Industrial Human-Robot Collaboration

Philipsen, Mark Philip; Rehm, Matthias; Moeslund, Thomas B.

Published in:

Proc. FAIM/ISCA Workshop on Artificial Intelligence for Multimodal Human Robot Interaction

DOI (link to publication from Publisher):

[10.21437/AI-MHRI.2018-9](https://doi.org/10.21437/AI-MHRI.2018-9)

Publication date:

2018

Document Version

Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Philipsen, M. P., Rehm, M., & Moeslund, T. B. (2018). Industrial Human-Robot Collaboration. In *Proc. FAIM/ISCA Workshop on Artificial Intelligence for Multimodal Human Robot Interaction* (pp. 35-38) <https://doi.org/10.21437/AI-MHRI.2018-9>

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.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

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

Industrial Human-Robot Collaboration

Mark Philip Philipsen^{1,3}, Matthias Rehm², Thomas B. Moeslund³

¹ Danish Technological Institute, Gregersensvej 9, 2630 Taastrup, Denmark,

² HMI Group, Aalborg University, Rendsburggade 14, 9000 Aalborg, Denmark,

³ VAP Group, Aalborg University, Rendsburggade 14, 9000 Aalborg, Denmark
mpph@create.aau.dk, matthias@create.aau.dk, tbm@create.aau.dk

Abstract

In the future, robots are envisioned to work side by side with humans in dynamic environments both in manufacturing and in societal contexts like health care, education, and commerce. Before this vision can be realized, robots must be socially accepted. Acceptance will have to be built through improvements in robot adaptability and with a gradual introduction, where robots learn on the job. It is our conviction that this can be achieved through a combination of human-robot interaction, multimodal signal processing and AI techniques. We seek to prove this in real world applications. Here follows a discussion on how we intend to utilize end-users in the continuous training and refinement of AIs and the challenges involved in building collaborative production cells.

1 Introduction

The premise for Human-Robot Collaboration (HRC) in the workplace and in contexts like health or elder care, education, commerce, is that robots are envisioned to work side by side with humans in highly dynamic environments. Collaborative production offers a “dream combination” of machine efficiency and human flexibility [Tan *et al.*, 2009]. And with many of the processes still left in human hands, not realistically being fully automated, at least not in the near term. However, in many cases the physically straining and repetitive tasks could be handled by a robot, leaving the more intricate tasks to human workers. The next step for industrial robotics is going to be robots and humans operating side-by-side in close collaborative relationship [Hayes and Scasselati, 2013]. For processes where full automation seems to be within reach, we would argue that humans will have to intervene from time to time, when edge cases are encountered. Our own experience has proven how difficult and time-consuming it can be to achieve the wanted performance for autonomous systems that must deal with significant biological variation. To speed up development we want to rely on the existing workforce for supervising and assisting in the training and refinement of future intelligent systems.

2 Collaboration Between Humans and Robots

Our initial focus is on one-to-one collaboration where one human worker interacts with a single robot. In reality, we envision that we will see all the following combinations of interactions:

- One-to-one: The robot is unable to perform all aspects of a task or has not learned to reliably perform a task. The human coworker performs the tasks the robot is unable to do and/or provide guidance.
- One-to-many: The robots are able to perform tasks autonomously. The human takes on a supervisor role, only interacting with the robots in case a problem occurs.
- Many-to-one/many: Human on human training in human-robot collaboration. Solving tasks that require multiple people and robots.

2.1 Challenges in Human-Robot Collaboration

Achieving close and effective human-robot collaboration is no easy task and many challenges have yet to be overcome [Chandrasekaran and Conrad, 2015; Hayes and Scasselati, 2013; Bauer *et al.*, 2008]. A particular challenge from a Human-Robot Interaction (HRI) perspective is the complex multimodal interplay of different interaction partners with asymmetric communication capabilities and preferences. This creates the need of a multilevel coordination between interaction partners for realizing successful interactions: (i) Coordination on a communicative level: e.g. turn taking, non-verbal alignment [Admoni *et al.*, 2016; St. Clair and Mataric, 2015; Jung *et al.*, 2013]; (ii) Coordination on a physical level: e.g. movements, trajectories, tools [Maeda *et al.*, 2017; Iqbal and Riek, 2016; Karami *et al.*, 2010]; (iii) Coordination on a social level: e.g. group dynamics, emotions [Medina *et al.*, 2017; Keebler *et al.*, 2012; Scheutz *et al.*, 2006]; (iv) Coordination on a task level: e.g. task interpretation, task efficiency, task reasoning [Li *et al.*, 2017; Nikolaidis *et al.*, 2017; Fiore *et al.*, 2016; Bagosi *et al.*, 2016].

