Smart Manufacturing and The Promotion of Artificially-Intelligent Human-Robot Collaborations in Small- and Medium-sized Enterprises

Megan Zimmerman and Jeremy A. Marvel
U.S. National Institute of Standards and Technology
Intelligent Systems Division
100 Bureau Dr., Stop 8230
Gaithersburg, MD 20899

Abstract

The U.S. National Institute of Standards and Technology (NIST) is developing new metrology toward the evaluation and assurance of collaborative robot performance in next-generation manufacturing. A significant research thrust along these lines involves the advancement of machine learning (ML) and artificial intelligence (AI) to enable the intuitive use and integration of collaborative robots in human-robot and robot-robot teams. This paper discusses a manufacturing-centric perspective of the evolution of human-robot interaction, tools of interest for advancement, and outlines some of NIST's relevant efforts toward imbuing robotic systems with ML and AI to improve operational performance, communication of task-relevant information, situation awareness, and ease of integration.

Introduction

With the rapid advancement of machine learning (ML) and automation technologies, small- and medium-size enterprises (SMEs) are realizing they must adopt "smart" automation technologies to remain competitive. Collaborative robots have revolutionized the robotics market, and are effectively lowering the barrier to entry in bringing automation to otherwise manual processes. Advertised as being safe to work with and around the human workforce, most collaborative robots are simply scaled-down versions of their traditional robotic predecessors. As such, they may still be difficult to integrate and use, re-task, and program to handle operations in increasingly flexible work environments.

To be truly impactful, next generation collaborative robots must not only be safe, but also intuitive to use, and responsive to uncertainty in the workcell, processes, and personnel with which they operate. In other words, they must be more intelligent. This, to some, may be seen as being both positive and negative. On one hand, increasingly intelligent systems are capable of adapting to uncertainty, or even possibly anticipating and preemptively steering away from it. Yet others may view this intelligence as a threat, and the technology as an usurper of the workforce.

This adaptability is not limitless, and at some point assistance from a human operator will be requested. Due to a shift in industrial ethos toward increasingly lean and adaptive technologies, greater emphases are placed on efficiency and effectiveness. The workforce is also evolving, and is increasingly technologically savvy, but also less patient (e.g., (Gursoy, Maier, and Chi 2008)). As reliance on artificial intelligence (AI) and ML increases, the probability that the experience-based, highly-tuned skills necessary to recognize, assess, and correct issues will become a scarce commodity within the factory also increases. At some point, a tipping point may be reached where even more advancements in AI will be required to compensate and replace the expertise of the workers long since disappeared. AI and collaborative robotics must be directed such that advancements of both are structured to enable and support the skills and capabilities of the workforce, and not to be an ad hoc crutch to promote productivity growth. The U.S. National Institute of Standards and Technology (NIST) is developing test methods and metrics to assess and assure the performance of collaborative robotic systems as part of human-robot teams in smart manufacturing environments. In this paper, we outline potential areas to which the application of AI to collaborative robotics would promote both workforce and industry growth. Areas of research include natural human-robot interaction (HRI), adaptation to uncertainty in the work environment, sensor fusion for manufacturing situation awareness, and intuitive integration of robot systems in smart manufacturing workcells.

Natural HRI

While sometimes considered a stop-gap solution bridging purely manual processes with full automation, collaborative operations are intended to leverage the strengths of both robots and human operators to compensate for the limitations of one another. The strength, precision, and repeatability of robots are complemented with the adaptability, intuition, and awareness of their human coworkers. The human-machine interfaces (HMI) connecting humans and robots, however, have been historically limited to hardware-focused task functionality and programming flexibility. The operator's abilities to command, program, and diagnose the systems have largely been secondary.

For the purposes of this paper we define "natural HRI" as HRI enabled with controls and interfaces appealing to intuitive mappings and mechanisms known already to hu-

Official contribution of the National Institute of Standards and Technology; not subject to copyright in the United States.

Figure 1: In addition to providing alternative interfaces for remote teleoperation and offline programming, technologies such as virtual reality provide feedback mechanisms useful for teach-by-demonstration interactions.

mans with limited training. This includes but is not limited to methods of teach by demonstration, natural language control, and systems tailored to human senses and comprehension. These tools are designed to enhance interactions between human robot teams and improve overall system flexibility. However the tools themselves are limited and the systems they are a part of can be improved with integrating AI/ML.

