Social Machines as an Approach to Group Privacy

Kieron O'Hara

Web and Internet Science Group
Electronics and Computer Science
University of Southampton
Highfield
Southampton SO17 1BJ
United Kingdom
kmo@ecs.soton.ac.uk

Introduction

Group privacy is an interesting topic made more salient in recent years by the growth of big data, enabling people to be targeted and understood via their personal attributes (on the basis of correlations between the group of people who possess those attributes and other phenomena), or alternatively via the properties of their networks (for instance, whether one is likely to default on a loan is correlated with one's social network). In each case, the dilemmas of privacy are thrown into sharp relief – visibility to one's network brings benefits, but compromises privacy. Furthermore, there is a distinct potential for injustice, as one may find oneself discriminated against on the basis of behaviour of *other* people in one's groups (Hildebrandt 2012).

We should separate out the injustice from the privacy, but there is a *prima facie* case for arguing that we could nip the injustice in the bud if groups as well as individuals had privacy rights. In the age of big data, data crunchers are not interested in the individual data points, so much as the mass – yet the crunching of data about the mass can have real-world implications for individuals.

One way of understanding this is to see it as one of the many ways in which data protection is an imperfect protection for privacy (O'Hara 2011, 7-11). Data protection requires an individual to be identifiable before data is classified as personal data so that it can only be processed with the subject's consent (of course there are many exceptions to this built into data protection legislation). Yet the notion of group privacy is consistent with something that we intuitively understand in the age of spam and junk mail – one need not be identified to have one's privacy invaded. The mere existence of a non-identifying profile of oneself, combined with a point of access such as an address, does not count as personal data, but is still an annoying invasion.

Group privacy as derivative?

However, couching the problem in this way still makes the privacy of the group derivative from the privacy of the individual – the individual's remedy for the invasion of his personal privacy is to insist on the privacy of a wider group of which he is an anonymous member. This seems to chime in with liberal ideas about privacy. The point of privacy, according to one influential analysis, is to support individual autonomy (Rössler 2005). Meanwhile, intrusion from the group itself has, since Mill, been seen as an serious threat to the individual (Mill 1859). And a number of theorists, for example feminists, have tended to see the privacy of small units (from the family up) as a means of concealing abuse, rather than of legitimately supporting the individual (MacKinnon 1989, 168, Allen 2003).

In a tradition that postulates the small group as the major threat to the individual, group privacy does not look like a serious runner, *unless* the group can be reconceptualised as an important support for the individual's autonomy – and so group privacy seems to derive its value from the needs of the individual, not the group itself.

Groups for social control

On the other hand, a conservative viewpoint is more ambivalent about the power and potential of an individual – for instance, Burke lauds the little platoons, and considers the individual as intrinsically unable to make consistent or wise moral judgments. Schoeman argues, *contra* Mill, that social control, far from being morally destructive, is an important factor in a valuable liberty. Our competence as rational agents depends on constructive adaptations of social control mechanisms in real-world contexts. Unpicking informal social control mechanisms in the name of autonomy, in Schoeman's view, actually deprives the individual of important social abilities, and "helps maintain both the integrity of intimate spheres as against more public spheres and the integrity of various public spheres in relation to one another" (Schoeman 1992, 157).

Adam Smith's view, with regard to the moral education of people in the newly emerging metropolises of the eighteenth century, is an interesting example of this kind of thought.

> A man of low condition, on the contrary, is far from being a distinguished member of any great society. While he remains in a country village his conduct may be attended to, and he may be obliged to attend to it himself. In this situation, and in this situation only, he may have what is called a character to lose. But as soon as he comes into the great city, he is sunk in obscurity and darkness. His conduct is observed and attended to by nobody, and he is therefore very likely to neglect it himself, and to abandon himself to every sort of profligacy and vice. (Smith 1994, vol.2, V.i.g.12, 795, footnote omitted)

The way to address this, thought Smith, was not more policing or the reduction of the private sphere of the 'man of low condition', but rather greater power for groups, specifically those that have an interest in the individual's moral conduct.

> He never emerges so effectually from this obscurity, his conduct never excited so much the attention of any respectable society, as his becoming the member of a small religious sect. He from that moment acquires a degree of consideration which he never had before. All his brother sectaries are, for the credit of the sect, interested to observe his conduct, and if he gives occasion to any scandal, if he deviates very much from those austere morals which they almost always require of one another, to punish him by what is always a very severe punishment, even where no civil effects attend it, expulsion or excommunication from the sect. In little religious sects, accordingly, the morals of the common people have been almost always remarkably regular and orderly; generally much more so than in the established church. (Smith 1994, vol.2, V.i.g.12, 795-6).

