generous in proportion. It looked like a delightful country home for a well-to-do family. Indeed, that seemed to be what was sought, judging by what the Orange executives said. But the house was not like any other. It had computers of various forms in all the main rooms, various sensors and control systems, a 'server farm' hidden in a cupboard the size of a kitchen, and a team of support engineers maintaining these systems remotely. On close inspection, it was like a stage, or rather, like a theatre with a stage and hidden, out of sight, the technologies that would somehow make the building come alive with intelligence.

As I have hinted, this infrastructure was intended to monitor, aid and support everything that the occupants did with a view to eventually modelling that behaviour so that, after a time, some of that labours undertaken in that behaviour could be taken on by the computer-controlled systems, with the building 'doing things for them'. Once patterns had been learnt, doors would automatically lock at predetermined times, for example; conversely, all would open when the first member of the family was due back – again, determined by frequency patterns. Movement sensors would be used to switch lights on and off as occupants moved around the house and these actions would likewise be used to make models of these movements that would eventually mean that lights would be turned on or off on the basis of prediction. Similarly with the heating systems: occupancy and preferences would be gradually modelled and then automated. In these ways the house would become 'intelligent'.

These goals were, in truth, thought of as low-lying fruit – easy to achieve, and likely to be successful. Automated heating systems of this kind have, indeed, become a commonplace today. Attempts to automate lighting have been less successful, with the costs and unreliability of such

systems seeming a waste when compared to the ease of use of traditional wall switches, by way of contrast. But other 'solutions' in the smart home struggled from the outset: for instance, contents in the fridge of the Orange@home were monitored (or at least registered as present or otherwise through users having to scan them in), so that when consumed they would be replaced by an automatic grocery order; routine needs were also modelled so that such things as 'white goods' – washing powders, toothpastes, toilet rolls and such – would arrive when needed. But the scanner did not recognise many of the items, and besides, created hassle when the fridge was being used. Most occupants stopped using the system, even though most thought it a 'nice idea'. Attempts to automate aspects of domestic leisure were also made: alert messages about the start of popular TV shows would be sent to users' mobile phones. These would summon them to the room with the TV at the right time. At first, this was expected to derive from modelling actual viewing in the house, but there was not enough data to create models and so these aspirations were soon replaced by reference to what was known to be popular: 'viewing by appointment shows' as they were called, shows that would muster millions each night: weekly soaps, Sunday night period dramas. Texts would be sent to the family members just before these started, wherever in the home or thereabouts they might be.

Examples of smart homes with similar (albeit similarly faltering) aspirations could be found elsewhere at the time: there was the 'Adaptive House' in Boulder, Colorado (2005), the House_n at MIT, Boston (2002) and, somewhat later, Georgia Tech's 'Aware Home' (2008, 3675-3680). Towards the end of the decade, Bill Gates had a team of Microsoft engineers transform his home in Redmond, Seattle, and this offered similar functions (though despite much larger budgets than Orange, some of the functions were achieved by having support staff: when guests came, they were welcomed by real butlers). In all of these, it was assumed that not everything would be successfully modelled, the impracticalities of capturing all items going in and out of a fridge being typical of this. But it was assumed that, with time, sufficient investment, and more data, intelligence in the building would be delivered. Doing so essentially required a view of domestic life as sets of behaviours, of movements in and between rooms, of incomings and outgoings, of actions that could be seen and sensed by technology. It was a kind of behaviourism, though there was no profound epistemology behind this as a view about what the technology of the time could capture.

What was the experience that resulted? In the case of the Orange@home, six families volunteered to live there for extended periods of time (up to two months). All accepted that the technology they were confronted with might not be able to deliver exactly what was hoped, namely perfect models of their behaviours, but nevertheless they were willing to tolerate infelicities in hope of something better for others later on. How they responded, though, was interesting. Their reactions expressed some pleasure and delight, but there was also an unease, a feeling that, if taken too far, the technology would be unsettling.

In particular, many of these families came to be intensely aware of how they liked their own autonomy, or rather, how they liked to think of themselves as doing things in their own family-like ways. For example, some of the families did not like having AI control the locks, as they felt that that was a responsibility that they ought to maintain. One pair of parents explained that they chose to lock the house not as a security measure but as a prelude to getting their kids to bed. Locking the door to the garden was about preparing these kids for the end of the day. And this would vary dependent on all sorts of factors – the day of the week, how tired the kids were (and how tired the parents were), the weather, even what was on the 'Telly' later that night. They explained that they did not think any computer system would ever come to know all these intersecting variables. These parents added that they thought they were unique in this practice. In their view, 'data' about other families would not help the computer systems better understand them; they were off the norm, so to say. In their view, the locking of doors expressed something about their particular manner of parenting. It was a moral matter, private to themselves and how they did things, judged on criteria beyond the understanding of a computer. However, these same parents also said that they had not recognised this before, the locking of doors to get the kids ready for bed not being something that they had thought about much, as it was something that had emerged as a kind of family habit. With the Orange@home technology trying to do it for them, they had started to think about it in new light and saw it in terms of identity. The experience of living in the smart home made them self-conscious in ways that had not occurred before.

