What Is the Robot Thinking? Embedded Feedback in the Intelligent Workcell Manufacturing Environment

Harrison Billmers

August 2017
CMU-RI-TR-17-49

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Daniel Huber, Advisor
Aaron Steinfeld
Allison Del Giorno

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Robotics.*

Abstract

Today's manufacturing industry relies on robots to complete tasks that require high speed, precision, and power beyond human capabilities. To ensure safety, these large robotic arms operate in isolated work cells separate from humans and are not designed to communicate with people. Some advanced safety systems, such as the Intelligent Workcell, monitor people's movement within a workcell, allowing people and robots to work together safely in collaborative manufacturing settings. During collaboration, the worker needs to know what the robot is thinking and what it plans to do. We hypothesize that embedded feedback within the environment can effectively communicate the robot's intentions. We developed, tested, and evaluated three visual feedback systems: a real-time display simulation of the virtualized workcell environment, arm-mounted embedded lighting, and an embedded floor display system. We conducted a user study employing a simulated manufacturing task to evaluate each feedback method for collaborative and overall performance. We found that feedback improves collaboration for human-robot assembly tasks by reducing average robot blockage time by at least 12% compared to the control group. Compared to the other feedback methods, the embedded floor display feedback allowed about 20% faster resolution of blockages when detected. Embedded arm-mounted light feedback helped participants preemptively avoid causing blockages. The embedded floor display was the most effective feedback modality evaluated, followed by partial performance gains from both the simulation and the embedded arm-mounted lights. We also provide recommendations for more extensive analysis of feedback systems in future work.

Acknowledgements

I'd like to thank my wonderful advisor, Dr. Daniel Huber,
for his endless guidance and encouragement throughout this project.

I'd like to thank Dr. Reid Simmons,
for his insight throughout the development of the procedures and feedback.

I'd like to thank my committee, Dr. Aaron Steinfeld and Allison Del Giorno,
for their constructive input while completing this document.

Finally, I'd like to thank my parents, Robert and Joanne Billmers, and my sister Carolyn,
for their unwavering support through even the hardest parts of this work.

This material is based upon work supported by
the National Science Foundation under Grant No. IIS-1208598.

Contents

1 Introduction

Existing manufacturing facilities were designed to ensure safety by keeping robots and people separate. Work cells using advanced safety systems that monitor people's movements can allow robots and people to work together. This makes possible a new type of collaboration between humans and robots, in which each worker completes tasks for which they are best suited. To effectively collaborate, however, the person needs to know what the robot is thinking or planning to do. We hypothesize that embedded feedback within the environment can effectively communicate the robot's intentions. This thesis describes the development, testing, and evaluation of several types of embedded and non-embedded feedback targeted at improving collaboration.

1.1 Background

Today, robots are being incorporated into factories around the globe at an increasing rate [1, 2]. In some cases, robots perform operations that are dangerous for humans to perform; in others, the tasks are beyond human capacities. Robots are faster, more precise, and more consistent than humans at certain jobs, such as welding, lifting, and drilling (Example workcell: Figure 1).

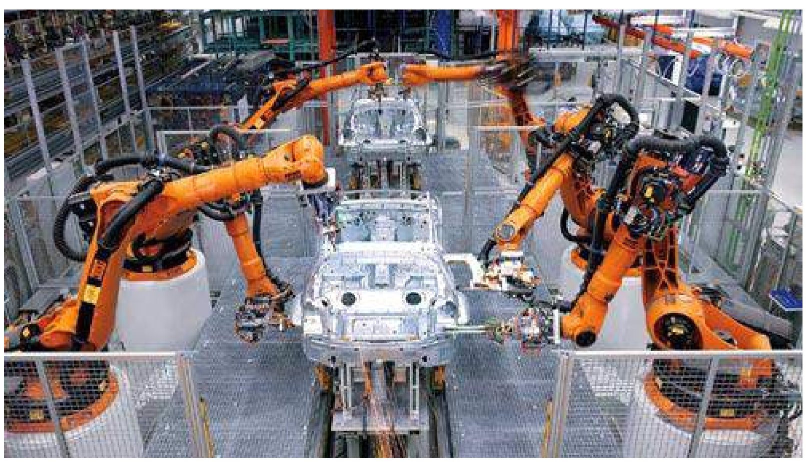

Figure 1: Industrial robots assembling a vehicle [3]

Small- and mid-size manufacturers have recognized that collaboration between human and robot workers on the assembly line allows both to perform tasks for which they are better suited, thereby increasing performance. For instance, robots are better at performing repetitive tasks without fatigue; humans are better at dexterous tasks in confined spaces. Collaborative use of robots on an automotive assembly line task has been proven to reduce human worker idle time by 85% [4].

Ensuring worker safety is critically important when humans and robots are collaborating. We can safeguard human coworkers either by building robots that are intrinsically safe or by building working environments that prevent unsafe events from occurring.

Compliant robots are intrinsically safe; they are built to do little or no harm in collisions with people (Figure 2). These lightweight robots have no pinch points, are back-drivable by pushing, and move slowly at speeds of 1 meter/second or less. Although they must conform to collaborative robot safety standards [5, 6], compliant robots can safely operate uncaged around human workers. However, these safety advantages come at a cost. Minimizing robot gross weight restricts their payload capacities to lightweight objects under 20 kg. Compliant robots have neither high-precision nor speed, limiting the scope of tasks that they can complete.

Figure 2: Rethink Robotics® Baxter compliant robot working on an assembly line [7]

Standard industrial robots are capable of handling heavy payloads and moving with both high-speed and precision. They are not intrinsically safe, requiring their surrounding environments to prevent contact with human workers. OSHA standards dictate that these robots must be separated from human workers either by physical boundaries, such as locked cages, or by electronic boundaries that trigger emergency stops. Limiting devices can also be used to curtail the robot's operating speed, acceleration, and reach [8, 9]. Their sheer size and mass (e.g., ABB's IRB 8700 [10] can reach 4 meters vertically and weighs almost 5 tons) mean that collisions with people can easily cause injury. Collaboration with these robots therefore requires an active safety system.

The Intelligent Workcell creates a safe environment for close-range collaboration with industrial robots. The system generates real-time dynamic and predictive safety zones around the human and the robot by fusing the output of multiple 3D sensors to generate a complete picture of the environment. This dynamic approach provides a collaboration-friendly alternative to physical or static sensor-based barriers between the human and the robot [11–13].

People naturally communicate using verbal and body language when collaborating on a task, but industrial robots working alongside humans do not communicate by either of these means. Verbal communication between robot and human is infeasible in the manufacturing environment, since auditory methods would need to compete with the often noisy surroundings. Industrial robots do not have human characteristics or extra movement freedom to express body language. Outfitting the worker with a communication link to the robot would require additional wearable equipment that can lead to failure resulting from loss, damage, and worker negligence. Without a reliable means of communication, the robot's plans are undisclosed to collaborating human workers. This "language barrier" can impede human-robot collaboration as much as a physical barrier.

Our research focuses on developing embedded channels through which the robot can communicate information about its environment to the worker. Embedded robot feedback could permit human team members to synchronize their movements to the robot's, reducing coworker interference. Embedded feedback could also alert human workers of unexpected changes or problems occurring in the environment. Overall, effective feedback could improve synergy between human and robot workers.

1.2 Research Question

We directed our research efforts toward discovering how embedded feedback can be provided to a human when collaborating with an industrial robot. Feedback can communicate many types of information, such as what the robot is presently doing, planning on doing, and how it is interpreting its environment. Feedback can either be embedded within the environment or non-embedded and exist external to the environment. The research questions that provided the basis for this study were: "How can embedded feedback be provided from the robot to the worker in the Intelligent Workcell environment? How can we objectively measure the collaboration? Is embedded feedback effective at improving collaboration or performance?"

To answer these questions, we developed, tested, and evaluated three methods of visual feedback using a collaborative assembly task. The methods were: a real-time display simulation of the virtualized workcell environment, robot-mounted embedded colored lighting, and a floor-based embedded colored display system. Each feedback modality uses a different delivery method and portrayal of system information. The only general requirement was that all methods provide the operational status of the safety system (actively moving ["running"], safety violation causing the robot to stop ["blocked"], or waiting on the participant to complete an action ["paused"/"waiting"]). We focused only on visual feedback since auditory, haptic, and wearable communication methods posed a number of inherent implementation shortcomings.

We conducted a user study for the evaluation of the three feedback modalities. The experiment consisted of a simulated manufacturing task designed to exemplify predicted collaboration scenarios. Participants were run with either a single feedback modality or no feedback as a control. We evaluated collaboration using a combination of factors, including interference between the worker and robot, idle time, and responsiveness to communication; performance was evaluated using task completion time.

2 Related Work

There has been limited research on communication between industrial robots and human operators. Robots used in manufacturing have historically been sequestered, so research into the area has not been considered impactful. Current HRI research has therefore focused mainly on anticipation of collision and subsequent robot reaction.

Common approaches in the current research literature for preventing human harm during interactive assembly tasks with robots include modalities designed to impact safety through modification of the hardware. In such cases, the mechanical design of the robot itself prevents or reduces collision trauma. Modifications include actuators which reduce impact momentum [14, 15], actuators with magneto-rheological fluid-based joints [16, 17] and robots with reduced mass and rounded, softer materials of construction [18]. Examples of these inherently safe compliant robots include Rethink Robotics® Baxter, KUKA LWR4+ and ABB's Dual Arm Concept Robot.

Other research areas focus on algorithms that mimic or predict the movements of a human worker in the environment. Mainprice and Berenson predicted potential collisions by anticipating the movements of human workers as they perform a task [19] and encoding the uncertainty of the robot regarding human actions. Shah et. al. employed human-robot cross-training methods, in which the human and robot iteratively switch roles, as a team training strategy [20]. In other work, this group encoded separate sets of human and robot actions and incorporated the data into specific rules to predict likely human task-related behavior during a reaching motion [21]. Both Shah and Sisbot have described the use of entropy rate as a risk-assessment tool [22, 23]. Henrich and Kuhn [24] developed an intuitive behavior model, which characterizes tasks according to free or guided motion and gross or fine interaction contexts.

Collision detection based on physical quantities obtained using the robot's existing internal proprioceptors can be used with appropriate algorithms to predict collisions [25]. DeLuca [26–29] and Lacevic and Rocco [30] used quantities such as generalized momentum and total kinetic energy to predict collision and provides directional information for robot reaction following impact. Kulic and Croft [31] generated a measure of danger during interaction or collision based on factors affecting the impact force, but often these methods are robot model specific. Lazzero et. al. developed a non-specific collision detection procedure using motor current values compared to a dynamic model, but this system cannot anticipate collisions [32]. Meguenami [33] developed a safety indicator that quantifies the degree of risk induced by a collision towards a human coworker based upon kinetic energy of the Kuka LWR4 robot.

Addition or modification of external sensing on the robot can provide data input for path planning and collision mediation software. Torque sensors at the joints permit active control of robot impedance and visual sensors on the robot can be used for rapid data analysis [34] and fast, accurate obstacle recognition. Ebert [35] presented an approach using a specialized tracking-vision-chip that triggers high speed emergency stops. Liu, Deng, and Zha developed a planning method based on mapping moving obstacles into the configuration space. The method included pausing to avoid a collision and sounding an alarm to the human worker [36]. However, these methods provide only local information about the moving manipulator arm.

A third technique uses camera systems mounted in the workspace to obtain information about the environment that can then be used to determine path movement in addition to predicting collisions [37]. Ebert and Henrich [38, 39] used four grayscale cameras to back-project directly to the manipulator's configuration space and look-up tables for fusing images and joint angles. Henrich developed a multi-camera 3D collision avoidance system based on difference images [40, 41] and used it to devise an algorithm enabling robots to perform safe pick-and-place operations [42]. Additional work by Henrich [43, 44] as well as DeLuca [26] presented general approaches for robot surveillance using multiple smart cameras and distributed processing. Roth used data from a 3D-PMD camera for fast obstacle recognition [45] to effect robot velocity changes, stops, and trajectory improvements. Kuhn [46] combined extended difference imaging and force torque sensor data to generate robot movements with a calculated maximum safe velocity. The SafetyEye system [25] uses stereo vision to detect moving objects.

Embedded and interactive communication systems designed to warn human workers of potential collisions have been explored. Ogorodnika [47, 48] proposed a wrist-mounted human warning system designed to alert a worker using vibrotactile and visual cues tied into a workcell safety monitoring system. The safety mode controller containing a danger index module receives protocols from the safe expert system monitoring each task. If the extent of anticipated hazard is great enough, signals are sent to the Human Awareness Interface, robot controller, and safeguarding system simultaneously. In 2006, a Human-Robot Interaction Operating System (HRI/OS) was designed that allows robots to ask questions of humans during tasks to improve efficiency of teamwork. The work scenario consisted of independently performed tasks linked together, so that actual physical interaction was not an issue [49]. Recent work by Shah et. al. [50] involving robot-human collaborative tasks allows verbal communication from human to robot when a phase of the task is completed and supply retrieval for the next stage is required. No physical interaction was required for the task, however.

Cardou et. al. recently documented a hybrid workspace which provides feedback about the human's position [51]. The human wears a safety helmet with an IMU and an indoor localization system (RSSI). The data is sent wirelessly to the OPC server and robotic controller to effect safe path planning. In other work [52], the human wears a tracking suit and a UWB localization system. The system guides the robot using visual servoing through a previously defined path. However, the human worker does not receive any cueing about the robot location or operation.

Research in the area of collaboration with standard industrial robots is limited. While workspaces that provide feedback do exist, there remains a need to develop methods that inform the human worker about the robot's movements and intentions in the environment.

3 Methods

In this section, we describe our approaches to communicating feedback, our proposed useful situations for feedback, and the task used for evaluating feedback performance. Before addressing feedback, we first provide a brief overview of the Intelligent Workcell, the platform which this research is built upon.

3.1 The Intelligent Workcell

The Intelligent Workcell [11–13] is constructed around an ABB IRB-2400 [53] robotic manipulator. This 375 kg manufacturing arm resides within a 4.0-by-2.7 meter floor area for assembly tasks. The area encompasses the range of operation of the robot with an additional margin (around half a meter) beyond its reach. The perimeter of the workcell is clearly delineated from the rest of the laboratory floor space (Figure 3).

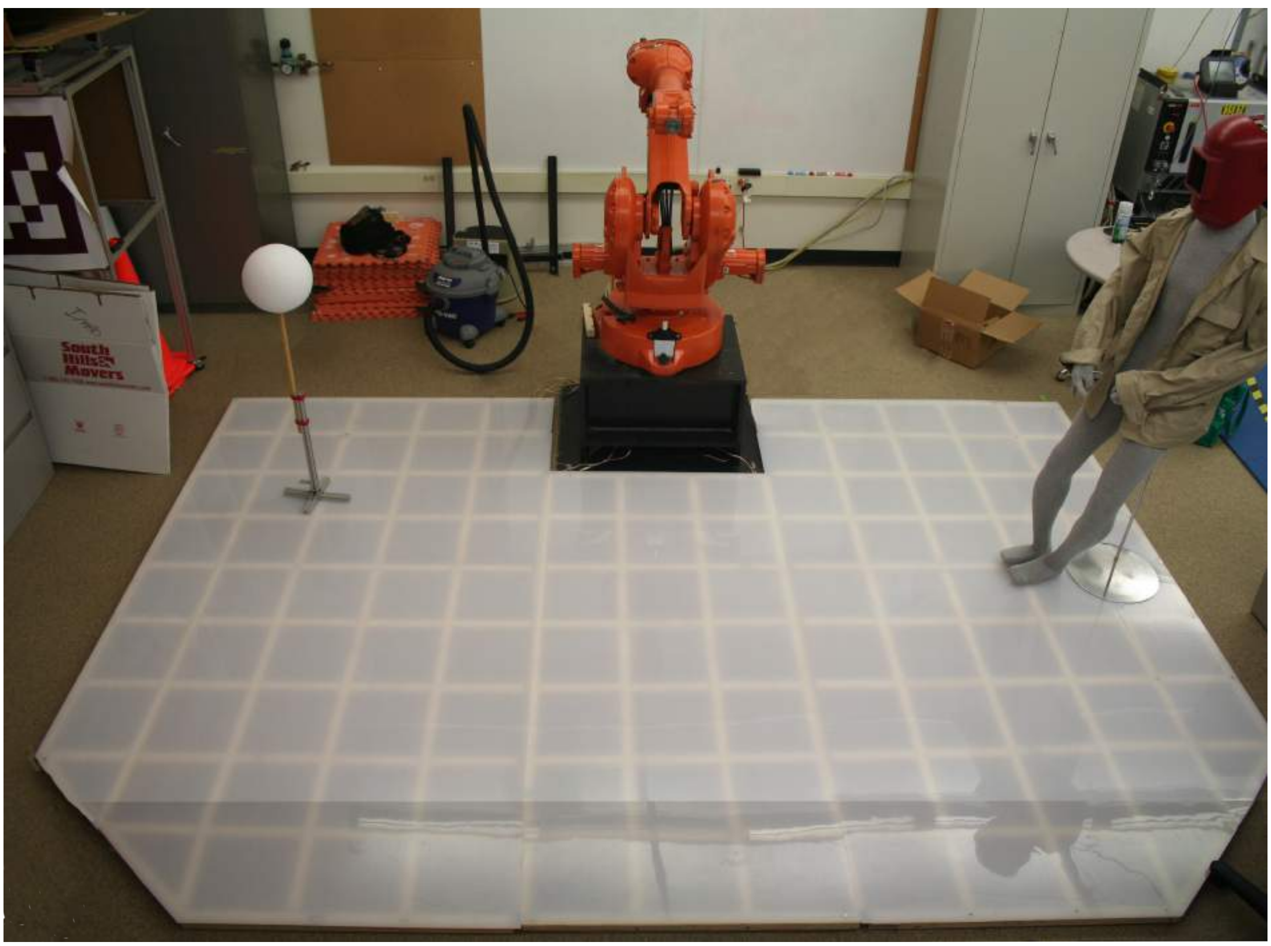

Figure 3: The Intelligent Workcell floorspace

The system uses two different types of 3D sensors to monitor the environment around the robot: Stereo cameras (TYZX Embedded Stereo Cameras, Figure 4a) and Flash lidar sensors

(Swissranger SR4000, Figure 4b). This hybrid approach of sensors permits detection of a variety of objects in the workcell. Stereo cameras detect sharp edges and contrasting surfaces, but tend to fail when detecting surfaces having only small amounts of color variation, contrast, or surfaces with repeating patterns. Flash lidar sensors, on the other hand, perform better at detecting large, matte surfaces with little or no texture or variation. The complementary functions of these two sensor types expand the range of detectable objects and help minimize the chances of sensor confusion or omission of objects from the environment. The system is easily capable of accepting more types of sensors (such as structured light), but these were not required for our purposes. Four sensors were used in this design because they provided adequate coverage of the environment space; the system does not require a specific number of sensors to operate, so long as the desired region is sufficiently covered. Previous studies [11] found through a simulation that the intuitive placement of sensors near the upper corners and directed toward the center of the environment operate well.

(a) TYZX embedded stereo camera

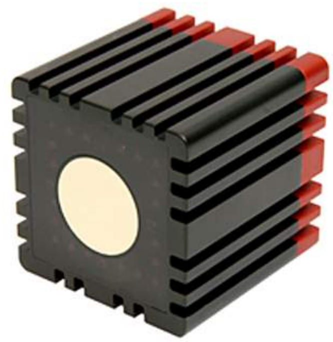

(b) Swissranger flash lidar sensor

Figure 4: Sensors used in the Intelligent Workcell

The software component of the Intelligent Workcell is outlined in Figure 5. The sensor fusion process begins by utilizing the joint angle information from the robot controller to define the orientation of a model of the robot in a virtualized environment. This is generated within the Open Robotics Automation Virtual Environment (OpenRAVE) [54]. The model is converted to voxels (3D pixels), providing a known location in the voxel space of the environment for every segment of the robot. In addition, the voxelized model is expanded through dilation to provide a buffer region surrounding the arm. This expanded grid is termed the danger zone. The system requires at least one to be generated, but can handle multiple (e.g., one that is expanded more can be used to trigger slowing, while another, less expanded region can be used to trigger stops).

The process for object detection within the workcell begins with analysis of the feed from each of the sensors. Each sensor uses a probabilistic model applied to a ray trace of each visible data point, with the result stored in a dense evidence grid. Originally, the grid is initialized to a 0.5 probability, indicating that occupancy is unknown. The probability of a given cell in the grid is reduced if the point is listed on the section of the ray between the sensor and the detected

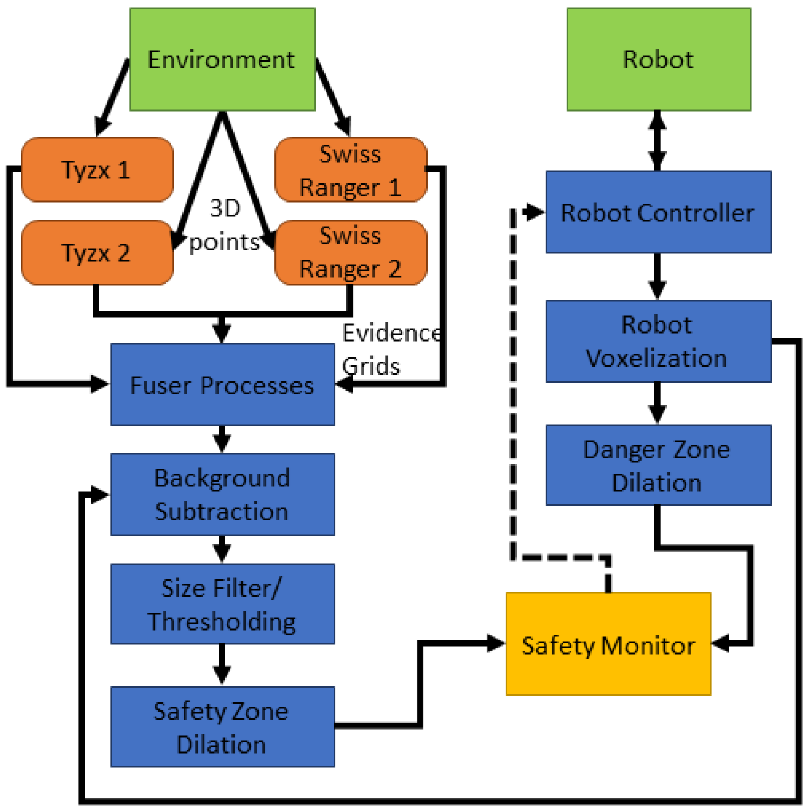

Figure 5: Flowchart of the Intelligent Workcell system

point, is increased as it approaches the detected point, and is left untouched beyond it. A lookup table of probabilities is utilized for accelerated performance. Fusion of individual evidence grids is performed by multiplying the probability across each of the sensors. To accelerate computation, the log-likelihood is stored, which permits summation of probabilities. The evidence grid is then filtered into the background subtraction process. This utilizes a specially generated grid that contains all the high-probability surfaces from when the background image was captured in addition to any points completely contained by the background. In this way, the background image effectively removes all areas of detection where people cannot be present and considers open areas on objects to be unsafe. The generated model of the robot is subtracted from the evidence grid, so that any remaining high-likelihood points should only be people or non-background objects. After background subtraction is complete, occupancy is calculated by thresholding the probabilities of each voxel cell. The system compiles the evidence grids from each sensor, and thresholds based on the fused value to determine occupied voxels. Connected components are used to determine the sizes of occupied regions for filtering, as small regions are considered sensor noise. The thresholding approach inherently forces the system to detect the largest bounding surface for a given set of viable detections. Next, the morphological dilation operator is applied across the remaining grid, expanding the region around any people or objects in the environment and providing a safety margin, resulting in what is termed the safety zone.

Finally, the system assesses the safety of the current operation by overlaying the safety zone and expanded robot zone. If an overlap is detected, the robot is signaled to either slow down or stop. Signals to the robot are updated as soon as a detection is made. A heartbeat signal is continuously sent to the robot controller, confirming that its connection to the host machine is still active. Two quad-core computers were used for the operation of the system, visualization, and data logging. All interconnection between the sensors and the computers were through gigabit Ethernet connections, allowing fast communication from any device to another.

3.2 Feedback Types

Collaboration in dynamically safe, shared manufacturing environments relies on an understanding of actions and goals between the worker and robot. Communication is the foundation of collaboration; however, large manufacturing robots are not designed to communicate intentions and operational states. Thus we needed to develop methods of feedback in order to provide possible interfaces between worker and robot.

When designing our feedback methods for collaboration, we considered different human worker-robot communication mechanisms for communicating robot state, position, and interpretation of the environment. Sound notifications, warnings, or speech from the robot could inform the worker of robot operation state or occurring safety violations. However, we anticipated two problems arising from the use of auditory feedback: First, it does not easily communicate spatial information, such as the location of safety issues or robot position. Second, sound is not a reliable communication method in manufacturing workcells, since manufacturing environments are often loud and space-optimized, leading to missed alerts and confused or garbled signals across adjacent cells. We also considered using wearable devices to serve as a communication medium between robot and worker. However, we were concerned that wearable devices would not be effective in an industrial setting. Wearable visual systems could be too distracting, either by restricting part of a worker's field of vision (if using augmented eyewear) or requiring both observation and a free appendage (if using wrist- or arm- mounted gear). Systems using haptic vibration warnings could be missed, since the worker is moving. Additionally, vibration can only communicate through a single bit (on/off), limiting the information which can be provided about robot state. Finally, and most importantly, wearables can be lost, broken, or neglected, which are additional hurdles we decided to avoid. We determined that visual feedback was best suited for communicating robot intent in a simulated manufacturing workcell. This method of communication inherently provides complicated information more quickly and with less effort than auditory cues or wearable devices. In addition, visual feedback has the ability to communicate spatial information as well as operational state.

We designed two different embedded feedback modalities and one non-embedded feedback modality for use in the Intelligent Workcell: a floor-based embedded display, arm-mounted embedded lights, and an external screen simulation (Figure 6). Each one provides, at minimum, a notification for the state of the robot (e.g. if it was "Waiting," "Blocked," or "Paused"). Each also provides additional information about the environment. The floor display shows a direct mapping of the voxel grids, providing the system's perception of the environment including the worker and robot with its coupled intentions and state. The arm lights directly highlight the interaction of the robot with the worker, focusing less on the overall environment and more on problems the robot is experiencing and its direction of movement. The non-embedded screen simulation communicates both the state of the arm and the overall environment, providing the most information to the worker. However, this information is comparatively isolated from the worker's location and requires the worker to spatially translate their position into the simulation they view on the screen. While other visual feedback modalities could be possible, we considered these three to be a good approximation of available categories of approaches which could reasonably be adapted to the manufacturing industry.

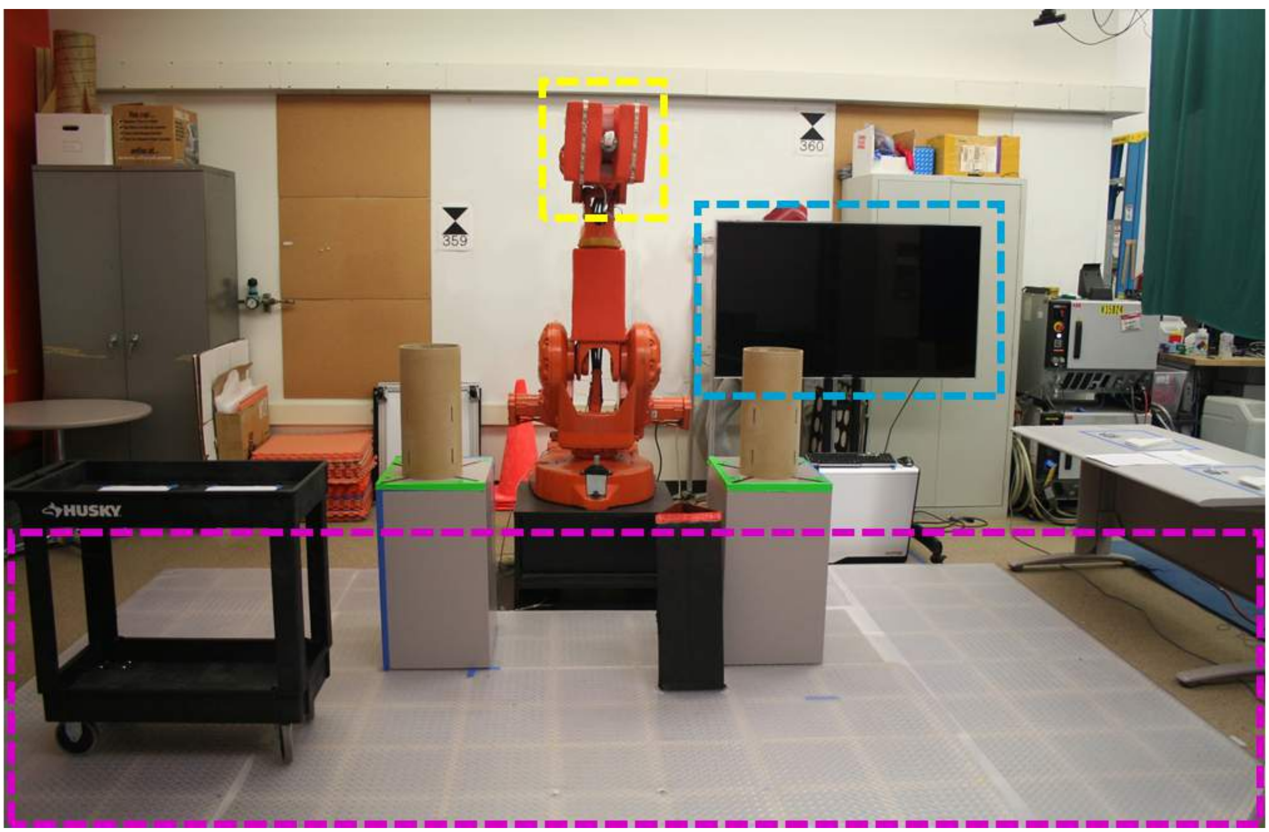

Figure 6: The different feedback modalities in the environment: The magenta box indicates the floor feedback modality, the yellow box indicates the arm-mounted embedded lights, and the cyan box indicates the screen simulation.

3.2.1 Embedded Floor Display

We designed the embedded floor display mode of visual feedback to provide spatially coupled information about both the workcell environment and the robot. This includes the interpretation of the environment from sensor data, expanded information from voxel growth, projected future positioning of the robot, and status of safety system blockage events. Therefore, the primary design requirement of the floor display was the ability to function as a display, requiring independently triggered, spatially separated elements capable of visually conveying multiple sources of information simultaneously. Influential factors included cost, availability of parts and materials, projected effectiveness, ease of operation and maintenance, portability, durability, and reliability in commercial applications.

3.2.1.1 Design Decisions Following our design requirements, we arrived at three possible approaches to a floor display: installing a ceiling-mounted projector system, purchasing a commercially-available dance floor, or developing and constructing a custom floor lighting system. Each approach has its own benefits and disadvantages.

The projector approach provides the benefits of simple and standardized inputs, commercial availability, and minimal space consumption. Since projectors normally utilize a standard connector and communication protocol, controlling the floor would be as simple as rendering the color pattern desired on screen. Projectors are also easily obtainable from many retailers, although ones designed for use in bright environments are generally more expensive. Additionally, since the design called for ceiling mounted projectors, this approach would not take up any space on the actual floor. The projector approach has a number of disadvantages as well: calibration, requiring additional ceiling mounts, beam blocking, and lack of true black. A projector-based system would require calibration between the environment and the projectors, meaning that, in addition to the sensors, the feedback also has the potential to be knocked out of alignment. Projectors also suffer from problems of occlusion, since the image projected to the floor can be blocked by both the person and the robot. This affects the output capabilities, which lead to problems with consistency as the location of robot and person change. Even though it is a minor issue in well-lit areas, projectors lack a true black display output resulting from light bleed-through. This would create a hazy appearance even when the projector is displaying nothing.

Purchasing a commercially available lighted floor was a second possible approach to creating the floor display. These displays are used in night clubs as dance floors and sometimes in public venues, such as museums or shopping malls, for interactive games. This route possesses advantages and disadvantages as well. Commercial availability means convenience of installation and reduced development and production time. The floors have high-resolution, full-color displays. They are durable, wear-resistant and touch-capable for use in interactive applications. The disadvantages of using a commercial floor for our floor display included expense, limited customizable shapes and sizes, unknown performance in bright light environments, and DMX control. Commercial floors are generally sold in 0.5 meter squares designed to be connected into the desired dimensions. The floors we examined were controlled through the DMX bus, which would require some type of add-on controller and dedicated computer interface for proper function. Online demonstrations only showed operation of the floors in dim lighting, which could not provide information about display intensity in bright industrial settings.

The final considered approach was building a custom floor display solution. Building our own floor display provides the benefits of customizing size and shape, cost and the ability to tailor the lighting system integrated control exactly to our needs. We have limited laboratory space in which to work, so control over shape and dimensions would allow us to efficiently utilize the available square footage surrounding the robot. Building our own floor display also gives us the ability to expand and/or modify the system. In the event of component failure, we also would be better equipped to repair the system. The greatest disadvantage to designing and constructing a floor display was the investment of time. This would require research into materials, components and software. It also provided the greatest potential learning experience.

We determined that a custom built floor display was the best option, since it would cost less than half the price of a prefabricated dance floor, would be more durable and dependable than projectors for longer term operation, and could be customized for exact shape and control.

3.2.1.2 Floor Architecture We envisioned a layered system, consisting of a bottom layer to contain components, a top plate, which provided a walking surface and translucent diffusing of the lights, and a support layer between to provide the space in which to mount the lights. We used full color LED strip lighting as the light source because it provides full spectrum, controllable bright light using minimal power and space. LED strip lights are strips of flexible circuit boards with LED's incrementally mounted throughout. The strips include integrated, inline Pulse Width Modulation (PWM) controllers (LPD8806) for holding and regulating colors, meaning that power and serialized RGB values were all that needed to be sent through the four connections. Based on experimentation with small samples of materials, we chose 3/8 inch thick polypropylene to serve as a top layer of the floor display. Polypropylene provides durability, flexibility, translucency, and handling ease when compared to other materials like polyethylene, acrylic, and glass. We built the middle, supporting layer of the floor display out of wood, because of its low cost and ease of use compared to metal or plastic. We determined that a 3.8 cm wide middle layer provided sufficient distance for diffusing LED light while still providing a low profile. Support beams run parallel to the light strips with perpendicular cross beams, which are slotted to allow threading of the strips. Strips were placed at four inch increments across the floor, with support structures forming 30 cm square boxes (Figure 7b). Strip lights were run along the short dimension of the floor, dividing the floor display into two 1.5 by 2.7 meter sections and one 0.9 by 1.8 meter section. This modular design enables transport or reconfiguration should the need arise.

The available floorspace and the workspace of our pedestal-mounted robot dictated the shape and size of the floor display. The maximum distance the robot arm can extend from the center of its base is 1.5 meters. We designed the floor display so that a person in the workcell would have sufficient space in which to operate plus a safe extended distance beyond the maximum reach of the robot arm (Figure 7a). This additional margin was set at 0.6 meters, requiring the floor to be at least 4.0 by 2.1 meters. To provide the potential for additional interaction points and to increase operational area, we set the floor dimensions at 4.0 by 2.7 meters, with the furthest two corners from the manipulator shortened diagonally by two feet. We do not consider the space behind the manipulator an operational area in the workcell, since it extends outside sensor coverage.

