Knowing and *not knowing* as system design imperatives

Robert J. Houghton & Karen Lancaster

Human Factors Research Group, Faculty of Engineering, University of Nottingham, UK

SUMMARY

We discuss the importance of "not knowing" as a design imperative in digital and automated systems with examples drawn from a range of different settings together with discussion of how this might be responsibly addressed based on analysis using E/HF methods. Reflection is also offered on situations where the temptation to ignorance should not be acted on in design terms - or simply ignored - but embraced as a sensitive heuristic tool for detecting wider system design challenges made salient by digitalisation.

KEYWORDS

Digitalisation, Epistemology, Knowledge Management

Introduction

In common with all applied science disciplines, and rational enterprises in general, Ergonomics/Human Factors (E/HF) places a strong emphasis upon knowing and finding out where gaps in knowledge are perceived to exist. Indeed, it might be argued that E/HF has a particular interest in probing the nature of knowing; a major part of practice of many ergonomists lies in teasing out knowledge from individuals and groups (i.e., knowledge elicitation) and making that sometimes-implicit knowledge concrete in analytic representations (e.g., hierarchical task analysis). In the area of macroergonomics human behaviour is often studied in situations defined by knowledge-seeking in conditions of scarcity of information (e.g., naturalistic decision making, sensemaking). At the same time, however, the technologies associated with Industry 4.0 (Schwab, 2017) and the general spread of ubiquitous computing (Greenfield, 2006) have introduced a step-change in the amount of data that are collected about people and situations. It is arguable that within a short period of time, we have gone from being data-hungry in these areas, to facing a potential digital deluge of both increasing amounts of data and increasingly powerful ways of extracting meaning from it (Sharples & Houghton, 2017). In several disparate studies in our laboratory, and looking more widely, we have noted several situations where tensions arise around unwanted knowledge and where, inherently perhaps, the natural urge is to want to ignore this. However, we were interested in taking seriously the idea that this is indicative of deeper tensions. The purpose of this article is to bring these issues together and reflect on their consequences and possible remedies.

Case studies

Table 1 contains some motivating cases encompassing encountered problems in this area; finding out things it is uncomfortable or disruptive to know, and the way in which it might recast existing rules in a less tolerable light.

Table 1. Example cases

Case	Challenge	Illustrative quote
Machine monitoring sensors *In the case of sensors designed to relay data about the condition of electro-mechanical systems, it became clear that as well as capturing the state of the physical assets, one side effect was to capture "off-book" (or at least strongly discouraged) maintenance shortcuts.*	Industrial politics made it impossible to act on this information or to disclose it was known.	"We don't want to have the conversation; surveillance wasn't supposed to be part of the package. We choose that we don't know this."
Domestic technologies *In prior work in our laboratory [Brown et al., 2015], we examined domestic activity logging in both the general and specific personalised scenarios. In the general case, nearly all respondents believed it was important for parents to know what children were doing outside their direct observation.*	In the specific case, few respondents wanted to know about what *their* children were doing once specific domestic context personal to them was present. Some expressed alarm that this would cause domestic tensions and arguments.	"If they know that you know, you have to do something. Sometimes it is better not to know; family life is complicated"
Over specified sensors *In a technical project in our laboratory we designed furniture that could detect posture through a range of pressure sensors and load cells. While pressure sensors give only qualitative readings, the load cells offer accurate measurement of load and as relatively cheap and robust sensors were used.*	The data stream collected when processed delivered the desired posture data; unfortunately, the load cell data stream weighed participants to a point of precision where digestive functions could be inferred.	"Cheaper, less accurate sensors do not exist, there is no point manufacturing worse products that would sell in fewer numbers for a higher price"
Video Assistant Referee controversy *VAR uses multiple cameras and digital video processing to identify violations of the offside law. Unfortunately, this has led to a great deal of controversy.*	The offside rule not designed with VAR in mind; this leads to calls that frustrate fans and players. One cannot change the technology nor the laws of the game (there is a principle that it should be the same played in the park as it is in the World Cup final). Compromise reached in wider onscreen bars; but neither law nor technology allow for margin of error, effectively undermining both.	"If you have a long nose, you are in an offside position these days… So our proposal will be – we will discuss this with our referees' division - that it is a tolerance of 10-20 centimetres…" [Aleksander Ceferin, President of UEFA]