Other families felt the same. The monitoring, pattern-seeking technology made them realise they were, perhaps, not as predictable and as regularised in their habits as they thought. Take the going to work and/or school, the leaving of the house, the unlocking and relocking of doors, the turning lights off when the exit had been made. Most families came to see that was

mostly done at the same time, but occasionally matters at work or school, even at home broke the pattern. There was a work offsite, say, meaning that departure could be a little later; or there was a dentist appointment that meant work had to be delayed whilst a child was delivered to the practice and then to school thereafter. As a result of such serendipities, automatic locking and heat timing didn't seem to fit lives that the occupants led – or at least those aspects of their lives the occupants became aware of. The behavioural models seemed not flexible enough. As with the family who used door locking to announce the end of the day for the children, other families began to feel that it would be better they stay in control of many of the things that they had thought AI might automate.

This self-awareness, not necessarily an aversion to what the technological sensors were doing, also led some of the families to see a kind of caprice in other aspects of their home life. When it came to food and feeding, for instance, some families explained that while their eating habits were often routine, and hence could be supported by a system that monitored what was in the fridge, and which would then replace food items as they were used up, they came to realise that their own preferences would cause trouble. They became aware that what they might want for dinner or lunch would alter. Desires were like moods, they began to think, and something different for dinner would be sought. Though there might be a need and justification for the replenishment of basic stuff - milk, tea, coffee, and bread – they began to think that choices about dinner were intrinsic to how they operated as families. After living in the Orange@home, many came to realise they did not like relinquishing control to the technology, not because the technology was bad or that it did a bad job at capturing their basic patterns, but because they, the family, were all too human in being resistant to being thought of as pattern bound.

Tessellating AI and People

The introduction of new technologies - and not just 'AI' labelled ones - often creates such a reaction: a relooking at the activities being supported in potentially new ways. Sometimes this revisiting is only passing, like the pause for breath before something new begins, but at other times it can lead to activities being seen more appositely. Most of the technologies in the Orange@home consisted of sensor-based systems – movement, heat, light, locks, etc. I shall come back to these shortly. But there was one tool that was process-based: a homework assistant for children. This was modelled on an expert system approach to learning – essentially a system

that allowed children to navigate through a course of known alternatives, with the system giving prompts for the right course or route (Randall, 2004, pp227-246).

I mention this now not because I want to pick up the distinction I drew between expert systems and data driven systems mentioned in Chapter 2. It is rather that both kids and the parents did not like the application and both disliked it for reasons that had nothing to do with the mechanisms of the technology. Nor was their dislike because the content did not match the curricula, though this was a factor. It was, instead, because of reasons to do with the social. Parents, for example, wanted to remain involved with the doing of the homework. For them, the technology took away what they thought was an important aspect of their role. They enjoyed giving help; homework justified 'spending time with the kids'. Doing so seemed to be a constantly sought-for aspiration, with other tasks taking parents away from the kids - work being one, housework another, just tiredness after a busy day being a third. With the homework tool (or assistant, as it was called) parents felt further divorced from their little ones, as if they were already too separated. Some parents had been suspicious of the tool on arrival for these very reasons, though others began to appreciate this when their children started trying to use it. It was then that what looked like being taken away suddenly showed its value. Meanwhile, the kids' complaints were differently nuanced, some saying that the assistant seemed to treat them as if they were in a test, rather than helping them through the process of learning. 'That's what my teacher does', one said; 'She helps me, but this thing just seems to want me to get it right all the time'. The reasons for the dislike of this 'education at home' ambition were, then, more about questions of sociability than content, or indeed interaction with that content. The assistant did not fail in doing what it was designed for, but through experiencing it, the users, both the kids and the parents, saw it as misconceived – though how only came to light when the families tried using it. They came to understand that, for them at least, learning at home was more than just a question of learning content; it was a social thing too.

Of course, it might have been that this need for sociability reflected the constrained nature of the educational tool, being based on the expert system model which, as we discussed in Chapter 2, did indeed present tasks, including learning ones, as route-like activities. The experience of this may have demanded a turn to the social, a seeking of route guidance from mum and dad. It may be that more recent approaches that are, as we saw, 'data driven', would have allowed greater flexibility in learning paths, and hence this might have led to less demand for

parents. But even so, there seems a moral here that suggests that turning to a data-driven approach would miss the point. It is not how effective a tool might be that is at issue: it is that that homework came to be recognised as an opportunity for being together, for mum and dad to spend time with their loved ones. In other words, the introduction of an AI tool made it clearer what was the phenomenon in question – the thing that intelligent machines were intended to enhance. It turned out not to be one that AI or any kind could support, whether in the automating manner of expert systems or the data driven form of NAI. The behaviour was not about learning; it was about human gregariousness, the being together of a family.

