

# Robots That Make Sense: Transparent Intelligence Through Augmented Reality

Alexandros Rotsidis\*  
University of Bath  
Bath, United Kingdom  
A.Rotsidis@bath.ac.uk

Andreas Theodorou\*  
Umeå University  
Umeå, Sweden  
andreas.theodorou@umu.se

Robert H. Wortham  
University of Bath  
Bath, United Kingdom  
r.h.wortham@bath.ac.uk

## ABSTRACT

Autonomous robots can be difficult to understand by their developers, let alone by end users. Yet, as they become increasingly integral parts of our societies, the need for affordable easy to use tools to provide transparency grows. The rise of the smartphone and the improvements in mobile computing performance have gradually allowed Augmented Reality (AR) to become more mobile and affordable. In this paper we review relevant robot systems architecture and propose a new software tool to provide robot transparency through the use of AR technology. Our new tool, ABOD3-AR provides real-time graphical visualisation and debugging of a robot's goals and priorities as a means for both designers and end users to gain a better mental model of the internal state and decision making processes taking place within a robot. We also report on our on-going research programme and planned studies to further understand the effects of transparency to naive users and experts.

## CCS CONCEPTS

• **Human-centered computing** → Mixed / augmented reality; Systems and tools for interaction design; • **Computing methodologies** → Artificial intelligence; • **Software and its engineering** → Software creation and management; Software design engineering; • **Social and professional topics** → Computing / technology policy.

## KEYWORDS

robots, mobile augmented reality, transparency, artificial intelligence

## 1 INTRODUCTION

The relationship between transparency, trust, and utility is a complex one. By exposing the inner 'smoke and mirrors' of our agents, we risk of making them look less interesting. Moreover, the wide range of application domains for AI and of the different stakeholders interacting with intelligent systems should not be underestimated. Therefore, What is effectively transparent varies by who the observer is, and what their goals and obligations are. There is however a need for design guidelines on how to implement transparent

\*We thank the EPSRC grant [EP/L016540/1] for funding Rotsidis and Theodorou. Both authors contributed equally to the paper.

systems, alongside with a 'bare minimum' standardised implementation [3]. In the end, the goal of transparency is should not be complete comprehension, that would severely limit the scope of human achievement. Instead, the goal of transparency is to provide sufficient information to ensure at least human accountability [7].

Still, the use real-time implementation can help users to *calibrate their trust* in the machine [13, and references therein]. Calibration refers to the correspondence between a person's trust in the system and the system's capabilities [12]. Calibrating of trust occurs when the end-user has a mental model of the system and relies on the system within the system's capabilities and is aware of its limitations. If we are to consider transparency as mechanism that exposes the decision-making of a system, then it can help users adjust their expectations and forecast certain actions from the system. This position about transparency is supported by Dzindolet et al. [8], who conducted a study where the participants decide whether they trust a particular piece of pattern recognition software. The users were given only the percentage of how accurate the prediction of their probabilistic algorithm was in each image. Yet, by having access to this easy-to-implement transparency feature, they were able to calibrate their trust in real time. Our own studies [discussed in 20], demonstrate how users of various demographic backgrounds had inaccurate mental models about a mobile robot running a BOD-based planner, Instinct [19]. The robot transmits a transparency feed to the real-time debugging software ABOD3 [7, 16]. The transparency display is customised for a high-level end-user display of the robot's goals and process towards those goals. Participants without access to the transparency software ascribe unrealistic functionalities, potentially raising their expectations for its intelligence and safety. When the same robot is used with ABOD3, providing an end-user transparency visualisation, the users are able to calibrate their mental models, leading to more realistic expectations, but interestingly a higher respect for the system's intelligence.

Yet, despite its effectiveness, there is a major disadvantage with the ABOD3 solution: a computer and display is required to run the software. One solution might be to port ABOD3 to run directly on robots with built-in screens, such as SoftBank Robotic's Pepper. Albeit that this is a technologically feasible and potentially interesting approach, it also requires that custom-made versions of ABOD3 will need to be made for each robotics system. Moreover, this is not a compatible solution for robots without a display.

Nowadays, most people carry a smartphone. Such mobile phones are equipped with powerful multi-core processors, capable of running complex computational-intensive applications, in a compact package. Modern phones also integrate high-resolution cameras, allowing them to capture and display a feed of the real world.