Emerging from these case studies we see a set of emerging themes. The first is that we can sometimes know more than we intended largely by accident. This might be because of confusion around what we said we wanted versus what we actually wanted (domestic technology) or because constructive knowledge has been generated inadvertently (machine monitoring) or because the technology is generating more data to a higher resolution than we anticipated (over specified sensors and VAR). The second is the effective discovery that rules intended to be taken seriously in an analogue or low information setting are now hard to live with in a digital or high information setting. These can be informal (such as domestic rules within a family) or formal and written (as in the case of the laws of football). Both these examples may also generalise to the workplace, where expectations of employees, once reified with data and sensing, become unexpectedly oppressive. This may also suggest that implicitly, but not explicitly, there was a tacit understanding there should be some leeway, or indeed that informal coping strategies were probably mutually expected. Third is the desire for constructive ignorance; choosing not to know what we do know or could know. This may be more challenging than it first seems. Managers responsible for machines are troubled by knowing damage is potentially done to them, but also troubled about both ignoring this, and trying to find an explanation of how they knew if they had to intervene. As noted in the illustrative quote from participants, parents identify the problem of not merely having to ignore information they do not want to know, but also plausibly being seen as not knowing it! In the VAR example controversy is regularly created through what is known and acted upon (ending an attacking move), and it seems likely pretending not to know offside was violated would be seen as equally incendiary by supporters of the opposing defending team. As sports journalists have noted, “The problem is not that the technology is not working. It is that it is working too well…This makes his proposed solution rather troubling… It hopes to blur the picture slightly, in the same way that it would if the VAR officials were forced to work with a 30-year-old black and white telly. But they won’t be working with old tech” (Wood, 2019).

The question then arises: how can we think about ignorance as anything other than a vice, is it acceptable and can E/HF contribute?

Design for Privacy

One significant step towards limiting risks around knowing lies is in implementing “Design for Privacy” (Hustinx, 2010). The UK Information Commissioner’s Office (ICO) states: “The UK GDPR requires you to put in place appropriate technical and organisational measures to implement the data protection principles effectively and safeguard individual rights. This is ‘data protection by design and by default’. In essence, this means you have to integrate or ‘bake in’ data protection into your processing activities and business practices, from the design stage right through the lifecycle” (ICO, 2023). The UK General Data Protection Regulation (UK GDPR, also known as the Data Protection Act 2018) itself emphasises rights around the protection, appropriate use of and minimisation of personal data. Recent standards have also emphasised that this covers all stakeholders to a given system (not just customers or users in a generic sense) and that a significant role lies in this process for human factors and ergonomics (ISO 31700-1:2023, 3.21, p5). Generically, the concept behind Design For Privacy lies in instituting concerns for privacy through the lifecycle of the system, from its initial concept through to its eventual disposal and that alongside the system lifecycle itself, there is a linked data lifecycle of collecting only what is required, holding and processing it securely and its eventual deletion and disposal. Manifest within this is a clear focus on personal information about people and the consequent risks this might pose in terms of unethical misuse, embarrassment, reputational damage, fraud and so on. The UK GDRP also notes that stronger legal protections are granted to information deemed particularly sensitive that includes protected characteristics and issues such as health, genetics, and biometrics.

However, reviewing this approach against the cases described in Table 1, we see it may not provide sufficient guidance as in these examples the technology is working as intended for the purposes for which they were intended. It is the case that over-specific sensors might fall under the rubric of data minimisation (i.e., record data on the device at a lower resolution from the outset) but this might not always be technically feasible, and may not necessarily be seen as desirable, as discussed earlier (business case, perceived need to know). Thus, while Design for Privacy does good service at addressing some risks, and as ICO emphasise, and is in any case a non-negotiable, it may not fully encompass the scope of the problem identified here. An analogy might be that while a university research ethics process might reduce risks of a given experiment or act of data collection, questions about whether the underlying venture itself is an appropriate one might arguably lie in a slightly different space, that of Responsible Research and Innovation (Von Schomberg, 2013). In other words, it is less about how things are found out or done than it is about whether they should be known or done at all, when considered in wider context. Avoiding a risk arising can be done from default design choices; deciding not to do something is an active positive decision that must be argued for and chosen.

Acceptability of not knowing

One objection to any state of ignorance is a general discomfort of scientists and engineers in deliberately choosing to be ignorant. Constructive ignorance is common to double-blind procedures (for example) and might also be argued for in bioethics scenarios. By contrast, many examples of deliberate ignorance or choosing to create ignorance such as those around harms to people or the environment from certain products and industries and are deservedly viewed as egregious. More widely, the business case for constructive ignorance may be felt to challenge the business case for digitalisation if there is a sense in which some actionable data is not collected, or actionable information not produced; why are we “leaving this on the table?” and limiting ourselves? Although a strict interpretation of Design for Privacy my militate against collecting personal data speculatively, it seems far more reasonable to amass archived industrial data for later use.

