

CHATTY FACTORIES: A VISION FOR THE FUTURE OF PRODUCT DESIGN AND MANUFACTURE WITH IOT

P. Burnap^{1}, D. Branson III², D. Murray-Rust^{3*}, J. Preston⁴, D. Richards⁵, D. Burnett⁵, N. Edwards¹, R. Firth⁴, G. Gorkovenko³, M. A. Khanesar², M. Lakoju¹, T. Smith², J Thorp⁵*

¹Cardiff University, Cardiff, UK

²University of Nottingham, Nottingham, UK

³University of Edinburgh, Edinburgh, UK

⁴University of Essex, Essex, UK

⁵University of Lancaster, Lancaster, UK

*corresponding author burnapp@cardiff.ac.uk

Keywords: IoT, Design, Manufacturing, Ethnography, Robotics, Cybersecurity, Artificial Intelligence

Abstract

Chatty Factories is a three-year investment by the Engineering and Physical Sciences Research Council (EPSRC) through its programme for New Industrial Systems. The project explores the transformative potential of placing IoT-enabled data driven systems at the core of design and manufacturing processes. The research focuses on the opportunity to collect data from IoT-enabled sensors embedded in products during real-time use by consumers, explores how that data might be immediately transferred into usable information to inform design, and considers what characteristics of the manufacturing environment might optimise the response to such data. The project also considers implications arising for skills development in the education sector as well as ethics in manufacturing. In this paper we provide a vision for future “Chatty Factories”.

1 Introduction

The Internet of Things (IoT) is expected to transform digital communication, with 29 billion devices expected to be connected by 2022. The Industrial Internet of Things (IIoT) is reported to be the largest investment sector for IoT¹, which raises the question of *how* future industrial systems can benefit from the vast volumes of data that will be created. This project has been conceived with a long-view towards developing factories of the future in 2030 and beyond. Therefore, we are not aiming to address manufacturing problems in factories as they currently exist, such as process optimisation and predictive maintenance. Rather, we are envisioning the opportunities provided by IoT and novel forms of data analytics that could exist in factories beyond 2030.

In current manufacturing systems, taking a product to market requires multiple discrete, highly specialised activities across research, design and manufacturing disciplines. The Chatty Factories vision for the manufacturing factory of the future is to take these traditionally discrete activities and collapse them into one seamless process that is capable of real-time continuous product refinement at an industrial scale. Part of achieving this requires us to dissolve traditional academic boundaries between the fields of ethnography, design, computer science, engineering and advanced manufacturing.

IIoT offer the opportunity for new insights drawn using real-time data harvested from sensors embedded in products. We see potential to accelerate product refinement through a radical product interruption of ‘consumer sovereignty’ based

around surveys and market research, to ‘use sovereignty’ and an embedded understanding of consumer behaviour – making products that are fit for purpose based on how they are used. Achieving continuous product refinement in response to real-time data on product use requires: data driven systems that provide an auditable, secure flow of information between all operations inside and outside the factory; digital tools that can interact with physical and virtual life; faster reconfiguration and reskilling of the human and robotic elements of the factory floor; seamless communication between product designers, robotic processes and factory floor production; and identification of new creative roles in the labour market to support new systems.

The aims for future Chatty Factories are therefore:

- Enabling new forms of agile engineering product development via “chatty” products that continue to relay their experiences—from the consumer world back to designers and product engineers—through the mediation provided by embedded sensors, IoT and data driven design tools.
- Enabling the manufacturing environment to dynamically respond to changes in terms of physical configuration and ethical reskilling of production elements such as robots and humans by utilising theory from exopedagogy and interpretable data analytics.
- Confronting the challenges of data volume, privacy and cybersecurity to develop an access- controlled manufacturing ecosystem in which product use data will be collated, analysed and disseminated across the factory floor, merging Operational Technology (OT) on the factory floor with Information Technology (IT) in the business.

¹ <https://www.i-scoop.eu/internet-of-things-guide/iot-spending-2020/>

In this paper we will outline the fundamental research challenges that the project will address in its lifetime. We do not expect to see full realisation of the Chatty Factories vision within the project, but aim to develop proof of concept and develop the foundational research while rallying the research community around these grand challenges to enable the UK to be the forefront of future industrial systems.

