Machine Intelligence

Alex S. Taylor
Microsoft Research
7 J. J. Thomson Avenue, Cambridge, CB3 0FB, UK
ast@microsoft.com

ABSTRACT

Under certain conditions, we appear willing to see and interact with computing machines as though they exhibited intelligence, at least an intelligence of sorts. Using examples from AI and robotics research, as well as a selection of relevant art installations and anthropological fieldwork, this paper reflects on some of our interactions with the kinds of machines we seem ready to treat as intelligent. Broadly, it is suggested that ordinary, everyday ideas of intelligence are not fixed, but rather actively seen and enacted in the world. As such, intelligence does not just belong to the province of the human mind, but can emerge in quite different, unexpected forms in things. It is proposed this opens up a new set of possibilities for design and HCI; examining the ways intelligence is seen and enacted gives rise to a very different way of thinking about the intersection between human and machine, and thus promotes some radically new types of interactions with computing machines.

Author Keywords

Intelligence, intelligent machines.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

For some time now I, and others I work with, have criticised the ideas of intelligence that are prevalent in technological imaginaries. We've given special attention to smart homes and how many of the underlying motivations sit uncomfortably with the kinds of practices we've observed in our various empirical investigations of home life [33]. We, and others far more qualified to do so (e.g. [26]), have reasoned that it is not only a massive and sometimes intractable technological undertaking to get intelligent technologies to work in the ways promised. The very promise has been fraught with theoretical and ontological uncertainties provoked by long-standing debates around what it actually means for machines to be intelligent (e.g. [6, 31, 34]).

As an alternative to building objects and environments with intelligence (as in the smart home), we've thus aimed to demonstrate through a range of examples that information and communication technologies can be designed so that they enable *us* to act intelligently. In other words, we've tried to re-direct the programmatic efforts of domestic ICT design away from the possibility of intelligent things and towards an idea that privileges human intelligence.

Reflecting on this position, one resounding problem that recurs to me is how the person on the street, if you will, regularly attributes life-like qualities to things and, in doing so, occasionally imbues them with intelligence. This is done without observable difficulties. We readily describe our household appliances as clever, smart and so on, and we are seldom slow to refer to objects as dumb or stupid. Although these phrases are not immediate equivalents to intelligence, they are suggestive of the idea,[1] and more importantly demonstrate our capacity and indeed willingness to attribute inanimate things with the ability to think.

Considering what I've increasingly come to see as the insightful points made by Edwards and Grinter [7] on the piecemeal adoption of technology into our homes, the aversion we've had in our research to the idea of intelligent things seems all the more restrictive. What I've begun to contemplate is the possibility of people gradually taking up seemingly inconsequential digital things that they see and treat as smart or intelligent—intelligent not like humans, but nevertheless intelligent in some casual, everyday sense of the word. Taking this possibility to its logical conclusion, it's plausible to imagine many of us living in environments suffused with an 'artificial intelligence', but with us having reached such a point without any clear purpose or intent. As Edwards and Grinter suggest, a home's smartness may come about incidentally as increasingly advanced technologies are incorporated piecemeal into our homes. Arguably, we've already started down this route. The readiness of youngsters and adults alike to care for their Tamagotchi-like devices and the less pronounced but perhaps more significant uptake of relatively simple robots like iRobot's vacuuming *Roomba* (more of which I'll discuss later), lend support to a vision of the most rudimentary 'thinking' machines seeping into everyday life and, significantly, our willingness to accept them as such.

What I want to do in this paper is ask how unremarkable

[1] That is, part of those 'determinative adjectives' that Ryle [29] associates with the "family of concepts ordinarily surnamed 'intelligence'." (p. 26)

CHI 2009, April 3–9, 2009, Boston, MA, USA.

things like toys and vacuum cleaners can be treated in these ways and what, if anything, this has to do with interactive system design. My intention is not to move from a position some might see as naively nostalgic to one propounding technological hyperbole. Rather, I hope to critically engage with what I believe to be the very real possibility of some significant transformations in computing technology, the consequences of which I will go onto argue have important implications for human-computer interactions.

Lessons from AI and Robotics

Thanks to the dogged persistence of one of my colleagues, I've become aware that some important lessons for this exercise in reflection can be had from casting an eye towards the developments in artificial intelligence (AI) and robotics over the last fifteen years or so. AI, as has been well documented, has undergone significant changes since the 1950s when Newell and Herbert [25] laid the groundwork for the project. Good Old Fashioned AI (GOFAI), as it is now known in certain circles [12], was based on the premise that human reasoning could be symbolically depicted, represented and modelled from the top down. Then at the heart of AI, this premise was a direct antecedence to the mechanical rule following Turing [31] ascribed both computing machines and humans (specifically the human mind). An exemplar of the GOFAI position was and still is the artificially intelligent chess 'player' designed at least in theory to incorporate a searchable model of all possible permutations within a certain range of play. This brute-force approach eventually resulted in 1997 in IBM's much heralded *Deep Blue II* (or *Deeper Blue*) beating the then World Chess Champion, Gary Kasparov (the win not without controversy). The hardware and computational demands were extreme for the time though, even with the clearly bounded rules and constraints of a chess game.

Suffice to say GOFAI has been heavily criticised not only for the substantial demands its approaches place on computing machinery but also for its inbuilt assumptions concerning the processes of human thought and action. Even in the subsequent but closely associated cognitive science programmes, serious doubt has been cast on the idea that thought is systematically and sequentially processed [6, 35], let alone the presumption that these processes occur solely in the brain [5, 14]. Those levelling such criticisms have argued that machines like Deep Blue and its offspring aren't so much indicative of a machine's capacity to simulate human intelligence, as they are illustrative of the hard problems AI programmers face in achieving anything like the dynamic, moment-by-moment qualities of human behaviour and reasoning observable in the world.[2]

In a similar vein, but from very a different perspective, Lucy Suchman's now well-cited book in HCI, *Plans and Situated Actions* [32], offered a thorough critique of intelligent systems. Through a set of *in-situ* experiments, she demonstrated how human behaviour, even when planned, is an emergent phenomenon worked out on an ongoing basis. In other words, Suchman showed how *a priori* representations of planned actions, on which GOFAI is predicated, poorly map onto people's real-world practices. I will return to this point in a moment. Here, I want to emphasise that the criticisms directed at GOFAI (as well as related shifts in funding) led to a good deal of introspection on behalf of the AI community (see [21]). Researchers and academic departments distanced themselves from ambitious AI programmes, a move exemplified by the adoption of less grand terms like 'machine learning' and the use of AI's tools to address more bounded problems such as image processing, computer vision and expert systems.