All of these coordination tasks require a high degree of contextual and situational adaptability and thus a tight coupling of multimodal signals and AI techniques for signal interpretation, behaviour generation, and learning. Solving these coordination tasks will advance HRC to a level, where

actual collaboration is possible between user and robot and will naturally lead to pro-active instead of reactive robots.

3 A Need for AI

Successfully building robots that can perform tasks with significant variation and complexity, while collaborating with humans will require the use of state-of-the-art methods from HRI, robotics, and machine learning. With the increased availability of sensing technology and advances in machine learning, more data and new methods are now available for renewed automation efforts. High complexity and a need for continued refinement makes conventional programming of control and decision software impractical. We think it is necessary to extensively employ machine learning methods for continuing the expansion of the uses for robots. We have a long way to go and from experience we know that systems relying on machine learning techniques can be significantly challenged when faced with reality. The training of an AI should not end the day it is deployed. The AI should be able to continuously adapt and improve while in use. This is necessary because the environment changes, the product changes and the coworkers change.

3.1 Machine Learning and Feedback

The training of AIs requires feedback. The amount and type of feedback that is available depends on the task. Feedback can be categorized as either; evaluation, e.g. good/bad and numerical rewards, or as correction, which typically comes in the form of annotations, correct answers and demonstrations [Cuayáhuil *et al.*, 2013]. Furthermore, feedback can either be directed towards the past, which is what is most widely used, or towards the future, where it can be seen as guidance [Thomaz *et al.*, 2006]. The often sparse and weak signal from evaluation type feedback is used with reinforcement learning methods, while the more direct correction type feedback can be used with supervised learning methods.

3.2 Transfer of Knowledge and Skills to Robots

An AI can be trained to solve complex tasks either through self-exploration, from expert feedback, or a combination of both. Training AIs to solve most real world tasks using only self-exploration and sparse feedback is not practical. A solution may be to incorporate expert demonstrations and thus circumvent much of the need for exploration encountered with standard RL. By using both demonstrations and typical RL exploration, the need for carefully engineered rewards can be reduced. The combination of the two also allows the AI to keep improving on the demonstrations that were presented to it [Večerík *et al.*, 2017]. The combination of Hindsight Experience Replay [Andrychowicz *et al.*, 2017] and demonstrations has been shown to produce impressive results in a block stacking task where rewards are relatively sparse [Nair *et al.*, 2017]. Expert behaviour can be modelled using generative adversarial networks, where a generator network is trained to produce behaviour that is indistinguishable from the expert [Ho and Ermon, 2016]. Even sparse human feedback providing only a weak yes/no signal, has been shown to provide sufficient signal for a RL algorithm to learn complex behaviour [Christiano *et al.*, 2017].

Robots should be able to learn on the job from people without expertise in ML. The training of AIs should move around the exploration-guidance spectrum to leverage human knowledge when available as well as do exploration and learning on its own. Furthermore, the training interaction can either be lead by the machine, as is the case with active learning where the machine learning algorithm queries the teacher, or it could be a human lead interaction ruled by how and what the human trainer wants to teach [Thomaz *et al.*, 2006].

4 Human Robot Collaboration in the Slaughterhouse

With the remainder of the paper we will present some examples from our work in relation to human robot collaboration in slaughterhouses. The slaughtering of animals differs from most other manufacturing in being a disassembly process where products with natural variation in size and composition must be deconstructed into the products that customers demand. Automation in slaughterhouses is characterized by specialized machines that rely on simple measurements and low hanging fruits have long since been harvested. What remains are tasks where automation has failed to compete with the flexibility and dexterity of human workers. For an overview of the major developments in meat industry since the 1950s look to [Kristensen *et al.*, 2014].

4.1 Cell Production

Current meat production follows a classical production line paradigm and focus has largely been on increasing throughput. However, as the limits of the people and machines along the lines are reached this path to improving competitiveness becomes less attractive. Running a production line at very high speeds increases the risk and impact of breakdowns as well as the workers' risk of developing musculoskeletal disorders. The way forward a transition to a cell based production paradigm where the focus is on flexibility and quality. Cell based production, where the efficiency of machines is combined with the flexibility of human workers, is the solution to rapidly changing marked requirements and the increasing demand for small made-to-order series [Tan *et al.*, 2009]. The idea of cell production in slaughterhouses is currently gaining traction through research projects such as the Norwegian Meat 2.0 project [Animalia, 2018] and the Danish Augmented Cellular Meat Production (ACMP) project [Danish Technological Institute, 2018].