Complexity

As we consider these issues in the age of big data, it is worth addressing the issue of whether group privacy will create greater complexity in policing and vigilance, and whether a right will be created which would go beyond existing expectations and preferences, and the needs of democratic societies. Individual privacy introduces a number of private spaces proportional (of course) to the number of citizens, whereas group privacy will be a correspondingly complex concept to enforce.

When we consider group privacy, if we think about the number of groups that people are willing to admit they are members of, and whose corporate privacy they wish to defend, the extra complexity probably grows in a linear fashion as population grows. On average, people might admit to membership of m groups (maybe m would be something between 10 and 100), while average membership of a group would be n people. Hence, for a population of x, the number of groups to be protected would be proportional to mx/n.

However, big data will change this. The point of big data is that data mining finds significance in correlations within groups that have no external significance – one might easily not know, or care, that one was a member of such a group (like, for instance, 26-35 year old males earning between £40k-£50k p.a. in households without children who have downloaded more than 5 unsolicited recommendations from iTunes in the last six months). Of course, for a population of x, the number of such potential groups is $2^x - 1$, but as big data crunchers do not consider the coherence or independent interest of such groups, it would be hard to single out which groups are worth protecting. This could create an extremely complex and difficult legal scene, with hard decisions to make about liability and the balance between social good and protection of rights.

So it is likely to be impractical to consider theoretically possible groups, whose number will grow exponentially with the population. The monitoring and policing of group privacy can more easily be kept tractable if we take into account those groups that individuals expressly understand themselves to be members of. In that sense, group privacy will remain derivative from individual privacy, but crucially in this case the value of the group's privacy is decoupled from the individual's privacy. In the conservative view of the world developed by Smith, the group's interests may actually precede those of the individual.

For example, consider what Nancy Rosenblum has termed the 'logic of congruence'. Participation in groups helps cultivate certain values and virtues in the members. Which ones are cultivated depends somewhat on the nature of the group in question. Membership tends to create individuals who are predisposed to internalise, uphold and perpetuate the values and virtues of that environment. Smith believed that this was inherently valuable.

Given that view, it may make sense to look at group privacy as a means of empowering the group to achieve its aims, and to see its protection as a means of institutionalising that empowerment. Of course this does not resolve the ethical question of when that is a good thing and when bad, but at least it gives us a rationale to do it. The final question I will consider in this paper is whether current conceptualisations of technology give us a handle on group empowerment. With that in mind, it may be worth exploring research into *social machines*.

Social machines

The world of big data has not, of course, been unaccompanied by other developments. In particular, as the amount of data that it is feasible to process has grown, so has the number of people that it is feasible to connect within a network. Figure 1, following David De Roure, gives a sense of different interaction modes of computing. Wherever there are more machines, to produce the big data paradigm at upper left, or more people, as in the social networking paradigm at lower right, distribution is inevitable, and hence Web or Web-like technologies are necessary to handle interaction at scale. The technological affordances have, over time, moved upward and toward the right, ultimately to reach the fourth quadrant.