That feed can be enhanced with the real-time superimposition of computer-generated graphics to provide Augmented Reality (AR) [1]. Unlike Virtual Reality that aims for complete immersion, AR focuses on providing additional information of and means of interaction with real-world object, locations, and even other agents.

In this paper we demonstrate new software, *ABOD3-AR*, which can run on mobile phones. *ABOD3-AR*, as its name suggests, uses a phone's camera to provide AR experience by superimposing the *ABOD3*'s tree-like display of Instinct plans over a tracked robot.

## 2 TOOLS AND TECHNOLOGIES FOR TRANSPARENCY

In this section we describe in some detail the tools and technologies used in our transparency experiments.

### 2.1 Behaviour Oriented Design

Behaviour Oriented Design is a cognitive architecture that provides an ontology of required knowledge and a convenient representation for expressing timely actions as the basis for modular decomposition for intelligent systems [5, 6]. It takes inspiration both from the well-established programming paradigm of object-oriented design (ODD) and its associated agile design [9], and an older but well-known AI systems-engineering strategy, *Behaviour-Based AI* [4].

BOD helps AI developers as it provides not only an ontology, addressing the challenge of 'how to link the different parts together', but also a development methodology; a solution to 'how do I start building this system'. It includes guidelines for modular decomposition, documentation, refactoring, and code reuse. BOD aims to enforce the good-coding practice 'Don't Repeat Yourself', by splitting the behaviour into multiple modules. Modularisation makes the development of intelligent agents easier, faster, reusable and cost efficient. Behaviour modules also store their own memories, e.g. sensory experiences. Multiple modules grouped together form a *behaviour library*. This 'library' can be hosted on a separate machine, for example in the cloud. The *planner* executing within the agent is responsible for exploiting a plan file; stored structures describing the agent's priorities and behaviour. This separation of responsibilities into two major components enforces further code reusability. The same planner, if coded with a generic API to connect to a behaviour library, can be deployed in multiple agents, regardless of their goals or embodiment. For example, the Instinct planner has been successfully used in both robots and agent-based modelling, while POSH-Sharp has been deployed in a variety of computer games [9, 19].

### 2.2 POSH and Instinct

POSH planning is an action-selection system introduced by Bryson [5]. It is designed as a reactive planning derivative of BOD to be used in embodied agents. POSH combines faster response times, similar to reactive approaches for BBAL, with goal-directed plans. Its use of hierarchical fixed representations of priorities makes it easy to visualise in a human, non-expert directed graph and sequentially audit.

Instinct is a lightweight alternative to POSH, incorporating elements from the various variations and modifications of POSH

released over the years [19]. The planner was first designed to run on low resources available on the ARDUINO micro-controller system, such as the one used by the R5 robot seen in Figure 2.

### 2.3 ABOD3

ABOD3 is a substantial revision and extension of ABODE (A BOD Environment), originally built by Steve Gray and Simon Jones. ABOD3 directly reads and visualises POSH, Instinct, and UN-POSH plans. Moreover, it reads log files containing the real-time transparency data emanating from the Instinct Planner, in order to provide a real-time graphical display of plan execution. Plan elements are highlighted as they are called by the planner and glow based on the number of recent invocations of that element. Plan elements without recent invocations dim down over a user-defined interval, until they return to their initial state. This offers abstracted backtracking of the calls, and the debugging of a common problem in distributed systems: race conditions where two or more sub-components constantly trigger and interfere with or even cancel each other. ABOD3 is also able to display a video and synchronise it with the debug display. In this way it is possible to explore both runtime debugging and wider issues of AI Transparency.

The editor provides a user-customisable user interface (UI) in line with the good practices for transparency introduced by Theodorou et al. [17]. Plan elements, their sub-trees, and debugging-related information can be hidden, to allow different levels of abstraction and present only relevant information to the present development or debugging task. The application, as shown in Figure 3, allows the user to override its default layout by moving elements and zooming the display to suit the user's needs and preferences. Layout preferences can be stored in a separate file. We have successfully used ABOD3 in both [20].