2. Chatty Factories – Research Objectives

In the lifetime of this Chatty Factories project, we have three strategic objectives that lay the foundations for the achievement of the longer terms aims described above. These objective are expanded in this section:

- Pioneer secure IoT data use to understand how products are used in the wild
- Transform the product design cycle by continually channelling and visualising product use data and improving design-manufacturing links inside the factory
- Deliver more fluid evolution of products on the factory floor through systems more capable of continuous change and informed by better human-robot learning partnerships.

2.1: Pioneering secure IoT data use to understand how products are used in the wild

2.1.1: Data collection

IoT products will be embedded with various sensors that can transmit live data about their use and environment. Sensors are already regularly embedded in a wide range of products, and we anticipate by 2030 they will be much smaller and measure many more attributes of the product experience, vastly widening the types of product that can be part of the Chatty Factories concept. A key challenge within the Chatty Factories project is how to gather and make use of this data in order to aid design and manufacturing research. Gathering data from use in the wild can limit the need for current product research studies, such as post-hoc lab studies and user experience surveys. Furthermore, in-situ research methods can capture the intricate and messy relationships between people and products, which may be difficult to assess through lab-studies [1].

The *Product Data* block in Figure 1 shows the overall vision of how data would be collected from products “in the wild” via IoT infrastructure - that will inevitably change over time - but for which we will develop foundational requirements. These include security and privacy policies, common data structures such that data can be annotated and reused, and means for two-way communication with consumers (see 2.1.2). We will also develop scalable infrastructure to handle the volume of data and reduce unnecessary data storage. Part of the reduction in noise is addressed in research relating to the use of the data itself (see 2.1.2).

2.1.2: Visualising and annotating

Sensor data can highlight issues with products and how use varies across locations, times, user needs and cultures. This knowledge can support the development of customised

products that cater to specific user groups. Within the Chatty Factories project, we propose a new method for gaining contextually grounded insights about product use in the wild through continuous sensor data gathering and limited but timely interactions with consumers.

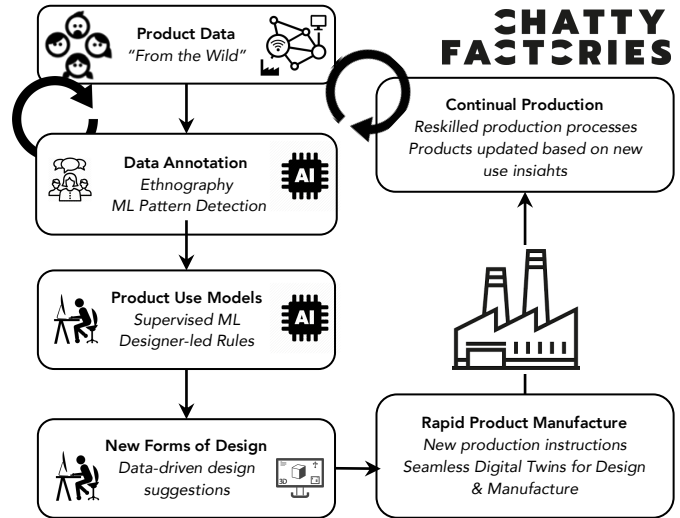


Figure 1 – The Chatty Factories Vision

This is shown in the *Data Annotation* block of Fig 1. The method will work by exploring *artefacts* within the data - clear points of interest in the product use that warrant further investigation. As suggested by Figure 1, artefacts may be identified (i) by product ethnographers - people who study the use data through visualisation and data exploration methods, and (ii) using artificial intelligence methods such as unsupervised machine learning (e.g. clustering and self-organising competitive learning methods) to identify trends and anomalies. Subject to privacy and security considerations, ground-truthing sensor data, and establishing appropriate modes for communication with users - we will enable product ethnographers to interact with consumers to conduct data-driven research into why and how these artefacts occurred, gaining insights into how the product could be better designed to fit its use. The circular arrow in Figure 1. shows how product use and ethnographic interaction will be cyclical and continue over the lifetime of the product to improve consumer experience through continual design change suggestions. A major reason for the ethnographer-consumer interaction is to assign some ‘meaning’ to the artefact (i.e. an explanation for the trend or anomaly in the data). Ethnographers will be able to label artefacts to attribute context later on in the process (see 2.1.3). Example labels include *drops*, *cracks*, *unusual environmental exposure*, *unexpected orientation* etc.

This new research method leverages the experiences and needs of consumers, but it also has the capacity to be invasive. To alleviate these ethical considerations, the project will investigate ways to gain meaningful insights without the use of sensors, such as GPS, cameras, and microphones. Instead we will ground truth the data through the proposed

interactions with users. Users can be queried about the activity that is of interest to designers, live as it is happening. The machine learning element will be tailored to identify novel artefacts but limit these interactions. This combination of live data, machine learning algorithms, and communication with consumers will enable designers and researchers to understand the motivations and issues around product use.

