Citation for published version:

Wortham, RH & Theodorou, A 2017, 'Robot transparency, trust and utility', *Connection Science*, vol. 29, no. 3, pp. 242-248. <https://doi.org/10.1080/09540091.2017.1313816>

DOI:

[10.1080/09540091.2017.1313816](https://doi.org/10.1080/09540091.2017.1313816)

Publication date:

2017

Document Version

Peer reviewed version

[Link to publication](#)

"This is an Accepted Manuscript of an article published by Taylor & Francis in *Connection Science* on 30th May 2017, available online: <http://www.tandfonline.com/10.1080/09540091.2017.1313816>."

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

To appear in *Connection Science*
 Vol. 00, No. 00, January 2017, 1–7

Robot Transparency, Trust and Utility

Robert H. Wortham^{a*} and Andreas Theodorou^a

^a *Department of Computer Science, University of Bath, Bath, UK*

(v1.1 released January 2017)

As robot reasoning becomes more complex, debugging becomes increasingly hard based solely on observable behaviour, even for robot designers and technical specialists. Similarly, non-specialist users have difficulty creating useful mental models of robot reasoning from observations of robot behaviour. The EPSRC Principles of Robotics mandate that our artefacts should be transparent, but what does this mean in practice, and how does transparency affect both trust and utility? We investigate this relationship in the literature and find it to be complex, particularly in non industrial environments where, depending on the application and purpose of the robot, transparency may have a wider range of effects on trust and utility. We outline our programme of research to support our assertion that it is nevertheless possible to create transparent agents that are emotionally engaging despite having a transparent machine nature.

Keywords: EPOR, transparency, agents, ethics, roboethics, robotics

Index to information contained in this article

1. Introduction
2. Theory of mind, trust and transparency
 - 2.1. Anthropomorphism and mental models of robots
3. Companions or tools: A false dichotomy?
4. Research programme
5. Summary and recommendations

1. INTRODUCTION

The EPSRC Principles of Robotics include a specific reference to transparency: “Robots are manufactured artefacts. They should not be designed in a deceptive way to exploit vulnerable users; instead their machine nature should be transparent.” (Boden et al., 2011). This initially appears to be a straightforward normative assertion, drawing on the commonly held idea that agents should not be deceptive, since deception generally leads to exploitation. This paper considers whether in fact transparency is really such a simple idea, and also whether making certain types of agents transparent reduces their utility. In considering this question, we must also address the relationship between transparency and trust.

In this paper, we use the term *agent* to mean an autonomous intelligent entity that acts in some world. We use the term *robot* as a subclass of agent, and mean it to be an

*Corresponding author. Email: r.h.wortham@bath.ac.uk

artefact that operates in the physical world.

What does it mean to trust a robot? We might initially simply assert that if a robot is more transparent, then we are able to trust it more, and therefore its utility increases. Here, we are using transparency to mean the extent to which the internal state and decision making processes of a robot are accessible to the user. We could also argue that trust is only required when a robot is not fully transparent, and therefore that increased transparency reduces the need for trust (Bryson & Rauwolf, 2017).

In the case of companion robots, trust is an important component of the purpose of the robot. The user must trust the robot in order to consider it a worthy companion. Thus, if the utility of a robot is measured by the degree to which it is trusted, then increasing transparency may reduce that utility.

So, we start to see that there is a complex relationship between the ideas of utility, transparency and trust. This relationship will depend on the purpose of the robot. In this paper we review the literature relating to transparency and trust, and we also describe ongoing practical research to investigate the proposal that it is indeed possible build an emotionally engaging yet transparent robot.

2. THEORY OF MIND, TRUST AND TRANSPARENCY

Although we may presuppose that communication between animals, and particularly between humans must be complex, in fact natural communication systems tend to exploit relatively simple and minimal signals, the meaning of which derives from extensive models (Wortham & Bryson, 2016) — that is to say, pre-existing shared paradigms and semantic relationships. In other words, evolution, or a shared phylogenetic history, provides adequate priors such that minimal data is required to communicate context.

Although some would argue otherwise (Gallagher, 2006), it is generally agreed that effective interaction, whether coercion or co-operation, relies on each party having some theory-of-mind (ToM) of the other (Saxe, Schulz, & Jiang, 2006). Individual actions and composite behaviours are thus interpreted within a pre-existing ToM framework. Whether that ToM is accurate is unimportant, provided that it is predictive in terms of behaviour.