 As I have suggested, reactions of the occupants to other technologies in the Orange@home were similar, even if these other technologies were less 'expert', being more behavioural through being sensor-based. Occupants began to look at their own conduct in terms of how they thought the sensor-based AI looked at it. Users thought they understood how the AI would produce models of their behaviours and then automate those activities that were both massively routine and automatable. But then, as they considered this, these same users thought that being understood in this way would not allow the space for them - and hence future users – to be 'themselves'. In their view, it was too gross a vision, too constraining of how they wanted to be understood.

I have wanted to suggest that this vision may be thought of as 'behaviourist'. This approach, though not a term the occupants of the Orange@home used, might have got a great deal right, but users didn't like being treated as reliable performers of patterns – they were more than behaviours, or somehow different. So, though the technology certainly created some excitement and, it had to be said, demonstrated that people were willing to try 'AI' in homes, experience with the technology that this consisted in altered what people thought the way forward. Families came to realise something about themselves as they experienced the technology. While they did do things in patterned ways, they often liked to break those patterns, or at least they thought they wanted the right to. Besides, some of the families were convinced that some of these patterns were unique to them, and that the way they broke those patterns was unique, too. In sum, experience of living in a smart home led many to conclude that their behaviours were unsuitable to AI. Through experiencing what AI of the time could 'capture', they came to valorise more greatly things that had hitherto been taken for granted: the moodiness that forced feeding habits couldn't be altered, the doing of things together that could be done apart.

The Agential Period

One could call the reaction of the Orange@home occupants an expression of agency that only began to emerge with the experience of the technology. The technology made people think more about what kind of relationship they wanted with machines that were said to be intelligent. They came to see that whatever relationship they had would say something about themselves. It led them to reconsider that what mattered was their own perceptions, not the functioning of the technology. It didn't matter whether the technology was expert system type or sensor driven; its existence made users self-aware, and with this awareness, came a change in behaviour. They wanted their agency preserved.

The word agency can obscure what was at issue, or rather encompasses a variety of forms. The complex serendipities of living could hardly be called purposeful, more the outcome of intersecting practicalities, as a case in point, and here agency emerges rather than steers outcomes. The caprices of mood, meanwhile, could be said to be a function of individuals, though again hardly an expression of their reasoned agency. A desire for being together could likewise said to be a pull of the heart, rather an intentional act of the will. In other words, who was being agential, even that they were being agential, was not always clear for those inside the Orange@home. Even so, agentiality, when it includes all these instantiations, gained importance when it appeared that the mere hint of it was under threat.

Coincidentally, at the same time the Orange@home was being lived in, and more radically immediately thereafter, a range of new technologies began to emerge that extended and transformed the ways people acted at home and broadened how this agency was expressed. These technologies had various manifestations, one set impacting how families dealt with the spaces inside the home in ways that distinguished them from spaces elsewhere - in the public domain, say, or in their neighbours' homes. As the example above about door locks showed, families acted in particular ways to express how they wanted to be seen as managing their family lives. Doors were locked to turn the attention of family members away from the world outside towards the world inside, the private experience of bathing, sleeping and being together as a family. After the external doors had been sealed, the room where the TV was, for some families, the next place - or step – in this transition. It was hence a room of especial importance, as the family would settle there, together, for a short while, musing on the day that was ending and the night about to begin. Here attention could be focused, calmed, a sensed of togetherness fostered

for a few moments before the rituals of bedtime commenced even as the BBC sought attention for its own content as the 'telly' played in the background. But technology and newly possible social practices of connection and communication altered this space, not only making these transitional moments more difficult to undertake but altering what 'watching telly' meant.

A good way of exploring those changes is by recalling the TV show 'Big Brother'. This was becoming very popular at this time. One of the most interesting things about this show was not its content – where strangers were forced to live together while under the continuous observation of the TV camera - so much as the way people interacted with those strangers and the narrative of 'living together' they made. Watchers could vote people off the show, and thus could engage in making the show's 'plot'. To this extent, the viewers were not passive but active. The term 'lean forward' came to be used to label this kind of engagement (See Taylor and Harper, 2003, pp115-126). Such lean forward behaviour was also suggestive of how the private space of home could be – perhaps was being - transformed, with the activities in one home becoming increasingly interactive with activities elsewhere, in other homes. In other words, lean forward labelled a new form of agentiality. Being in the living room, watching TV, was no longer something that was private by dint of where it was done; it could be a place where a new kind of social and hence public behaviour occurred. Being in the living room of one's own home watching the TV with all the family could now mean being in lots of living rooms at the same time, at least 'digitally'.