Robot Learning by Demonstration (LbD) is a method for training robots by manually demonstrating the task for the robot. This can be done by either directly moving the robots arms or with some other control system (fig.1), rather than programming motion primitives. LbD uses ML algorithms to learn specific tasks for the robot to accomplish. The task flexibility desired by SMEs requires that many tasks are able to be learned quickly. The resulting programs are not guaranteed to be optimal for long-term production efficiency or quality, but are sufficient for short-term, high-turnover manufacturing processes. ML applications capable of learning more than one task very well would be able to accommodate these needs, but are not yet easily applied.

Natural language processing (NLP) is a mechanism by which humans can intuitively interact with the environment around them. A variety of NLP applications are already present in daily life in mobile and smart-home technologies. These tools are used as control systems interfacing with our physical environment with Internet of Things enabled devices, and are growing in popularity. These same NLP technologies have the potential to improve the cohesiveness of human robot collaboration across a variety of fields, one of the less recognizable being manufacturing. Notable and thorough evaluation of NLP as an interface mechanism may be necessary to justify the adoption of the technology by manufacturers. This evaluation and uptake would be hastened by the advancement of integrated NLP systems capable of evaluating physical intent from casually-spoken sentences.

To promote the development and advancement of HRI and HMI, NIST is developing a framework for evaluating the effectiveness of human-robot collaborations in manufacturing environments. Drawing from the established literature, the framework seeks to address both the objectively quantitative aspects of HRI/HMI, as well as the subjectively qualitative assessments of major stakeholders of manufacturing processes. These metrics, in turn, drive design recommendations to optimize the utility of collaborative robots in human teams.

Adaptation to Uncertainty

As manufacturing practices move away from the monolithic infrastructures of rigid workcells and workflows, and adopt more agile and lean paradigms where processes, products, and logistics are constantly fluctuating, the need for intelligent adaptability becomes increasingly paramount. To many skilled workers, this flexibility is intuitive. Automation technology, however, is less amenable to change. In human-robot teams, a certain percentage of this uncertainty originates from the human operator, to which the robot must adapt on a constant basis.

Minor position and orientation errors may be corrected using motion compliance, but large deviations from preplanned trajectories will require the robot to parameterize and adjust its plans dynamically. While process parameters can be tuned with relative ease using statistical designs of experiments, re-tuning is disruptive to the workflow, and many optimized adjustments will be short-lived. Systems must be able to not only identify when their current parameters are no longer optimal, but also be able to automatically initiate and complete the optimization process in real time to maintain pace with their human coworkers.

Different forms of AI and ML may be applied to different aspects of the optimization process. One potentially effective application of ML in self-optimization, for example, is the automatic generation of models to characterize a system's performance. Such models may be used to capture, and subsequently predict, the manufacturing performance (e.g., assembly time, incurred forces, or the number of successful/quality welds) resulting from different sets of parameters. This enables the system to reject parameters that don't perform well (Marvel et al. 2009), or may be hazardous to human coworkers. Similar approaches are being explored by NIST as metrology tools, where even being able to measure performance in the face of uncertainty has significant impacts on verification, validation, and system safety.

Intelligent Process Sensing

Any robot system must maintain self awareness to ensure continued safe operations. System feedback in the form of collision detection, proprioceptive feedback, and control responsiveness contribute to the robot's physical state. Purpose-driven robot systems also provide additional process monitoring and feedback to ensure workcell functionality. In HRI, such monitoring will necessarily include the human factor in which an observer system tracks operator poses, potential safety violations, the motions of workpieces, and state machine transitions such that the robot may respond to the operator's actions accordingly.

Training sensor-driven observer systems is long, difficult, and expensive; more-so in flexible factory environments, where parts are no longer fixtured or fully controlled. Com-

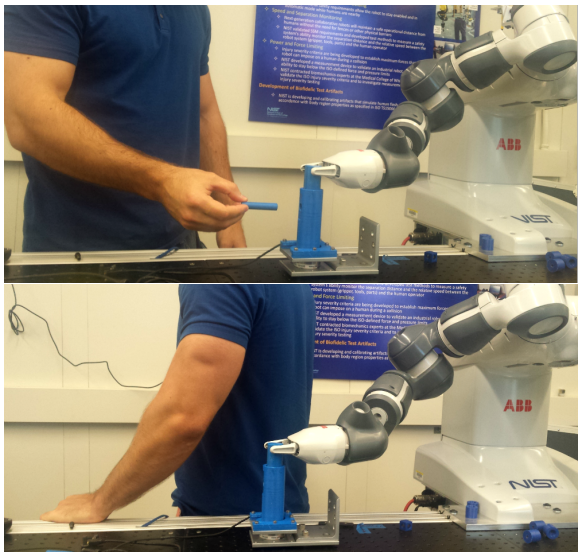

Figure 2: To simultaneously maintain safety and usability robots in human-robot teams need to distinguish between scenarios involving humans. Safety protocols are different in times where the human is in proximity and intends to interact (top), from times when the human is in proximity without intention to interact(bottom).

mon systems like camera-based identification and localization require nuanced expertise for adjusting lighting, camera, lens, and presentation settings to maximize the likelihood sensing with consistent quality. Given this difficulty, commercial vision systems are typically limited to unstructured two-dimensional (2D) scenes, or, at best, structured 3D scenes.