4.2 Augmented Cellular Meat Production

With ACMP we seek to develop collaborative production cells where butchers are assisted through decision support and by having robots perform the most physically straining tasks. The concept will be demonstrated by transforming three slaughterhouses processes that each requires extensive collaboration between human and machine.

Process 1: Picking unwanted elements

We want to assist human quality control workers by augmenting their vision with information from an x-ray machine. This will help identify and remove unwanted elements such

as pieces of bone that may exist on or below the surface of the product. Simultaneously, we want to relieve workers by introducing a coworker robot for performing the pick-and-place task alongside the human workers.

Process 2: Processing pork bellies

The processing of pork bellies includes heavy lifting and requires extensive experience to perform well. Pork bellies must be cut and trimmed to a customer given specification. The most challenging task involves determining how deep to trim a layer of fat to leave it at a specified thickness. Information from a computed tomography scan of the product will be used to produce a visualization showing the butcher where and how much to trim. A robot will take over the task of manipulating the meat such that the butcher can work on every surface. To improve efficiency and ergonomics the robot must take into account the butcher's preferences for the presentation of the product.

Process 3: Multi-functional robot cell

A range of processes can be automated to a large degree. However, because of natural variation or because of the occasional mistakes made in prior processes, the automation is bound to encounter cases that are outside of its capabilities. For these types of processes we envision collaborative production cells where robots execute a pool of tasks such as; tenderloin extraction, cutting of toes and ears, and removal of the head. In problematic cases the robots must allow a human supervisor to guide the robot to a solution or manually perform the operation.

5 Virtual + Physical Robot Cell

We are building a physical demonstration cell with a virtual twin to facilitate the development and evaluation of learning methods. The virtual twin will allow for the collection of feedback and demonstrations, by means of e.g. kinesthetic teaching, through virtual reality. Figure 1 shows an overview of the system.

Figure 1: (Left) Virtual cell, (Right) Physical cell.

For experimentation in HRC we are building a realistic slaughterhouse simulation where participants can interact with virtual meat, robots and tools. See Figure 2.

Figure 2: Human-robot collaboration simulator.

6 Conclusion

In this position paper we have discussed the value of using end-users in the training of robots. We have highlighted some focus areas from our research on Human-Robot Collaboration. Specifically, the requirement for dynamic and flexible coordination between human and robot on several interaction layers (communicative, physical, social, task). We have given some examples of our work in an industry context, where we will rely on AIs for controlling robots in collaborative cells.

References

- [Admoni *et al.*, 2016] Henny Admoni, Thomas Weng, Bradley Hayes, and Brian Scassellati. Robot nonverbal behavior improves task performance in difficult collaborations. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction, HRI '16*, pages 51–58, Piscataway, NJ, USA, 2016. IEEE Press.
- [Andrychowicz *et al.*, 2017] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In *Advances in Neural Information Processing Systems*, pages 5048–5058, 2017.
- [Animalia, 2018] Animalia. Meat2.0. <https://www.animalia.no/no/animalia/om-animalia/arsrapporter-og-strategi/aret-som-gikk-2017/forsker-pa-framtidens-slakterier/>, 2018.
- [Bagosi *et al.*, 2016] Timea Bagosi, Koen V. Hindriks, and Mark A. Neerinx. Ontological reasoning for human-robot teaming in search and rescue missions. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction, HRI '16*, pages 595–596, Piscataway, NJ, USA, 2016. IEEE Press.
- [Bauer *et al.*, 2008] Andrea Bauer, Dirk Wollherr, and Martin Buss. Human–robot collaboration: a survey. *International Journal of Humanoid Robotics*, 5(01):47–66, 2008.
- [Chandrasekaran and Conrad, 2015] Balasubramaniyan Chandrasekaran and James M Conrad. Human-robot collaboration: A survey. In *SoutheastCon 2015*, pages 1–8. IEEE, 2015.