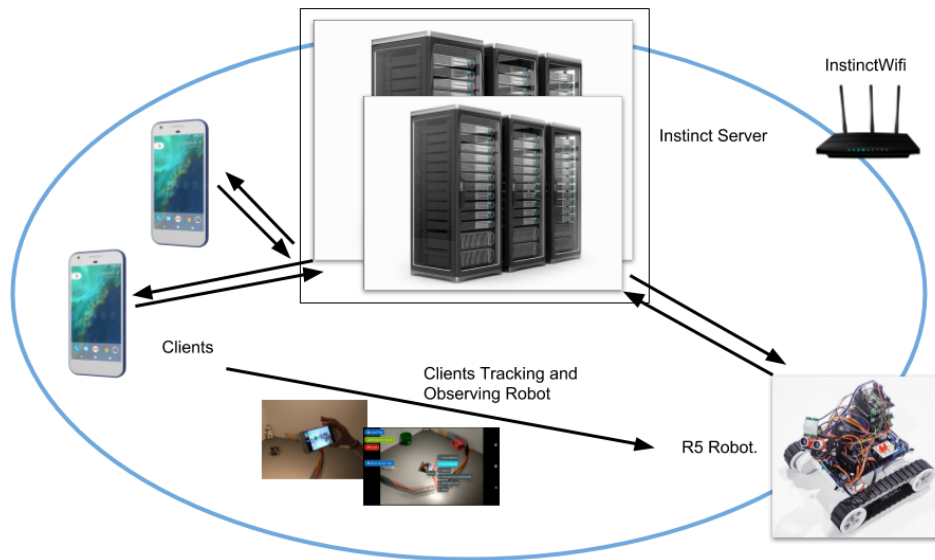
### 2.4 ABOD3-AR

ABOD3-AR builds on the good practice and lessons learned through the extended use of ABOD3. It provides a mobile-friendly interface, facilitating transparency for both end users and experts. In this section, we not only present the final system, but also look at the technical challenges and design decisions faced during development.

**2.4.1 Deployment Platform and Architecture.** The Android Operating System (OS) <sup>1</sup> is our chosen development platform. Due to the open-source nature of the Android operating system, a number of computer vision and augmented reality (AR) libraries exist. Moreover, no developer's license is required to prototype or release the final deliverable. Android applications are written in Java, like ABOD3, making it possible to reuse its back-end code. Unlike the original ABOD3, ABOD3-AR is aimed exclusively for embodied-agents transparency. At the time of writing, Instinct (see Section 2.2) is the only supported action-selection system.

Our test configuration, as seen in Figure 1, includes the tried-and-tested R5 robot. In the R5 robot, the callbacks write textual data to a TCP/IP stream over a wireless (WiFi) link. A JAVA based Instinct Server receives this information, enriches it by replacing element IDs with element names and filters out low-level information,

<sup>1</sup><https://www.android.com/>



**Figure 1: R5 uses a WiFi connection to send the transparency feed to the Instinct Server for processing. Smartphones, running ABOD3-AR, can remotely connect to the server and receive the processed information.**

sending this information any mobile phones running ABOD3-AR. Clients do not necessarily need to be on the same network, but it is recommended to reduce latency. We decided to use this ‘middle-man server’ approach to allow multiple phones to be connected at the same time.

**2.4.2 Robot tracking.** Developing an AR application for a mobile phone presents two major technical challenges: (1) managing the limited computational resources available to achieve sufficient tracking and rendering of the superimposed graphics, and (2) to successfully identify and continuously track the object(s) of interest.

**2.4.3 Region of Interest.** A simple common solution to both challenges is to focus object tracking only within a region of the video feed, referred to as the Region of Interest (ROI), captured by the phone’s camera. It is faster and easier to extract features for classification and sequentially track within a limited area rather than over the full frame. The user registers an area as the ROI, by expanding a yellow rectangle over the robot. Once selected, the yellow rectangle is replaced by a single pivot located at the centre of the ROI.

**2.4.4 Tracker.** Various solutions were considered; from the built-in black-box tracking of *ARCore*<sup>2</sup> to building and using our own tracker. To speed-up development, we decided to use an existing library *BoofCV*<sup>3</sup>, a widely-spread Java library for image processing and object tracking. *BoofCV* was selected due to its compatibility with Android and the range of trackers available for prototyping.