"Deep learning" techniques may be leveraged to automatically compensate for the sensing challenges, and to produce manipulation and control strategies with little knowledge of the robot's capabilities or kinematics (Levine et al. 2016). Given the resources necessary to train these networks, such approaches are impractical for a majority of manufacturing tasks and environments, and all but impossible for most SMEs. However, such research demonstrates the potential for the technology. Paired with even basic image processing algorithms (e.g., edge detection, color segmentation, etc.), signal processing algorithms (e.g., independent component analysis), and statistical estimations of noisy data (e.g., Kalman filters), it is expected that black-box neural network implementations could easily provide reasonable performance for object detection, localization, acquisition, and inspection.

Process sensing extends to the tracking and sensing of humans for human-robot teams, and for non-collaborative tasks in environments where humans are present. Operator position and intent may be tracked for safety and interaction purposes (fig.2). Here, the identification of "intent" may be characterized as objects of attention, travel destinations, and potential actions or tasks within an established manufacturing process. NIST is exploring the technologies and applications to which process sensing contribute to situational awareness, and enhancing flexibility and process performance within human robot teams.

Intuitively-Integrated Robots

With any sufficiently agile HRI application, some frequency of reconfiguration and tuning may be expected. Installing, removing, and re-tasking equipment should be simplified to reduce the workload on the operator without compromising the confidence in the technology performing in human-occupied spaces. This burden encompasses not only the difficulty of calibrating, registering, and reprogramming robotic systems, but also necessarily includes an element of safety, as a full risk assessment and reduction process would be required to ensure no new hazards have been introduced.

Ultimately, humans operators must effectively use and interact with these robots when working in collaborative teams. If a robot cannot be easily integrated into the manufacturing process or team, there is less rationale to do so. Multi-robot teams with differing makes that are unable to easily integrate, will not be adopted. And robots with limited safety and HRI capabilities will not be integrated into human heavy workplaces like SMEs. To enable the adoption of many of these robotic technologies, adaptive interfaces, capable of connecting both robot-robot and human-robot teams, must be prioritized.

One challenge with integrating robots into existing workcells is the post hoc registration of the robots with machine tools, fixtures, and other robots (Marvel et al. 2015). It has been shown that using even simple, unsupervised ML algorithms for localized registrations can provide striking reductions in registration uncertainty (Van Wyk and Marvel 2017). Registration and calibration are only small parts of the greater integration problem; integrated equipment must be able to communicate and coordinate. By default, there is little support across manufacturers for protocols for interfacing robots and systems from different companies. Open interfaces and standards for industrial applications for command and feedback such as ROS-I (Robot Operating System Industrial, (Edwards and Lewis 2012)) and MTConnect (Vijayaraghavan et al. 2008), respectively, may be successfully leveraged to bridge this gap and facilitate this flexibility. Such systems, however, may prove challenging for novice users, particularly when dealing with equipment with little or no system support. An alternative approach is to provide robot-agnostic high-level planning interfaces such as NIST's XML-based Canonical Robot Command Language (Proctor et al. 2016).

Perhaps a more preferable take on the issue would leverage open technologies in adaptive virtual bridges that provide more transparent, plug-and-play functionality between robots, programmable logic controllers, and machine tools. NIST is currently investigating this potential by using statistical analysis of robot state messages (e.g., Cartesian pose of the attached tooling, or joint values) paired with external observer systems (as discussed in the previous section) in an interface that attempts to identify, parse, and validate robot states for process and system situation awareness. While the information exchange is still effectively unidirectional, such

a system would also facilitate natural plug-and-play functionality through the automatic detection of *a priori* defined interfaces. Additionally, NIST is developing metrics for usability evaluation of robotic systems and the "intuitiveness" of already integrated systems.