From a broader philosophical perspective, it has been argued that there is a pervasive tendency to suffer from a so-called ‘epistemic anxiety’ – characterised as a need or desperation to know something (Nagel 2010; Hookway 2011). The primary role of epistemic anxiety is that it alerts us to the fact that we lack knowledge in a particular domain (Hookway 2011, p. 36). In cases where the stakes are high, and not knowing something could cause serious problems, the anxiety is heightened (Nagel 2010, p. 408). There is a pressure – whether from ourselves or from society at large – to constantly improve our epistemic standpoint by increasing our knowledge. We should not, for example, risk being dogmatic and closed-minded, since these are epistemic vices (Cassam 2016). Under normal circumstances, it is epistemically virtuous to pursue knowledge and be open to new forms of information; it is often seen as epistemically vicious to be wilfully ignorant and closed-minded to new information. ‘Active ignorance’ – “the kind of ignorance that is deeply invested in not knowing” (Medina, 2016, p. 182) or ‘motivated ignorance’ – ignorance “that results from a desire not to know” is, in garden variety cases, an epistemic vice which is arrogant, where one indulges in the ‘luxury’ of refusing to examine potentially unpalatable information (Tanesini, 2006 p. 60). Peels (2023) distinguishes several varieties of ignorance; one variety is ‘undecided ignorance” – where one has simply not thought about a matter; another is ‘strategic ignorance’ – where one chooses (intentionally) to remain ignorant of a matter. In the cases we outline, prior to the technology gathering the data, people were in a state of undecided ignorance – they had never thought about the matter in hand (and even if they had, they would have remained agnostic because no such information was available to them). But once the information is gathered by technology and is available to them, they are presented with a choice – to remain (strategically) ignorant, or to know. In the cases we outline, people preferred the (forced) agnosticism of the time when such information was simply never gathered or available to them. Consequently, this suggests that such difficult decisions about what information will be gathered might be better made from the beginning (echoing the lifecycle approach taken in Privacy by Design and the suggestion that this should not be an ‘add on’ but inherent within the venture) and would benefit from sensible and supportable justifications.

Can E/HF methods help?

As with any applied science, E/HF derives from a largely positivist approach to science and, as discussed earlier, methodologically demonstrates a particular interest in eliciting and making concrete the knowledge held by workers and users and in provision of information to them. While a complete review of E/HF methods is outside the scope of this article, some examples and suggestions can be made. First, to the extent E/HF draws on experiments (and particularly in some areas cognitive psychology experiments), experiments are normally designed with experimental controls that restrict the behaviour of the participant to elucidate underlying cognitive mechanisms. At the same time, however, they may show how reasonable performance can be achieved with surprisingly little information. For example, the ability to infer sophisticated intention and affect from point-light displays of biological motion absent any figural representation of people at all (see Blakemore & Decety, 2001) suggests some CCTV tasks, particularly in sensitive settings, can be performed without ever needing to share high resolution imagery with operators, even if this represents a positive decision to work with less information than is actually available. In terms of knowledge elicitation techniques, we note that while many popular techniques, such as those associated with Cognitive Task Analysis, ask about information needs, information in use – and indeed what people wish they had known but did not – it is less common to ask what information or knowledge a person wishes they had *not* been aware of, or which made the task more challenging (Crandall, Klein & Hoffman, 2006). Finally, in terms of representational methods, one way to make concrete what we do not want to know about is to, broadly, expose and discuss the negative complement. We show this in generic form in Figure 1, but this could be applied to other methods in the sense their diagrams could in principle be reduced to connected/not connected nodes (e.g., Ramussen's Decision Ladder) where attention might also be explicitly focused on the absence of connections and what we intend by those absences, which is typically left unclear.

These actors know or need to know….		Information				
		Rate of process	Stage of process	Condition of assets	Performance	Error states
Actors	Supervisors	■	■	■	■	■
	Maintainers			■	■	■
	Operators	■	■	■	■	■
	Users		■			
	Emergency staff			■	■	■

These actors do not know or *should not* know?		Information				
		Rate of process	Stage of process	Condition of assets	Performance	Error states
Actors	Supervisors					
	Maintainers	■	■			
	Operators					
	Users	■		■	■	■
	Emergency staff	■	■			

Figure 1. Negative complement for actor/information allocation

Temptation to ignorance as a formative event

In the above examples drawn from work in our laboratory, we also note just how tempting it is to overlook situations in which unwanted knowledge has become present through digitalisation. The risk inherent in such situations is one of drift where a degree of 'peeking' slowly becomes normalised. We suggest that rather than treating this as an inconvenience, rather, we should recognise that this may be one of the most sensitive heuristic tools available to us for detecting E/HF challenges in work systems, as in reality we have vividly exposed a situation where "work as imagined" is newly revealed as colliding with "work as done". In our earlier motivating examples, it is clear that they indicate cases for action. The overspecified sensor can be dealt with perhaps most straightforwardly through implementing Design for Privacy and more actively taking steps (perhaps in hardware or low-level software) to deliberately reduce its resolution to just that required for the purposes it was bought for. This would require us to accept that if future work requiring higher resolution data were required, based on reasonable justifications, more data collection would need to take place.