The robot’s transparency model does not define the ToM employed by the human user, but it is the transparency model that we can directly adjust and this is therefore the focus of this paper.

It is well known that observable behaviour can communicate the internal mental states of the individual. Breazeal, Kidd, Thomaz, Hoffman, and Berlin (2005) found that implicit non-verbal communication improves transparency over that of only deliberate non-verbal communication. Here, *implicit* is defined as conveying information inherent in behaviour but which is not deliberately communicated by the robot designer. People have strong expectations for how implicit and explicit non-verbal cues map to mental states.

Breazeal et al. also found that transparency reduces conflict when errors occur, particularly when a joint task is being attempted. Reduced conflict implies that when an error occurs during task execution, recovery is still possible with less apportionment of blame. Breazeal et al. term this reduced conflict *robustness*, and this robustness is one effective measure of utility.

2.1. *Anthropomorphism and Mental Models of Robots*

Humans have a strong predisposition to anthropomorphise not only nature, but anything around them (Dautenhahn, 2007). The Social Brain Hypothesis (Dunbar, 1998) may explain this phenomenon, however humans do not treat robots identically to humans, for example with respect to moral standing (Kahn, Ishiguro, Friedman, & Kanda, 2006). Although there is significant debate about the ontology of robot minds versus human minds, what is of more practical importance is how robot minds are understood psychologically by humans, i.e. what is the perceived, rather than actual, ontology.

Stubbs, Hinds, and Wettergreen (2007) consider it essential to form a mental model of robots in order to build common ground — which we might also interpret as the basis for human trust. Stubbs et al. also found that this common ground can be effectively established via an interactive dialogue with the robot. Although this study primarily considered remote robots working in an industrial or exploratory setting, rather than robots operating in domestic environments, we should take note of the importance of dialogue in establishing trust. Indeed Mueller (2016) sees *dialogue* as one of the three main characteristics of transparent computers, the others being *explanation* and *learning*.

Meerbeek, Hoonhout, Bingley, and Terken (2006) investigates the relationship between a robot’s perceived personality and the level to which the user feels in control during the interaction. In order to be believable, Meerbeek et al. found that the personality expression should be linked to an internal model that deals with the behaviour (e.g. decision making) based on personality and emotion. More expressive, informal behaviour is associated with a higher perception of user control.

Non-specialist humans either have little ToM for robots, or may have a model based on contemporary science fiction, and therefore interpret behaviours using a default *other agent theory*, which assumes the agent to share human-like motivations. This can be understood in evolutionary terms through our ancestors’ need to rapidly categorise proximal activity as either neutral (the rustling of leaves in the wind), friendly (the approach of a tribe member) or hostile (the approach of a predator or foe). When sensory information is uncertain, in an environment where one is frequently the prey, not the predator, then evolving a bias towards an assumption of both agency and hostility is selective for individual longevity. Even in our technological environments we often experience fake agency, such as robotic dialing of sales calls, automated twitter postings and auto-generated, personalised spam emails. These are complex matters and require further research, particularly to investigate how folk psychology and science fiction affect the ToM humans use to interpret robot behaviour.

In a study conducted in 2006 in a community hospital in the USA, the nursing staff were constantly searching for reasons why the robots acted as they did. They would ask themselves and others, “*What is going on here? Is the robot supposed to do this or did I do something wrong?*”. This research asserts that low levels of transparency led people to question even the normal behaviours of the robot, sometimes even leading people to think of correct behaviours as errors (Kim & Hinds, 2006).

Hancock et al. (2011) assert that if we cannot trust our robots, we will not be able to benefit from them effectively. However, given that we happily interact in society with others whom we do not completely trust, and increasingly we interact with computers knowing that their recommendations maybe faulty, it may be that Hancock et al. are over simplifying for the general case. Perhaps what is more important is that humans are able to correctly calibrate their level of trust appropriately (Dzindolet, Peterson, Pomranky, Pierce, & Beck, 2003).

3. COMPANIONS OR TOOLS: A FALSE DICHOTOMY?

Principle One of the Principles of Robotics asserts that robots are tools. Within industrial and engineering environments this is fairly clear, in the sense that a human uses the robot to complete a technical task. The designer and user of the robot share the goal of the robot: to complete the task. However, within domestic and healthcare environments, robots may have rather a different relationship with those they interact with. They may be intended to provide companionship and simultaneous covert monitoring of patient well-being. They may be tools for the healthcare professional, but for the patient they are companions. In such an environment the utility may be negatively affected by increased transparency.