For any new technology to be adopted in manufacturing, a repeatable level of performance assurance must be given. Adopting new technology is an expensive and risky venture, and prototype systems may only be given a single opportunity to demonstrate their effectiveness. Experimental technologies are rarely even given this chance. Time spent training and retraining neural networks, exploring suboptimal parameters for assembly that result in wasted cycles, or ML algorithms that result in damaged parts, tooling, or robotic equipment are all examples of waste that may signal the premature demise of AI and ML in a given factory. Once spent, such opportunities for AI may not be granted again for several years. One possible mechanism to alleviate concern regarding the effectiveness (and appropriateness) of applying AI/ML to an application is to provide certifications attesting to the instantiations' performances. Currently, there are relatively few standards regarding the certification of software. Those standards that have been established are focused primarily on the safety of mission-critical platforms for military (e.g., (MIL-STD-882E 2000)) and medical (e.g., (IEC 62304 2006)) applications. There are no general-purpose, safety-rated software standards, let alone specifically for manufacturing. Providing standardized test methods and metrics for ML/AI would leverage the rigorousness of the standardization process, and provide tools for the assessment and assurance of system performance. The result would increase confidence in AI/ML technologies, and may have the added benefit of alleviating concerns regarding the dangers of AI (e.g., (Helbing et al. 2017)).

To assist with the certification of AI and ML, functional unit tests (e.g., per (IEC 61508 2010) for safety applications) that evaluate for fault tolerance would be required. However, it is not immediately clear what such unit tests would be. Supervised learning-based systems are relatively easy to evaluate for performance given known, targeted results in a validation set. While such evaluations test only for accuracy, they would also benefit the response characterization in complex environments by identifying performance outliers. Testing for system reliability, however, is significantly more challenging. Most automated manufacturing applications require consistently good performance with $\geq 99.95\%$ uptime (i.e., less than one minute per day of lost productivity). An overwhelming majority of AI/ML systems are not evaluated under continuous, uncertain operating conditions long enough to even measure their reliability.

More importantly, there is reason for concern of operator safety in any AI/ML-driven robot platform. AI is nondeterministic, so overriding functions must be in place to ensure that safety is the highest priority, possibly even so far as to have the robot shut itself down. Although negatively impacting process performance, such actions would help maintain a safe operational environment.

Disclaimer

Certain commercial equipment, instruments, or materials are identified in this paper to foster understanding. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

References

Edwards, S., and Lewis, C. 2012. Applying the robot operating system (ros) to industrial applications. *Proceedings of the IEEE International Conference on Robotics and Automation: ECHORD Workshop*.

Gursoy, D.; Maier, T. A.; and Chi, C. G. 2008. Generational differences: An examination of work values and generational gaps in the hospitality workforce. *International Journal of Hospitality Management* 27(3):448–458.

Helbing, D.; Frey, B. S.; Gigerenzer, G.; Hafen, E.; Hagner, M.; Hofstetter, Y.; van den Hoven, J.; Zicari, R. V.; and Zwitter, A. 2017. Will democracy survive big data and artificial intelligence. *Scientific American. Feb* 25.

2010. IEC 61508 Functional safety of electrical/ electronic/ programmable electronic safety-related systems. Standard, International Electrotechnical Commission.

2006. IEC 62304 Medical device software Software life cycle processes. Standard, International Electrotechnical Commission.

Levine, S.; Pastor, P.; Krizhevsky, A.; and Quillen, D. 2016. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *CoRR* abs/1603.02199.

Marvel, J. A.; Newman, W. S.; Gravel, D. P.; Zhang, G.; Wang, J.; and Fuhlbrigge, T. 2009. Automated learning for parameter optimization of robotic assembly tasks utilizing genetic algorithms. *Proceedings of the IEEE International Conference on Robotics and Biomimetics* 179–184.

Marvel, J. A.; Messina, E.; Antonishek, B.; Van Wyk, K.; and Fronczek, L. J. 2015. Tools for robotics in sme workcells: Challenges and approaches for calibration and registration. *NIST Interagency/Internal Report (NISTIR)-8093*.

2000. MIL-STD-882E Department of Defense Standard Practice - System Safety. Standard, U.S. Department of Defense.

Proctor, F.; Proctor, F.; Balakirsky, S.; Balakirsky, S.; Kootbally, Z.; Kootbally, Z.; Kramer, T.; Kramer, T.; Schlenoff, C.; Schlenoff, C.; et al. 2016. The canonical robot command language (crcl). *Industrial Robot: An International Journal* 43(5):495–502.

Van Wyk, K., and Marvel, J. A. 2017. Strategies for improving and evaluating robot registration performance. *IEEE Transactions on Automation Science and Engineering*.

Vijayaraghavan, A.; Sobel, W.; Fox, A.; Dornfeld, D.; and Warndorf, P. 2008. Improving machine tool interoperability using standardized interface protocols: Mtconnect. *Proceedings of the International Symposium on Flexible Automation*.