Some providers of TV content, such as the BBC, imagined that this kind of lean forward, digitally enabled behaviour would evolve into more than just voting, entailing the active creation of more complex story lines and narrative structures. People in every home would be able to create content, and the relationship between different living rooms and broadcasters would alter. So too would the relationship between the families engaging with these new broadcast forms. What this new relationship would be was yet unsure, a greater problem being how to make such new experiences available.

The emergence of light weight mobile messaging and social media platforms offered a way forward. But these technologies shifted what lean forward meant, or rather altered the intention behind it. In particular, they altered the relationship between watchers more than the relationship between broadcasters and their watchers. With social media (accessed through smart phones), people began to share their views about shows like Big Brother with other families and friends as

they were watching the same shows. All may have voted on some show, but messaging between families and friends began to offer enchantment through making the domestic space – at least where the telly was - more social. Indeed, 'posting' to social media about TV came to be part and parcel of the experience of watching TV and not just when Big Brother was on. People began to develop richer ways of communicating, such as taking images of content to share. For a while, there was something improvised and furtive about these practices, the images being poor, the meaning of the sharing not always well expressed. Since then, of course, this kind of sharing has developed and become more standardised, with new platforms such as Instagram and TikTok enabling such sharing to be a primary currency of their users.

Be that as it may, the point is that, with these practices, so the meaning of spaces in home altered. If, before, the home might have been a space families retired to, and in a sense used to withdraw from the world (except for very occasional celebrations and events such as birthdays and Christmas), now families started to go home to do things with close friends and other family elsewhere, in other homes, similarly connecting, and all jointly sharing and making the experience of watching at home somehow a performance across homes, even elsewhere, with pubs and other public places. Sitting at home, 'switching on to switch off' so to say, turned from an activity for one family wanting to conduct itself in its own ways, into a social process, one that delivered a sense of togetherness with other families that was only a screen-tap away.

This was an altogether much more powerful expression of agency than had been possible before. The geography of the home was being transformed by (and partly into) digital spaces (hence the title of my second book on home life: *The Connected Home* (2011)[5]. Whereas the initial smart home agenda had sought to use technology to take things away from the home occupant, to reduce, as it were, their agency, smart phones and the mobile internet led people to start thinking about technology as tools to do more. This is why I am calling this period the agential one – not because agency did not matter before but because 'being in control' through technology came to be the leitmotif of this period.

Space, Agency, Identity

Above I remarked on the complex behaviours that are encapsulated under the term agency, drawing particular attention to how emotion and desire might drive action rather than reasoning and thoughtfulness, and how the patterns a family makes (and which in theory technology based

on sensors can 'see') can be an emergent phenomenon, the aggregation of practical contingencies and individual purpose. With the spread of social media, mobile internet and so forth, just who could act 'agentially' was opened-up, adding to the complexity that already made-up home life. It wasn't just mum and dad posting comments from the living room couch as it could be teenagers and adolescents too.

Not everyone thought this a good idea. Many social theorists of the time argued that the digital was dissolving important boundaries between the public and the private and even threatening the family as a unit of safety. In their view, this was most obvious in the ways the protective space of homes could be violated by digital predators. Facebook was critiqued by some as a technological open door for ill-motivated strangers who would access the very young, getting into their private bedrooms without parents knowing. Agency could be the wrong person's agency.

Although there was some merit in these concerns, the way the digital was being used to reshape space was more subtle than these arguments implied, however. The digital did not dissolve the social power of 'bricks and mortar' in quite the way that was suggested, as it allowed new nuances as regards what bricks and mortar symbolised and enabled. As I reported in my book, *Texture* (2013), the take-up of platforms like Facebook amongst adolescents and teens was not to be seen as enabling strangers to contact and threaten them in new ways as it was more accurately judged as the means through which adolescents and teens made the relations they already had stronger and with more clearly defined boundaries. Though there were instances of digital predators, the research showed that the young (or at least of school age) used Facebook to do things like homework with their best friends, and so extend their friendship beyond the physical domains of school where they were given that homework. There was little likelihood of digital stalking as they only networked with those they knew.

This research also showed that they used Facebook as a way of controlling their own domestic space. With Facebook (and other social media platforms, some of which are now forgotten such as Myspace), adolescents and teens could make their bedrooms more private within the confines of family life. With Facebook especially, teenagers found that they had a 'place' – albeit a digital one - that mum and dad could not visit. They learnt that they did not have to give their parents access to their Facebooks. Virtual doors turned out to be more effective than real doors. Teenagers knew that bedroom doors were all to easily opened. They discovered that

the door to their Facebook was entirely in their control. What they perceived as security was enhanced, not undermined by the digital.