*BoofCV* receives a real-time feed of camera frames, processes them, and then returns required information to the Android application. A number of trackers, or *processors* as they are referred to in *BoofCV*, are available. We narrowed down the choice to the

*Circulant Matrices* tracker [10] and *Track-Learning-Detect* (TLD) tracker [11].

The *Track-Learning-Detect* tracker follows an object from frame to frame by localising all appearances that have been observed so far and corrects the tracker if necessary. The learning estimates the detector’s errors and updates it to avoid such errors, using a learning process. The learning process is modelled as a discrete dynamical system and the conditions under which the learning guarantees improvement are found. However, the TLD is computationally intensive. In our testing we found that when TLD was used the application would crash in older phones, due to the high memory usage.

The *Circulant Matrices* tracker is fast local moving-objects tracker. It uses the theory of *Circulant matrices*, *Discrete Fourier Transform* (DCF), and linear classifiers to track a target and learn its changes in appearance. The target is assumed to be rectangular with a fixed size. A dense local search, using DCF, is performed around the most recent target location. Texture information is used for feature extraction and object description. However, as only one description of the target is saved, the tracker has a low computational cost and memory footprint. Our informal in-lab testing shown that the *Circulant* tracker provides robust tracking.

The default implementation of the *Circulant Matrices* tracker in *BoofCV* does not work with coloured frames. Our solution first converts the video feed, one frame at a time, to greyscale using a simple RGB averaging function. The tracker returns back only the coordinates of the centre of the ROI, while the original coloured frame is rendered to the screen. Finally, to increase tracking performance, the camera is set to record at a constant resolution of 640 by 480 pixels.

**2.4.5 User Interface.** ABOD3-AR renders the plan directly next to the robot, as seen in Figure 2. A pivot connects the plan to the centre