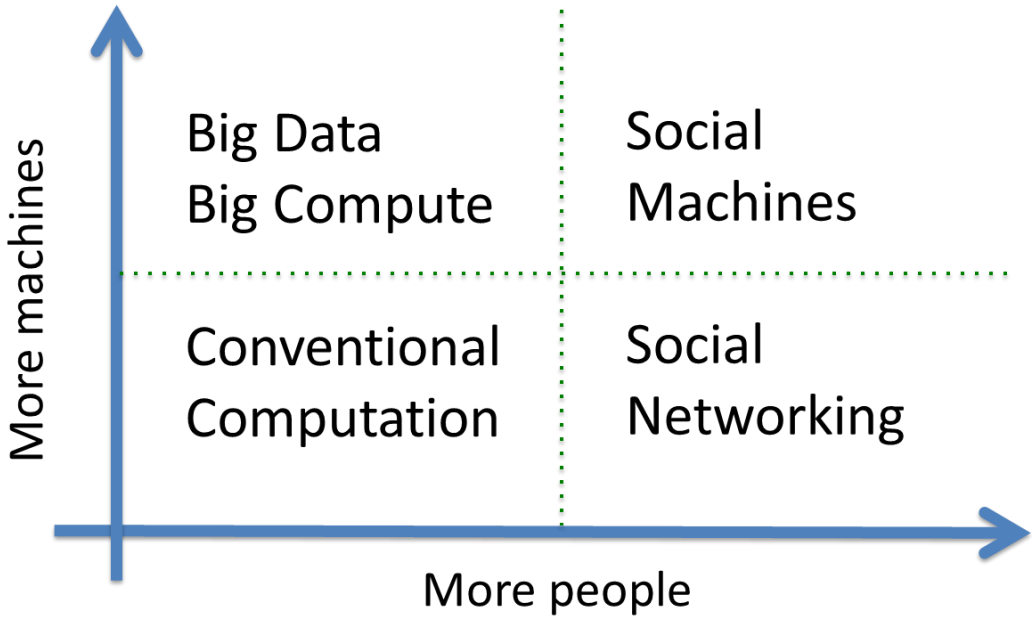

Figure 1: A matrix showing the affordances of scale (adapted from De Roure 2013)

This fourth quadrant is termed 'social machines' (Berners-Lee 1999, Hendler & Berners-Lee 2010, Shadbolt et al 2013, O'Hara et al 2013), which are a nascent focus of computing research (Bernstein et al 2012). 'Programming the global computer' or 'global ubiquitous computing' has been recognised as a grand challenge for computing (Kwiatkowska et al 2004), while peer-to-peer technologies flexibly link people and computers, as explored in projects such as SOCIAM (http://sociam.org/), OpenKnowledge (http://www.openk.org/) and the Social Computer community (http://www.socialcomputer.eu/). As we unravel the mysteries of scale and control, we will need not just to understand the emergent phenomena, but to develop means, methods and tools for controlling large-scale phenomena, at least partially (O'Hara et al 2013). The problem is sharpened by the desideratum that 'programming the social computer' must be achievable from *within* the social computer – research here should democratise control by allowing people to develop social machines to achieve their own smaller-scale, local, idiosyncratic purposes.

If we unpack that image as Figure 2, we see the potential space for advancement in more detail. We see conventional computation, even highly complex domains such as air traffic control and climate modelling, on the left hand side, where social complexity is low even if computational complexity is high. Current systems with

high social complexity still involve relatively low computational complexity. Crowdsourcing systems, such as the citizen science initiative Galaxy Zoo (Lintott et al 2008) have a relatively low level of social complexity as well. More complex social arrangements can be found in the co-creation of content, e.g. Wikipedia, and social networking. However, greater complexity can be found, for example, when social network acts as platforms for crowdsourced co-creation of content, as recently happened with the Ushahidi map of election violence in Kenya in 2007 (Okolloh 2009), or the reuse of Ushahidi software to create a post-earthquake map of Port-au-Prince in Haiti in 2010 (Morrow et al 2011). As we explore this space of social computation, to address perceived issues where there are collective action problems, as with public health, transport or crime, we would expect to find solutions with small impacts locally which will be magnified at scale, as long as the requisite social infrastructure (including Web technologies) is in place.