Our sense of companionship is related to the measure of agency we project onto the robot. If we are able to understand the workings of the intelligence does it inherently appear to become less intelligent in the folk sense, such that we then project less agency, and as a result experience less benefit from the robot? We might compare this with television. We know it has no agency, but its presence in the corner of our sitting room does provide companion like benefits. Maybe this has to do with the conscious suspension of disbelief, or maybe we have an unconscious agency detector which is more easily fooled by technology.

Common-sense notions of intelligence are conflated with folk psychology ideas of agency and also of living. Things that are intelligent are alive, in the sense that they have their own beliefs, desires and intentions that we understand are fundamentally self serving, or selfish. We implicitly recognise selfishness as a fundamental characteristic of all life (Dawkins, 1976). If such a selfish agent engages with us then it considers us to be important in the pursuit of these selfish objectives. Such agents are worthy of becoming our companions because they ascribe true value in their relationship with us, and this increases our value in society.

Conversely, agents who have no self-serving agency are not worthy of our attention because they convey no social value. Perhaps therefore, artificial agents whose sole purpose is companionship and are truly transparent in this respect are thus disqualified from being worthy companions. In some situations robot transparency may therefore be at odds with utility, and more generally it may be orthogonal rather than beneficial to the successful use of the robot.

This argument is at present merely hypothesis; the processes leading to a robot achieving companion status need to be more thoroughly understood. Whilst we may invent scenarios and continue to discuss the theoretical and philosophical interplay between transparency, trust and utility, as scientists we await the outcome of our experiments.

4. RESEARCH PROGRAMME

We are beginning a programme of practical research to investigate the transparency, trust, utility triangle. Initially using the non-humanoid R5 robot, we are conducting experiments to determine the effect of various types of robot transparency on mental models formed by humans. We are also interested in the consequential effect on human emotional response. At the heart of our experiments we are using reactive planning techniques to build autonomous agents. We have developed the *Instinct* reactive planner (Wortham, Gaudl, & Bryson, 2016) based on Bryson's POSH (Parallel Ordered Slipstack Hierarchy) reactive planner and Behaviour Oriented Design (BOD) approach (Bryson, 2001). A POSH plan consists of a *Drive Collection* containing one or more *Drives*. Each

Drive has a priority and a releaser. When the *Drive* is released as a result of sensory input, a hierarchical plan of *Competences*, *Action Patterns* and *Actions* follows. The Instinct Planner reports the execution and status of every plan element in real time, allowing us to implicitly capture the reasoning process within the robot that gives rise to its behaviour. Our experiments will investigate and demonstrate how this transparency data from the planner can be used to make the behaviour of a robot more understandable. Initially we are primarily interested in making the behaviour transparent for a robot designer, since robots with complex plans are typically very hard to design and debug. However, these initial experiments may also improve transparency for non-specialist observers.

We will subsequently investigate how we can harness the transparency mechanism embedded with the Instinct Planner to produce a more effective domestic robot. The research will investigate whether transparency makes people feel more or less bonded to their robot, and whether they are more or less able to accurately assess the needs of the robot, as it works to achieve its goals.

It is anticipated that these trials should take place within a domestic or near-domestic environment, such as a retirement home. We must gain feedback from non-specialist observers/users about the qualitative level of intelligence of the robot, and also about how comfortable they would be to have such a device in their home environment. The research will attempt to assess initial levels of fear, anxiety, mistrust of AI and robots in general, and of domestic robots in particular. Having established a reference position, transparency of the robot must be enabled by providing feedback to the user based on the real time execution within the reactive planner. The methods we currently envisage are:

- Real-time presentation of textual statements relating to plan execution.
- Graphical real-time visualisation of plan execution.
- Audio (i.e. verbal) statements relating to robot plan execution.

For each of these methods the transparency information could either be presented on/from a remote device, or on/from the robot itself. There are thus six possible combinations. Of course additional transparency fusion, such as audio combined with graphical, could also be tested based on the success or failure of initial experimental results.

As the literature indicates that dialogue is important in establishing trust, this research should give some consideration to the possibility of accepting speech input, albeit restricted to simple commands, as a means for users to inquire of the robot what it is doing, and to have the robot respond appropriately.