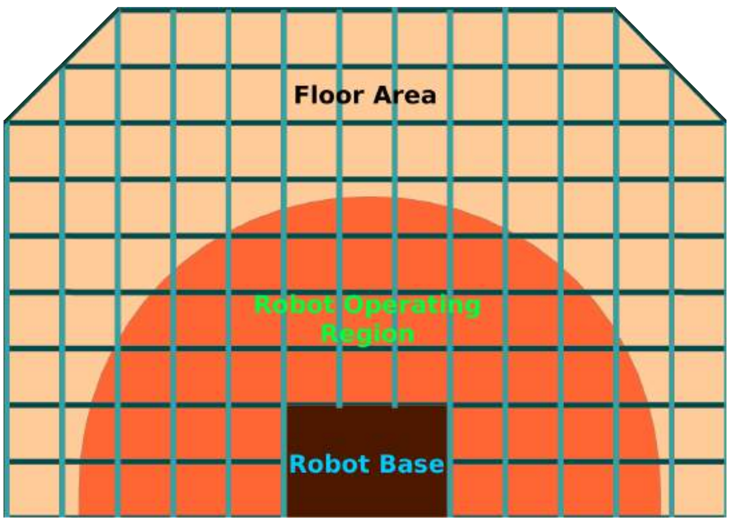

(a) Layout of the floor

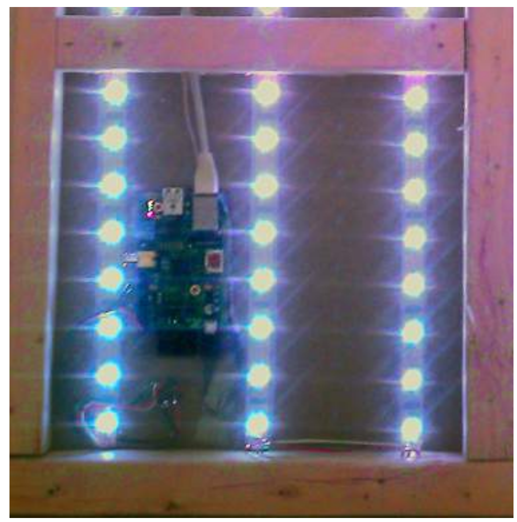

(b) Single exposed floor cell

Figure 7: Design and construction of the floor display

3.2.1.3 Control Electronics The LED strip lights are capable of handling 21 bits of color, transmitted through a Serial Peripheral Interface (SPI) bus. We chose to use three Raspberry Pi embedded Linux computers to drive the three sections of the floor. The Raspberry Pis are able to dump a section of RAM directly to the SPI bus, giving a floor refresh rate in excess of 20 Hz. Since each section of the floor display has its own Raspberry Pi, the floor is completely modular. The floor display also runs using the lab's existing Gigabit Ethernet network.

Each Raspberry Pi is able to directly control flashing, so that a slower refresh rate from the primary computers does not cause visual stuttering. This flashing is achieved locally on each Raspberry Pi by interpolating between two color buffers for each color element, with the proportion controlled by a global sinusoidal function relying on a frequency multiplier (set to 1.5 Hz) and wall time, which was synchronized between Rasberry Pis. If both buffers are green, the floor element displays a constant green. If one buffer is green and the other red, the floor element smoothly alternates between green and red. If one buffer is green and the other black, the floor element alternates between green and off, generating a smooth flash. Two 100 amp, 5 volt switching power supplies provide power to the floor display and the Rasberry Pis.

3.2.1.4 Algorithm and Color The floor display communicates multiple types of information to the worker, including robot run state, current and future position of the manipulator, locations of movable objects in the environment (including workers), and collision locations with identification of the offending object. The floor display algorithm takes in three different sources of information: the danger voxel grid, the future robot voxel grid, and the safety voxel grid. The danger voxel grid indicates the current robot position in the environment, through a grown, voxelized model of the robot with the current joint-angles. The future robot voxel grid indicates the future robot movement, generated by voxelizing and growing the robot model at a few positions interpolated from the current position and upcoming positions in the movement script. The safety voxel grid indicates the environmental location of movable objects, including workers. To collapse the voxel grids into a 2D interpretation for use by the floor display, the system iterates through each x/y location and (for each pair) examines the values of the grids throughout the z-axis. This examination utilizes a state machine to determine the label for the column, listed in increasing order of precedence:

lowest
- Empty
- Future robot position
- Object or person
- Robot danger zone
- Both safety and danger zone
- Collision between robot and object/person

highest

This provides us with a single, collapsed, 2D interpretation of all of the voxel grids, which we call the label map. The first thing to note is that collisions between robot and worker do not mean that the robot is physically in contact with the object/person, but rather than their expanded detection regions are intersecting with at least one voxel. The second is that having both a populated safety and danger zone voxel in the column is not a sufficient condition for a collision, since they can be present at different heights without causing the robot to stop. However, it is important to treat this overlap-at-different heights condition with special consideration so that the worker is not confused by conflicting results being displayed. Coupled with the running state of the robot, this 2D interpretation generates the final output for the display. The floor LEDs are mapped to pixels in an image that lets us render a 27x39 pixel image across the floor. However, the floor display does not have a resolution which is an integer multiple of the voxels in the environment. To accommodate for this difference, each coordinate in the label map is mapped to the pixel it should be centered over. Since multiple labels can map to the same pixel, the highest precedence label is chosen.

The robot run states are defined as “Running,” “Blocked,” and “Waiting.” “Running” means the robot is currently unhindered and is moving to a position in the movement script. By definition, being in this state indicates that there are no collision labels in the labeled map. “Blocked” means the robot is currently hindered by a movable object and thus stopped while trying to move to a position in the movement script. Contrary to the “Running” state, the “Blocked” state indicates there is at least one collision label in the label map. “Waiting” means the robot does not currently have a destination position, and is waiting for an event to trigger another movement sequence. All labels are allowed while the robot is in the “Waiting” state; however, collisions are changed to object/person labels, since the robot cannot collide while waiting. We selected display colors and patterns based upon informal evaluation by a small group of graduate students.

While the robot is in “Running” state (Figure 8a), the floor displays solid red for the danger zone of the robot, solid orange for the future robot position, and solid blue for the movable objects and people. All the robot lights are warm colors to indicate that these regions should be avoided. The assigned color are: red for the danger zone, orange for the future robot position, blue for movable objects in order to provide a sharp visual contrast to the warning colors, and purple for the “Both” label, as it is a combination of red and blue. We did not use the color green for objects, since one of its composite colors would be yellow, which is used elsewhere.

While the robot is in the “Blocked” state (Figure 8b), the floor flashes as a global warning that something has occurred within the environment. The display flashes red for the robot, orange for the future position, blue for people and white/red for the collision region. Additionally, a flashing white/red bounding box is wrapped around the offending object and the collision zone. This box is designed for fast and easy location in the event the blockage is obscured from view. A flashing global state of the system communicates blockage to a worker without the need for him/her to examine a particular part of the floor.

While the robot is in the “Waiting” state (Figure 8c), the floor displays flashing yellow for the robot, and solid blue for objects. Yellow light indicates caution, to show that the robot is still present but not dangerous. Blue still indicates objects and workers for consistency between robot states. Collisions and overlaps during a “Waiting” state are changed to blue as well, because the robot’s space does not need to be avoided by workers at such times.

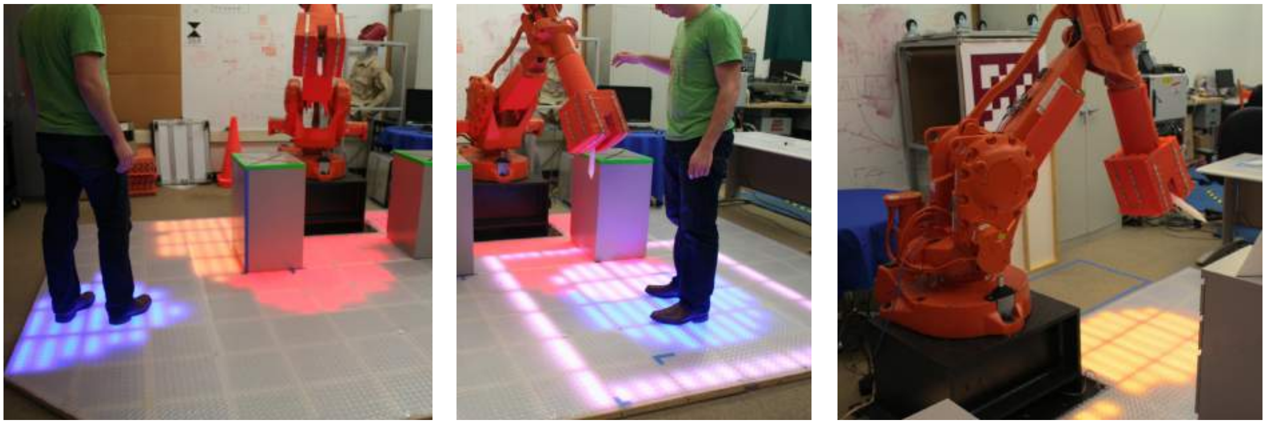

(a) (b) (c)

Figure 8: Floor display (a) running, (b) blocked, and (c) waiting

3.2.2 Arm-Mounted Lights

3.2.2.1 Overview We designed a second embedded feedback system composed of arm-mounted lighting that provides contextual information about the arm and its status. The information includes the robot's direction of motion, direction of blockages, and operational state of the manipulator.

3.2.2.2 Design We mounted full color LED light strips on the outside of the robot wrist joint. This is the furthest-reaching and most visible portion of the arm. The wrist is surrounded by a foam box, for additional safety during experimentation. The box is large enough to allow mounting of two strips of LED lights on each of the five available sides (the back of the box is occupied by the forearm of the manipulator).

We mounted a Raspberry Pi microcontroller on the first segment of the arm to drive the light display. The LED strips provide a direct, highly visible means of relaying information about the robot to the worker. We treat each side of the box as a single display element, meaning that all LEDs on a side always show the same color.

3.2.2.3 Control The robot arm lights communicate several types of information to the worker including robot run state, direction of movement and blockage direction. To handle these communication channels, the arm lights require input of the robot run state, robot position, and the safety voxel grid. To generate signals that depend on the direction of motion, we calculate a delta vector for the position of the wrist joint from the OpenRAVE simulation environment. This vector is then rotated into the frame of the wrist joint, again utilizing OpenRAVE. We then use the normals of each face (pointing outward) and the rotated delta vector to calculate the dot product for each face. Faces with a positive dot product are pointing in the direction of robot motion, while faces with negative dot products are pointing opposite to the direction of robot motion. Direction of collision is calculated by a virtual box that is wrapped around and constrained to the wrist joint of the manipulator model in OpenRAVE. Each side of the box checks for collision with the safety grid: if a side has collided with an occupied safety grid voxel, that side of the box is marked as a direction of blockage. We define robot run state as "Running," "Blocked," or "Waiting." Colors are primarily used for indicating direction, whereas flashing is used primarily to indicate state.

While in the "Running" state (Figure 9a), the robot is moving to locations defined in the movement script. It displays solid green in the direction of travel, and solid white on the other sides. Green is used as a "headlight" to indicate active motion toward the worker. Solid white is used for the running state.

While in the "Blocked" state (Figure 9b), the robot is attempting to move but is hindered by an object or worker. It displays flashing red in the direction of the blockage, flashing green in the direction of travel, and flashing white elsewhere. Conflicts are resolved by giving red priority. The other sides of the box flash white during blockages. Blockages caused by close proximity to the robot base alone will not cause any red illumination.

While in the "Waiting" state (Figure 9c), the robot is not attempting to move and is waiting on the next event to trigger movement. In this state, every side flashes yellow. This form of visual feedback requires the workers to observe the robot's end effector for feedback information.

(a) (b) (c)

Figure 9: Arm-mounted lights (a) running, (b) blocked, and (c) waiting

3.2.3 Screen Simulation

3.2.3.1 Overview The third mode of visual feedback we developed is a non-embedded screen simulation. This method communicates the status of the environment on a screen; the information includes the running state, the robot position and future locations, and the locations of blockages and objects in the environment.

3.2.3.2 Software Design Since we already used the OpenRAVE environment to display voxel grids for testing sensors and voxel data, it was logical to extend the same environment to the screen display. The OpenRAVE simulation displays the robot and the environment, as well as the unexpanded safety voxel grid, the intersection of the safety voxel grid, and the danger voxel grid. The grown safety voxel grid indicates where people and non-background objects are located; utilizing the grid before expansion removes excessive occlusion of the environment, and makes the grids more realistic in size. The collisions between the safety grid and danger grid serve as safety violation regions, showing workers where collisions are occurring. Some additional programming was required to allow the robot to display future positions in the movement script, progression into the movement script, and robot run state.

3.2.3.3 Position Interpolation To show the path the physical robot would be taking, an interface between the movement script (which only set key points for actions) and the simulation was required. OpenRAVE was already capable of placing the virtual manipulator at specified joint-angle positions, but we did not have the capability to show a set number of seconds of the robot's future path. For this, linear interpolation is used between the different joint angle positions in the movement script. The velocity between key points is defined by the maximum angular movement speed per joint; whichever joint needs the most time to obtain its goal position at maximum angular velocity defines the movement time for that segment of motion. All other joints then calculate their velocity based on the movement time and their delta angle.

For every motion, the current robot position (reported from the controller) is used as the initial position. The current line number in the script is reported to the simulation by the robot-controller interface (responsible for sending out the next position to the manipulator controller). It then becomes possible to calculate the next 10 seconds of movement and render it at a 4-fold speed increase, allowing an up-to-date display of the future motion of the robot. Since the position is based on time, the simulation knows how far through the future path the display has progressed; this information is communicated to the worker through a progress bar.

3.2.3.4 Progress/Status Bar The progress bar indicates the approximate amount of future action remaining to be displayed and provides a temporal reckoning of robot arrival at a particular position. Placed across the entire width of the screen and about a sixth of the height, the progress bar also indicates the current state of the system using both words ("Running," "Blocked," or "Paused,") and colors (Green, Red, or Yellow, respectively). At the end of the future path display, the entire simulation environment turns black for approximately 0.25 seconds while the status bar and robot path are reset to the initial positions. This blink provides a visual indicator that the path display is refreshing.

3.2.3.5 Operation

For all of the different robot run states, the simulation displays the unexpanded safety voxel grid. Safety voxel grids are displayed as green, to indicate that they are "friendly" or "safe" objects. However, the progress bar does change between states. While the robot is in the "Running" state (Figure 10a), the status bar is set to green. Additionally, black text inside the progress bar reads "Running." In this state, the future action of the robot is displayed.

During a "Blockage" state (Figure 10b), the status bar turns red and displays the text "Blocked," showing that the system is stopped by something the worker needs to correct. The offending safety/danger voxel grid collision region is illuminated in red voxels on the display. The future action of the robot is still displayed, so that the worker can see the robot's intended path.

When the robot is in the "Waiting" state (Figure 10c), the status bar turns yellow and displays "Paused." This indicates that the robot is not trying to move, but that the person is still responsible for completing a task. Any voxel grids (green for people or red for collisions) remain on the display screen, so that participants still view the simulation as a live feed of the system.

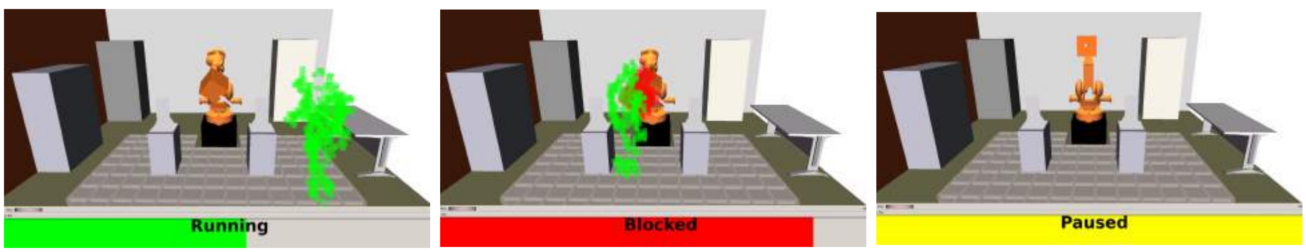

(a) (b) (c)

Figure 10: Simulation output (a) running, (b) blocked, and (c) waiting

3.3 Most Useful Feedback Situations

In general, any form of human-robot communication within a manufacturing environment can be useful; however, some instances of interaction can provide greater impact on overall performance through collaboration than others. To explore the utility of feedback, we propose a few different situations in which feedback from the robot could allow the worker to increase performance or ease of operation. We present three situations that potentially have the greatest impact.

3.3.1 Target Situation A: Close Worker Proximity but Independent Tasks; A "Corridor Situation"

One situation in which feedback might prove useful is when a robot and human worker are completing independent tasks but have partially or fully overlapping task regions. The overlapping region could lead to task interference between worker and robot, and leads to what we term a corridor situation. This occurs when the human worker needs to operate in an area in close proximity to the robot's working location, leaving a small region or "corridor" between the two (Figure 11). The challenge for the worker is to determine if this corridor is large enough to work in without hindering the robot, and vice versa.

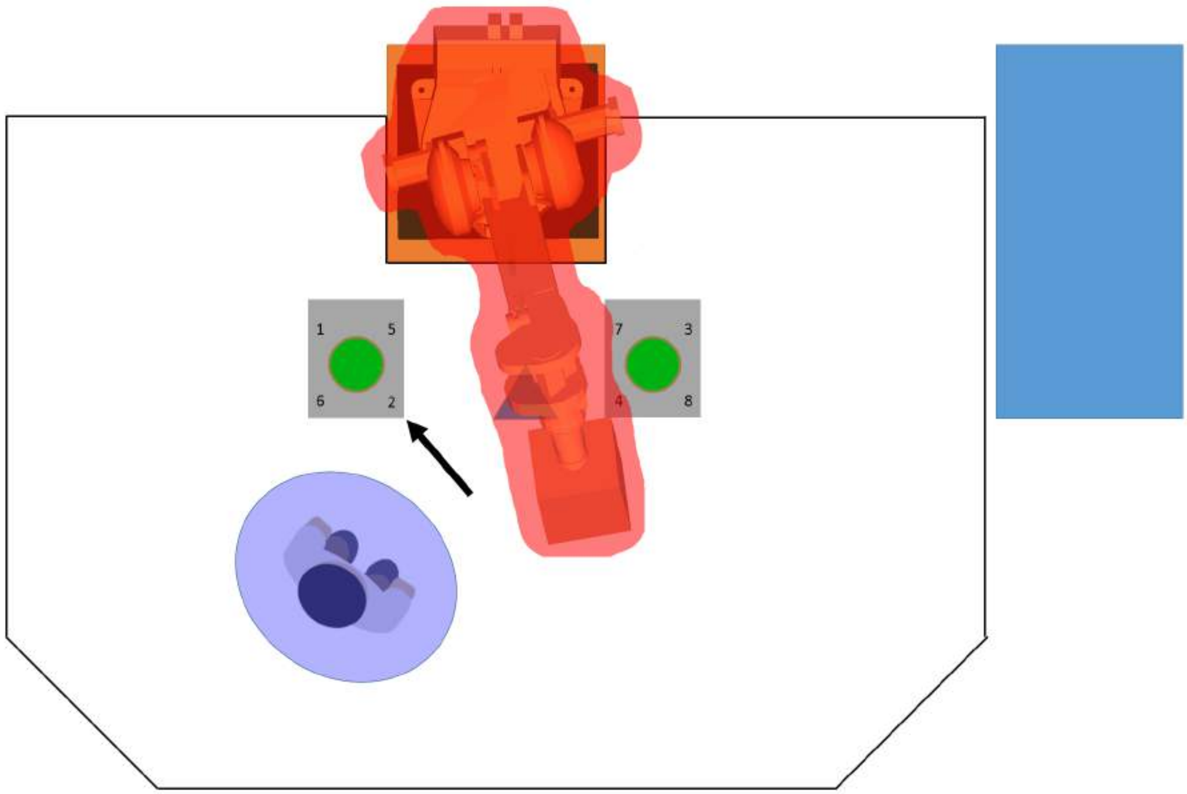

Figure 11: In this figure, a corridor situation can arise when the worker is attempting to assemble the part on the left while the robot works on the left side of the right part.

In the Intelligent Workcell, clearance around the robot is not intrinsically communicated, since the safety boundaries are virtual and larger than the robot. Without a physical representation of the robot's range, it is difficult for a worker to determine when exactly the robot will and will not be blocked by their actions. Explicitly communicating boundaries or related information in this setting could improve collaboration and performance, and potentially even preemptively avoid blockage. Effective use of the shared workspace could mean reduced idle times for both robot and worker. Feedback information in this situation might give workers an idea of when and/or where a blockage of the robot will occur. In manufacturing, this situation can occur if a worker needs to re-stock parts, move parts from one location to another, or complete a task in close proximity to the moving robot.

3.3.2 Target Situation B: Dynamic Scheduling of Tasks for the Robot

A second possible useful feedback situation is when the robot varies from its default trajectory, potentially conflicting with the collaborating worker's expectation. As manufacturing processes become more intelligent and dynamic, it will be possible for manufacturing robots to schedule their own actions based on sensor detections in the environment, as opposed to following a static regimen known verbatim by the worker. For proper collaboration, changes in action need to be clearly communicated to the human worker, so they do not unwittingly interfere with robot operations and cause delays. Effective feedback could allow such dynamic behaviour without negatively impacting the overall performance of the system.

A current application of this situation is in manufacturing lines that produce dynamic configurations requiring slightly different assembly procedures. Different instruction sets could sometimes confuse workers, but the addition of feedback could allow for better collaboration and thus more agile manufacturing.

3.3.3 Target situation C: Objects Obstructing Operation

The third situation can occur when an out-of-place object is present in the workcell. Unexpected tools, parts, or pieces of equipment left behind could cause a safety system to trigger.

There are two primary areas of interest based on feedback in this situation. First, the action of the worker: workers instinctively understand that moving radially away from a fixed robot removes them from its reach; therefore, moving away is always a viable safe path. However, this does not necessarily hold true for objects; object-robot proximity and orientation are not necessarily apparent to the worker, so additional information might help resolve the blockage. Second, the speed of recovery; we theorize that object blockages are the least expected and potentially the least noticeable blockage type. Thus, it is important for a worker to be able to quickly notice and resolve outstanding violations. Feedback in this situation could allow fast detection and resolution of blockages caused by out of place objects that could otherwise cause unexpected drops in productivity.

3.4 Experimental Evaluation of the Feedback Modalities: User Study

With our visual feedback developed, we devised a method for testing their performance in a collaborative human-robot environment. Utilizing the Intelligent Workcell manufacturing environment, we created a simulated manufacturing procedure for feedback evaluation. The procedure requires assembly of two "rockets" utilizing both a human participant and a robot for different tasks. The assembly procedure requires four fins to be placed and secured onto each of the two rocket cylinders by the worker. After the participant secures each fin, the robot performs a simulated "weld" to permanently attach the fin.

We designed the task to require close proximity work between the robot and participant, presenting a number of situations in which feedback could potentially improve performance. Our experimental Intelligent Workcell environment consists of the robotic manipulator, two work stations, a loading station, a cleaning station, and a parts cart (Figure 13).

Assembly of the two rockets at the work stations occurs in parallel, requiring the participant to alternate between cylinders after every two fin attachments, following the order shown in Figure 12.

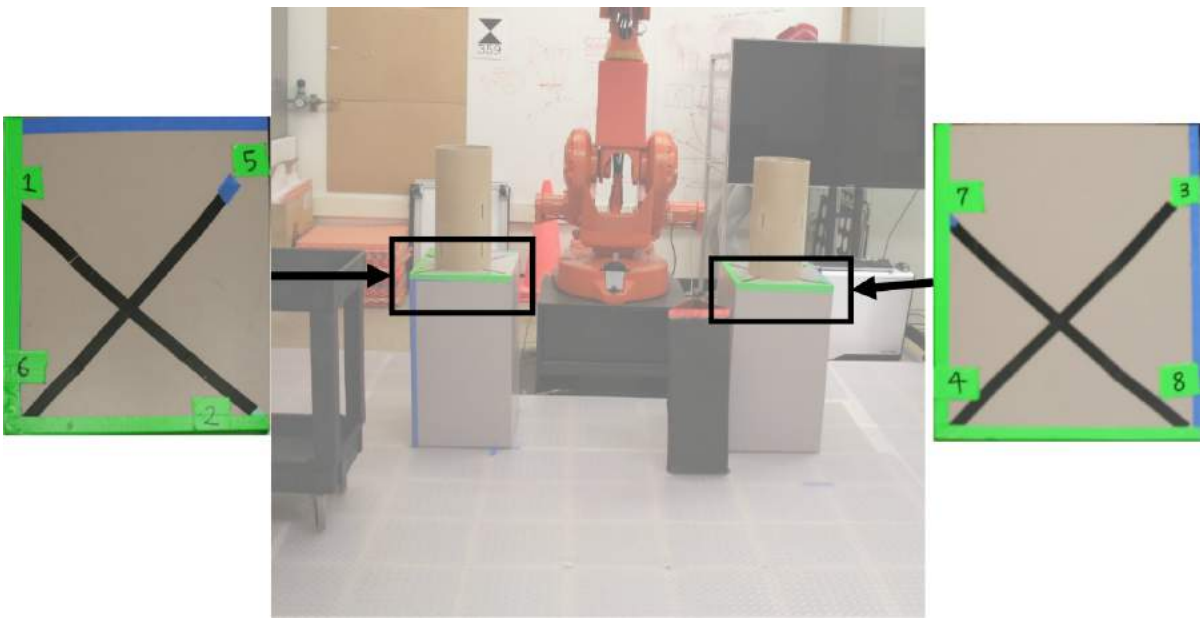

Figure 12: Order of assembly for the fins on each of the cylinders

We instructed participants to treat the fins as though they were heavy parts. OSHA standards require parts over a specific weight to be transported using equipment such as carts, dollies or hoists [55]. Having "heavy" parts allowed us to require a cart (Figure 14) for transportation, creating situations similar to Target Situation C (Section 3.3.3). We instructed participants to transfer fins one or two at a time on the cart through the environment.

Fin attachment consists of sliding the fin into a slot on the cylinder, pushing the fin down to lock it in, and adding a bolt and nut to the fin (Figure 15 & 16).

After the participant attaches each fin to a cylinder, the robot moves in and performs a simulated weld on the part (Figure 17). The weld action consists of the robot moving its foam tip along the fin's seam at a decreased movement rate. Assembly completion of each fin by the participant is manually detected by the experimental operator, who informs the system when each fin is ready for welding.

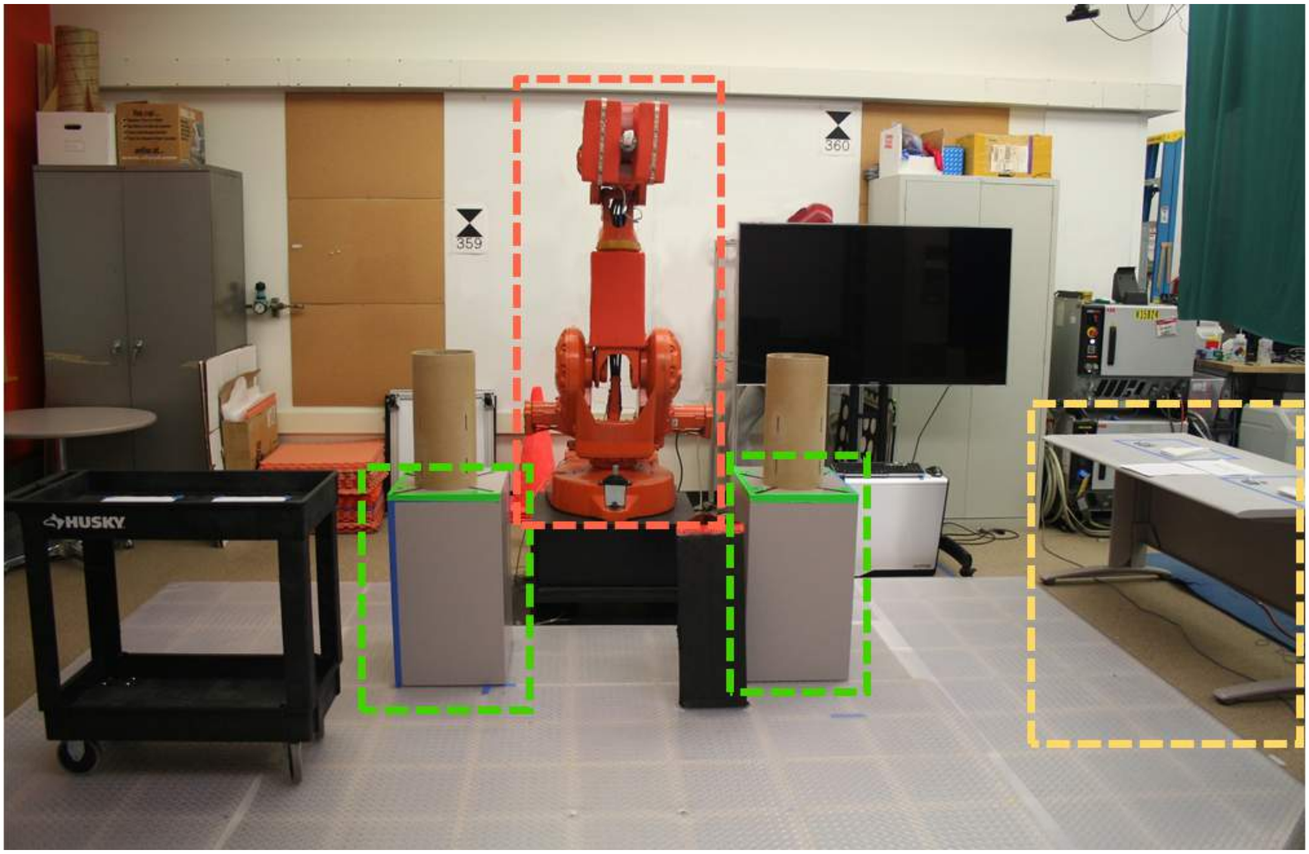

(a) Manipulator, welding stations, parts table

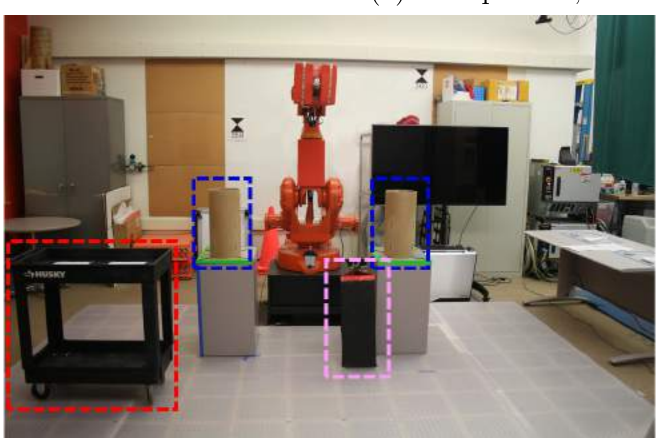

(b) Cleaning station, cylinders, cart

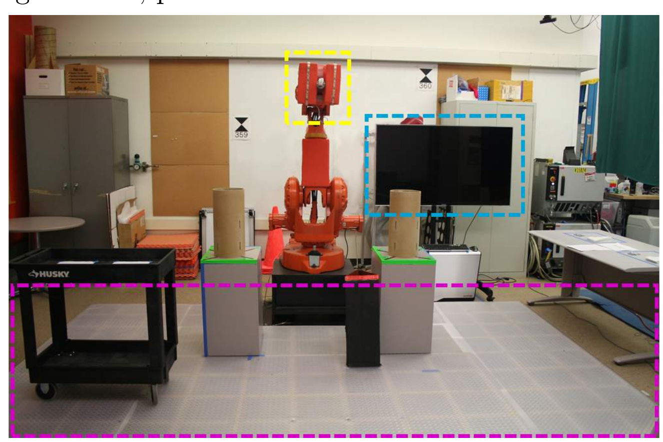

(c) Feedback locations

Figure 13: Intelligent Workcell environment

Figure 14: Cart used for transporting parts in the environment during the experiment

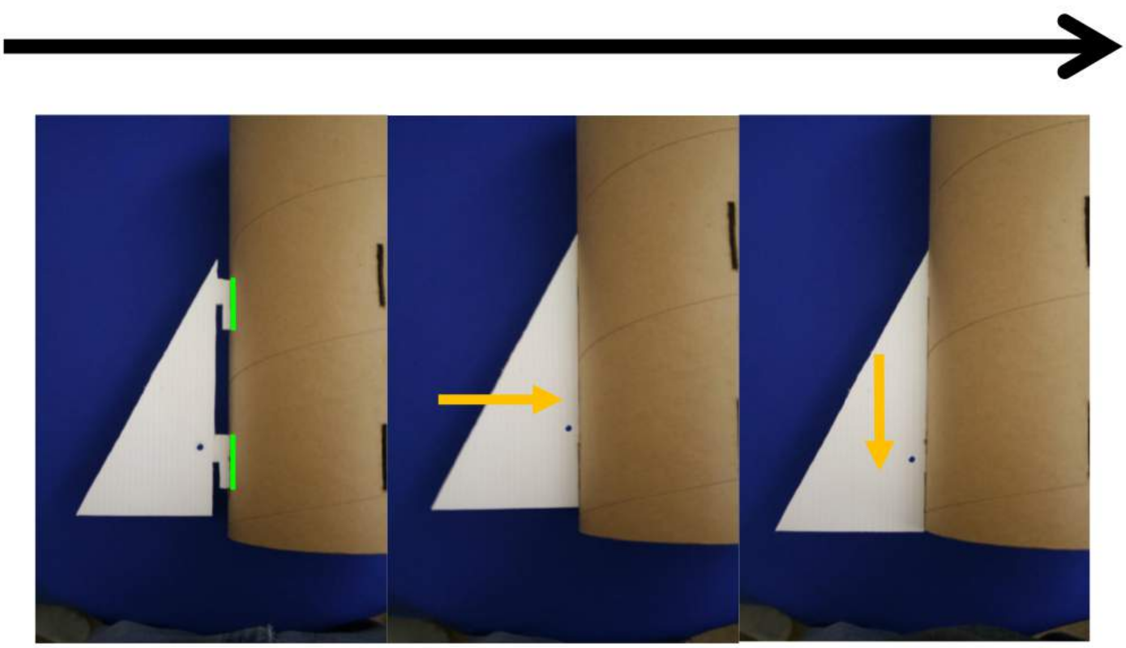

Figure 15: Steps for installing a fin: insert the fin, slide fin down to lock

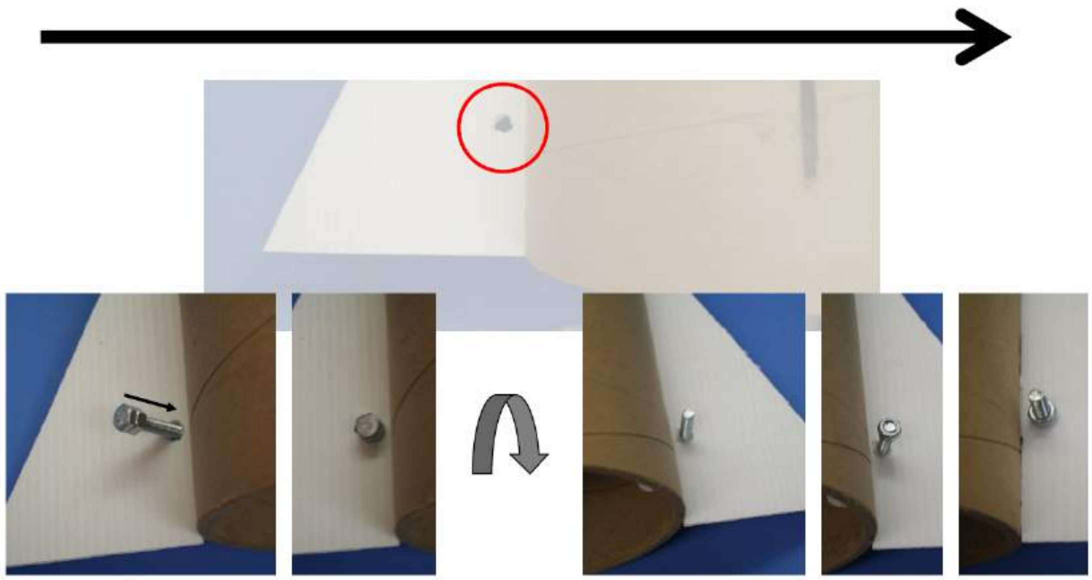

Figure 16: Steps for installing a bolt on the fin: insert bolt through hole, add nut, and tighten

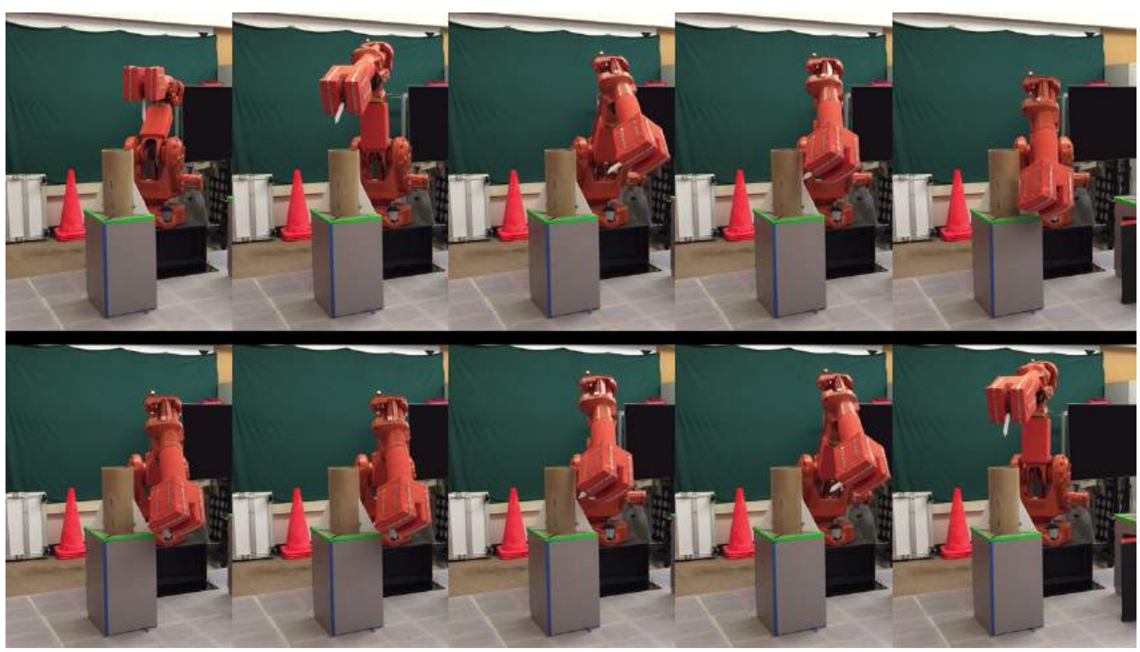

Figure 17: Robot completing its simulated weld: runs the foam welding "tip" along the fin seam at reduced speed

At "unplanned" points in the assembly procedure, the robot moves to the cleaning station and simulates cleaning its welding tip. We informed participants of this cleaning task, but not of its scheduling. For consistency, times of the cleaning task were uniform for all conditions. We used the cleaning to simulate an unexpected task, as in Target situation B (Section 3.3.2). The central location of the station provided Target situation A (Section 3.3.1) interactions as well, since the participant and robot could potentially work in close proximity during portions of the task.