From this period that some rather poetically called AI's winter, a "New AI" has taken shape (in some cases, explicitly building on Suchman's ideas [1]). The progenitor of this turn, Rodney Brooks, has devoutly pursued *emergence* as a central tenant to machine intelligence [4]. Broadly, Brooks and his successors at MIT have sought to build systems with initially relatively simple computational underpinnings that adapt and learn using bottom-up processes so that intelligence 'emerges', as it were, when the technology is put into the wild. The key difference from GOFAI has been to eschew any effort to represent the world, *a priori*, and instead have the machines react to as well as learn from the situations they are placed in. Brooks originally posed his ideas in terms of primitive life forms, robotic insects, etc. [3] although, unsurprisingly perhaps, they have gone onto find form in human-like robots. Inspired by Brooks' thinking, Kismet [2], for instance, offers a mechanically sophisticated although visually crude example of a humanoid robot designed to learn and respond to emotional cues.

Alongside these ambitious works, an assortment of smaller scale projects has similarly built on the idea of emergence. Yet distinctive amongst at least some of these projects has been a move away from simulating the kinds of thinking, actions or motivations we would normally associate with being human. For instance, Frederic Kaplan and his colleagues at Lausanne's EPFL robotics lab have experimented with the intrinsic motivation of robots [17, 27]. Using Sony's robotic dog AIBO as a platform, Kaplan demonstrates how robots can independently develop their own behaviours when programmed with simple learning algorithms. Collaborating with interaction designers, Kaplan exhibited examples of AIBO (stripped of its commercial programming) gradually learning to move, 'cycle' (on a custom built bike) and 'swim' (in a specially designed swimsuit), and appearing to do so in idiosyncratic ways. Beyond the amusing results, a crucial aspect to work such as Kaplan's is the difference in thinking about machine intelligence: the work indicates how machines can exhibit seemingly intelligent behaviours without having to replicate the mental processes of humans, or animals for that matter.

[2] The terms *AI-complete* and *AI-hard* are used to refer to the foundational problems associated with building computers to be as intelligent as humans. This, incidentally, reflects one of the underlying motivations in AI (and, in particular, GOFAI).

A newly found pragmatism in AI has further contributed to this apparent rethinking. Sidestepping Brooks' notoriously dogmatic position against representation [3] and the quagmire of arguments around intelligence, a stance has been adopted that combines elements of GOFAI and Brooks' adaptive learning. Seen not as mutually exclusive, the different approaches are used as alternative resources for design to produce 'hybrid systems'. The consequence has been to promote a more experimental use and application of machine intelligence, again, where new rather than rehearsed ideas of behaviour and intelligence are pursued.

The point to this potted history is that there are some interesting lessons to be learnt from the circuitous route travelled between the brute-force machines of the GOFAI era to the robots and agents we see today. What I want to suggest is that the developments in AI have been instructive on two broad counts. First, and putting to one side the ambitious projects from the likes of Brooks and his successors, AI and robotics research appears to be offering a path away from a rigid and restrictive notion of intelligence in machines and towards alternative and in some cases new possibilities. To reiterate, no longer is it assumed that the sort of intelligence to be attained in machines should simulate human intelligence. More subtle ideas are finding a foothold, leading to machines suggestive of different forms of intelligence.

The second lesson comes more from outside AI (from an active and maturing critique of the developments in AI and robotics), and provides an interesting possibility for (re)-framing the ideas we have of intelligence. As it happens, Suchman's *Plans and Situated Actions* has played an important if not key role in this critique [30]—it criticised assumptions implicit in intelligent or expert systems, particularly those being designed into the photocopiers her then organisation manufactured. The primary lesson HCI has drawn from this work has been to recognise the situated character of human action—that even planful behaviour is contingent on the contexts in which we find ourselves. Suchman though had another major theme in her work that was given far less attention (although one she has since pursued and elaborated on [32]). As part of her critique, Suchman articulated an argument that foregrounded the constitutive nature of the human-machine intersection. That is, how both the ideas of human and machine are manifest features of real-world interaction. She suggested that even the idea of human-computer (inter-)action implied assumptions around what humans and machines are.

Suchman's arguments parallel and in some cases build on a programme of work—if it can be called that—that has recast or reconfigured the ontological distinctions made between humans and things. Influences include actor network theory [18, 19], the theory of extended mind [5] and feminist techno-science [11, 13]. Like the work from within AI, these positions although by no means unified, provide a way of re-imagining what machines might do. Importantly, for this paper, they also encourage us to rethink intelligence itself. They have us reflect not only on the "figurations" of human intelligence, but also to take seriously the prospect of alternative and possibly new conceptions of intelligence.

INTELLIGENCE AND HCI

In this paper, then, I will draw on these two themes to reexamine the critical position I have previously taken against intelligent technology and smart homes. I'd like to build on the themes to take the possibility of machine intelligence seriously in a way that I—and I'd argue much of HCI research—have chosen to avoid. I thus see the following as a thought piece, taking as its starting point a determination to rethink intelligence in HCI.