Put differently, the digital allowed the social space of the home to echo and make more concrete some of its real 'bricks and mortar' dimensions. As it did so, how families were complex social entities became clearer, too. The technology in smart homes had presupposed that behaviours could be seen and hence modelled, not that there were political differences in the status of the bodies seen. There had been an oversimplification of who lived in homes, making families seem to be always of one kind: the idealised young family with happy parents and two giggling kids. In this scenario, there was no likelihood of one member of the family saying to another, 'keep away'. iPhones, social media and such had the effect of accentuating the importance of how families can be diverse, with individuals inside families having their own trajectories. By the time I wrote *Texture*, it had become clear that, if one wanted to design technology for the home, one should not automatically think about the occupants as individuals who only come together to celebrate their love and affection. One needed to think of 'family' (or families) as a category that could include individuals who might have very different agendas. For though it might be the case that when individuals meet and fall in love they make a home together, at some later time those same individuals may fall out of love and seek separation, and if married a divorce. Couples falling apart is as much a consequence of agency in homes as is the opposite. The home couples occupy will be, in some respects, the physical embodiment of the idea that they have of each other and their relationship in time, whatever it might be: at one moment, homes to be made and at another to be broken and divvied-up. Besides, it is not just the purposes of couples that matter in home life: children grow up too, and, as they move towards adulthood, seek to *leave* home. As a consequence, the trajectories of different members of a home can have almost opposite concepts of what a home is: for some, it is the place to end up; for others, the place to escape from.

The Period of Natural Interaction

What I have been showing is that, in reaction to early attempts at bringing AI technologies to the domestic sphere, people inside smart homes came to baulk at the way their behaviour was seen by that technology. This was not because the technology was bad at getting data, at representing that data in terms of a model. It was that the individuals whose behaviour was being captured did not

like what was implied. They wanted to reshape, if you like, the view of technology (at least as they understood it), claiming that the values they gave to some actions were not always proportionate to frequency and-or likelihood, and hence hard to see from the technological perspective. People wanted to draw attention to one-off behaviours, as if this was how they expressed themselves. By the time the iPhone came along (and partly to do with the opening-up of the internet and mobile social media that device enabled) individuals found technology offered them new ways of expressing, of being who and what they wanted, and this in turn altered what homes meant as places where this newly sought for agentiality could be found. Instead of being places where machines might take on some of the drudgery and labour, homes came to be places where people could network with friends and family, posting content, sharing views and engaging with the digital world outside that home. Though aware that AI-type technologies might well be part of the infrastructures that enabled this agency, the enchantment with what the technologies allowed them to do meant that less interest was given to how the technology worked, or whether the things it saw misrepresented how they wanted to be seen. The various technologies constitutive of AI at the time were allowed to disappear into the screens and the keyboards as people took more interest in what *they* were doing, not in what *it* was doing.

The technology did not disappear, of course, even if it slipped from the perceptual field of many. On the contrary, it has come to remanifest itself in homelife in quite striking ways, even if those ways don't seem to provoke much self-reflection on the part of users as did early instantiations of AI in the home. Key here has been the development of so-called natural interaction interfaces, such as voice-based command and control systems, and associated digital assistants – Siri, Alexa and similar. This is the last and still the current phase of smart home technology.

Natural interaction includes systems that use vision, such as Microsoft Kinect, as well as those that used auditory control such as Alexa. I won't say anything about vision systems as they have not flourished in the ways that voice-based ones have, one reason being that body movement offers a quite limited set of meanings for a camera-based system to catch, and as a result, users of such systems have to learn a set of movements that the computer recognises. The vocabulary of interaction has thus been limited; the behaviours too can seem awkward, making for embarrassment. Voice-based systems, especially latterly, have been much more able, and the vocabulary they can parse much greater, and hence have received much wider acceptance.

Besides, there is a tendency for people to be simply impressed by the technology: it appears to listen and then it speaks back appositely, as if it were human. When first used, many users are left in awe[6].

The techniques used to enable these systems were being developed prior to the LLMs we discussed in Chapters 2 and 3. The devices and applications that deploy them have gradually been introduced into home settings over the past decade, and are now quite commonplace, and, as I say, often well received. In addition, and of relevance to arguments above, these systems have not provoked the kind of resistance that early smart home technologies did. One reason for this is that voice-based systems model phrases and commands, not what people do in the home. For many users, there is no offence in their spoken words being modelled, as the output of that modelling is understanding what they mean. In contrast, modelling of behaviours does not always deliver the meaning that is desired. As we saw, people can dispute how to interpret their actions. When they use words, in contrast, they imagine they control what is understood. At the same time, other aspects of the spoken mode of interaction seem self-evidently beneficial to the user – no keyboards, no touch, sometimes even, no screens. Hence the term, 'natural' to label the interaction. This is not to say that there are no complaints, with many users whinging at the time it can take for voice-controlled systems to learn their accents. There can be complaints too at the way children mock these systems through offensive language and trick questions. Once the 'learning' is complete, however, many find such systems allow the control of their consumption of, for example, music to such an extent and with such apparent ease that one might say they enable 'self-curation'.