2.1.3: Model building

Once data use artefacts have been identified and assigned a label by ethnographers, we will use the labels and associated data to train supervised machine learning algorithms. This is shown in the *Product Use Models* block of Figure 1. The rationale behind this is to automatically identify points of interest in data as it streams into the factory. If we train models to detect *drop*, *crack* or *extreme temperature exposure* instances, this can be flagged, and the product designers can be made aware of how regularly these things happen. In addition to the supervised machine learning, designers may use this information to drive the placement of sensors on products with a view to design innovations. Take the *crack* scenario, embedded sensors may detect this, but the designer may wish to know how this affects the product. In 2.2 we will discuss Digital Twins as a method to utilise product data. We propose to overlay product use data onto a Digital Twin of the product design and map sensor data from IoT onto the Twin. Acoustic sensors are an example of how the *crack* artefact can drive new sensor placement on the product based on regular events detected in incoming data. Future alerts to *crack* artefacts can then be mapped back to the Twin and design tweaks can be made to reduce these in future.

2.1.4: Cyber Security

Cyber security is a critical issue surrounding the implementation and adoption of innovative products and services. Threat actors will capitalise on the evolving nature and complexities that come with technological innovations such as the Internet of Things (IoT). Unless there is a clear focus from the start on understanding new vulnerabilities and mitigating them by design, threat actors can disrupt safety critical control systems by manipulating vulnerabilities in software, communication channels, security policies, and even social engineering.

We will conduct a thorough risk assessment of the Chatty Factories vision; examining data flows within and between systems, and identifying potential threats to confidentiality, integrity and availability (CIA). The project will explore the implications of opening the factory floor to IoT and develop a security architecture using a defence-in-depth approach, which employs industry best practices and standards (e.g. NCSC² and the Purdue model [2]). For example, leveraging Intrusion Detection Systems (IDS), access control and secure communication channels, and developing methods to detect malicious activity in the pipeline from product to factory floor to identify activities that compromise the integrity of the data flow from end to end.

Risk factors additional to CIA that also require management relate to transparency and trust in the system. Collapsing the previously distinct design and manufacturing process to rapidly push design changes through to manufacturing will require a clear pathway of decision making and potentially the ability to expose the logic of the decision should anything go wrong in the end product. For instance, if data drives a new design change, which subsequently fails at the production stage, how do the design and manufacturing team reverse engineer the decision and roll it back to a previous successful state?

2.2: Transforming the product design cycle by continually channelling and visualising product use data and improving design-manufacturing links inside the factory

2.2.1: Data to design

Applying real-world product usage data to improve product designs is not in-itself novel. Formula one cars already contain sensors that provide automotive designers and engineers with data to improve racing performance. However, the ability to utilise real-world usage insights in the *consumer market at scale* will enable two significant opportunities for innovation in data-driven product design [3].

The first opportunity is to enable *continuous product refinement* to cut manufacturing costs, and improve product lifespan, sustainability, and user satisfaction. In this situation designers will be able to better understand real use of products, refine design solutions, and manufacturing instructions accordingly. For example, reinforcing parts of designs that are most susceptible to failure, and iteratively testing micro-product variations via real-world AB testing ‘in the wild’ to fine-tune designs. Notably, over time this could result in product speciation [4], whereby standard design solutions are replaced by an explosion of data-driven tailored solutions to suit different customers, use cases, and/or geographic locations.

The second opportunity is to enable *new forms of disruptive innovation*, allowing businesses to leverage usage insights of current products to identify trends and new value propositions that can be fed into early stage concept design processes to help inspire new ideas for future products. Giving them the advantage of being first to market in this new space.

2.2.2: Visualising and filtering

Critically, to deliver continuous product refinement and disruptive innovation designers must be able to create and test new ideas. As per the *New Forms of Design* block of Figure 1, the capability to ask questions of the data, reveal new insights, and adapt product sensors to answer difficult questions will be core to enabling new forms of data-driven design.