We required participants to complete the procedure once with one of the three feedback modalities or with a control case (without any feedback). Each participant was trained to perform the assembly procedure and to utilize their assigned feedback modality (if required). Examples of the assembly and feedback manuals can be found in the Appendix (Sections 6.1, 6.2, 6.3, and 6.4)

3.4.1 Safety of the Participants

Participant safety was a primary concern in running this experiment. We aimed to minimize risk as a whole. The Intelligent Workcell acted as our primary safety system, and was considered sufficient protection from undesired events. To help guarantee that our study was classified as "minimal risk," we included three additional safety measures in the experiment: manipulator padding, an attentive safety operator, and limited manipulator movement speed.

We added 2 inch foam padding around the end effector and forearm of the robot and along the front of the first arm segment. These regions were theorized to be the fastest moving and most exposed to the participants. Should contact occur with the arm, even if the robot is motionless, the added padding makes sure that contact would be soft.

A safety operator was present while the participants were within the operational reach of the arm, watching and ready to trigger an immediate emergency stop of the robot. This provided a fail-safe in the event an unexpected risk presented itself during the experiment.

Finally, the arm's movement speed was limited. Although we wanted to complete the experiment with the robot at normal operational speed, this reduction insured safety in two ways: first, any potential impact of robot would occur with reduced momentum, and second, it provided the safety operator with additional reaction time in which to trigger an emergency stop override.

3.5 Training

For all of the modalities, we used the same basic training procedure. The training familiarized the participants with the system before they entered the environment, and provided a brief experience with the robot before they attempted the assembly procedure. Before each participant performed the procedure, we wanted to give them:

- Familiarity, comfort, and a sense of safety while working with the manipulator.

- Time and practice handling the parts used in the assembly procedure.
- First-hand experience in the environment with the robot in operational mode.
- Practice handling the cart.
- Verification of their understanding of the assembly procedure.

3.5.1 Overall Procedure

Our process for running a participant through the experiment was:

- Describe briefly the Intelligent Workcell, its purpose, and how it functions.
- Have the participant read the Assembly Instruction manual.
- Clearly tell the participant:
  - Do not turn the cylinders during the procedure.
  - Try not to block the robot during the procedure. Continuing to work while the robot is in motion is fine, but try not to block the robot.
- If Feedback is enabled:
  - Have participant read corresponding feedback modality manual.
  - Have participant watch corresponding feedback modality video.
- Allow the participant to perform the sweeping tutorial in the environment.
- Allow the participant to perform the first fin assembly task.
- Confirm that the fin assembly task was performed properly.
- Robot demonstrates the first weld task.
- Answer any remaining questions.
- Run full assembly procedure with participant.
- Administer exit survey and compensate participant.

3.5.1.1 Assembly Procedure Manual For the training process, we provided the participant with a manual detailing the assembly procedure (Appendix 6.1). This included details on how to assemble the parts, general guidelines, locations/names of parts, and the ordering of the fin assembly. The corresponding fin assembly numbers are also labeled on each of the assembly stations. The assembly manual advises participants not to block the robot but that it is acceptable to continue working while the robot is in motion. We intended these guidelines to encourage cooperation between participant and robot and to discourage participants from avoiding the robot throughout the assembly procedure.

3.5.1.2 Feedback training After they reviewed the assembly procedure manual, participants operating with a feedback modality were given a manual specifically tailored to the feedback, followed by a demonstration video (See Appendix 6.2, 6.3, and 6.4 for the arm, floor, and simulation manuals, respectively). We used a recorded video rather than a live demonstration of the feedback to keep the instruction as uniform as possible across conditions. Each of the different videos included a demonstration of the robot in its "waiting," "blocked," and "running" states. (Video screenshots: floor display, Figure 18; arm-mounted lights, Figure 19; screen simulation, Figures 20, 21, and 22)

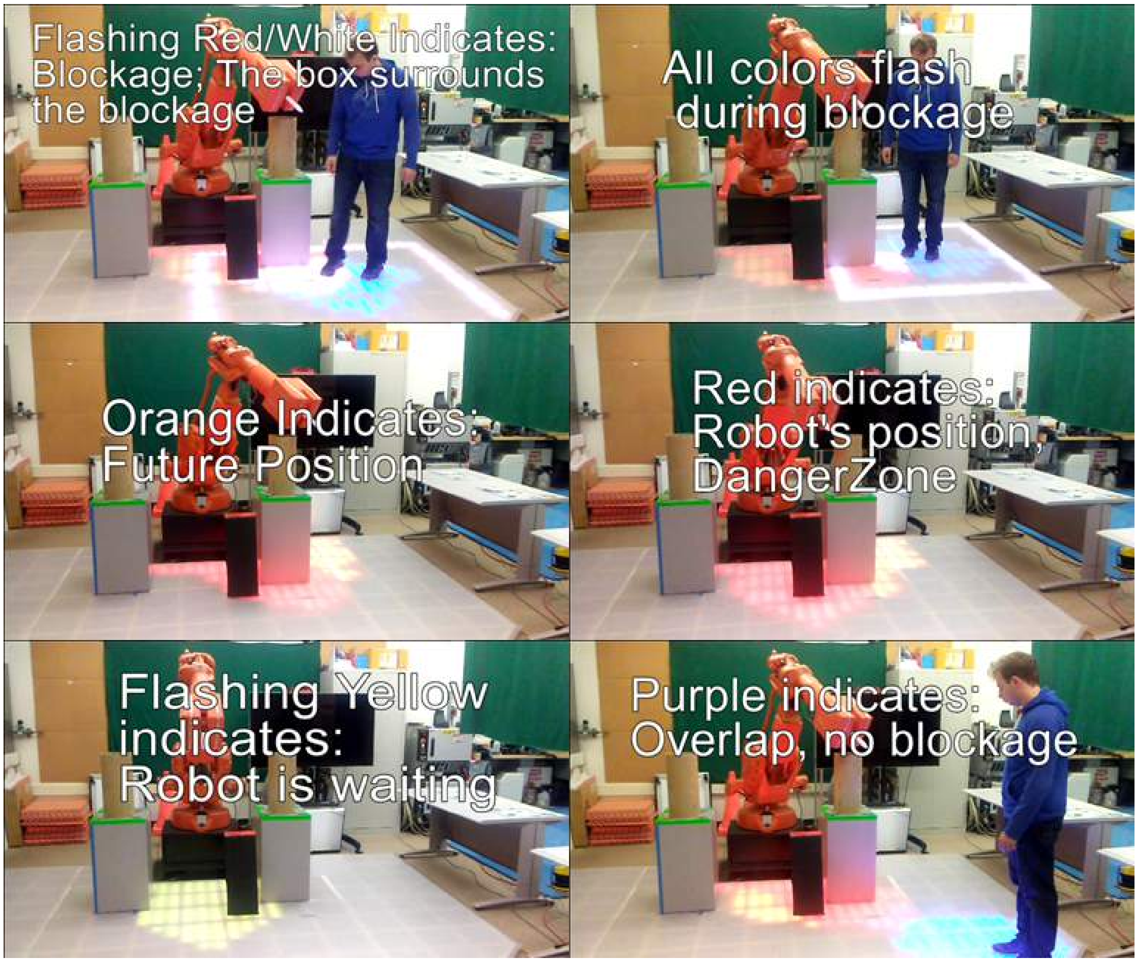

Figure 18: Floor display video tutorial. Describes the different states and coupled information: blocked state along with the blockage localizing box and overlaps, running state with the robot danger zone and future position, and the passive waiting state.

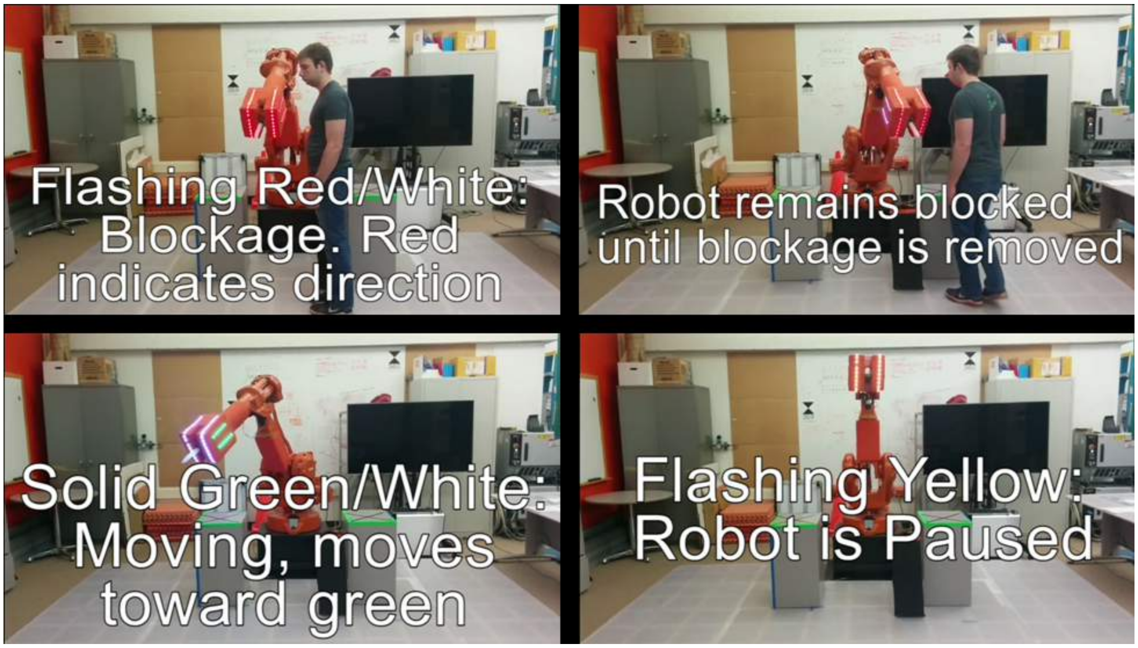

Figure 19: Arm-mounted lights video tutorial. Describes the different states: blocked state and associated direction of blockage, running state with movement direction indication, and the passive waiting state.

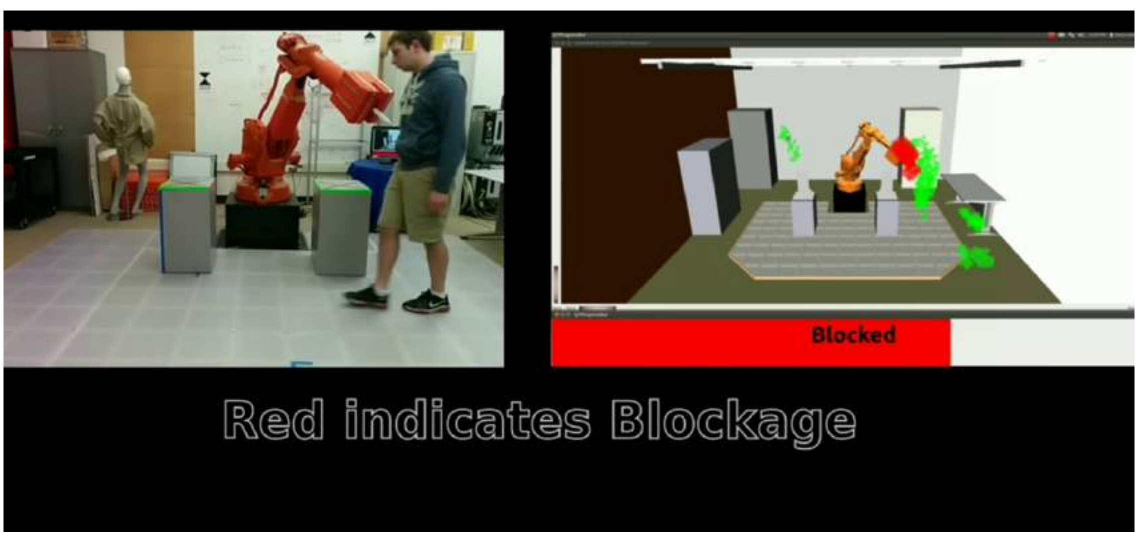

Figure 20: Simulation video tutorial part 1: blocked state. Shows blockage stopping the robot in red, the position of the worker in green, and the progress into the future action along the bottom.

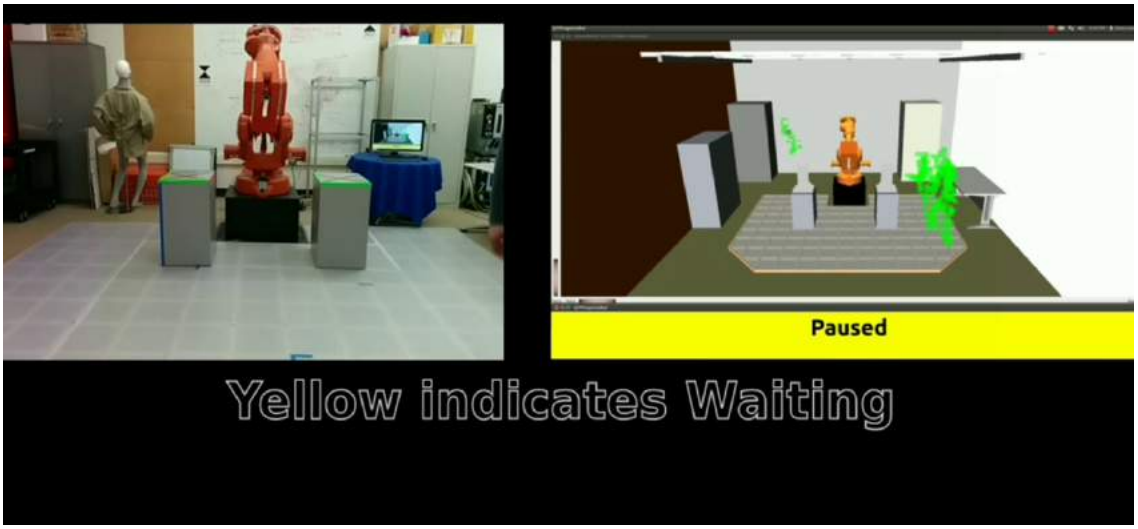

Figure 21: Simulation video tutorial part 2: waiting state. Shows the passive waiting state of the robot.

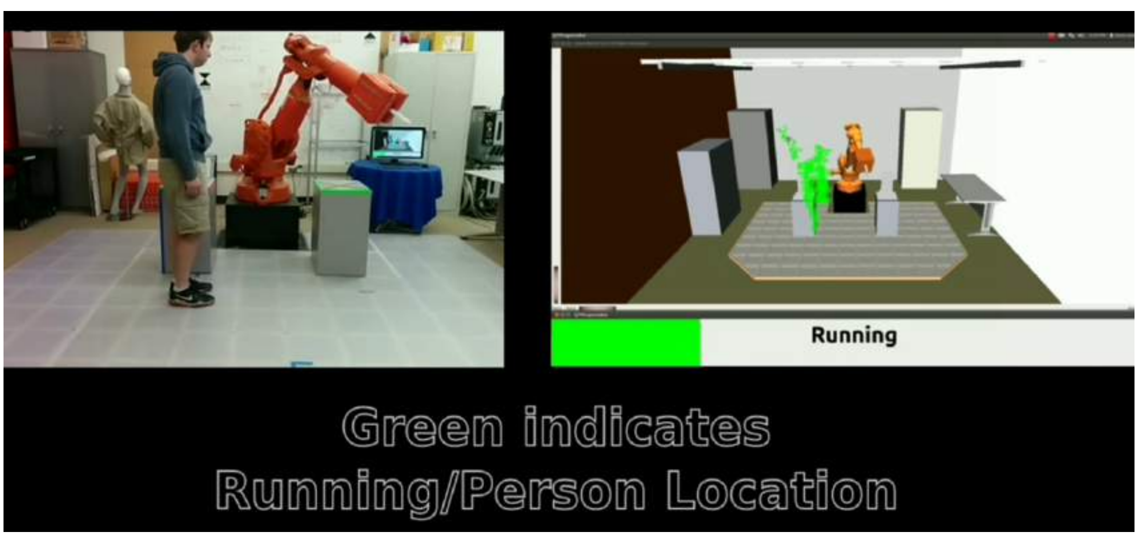

Figure 22: Simulation video tutorial part 3: running state. Shows the running state of the robot. Note that the person is still labeled green, the color and completion of the status bar reflects the running state and the progress into the future position of the robot, respectively.

3.5.1.3 Tutorial Task Once all the media was reviewed, we addressed any questions that the participant had at the time. Then we allowed them into the environment to experience working with the robot and safety system using a slow-moving tutorial (Figure 23). This tutorial simply involved the robot sweeping slowly back and forth in the environment; participants were instructed to get close to the robot to block it and then move around to see the robot resume. Feedback systems were turned off during training so no learning bias would be introduced.

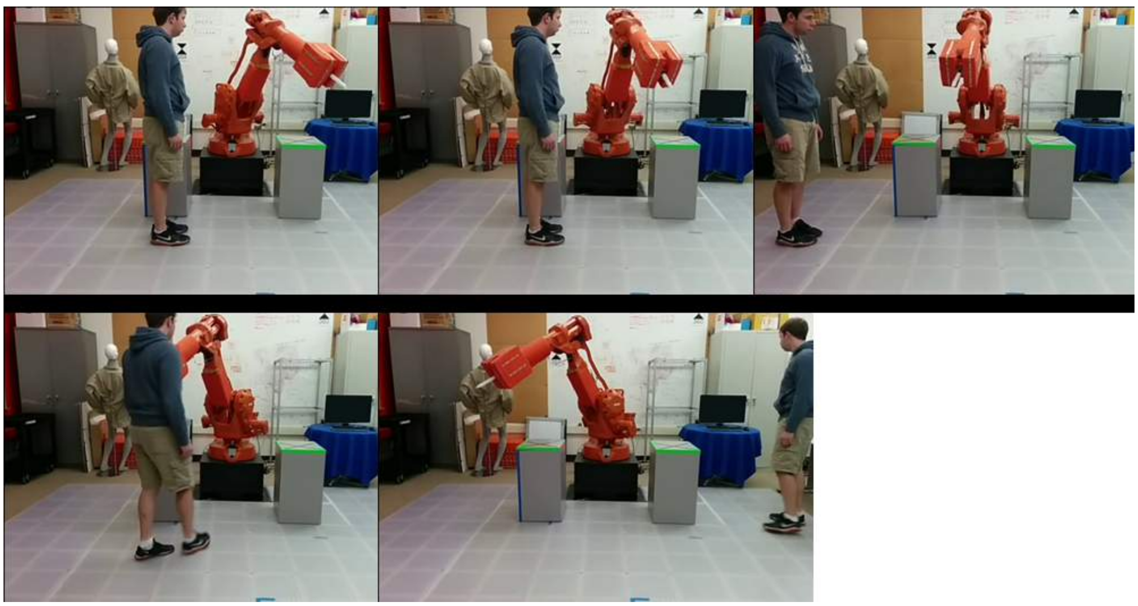

Figure 23: The robot sweeps back and forth at a low speed and near full extension, allowing the participant to get personal experience with the safety system

3.5.1.4 Practice Fin After gaining some familiarity with system in operation, participants performed the first fin assembly task with the robot disabled. This dry run ensured that participants understood how they were expected to perform the task, and gave them a chance to handle the parts and cart before performing their actual procedure. Here we would check their work and confirm they were doing it properly, or rectify causes of error.

3.5.1.5 Demonstrative Weld With the participant out of the environment, the manipulator performed the first weld action, at the same speed as it did during the assembly procedure (Figure 24). This gave the participants a chance to see the robot perform a weld, the run speed of the robot for the assembly procedure, as well as an approximate time for completing each fin. It also showed the path the robot would take while transitioning between cylinders.

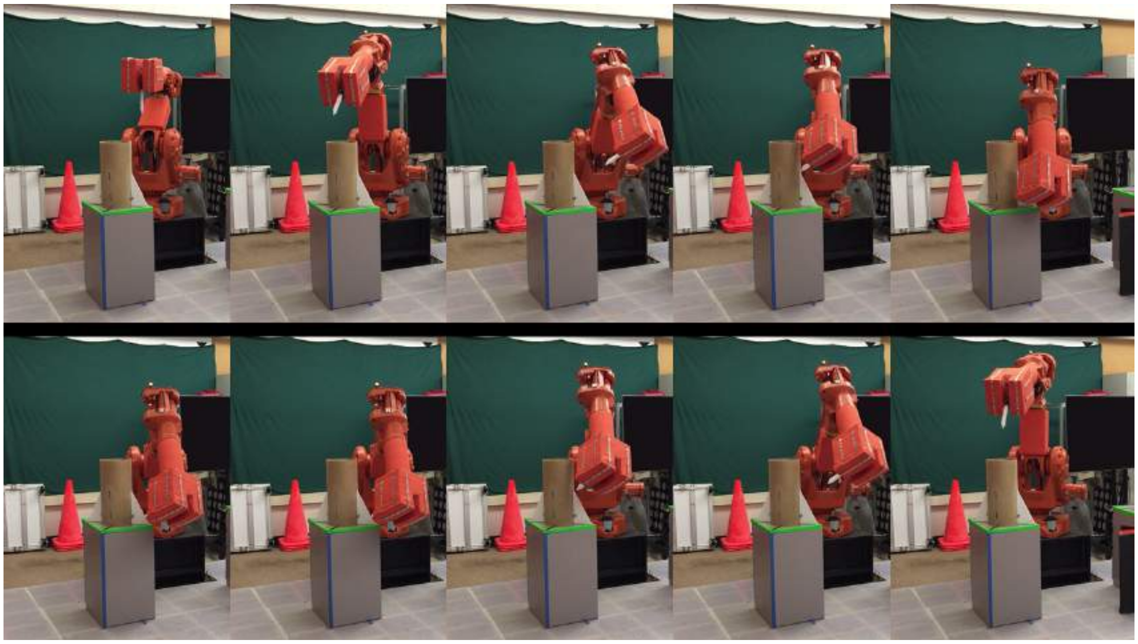

Figure 24: Demonstrative weld action. The robot moves from a neutral position to the fin, slows down, and runs along the seam. Then it returns to the neutral position.

3.5.1.6 Experimental Assembly Procedure After resetting the first fin and reiterating the general guidelines, the coordinator answered any remaining questions about the procedure. At this point, the experimental assembly procedure was carried out.

3.5.1.7 Post Experiment After completing the assembly procedure, we administered an exit survey. The survey investigated the participant's comfort level with each of the different feedback modalities and their comfort levels regarding the robot in general. There were also three open-ended questions, permitting commentary on the system. If the participant had any final questions, concerns, or comments, they were addressed.

4 Results

To understand the effectiveness of the feedback modalities, we examined the results from our collaborative feedback user study. The study included a total of 74 participants and was conducted over a period of 10 weeks. We ran four different conditions: with arm-mounted lights, with the floor display, with the screen simulation, and without feedback (our control group). We targeted a minimum of 15 participants for each of the four conditions. We could not use data from 5 participants, because of feedback malfunctions or data logging errors. We also removed one of the participants during post-processing, as the participant was deemed non-compliant for performing the assembly procedure improperly. Optional video recording consent was obtained for 65 of the participants; the remaining three participants declined on the IRB consent form. We used the video recordings in analysis, thus we disqualified the three participants who declined permission. Our final count totaled 65 participants, with 15 participants for the arm lights, 16 for the floor display, 18 for the simulation, and 16 for the control.

The following questions guided our analysis of the different modalities:

- Did performance of the process directly improve with feedback when compared to the control?
- Is there more collaboration with the feedback modalities compared to the control?
- Was there a change in performance across different feedback modalities?
- Did worker/robot collaboration change across different feedback modalities?

In general, we used time measurements to address performance changes. We examined movement patterns for possible additional information on human-robot interaction within the workcell environment. We measured collaboration by examining different human-robot interaction types, periods, and frequencies. Overall, we compared total trial times, examined heatmaps of participant placement in the environment, categorized statuses of blockages and waits, and analyzed survey results.

For clarity in the document, we define a few terms here. Blockage events refer to periods when the safety system was triggered, causing the robot to stop until the safety violation had been cleared. Wait periods refer to periods when the robot's next task was not ready, forcing it to wait idly. Specifically, all wait periods were triggered when participants failed to have the next fin ready for robot "welding."

To evaluate performance of participants, we examined the overall runtime of the robot (Section 4.1). Collaborative performance required more in-depth analysis of data including location of workers over time (Section 4.2), wait periods (Section 4.4), and exit survey results (Section 4.6). We also investigated collaborative performance through classified blockage reactions (Section 4.3) that focused on:

- Overall blockages (Section 4.3.3)

- Reaction type percentages (Sections 4.3.4, 4.3.5, 4.3.6, 4.3.7, and 4.3.14)
- Combined reactions (Sections 4.3.8 and 4.3.11)
- Time-per-occurrence analysis (Sections 4.3.9 and 4.3.13)
- Cart and person data (Section 4.3.10)
- Filtering reactions (Section 4.3.12)

4.1 Overall Performance of Participants

One way to measure the performance of a process is to look at the time it requires for completion; the faster the completion time, the more throughput a process can handle. We therefore examined total runtimes by condition.

We measured time from when the participant began moving at the start of the assembly procedure to when the robot reached the final position in the movement script. The fastest theoretical robot completion time, (without any waits or blockages) was approximately 220 seconds. The fastest theoretical assembly procedure completion time, including the participant's portion of the task (they are required to complete one fin assembly before the robot begins moving) was approximately 250 seconds.

The average runtimes for the three feedback conditions and control are shown below (Figure 25). We can see that all of the average runtimes were much greater than the theoretical minimum. Overall, it appeared that our participants either were slower at completing their tasks than the robot, or that the points at which they chose to collaborate with the robot did not reward them with better runtimes.

Average runtimes ranged from around 415 to 450 seconds. Feedback of all types had faster average assembly procedure runtimes by at least 20 seconds when compared to the control group, and thus had increased performance. However, the figure shows that every condition also experienced large variation in total runtimes. This leads us to believe that the variations in time are systemic, caused by differences in individual participant approaches to completing the trial and not by any of the feedback systems. For example, some participants actively sought to remove themselves from the vicinity of the robot whenever it was in motion, regardless of what task they were currently performing. While the time the robot was blocked decreased drastically for these cases, often the wait time increased. Other participants ignored all feedback while working (in some cases even after observing the state of the robot from feedback), causing a large number of robot blockages and few wait periods. Many of the participants performed the assembly procedure using a combination of approaches. Overall, feedback sped up the runtime but further experimentation is needed to determine whether the performance increase is statistically significant.

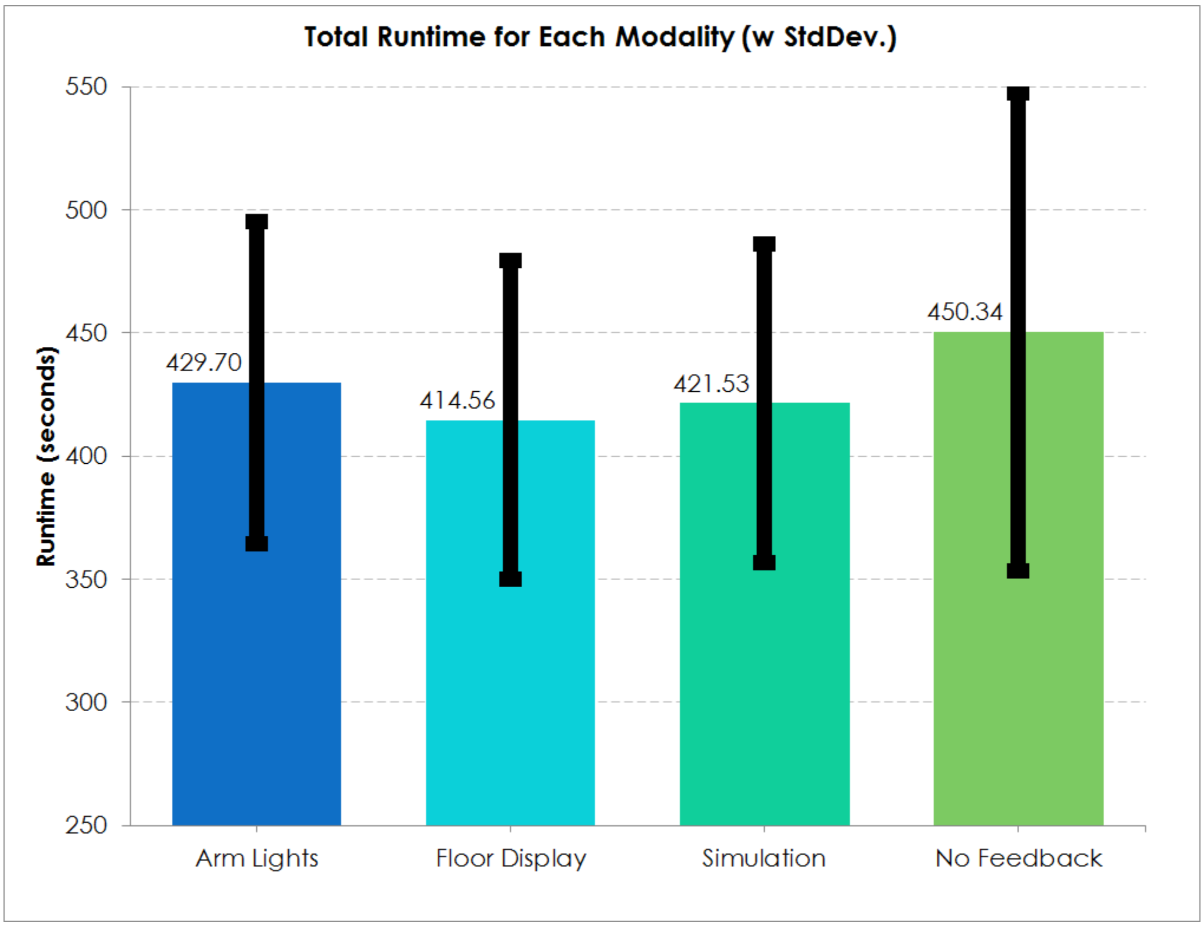

Figure 25: Overall time for each feedback type

4.2 Positional Information for People and Objects

One way to evaluate how people are interacting with the robot is to look at where they stand and move in the workcell during completion of the assembly procedure. We relied on heatmaps generated as a projection of the voxel grids from both the safety and the danger grids to supply us with a colored narrative showing where the participant moved during the procedure. We summed the occupied voxels along the vertical axis of the environment, providing a higher concentration at coordinates where more voxels were occupied between the min and max height. This created a projection for each trial. Then, all of the projections for single trials were summed and normalized by the largest value, creating a composite heatmap for each participant in a condition. We generated the final heatmaps by summing all of the individual composites for participants in a condition into an overall composite, which can be seen in Figure 26.

Examining the final composite heatmaps, we see no substantial difference in the pattern that participants traversed during their time in the environment. The floor display does show longer occupancy on a path closer to the cleaning station, but this could be caused by the final standing location of participants in the floor modality. This most likely means that the locations traveled by participants depended more on the participant task completion style than on the feedback modality. Without isolated locations as part of the task, different amounts of time spent on different tasks will not necessarily be clearly articulated or separable.