One can paraphrase what a user might say to achieve this: 'Alexa, play X!' and hence a piece of music is heard that says something about the taste of the user, the mood they are in or even the mood they want to make. Their music selection curates not just the music, but themselves too. As this happens, the role of the technology, its mechanics, its centrality to what is being afforded, gets lost. The self becomes the topic, not how technology comes to understand. Hence, later, when the same user may want to alter their mood through what is played, they issue a different command, resulting in different music: they are once again curating their identity and not thinking about the technology at all. Saying 'Alexa, play Y!' is less of a command to a computer, as it says something about the one making the commands. In this sense, the AI that

enables voice-based systems seems ever more supportive of who users want to be and, in the process, makes itself invisible.

Self-Curation as Data for Platform Capitalism

This comes at a price, however. Most obviously, the actions of the user, their choices, how long they listen, how often, and in what order, becomes part of the data that the voice systems offer up to the supplier of that device (if I can treat the term 'supplier' as sufficiently comprehensive to cover a multitude of relations between hardware maker, device retailer and the software platform organisations behind them). As we saw in Chapter 3, what users do in the digital doesn't just create a footprint but is linked in diverse ways with other traces of action, other footprints, and through computational prediction, used to create outputs that shape other, perhaps quite removed digital contexts and behavioural opportunities. A player of music at home might find, as a result of some play selection, their Facebook page suffused with adds or click-throughs to matters that seem quite removed from that particular piece of music. But the connection will have been identified and offered up as a result of how other listeners to that same music have strong preferences for certain kinds of clothes, a brand or style, which the AI has identified. What started as the curation of some music for the self, and, through extension, in effect, to self-shape through that selection (as in 'I am the person who likes this music' and hence will curate my listening habits accordingly) leads to self-entrapment, if you like, when the self becomes expressed as a point in the murmurations of thousands of others.

This is to sketch the way connections are made, of course. Many individuals might find nothing sinister about such connections, and if asked might respond by saying that this seems a modest price to pay for the delights of self-curation: personalized selection is leading to personalized adds. But what I want to draw attention to is what this implies about agentiality. For this is suggesting that what feels like agentiality in one dimension results in agentiality being subtly constrained and directed in another. The sense of freedom a person might have at being able to select music with a single phrase comes at the cost of narrowing their freedoms elsewhere – when they are on the web, or when they are wanting to connect with friends via social media. The confusions and opacities that result were discussed in Chapter 3, so I will not say anything more on them now.

What I do want to draw attention to is a different price that users are likely to pay with voice-based interaction. This is more local, pertinent to their actions rather than to some distant point in digital vectors on the web. This price is paid when people try and use voice-based systems to engage in what might be matters of a search of some kind. For, with voice-based systems, it is likely that users avoid extended interaction over several iterations. If they ask a voice assistant what might be a good choice for some product, they might listen to one, two or three items, but not a list of many more. They do so because listening takes time, certainly more than would be the case if a search had been done through a screen and the results presented as a list to be looked at. A feature of a visual field is, after all, that it can be seen as a whole and at a glance. If one uses Google on a PC, for example, one can see how many items are listed on the first page of the search results page. One can survey this rapidly and then choose which one to inquire further into. Items on a list presented through sound, in contrast, that is to say through a speech-based interface, have to be listened to in sequence. Of course, one could listen to them all but doing so would seem to contradict one of the appeals of such systems - that they are quick to use. The upshot of this is that a user of voice-based search is likely to select one of the first items offered, not one of the later ones. This might not matter if the list is constructed strictly on the basis of objective search. But as any user of current search engines knows, the ordering of items on search query pages is typically driven by advertising, and not just relevance. In the case of some suppliers, products listed at the top are the ones being promoted, not necessarily the ones best suited to a searcher's needs. When this is combined with the likelihood that the user will opt for a quick-to-select option, the use of voice-based systems has strong advantages for the supplier of content to those systems: they can predict what items are selected by dint of the position they have in the order of presentation. Needless to say, the same principle applies to search via WIMP systems: where an option is placed in a list will affect the likelihood of its selection. What I am saying here is that with spoken interaction this likelihood is stronger, making placement even more valuable to the vendor.

I am also wanting to suggest, however, that this begs further questions about agency. For while a user might think that voice-based systems allow them to self-curate in ways that is infinitely more refined and moment by moment than is possible with other forms of computing, the kinds of agential acts that they undertake have consequences for the scope of that agency. In the case of a search, they will find the patterning of their actions taken advantage of by the

supplier, by the owner of the platform that operates the speech-based system. Though they might imagine they have more agency than ever, in this regard, the real agency is someone else's.