A challenge in enabling IoT-based data-driven design is the designer’s ability to make sense of, and meaningfully query the data. Referencing the *crack* scenario discussed in 2.1.3, if a designer is trying to eliminate cracking in a product then they might like to ask the following questions of the data. Where are the cracks emerging? Are there any specific events

² <https://www.ncsc.gov.uk/guidance/security-industrial-control-systems>

that often precede this type of failure? If it is caused by an impact, what is the location of the impact? What is the ambient temperature when this cracking emerges? How does this mode of failure impact the structural integrity of other parts of the design? Some of these questions might be easily answered by looking at the existing data collected from product sensors (e.g. location of cracks). Some might require additional filtering, processing, and/or prediction based on available data to reveal insights (e.g. what events precede the cracking). Others might demand new, additional, and/or replacement of sensors (e.g. what is the temperature during cracking).

2.2.3: *New types of Digital Twins*

A Digital Twin [5] is an electronic replica of a product, factory, smart city or even the human body from which real time data fed-back from installed sensors can be visualized and/or predictions made. It emulates the real world by replicating physical laws that impact a process. Digital twins of manufacturing systems enable evaluation of the processes to correct production lines before implementation in the real world [6], thereby saving costs, ensuring quality and informing design at an early stage.

In the Chatty Factories model, required changes in the product and consequent changes on the factory floor will first be tested in their respective digital twins and results fed back simultaneously to multiple stakeholders. This enables complete inspection of the production process and its effect on the product, saving cost resources in time, components and energy. Furthermore, such digital twins can also be utilised to train personnel in virtual environments, while providing valuable feedback from them on the production process from an expert user standpoint.

While digital twins are common within *high-value* manufacturing domains, they are rarely if ever used within *high-volume product design* domains. Technologies that utilise real-time IoT usage data and manufacturing constraints directly from the factory to inform product design at scale do not yet exist [7], yet will likely provide entirely new opportunities for design-based SMEs.

To enable data-driven product design at scale, the Chatty Factories model requires design-based digital twins that can deal with populations of live products, and that can combine usage data with manufacturing factory-based twins to make informed design changes that can be passed to the factory floor.

2.2.4: *Communication between design and production twins, through to factory floor*

Factors that govern the quality of products (e.g. unit cost, time to produce) are the result of both design and manufacturing decisions. But in reality, there is often a significant disconnect between the decisions made at each point. For example, a designer may suggest a change to a digital CAD model to cut production costs, yet such changes may require an additional assembly step during manufacturing that eliminates the intended saving altogether.

Sophisticated data-driven design at scale will require an understanding for *how design decisions impact both real-time customer use, and manufacturing operations*. This will enable two significant benefits. Firstly, it will allow coordinated innovation throughout the process. In this situation, product use data might be used to increase speed of production in parallel with improving the consumer experience of the product. Secondly, a close connection between design and manufacturing will be required to rapidly push updated product solutions into production on the factory floor.

To reduce the cost of iterations between design and manufacturing, communication between their corresponding digital twins is necessary. This means that the manufacturing process can be tested against constraints imposed by the capabilities and accuracies of production elements. Information on such constraints can be fed to designers to inform their proposed changes before production requests are sent. Moreover, it may be necessary to make the production line more flexible and powerful to grant product modifications. The costs and benefits of such modifications may also be predicted virtually. This reduces the time to production and ensures greater flexibility to the process.

2.3: *Deliver more fluid evolution of products on the factory floor through systems more capable of continuous change and informed by better human-robot learning partnerships.*

2.3.1: *New theories of machine learning and applying these to training and collaborative reskilling of humans and robots*
Traditionally robots and human worked in separate environments with secure cage barriers between them. The more recent introduction of human-safe robots is making it possible for the robots and humans to work collaboratively in the same space while maintaining a safe working environment [8]. This creates an opportunity to benefit from the flexibility, intelligence and dexterity of human workers, while maximizing use of automated systems in highly repetitive tasks. Machine learning techniques including natural language processing, image processing and data analysis as well as artificial intelligence approaches such as optimisation algorithms are now viable methods to deal with challenges in the shop floor. Examples of such challenges collaborative systems may encounter are to determine who does what and when; improve the ability of robot to communicate with, and respond to, expert human collaborators more effectively; and improve robustness of the automated systems to changes in the production floor environment and/or product design and manufacture.

Chatty Factories therefore focuses on human-robot collaborative manufacture enhanced with advanced machine learning, pedagogical theory, interpretable data analytics and direct demonstration to move beyond staged manual refinement and increase productivity. A pedagogical environment (See 2.3.2) will support new ways to transfer knowledge between the human and robot. This is particularly important as human experts may have different methods to perform the same task that must be effectively transferred to

the robot. We therefore seek to develop new machine learning methods to accommodate these variances.