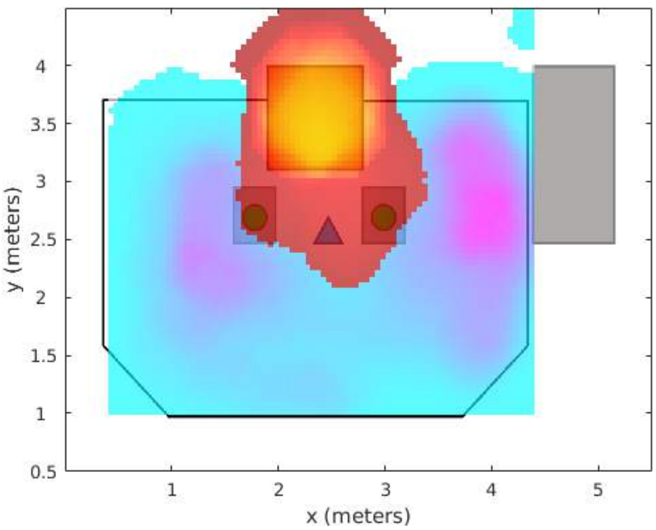

(a) Arm-mounted lights

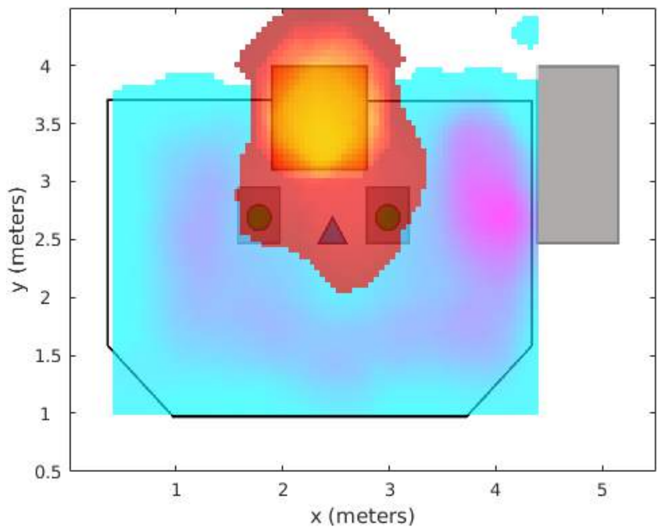

(b) Floor display

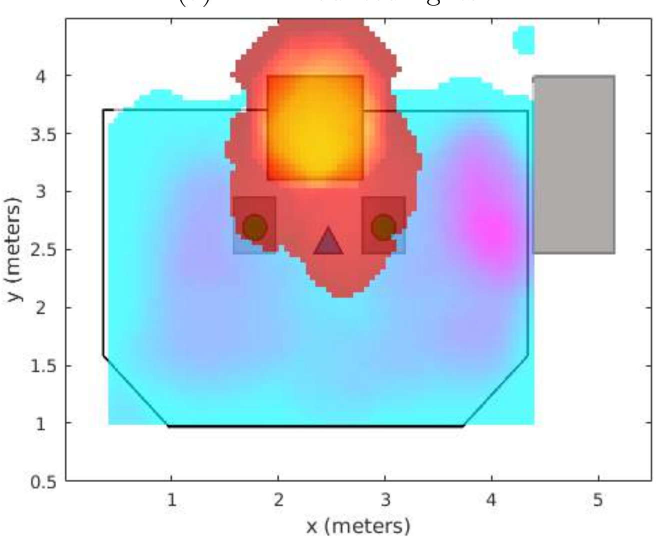

(c) Simulation

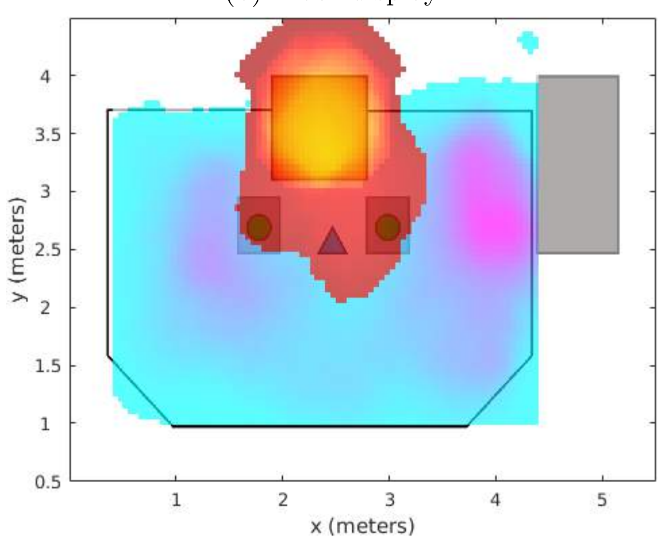

(d) No feedback

Figure 26: Composite heatmaps for each condition

4.3 Blockage Events and Participant Reactions to Blockages

In order to measure collaboration between the participant and robot in our experiment, we analyzed human-robot interactions; specifically, times during which blockages of operation occurred.

Interference could be caused by one team member working too slowly (or quickly) compared to the other or by physically occupying a shared workspace needed by the other. Either case results in idle time, which could decrease performance. We analyzed blockage data by number of occurrences, length (time) of events, and ratios. The system recorded time information for blockages and wait periods. For a more fine-grained analysis, we decided to add classification labels to the data.

4.3.1 Annotating the Videos

In order to more closely examine the collaboration between the robot and our participants, we looked at the video record for each participant. We recorded every blockage's start and end time, and classified each according to origin of blockage event, participant observation of feedback, and whether or not corrective action was taken. We focused on blockages caused at least in part by the participant; cart-only blockages often occurred only at the end of runs and varied greatly in occurrences and time. Our video annotation enabled us to identify and label four different categories of blockage reactions as shown in Table 1.

	No Change	Course Change
Misses Feedback	Oblivious	Independent
Notices Feedback	Dismissive	Responsive

Table 1: Different blockage reaction classifications

For reactions in which the participant noticed the feedback, we also recorded the notice time. If the participant could be seen actively observing the feedback after the beginning of the blockage event, we recorded this point of observation as the notice time. If it was not clear when the participant noticed the feedback, the notice time was recorded as when the participant first made a physical motion (e.g. change of posture, footing, weight shift) to clear the blockage.

4.3.1.1 Blockage classifications When participants take no notice of the feedback and do nothing to change their actions, we classify the blockage as an Oblivious Reaction. Here, the participant is unaware of the robot's situation, or sees the robot (but not the feedback), and continues to work (Figure 28).

Independent Reactions are defined as occurring when participants block the robot while it is running, notice the blockage and change their action to clear the blockage, all without observing the feedback (Figure 27). For example, the participant might see the robot approach them using peripheral vision and change their actions. Independent Reactions figured prominently in the control conditions, where feedback was unavailable.

Dismissive Reactions are defined as instances when participants notice the feedback but either do not change their actions, or try to correct (e.g., by moving into a slightly different position) unsuccessfully before moving on to the next step in their assembly list (Figure 29). Dismissive Reactions occurred in all conditions. Cases of Dismissive Reactions which occurred in the control condition were special cases of Oblivious Reaction in which the participant attentively watched the robot come to a halt in the middle of an action, then proceeded to continue blocking the robot while completing a part of their task.

Responsive Reactions occur when participants notice the feedback, heed the blockage notification and change their actions to clear the blockage based on the information provided (Figure 30). Responsive Reactions are considered a desired response for the experiment.

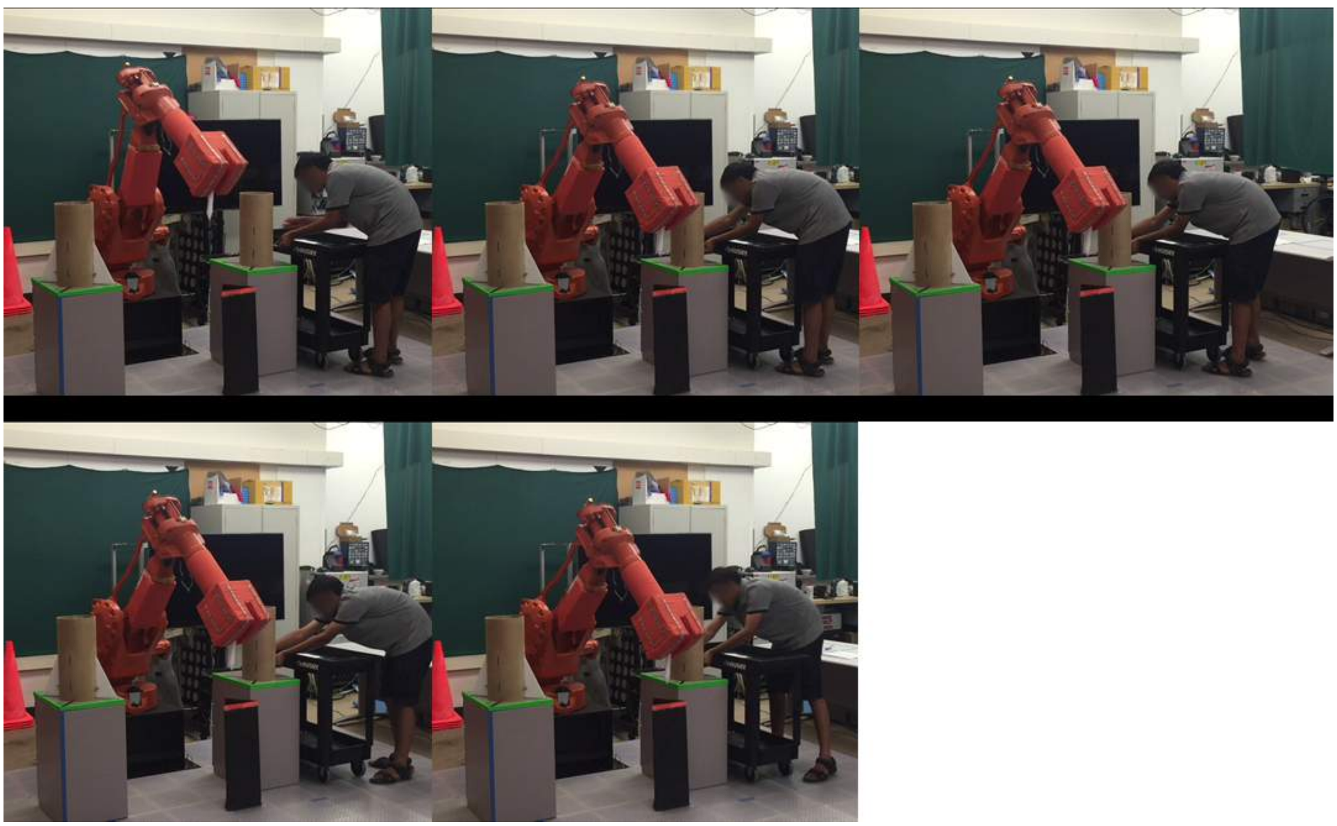

Figure 27: Independent Reaction (shown with the control case): the robot and person are working. A blockage occurs. The participant notices the blockage event without the use of feedback. The participant halts work and changes position to unblock the robot. Finally, both continue working. Feedback is never observed, however something else prompts a reaction from the participant.

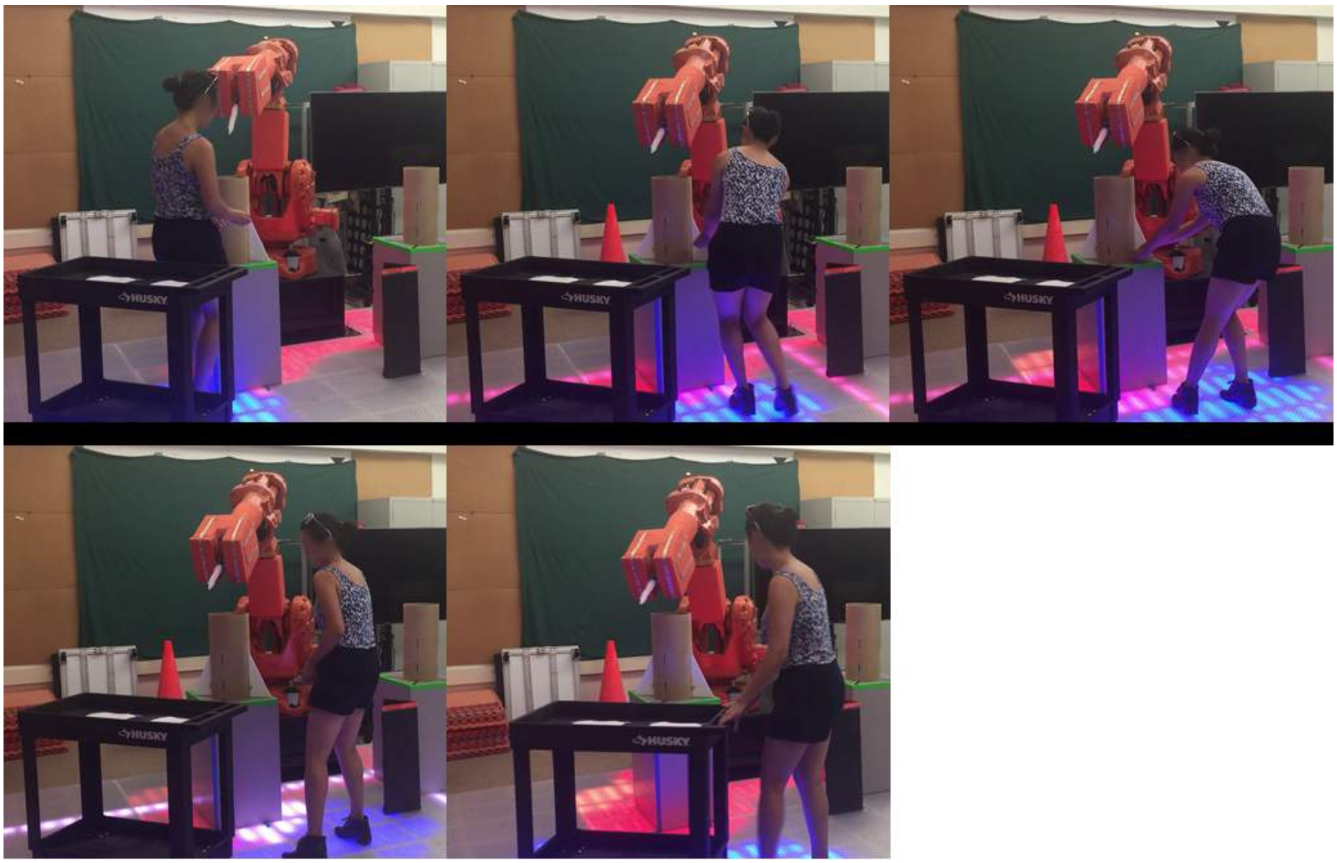

Figure 28: Oblivious Reaction (shown with the floor lights): person and robot are working together. Robot becomes blocked. Worker does not notice the feedback or that the robot is blocked, and keeps working. After the worker's task is complete, they move away. Robot resumes action. Feedback is never observed.

Figure 29: Dismissive Reaction (shown with the arm lights): robot and worker are working together. Robot becomes blocked. Worker observes the feedback indicating a blockage. Worker continues to work until their task is complete. Worker then moves away and clears the blockage. Robot resumes its tasks. Feedback is observed, but no corrective action is taken.

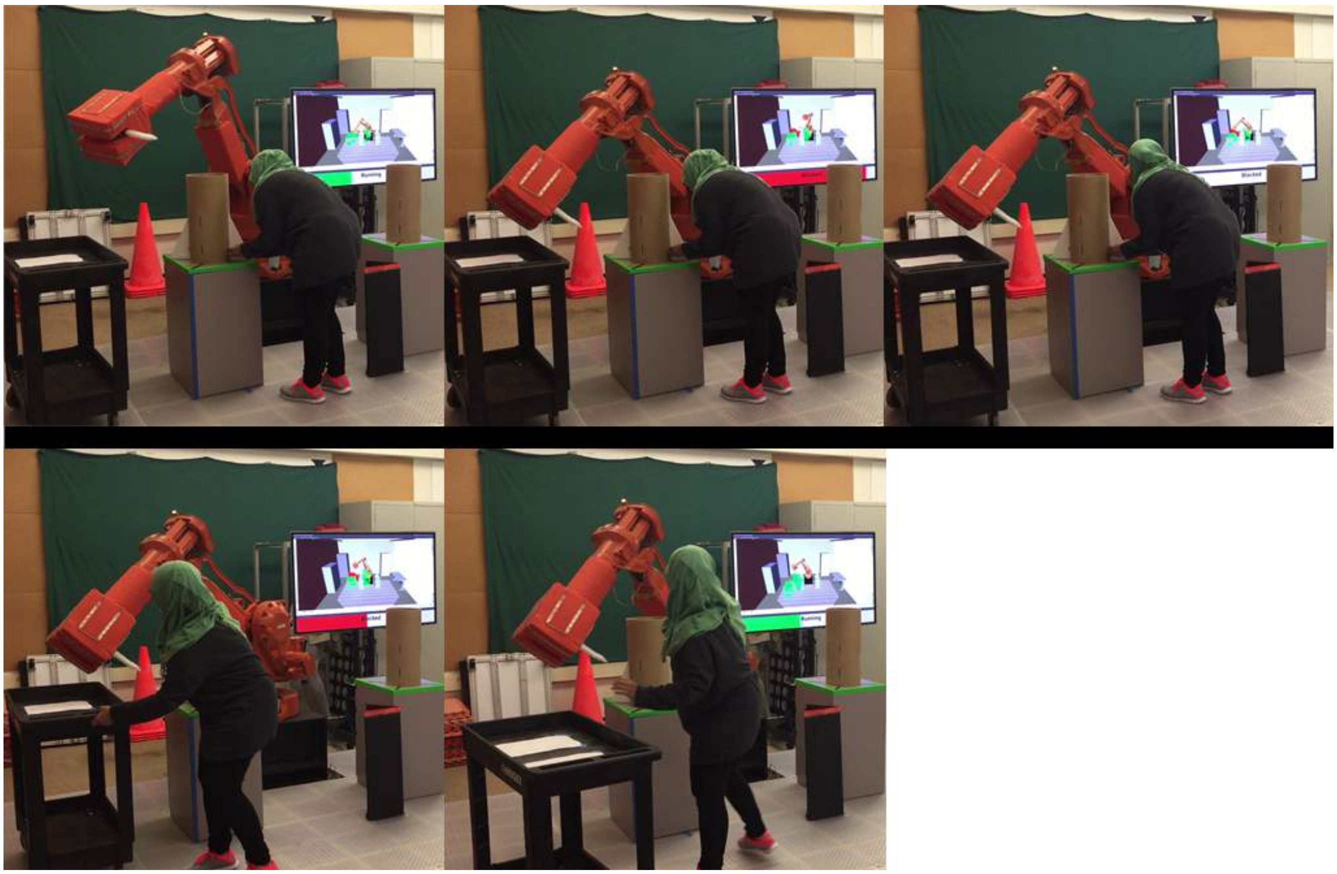

Figure 30: Responsive Reaction (shown with the simulation): robot and worker are completing tasks in parallel. Robot becomes blocked. Worker observes the feedback. Worker takes immediate or near immediate action to correct the blockage. Robot and worker both resume tasks. Feedback is observed, and corrective action is taken.

Since the control condition obviously did not have feedback, a comparison of the three feedback modalities to the control require the data to be generalized slightly. The definitions of different categories of blockages do not correlate between feedback and control condition; Dismissive and Responsive Reactions are undefined for the control. An equitable comparison can be made by summing along the columns in Table 1, merging Responsive with Independent, and Oblivious with Dismissive. It therefore becomes possible for us to examine total occurrences of each type of event (for comparison between the three feedback modalities) and the compositions of different modes (for comparison against the control group).

In addition to blockages, we annotated the wait periods for the robot. These were times when the robot's next task was not ready, preventing it from moving on. We measured wait periods from the stoppage of the robot to fin task completion by the participant. Any amount of time the robot was forced to spend waiting beyond task completion (e.g., for the participant to move out of the robot's way), was recorded as a blockage.

4.3.2 Acronyms

For convenience, acronyms are used to describe the relationships between data sets used in graphs of blockages and are defined here.

- Time per Occurrence (TpO): the total time for all occurrences of a single type of blockage for a participant divided by the number of occurrences.
- Average Time per Occurrence (ATpO): the average of all TpOs for a given condition. By focusing on the average for a condition, ATpO removes the weighting based on number of occurrences from each participant.
- Total Time per Occurrence (TTpO): the total time for all occurrences of a single type of blockage for a condition type divided by the number of occurrences. Unlike ATpO, TTpO weights all occurrences of a blockage type equally.
- Ratio to All Blocks (RtAB): the percentage of a single type of blockage divided by the total of all blockages. This percentage can be based on number of occurrences or total time of events.

4.3.3 Total Blockages

The total occurrences of blockage events is shown in Figure 31, combining Independent, Oblivious, Dismissive, and Responsive blockages. This figure indicates a tally of all types of blockages for each participant. We can see that the number of blockages varies widely from as few as one to at least eleven per modality per participant. The simulation modality was, on average, blocked at least 2.2 occurrences more than the other three conditions. On average, both of the embedded feedback modalities had fewer blockage occurrences than the non-embedded simulation modality.

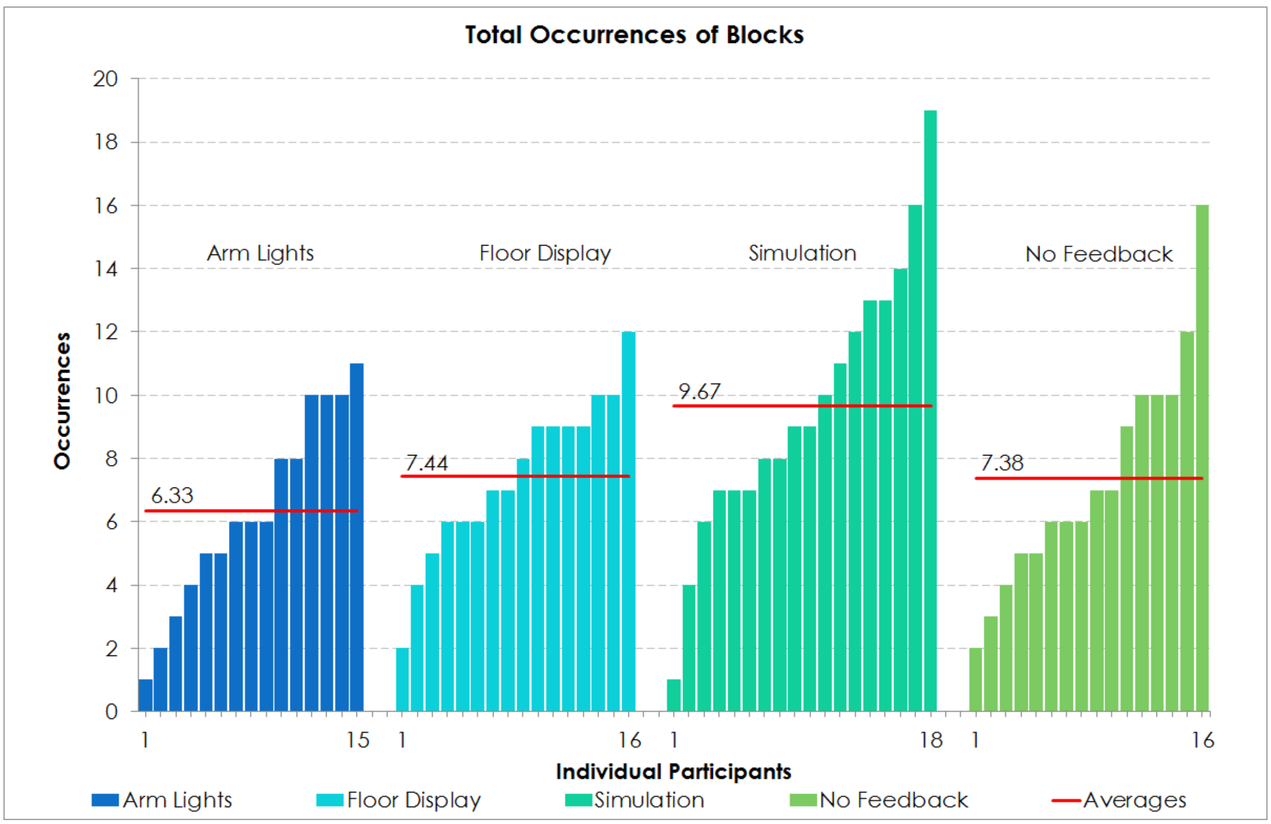

Figure 31: Total number of blocks in all situations, occurrences per participant. Averages per modality are shown in red.

The total time of blockage events is shown in Figure 32, again combining all four blockage types. Here, time from every event is summed for each participant, and then averaged for each modality. We can see that control participants averaged at least 12 seconds longer for blockages than those using feedback, which suggests that participants without benefit of feedback had reduced cooperation between participant and robot. Additionally, we see that the embedded feedback modalities perform on average at least 23% better than the control condition.

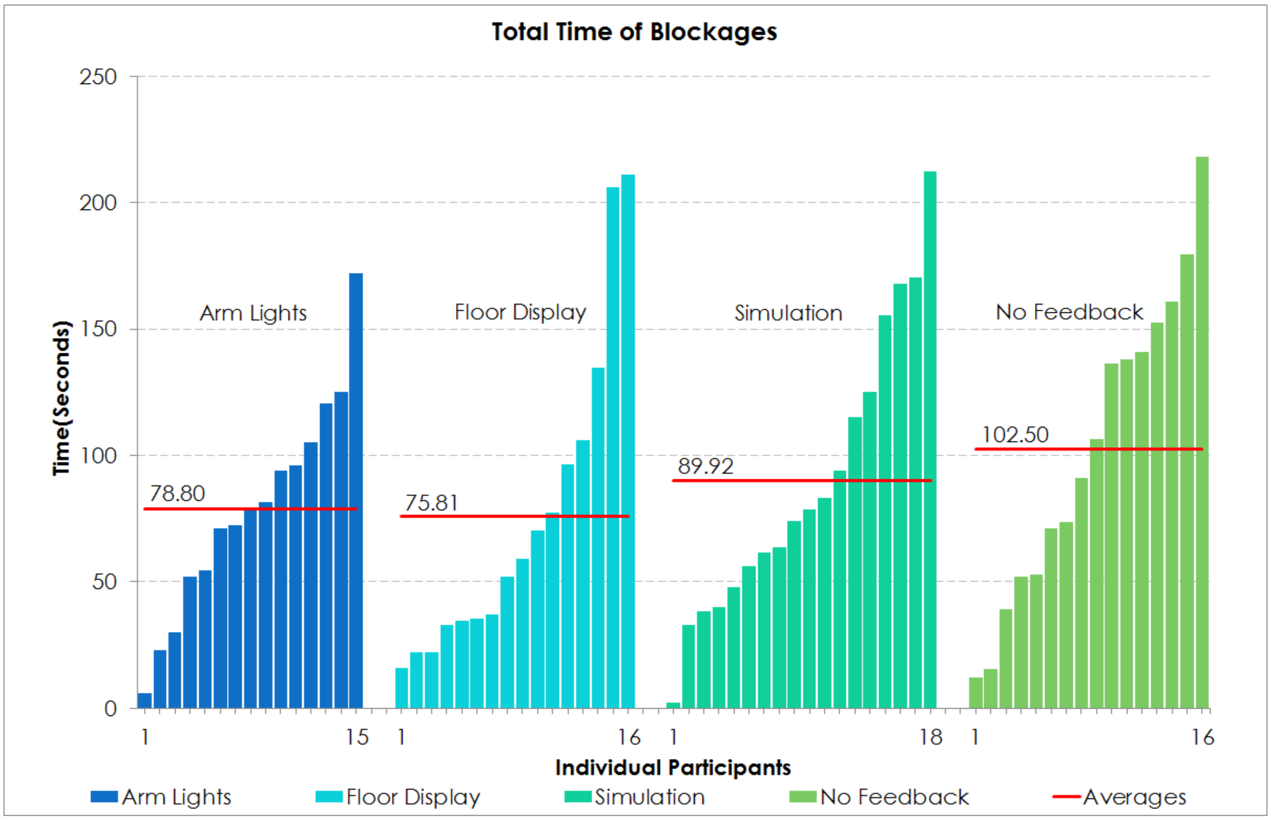

Figure 32: Total combined time of all block situations per participant. Average times per modality are shown in red.

4.3.4 Responsive Reaction Proportion

We constructed charts of the four different types of blockage reactions based on their overall proportion of occurrences per participant for each feedback modality (RtAB), allowing us to consider the relationship between feedback modality type and blockage type. Proportional analysis helps determine rate of common situations across different modalities.

Responsive Reactions are defined as events in which the participant notices the feedback, heeds the blockage notification and changes their actions based on the information provided. Figure 33 shows the ratios of Responsive Reactions relative to the total blockage occurrences. Participant response to feedback varied widely, from 0% to 100% of responses using feedback. We can also see that, while the simulation and floor display average within 5% of each other, the arm lights average more than 10% lower for responsive reactions. Most importantly, when feedback of any kind is provided, at least 30% of the blockages on average caused by the participant were resolved using the feedback information.

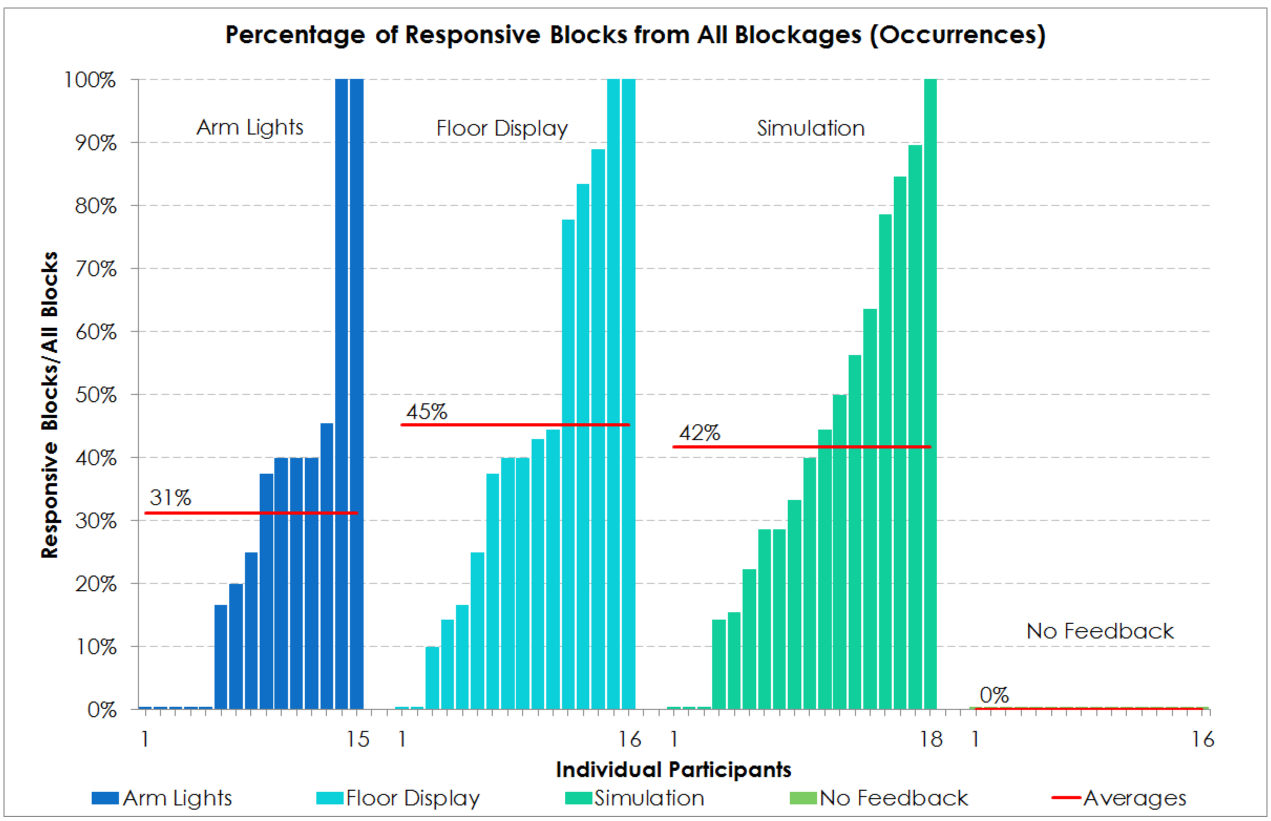

Figure 33: Responsive Reactions, ratio of occurrences per participant. Average occurrences per modality are shown in red.

4.3.5 Dismissive Reaction Proportion

Dismissive blockages are defined as blockages in which participants notice the feedback while blocking the robot but do not change their course of action. Figure 34 displays the ratios of Dismissive Blockages relative to the total block occurrences.

The figure shows that the proportion of Dismissive Reactions using arm lights feedback averages 18% lower than the floor lights and 14% lower than the simulation. This indicates that the embedded floor lights have the highest average proportion of observed and ignored feedback, whereas the arm lights have the lowest average proportion. Coupled with the same pattern seen in the previous section, these results lead us to conclude that the floor lights were the most commanding of participant attention of the three modalities. Although there are a few occurrences of Dismissive Reactions in the control condition, these special cases reflect situations in which the participant watched the robot approach them, directly blocked it, and then continued working.

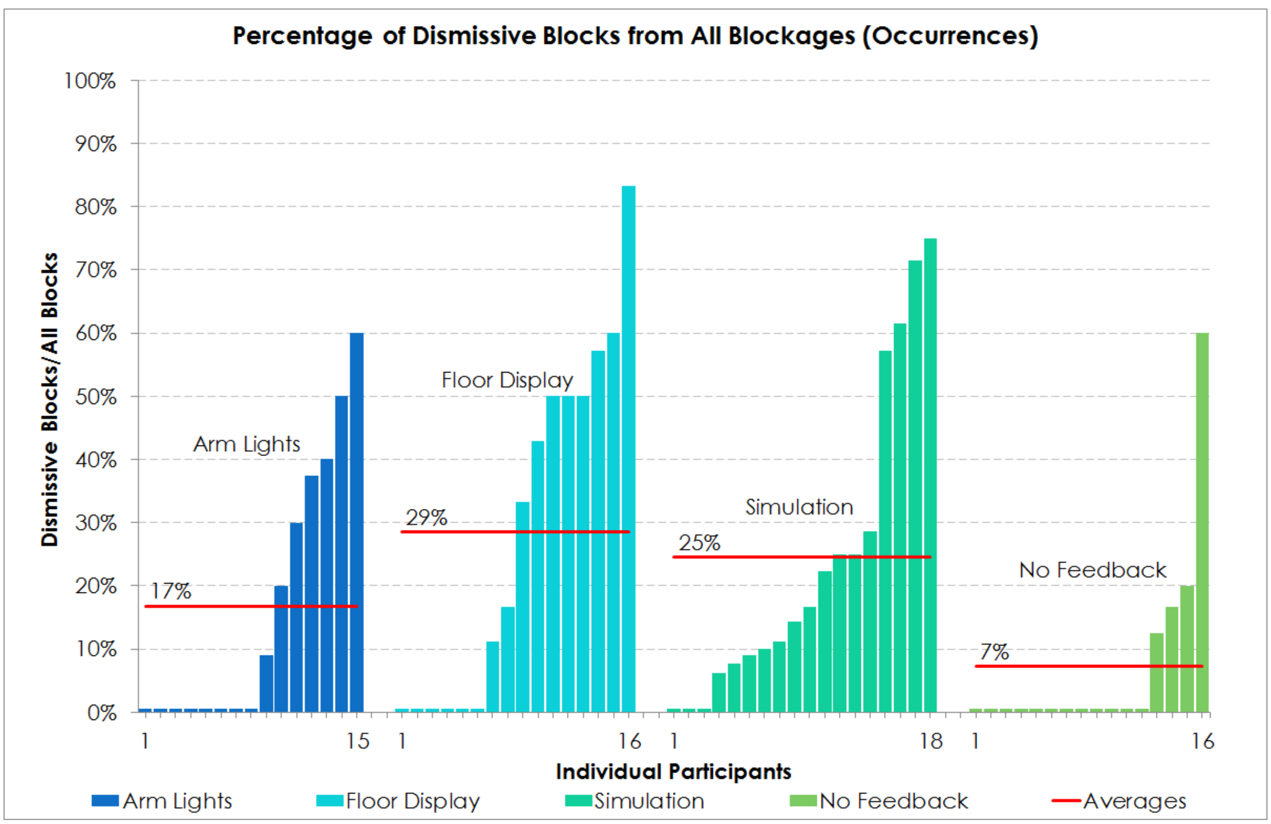

Figure 34: Dismissive Reactions, ratio of occurrences per participant. Averages per modality are in red.