Let me put this another way: an unwillingness to engage in lengthy spoken interaction, perhaps combined with a temptation to be awestruck by such systems, can result in users losing a sense of what they are about when they 'speak' to their 'smart speaker' or whatever assistant they use. They might come to neglect noticing how a natural interface, a natural language one, might be hiding what others, such as those behind the 'speech bots' in their smart speakers, are wanting to do when their technologies get used. The company that supplies the technology might have its own agenda as regards what comes with that technology, using it to increase the sales of preferred products and services in ways that is not visible or at least not noticed by the user. It might be the intelligence of companies like Amazon and Apple comes to tacitly operate behind the scenes of the home with natural interaction interfaces, seeking to do things that might not be in the service of the user's own interests. It is not the intelligence of machines that is consequential, then, but that of organisational actors elsewhere and that comes at the expense of the intelligence of home occupants. This shifting of whose intelligence matters is hidden by the charms of computers that listen and speak in return and about which we use the term intelligence. But that very use distracts from where the intelligence now is: not one in a machine at all, but inside corporate walls, inside 'platform capitalists,' in the marketing departments of the Silicon Valley corporate megaliths.

Conclusion

As it happens, much of what I am suggesting is not as consequential as one might fear. Though voice-based systems controlling smart speakers and the like are often well received at first, there is a tendency for people to get bored with them, preferring to interact with computer systems that offer richer modes of interaction - tablets, for example, as they are able to both control audio speakers via online music services, while affording a screen for search (and hence the visual), as well as a virtual keyboard for posting comments on social media. They can do all this while they offer an option for voice-based interaction when desired. At the same time, the vendors of the voice-based systems, though hopeful of the benefits described above, have found that their systems suffer from what the suppliers themselves call 'poor service discoverability'. By this is meant that these voice-based systems are often designed to allow users to access many different services and products and not just music content nor access to search tools, nor to one product

supplier with its own intranet search functionality, such as Amazon. Each of these other services would offer further sources of revenue, and this was expected to improve the cost margins on the manufacture of the products including the development of the software engines that run the speech interfaces. But users do not find discovering these other services easy, and so they remain hidden, and the sought-for revenues are not realised. Even the most popular of these products are gradually being withdrawn from the market[7].

My concern, however, is not with the fate of these products as it is to use this as an example of the relationship people have with AI-like products in the home setting. These most recent relations connect instances of natural interaction devices that are in many ways very different from the sensors-based AI typically used in early smart homes. The occupants of the Orange@home would have been amazed, even incredulous at Siri and Alexa. As it happens, both the early technologies and the latest natural interaction ones use frequency and likelihood engines, though the data that drives those engines are different as are the ways that the data are aggregated. But if AI technologies have changed their processing tools over the years and so too their interaction modalities, then so too have the behaviours of users that attend these technologies. I have been showing how people react to AI varies as much as do the computers that are labelled AI. There is certainly no single thing that looks like AI in the home nor yet is there a single thing, 'the user' in the home, nor yet a singular reaction to AI in that setting. There are certainly observable fashions in those behaviours, reactions to the technologies in question which shape how the technologies gets used and which in turn shape user action – I have listed three broad eras to denote this: the behavioural, the agential, the natural. Doubtless, one could determine more, albeit the term era will lose its valence if it becomes too detailed. The point, though, is variation and the systematicity in this.

From inside the world of computing, especially when the smart homes agenda emerged in the 1990s, that there would be such variation was not so obvious. It was commonly assumed that technology of the first period, for monitoring, would be straight-forward to build and use, the gathering and building up data only needing evidence about movement. Users of the time came to think otherwise: movement was not a good index, at least if intelligent machines were going to 'understand them'. As we saw, AI wasn't really seeking understanding, so much as automating, the real question being what could be automated: door locks could be but were rejected, automatic lighting likewise; heating control on the other hand has become a success.

What does one learn about AI? Take the last instance of technology: one can hardly say that managing heating requires intelligence; a clock is needed with a register of comings and goings. And that, indeed, is all that contemporary AI technologies like Nest are, their advertising claims notwithstanding: simple clock counting devices linked to movement monitors. Certainly, they aggregate patterns but how they do so and what they make with those patterns hardly seems to justify the term intelligence (artificial or otherwise). They are more like plants, reacting to light and heat as if governed by laws; they are not like animals judging weather and adjusting their behaviours intentionally. There is no sentience governing alternate paths of action. For Nest it is 'this' or 'that' given a prior and the prior is mechanically gathered data; not matters of judgement. What value does the term AI give here? Does the linking of that technology to other types of AI under the banner of NAI increase its value? It certainly might increase the price people are willing to pay for it. Whether it leads to users better appropriating the technology once in their home, that is another matter.