With this in mind, the body of the paper will examine a range of ideas and projects, and aim to illustrate that there are some nuanced ways of thinking about intelligence relevant to HCI. The ideas that have been influential for me have their roots in philosophical discussions—beginning in the 1930s—that centred on the mind and intelligence (contributed to by such notables as Ryle, Wittgenstein and Turing). Also influential have been contemporary arguments from critical theory (e.g. Haraway, Hayles, etc.) philosophy and the social sciences (e.g. Clark, Dreyfus, Latour, Suchman, etc.). The projects I will use to elucidate some of these ideas range from art installations, to anthropological fieldwork, to AI and robotics research. I've grouped these projects into three 'thought exercises' demonstrating how they have been, more than anything, an aid for myself in thinking through the many relevant areas and in helping to organise my thoughts around the possibilities of machines that are not exactly like us, but are nevertheless something we might imagine to be intelligent.

Towards the end of the paper, I consider a number of implications the three thought exercises raise for HCI and pose questions for those of us in HCI who may want to consider the design of intelligent machines. Intentionally, this discussion is kept broad in the hope that it opens up areas for discussion and future research.

Related work in HCI

Before I delve into the three exercises, I should emphasise that the thoughts above have aspects in common with several reported HCI projects. Perhaps the most obviously relevant have been a small number of field studies of domestic robots. I will return to a study of the Roomba from Forlizzi and DiSalvo [8] below, but what has intrigued me about these works are the ways households develop relations with their robots over time and how this can have a bearing on how machine intelligence is understood.

Other examples in HCI have shown that such relations are not just limited to mechanical robots; devices that display even the slightest capacity for 'interpreting' behaviour also appear to provoke a level of curiosity about their 'inner' workings. Two systems, one the *Home Horoscope* [9] from the Equator Project, and the other the *Tableau Machine* [28] from Georgia Tech and UC Santa Cruz, demonstrate the lengths household members will go to make sense of machines that apparently respond to daily rhythms. Evoking a similar curiosity, a project by the *Future Applications Lab* [22], in Sweden, relies on AI programming to project changing wallpaper patterns onto a home's walls (the patterns based on photos sent via mobile phone).

What I find particularly compelling about the above examples is how some relatively simple yet seemingly independent machines can produce quite different and novel kinds of interactions. Indeed, they align closely with some of the themes I have discussed above and, in doing so, contrast with a growing body of research being undertaken in HCI and emerging areas like *Human-Robot Interaction*. While some important results are being produced in this research, the focus remains largely wedded to gauging and designing human-machine interactions in more conventional terms. Work centres on, for example, the recognition and production of facial expressions, gestures, body orientation, etc., as well as explicit efforts to replicate human qualities in machines (e.g., [16, 37]). Again, I have no doubt this work is worthwhile. It does seem, however, to adopt a limited perspective on human-machine interaction and runs into the same old problems of whether replicating anything close to human qualities in machines is achievable. As I've noted, my hope here is to explore whether there is scope to move away from this restrictive debate.

Finally, two influences from HCI deserve particular mention. One is a recent paper from Leahu et al. [20] presented at the annual conference on Ubiquitous Computing. In brief, the paper's authors argue that by borrowing on some of the lessons learnt in AI, ubicomp might refine its outmoded position on machine representations (e.g., those incorporated into smart home systems). In particular, emphasis is placed on insights from what they refer to as an 'interactionist AI' (a more refined perspective, but broadly similar to the New AI characterised above). The paper's overall objective is thus to promote a shift in thinking in ubicomp systems development. In many ways my hope is that the following will be seen as an addendum to this work. My specific addition is to suggest that AI's developments give cause for a renewed interest in the idea of intelligence in HCI, where intelligence itself becomes the design focus.

A second influence I wish to raise has much in common with this position as well as the overall thrust of this paper. Inspired by Michael Mateas' notion of *Expressive AI* [23] and further developed with the idea of *Alien Presence* [28], members of Georgia Tech along with Mateas have experimented with using AI to design machines that behave in coherent but at the same time non-human, 'alien' ways. Again, the ideas that follow are seen very much as a progression of this parallel work; like Mateas and his colleagues, I wish to consider the possibility of machines that are unlike humans but that are seen, in a fashion, as intelligent. What I want to add to this a further explication of how intelligence is seen in the world and, as a consequence, to reflect on what else, *besides* AI, might be considered in designing intelligent machines.

With regards to intelligence, and to make one last point before moving on, I must note an important caveat: so far I've muddled the words intelligence, smart, clever, etc. Of course, research into AI and robotics have produced a litany of themes, categorisation schemes and even laws that are meant to somehow characterise our interactions and relations with intelligent machines (somewhat bizarrely the best known, *The 3 Laws of Robotics*, is from Asimov's science fiction story *Roundabout*). In the following, I've intentionally chosen not to get bogged down in this rule or law making, but instead to treat intelligence in everyday terms. In this respect, my interests thus focus on how we come to use intelligence and words like it in *ordinary* ways (see [10]).

THREE EXERCISES:

Seeing intelligence

In the first of the three exercises, I've used examples from two artists to reflect on how we come to see intelligence in machines. My interests here are in the active role of the machine and how, under some circumstances (and not others), an intelligence is *made* seeable.

The *Senster* is a sculpture I've repeatedly returned to in my thinking about intelligent machines. On show in Eindhoven between 1970 to 1974, Senster was one of the earliest computer-controlled, robotic sculptures to be publicly exhibited (Fig. 1). Built by the late sculptor Edward Ihnatowicz, Senster comprised a "sensing head" attached to a 4 meter-long, articulated aluminium "arm", manipulated using electro-hydraulics. The sculpture sensed sound and movement through microphones and radar receivers mounted on the sensing head. Persistent sound below a certain threshold would lead the head and then arm to move towards the apparent source, whereas loud noises or sudden movement would make them recoil.[3]

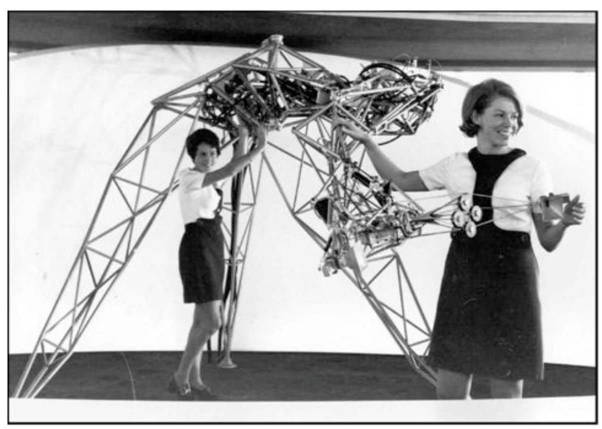

Figure 1. Senster (with permission from Alex Zivanovic).