4.3.6 Independent Reaction Proportion

Independent Reactions are defined as reactions in which the participant clears a blockage by changing paths without using the feedback. Figure 35 shows the proportion of total blockages of the Independent Reaction type across the different feedback modalities. We can see that the largest proportional occurrences of Independent Reactions appear in the control condition. This relationship is expected because only two reaction options are really possible for the control condition; without feedback, either the person reacts to clear the blockage (Independent Reaction) or continues to work through the blockage (Oblivious Reaction).

The graph shows that approximately half of the participants utilizing the simulation condition have at least one occurrence of an Independent Reaction. It appears that unless the participant makes the effort to look directly at the simulation screen, which is located to the side/back of the workcell, the feedback goes unnoticed. This extra effort could cause participants to sometimes rely on instinct instead of feedback to clear blockages, and indicates that non-embedded feedback is used less than either embedded method. As a side note, the arm lights modality average may in fact be larger than calculated; our annotation method did not allow us to discern between looking at the robot and looking at the feedback in this case.

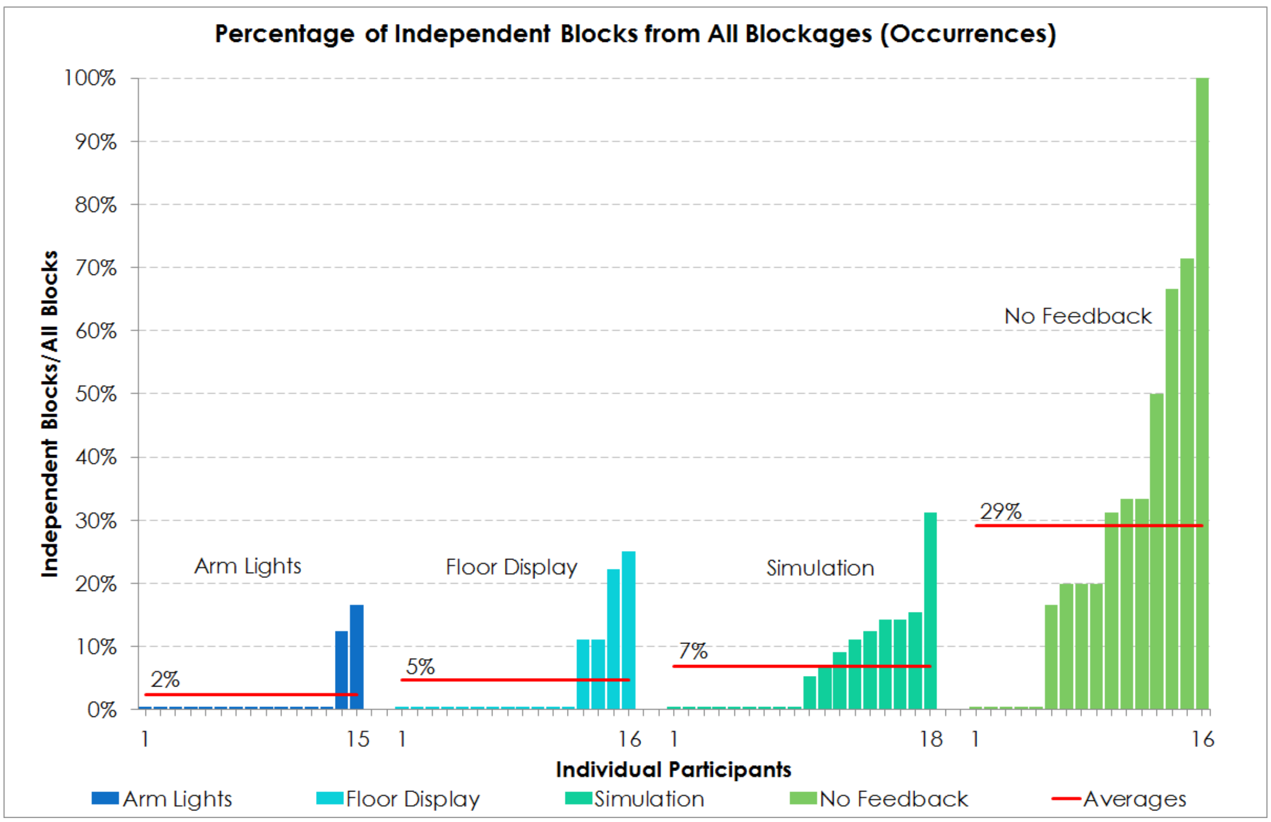

Figure 35: Independent Reactions, ratio of occurrences per participant. Averages per modality are in red.

4.3.7 Oblivious Reaction Proportion

Oblivious blockages occur when the robot is blocked and the participant fails to observe the feedback or perform corrective action. In Figure 36 the ratio of Oblivious Reactions to all blocks (occurrence-based) is illustrated. The control condition has the highest average percentage of Oblivious Reactions, as expected, since its participants could not have feedback-based resolutions. Examining the averages from the feedback conditions, we see that the arm lights modality has approximately twice the average percentage of Oblivious Reaction occurrences when compared to the simulation and floor display. This indicates that when blockages occurred, participants were less likely to pay attention to the arm lights feedback than to either the floor display or simulation screen feedback. Four of the sixteen arm lights participants show that Oblivious Reaction were 100% of the blockage types, indicating that the feedback was missed during every blockage. This result was surprising. Given the fact that the mounted arm lights were quite literally in the participants' faces, it was expected that this modality would be the most difficult one to miss.

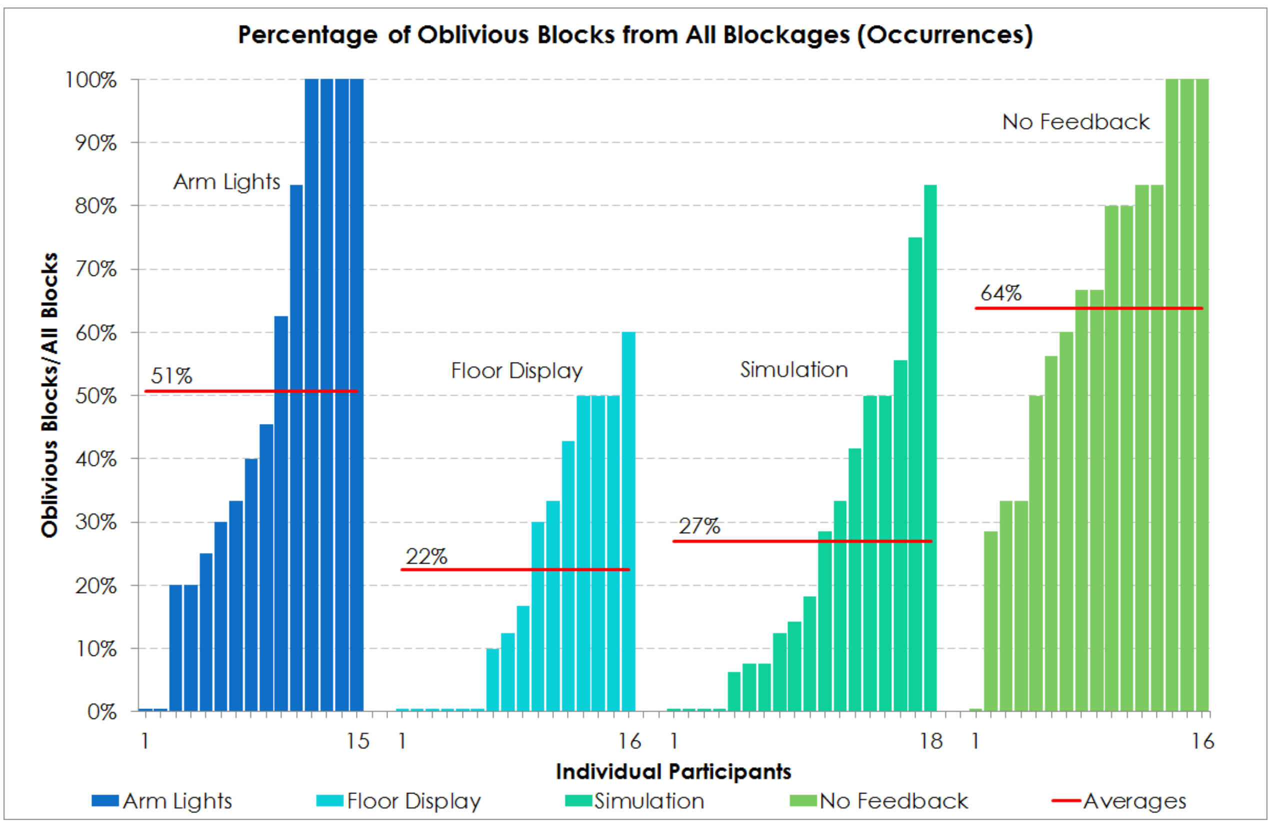

Figure 36: Oblivious Reactions, ratio of occurrences per participant. Averages per modality are in red.

4.3.8 Combined Feedback-Unaware Reactions

The surprising results from Oblivious Reactions leads us to investigate the non-feedback dependent reactions (specifically Oblivious and Independent) across the modalities. We believe that since both arm lights and simulation rely on the participant looking to a particular location in the environment for feedback, the two modalities might provide similar performance.

In Figure 37, we compare the occurrences of combined Oblivious and Independent Reactions between the different modalities. We can see that the occurences of the screen simulation and the arm lights differ by 11% on average.

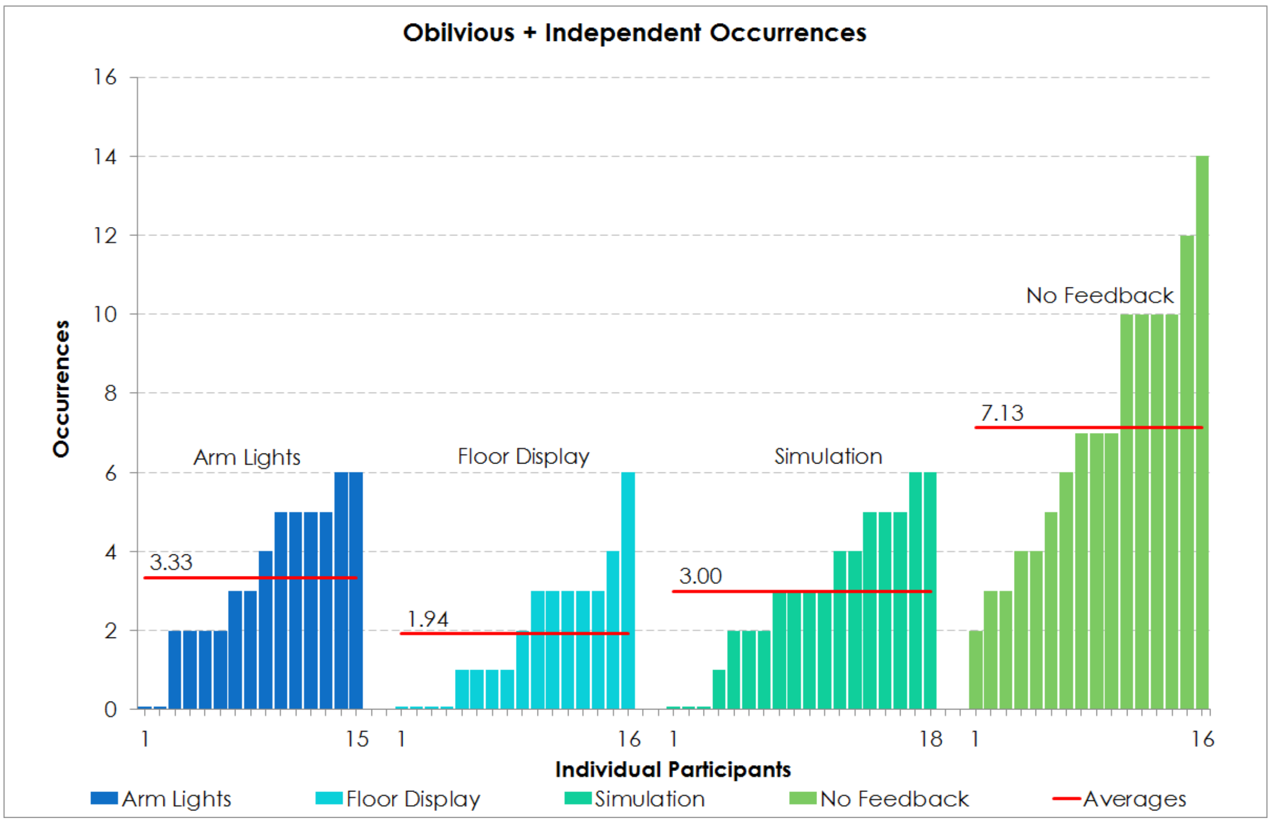

Figure 37: Total occurrences per participant of blocks in situations where feedback is not used for reaction. Average occurrences per modality are shown in red.

In Figure 38, the time per participant for Oblivious and Independent responses is compared between the different modalities. We see that there is an approximately 13 second difference in average time of feedback-unaware blockage reactions. This difference is enhanced by the single large peak in the arm light modality and the previously shown increase in occurrences. Thus, the arm lights on average have an increase in feedback-unaware blockage reactions when compared to the floor display and the simulation.

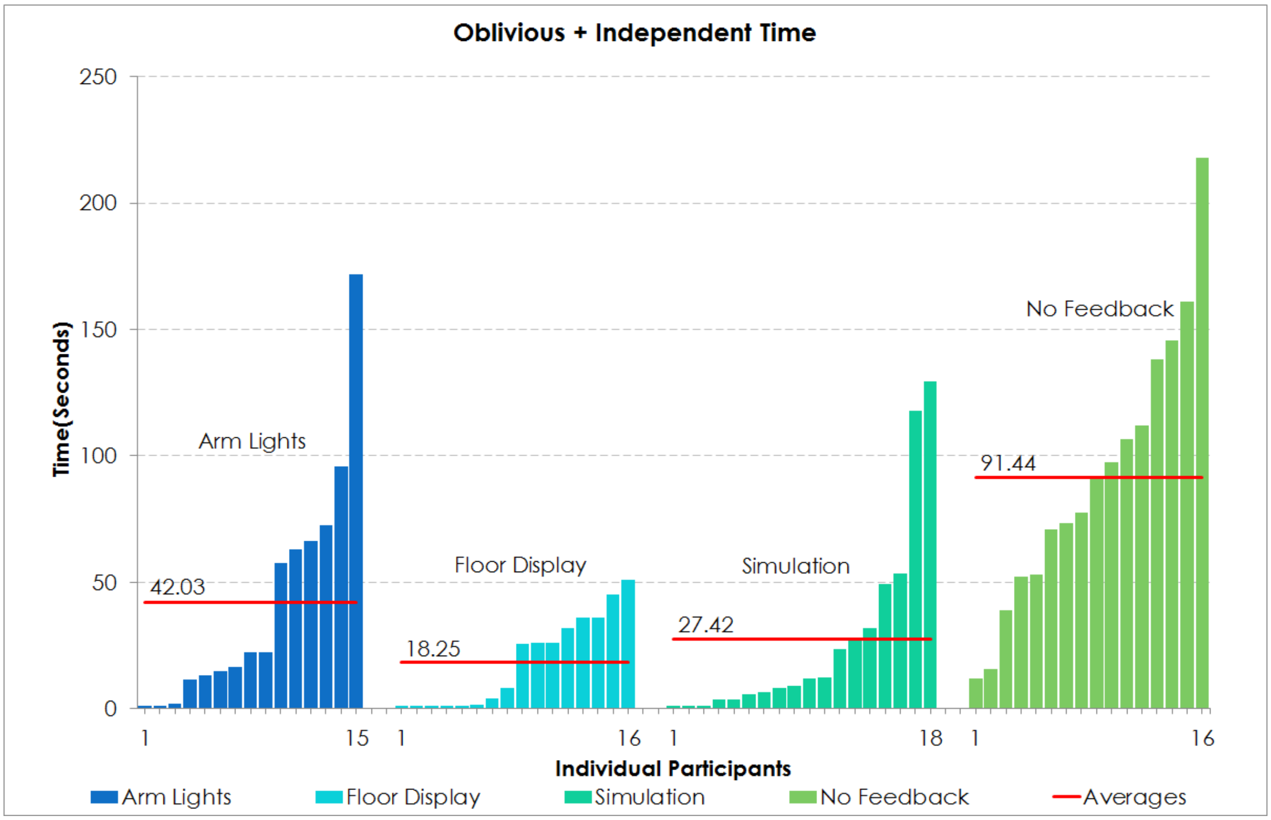

Figure 38: Total time per participant of blocks in situations where feedback is not used for reaction. Average occurrences per modality are shown in red.

4.3.9 Responsive Reaction Time per Occurrence

Average time per occurrence (ATpO) reports the total time taken for a participant to notice feedback and perform a corrective action to clear the blockage. ATpO is considered as a potential differentiation between the modalities because it can be used as a measure of overall ease of clearing blockages. Note that ATpO can not be and is not factored in for participants with zero occurrences. We considered the relationship between time per occurrence of feedback-based corrections and feedback modality (Figure 39). The figure shows that there is variation in time per occurrence for Responsive Reactions between participants, even within a single modality. This variation indicates that participants responded to the same type of feedback in different ways. The data shows that participants on average correct blockages with the embedded floor display more than a second faster than when using the arm lights or simulation. The worst case for the floor lights is also 2 seconds per occurrence less than the worst case of both the arm lights and simulation. Note also that the arm lights and simulation had overall average Responsive blockages within 1% of each other.

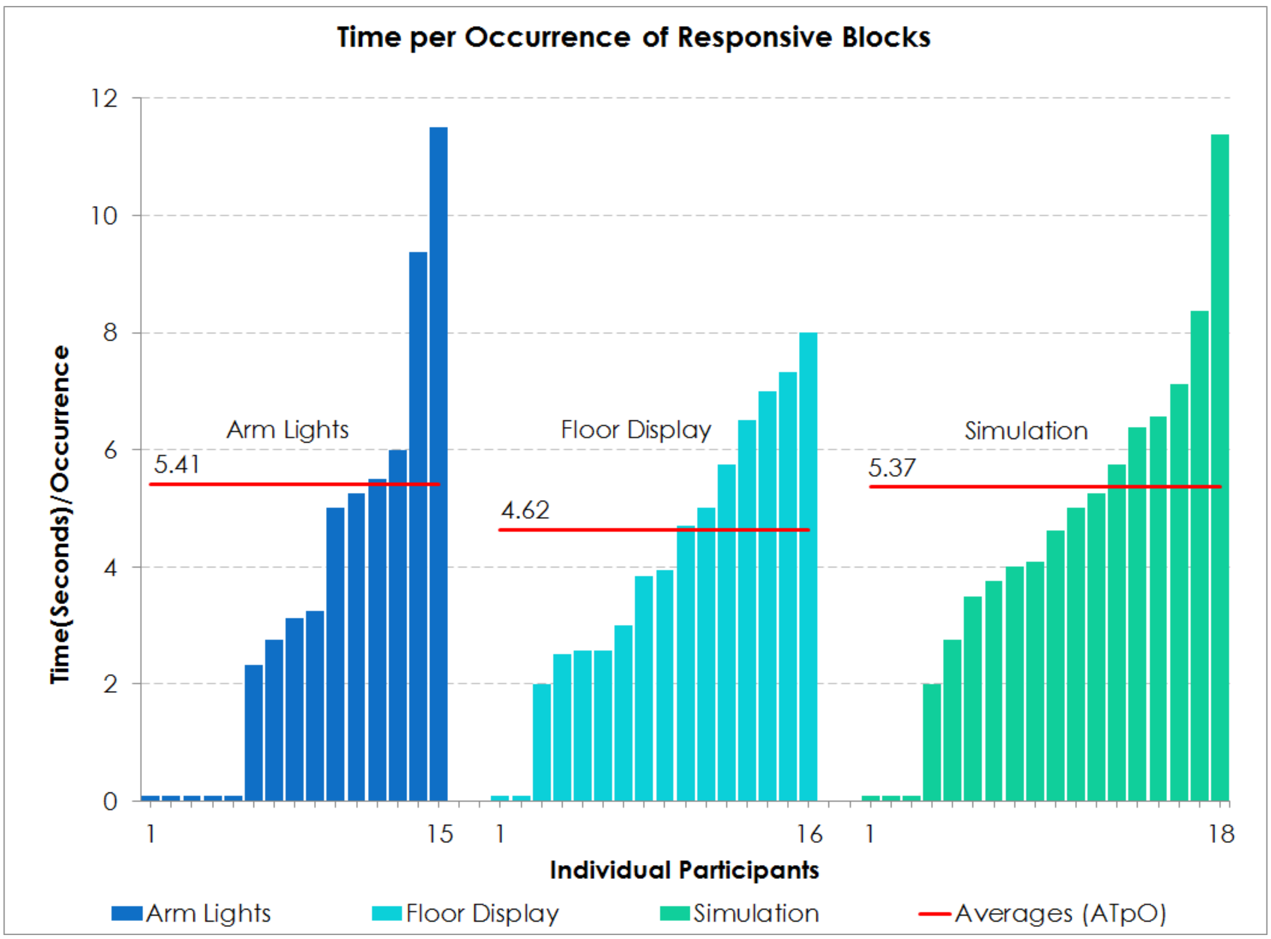

Figure 39: Responsive time per occurrence per participant. Overall averages per modality are in red.

4.3.10 Cart-Based Blockages

A Responsive Reaction for the cart is defined as an event in which the cart is causing a blockage and feedback is utilized to change action and clear the problem (i.e., move the cart). A Responsive Reaction for the person occurs if the person causes the blockage and clears it using feedback. Cart and person Responsive Reactions do not need to be separate; if both are blocking, both are included. We examined the average time per occurrence of the Responsive Reaction for both the cart and the person (Figure 40).

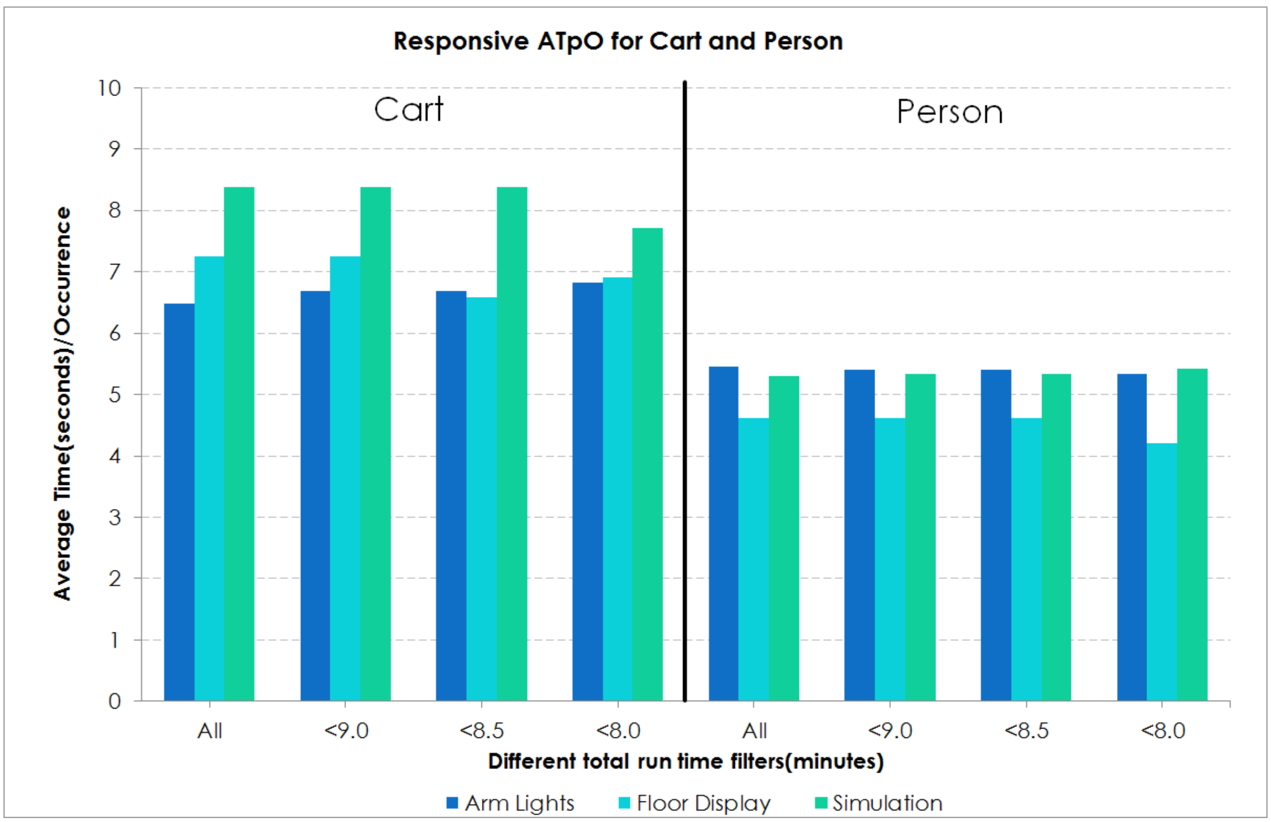

Figure 40: Responsive average time per occurrence for both cart-based and person-based blockages

The horizontal axis shows data filtering based on total runtime in minutes. Our intention for data filtering is to remove the longest runtimes of each condition and isolate only faster participants, a process which might remove outliers resulting from variations in participant performance.

The figure shows that on average participants responded to cart blockages about a second slower than personal blockages, regardless of total assembly procedure completion time. During the annotation process, we observed many cart blockages occurring near the end of the assembly procedure, specifically during assembly of fins 7 and 8. At this point, participants would return the cart to its starting position near the parts table and then step aside. This placement commonly caused a cart blockage with the robot arm. Since the participants had already stepped away from the cart, they needed to spend extra time walking back to unblock it. The figure also shows that the slowest average cart responses occurred when using the screen simulation feedback. We think this is due to participants moving out of the immediate assembly area where the non-embedded screen simulation could not attract their attention. We believe the two embedded feedback modalities are more effective at alerting the participant to blockages regardless of their standing position in the workcell.

4.3.11 Combined Feedback-Aware Blockages

We investigated how often each type of feedback was noticed during blockages by plotting the percentage of Dismissive plus Responsive Blockage reactions across all the participants. We used the combination of Dismissive and Responsive Reactions because both require observation of the feedback. This information provides a measure of how well a feedback modality commands the participant's attention, rather than its overall utility in correcting blockages.

Figure 41 indicates, for each participant, the percentage of blockages in which feedback was observed. In effect, this is the observation rate of the different feedback modalities. We can see that the arm lights modality has the smallest average number of feedback observations, 18% fewer than the simulation. The floor display was most observed, with an average 7% higher than the simulation. These results show that embedding feedback does not automatically increase the average observation rate during blockages, as the non-embedded feedback fell in the middle of the range.

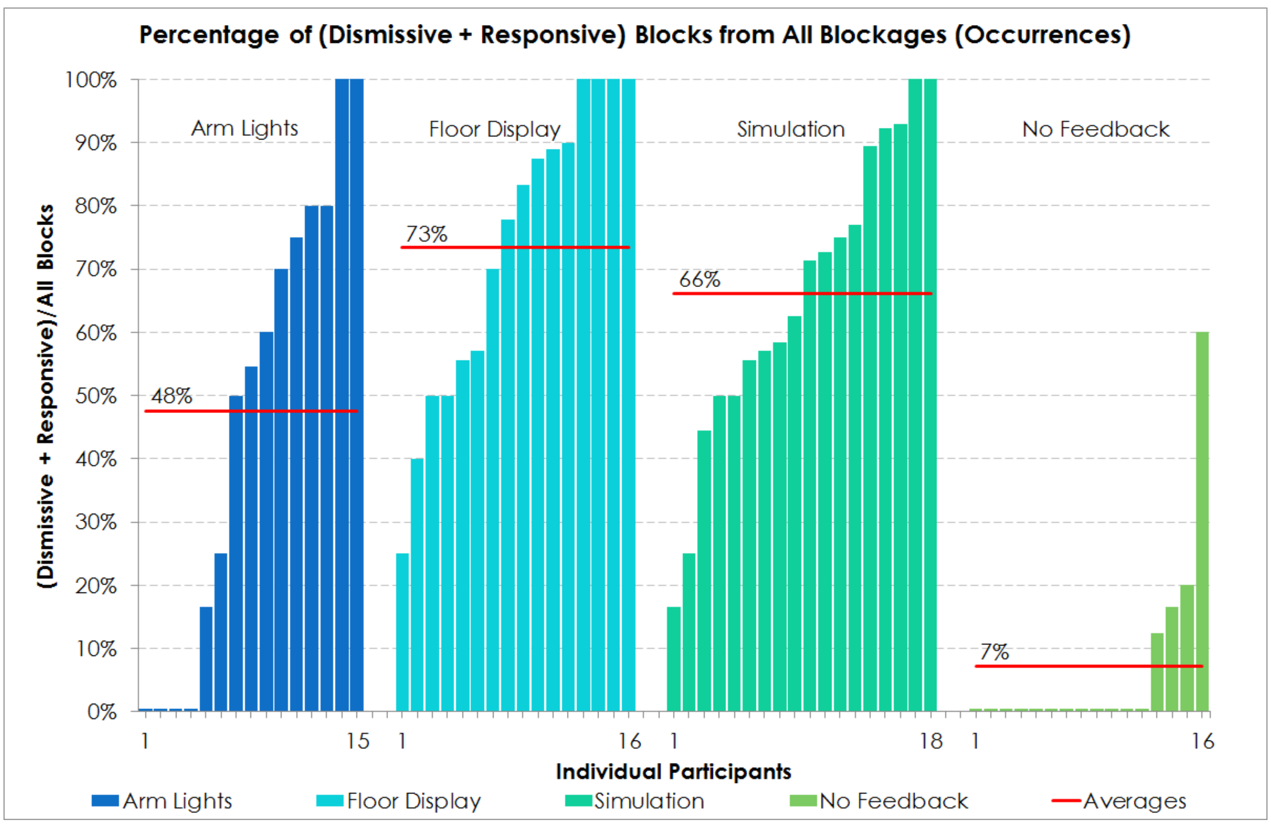

Figure 41: Combined occurrences of Dismissive and Responsive Reactions (both where feedback is noticed) ratio to all blocks per participant. Averages per modality are in red.

4.3.12 Top Seven Analysis

To focus on situations where feedback was observed consistently, we examined the percentage of occurrence of Responsive Reactions for the seven participants from each feedback modality with the highest notice rate of feedback (percentage of Dismissive + Responsive Reactions to all Reactions) (Figure 42). We selected the top seven participants, since this removes those with less than 60% of the blockage reactions occurring while observing feedback. We can see that even for participants who were most responsive to the feedback from the arm lights, the average percentage of Responsive Reactions remains at least 20% lower than that of the floor and simulation modalities. In this filtered data, the average floor lights Responsive Reaction percentage is 8% greater than the simulation modality, indicating that the floor lights had on average the highest Responsive rate.

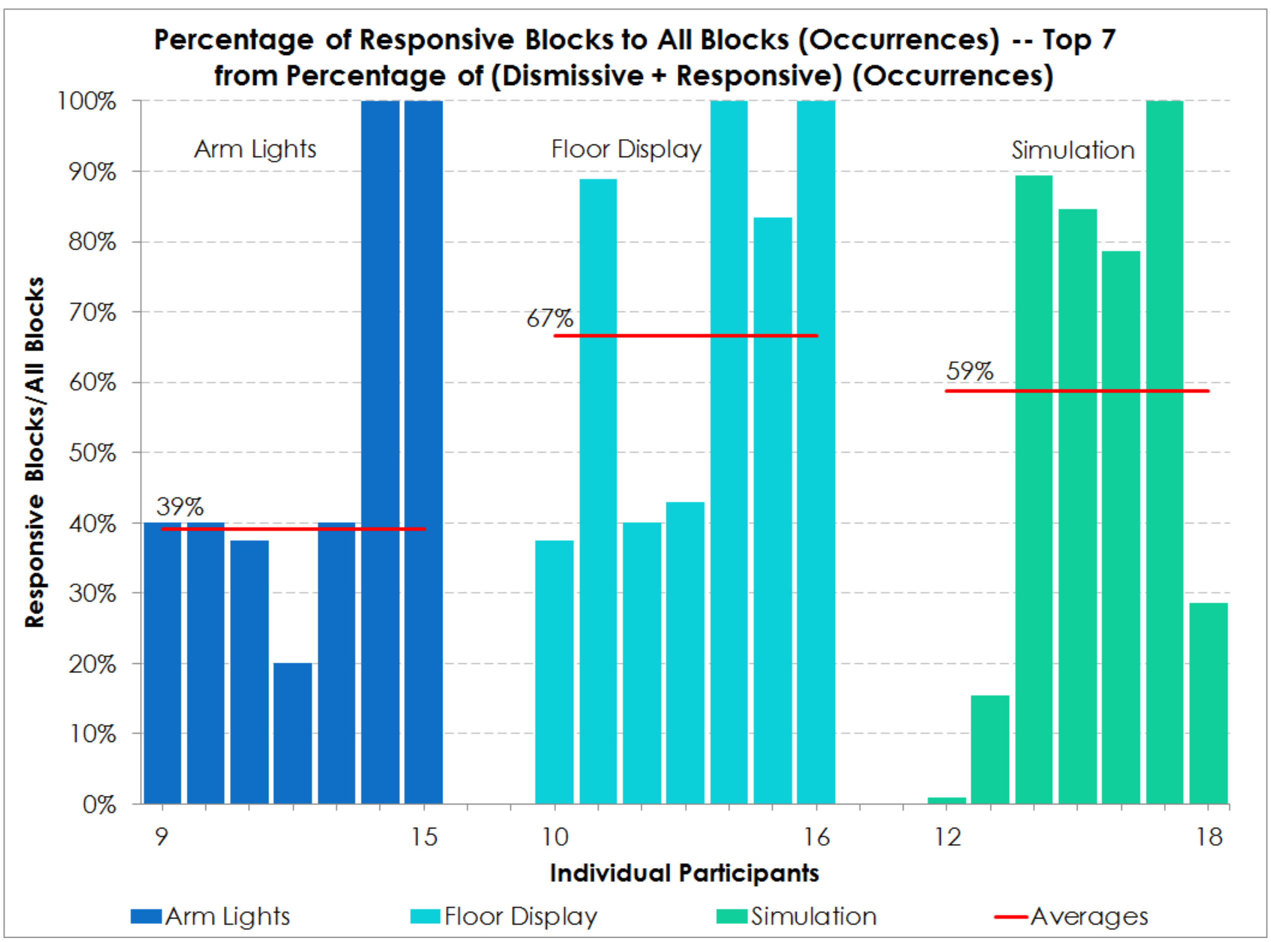

Figure 42: Responsive Reaction RTaBs for participants with the top seven feedback observation percentages per modality. Feedback observation percentage from Figure 41. Averages per modality are in red.