A key challenge is that the kinematic and dynamics of robots are different from those of humans. For instance, in the case of a Snake robot, the movement methods are completely different. Hence, we further seek to optimise automated assistance and transfer these skills to the human knowledge use side. Such dynamic reskilling of production elements on the factory floor, with integrated Digital Twins, will minimise training time to ensure maximum flexibility and efficiency. Results will then enable manufacturing ecosystems that can continuously reskill and reorient the human and robot production elements in response to rapid design changes based on use data - while ensuring compliance within any regulatory structure.

2.3.2: *New theories of pedagogy*

One purpose of this project is to think of new ways in which humans and robots might work and learn together and to consider the implications of this. A starting point is to analyse the limitations of interdisciplinary literatures on Human-Robot Interaction (HRI) and the related field of Human-Computer Interaction (HCI), and Machine learning. A critique of the literature on HRI / HCI is that it is limited in the way in which it considers pedagogical relationships between entities.

Despite their focus on relationships HCI / HRI literature often treats the human and robot as separate and autonomous entities. Other literatures hold alternative views on the best ways to conceptualise humans, robots, and the relationship between them. For example, we could see very different perspectives on human-robot interaction from the application of Actor Network Theory (ANT), which pays attention to the ways in which people and objects learn together in assemblages [9], and Object Oriented Ontology (OOO), which views humans and robots as distinct objects, yet they may also create emergent properties which produce systems that are larger than the sum of their parts [10].

From an *exopedagogical* perspective [11] and from the perspective of other theories such as Donna Haraway's *Cyborg Manifesto* [12], a critique can be drawn of the essentialist assumptions of human nature in the HCI/HRI literature. The cyborg critique of mainstream theory is that the human is perceived to have various human qualities which he/she essentially shares with all other humans; thus it is possible to have a field called 'learning science', which creates models of the ways in which *all* humans learn. Robots are perceived to have robot qualities, which may draw upon these ideas about universal human learning. Exopedagogy looks critically at the 'human type entity', showing how many aspects of the human cannot be captured by human discourses of pedagogy or 'the human' as a category – there is always an excess. This might open up new possibilities for learning with robots, based on an ethos of creativity rather than instrumentalism.

All of these theories suggest future directions for human-robot interactions and pedagogies which are not usually considered in the HCI/HRI literature. This may include

looking at assemblages and emergent properties (collective learning) and cyborg / alien methods of learning which are currently not part of learning science. This involves a conceptual leap in how we are looking at learning and training. As there are multiple theories of 'post-human', or otherwise non-essentialist, pedagogy we will create a typology to conceptualise human-machine learning. We will use the typology to categorise learning in manufacturing environments with different organisational forms in the United Kingdom and United States. This will be the basis of the *socio-technical framework*.

The socio-technical framework and research findings will be used in an iterative process to challenge current ways of thinking around new forms of machine learning that underpin human-robot collaboration (see 2.3.1).

2.3.3: *Effects of changing designs on the manufacturing process*

Changing product designs pose significant challenges to the overall manufacturing process. In traditional manufacturing task assignments to workers are typically rigidly decided prior to production, with limited ability to reorientate factory floor elements. In a continually evolving product design environment though this will hinder the capability of the factory floor to adapt to changes, increasing production costs and time.

In the Chatty Factories model, manufacturing tasks will be completed within a collaborative team composed of multiple element types (e.g. human, robot, automated processes, etc.), with production indicators collected to dynamically quantify capabilities and quality of production across work timelines. Data from these individualized production elements will be used to optimise the full production process by creating a dynamic production plan that obeys regulatory restrictions while adapting to production variations. This starts with an assembly plan automatically generated during the design stage as currently happens. But the plan will then be updated based on quantifiable production data, to enable a fluidly evolving production environment. With the further integration of digital twins (Section 2.2.3), new theories of machine learning (Section 2.3.2), and more efficient software interfaces to the designers (Section 2.2.1) we will speed up the time to launch, and respond to design changes in dynamic production requirements [13].

2.4 *Effects of Chatty Factories on the future labour force and society*

The socio-technical framework (2.3.2) will not only consider process and pedagogy but also values and politics. Through focus groups with industry, government and other stakeholders we will consider the implications of the Chatty Factory model for skills and society. We will examine the implications of the Chatty Factory for human/machine symbiosis, for training and skills and for questions of value creation and distribution.