In the case of the machine monitoring, the finding is that the wrong maintenance method has been made easy, and the correct method hard. This may indicate a need for clearer training and ostensible supervision and most likely indicate redesign to make the wrong method impossible and the right method easier. It may also be the case that the organisation needs to review its motivational policies towards the right thing rather than the quick/profitable thing (as safety scientists have previously identified in the context of perceived conflicts of safety versus production, e.g., Hollnagel, 2004). More radically, it might want to seriously examine whether workers have in fact made the right choice in shopfloor practice versus management theory, and take steps to facilitate rather than oppose their approach (Dekker, 2014). In the case of domestic technology, it suggests a failure of the requirements capture process which has failed to capture nuances of the family environment. One strategy might be to investigate sociological and ethnographic accounts of family life and tensions to inform sensitive design choices. Here, implicit requirements are found in tension with explicitly stated requirements.

Finally, in the case of workplace situations where rules come into conflict with data, it may suggest again a degree of acceptance of a previously denied complexity ranging from whether optimal workplace behaviour is really achievable to social and work-life balance issues. In the specific case of VAR, it may boil down to undertaking the difficult task of reopening the rulebook and deciding how people truly want football matches to be played, and whether attack or defence should be advantaged or disadvantaged.

Conclusion

As we have seen in the above cases, the main impact of digitalisation and the perceived requirement to know is that it can shine and strong and perhaps unexpected light on elements of human and system behaviour that had previously lain hidden, unacknowledged or otherwise fudged. The challenge, then, ultimately, is not really one of digitalisation alone, rather also engaging with what it reveals through a resubscription to the principles of E/HF and people-led design of work to understand the systems and contexts in which work and wider life exist.

Acknowledgement

This work was supported by Made Smarter Innovation – Centre for People-Led Digitalisation (EP/V062042/1)

References

Blakemore, S-J., Decety, J. (2001). From the perception of action to the understanding of intention. Nature Reviews Neuroscience, 2: 61–567

British Standards Institution (2023). ISO 31700-1:2023 Consumer protection – privacy by design for consumer goods and services. BSI Standards Limited.

Brown, M., Coughlan, T., Blum, J., Lawson, G., Houghton, R.J., Mortier, R., Goulden, M., & Arunchalam, U. (2015). Tailored scenarios: A low-cost online method to elicit perceptions of home technologies using participant-specific contextual information. Interacting with Computes, 27, 1, 60-71.

Cassam, Q. (2016). Vice epistemology. The Monist, 99 (2), 159-180.

Crandall, B., Klein, G., & Hoffman, R.R. (2006). Working minds. Cambridge MA: MIT Press.

Data Protection Act (2018). Available at: https://www.legislation.gov.uk/ukpga/2018/12/contents/enacted [Accessed December 2023].

Dekker, S. (2014). Safety differently: Human factors for a new era. CRC Press.

Greenfield, A. (2006). Everyware: The dawning age of ubiquitous computing. New Riders.

Hollnagel, E. (2004). Barriers and accident prevention. Aldershot UK: Ashgate.

Hookway, C. (2011). James's epistemology and the will to believe. European Journal of Pragmatism and American Philosophy, 3 (1), 30-38.

Hustinx, P. (2010). Privacy by design: Delivering the promises. Identity in the Information Society, 3, 2: 253-255.

Information Commissioner's Office (2023). Data protection by design and default. Available at https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/accountability-and-governance/guide-to-accountability-and-governance/accountability-and-governance/data-protection-by-design-and-default/ [Accessed December 2023].

Medina, J. (2016). Ignorance and racial sensitivity. In R. Peels and M. Blaauw (Eds.) The epistemic dimensions of ignorance. Cambridge UK: Cambridge University Press.

Nagel., J. (2010). Epistemic anxiety and adaptive invariantism. Philosophical Perspectives, 24 (1), 407-435.

Peels, R. (2023). Ignorance: A philosophical study. Oxford UK, Oxford University Press.

Schwab, K. (2017). The Fourth Industrial Revolution. Penguin.

Sharples, S. & Houghton, R.J. (2017). The field becomes the laboratory? The impact of the contextual digital footprint on the discipline of E/HF. Ergonomics, 60 (2), 270-283.

Tanesini, A. (2020). Ignorance, arrogance and privilege. In J. Kidd, H. Battaly and Q. Cassam (Eds.) Vice epistemology. Taylor & Francis.

Wood, G. (2019). Aleksandr Ceferin's 'long nose' issue highlights the VAR offside conundrum. The Guardian, Sunday 8th December, 2019.

Von Schomberg, R. (2013). A vision of responsible innovation. In R. Owen, M. Heintz and J. Bessant (Eds.) Responsible Innovation. London: John Wiley.