What I find particularly compelling about Senster is its evocative motion. Ihnatowicz was either incredibly successful or very fortunate to capture something that imbued the

[3] A low-resolution but evocative video of an audience interacting with Senster is available online: <http://www.senster.com/ihnatowicz/senster/sensterphotos/senster_long.mpg.>

structure with some sense of it thinking. Despite its primitive and limited mechanical movements, many of those who actually saw the piece exhibited were struck and apparently often mesmerised by its life-like quality. As Ihnatowicz [15] subsequently recounted:

> The atmosphere around the sculpture was much more like that of a zoo than an art exhibition and I am sure that the majority of visitors would have been surprised to learn that the constructor of this machine had any pretensions to being an artist. p. 32

Putting to one side his self-professed bias towards movement, Ihnatowicz raises a particularly interesting way of thinking about intelligence. What his reflections on Senster helped him articulate is how intelligence is something seen in particular doings in the world—in Senster's case through its visible actions and reactions to a museum audience. It's quite possible that Ihnatowicz was drawing on philosophical discussions on the theory of mind that were prominent at the time and specifically Gilbert Ryle's thoughts on intelligence. Whatever the case, aspects raised in Senster's design and this philosophical perspective have some commonalities. In his seminal text, Ryle [29] eschewed the "ghostly" processes supposedly a prerequisite of intelligent behaviour, and suggested people (and we might surmise machines) are judged for their intelligence in what they do, and how they are seen to do it:

> To find that most people have minds… is simply to find that they are able and prone to do certain sorts of things, and this we do by witnessing the sorts of things they do. Indeed we do not merely discover that there are minds; we discover what specific qualities of intellect and character particular people have. (p. 59)

Both Ryle and Ihnatowicz thus drew attention to the ordinary ways in which intelligence is displayed and seen *in the world*. Ryle contested any idea of an "essence" of intelligence, and maintained it to be observable in and through everyday, "overt doings". Ihnatowicz showed through Senster how such mastery can be made visible through motion by imbuing the structure with an apparent sensitivity to its changing surroundings and its own movements.

Figure 2. *Motores, aluminio, zapato* (with permission from artist).

A playful contemporary project from the artist Adriana Salazar is suggestive of something similar, although through somewhat different, less extravagant means. Her playful range of mechanically driven objects are designed to repeatedly perform distinctively human behaviours such as tying and untying shoelaces (Fig. 2) or inhaling and exhaling smoke from a cigarette. Isolating the behaviours and reducing them to mechanical movements, Salazar cleverly juxtaposes ordinarily purposeful accomplishments with a complete detachment from the surrounding environment. Her choice of behaviours, ever-so mundane, give emphasis to the contrast; making actions like lace tying or smoking visibly mechanical, emphases their separation from their bodily performances in the world. Her mechanical objects are the epitome of mindless repetition, imbuing a certain absurdity through a blindness to the world. In common with Ihnatowicz (and arguably Ryle), what Salazar's work expresses is how there might be a perhaps counterintuitive way for us to think about independent motion and intelligence. It seems it isn't an independence from the material environment that has us imagining intelligence in things. Salazar's objects are, after all, moving independently. Rather the independence or *autonomy*—to use the AI/robotics jargon—of movement must be seen in some discernible way to relate to the world and to be bound to it. This is interesting because it appears to place constraints on autonomy as one of the tropes of AI and robotics. Machines that move independently appear curious to us, but not necessarily thinking. Our *seeing*, if you will, of intelligence in machines doesn't appear to be purely a consequence of a machine's autonomy of movement, but to do, equally, with its autonomous interactions in and with the world around it.

Emergent intelligence

An interesting proposition for intelligent machines arises when we consider the way people anthropomorphise and attribute a degree of agency to inanimate objects. To develop a line of thinking around this, I want first to focus on the general issue of attribution of agency and then to consider, more closely, intelligence.

A study using a unique tack to examine anthropomorphism (and offering an appealing contrast to all-out robotic research) was undertaken by the anthropologist Vidal [36]. With his son in tow, Vidal observed people interacting with the wax figures at London's Madam Tussauds. His own interest was in reflecting on the uncanny valley hypothesis [24] that postulates people are unsettled, even repulsed, by humanoid robots close to humans in appearance and movement, but somehow not quite right, somehow uncanny. Vidal used his observations at Madame Tussauds to demonstrate that, with inanimate figures at least, people are quite ready to engage with the uncanny. Despite their waxy appearance (not dissimilar to humanoid robots) and obvious absence of 'natural' movement, the wax figures are repeatedly lent on, touched, groped, etc.—something evident in the photos taken and displayed on sites like flickr (Fig. 3). Vidal argues that the figures (and presumably the place itself) invite these interminglings, offering the punters, if you will, a very particular range of interactions. What he exemplifies is the willingness people have to anthropomorphise, to treat in a fashion the inanimate as real (even if only momentarily). He develops his position to suggest that visitors actively parody the characters the figures represent. The pictures taken are of people doing the unbelievable: kissing Marilyn, saluting Hitler and so on.

I think Vidal's points are interesting because they draw attention to how very particular qualities can be ascribed to objects under specific conditions. In Madame Tussauds, we see that there are kinds of activities that allow the wax figures to be treated as having a kind of intent. Key here is that this intent is not conjured up in any real sense, but rather that there is a willingness on the part of Madame Tussauds' visitors to attribute a limited set of characteristics to the figures so that they are made to 'work' in the context of an exhibit. Perhaps inadvertently, Vidal's ideas thus have us ask in what ways we are willing to ascribe or attribute lifelike qualities to things, if only for a moment. It appears it is not a matter of all or nothing, but instead that people are willing to ascribe lifelike qualities in different and sometimes partial ways.