The Chatty Factories model has particular implications for how we think about skills and vocational education. Rather than vocational education being front loaded, the dynamic factory would involve real time training and reskilling where

distinctions between human, robot and A.I. learning break down. This could have dramatic implications for models of apprenticeship where dynamism would mean that there is no formal end to processes of human training, robots (who would acquire knowledge as production continues) could also be apprentices, and firm-level or industry A.I. / distributed ledger technology might replace the role of government as regulator of training. On the other hand, such systems might lead to co-operative skill share methods in manufacturing.

3 Expected Outcomes

Outcomes from the programme of research include:

- IoT-enabled methods to achieve a real-time view of the way products are used
- Novel computational intelligence and mathematical methods to reduce noise in product sensor data and identifying ‘useful’ data artefacts that inform an understanding of how products are used.
- Novel ethnographic methods to make use of this data in the design process, allowing product designers to quickly and efficiently respond to data as it comes in, creating a more responsive product development cycle and data-driven design innovation.
- Software tools and visualisation methods that will allow creative designers to better understand how products are being used “in the wild” and to exploit these insights to accelerate product innovation
- Software methods for automatically compiling digital designs into “hybrid” manufacturing instructions along with methods for validating the instructions (two-way design/manufacture digital twin communication)
- An understanding of the practical implications of opening up the factory floor by investigating new data security and the risks of converging IoT with Information Technology (IT) and Operational Technology (OT)
- Methods to enable dynamic reskilling of production elements on the factory floor to minimise training time and ensure maximum flexibility and efficiency when fabricating products that are varying in their design.
- Mapping how assumptions about human learning inhibit organisations from developing these values

4 Conclusions

With Chatty Factories, we see a future that empowers UK manufacturing companies with the capability to harness AI and ‘Big Data’ (generated from sensors embedded in various types of products) in real-time, integrating this as part of the design and manufacturing process by feeding information that reshapes not only the design but also training the robots and humans on the factory floor – ultimately cutting long-term R&D costs and optimising the production process all wrapped with an intuitive and adaptive IoT/IT/OT Security Architecture.

5 Acknowledgements

This research was funded by the EPSRC under the New Industrial Systems theme: Chatty Factories, EP/R021031/1.

6 References

- [1] Crabtree, A., Chamberlain, A., Grinter, RE, et al. Introduction to the Special Issue of “The Turn to The Wild.” *ACM Trans Comput Interact* 2013, (20), pp. 1–4.
- [2] Green, B., Krotofil, M., Hutchison, D., Achieving ICS Resilience and Security through Granular Data Flow Management. *Proc 2nd ACM Work Cyber-Physical Syst Secur Priv - CPS-SPC ’16*, 2016, pp. 93–101.
- [3] Tao, F., Cheng, J., Qi Q, et al. Digital twin-driven product design, manufacturing and service with big data. *Int J Adv Manuf Technol* 2018;94:3563–3576.
- [4] Sterling, B., *Shaping things*. (MIT Press; 2005).
- [5] Grieves, M., Vickers, J., *DigitalTwin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems (Excerpt)*, 2016
- [6] 'Coppinger, R., E&T', <https://eandt.theiet.org/content/articles/2016/11/digital-twins-cad-design-through-the-looking-glass/>, accessed 27 November 2016
- [7] Lee, J., Bagheri, B., Kao H-A. A Cyber-Physical Systems architecture for Industry 4.0-based manufacturing systems. *Manuf Lett*, 2015, (3), pp. 18–23.
- [8] Bloss, R., Collaborative robots are rapidly providing major improvements in productivity, safety, programing ease, portability and cost while addressing many new applications. *Ind Robot An Int J*, 2016, 43, pp. 463–468.
- [9] Latour, B., *Reassembling the social: An introduction to actor-network-theory*. (Oxford University Press, 2005) .
- [10] Harman, G., *Object-Oriented Ontology*. In: *The Palgrave Handbook of Posthumanism in Film and Television*. London: (Palgrave Macmillan UK, 2015.), pp. 401–409.
- [11] Lewis, TE, Kahn, R V., *Palgrave Connect (Online service)*. *Education out of bounds : reimagining cultural studies for a posthuman age*. (Palgrave Macmillan, 2010).
- [12] Haraway D. *Manifestly Haraway: The Cyborg Manifesto; The Companion Species Manifesto & Companions in Conversation*. (University of Minnesota Press, Minneapolis, 2016).
- [13] 'Brown D, Swift A n. d. - Understanding data, analytics and decision making' , <http://managementdisrupted.com/resources/understanding-data-analytics-and-decision-making/>, accessed 20 February 2019