4.3.13 Average per Modality: Overall Time per Occurrence

Since the approach and performance of participants varied widely within each modality, examining specific participant information did not completely address the complexity of the data. Therefore, we decided to compare the different blockage reactions across conditions for the most complete analysis available. Figure 43 shows the ATpO from each different feedback modality type as well as the control group across each type of blockage reaction.

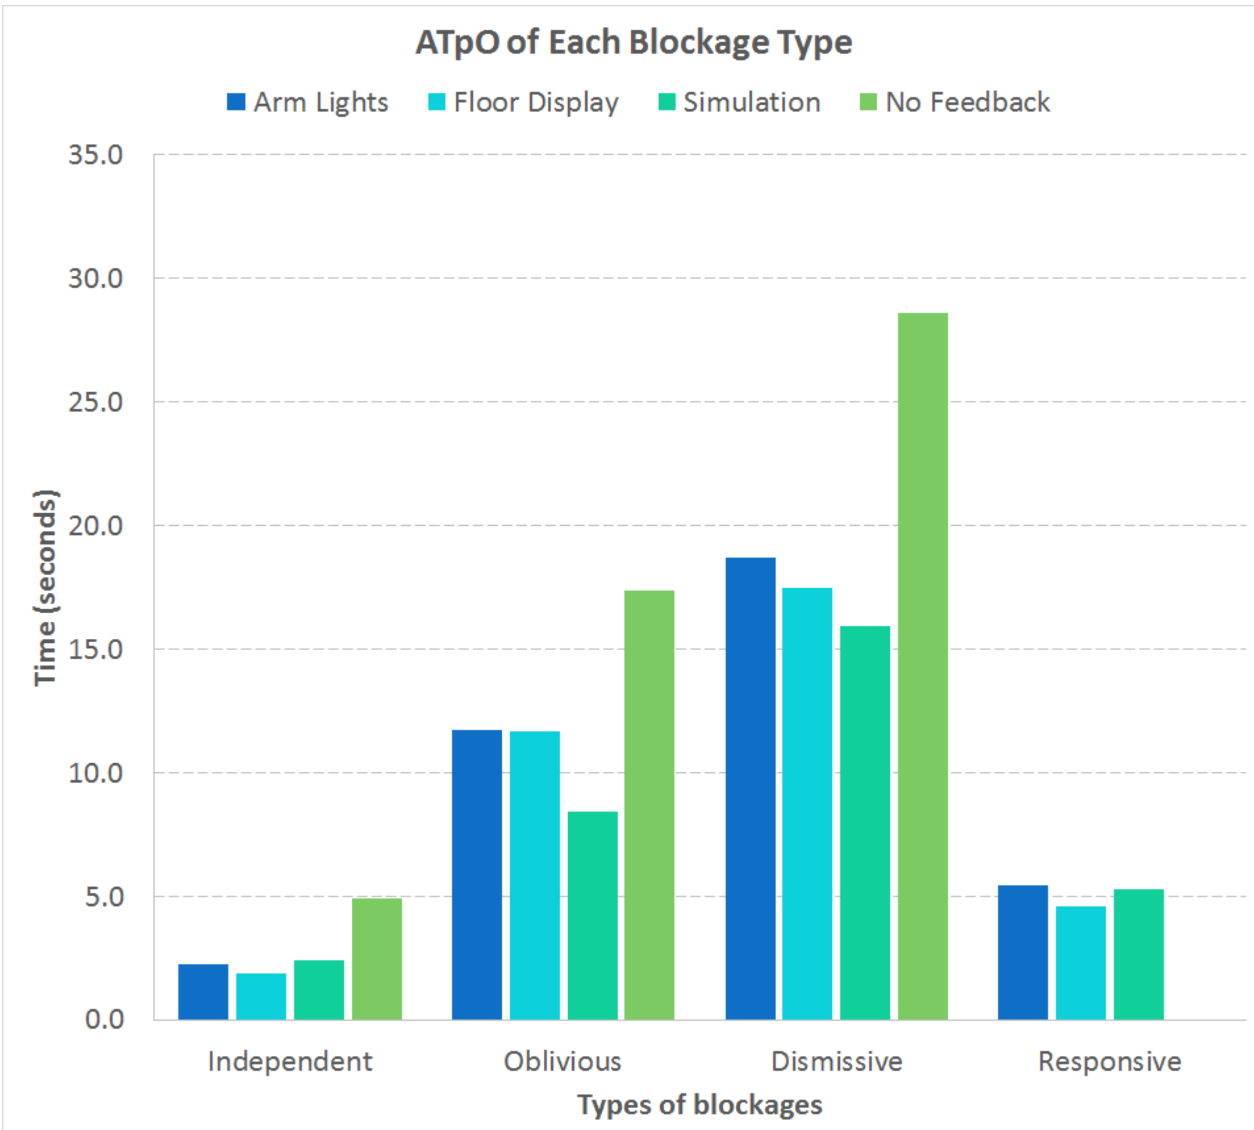

Figure 43: Average time per occurrence for all types of blockages and modalities

An interesting highlight of the ATpO is that the Dismissive cases last on average at least 10 seconds longer than the Oblivious case for every modality. We think this is due to differences in the actions in progress when the blockage is triggered. In Oblivious cases, the blockages often occur when participants are already working. On the other hand, Dismissive cases occur when participants assume the right-of-way and intentionally block the robot in order to complete a task. This difference in behavior appears to make dismissive cases overlap the entirety of a task whereas the Oblivious cases only contain a portion.

The figure shows that Independent blockage types have the shortest average duration out of all reactions. Participants in these cases are anticipating a blockage, without using the feedback, accelerating responses and eliminating reaction time.

Responsive and Independent average times are shorter relative to the other two types because participants performing these kinds of actions do so immediately. We think Oblivious and Dismissive average times are longer due to participants working through the blockages rather than acting directly act to clear them, extending the average time for blockages with these types of reactions.

It should be noted that there are Dismissive Reaction cases for the control group, even though there is no feedback to be dismissed in this condition. These situations occurred when the participant explicitly and attentively watched the robot come to a halt in front of them, and then began or continued to work on the part. Dismissive Reactions in the control condition were on average longer than the average for any of the feedback modalities by about 10 seconds. We think this could be caused by both the rarity of Dismissive cases in control condition and a few cases where control condition participants blocked the robot through multiple fin assemblies, which would artificially inflate the time for a single blockage event.

4.3.14 Average per Modality: Ratios to All Blocks

Figures 44 and 45 show the occurrences and times of each blockage reaction type expressed as a percentage of the total. Combinations of reaction types like “Independent plus Responsive” and “Oblivious plus Dismissive” are used for comparison against the control data, as described in section 4.3.1.1. The combination “Independent plus Oblivious” collects all situations in which the feedback was not noticed for a blockage. The combination “Responsive plus Dismissive” collects all situations where the feedback was noticed.

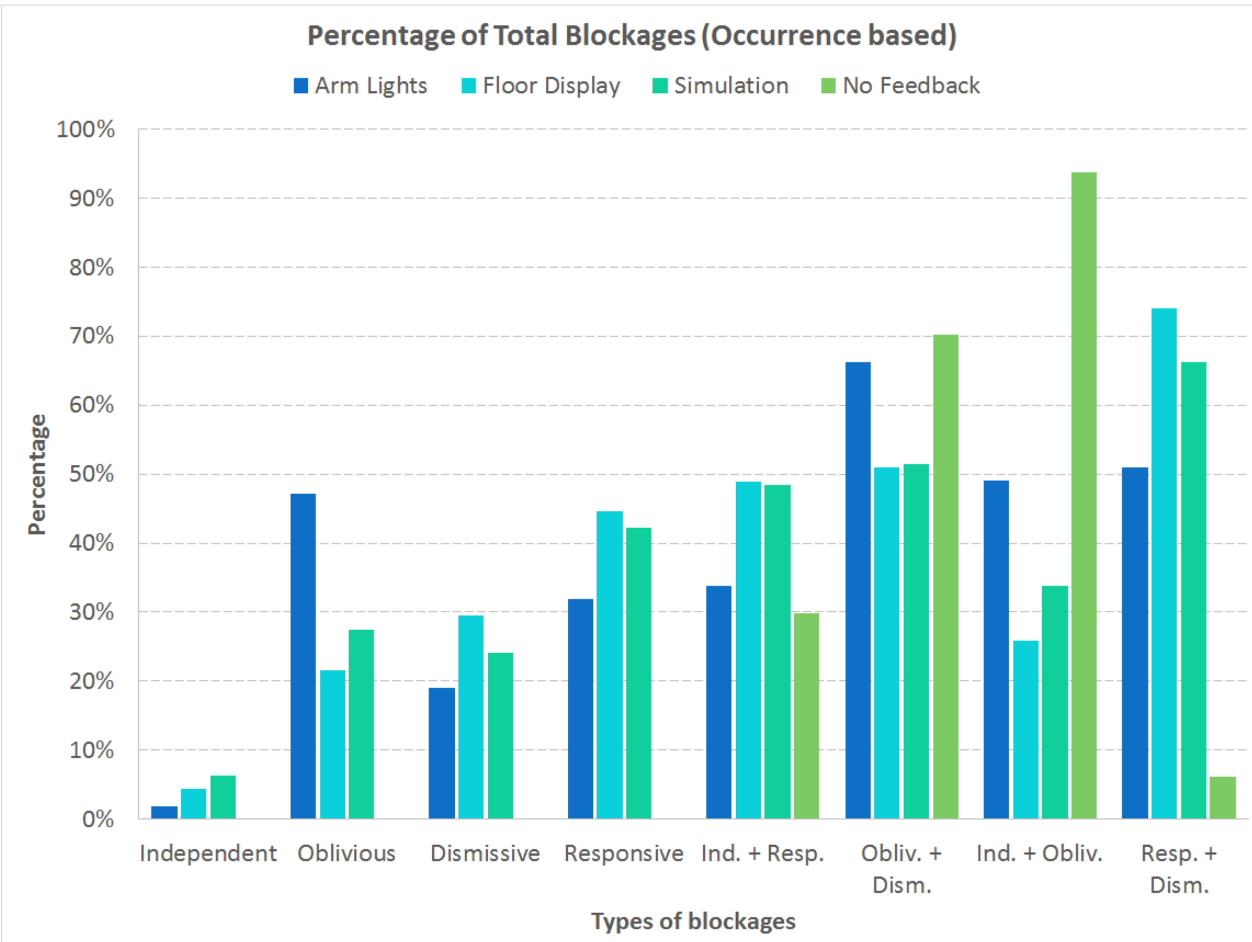

Figure 44: Percentage of blockage type based on occurrence

Comparing feedback against the control condition, we see the control has a lower average proportion of corrective-action reactions (Independent plus Responsive) than any of the feedback modalities, at 3% less than the arm lights and over 15% less than the simulation and floor display. This leads us to the conclusion that any type of feedback increases collaboration, as feedback blockage reactions had an overall average increase in proportion to blockages with corrective action taken.

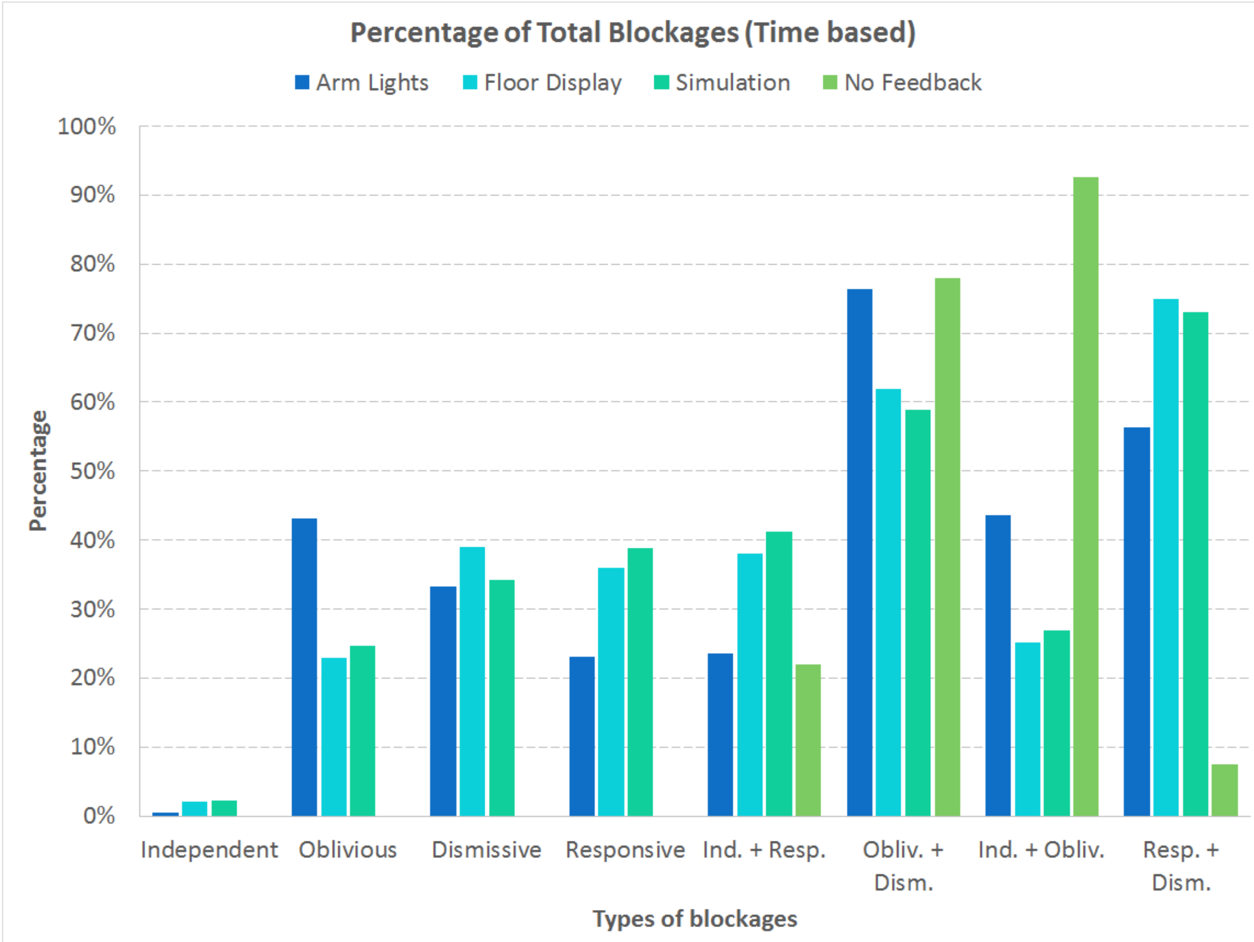

Figure 45: Percentage of blockage type based on time

Of the three feedback methods, the arm light modality has the lowest percentage of Responsive Reactions. However, when participants notice the arm lights, they exhibited similar speed in corrective action compared to the simulation (Section 4.3.9). Coupled with the smaller average total blockages relative to the simulation by over 30% (Section 4.3.3), we came to two possible conclusions. Either the arm lights were effective at communicating feedback (contributing to the overall lower average blockages compared to the other feedback modalities) or the arm lights were not effective at communicating (contributing to the heightened average level of Oblivious Reactions). We theorize that actually both are occurring; the arm lights modality is effective for helping collaboration when observed, but is often missed depending on participant location in the work cell.

4.4 Wait Periods

In addition to blockages, we examined the wait periods for all of the conditions. Waits are periods of time when the robot's next task is not ready, preventing it from continuing work. Wait periods begin when the robot stops moving and ends at the fin task completion.

Figure 46 shows the number of wait periods per participant by modality. Only four of the participants within the control condition had any wait periods, two of which had the maximum number. This means that two of the control participants never successfully completed a fin before the robot required it, which means no collaboration could occur.

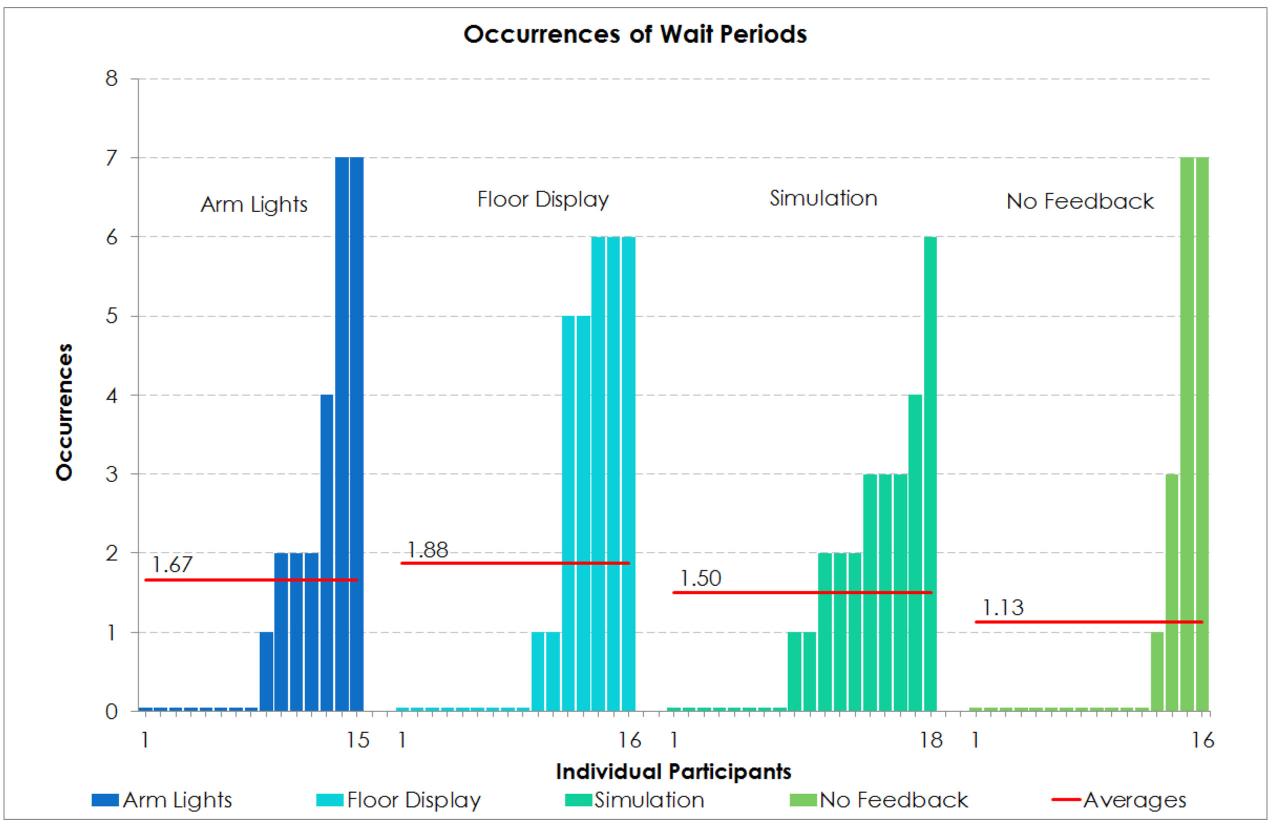

Figure 46: Wait periods, occurrences per participant. Overall averages per modality are in red.

We can see from Figure 47 that participants without feedback experience longer wait periods compared to participants with feedback. The feedback modality participants generally have a greater number of waits (Figure 46), but they are generally of shorter duration. For the total runtime to stay the same, wait time plus blockage time must equal a constant. This means that if wait time goes down, blockage time must increase. We interpret more blockage time to indicate less collaboration, since the robot's operation is being neglected.

We also consider an increase in wait periods and a decreasing time per period an indication of improved cooperation between the participant and robot as this indicates that the participant is working alongside the robot. Increases in average number of wait periods indicates that the participant could be paying close attention to the robot when it is blocked, yielding the right-of-way frequently and choosing instead to make the robot wait if needed. Decreases in average wait periods indicates that the participant could be attempting to work alongside the robot, partially or fully completing tasks ahead of the robot's schedule.

The floor lights modality was most effective at increasing collaboration compared to the control group, since it had the largest increase in average occurrences of wait periods and the largest decrease in average wait period duration. The arm lights were less effective than the floor display for collaboration, since the arm lights had similar numbers of average wait occurrences to the floor display (within (12%), but almost no decrease in average time per occurrence (3 seconds decrease for the arm lights vs 11 second decrease for the floor display). The simulation was also less effective than the floor display for collaboration, since the simulation had fewer average wait occurrences (about 20% fewer), but was effective at decreasing the average time per wait (within 2 seconds of the floor display). These results indicate that participants in the simulation condition did not give the robot right-of-way as often as participants using the floor lights, and that participants in the arm lights condition did not operate alongside the arm as often as participants with the floor lights.

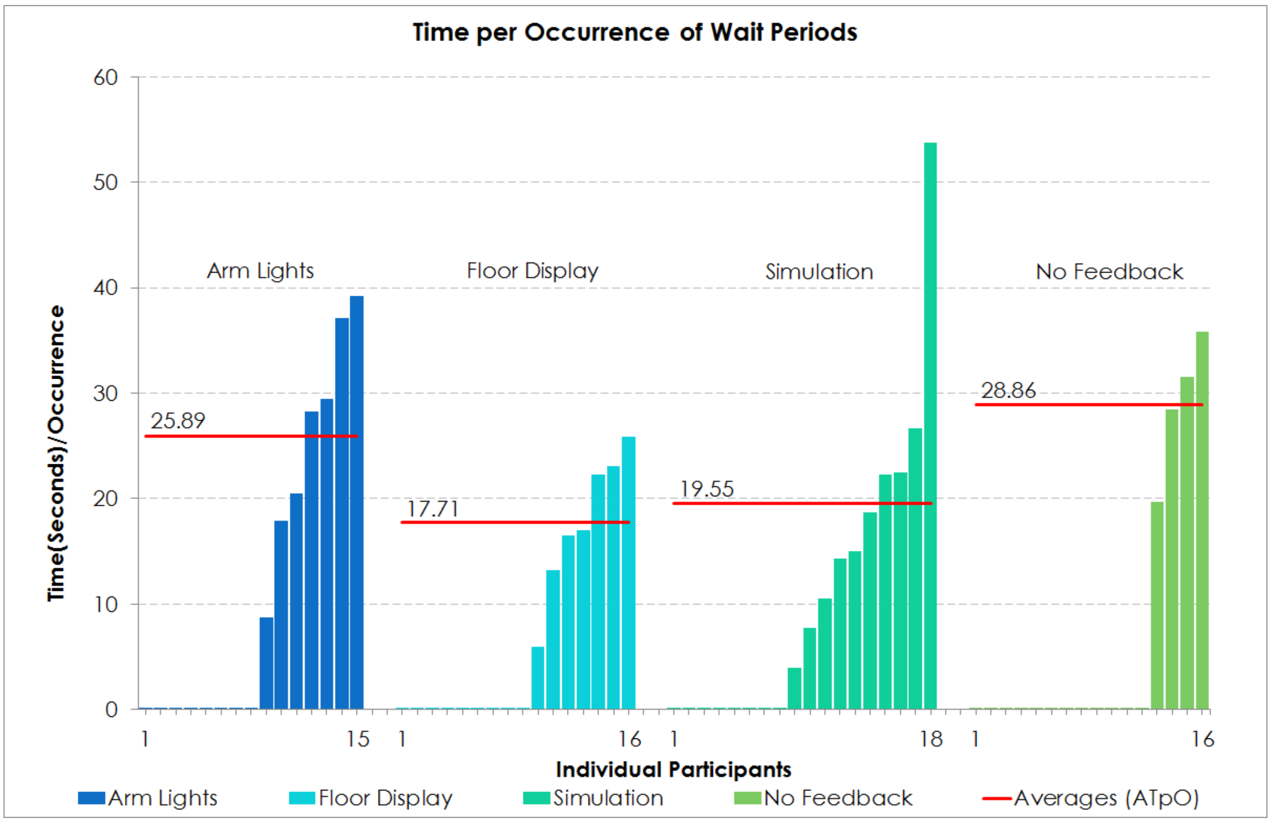

Figure 47: Wait periods, time per occurrence. Red indicates the average by modality

4.5 Summary of Assembly Procedure Data

What follows is a brief summary of our findings comparing human-robot collaboration using the different feedback modalities (Overall comparison table: 2) All three feedback conditions perform better than the control in almost every category. In addition, the embedded floor display has the best performance in almost every category. The embedded arm lights and the non-embedded simulation modalities both show improved performance relative to the control. When compared to each other, the arm lights and simulation modalities approximately match performance, with each performing better in some categories.

Average Per Modality	**Arm Lights**	**Floor Display**	**Simulation**	**Control**
Runtime (s)	429.7	**414.6**	421.5	*450.3*
Number of blockages	**6.3**	7.4	*9.7*	7.4
Total time blocked (s)	78.8	**75.8**	89.9	*102.5*
Response time* (s)	5.4	**4.6**	5.4	
Action Taken (%)	39%	**49%**	48%	*30%*
Feedback Noticed (%)	51%	**74%**	66%	
Number of wait periods	1.67	**1.88**	1.50	*1.13*
Wait period duration (s)	25.9	**17.7**	19.6	*28.9*

Table 2: Comparison of results. Bolded numbers are best performance, italicized are worst. *indicating the time-per-occurrence of responsive-reaction type blockages

4.6 Survey Results

Each participant completed an exit survey in order to acquire subjective appraisals of the robot workspace and feedback system operation. Questions involved participant feelings of comfort and safety around the robot and robots in general, ease of feedback use and level of distraction caused by the feedback modality (See Table 3; an example survey can be seen in the Appendix, section 6.5). Several general conclusions can be drawn from the data.

- The participants felt that all three feedback systems were helpful in task completion.
- They did not feel that the systems were distracting.
- We observed a surprising result for our question "It was hard to determine the robot's future intentions or future movements"; participants with feedback tended to agree more with the statement than those without.
- Our understanding of this result is that the feedback systems heightened participants' awareness of the robot's actions; specifically, participants felt the need to think about the robot's future actions and movements.
- All three feedback modalities increased the participants' sense of connectedness with the robot versus no feedback ("The robot was aware of my presence.").
- All of the feedback modalities helped participants find a path out of a blockage.

	Modality			
Question: How True? Scale of 1-7 (7 being most true)	**Control**	**Screen**	**Arm Lights**	**Floor Lights**
I am anxious around robots in general.	(A) 2.50 (M) 2 (SD) 1.72	(A) 3.11 (M) 2.5 (SD) 1.8	(A) 2.00 (M) 2 (SD) 1	(A) 2.93 (M) 2 (SD) 2.0
The feedback was easy to notice.		(A) 5.39 (M) 6 (SD) 1.6	(A) 5.88 (M) 6 (SD) 1.3	(A) 5.47 (M) 5 (SD) 1,1
When the robot was blocked, the feedback made it easy to find an available position to move into.		(A)5.67 (M) 6 (SD) 1.6	(A) 5.47 (M) 6 (SD) 1.2	(A) 5.13 (M) 5 (SD) 1.6
It was hard to determine the robot's future intentions or future movements.	(A) 2.50 (M) 2.5 (SD) 1.39	(A) 3.56 (M) 3.5 (SD) 1.8	(A) 3.65 (M) 4 (SD) 1.6	(A) 3.47 (M) 3 (SD) 1.5
The feedback system was helpful.		(A) 5.94 (M) 6 (SD) 1.1	(A) 6.29 (M) 6 (SD) 0.7	(A) 5.73 (M) 6 (SD) 1.1
The feedback system was distracting.		(A) 1.89 (M) 2 (SD) 1.0	(A) 1.82 (M) 2 (SD) 1.0	(A) 2.47 (M) 2 (SD) 1,2
The robot was aware of my presence.	(A) 6.00 (M) 6 (SD) 1.5	(A) 6.50 (M) 7 (SD) 0.9	(A) 6.82 (M) 7 (SD) 0.5	(A) 6.40 (M) 6 (SD) 0.6
During the experiment, I felt the robot was safe to work alongside.	(A) 6.56 (M) 7 (SD) 0.7	(A) 6.17 (M) 6 (SD) 0.8	(A) 6.65 (M) 7 (SD) 0.6	(A) 6.20 (M) 6 (SD) 0.9
(A) = Average				
(M) = Median				
(SD) = Standard Deviation				

Table 3: Results from post-survey questions

5 Conclusion

The research questions which provided the basis for this study were: "How can embedded feedback be provided from the robot to the worker in the Intelligent Workcell environment? How can we objectively measure the collaboration? Is embedded feedback effective at improving collaboration or performance?" Results from our experiment have led to the following conclusions:

Conclusion I: Feedback improves collaboration for human-robot assembly tasks.

Human-robot collaboration appears to be improved using embedded feedback. Participants completing the task with the benefit of feedback exhibited a greater average number of wait periods (by at least 30%), shorter average wait periods (by at least 10%) and fewer robot blockages (by at least 12%). We interpret this to mean that participant responsiveness to the robot (our metric of "collaboration" between the robot and worker) is lower overall without feedback; conversely, every form of feedback improves collaboration.

Conclusion II: Compared to the other feedback methods, floor display feedback allowed faster resolution of blockages when detected.

Participants using the floor display cleared Responsive blockage events faster on average (about 1 second or 20%) compared to those using the other feedback modalities. We interpret this to mean that the floor display communicates necessary information in a manner more quickly processed, perhaps more visible, by participants.

Conclusion III: Arm-mounted light feedback helped prevent participants from causing blockages in some situations.

Arm-mounted lights had a lower proportion of feedback-based blockage responses (Responsive plus Dismissive, by at least 18% on average) than other feedback modalities. However, the total non-feedback based blockage events (Independent plus Oblivious) were similar to the screen simulation modality (within 10%). We interpret this to mean that the arm-mounted lights allowed participants to preemptively avoid the robot.

Overall conclusions:

The floor display performed the fastest with an average runtime 35 seconds (7.5%) faster than the control. It also had the highest degree of collaboration, since it had:

- the highest proportion of observation of the feedback (73%) and responsiveness to feedback (45%)

- the largest decrease in average total time blocked (27 second (26%) decrease)
- the largest increase in average wait occurrences (0.75 occurrence (66%) increase)
- the largest decrease in average wait time-per-occurrence (11 second (39%) decrease)

All of the data indicate that the embedded floor display provided the greatest benefit to collaboration of any of the feedback modalities tested, for both robot awareness and for concurrent work.

The simulation feedback provided the second fastest average runtime at 28 seconds (6.2%) faster than the control. This feedback increased collaboration but demonstrated the smallest degree of robot awareness of any feedback, since it had:

- the highest average number of blocks of any condition (9.7 occurrences)
- the smallest decrease in average total blockage time (12.5 second (12%) decrease)
- the smallest increase in average wait occurrences (0.37 occurrence (32%) increase)
- the second largest decrease in average wait time-per-occurrence (9.3 second (32%) decrease)

The increased blockages and high blockage time indicate that the participants were less aware of the robot. However, the large decrease in average wait period indicates that participants were still trying to work concurrently with the robot. Performance may have been slowed by the non-embedded simulation screen because using it required the participants to actively watch it for several seconds; participants in certain parts of the environment sometimes needed to stop their task in order to observe feedback, breaking their task flow.

The embedded arm-mounted lights feedback provided the least improvement in runtime with an average of 20 seconds (4.6%) faster than the control. Regarding collaboration, the arm-mounted lights increased robot awareness but made the smallest impact on concurrent work, since it had:

- the fewest average occurrences of blockages (6.3 occurrences)
- the second largest decrease in average total blockage time (23.5 second (23%) decrease)
- the second largest increase in average wait occurrences (0.37 occurrence (32%) increase)
- the smallest decrease in average wait time-per-occurrence (3 second (10%) decrease)

The fewest average total blockage occurrences indicates that participants were very aware of the robot, as does the increase in average wait occurrences. The minor decrease in wait time-per-occurrence indicates that participants were not successful at working concurrently with the robot. This is supported by the minimal acceleration in runtime, which should improve with concurrent work. Note that the simulation modality and the embedded arm lights showed similar average Responsive blockage reaction durations (5.41 seconds for the arm lights vs 5.37 seconds for the simulation)

We conclude that the embedded floor display provided the highest degree of collaboration with the robot, and performed well in every category. The embedded arm lights feedback improved participant awareness of the robot, but was least effective at encouraging concurrent work. This embedded feedback could be useful in applications where concurrent work carries less importance than alerting the worker of robot status. Modifying the lights to cover the entirety of arm might also help improve concurrent performance. The simulation feedback improved participant concurrent performance, but did not improve robot awareness (actually increased the number of blockages). While the simulation is still useful over no feedback, it would most likely be best suited in an environment where robot awareness is of lesser concern.

Possible sources of error and future directions

Data derived from the user study showed a wide variation in all of the measurements we used for assessment purposes. It became obvious that our participants, although provided with the same initial training regimen, exhibited scattered performance abilities when faced with the assembly procedure in the Intelligent Workcell. We have postulated several possible sources of variation and suggest potential solutions below.

First, our participant pool was not given sufficient time to optimize their performance in the environment. Longer runtimes, multiple attempts at the assembly operation, or extended practice could remove or minimize mistakes from participants.

Second, our participants were not required to have any skills beyond understanding the assembly procedure. We did not assess the basic mechanical ability of any of our participants or define a control skill set. Participant skill at the assembly procedure itself introduced performance variation. We believe adding a restriction in participant selection based on their experience and/or mechanical dexterity could help minimize this performance variation.

Third, the experimental task did not directly mimic existing tasks in manufacturing environments. Our experiment consisted of a simulated procedure, designed to be completed with minimal training and assembly knowledge. Using an experienced participant pool could permit data collection based upon a more difficult task, which might better highlight the utilization and effectiveness of feedback.

Overall, our work using the Intelligent Workcell demonstrates that collaboration with industrial manipulators can be augmented by utilization of feedback. The embedded floor display was the most effective modality used, followed by partial performance gains from both the simulation and the embedded arm-mounted lights. In addition, our task did not reward collaboration with better runtimes. Further work is needed to determine the extent of benefits possible from embedded feedback-enhanced collaboration with standard industrial robots.

6 Appendix

6.1 Training Manual

Rocket Assembly Instructions

1

Assembly Overview

Your goal is to work with the robot to assemble two rockets.