Figure 3. Woman and wax figure at Madame Tussauds (flickr.com).

To return to the issue of intelligence, let me refer back to a paper I noted earlier in which Forlizzi and DiSalvo reported on several households' purchase and use of iRobot's Roomba vacuum cleaner. In the early stages of their paper the authors recount the disappointment households had with the robotic vacuums. They relay the expectation participants had of the vacuum learning about its environment, 'remembering' the distinctive physical geographies of the homes they were used in. From this, a lack of intelligence was attributed to the Roombas. One householder pronounced her machine to be "dumb", if a little "pathetic". Leaving for the moment the apparent preparedness of the participants to anthropomorphise, what the reader discovers in the latter stages of the Forlizzi and DiSalvo paper is how the Roombas were worked on by households, coaxed as it were, to get the cleaning done. Not only were obstacles lifted and moved for Roomba, they were also employed to orient the machine and cordon off areas to be vacuumed. Interestingly, it is through this orchestration between user, machine and environment that something unexpected happens: it appears some of the households begin to construe an inkling of intelligence in their robotic vacuums. Having named their Roomba *Manuel*, after the incompetent butler in *Faulty Towers*, a woman in one household explains:

> "We named it because it has a personality, I mean well, it's doing the work of a person may be a part of it, and it seems to be sort of intelligent, has a little bit of intelligence in it." (p. 263)

Of course, I'd be wrong to attribute too much to the no doubt passing comments of a participant. I think though that they (and the points from the previous example) draw attention to two issues worth considering. First, the examples illustrate how our relationships with things are engineered, and actively so as our interactions with them unfold. In Madame Tussauds, visitors mock the realism of the figures by playfully engaging with them as exhibits on the one hand and historical/contemporary personalities on the other. A pose is struck just long enough and well enough for the parodying to work. Indeed, the parody is reliant on both seeing the staged pose and recognising it as implausible at one and the same time. With the Roomba, the relationship is engineered over time. Roomba owners find themselves taming the machine, so to speak, by controlling the environment in which it operates. The taming is not one purely of human over machine, however. The achievement demands that the person learns to see how the machine works and, in turn, for the machine to respond to the changing conditions.

A second issue adds an interesting dimension to this interplay. What we begin to see are circumstances under which particular lifelike qualities can be worked-up in things. In Madame Tussauds, Vidal demonstrates how, in regularly performed stagings by visitors, the wax figures are attributed with an intent not immediately apparent in their mere physical form. In the case of the Roomba vacuum cleaners, an intelligence emerges. As Forlizzi and DiSalvo report, a 'little bit of intelligence' comes about in working out the machine's methods of operation and then orchestrating the conditions under which it does its work. It is not just an object's capacity to remember or learn that allows us to attribute intelligence. There appear to be other characteristics such as the emergence of a relationship.

For the purposes of this paper, I want to highlight this apparent capacity we have to actively construe a kind of intelligence in things, and to do so in what seems a largely matter-of-fact way. To my mind, this has particular relevance to intelligent machines because it points to an openness in how we actively engage with things and mutually construct ideas of intelligence. Yes intelligence is seen in the world, but it is also enacted in it through the ordinary, unremarkable relationships we have with things. Strict ideas of intelligence are not rigidly adhered to, but rather fluid and configured in our daily doings. So, yes, the Roomba lacks intelligence because it cannot learn. Intelligence of a different sort appears to emerge, however, as a product of the coordinated efforts of both the person and the machine.

New Assemblies

In this third and final thought-piece, I want to touch on two areas of research that I've come across in searching for machines that deviate from human form. Together the two areas will hopefully illustrate how there are ongoing research programmes that raise questions about what constitutes intelligence in machines.

The first area of research I'd like to consider is one associated with robots that are able to produce and supply their own power. Specifically, I want to reflect on a project based at the Bristol Robotics Lab (BRL) in the UK centred on the *EcoBot*, a robot that is able to independently power itself (what the BRL call *energy autonomy*). The project, that began in roughly 2002, has sought to build robots that use organic materials as a source of energy. For example, the first of the project's working robots, *EcoBot I* (Fig. 4a), used simple refined fuels (i.e., sugar) to produce power, whereas the subsequent *EcoBot II* and *III* have been able to utilise more complex raw materials such as fruit or flies.

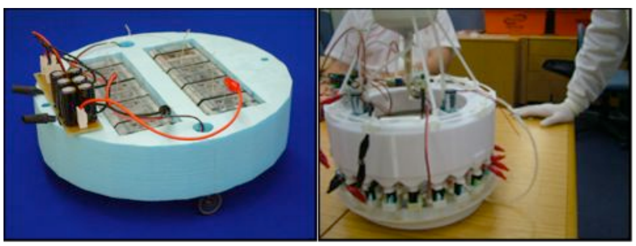

Figure 4. (a) EcoBot I and (b) EcoBot II (with permission from Bristol Robotics Lab).

The successive incarnations of the EcoBot have also seen developments in how different raw materials are converted into energy and then used. *EcoBot I* relied on sugar being manually supplied or 'fed' into an array of eight microbial fuel cells (MFCs) incorporated into the robot's foam body (Fig. 4a). The electric current produced using the cells was subsequently used to (very slowly and erratically) move the robot towards light (to demonstrate its capacity to perform a simple task). *EcoBot III* has not only enabled the use of raw organic materials, it also incorporates a synthetic digestive system that collects food (using a Venus flytrap like mechanism), channels it to the array of fuel cells, and either reuses or evacuates the waste (Fig. 4b).

One motivation for the EcoBot project has been to build machines that are autonomous not only in sensing and acting, but also in how they power themselves. For the purposes of building fully independent robots, this motivation is understandable. However, and as the project's members explain on their website, there are arguably more significant implications: incorporating the search for and consumption of 'food' into a robot's "behavioural repertoire" introduces "a paradigm shift in the way action selection mechanisms [in robots] have been designed so far." In other words, the prospect of robots locating their own food and consuming it, autonomously, has the potential to radically alter the design of robots and how they operate.