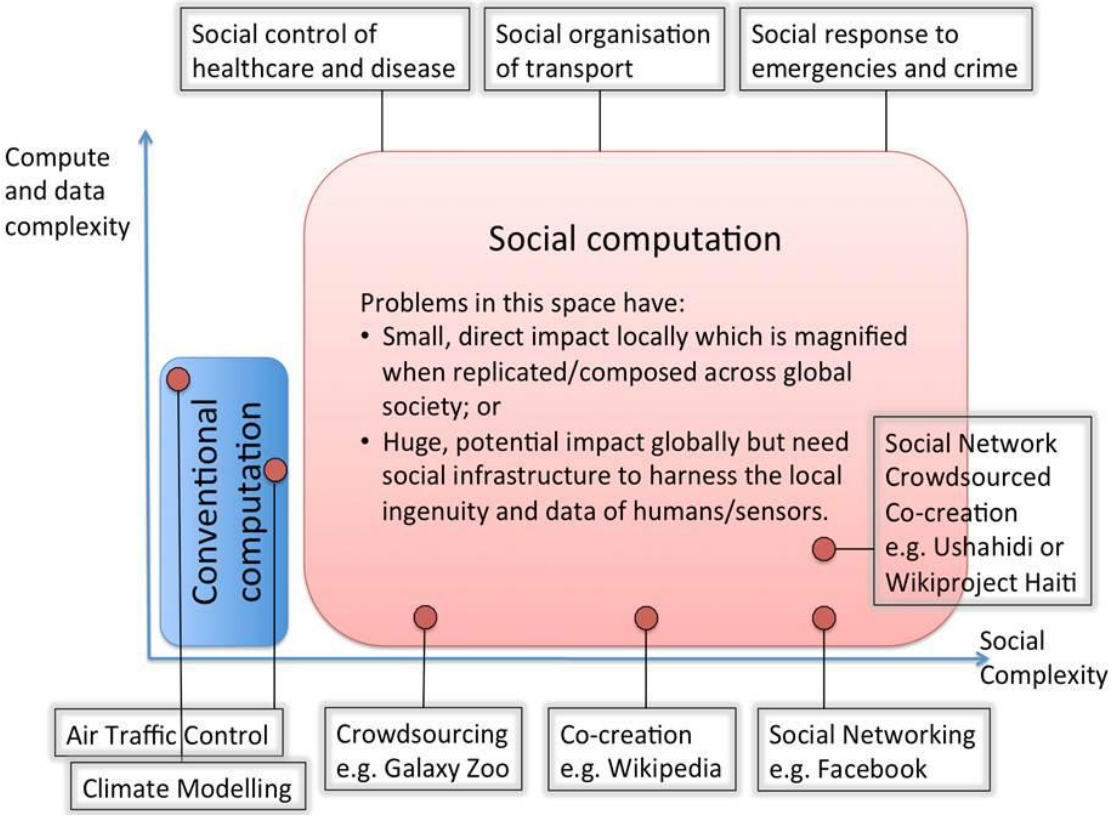

Figure 2: The space of social machines (O'Hara et al 2013)

The idea of a social machine has been implicit throughout the history of the Web. As Berners-Lee put it in 1999:

> Real life is and must be full of all kinds of social constraint – the very processes from which society arises. Computers can help if we use them to create abstract *social machines* on the Web: processes in which people do the creative work and the machine does the administration. (Berners-Lee 1999, 172, Berners-Lee's emphasis)

Many social machines are built on SNSs such as Facebook, in which human interactions from organising a birthday party to interacting with a Member of Parliament are underpinned by the engineered environment. Another type of example is a multiplayer online game, where a persistent online environment facilitates interactions concerning virtual resources between real people. A third type is an

online poker game, where the resources being played for are real-world, where the players may be human or bots, and where the environment in which the game takes place is engineered around a relatively simple computational model. In such systems, (some of) the social constraints that Berners-Lee talks about, currently norm-driven, are administered by the architecture of the programmed environment.

A generalised definition of a social computation is provided by (Robertson and Giunchiglia 2013):

> A computation for which an executable specification exists but the successful implementation of this specification depends upon computer mediated social interaction between the human actors in its implementation.

In such an environment, self-organisation (partial or full) becomes viable and scalable, while physical objects, agents, contracts, agreements, incentives and other objects can be referred to using URIs. 'Programming' the social computer (as opposed to simply supporting and directing interactions on an engineered environment) and integrating larger numbers of people and machines will become increasingly feasible.

As a small example of a social machine, consider reCAPTCHA (Von Ahn et al 2008). A CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart), invented by Louis Von Ahn, is the distorted sequence of letters that someone has to type in a box to identify him- or herself as a human (e.g. to buy a ticket online, or to comment on a blog). This is a task that computers cannot do, and so the system stops bots buying thousands of tickets for a concert or sporting event for later resale, or for a spambot to leave spam messages as comments to blogs (Von Ahn et al 2003).

Von Ahn extended the idea of the CAPTCHA to create the reCAPTCHA, which socialises the same principle to solve another problem. Google (which acquired reCAPTCHA in 2009) uses it to scan older books automatically. The original CAPTCHA device was being used over 200m times a day, about half a million person-hours of effort. reCAPTCHA puts these person-hours to more productive use, presenting the user who wishes to identify him- or herself as a human with two words, not one. The first is a normal CAPTCHA, and the second is a word from an old book that Optical Character Recognition had failed to identify. If the person succeeds with the first CAPTCHA, then he or she is known to be a human. As humans are reliable at word recognition, the response to the second word as a plausible suggestion of what it is. Presenting the same word to multiple users allows a consensus to emerge. The goal of the social machine is to digitise books – people's needs to prove themselves human provides the mechanism.