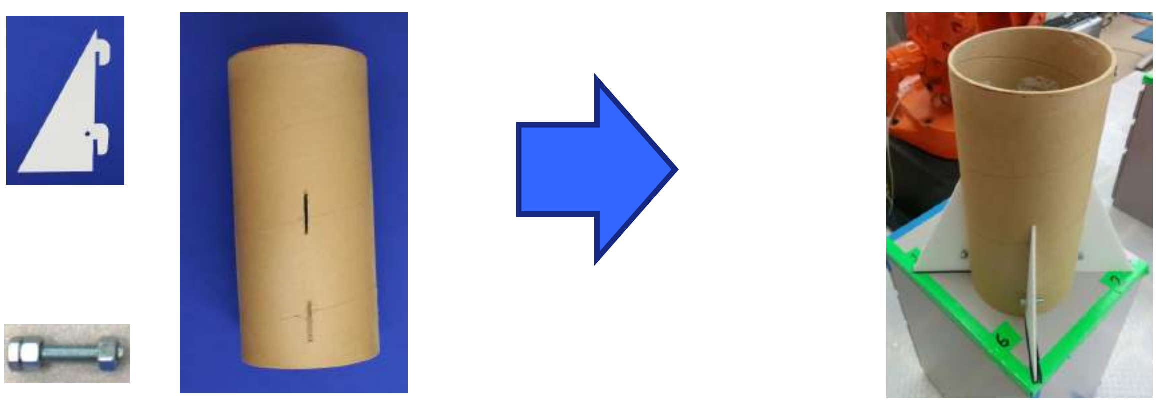

Rocket Parts

Assembled Rocket

Parts List

Fins (8)

Bolts (8)

Cylinders (2)

3

Assembly Environment

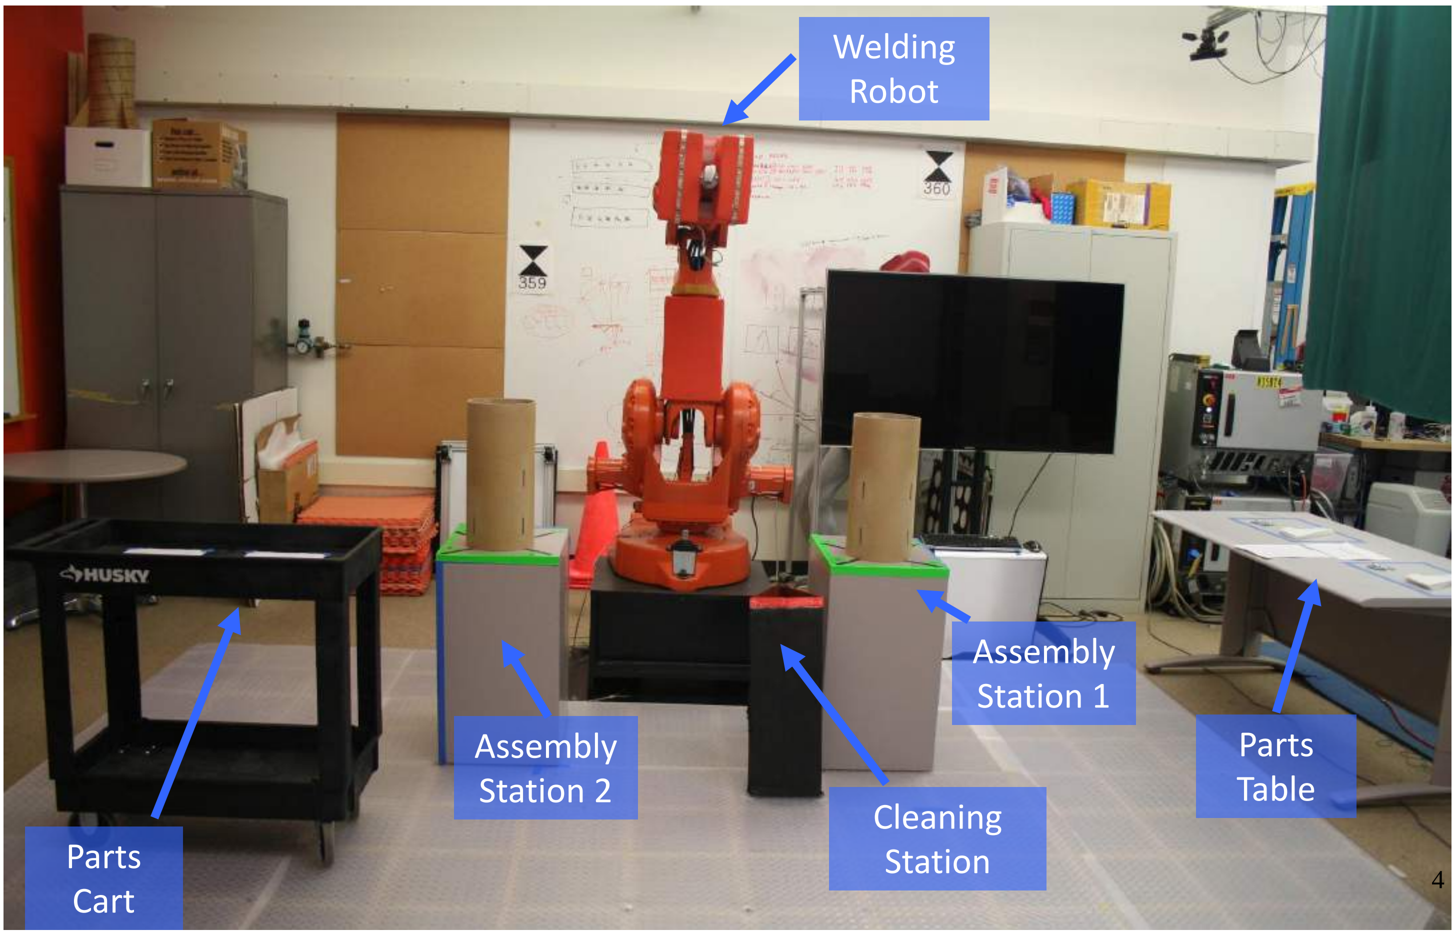

Assembly Procedure

- Add a fin to a cylinder, starting with fin #1.
- The robot will then "weld" the fin.
- Repeat until rockets are complete.

5

Rules and Guidelines

- The parts are “heavy.”
  - You must transport fins/bolts using only the part cart.
  - You can only hold one fin at a time.
- Your task of adding a fin must be complete before the robot can weld. The robot may pause to wait for you to finish.
- The robot welds in the sequential order listed on the assembly stations.
- You do not have to wait for the robot before continuing.
- Try not to block the robot.

6

Parts Table

8 Fins (4 per rocket)
8 Bolts (4 per rocket)

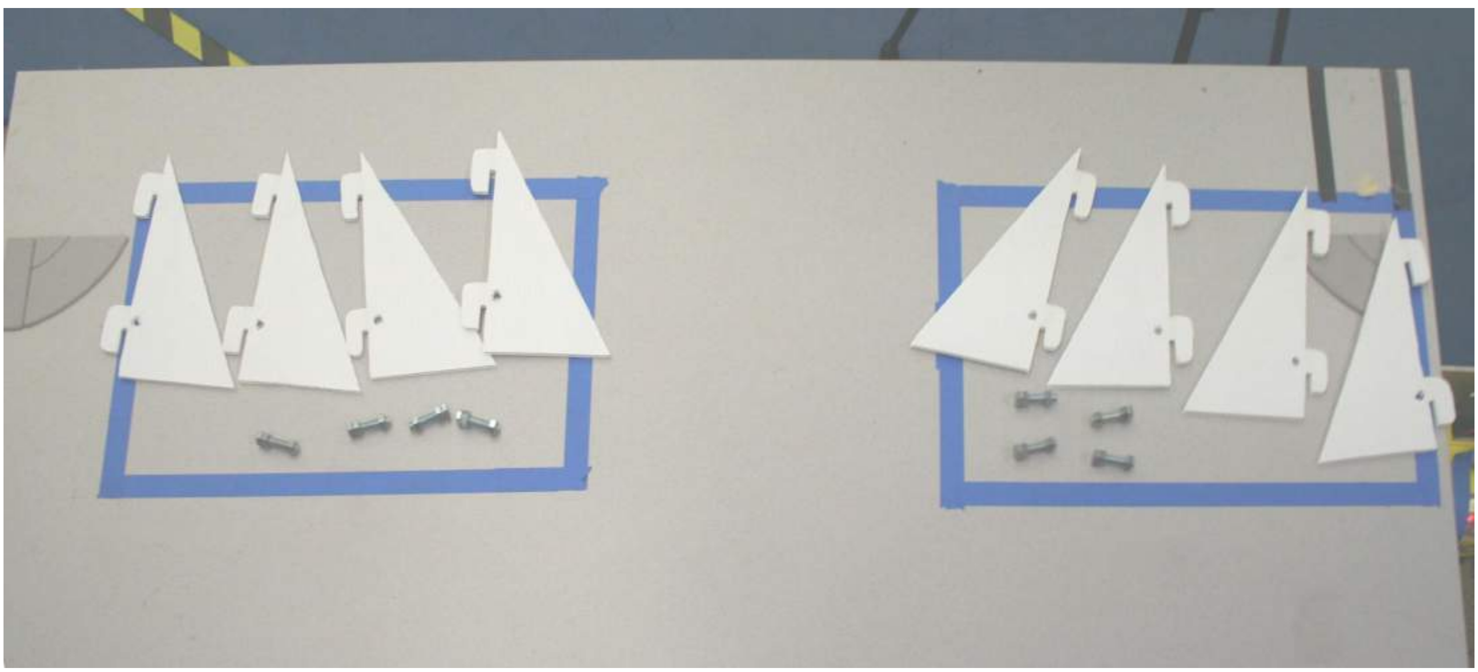

7

Parts Cart

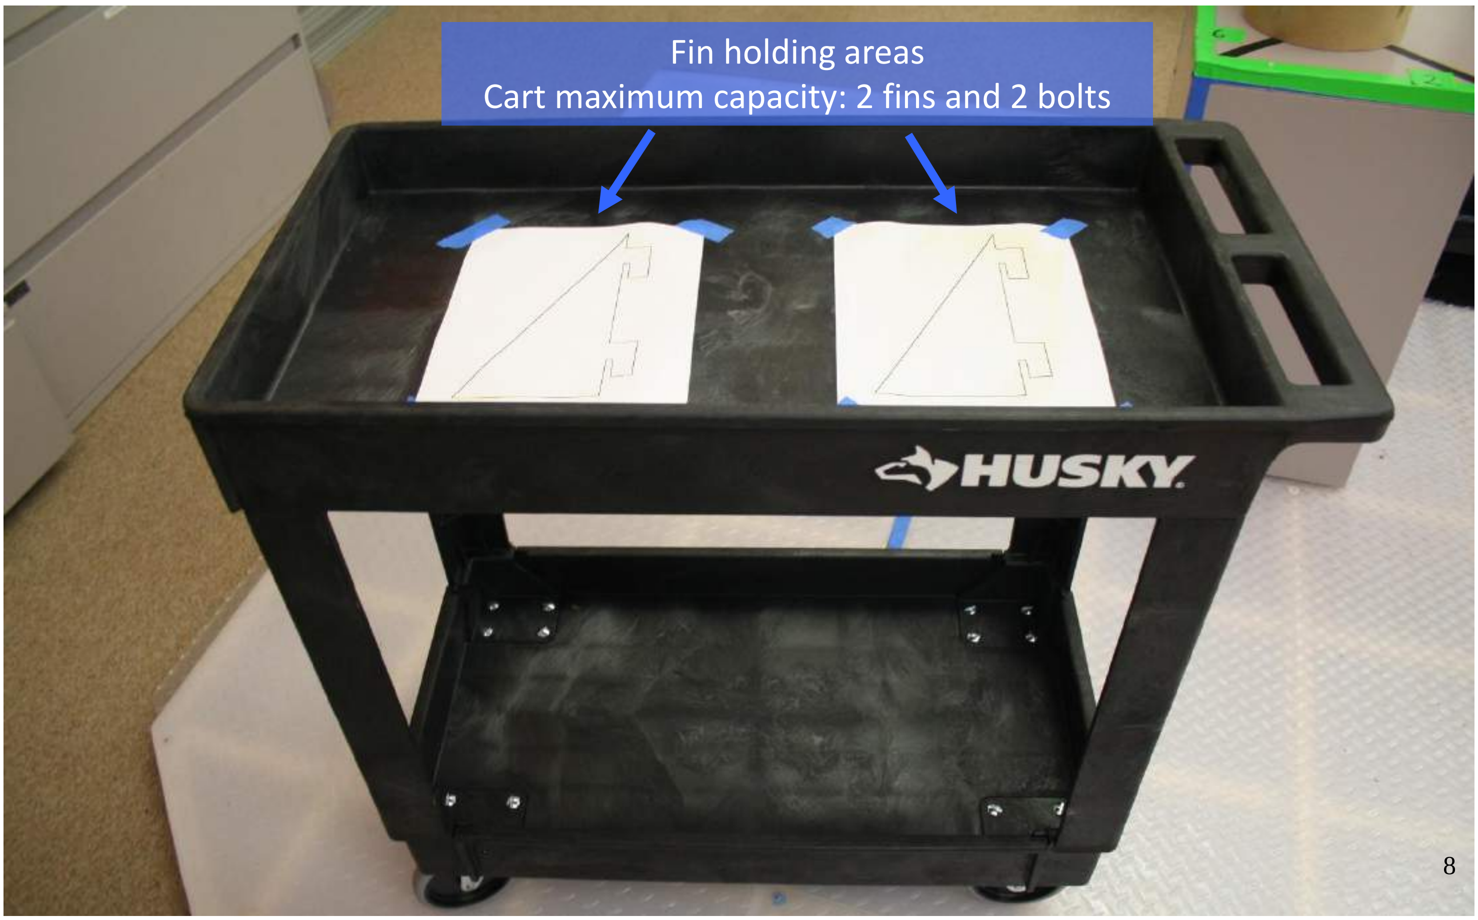
Fin holding areas
Cart maximum capacity: 2 fins and 2 bolts
HUSKY
8

Cart Placement

Keep part cart adjacent to loading/unloading areas.

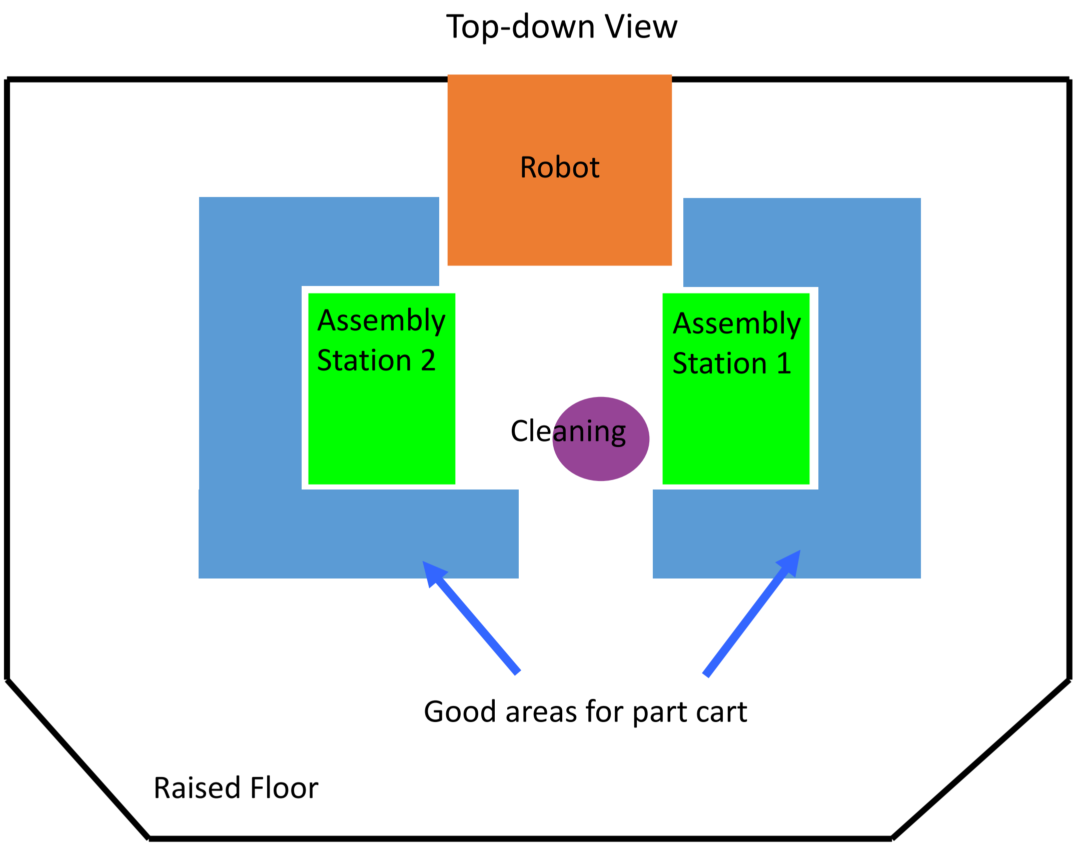

9

Adding Fins (1 of 2)

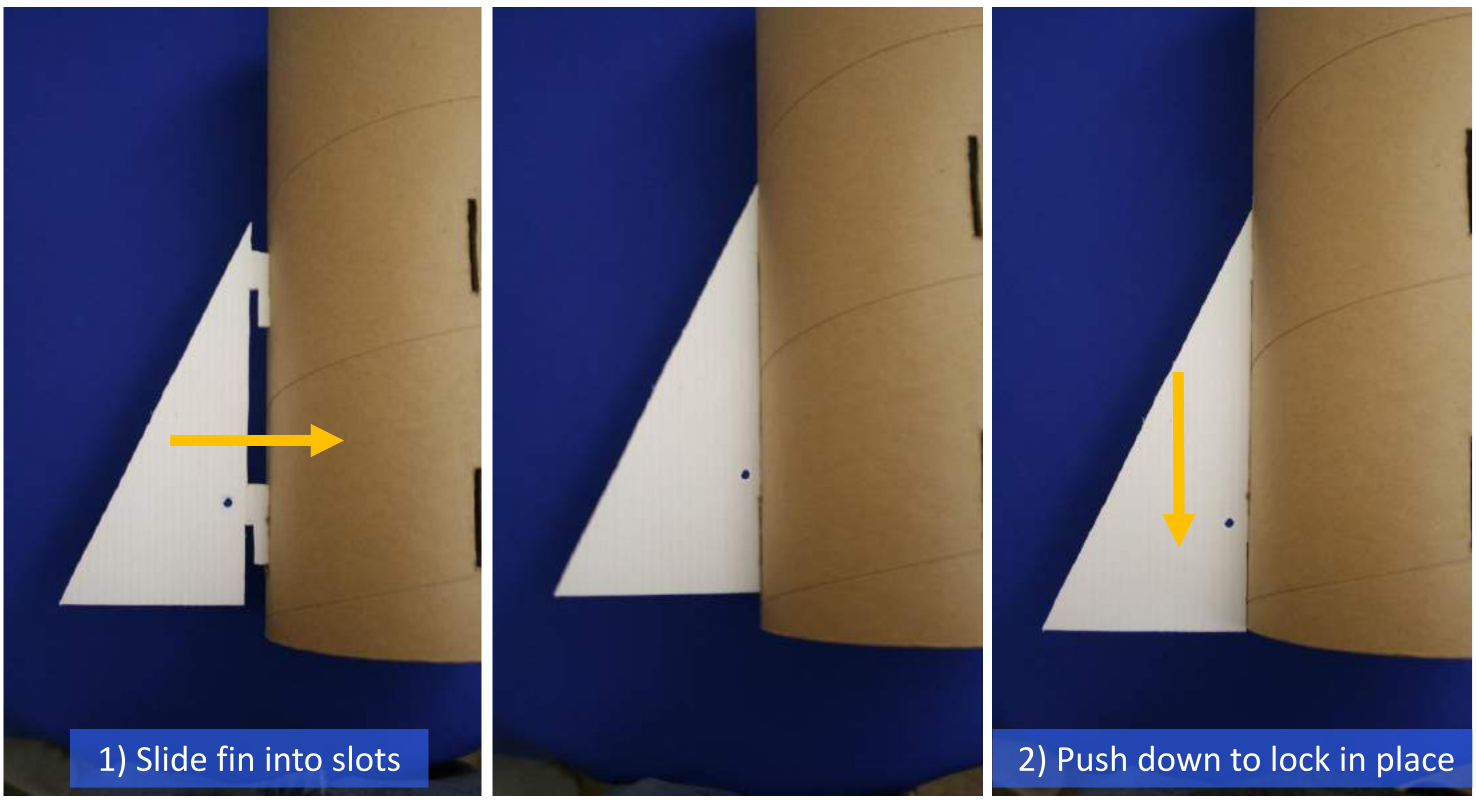

10

Adding Fins (2 of 2)

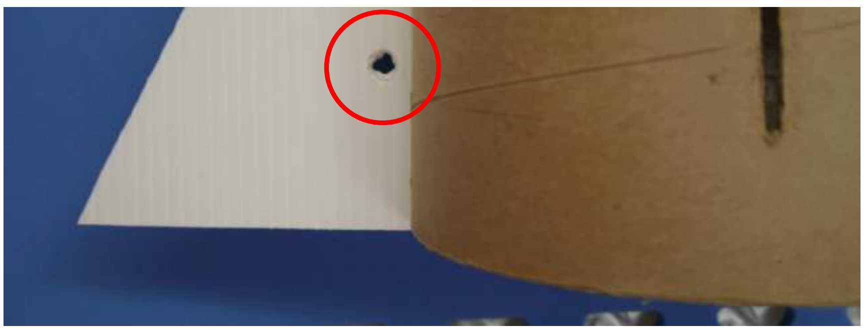

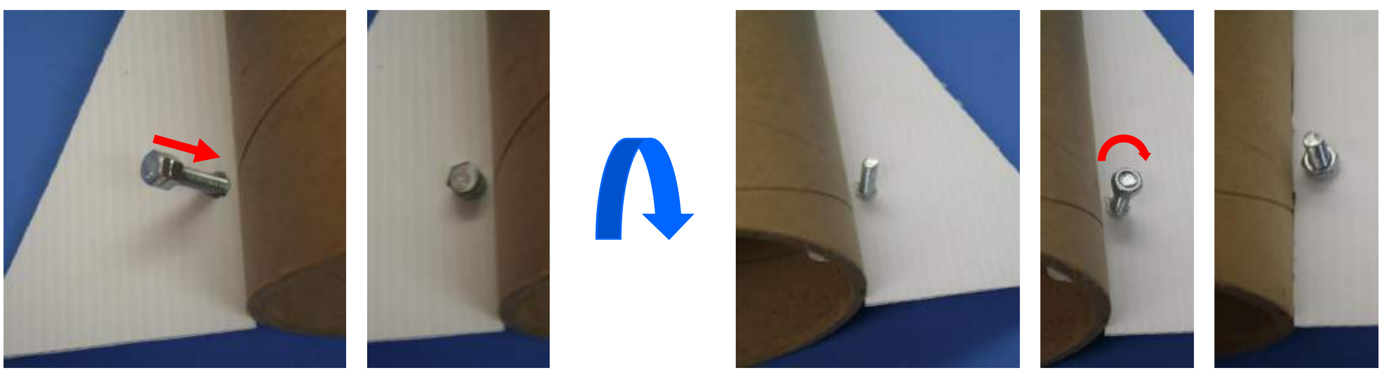

1) Insert bolt into hole on fin.

2) Add nut on other side of fin.

3) Tighten until nut is against fin.

11

Lining Fins Up for Welding

Place the tip of the fin on the end of the black line.
Line up the rest of the fin with the line.

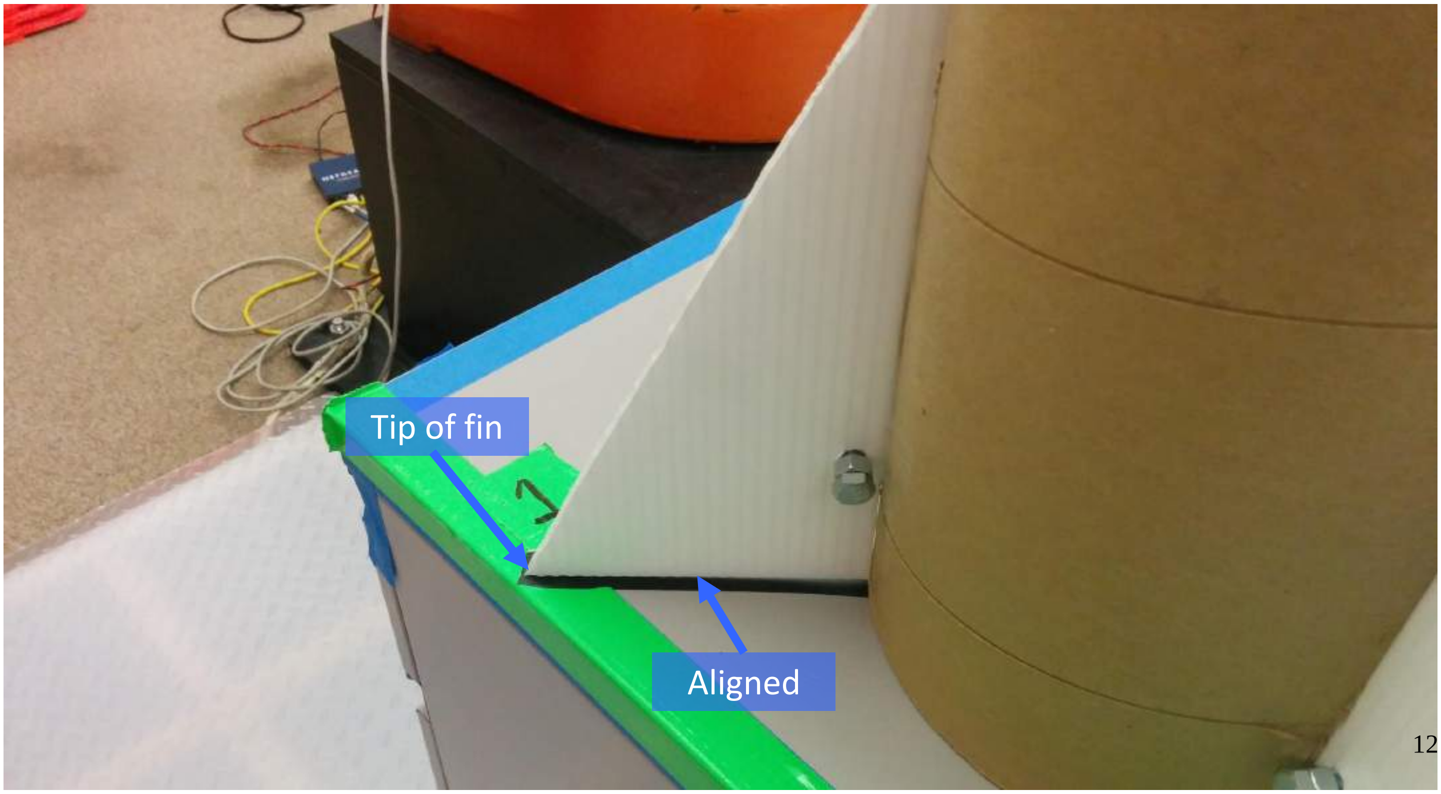

Fin Welding Order

The robot welds the fins in numerical order as shown.

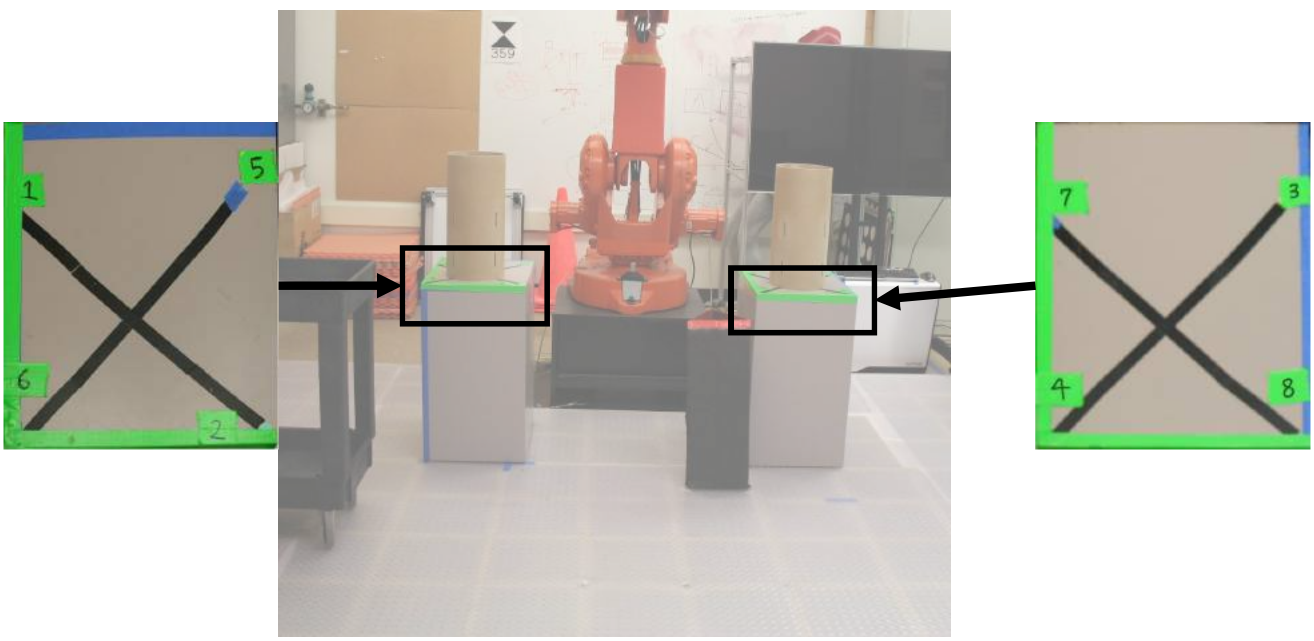

Cleaning Procedure

Sometimes the robot will clean its welding tip at the cleaning station (orange cone).

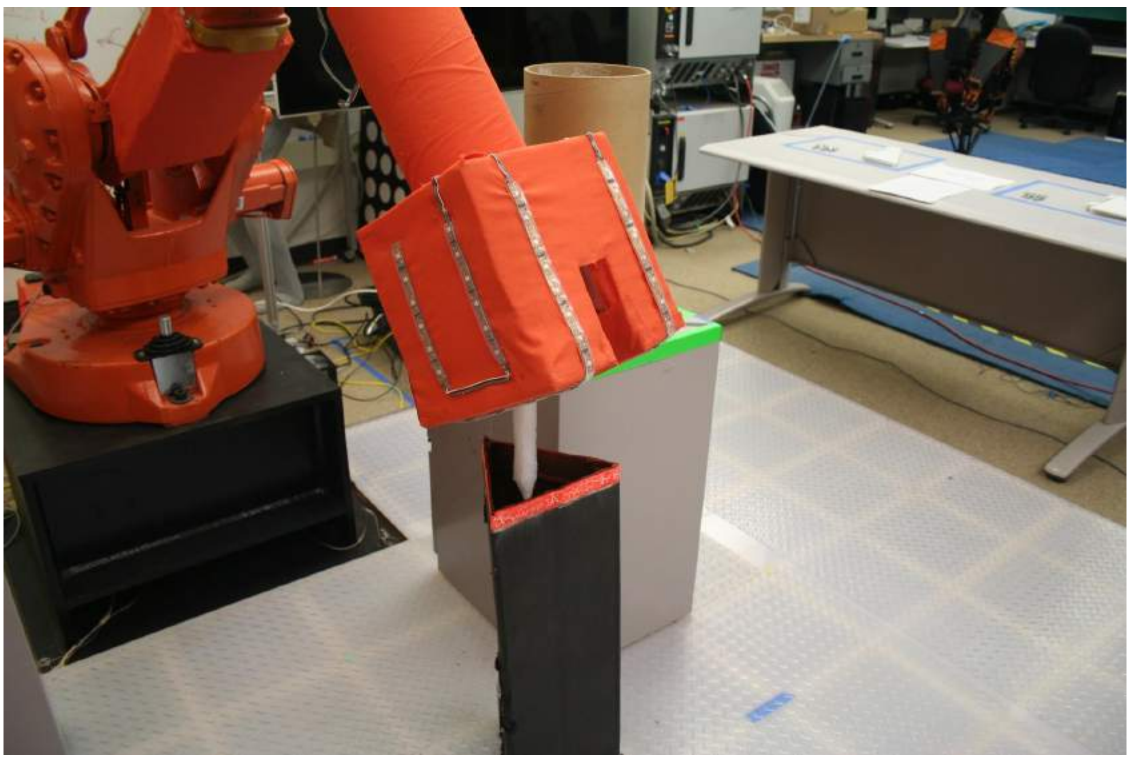

14

Safety Notes

- The system is designed to detect a human's presence. The robot will stop if you get in its way.
- In the unlikely chance that the robot does not detect you, the safety operator will stop the robot.
- The robot is also equipped with padding.