<sup>2</sup><https://developers.google.com/ar/>

<sup>3</sup><https://boofcv.org/>

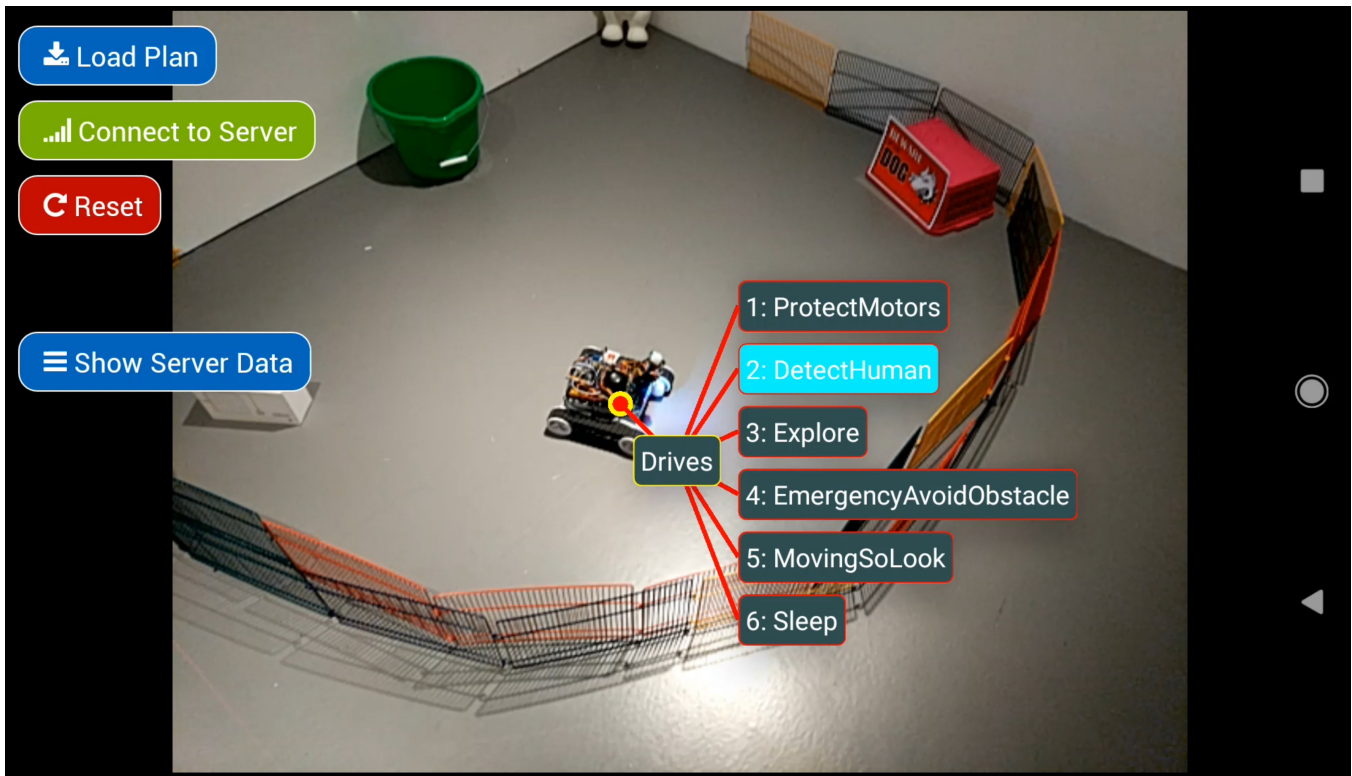


Figure 2: Screenshot of ABOD3-AR demonstrating its real-time debugging functionality. The plan is rendered next to the robot with the drives shown in a hierarchical order based on their priority. The robot here is executing one of its predefined action - detecting for humans and lighting up its LEDs



Figure 3: The ABOD3 Graphical Transparency Tool displaying a POSH plan in debugging mode. The highlighted elements are the ones recently called by the planner. The intensity of the glow indicates the number of recent calls. ABOD3 (used as an IDE for the entire hierarchical cycle) show everything ABOD3-AR uses parts only (2 levels only)

of the user-selected ROI. The PC-targeted version of ABOD3 offers abstraction of information; the full plan is visible by default, but the

user has the ability to hide information. This approach works on the large screens that laptops and desktops have. Contrary, at time of writing, phones rarely sport a screen larger than 15cm. Thus, to accommodate the smaller screen estate available on a phone, ABOD3-AR displays only high-level elements by default. Drives get their priority number annotated next to their name and are listed in ascending order. ABOD3-AR shares the same real-time transparency methodology as ABOD3; plan elements light up as they are used, with an opposing thread dimming them down over time.

Like its ‘sibling’ application, ABOD3-AR is aimed to be used by both end users and expert roboticists. A study conducted by Subin et al. [15] demonstrates how users of AR applications aimed at developers that provide transparency-related information require an AR interface that visualizes additional technical content compared to naive users. These results are in-line with good practices [17] on how different users require different levels of abstraction and overall amount of information. Still, we took these results into consideration by allowing low-level technical data to be displayed in ABOD3-AR upon user request. A user can tap on elements to expand their subtree. In order to avoid overcrowding the screen, plan elements not part of the subtree ‘zoomed in’ become invisible. Subin et al. [15] shows that technical users in an AR application

Question	Group 1 ( $N = 23$ )	Group 2 ( $N = 22$ )	$p$ -value
<b>Dead - Alive</b>	2.39 ( $\sigma=0.988$ )	3.27 ( $\sigma=1.202$ )	<b>0.01</b>
<b>Stagnant - Lively</b>	3.30 ( $\sigma=0.926$ )	4.14 ( $\sigma=0.710$ )	<b>0.02</b>
Mechanical - Organic	1.91 ( $\sigma = 1.276$ )	1.45 ( $\sigma = 0.8$ )	0.158
Artificial - Lifelike	1.96 ( $\sigma = 1.065$ )	1.95 ( $\sigma = 1.214$ )	0.995
Inert - Interactive	3.26 ( $\sigma = 1.176$ )	3.68 ( $\sigma = 1.041$ )	0.211
Dislike - Like	3.57 ( $\sigma = 0.728$ )	3.77 ( $\sigma = 1.02$ )	0.435
<b>Unfriendly - Friendly</b>	3.17 ( $\sigma=1.029$ )	3.77 ( $\sigma 0.869$ )	<b>0.041</b>
Unpleasant - Pleasant	3.43 ( $\sigma 0.788$ )	3.77 ( $\sigma 1.066$ )	0.232
Unintelligent - Intelligent	3.17 ( $\sigma=0.937$ )	3.14 ( $\sigma=1.153$ )	0.922
Bored - Interested	3.80 ( $\sigma=0.834$ )	4.19 ( $\sigma=0.680$ )	0.110
Anxious - Relaxed	4.15 ( $\sigma=0.933$ )	3.81 ( $\sigma=1.167$ )	0.308