A corollary to this (and a point I will suggest has implications for HCI), is that this shift has the potential to transform how machines observably behave and interact in the world. If the Ecobot is anything to go by, these robots will emerge from curious hybrids where computation, electronics and material—both naturally occurring and fabricated—intermingle in novel ways. Indeed, what strikes me about *EcoBot III* as such an assemblage is how different it is from anything we now know as robotic. Although it is loosely composed of all the necessary ingredients, it substantially diverges from the much touted humanoid robot. Instead, it suggests strange and new intersections between bio-based matter, machines and ourselves. The EcoBot's 'feeding cycle', for instance, demands that one tend to it: its sludge must be kept alive, its tubes free of blockages, and its evacuated waste tidied. Moreover, its mechanical outputs (e.g., movement) occur in short bursts to allow sufficient charge to be accumulated via the MFC arrays. Admittedly, the EcoBot is only a proof of concept and any workable solution is likely to take a very different form. Nevertheless, as it stands, the robot expands our imaginations to contemplate something altogether different *vis-á-vis* machines.

Returning to the themes developed so far, I want to suggest that the curious assembly of parts in EcoBot offers up an opportunity to begin re-thinking intelligence. EcoBot, for instance, begins to impinge on the kinds of issues I raised in discussing Senster. It introduces, however, quite a different idea of a machine having a sense of itself. No longer is our attention drawn, solely, to physical movement; the machine's sense of its own energy consumption and expenditure is also brought into play. Similarly, as with Roomba, the kinds of interactions suggested in the EcoBot's design open up the possibilities for different enactments of intelligence; because of its bio-based digestive system, there is the inkling of the EcoBot incorporating an autonomy that is not programmed into it, *per se*, but is, in a sense, a 'naturally occurring' phenomenon (i.e., organic decomposition).

My point here is not to make any strong claims about the intelligence or even lifelike qualities of a machine like EcoBot or its offspring. Indeed, I appreciate it takes some (if not a good deal of) imagination to see how the topic of intelligence relates to EcoBot as it stands. What I believe such examples help us to consider are the properties a machine might exhibit for us to begin wondering about intelligence or for us to construe some possibly new idea of it.

To expand on this point, let me move onto a second area of research that focuses on the self-regulated interactions between machines. The work I want to draw on here can be traced back to the ideas set out by Rodney Brooks and his colleagues. My interest though is in those primitive machines that might not be immediately thought of as intelligent. In a publication in this realm, Kaplan [17] investigates the ways large numbers of (virtual) agents can communicate their states to one another and how different methods of coordination can affect the agents' collective behaviours.

My primary reason for using Kaplan's example is to illustrate how machines, in principle, have the capacity not only for independent behaviour, but also to develop self-regulated, co-ordinated interaction. What we find is just how little it takes to get groups of machines to talk amongst themselves, as it were, and to behave in a co-ordinated fashion. Indeed, Kaplan states an interest in how such primitive models scale in terms of language processes and specifically the emergence of language between agents. It's this idea of machines having the capacity to organize their co-ordination, no matter how simply, that intrigues me. I don't want to get too worried at this stage about what is meant by language. The interesting proposition is the enabling of a primitive machine 'language' that allows for what appears to be reasonably sophisticated machine interactions.

As well as building the EcoBot, the Bristol Robotics Lab have taken up just this agenda by proposing a research programme investigating what they have termed, perhaps over zealously, 'robot societies'. Their efforts have so far been limited to co-ordinating the motion of robotic devices and having robots collectively perform simple tasks. As with EcoBot, however, the specific technical workings of such a system are not the only issues of significance at this stage. As I see it, a fruitful line of inquiry might be one that examines how we are able to see something like language arise when we bring together a simple set of algorithms, a primitive mechanical apparatus, an observable collection of interactions, and so on. We might also go onto ask under what circumstances machines that communicate with one another—using their own learnt vernacular—could be construed as indicative of a kind of intelligence. Like Roomba, given the right context, might we be willing to learn to see intelligence in such 'performances'?

I've used these last two examples, then, to build on the points made earlier. I've hopefully reiterated that intelligence is something 'seeable', but also something enacted—emerging from those particular details of a setting. Beyond this, what I've tried to emphasise is how the developments in computing technology (including AI and Robotics) have the capacity to dramatically transform the different ways of seeing and enacting intelligence in things. It's not that the examples I've offered succeed in this transformation. I believe, though, that they begin to be suggestive of something; they hint at the possibility of different ways of imagining intelligence in machines and thus very different kinds of interactions with them. Indeed, to return to Suchman's thesis [32], they invite us to wonder what it is to be human:

> Not only do these experiments [with new assemblies] promise innovations in our thinking about machines, but they open up as well the equally exciting prospect of alternate conceptualizations of what it means to be human. (p. 281)

Implications for Design

Mulling over the kinds of innovations discussed so far, as well as many others, it occurs to me what exciting times we are in. Even judged superficially, the realisation of robotic vacuum cleaners, ecobots, robot collectives and so on reiterate how narrow my criticisms of technology have been and how closed I've been to alternative ideas of intelligence, especially in machines.

Above, then, I've drawn out two general lessons from AI to expand on how we might think of machines and the kinds of intelligence we encounter when we interact with them. As a reminder, the first of these lessons ties back to a shift in thinking about intelligence that has taken place in AI, a shift characterised by the pragmatic rejection of any strong position for or against *a priori* representation, and in some cases by work on hybrid systems. In the examples above, I've tried to show how developments related to this shift continue to disrupt our ideas of machines and what they are capable of. The second lesson builds on the largely constructive critique of AI that presents the possibility of new and emergent intersections between humans and machines. This critique underlies my above efforts to reconsider where intelligence resides in machines when we interact with them. In my closing remarks, I want to consider how these positions could be pursued in the context of HCI. To do so, I've drawn out three broad points tying back to the exercises above.