5. SUMMARY AND RECOMMENDATIONS

We have seen that unpacking transparency and trust is complex, but can be partly understood by looking at how humans come to understand and subsequently trust one another, and how they overcome evolutionary fears in order to trust other agents, through implicit non-verbal communication. Unacceptable levels of anxiety, fear and mistrust will result in an emotional and cognitive response to reject robots.

Finally, there may be applications where transparency is at odds with utility. Our ongoing programme of research is intended to validate our hypothesis that we can indeed create transparent robots that are nevertheless emotionally engaging and useful tools across a wide range of domestic and near-domestic environments. Meanwhile, there remains a great of work to be done to unpack the relationship between transparency, utility and trust.

References

- Boden, M., Bryson, J., Caldwell, D., Dautenhahn, K., Edwards, L., Kember, S., . . . Winfield, A. (2011, April). *Principles of robotics*. The United Kingdom's Engineering and Physical Sciences Research Council (EPSRC). (web publication)
- Breazeal, C., Kidd, C., Thomaz, A., Hoffman, G., & Berlin, M. (2005). Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *2005 ieee/rsj international conference on intelligent robots and systems* (pp. 708–713). Alberta, Canada: Ieee. Retrieved from <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1545011>
- Bryson, J. J. (2001). *Intelligence by Design: Principles of Modularity and Coordination for Engineering Complex Adaptive Agents* (Doctoral dissertation, MIT). Retrieved from <ftp://111.ai.mit.edu/people/joanna/phd-tr.pdf>
- Bryson, J. J., & Rauwolf, P. (2017). Trust, Communication, and Inequality. In *The 38th annual meeting of the cognitive science society*. Philadelphia, USA.
- Dautenhahn, K. (2007). Methodology & themes of human-robot interaction: A growing research field. *International Journal of Advanced Robotic Systems*, 4(1 SPEC. ISS.), 103–108.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.
- Dunbar, R. I. M. (1998). The Social Brain Hypothesis. *Evolutionary Anthropology*, 178–190.
- Dzindolet, M. T., Peterson, S. A., Pomranky, R. A., Pierce, L. G., & Beck, H. P. (2003). The role of trust in automation reliance. *International Journal of Human Computer Studies*, 58(6), 697–718.
- Gallagher, S. (2006). The narrative alternative to theory of mind. In R. Menary (Ed.), *Radical enactivism: Intentionality, phenomenology, and narrative* (pp. 223–229). Amsterdam: John Benjamins. Retrieved from <http://www.researchgate.net/>
- Hancock, P. a., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., & Parasuraman, R. (2011). A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 53(5), 517–527.
- Kahn, P. H., Ishiguro, H., Friedman, B., & Kanda, T. (2006). What is a human? - Toward psychological benchmarks in the field of human-robot interaction. *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, 3, 364–371.
- Kim, T., & Hinds, P. (2006). Who should I blame? Effects of autonomy and transparency on attributions in human-robot interaction. *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, 80–85. doi:
- Meerbeek, B., Hoonhout, J., Bingley, P., & Terken, J. (2006). Investigating the relationship between the personality of a robotic TV assistant and the level of user control. *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, 404–410. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4107841
- Mueller, E. T. (2016). *Transparent Computers: Designing Understandable Intelligent Systems*. San Bernardino, CA: Erik T. Mueller.
- Saxe, R., Schulz, L. E., & Jiang, Y. V. (2006, jan). Reading minds versus following rules: dissociating theory of mind and executive control in the brain. *Social neuroscience*, 1(3-4), 284–98. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/18633794>

- Stubbs, K., Hinds, P. J., & Wettergreen, D. (2007). Autonomy and Common Ground in Human-Robot Interaction: A Field Study. *IEEE Intelligent Systems*, 22(2), 42–50.
- Wortham, R. H., & Bryson, J. J. (2016). Communication. In *Handbook of living machines {accepted for publication}*. Oxford: Oxford University Press. Retrieved from <http://www.cs.bath.ac.uk/~jjb/ftp/Communication15.pdf>
- Wortham, R. H., Gaudl, S. E., & Bryson, J. J. (2016). Instinct : A Biologically Inspired Reactive Planner for Embedded Environments. In *Proceedings of icaps 2016 planrob workshop {accepted for publication}*. London, UK.