**Table 1: ABOD3-AR Experiment: Means (SD) of the ratings given by each group at various questions. The results show that participants in Group 2 perceive the robot as significantly more *alive* if they had used ABOD3-AR compare to participants in Group 1. Moreover, participants in the no-app condition described the robot as more *stagnant* compare to the ones in Group 2. Finally, in the ABOD3-AR condition, participants perceived the robot to be *friendlier* than participants in Group 1.**

prefer to have low-level details. Hence, we added an option to enable display of the Server data, in string format, as received by ABOD3-AR.

### 3 USER STUDY

A user study was carried out to investigate the effectiveness of ABOD3-AR. This took place at the University of Bath in an open public space. The study ran over five days. The principle hypothesis of this experiment is that observers of a robot with access to ABOD3-AR will be able to create more accurate mental models. In this section, we present our results, and discuss how ABOD3-AR provides an effective alternative to ABOD3 as a means to provide robot transparency. Moreover, we argue that our results demonstrate that the implementation of transparency with ABOD3-AR increases not only the trust towards the system, but also its likeability.

The R5 robot is placed in a small pen with a selection of objects, e.g. a plastic duck. The participants are asked to observe the robot and then answer our questionnaires. The participants are split in two groups; Group 1 used the AR app and Group 2 did not use the app. Participants are asked to observe the robot for at least three minutes. A total of 45 participants took part in the experiment ( $N = 45$ ). The majority of users were aged 36 to 45. Each group had same number of females and males. Although they worked regularly with computers, most of them did not have a STEM background – This was the main difference with participants in previous research [18].

The *Godspeed questionnaire* by Bartneck et al. [2] is used to measure the perception of an artificial embodied agent with and without access to transparency-related information. These are standard questions often used in research regarding Human Robot Interaction (HRI) projects, and also used in similar research [14]. We used a Likert scale of 1 to 5 as in Bartneck et al. [2].

#### 3.1 Results

Individuals who had access to ABOD3-AR were more likely to perceive the robot as *alive* ( $M = 3.27$ ,  $SD = 1.202$ ) compare the ones without access to the app;  $t(43) = -0.692$  and  $p = 0.01$ . Moreover,

participants in the no-transparency condition described the robot as more *stagnant* ( $M = 3.30$ ,  $SD = 0.926$ ) compare to the ones in Group 2 ( $M = 4.14$ ,  $SD = 0.710$ ) who described the robot as *Lively*;  $t(43) = -3.371$ ,  $p = 0.02$ . Finally, in the ABOD3-AR condition, participants perceived the robot to be *friendlier* ( $M = 3.17$ ,  $SE = 1.029$ ) than participants in Group 1 ( $M = 3.77$ ,  $SE = 0.869$ );  $t(43) = -2.104$ ,  $p = 0.041$ . No other significant results were reported. These results are shown in Table 1.

#### 3.2 Discussion

We found a statistically significant difference ( $p$ -value  $< 0.05$ ) in three Godspeed questions: *Dead/Alive*, *Stagnant/Lively*, and *Unfriendly/ Friendly*. The R5 has connecting wires and various chipsets exposed. Yet, participants with access to ABOD3-AR were more likely to describe the robot as *alive*, *lively*, and *friendly*. All three dimensions had mean values over the ‘neutral’ score of 3. Although not significantly higher, there was an indicatively increased attribution of the descriptors *Interactive* and *Pleasant*; again both with values over the neutral score. At first glance, these results suggest an increase of anthropomorphic – or at least biologic – characteristics. However, transparency decreased the perception of the robot being *Humanoid* and *Organic*; both characterizations having means below the neutral score.

Action selection takes place even when the robot is already performing a lengthy action, e.g. moving, or when it may appears ‘stuck’, e.g. it is in Sleep drive to save battery. These results also support that a sensible implementation of transparency, in line to the principles set by Theodorou et al. [17], can maintain or even improve the user experience and engagement.