However, reCAPTCHA is purely exploitative, as the goal of the machine is independent of the requirements of its human 'components'. As another example, (Robertson and Giunchiglia 2013) use the DARPA balloon challenge of 2009, in which the aim was to find ten weather balloons placed randomly around the US (in nine different states from California to Delaware). The rules of the challenge were intended to support the growth of a network of people taking part in the search, enabling a crowdsourced solution. The means of doing this in the winning solution (from Sandy Pentland at the Massachusetts Institute of Technology) was to set out financial incentives according to a Query Incentive Network Model (Kleinberg & Raghavan 2005), in which people were incentivised both to look for the balloons and to add more people to the network. Pentland's team began with 4 people, and using

social media had recruited over 5,000 at the point of completion, which took under ten hours (Pickard et al 2010).

reCAPTCHA and the DARPA challenge were designed to solve a particular problem, but social machines can, and indeed should (O'Hara 2012), solve the problems of the people who constitute them. In such cases, the incentive of the participants is that the machine's smooth functioning is in their own interests. One could imagine, for instance, a set of computer-mediated interactions enabling a community to provide a social response to problems of crime (such as BlueServo, http://www.blueservo.net/, which crowdsources the policing of the Texas-Mexico border), or enabling those suffering from a particular health care problem to pool resources and to offer support and advice to fellow sufferers (such as curetogether.com, http://curetogether.com/). There is a growing number of health social machines, as surveyed in detail in (Van Kleek et al 2013). It will be obvious from these examples, particularly BlueServo, that such efforts will not always be uncontroversial. Attempts to crowdsource the identities of the bombers of the Boston Marathon in 2013 bordered on farce, and, although the countercultural website 4chan was prominent in the home-made policing efforts with its so-called '4chan Think Tank', its lamentable efforts were soon parodied elsewhere on the same site (Walker 2013). Trust will be a major factor in the success of such machines (O'Hara 2012).

Social machines as a way to approach group privacy

Suppose methods and tools were available to enable and empower communities to use data and networked communications to solve self-identified problems. In that event, the social machine would have certain functional requirements. The 'program' for the machine – i.e. the computations that it would carry out to transform the state of the world – can be written down in abstract terms, independent of whether the computing was being done by machines, people or groups of people.

Such an abstract specification will need to be 'filled in' by accounts of other key factors for the machine to function – for instance, the decisions by actors to engage with the machine, the knowledge that actors bring to the computation, the ontology of annotations for the interactions which render them understandable, data management methods, the incentives to participate, the restrictions on who can participate and so on.

One of the issues alluded to above is the actors' trust of the system, which will include their needs for privacy. Modelling a sociotechnical interaction in this way might enable the privacy requirements of participants to be specified. For instance, one might say that a particular computation $n(S_1,S_2)$, which transforms the world from state S_1 to state S_2 would only be possible if the actor(s) involved, who carry out n, can freely exchange personal data with each other, while ensuring that it does not spread beyond their circle. Or it might be that the actors require access to the personal data of all participants in the social machine (including of those not involved in that particular computational step). Or it might be that external services might be required which need access to personal data, or anonymised personal data, of participants in the social machine.

A statement in terms of the information-processing needs of the social machine would help make the demands made on information about the group as a whole explicit in terms of the goals which it wishes to achieve. In that way, we might find ways of explaining the functional value of privacy for a particular group, independently of

moral generalisations about group privacy rights – though not, of course, independently of the moral question of whether a group should be empowered to achieve its particular goals.

References

Anita L. Allen (2003). 'Privacy isn't everything: accountability as a personal and social good', *Alabama Law Review*, 54.

Tim Berners-Lee (1999). *Weaving the Web: the Original Design and Ultimate Destiny of the World Wide Web*, New York: HarperCollins.

Abraham Bernstein, Mark Klein & Thomas W. Malone (2012). 'Programming the global brain', *Communications of the ACM*, 55(5), 41-43.

David De Roure (2013). *Social Machines of Science*, powerpoint presentation, Bangalore: Infosys, https://dl.dropboxusercontent.com/u/15772302/SocialMachinesOfScience.pptx.

James Hendler & Tim Berners-Lee (2010). 'From Semantic Web to social machines: a research challenge for AI on the World Wide Web', *Artificial Intelligence*, 174(2), 156-161.

Mireille Hildebrandt (2012). 'The dawn of a critical transparency right for the profiling era', in Jacques Bus, Malcolm Crompton, Mireille Hildebrandt & George Metakides (eds.), *Digital Enlightenment Yearbook 2012*, Amsterdam: IOS Press, 41-56.

Jon Kleinberg & Prabhakar Raghavan (2005). 'Query incentive networks', in *Proceedings of the 46th Annual IEEE Symposium of Foundations of Computer Science (FOCS'05)*, Pittsburgh, 132-141.