A sensitising to non-human intelligence—Some of the examples above indicate, I hope, how intelligence can be seeable in machines under certain circumstances. They suggest, as well, that a machine's intelligence is not only judged by how well it measures up against human talk, movement, expression or even thought. The ordinary ways intelligence is used appear less restrictive than that.

This raises an opportunity for designing human-computer interactions. It opens up the possibility for designing interactions that might locate machines in curiously new seeable categories of intelligence (as well as familiar ones). Reflecting on an interaction design of this kind, I also find myself asking what, exactly, are the interactional circumstances—to coin an awkward phrase—under which things are seen as intelligent? So we might wonder whether there are certain qualities of interaction, situated in one place or another, that help us to see different kinds of intelligence.

Taking these possibilities seriously in systems design might offer quite different ideas of what to design for and what resources to design with. To take an example, what if we were to make more of actuators in designing human-machine interactions? The tendency has been to build systems that simulate human or animal movement. Could meaningful relations be struck, however, with different kinds of autonomously moving machines and under what circumstances might they begin to be seen as intelligent? More generally, could the ways intelligence is seen in things (and not just humans) become a resource for design?

Designing an emergent intelligence—Above, I've tried to illustrate how different ideas of intelligence are enacted through our interactions with things. As well as seeing an intelligence in things, we routinely fashion our ideas of what intelligence is. As we've seen, Roomba presents a nice example of this because it shows how a machine intelligence, of sorts, can emerge through ongoing and wholly ordinary interactions with the vacuuming robot.

An interesting problem for HCI is whether such enactments can be designed for and what they might look like in practice. Something I see to be of particular interest is whether design might be used to provoke new human-machine interactions where unexpected ideas of intelligence arise. In Forlizzi and Disalvo's paper, it's no coincidence that the families resorted to seeing their Roombas as butlers or servants (incompetent or otherwise); arguably, the tensions around servitude or control versus autonomy are the most predictable when it comes to (domestic) robot design. What else is there though between humans and machines?

Returning to the use of actuators, Senster is instructive. As noted, intelligence seemed to come about through the audience seeing an apparent mastery over graceful movement and discovering a certain sympathy with the sculpture. In HCI, could we investigate a similar range of interrelations with machines and how intelligence might be construed, differently, through them? Might we consider, for example, how intelligence is manifest by designing acts of corroboration, mutual learning, sympathy, or compassion?

New assemblies, new intersections—In the final exercise above, I presented research that combines the mechanical, computational, biological, social, etc. in novel ways. I've not wanted to suggest these assemblies are in any way more intelligent than conventional machines. However, as I see it, they further introduce the possibility of new human-machine intersections and thus potentially new ways of seeing and enacting machine intelligence.

As I've said, descriptive accounts of such reconfigurations and transformations abound in the studies of science and technology literature. Relevant for HCI is what role the practical business of design has in this figuration business. To my mind, and if one is to take its title seriously, HCI has a responsibility for understanding how the intersections change with emerging technologies like the above. Might HCI thus aim to purposefully re-design the human-machine intersection and consider how, if at all, these purposeful interventions transform everyday ideas of intelligence?

HCI could begin to reflect on how the design of hybrid assemblies reconfigures the human-machine intersection. With respect to intelligence, questions arise around the different kinds of intelligence introduced by innovative assemblies. If it is an autonomy of 'life' (in technologies such as Ecobot), or the capacity for collective reasoning and action that have us imagine an intelligence in machines, then what would it mean to design for and expose these characteristics in human-computer interactions and, indeed, expose the fragilities that come along with them?

Conclusion

Overall, the issues discussed in this paper can be summarised using four overarching themes. First, as a general theme, I've made a case for rethinking intelligence as both a topic of inquiry and a resource for design in HCI. Second, I've shown that the ways intelligence is seen and enacted in things can profoundly influence the interactions between human and machine. Third, I've proposed that various innovations in computing machines introduce new ways of thinking about intelligence and, consequently, challenge us to reconsider the human-machine interface. Fourth, I've suggested that reflecting on intelligence in these ways opens up a new and, perhaps, radically different set of possibilities for interactive design.

These broad points are undoubtedly difficult ones to address, particularly in such abstract terms. Not only do they raise some challenging questions concerning how we think of our practice and where it is headed; somehow, we must assess whether they help us in practical terms—in building things relevant to people's everyday needs, experiences and values. With these concerns in mind, my recent efforts—with much help from others—have included attempts to explore this area in more concrete ways. Building different probes or prototypes (some working, some not), we have sought to introduce aspects of machine behaviour that might be unfamiliar or even disruptive. Our emphasis has been on provoking questions around where intelligence is seen to reside in new assemblies and to what extent people are able to re-imagine intelligence in their machine interactions. There is too little space to go into the detail of each of these projects. My overarching hope has been, however, to use this work to start asking what design has got to do with the interactions between people and the kinds of machines they treat, casually, as intelligent.

In conclusion, what I hope to have shown above is the value of reflecting on intelligence and its role in human-computer interaction. By continuing to consider the prospect of intelligent machines, my hope is that I and possibly others in HCI might start to develop an informed way of thinking about the subject. The above is merely a partial attempt to begin this process—to engage in a re-thinking of 'artificial intelligence' (in the loosest sense of the phrase) and the idea of intelligence more broadly. Perhaps this will allow some of us to reach a point where we can imagine interactions between people and very different kinds of machines, and where we could offer some useful answers to what it might be like when intelligent machines are an ordinary, unremarkable feature of everyday (inter-)action.

ACKNOWLEDGEMENTS

I am indebted to many friends and colleagues for reading and commenting on the numerous drafts of this paper. Thanks must go to the members of my group at MSR, especially Richard Harper, Shahram Izadi, Dave Kirk and Abi Sellen. There are, as well, others from further afield including Barry Brown, Beki Grinter, Anab Jain, Lucian Leahu, Sara Ljungblad, Susanne Seitinger, Phoebe Sengers and Amanda Windle. Finally, Laurel Swan deserves special thanks for pointing me in the direction of robotics/AI research and in acting as a patient sounding board.