An explanation for the high levels of *Interest* (3.8 mean for Group 1 and 4.19 mean for Group 2) is that embodied agents – unlike virtual agents – are not widely available. Participants in both groups may have been intrigued by the ideal of encountering a real robot. Nonetheless, our findings indicate that transparency does not necessary reduces the utility or ‘likeability’ of a system. Instead, the use of a transparency display can increase the utility and likeability of a system.



There are several characteristics of augmented reality that makes it a promising platform to provide transparency information for both industrial and domestic robots. These include the affordability of AR enabled devices, its availability on multiple platforms such as mobile phones and tablets, the rapidly increasing progress in mobile processors and cameras, and the convenience of not requiring headsets or other paraphernalia unlike its competitor virtual reality.

#### 4 CONCLUSIONS AND FUTURE WORK

In this paper we presented a new tool, ABOD3-AR, which runs on modern mobile phones to provide transparency-related information to end users. Our tool uses a purpose-made user interface with augmented-reality technologies to display the real-time status of any robot running the Instinct planner.

As far as we are aware this is the first use of mobile augmented reality focusing solely on increasing transparency in robots and users' trust towards them. Previous research regarding transparency in robots relied on screen and audio output or non real-time transparency. Building upon past research, we provide an affordable, compact solution, which makes use of augmented reality.

The results from a user study presented in this paper demonstrate how ABOD3-AR can be successfully used to provide real-time transparency to end users. Our results demonstrate how naive users calibrate their mental models and alter their perception of a robot as its machine nature is made cleared. Moreover, they indicate that participants with access to ABOD3-AR has higher interest to the system; potentially increasing its utility and user engagement.

The work presented in this paper is part of a research programme to investigate the effects of transparency on the perceived expectations, trust, and utility of a system. Initially this is being explored using the non-humanoid R5 robot and later we plan to expand the study using the *Pepper* humanoid robot manufactured by Soft-Bank Robotics. We argue that humanoid appearance will always be deceptive at the implicit level. Hence, we want see how explicit understanding of the robot's machine nature effects its perceived utility. Moreover, if transparency alters trust given to the machine by its human users.

Planned future work also aims at improving the usability of the application further. Currently, the robot-tracking mechanism requires the user to manually select an area of ROI which contains the robot. Future versions of ABOD3 - AR would skip this part and replace it with a machine learning (ML) approach. This will enable the app to detect and recognize the robot by a number of features, such as colour and shape. The app will also be enhanced to be able to retrieve the robot type and plan of execution from a database of robots.

#### REFERENCES

[1] Ronald T. Azuma. 1997. A Survey of Augmented Reality. *Presence: Teleoper. Virtual Environ.* 6, 4 (Aug. 1997), 355–385. DOI: <http://dx.doi.org/10.1162/pres.1997.6.4.355>

[2] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. *International Journal of Social Robotics* 1, 1 (01 Jan 2009), 71–81. DOI: <http://dx.doi.org/10.1007/s12369-008-0001-3>

[3] Margaret Boden, Joanna Bryson, Darwin Caldwell, Kerstin Dautenhahn, Lilian Edwards, Sarah Kember, Paul Newman, Vivienne Parry, Geoff Pegman, Tom Rodden, Tom Sorrell, Mick Wallis, Blay Whitby, and Alan Winfield. 2017. Principles of robotics: regulating robots in the real world. *Connection Science* 29, 2 (2017), 124–129. DOI: <http://dx.doi.org/10.1080/09540091.2016.1271400>

[4] R. A. Brooks. 1991. New Approaches to Robotics. *Science* 253, 5025 (1991), 1227–1232. DOI: <http://dx.doi.org/10.1126/science.253.5025.1227>

[5] Joanna J. Bryson. 2001. *Intelligence by Design : Principles of Modularity and Coordination for Engineering Complex Adaptive Agents*. Ph.D. Dissertation.

[6] Joanna J. Bryson. 2003. Action Selection and Individuation in Agent Based Modelling. In *Proceedings of Agent 2003: Challenges in Social Simulation*, David L. Sallach and Charles Macal (Eds.). Argonne National Laboratory, Argonne, IL, 317–330.