If we look at natural interaction devices and hence another 'era' we might conclude something different. Here, though there might have been some intelligent engineering in the machines that support spoken interaction, where intelligence turns out to be in the 'ecosystem' of these technologies is more consequential. I am using the term ecosystem to point towards interdependencies: between the speech engine, the users' use of the same, the desire of those users to self-curate through the light touch of the spoken word, and the sifting of that action into aggregates by what I called the platform capitalists: these are elsewhere, or rather partly in the home and partly outside. I have suggested that it is these organisations that use these aggregates to shape the paths of not just the single user in their home but those of myriad others users elsewhere, outside of homes. This aggregation will use various machine learning techniques to identify patterns and these will be made available to these platforms to judge how they might steer the actions in question. But is it artificial intelligence that makes those judgements? Or is the machinic processing subordinate to the intentions of the organisations in Silicon Valley? It certainly is organizational intelligence that seems afoot, an intelligence intending to make profits. But when one uses that term, 'organizational intelligence', suddenly the meaning of the second word alters: now the intelligence seems ruthless, cynical, money-seeking. When I read Suleyman's book, The Coming Wave (2023), in this light, I don't think of artificial intelligence and the magical power of creative machinery. I think of capitalism. Is that what NAI does? Does it hide this

interpretation by leading our thoughts towards the ineffable? Whether it does so is perhaps beyond knowing with the evidence at hand. But that it might ought to give us pause for thought.

What is sure is that the very different AI technologies, that I have claimed reflected different eras, all sought patterns. The relationship between people and these patterns was different in each case. In the first, they wanted to own those patterns, insofar as they felt it important that they steer how the captured patterns were interpreted. As we saw, though, it wasn't that the AI got the patterns wrong. Rather, in being aware of these been captured, the subjects seen in those patterns came to adjust how they wanted to be understood, even to the extent of saying that 'patterning' was not applicable to them. Some of the families observed came to insist that their behaviours could not be seen as typical of families. Whether this was true or not wasn't the issue, it was their claiming the right of self-definition.

In the era of natural interaction, one does not find the same concern – the patterns inferred through practices of self-curation do not seem to be defended in terms of who owns the meanings that derive. Family turns out to be done through self-awareness, reflection, self-definition. Being an individual, on the other hand, does not seem so political in the home setting. The management of what a family is thought to be appears to matter more than what being an individual is thought to be, at least in the particular geography. Of course, and indeed as we shall see, in other geographies what being an individual is might be as political.

Performance, politics, the suggestion that the doing of family is an expression of a kind of mental acuity – this is the sum of my analysis. Do they add up? I think they do. We have seen that members of families seek to shape how they are seen and understood. They pro-actively create what family means. And we have seen that what family means can also be intimately linked to the places that families call home – even a temporary home, like the Orange@home. Homes are not static, they are made (and they can be unmade too). The making is not in their bricks and mortar as in how the spaces are used, talked about and made again through the 'practice rhythms' of domestic life. Homes and families are doings; doings that make certain geographies what they are: spaces of special meaning. And what we have seen, the most important finding, is, I think, that the kind of meaning people seek is not something that can be made by AI of any kind. This is not because the technology might misrepresent patterns and hence express the wrong meaning, as it is not the business of machines to do this. People want to keep this for themselves; it is part of the geography of what they are. It is a measure of them, of their intelligence.

---

[1] That AI tools tend towards behaviorising observed action has been a concern for some time. See, for example, Rouvroy (2012); also Hilderbrant (2006).

[2] I could explore this term and whether it makes sense – speaking to a machine is not the equivalent of speaking to another person, even if the sound shapes are similar. Its grammar of action is different, to put it in terms used in Chapter 3.

[3] See Aldrich 'Smart homes: past, present and future', in Harper (Ed) (2003: 17-36).

[4] This labelled the failure of 'IT' to enhance productivity between the late Seventies and the end of the Nineties. See Brynjolfsson (1993), https://dl.acm.org/doi/10.1145/163298.163309; and Dewan and Kraemer (1998), https://escholarship.org/uc/item/3dd2j1gd. For further discussion, https://en.wikipedia.org/wiki/Productivity_paradox (accessed 02 February 2024).

[5] Harper (Ed) (2011).

[6] The same could be said about the first WIMP systems, of course, and indeed about many technologies when first new: that they work can seem so amazing that it seems like magic, to paraphrase the great sci-fi writer A.C Clark. He actually wrote: 'Any sufficiently advanced technology is indistinguishable from magic.' See https://en.wikipedia.org/wiki/Clarke%27s_three_laws, accessed 18 March 2024.

[7] In the case of Amazon and its echo product (where Alexa is the name of the assistant), though initial volumes of product sales were high, revenues through the device soon abate, leading to the likelihood of the product being withdrawn. This naturally creates considerable interest in technology blogs. See https://www.thestreet.com/investing/amazon-alexa-is-everywhere-but-revenue-is-elusive, (accessed 02 February 2024).

Chapter 5

Sample from R. Harper, *The Shape of Thought: Reasoning in the Age of AI*
Forthcoming, McGill-Queens Press, Montreal, Canada