REFERENCES

1. Agre, P. & Chapman, D. What are plans for? In Maes, P. (Ed.), *Designing Autonomous Agents.* MIT Press, Cambridge, MA, 1990, 17-34.
2. Breazeal, C. Emotion and sociable humanoid robots. *Int'l Jn'l Hum. Comp. Std. 59*, 1-2 (2003), 119-155.
3. Brooks, R. A. Intelligence without representation. *Artificial Intelligence. 47*, 1-3 (1991), 139-159.
4. Brooks, R. A. *Flesh and Machines: How Robots Will Change Us.* Pantheon, New York, 2002.
5. Clark, A. & Chalmers, D. J. The extended mind. *Analysis. 58* (1998), 7–19.
6. Dreyfus, H. L. Why Heideggerian AI failed and how fixing it would require making it more Heideggerian. *Artificial Intelligence. 171*, 18 (2007), 1137-1160.
7. Edwards, K. & Grinter, R. At home with ubiquitous computing: seven challenges. In *Proc. Ubicomp '01*, Springer (2001), 256-272.
8. Forlizzi, J. & DiSalvo, C. Service robots in the domestic environment: a study of the Roomba vacuum in the home. In *Proc. HRI '06*, ACM (2006), 258-265.
9. Gaver, W., Sengers, P., Kerridge, T., Kaye, J., and Bowers, J. Enhancing ubiquitous computing with user interpretation: field testing the home health horoscope. In *Proc. CHI '07*, ACM (2007), 537-546.
10. Hanfling, O. *Philosophy and Ordinary Language.* Routledge, New York, NY, 2000.
11. Haraway, D. J. *Simians, Cyborgs, and Women: The Reinvention of Nature.* Routledge, New York, 1991.
12. Haugeland, J. *Artificial Intelligence: The Very Idea.* MIT Press, Cambridge, MA, 1985.
13. Hayles, N. K. *How We Became Posthuman.* University of Chicago Press, Chicago, Ill, 1999.
14. Hutchins, E. *Cognition in the Wild.* MIT Press, Cambridge, 1995.
15. Ihnatowicz, E. Towards a thinking machine. In Leavitt, R. (Ed.), *Artist & Computer.* Harmony, 1976, 32-34.
16. Kahn, P. H., Ishiguro, H., Friedman, B., & Kanda, T. What is a human? Toward psychological benchmarks in the field of human-robot interaction. In *Proc. RO-MAN '06*, IEEE (2006), 364-371.
17. Kaplan, F. & Oudeyer, P.-Y. Intrinsically Motivated Machines. In Lungarella, M., Iida, F., Bongard, J., & Pfeifer, R. (Eds.), *Intrinsically Motivated Machines.* Springer-Verlag, Berlin, 2007, 304–315.
18. Latour, B. *We Have Never Been Modern.* Harvester Wheatsheaf, Hemell Hempstead, 1993.
19. Law, J. & Hassard, J. *Actor Network Theory and After.* Wiley-Blackwell, 1999.
20. Leahu, L., Sengers, P., & Mateas, M. Interactionist AI and the promise of ubicomp, or, how to put your box in the world without putting the world in your box. In *Proc. Ubicomp '08*, (forthcoming).
21. Lighthill, J. *Artificial Intelligence: A General Survey, in Artificial Intelligence.* Science Research Council, a paper symposium, 1973.
22. Ljungblad, S. & Holmquist, L. E. Transfer Scenarios: Grounding innovation with marginal practices. In *Proc. CHI '07*, ACM (2007), 737-746.
23. Mateas, M. Expressive AI: A hybrid art and science practice. *Leonardo. 34*, 2 (2001), 147-153.
24. Mori, M. (Trans. MacDorman, K. F. and Minato, T.). The Uncanny Valley. *Energy. 7*, 4 (1970), 33-35.
25. Newell, A. & Simon, H. A. GPS: A Program that simulates human thought. In Feigenbaum, E. A. & Fielman, J. (Eds.), *Computers & Thought.* McGraw-Hill, New York, 1963.
26. Norman, D. & Dunaeff, T. *Things That Make Us Smart.* Basic Books, 1994.
27. Oudeyer, P. Y., Kaplan, F., & Hafner, V. V. Intrinsic motivation systems for autonomous mental development. *IEEE Trans on Evolutionary Computation. 11*, 1 (2007), 265-286.
28. Pousman, Z., Romero, M., Smith, A., & Mateas, M. Living with tableau machine: a longitudinal investigation of a curious domestic intelligence. In *Proc. UbiComp '08*, ACM (2008), 370-379.
29. Ryle, G. *The Concept of Mind.* Peregrine Books, Harmondsworth, UK, 1949.
30. Sengers, P. The autonomous agency of STS: Boundary crossings between STS and Artificial Intelligence. *Soc. Stud. of Science.* (forthcoming).
31. Shanker, S. *Wittgenstein's Remarks on the Foundations of AI.* Routledge, Abington, 1998.
32. Suchman, L. A. *Human-machine reconfigurations: Plans and situated actions* (2nd ed.). Cambridge University Press, Cambridge, 2007.
33. Taylor, A. S., Harper, R., Swan, L., Izadi, S., Sellen, A., & Perry, M. Homes that make us smart. *Pers. & Ubi. Comp. 11*, 5 (2007), 383-393.
34. Turing, A. M. Computing machinery and intelligence. *Mind. LIX*, 236 (1950), 433-460.
35. van Gelder, T. J. The dynamical hypothesis in cognitive science. *Behavioral & Brain Sciences. 21* (1998), 1-14.
36. Vidal, D. Wax models and robotics: another approach to anthropomorphism. In *Proc. Practices of anthropomorphism: From automata to robotics*, Maison Française d'Oxford (2007).
37. Yamazaki, K., Kawshima, M., Kuno, Y., Akiya, N., Burdelski, M., Yamazaki, A. *et al*. Prior-to-request and request behaviors within elderly daycare: Implications for developing service robots for use in multiparty settings. In *Proc. ECSCW '07*, Springer (2007), 61-78.