[7] J. Bryson, Joanna and Andreas Theodorou. 2019. How Society Can Maintain Human-Centric Artificial Intelligence. In *Human-centered digitalization and services*, Marja Toivonen-Noro, Evelina Saari, Helinä Melkas, and Mervin Hasu (Eds.).

[8] Mary T. Dzindolet, Scott A. Peterson, Regina A. Pomranky, Linda G. Pierce, and Hall P. Beck. 2003. The role of trust in automation reliance. *International Journal of Human Computer Studies* 58, 6 (2003), 697–718. DOI: [http://dx.doi.org/10.1016/S1071-5819\(03\)00038-7](http://dx.doi.org/10.1016/S1071-5819(03)00038-7)

[9] Swen Gaudl, Simon Davies, and Joanna J. Bryson. 2013. Behaviour oriented design for real-time-strategy games: An approach on iterative development for STARCRAFT AI. *Foundations of Digital Games Conference* (2013), 198–205.

[10] João F Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista. 2012. Exploiting the circulant structure of tracking-by-detection with kernels. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 7575 LNCS. 702–715. DOI: [http://dx.doi.org/10.1007/978-3-642-33765-9\\_50](http://dx.doi.org/10.1007/978-3-642-33765-9_50)

[11] Zdenek Kalal, Krystian Mikolajczyk, and Jiri Matas. 2011. Tracking-Learning-Detection. *IEEE transactions on pattern analysis and machine intelligence* 34, 1 (2011), 1409–1422. DOI: <http://dx.doi.org/10.1109/TPAMI.2011.239>

[12] John D. Lee and Neville Moray. 1994. Trust, self-Confidence, and operators' adaptation to automation. *International Journal of Human - Computer Studies* 40, 1 (jan 1994), 153–184. DOI: <http://dx.doi.org/10.1006/ijhc.1994.1007>

[13] Joseph B Lyons. 2013. Being Transparent about Transparency : A Model for Human-Robot Interaction. *Trust and Autonomous Systems: Papers from the 2013 AAAI Spring Symposium* (2013), 48–53.

[14] Maha Salem, Gabriella Lakatos, Farshid Amirabdollahian, and Kerstin Dautenhahn. 2015. Would You Trust a (Faulty) Robot?: Effects of Error, Task Type and Personality on Human-Robot Cooperation and Trust. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '15)*. ACM, New York, NY, USA, 141–148. DOI: <http://dx.doi.org/10.1145/2696454.2696497>

[15] E. K. Subin, Ashik Hameed, and A. P. Sudheer. 2017. Android based augmented reality as a social interface for low cost social robots. In *Proceedings of the Advances in Robotics on - AIR '17*. ACM Press, New York, New York, USA, 1–4. DOI: <http://dx.doi.org/10.1145/3132446.3134907>

[16] Andreas Theodorou. 2017. ABOD3: A graphical visualisation and real-time debugging tool for bod agents. In *CEUR Workshop Proceedings*, Vol. 1855. 60–61.

[17] Andreas Theodorou, Robert H. Wortham, and Joanna J. Bryson. 2017. Designing and implementing transparency for real time inspection of autonomous robots. *Connection Science* 29, 3 (jul 2017), 230–241. DOI: <http://dx.doi.org/10.1080/09540091.2017.1310182>

[18] Robert H Wortham, Andreas Theodorou, and Joanna J Bryson. 2017. Improving robot transparency: real-time visualisation of robot AI substantially improves understanding in naive observers. <http://www.ro-man2017.org/site/> IEEE RO-MAN 2017 : 26th IEEE International Symposium on Robot and Human Interactive Communication ; Conference date: 28-08-2017 Through 01-09-2017.

[19] Robert H. Wortham, Swen E. Gaudl, and Joanna J. Bryson. 2018. Instinct: A biologically inspired reactive planner for intelligent embedded systems. *Cognitive Systems Research* (2018). DOI: <http://dx.doi.org/https://doi.org/10.1016/j.cogsys.2018.10.016>

[20] Robert H. Wortham, Andreas Theodorou, and Joanna J. Bryson. 2017. Robot transparency: Improving understanding of intelligent behaviour for designers and users. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 10454 LNAI (2017), 274–289. DOI: [http://dx.doi.org/10.1007/978-3-319-64107-2\\_22](http://dx.doi.org/10.1007/978-3-319-64107-2_22)