- [Christiano *et al.*, 2017] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, pages 4302–4310, 2017.
- [Cuayáhuitl *et al.*, 2013] Heriberto Cuayáhuitl, Martijn van Otterlo, Nina Dethlefs, and Lutz Frommberger. Machine learning for interactive systems and robots: a brief introduction. In *Proceedings of the 2nd Workshop on Machine Learning for Interactive Systems: Bridging the Gap Between Perception, Action and Communication*, pages 19–28. ACM, 2013.
- [Danish Technological Institute, 2018] DTI Danish Technological Institute. Augmented cellular meat production. <https://www.teknologisk.dk/ydelser/intelligente-robotter-skala-fastholde-koedproduktion-i-danmark/39225>, 2018.
- [Fiore *et al.*, 2016] Michelangelo Fiore, Aurélie Clodic, and Rachid Alami. *On Planning and Task Achievement Modalities for Human-Robot Collaboration*, pages 293–306. Springer International Publishing, Cham, 2016.
- [Hayes and Scassellati, 2013] Bradley Hayes and Brian Scassellati. Challenges in shared-environment human-robot collaboration. *learning*, 8(9), 2013.
- [Ho and Ermon, 2016] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pages 4565–4573, 2016.
- [Iqbal and Riek, 2016] Tariq Iqbal and Laurel D. Riek. Human coordination dynamics with heterogeneous robots in a team. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, HRI ’16, pages 619–620, Piscataway, NJ, USA, 2016. IEEE Press.
- [Jung *et al.*, 2013] Malte F. Jung, Jin Joo Lee, Nick DePalma, Sigurdur O. Adalgeirsson, Pamela J. Hinds, and Cynthia Breazeal. Engaging robots: Easing complex human-robot teamwork using backchanneling. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, CSCW ’13, pages 1555–1566, New York, NY, USA, 2013. ACM.
- [Karami *et al.*, 2010] Abir-Beatrice Karami, Laurent Jeanpierre, and Abdel-illah Mouaddib. Human-robot collaboration for a shared mission. In *Proceedings of the 5th ACM/IEEE International Conference on Human-robot Interaction*, HRI ’10, pages 155–156, Piscataway, NJ, USA, 2010. IEEE Press.
- [Keebler *et al.*, 2012] Joseph Roland Keebler, Florian Jentsch, Thomas Fincannon, and Irwin Hudson. Applying team heuristics to future human-robot systems. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI ’12, pages 169–170, New York, NY, USA, 2012. ACM.
- [Kristensen *et al.*, 2014] L Kristensen, S Støier, J Würtz, and L Hinrichsen. Trends in meat science and technology: The future looks bright, but the journey will be long. *Meat science*, 98(3):322–329, 2014.
- [Li *et al.*, 2017] Yanan Li, Keng Peng Tee, and Shuzhi Sam Ge. *Switchings Between Trajectory Tracking and Force Minimization Force minimization in Human-Robot Collaboration Human-robot collaboration (HRC)*, pages 65–81. Springer International Publishing, Cham, 2017.
- [Maeda *et al.*, 2017] Guilherme J. Maeda, Gerhard Neumann, Marco Ewerton, Rudolf Lioutikov, Oliver Kroemer, and Jan Peters. Probabilistic movement primitives for coordination of multiple human–robot collaborative tasks. *Autonomous Robots*, 41(3):593–612, 2017.
- [Medina *et al.*, 2017] José Ramón Medina, Tamara Lorenz, and Sandra Hirche. *Considering Human Behavior Uncertainty and Disagreements in Human–Robot Cooperative Manipulation*, pages 207–240. Springer International Publishing, Cham, 2017.
- [Nair *et al.*, 2017] Ashvin Nair, Bob McGrew, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Overcoming exploration in reinforcement learning with demonstrations. *arXiv preprint arXiv:1709.10089*, 2017.
- [Nikolaidis *et al.*, 2017] Stefanos Nikolaidis, Swaprava Nath, Ariel D. Procaccia, and Siddhartha Srinivasa. Game-theoretic modeling of human adaptation in human-robot collaboration. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, HRI ’17, pages 323–331, New York, NY, USA, 2017. ACM.
- [Scheutz *et al.*, 2006] Matthias Scheutz, Paul Schermerhorn, and James Kramer. The utility of affect expression in natural language interactions in joint human-robot tasks. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-robot Interaction*, HRI ’06, pages 226–233, New York, NY, USA, 2006. ACM.
- [St. Clair and Mataric, 2015] Aaron St. Clair and Maja Mataric. How robot verbal feedback can improve team performance in human-robot task collaborations. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI ’15, pages 213–220, New York, NY, USA, 2015. ACM.
- [Tan *et al.*, 2009] Jeffrey Too Chuan Tan, Feng Duan, Ye Zhang, Kei Watanabe, Ryu Kato, and Tamio Arai. Human-robot collaboration in cellular manufacturing: Design and development. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 29–34. IEEE, 2009.
- [Thomaz *et al.*, 2006] Andrea Lockerd Thomaz, Cynthia Breazeal, et al. Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance. In *Aaai*, volume 6, pages 1000–1005. Boston, MA, 2006.
- [Večerík *et al.*, 2017] Matej Večerík, Todd Hester, Jonathan Scholz, Fumin Wang, Olivier Pietquin, Bilal Piot, Nicolas Heess, Thomas Rothörl, Thomas Lampe, and Martin Riedmiller. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv preprint arXiv:1707.08817*, 2017.