Marta Kwiatkowska, Robin Milner & Vladimiro Sassone (2004). 'Science for global ubiquitous computing', *Bulletin of the European Association of Theoretical Computer Science*, 82, 325-333, http://eatcs.org/images/bulletin/beatcs82.pdf.

Chris J. Lintott, Kevin Schawinski, Anže Slosar, Kate Land, Steven Bamford, Daniel Thomas, M. Jordan Raddick, Robert C. Nichol, Alex Szalay, Dan Andreescu, Phil Murray & Jan Vandenberg (2008). 'Galaxy Zoo: morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey', *Monthly Notices of the Royal Astronomical Society*, 389(3), 1179-1189.

Catherine A. MacKinnon (1989). *Toward a Feminist Theory of the State*, Cambridge MA: Harvard University Press.

John Stuart Mill (1859). *On Liberty*, London: John W. Parker & Son.

Nathan Morrow, Nancy Mock, Adam Papendieck & Nicholas Kocmich (2011). *Independent Evaluation of the Ushahidi Haiti Project*, Development Information Systems International, http://ggs684.pbworks.com/w/file/fetch/60819963/1282.pdf.

Kieron O'Hara (2011). *Transparent Government, Not Transparent Citizens: A Report for the Cabinet Office*, London: Cabinet Office, https://www.gov.uk/government/publications/independent-transparency-and-privacy-review.

Kieron O'Hara (2012). 'Trust in social machines: the challenges', in *Proceedings of the AISB/IACAP World Congress 2012: Social Computing, Social Cognition, Social*

Networks and Multiagent Systems (SOCIAL TURN/SNAMAS), http://eprints.soton.ac.uk/339703/.

Kieron O'Hara, Noshir S. Contractor, Wendy Hall, James A. Hendler & Nigel Shadbolt (2013). 'Web Science: understanding the emergence of macro-level features on the World Wide Web', *Foundations and Trends in Web Science*, 4(2/3), 103-267.

Ory Okolloh (2009). 'Ushahidi, or "testimony": Web 2.0 tools for crowdsourcing crisis information', *Participatory Learning and Action*, 59(1), 65-70.

Galen Pickard, Iyad Rahwan, Wei Pan, Manuel Cebrian, Riley Crane, Anmol Madan & Alex Pentland (2010). *Time Critical Social Mobilization: The DARPA Network Challenge Winning Strategy*, arXiv.org 1008.3172v1, http://hd.media.mit.edu/tech-reports/TR-660.pdf.

David Robertson & Fausto Giunchiglia (2013). 'Programming the social computer', *Philosophical Transactions of the Royal Society A: Mathematical Physical and Engineering Sciences*, 371(1987).

Beate Rössler (2005). *The Value of Privacy*, Cambridge: Policy Press.

Nancy L. Rosenblum (2000). *Membership and Morals: The Personal Uses of Pluralism in America*, Princeton: Princeton University Press.

Ferdinand David Schoeman (1992). *Privacy and Social Freedom*, Cambridge: Cambridge University Press.

Nigel Shadbolt, Daniel Smith, Elena Simperl, Max Van Kleek, Yang Yang & Wendy Hall (2013). 'Towards a classification framework for social machines', in *Proceedings of SOCM2013: The Theory and Practice of Social Machines*, Rio, http://eprints.soton.ac.uk/350513/.

Adam Smith (1994). *An Enquiry into the Nature and Causes of the Wealth of Nations*, 2 volumes, Indianapolis, IN.: Liberty Fund.

Max Van Kleek, Daniel Smith, Wendy Hall & Nigel Shadbolt (2013). '"The crowd keeps me in shape": social psychology and the present and future of health social machines', in *Proceedings of SOCM2013: The Theory and Practice of Social Machines*, Rio, http://eprints.soton.ac.uk/350511/.

Luis Von Ahn, Manuel Blum, Nicholas J. Hopper & John Langford (2003). 'CAPTCHA: using hard AI problems for security', in Eli Biham (ed.), *Advances in Cryptology: EUROCRYPT 2003*, Berlin: Springer-Verlag, 294-311.

Luis Von Ahn, Benjamin Maurer, Colin McMillen, David Abraham & Manuel Blum (2008). 'reCAPTCHA: human-based character recognition via Web security measures', *Science*, 321, 1465-1468.

Peter Walker (2013) 'Boston bombing identification attempts on social media end in farce', *The Guardian*, 19th April, 2013, http://www.guardian.co.uk/world/2013/apr/19/boston-bombing-suspects-reddit-social-media.