15

6.2 Arm Light Training Manual

Robot Arm Lights

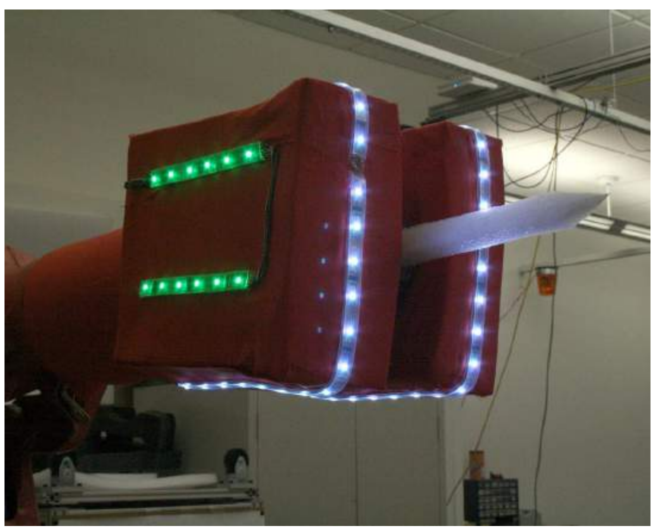

Robot is Moving

Solid lights

Green = moving direction

White = elsewhere

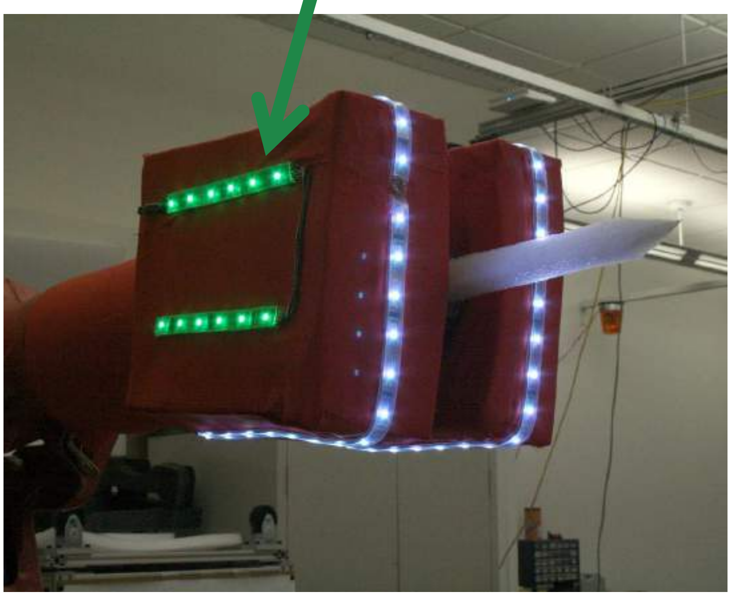

Robot is Blocked

Flashing lights

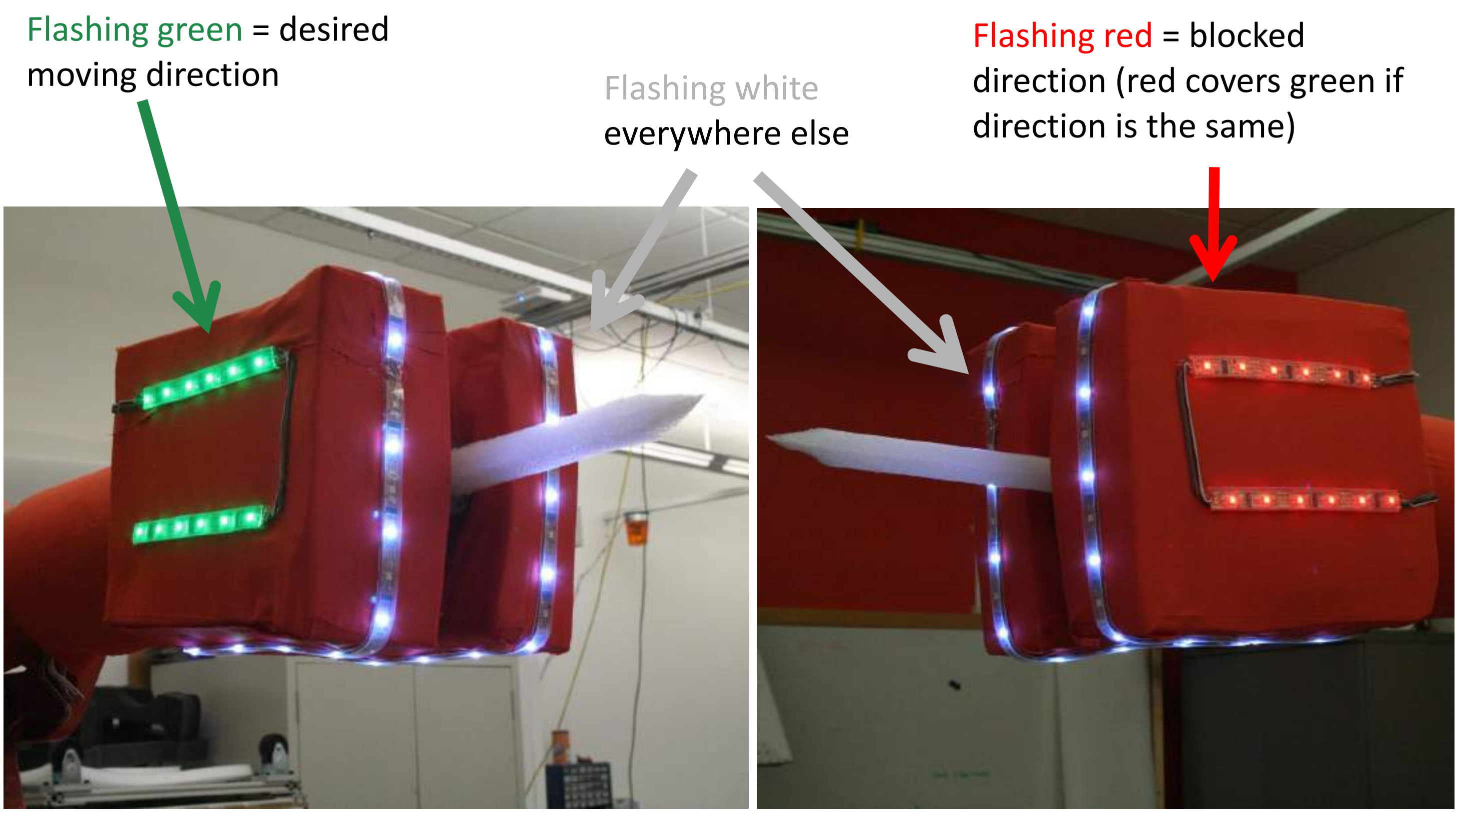

Robot is Paused

Flashing yellow lights

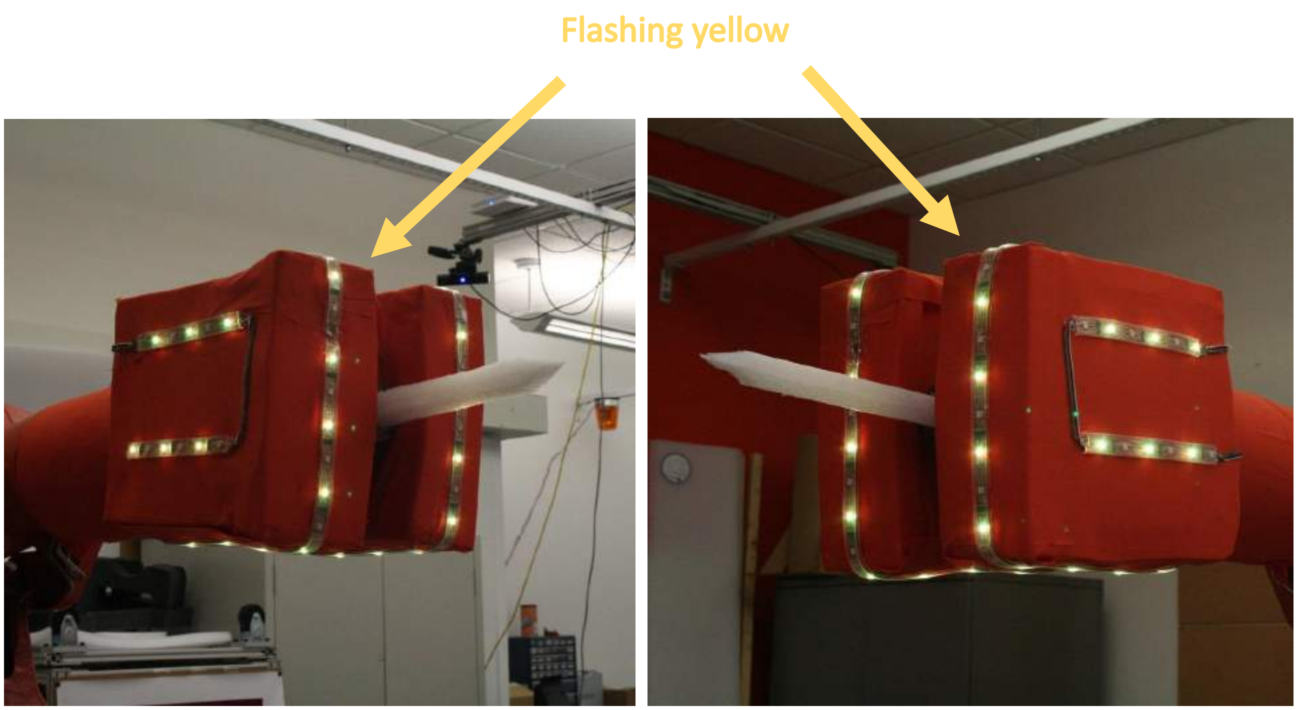

Legend

Green	Moving direction
Red	Blocked direction
Solid lights	Robot is moving
Flashing **red**/white /green	Robot is blocked
Flashing yellow	Robot is paused

6.3 Floor Light Training Manual

Floor Lights

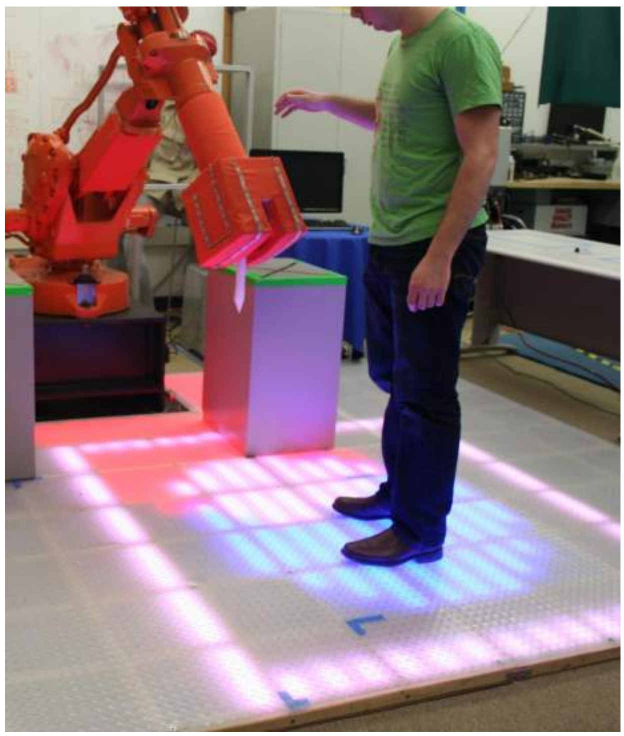

Robot is Moving

Solid lights

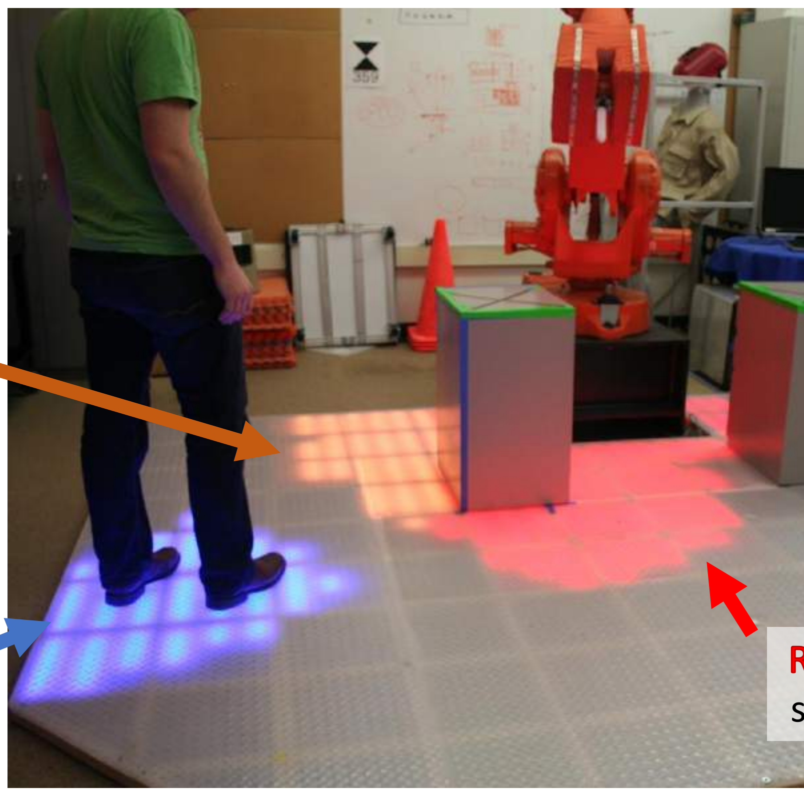

Orange = future position of robot

Blue = person or unknown object

Red = danger zone surrounding robot

Robot is Blocked

Flashing collision area (Red/White) + box

Flashing red/white = collision area

Purple (not shown) = object and robot overlap, but at different heights

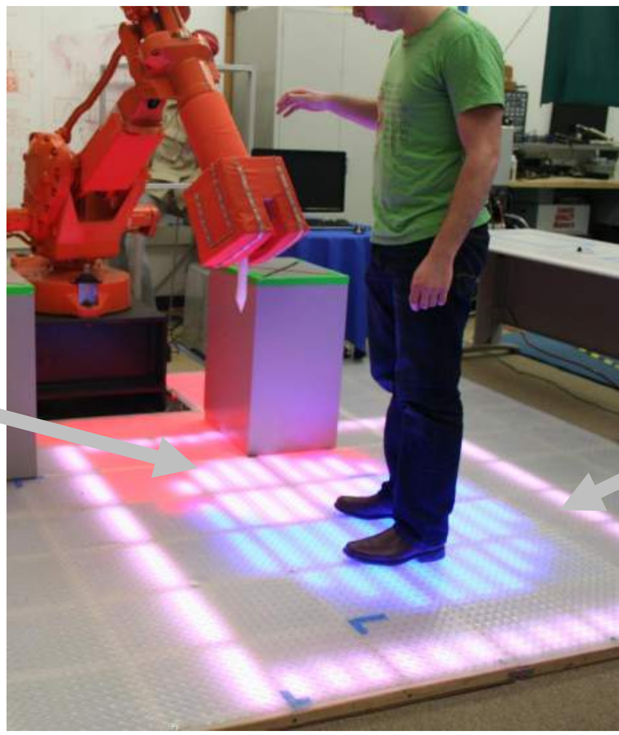

Box indicates position

Robot is Paused

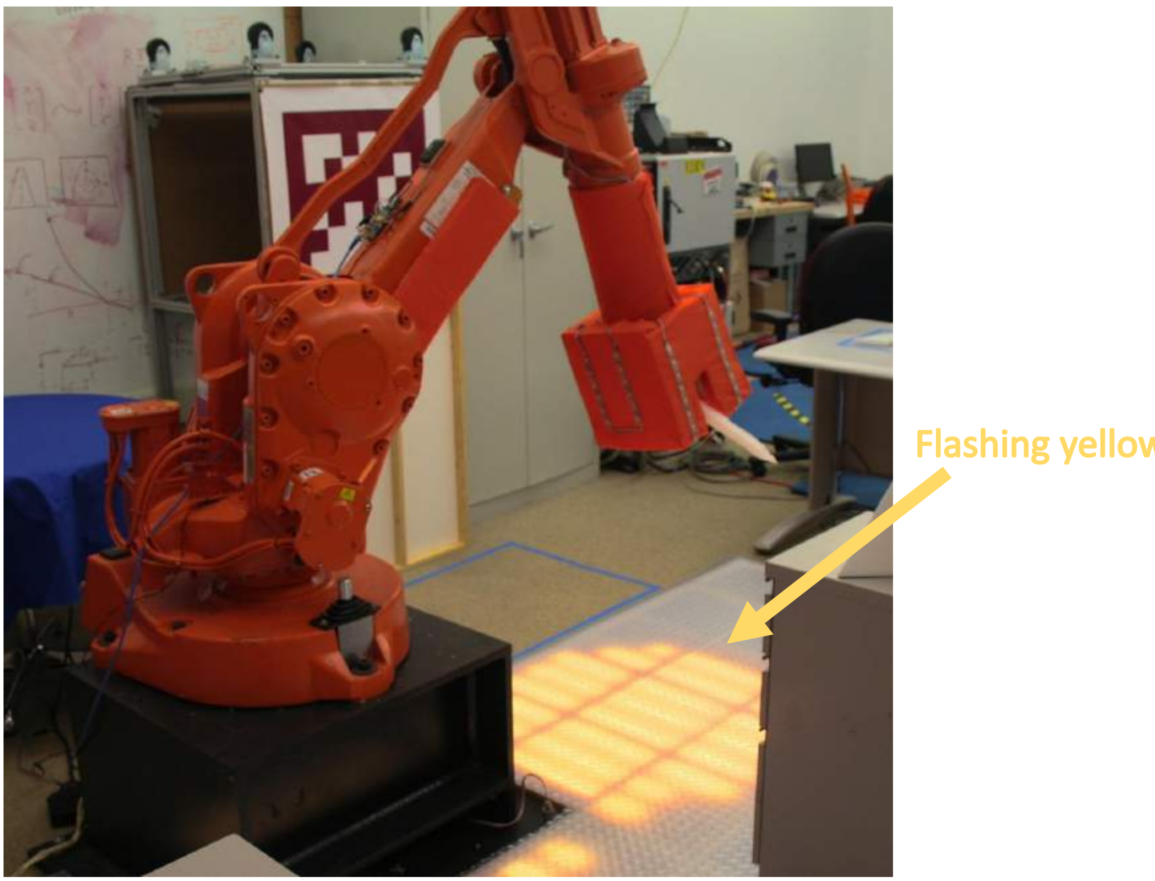

Legend

Blue	Person or unknown object
Red	Danger zone surrounding robot
Orange	Future position of robot
Flashing red/white	Collision area. Box indicates location
Purple	Object and robot overlap, but at different heights
Flashing yellow	Robot is paused

Robot Simulator

Shows planned future motion of robot on monitor

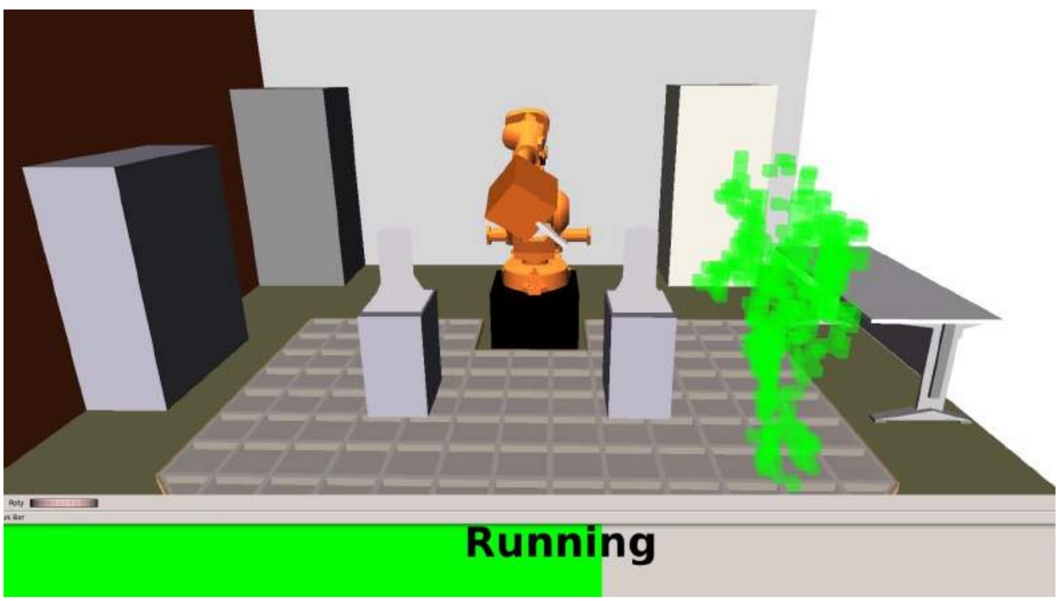

Robot is Moving

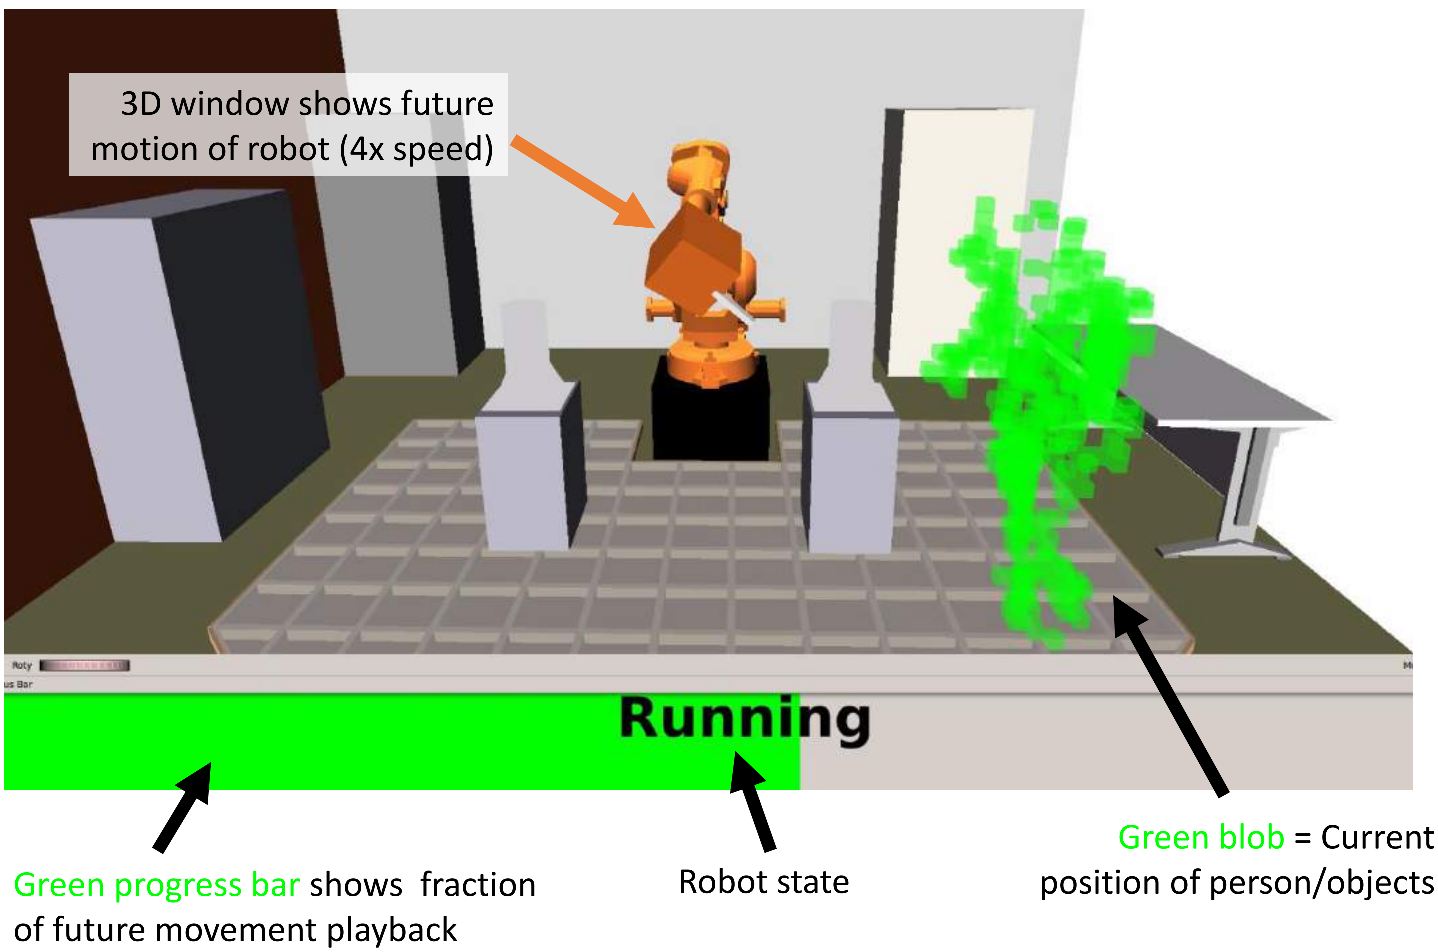

Robot is Blocked

Red blob = collision region

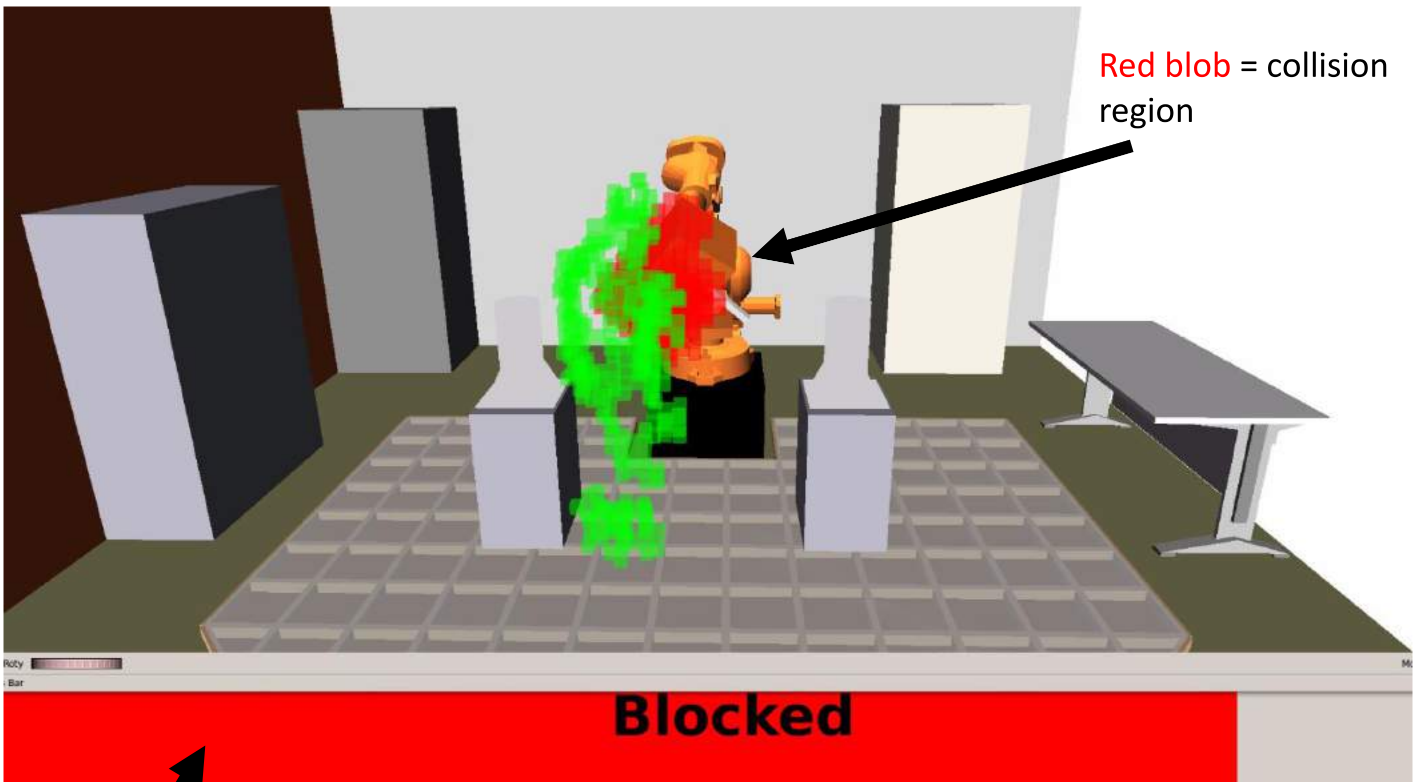

Progress bar turns red

Robot is Paused

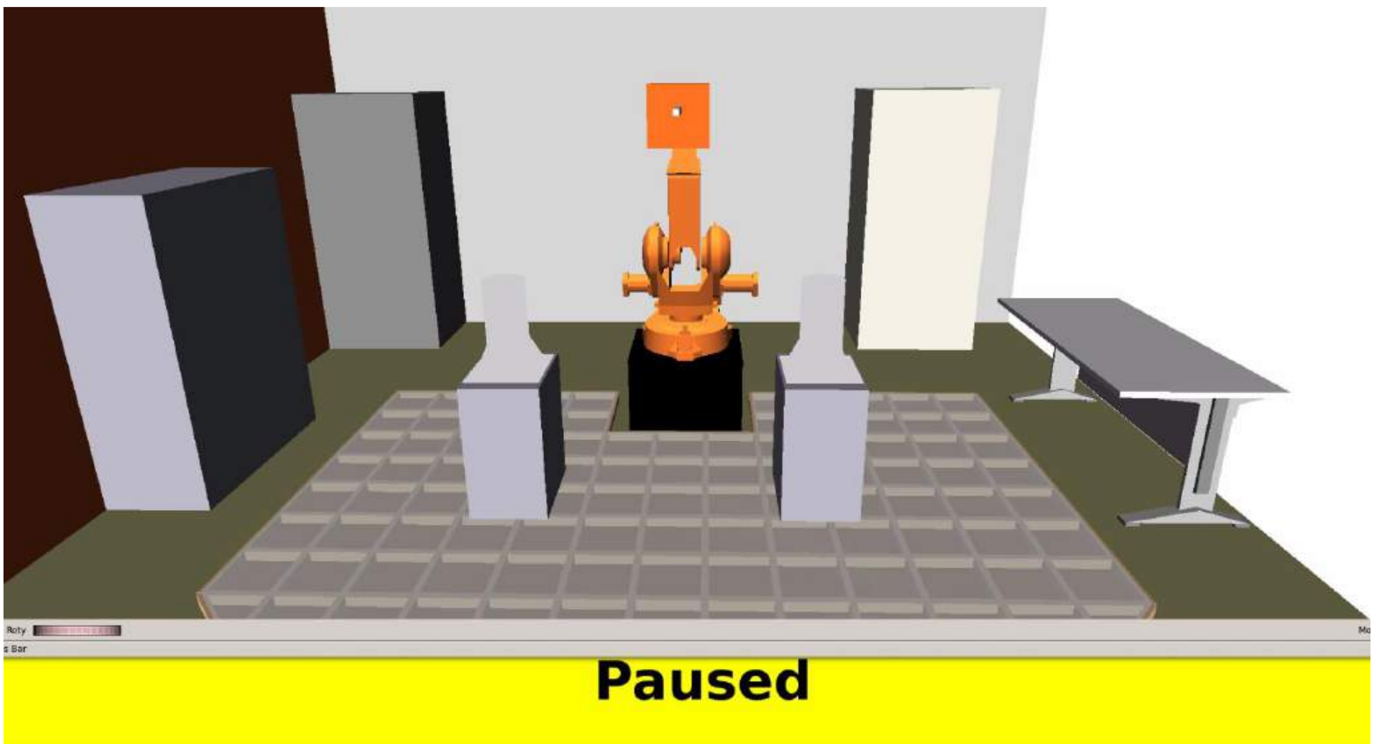

Progress bar turns yellow

Legend

3D animation	Shows future motion of robot
3D green blob	Shows current location of person/objects
Colored progress bar	Shows fraction of future movement playback
Green	Robot is moving
Red	Robot is blocked
Yellow	Robot is paused

6.5 User Study Exit Survey

Intelligent Workcell Feedback User Study: Exit Survey

Participant ID:________________________

Please circle the degree to which you agree with each of the following statements on a seven point scale ranging from 1 (strongly disagree) to 7 (strongly agree). For open-ended questions, please write your answer.

Background:

1) I am in the age range:

(18-25) | (26-30) | (30-35) | (36-40) | (41-45) | (46+)

2) I have experience around robots.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

3) I have a high level of training focused on manufacturing robots.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

4) I am anxious around robots in general.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

Assembly task:

5) The feedback was easy to notice.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

6) When the robot was blocked, the feedback made it easy to find an available position to move into.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

7) It was hard to determine the robot's future intentions or future movements.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

8) The feedback system was helpful.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

9) If possible, please describe what would make the feedback system more helpful: ________________________

____

____

____

____

(continued on back)

10) The feedback system was distracting.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

11) If possible, please describe what would make the feedback system less distracting: ________________________

____

____

____

____

12) The robot was aware of my presence.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

13) During the experiment, I felt the robot was safe to work alongside.

1 (strongly disagree) | 2 | 3 | 4 | 5 | 6 | 7 (strongly agree)

14) If possible, please describe what would make the system feel safer:______________________________________

____

____

____

____

References

[1] International Federation of Robotics. (2016). World Robotics Report 2016, [Online]. Available: https://www.ifr.org/img/uploads/2016-09-29_Press_Release_IFR_World_Robotics_Report_2016_ENGLISH.pdf (visited on 07/04/2017).

[2] International Federation of Robotics. (2016). September 29th 2016 IFR Press Conference, [Online]. Available: https://www.ifr.org/img/uploads/Presentation_market_overviewWorld_Robotics_29_9_2016.pdf (visited on 07/04/2017).

[3] Mark Allinson (Robotics and Automation News). (2017). Automotive industrial robot market forecast to reach $8 billion within four years, [Online]. Available: https://www.roboticsandautomationnews.com/2017/01/18/automotive-industrial-robot-market-forecast-to-reach-8-billion-within-four-years/10664/ (visited on 07/04/2017).

[4] V. V. Unhelkar, H. C. Siu, and J. A. Shah, "Comparative performance of human and mobile robotic assistants in collaborative fetch-and-deliver tasks", in *Proceedings of the 2014 ACM/IEEE International Conference on Human-robot Interaction*, ser. HRI '14, Bielefeld, Germany: ACM, 2014, pp. 82–89, ISBN: 978-1-4503-2658-2. DOI: 10.1145/2559636.2559655. [Online]. Available: http://doi.acm.org/10.1145/2559636.2559655.

[5] "Industrial Robots and Robot Systems – Safety Requirements", American National Standards Institute, ANSI/RIA R15.06-2012.

[6] "Robots and Robotic Devices – Safety Requirements for Industrial Robots", International Organization for Standardization, ISO 10218:2011.

[7] R. Robotics. (2017), [Online]. Available: http://www.rethinkrobotics.com/build-a-bot/ (visited on 07/04/2017).

[8] "Selection and Use of Work Practices", Occupational Safety and Health Administration, OSHA 29 CFR 1910.333.

[9] "Lockout/Tagout", Occupational Safety and Health Administration, OSHA 29 CFR 1910.147.

[10] ABB. (2017). IRB 8700 - Industrial Robots From ABB Robotics, [Online]. Available: https://www.new.abb.com/products/robotics/industrial-robots/irb-8700 (visited on 07/04/2017).

[11] P. Rybski, P. Anderson-Sprecher, D. Huber, C. Niessl, and R. Simmons, "Sensor fusion for human safety in industrial workcells", in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2012, pp. 3612–3619. DOI: 10.1109/IROS.2012.6386034.

[12] P. Anderson-Sprecher, "Intelligent monitoring of assembly operations", Master's thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Jun. 2011.

[13] P. Anderson-Sprecher and R. Simmons, "Voxel-based motion bounding and workspace estimation for robotic manipulators", in *International Conference on Robotics and Automation*, Pittsburgh, PA, Apr. 2012.

[14] G. Tonietti, R. Schiavi, and A. Bicchi, "Design and control of a variable stiffness actuator for safe and fast physical human/robot interaction", in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, Apr. 2005, pp. 526–531. DOI: 10.1109/ROBOT.2005.1570172.

[15] M. Laffranchi, N. G. Tsagarakis, and D. G. Caldwell, "Safe human robot interaction via energy regulation control", in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2009, pp. 35–41. DOI: 10.1109/IROS.2009.5354803.

[16] M. R. Ahmed and I. Kalaykov, "Semi-active compliant robot enabling collision safety for human robot interaction", in *2010 IEEE International Conference on Mechatronics and Automation*, Aug. 2010, pp. 1932–1937. DOI: 10.1109/ICMA.2010.5589111.

[17] M. R. Ahmed and I. Kalaykov, "Static and dynamic collision safety for human robot interaction using magneto-rheological fluid based compliant robot manipulator", in *2010 IEEE International Conference on Robotics and Biomimetics*, Dec. 2010, pp. 370–375. DOI: 10.1109/ROBIO.2010.5723355.

[18] V. Duchaine, N. Lauzier, M. Baril, M. A. Lacasse, and C. Gosselin, "A flexible robot skin for safe physical human robot interaction", in *2009 IEEE International Conference on Robotics and Automation*, May 2009, pp. 3676–3681. DOI: 10.1109/ROBOT.2009.5152595.

[19] J. Mainprice and D. Berenson, "Human-robot collaborative manipulation planning using early prediction of human motion", in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov. 2013, pp. 299–306. DOI: 10.1109/IROS.2013.6696368.

[20] S. Nikolaidis and J. Shah, "Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy", in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Mar. 2013, pp. 33–40. DOI: 10.1109/HRI.2013.6483499.

[21] C. Pérez-D'Arpino and J. A. Shah, "Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification", in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 6175–6182. DOI: 10.1109/ICRA.2015.7140066.

[22] R. Alami, R. Chatila, A. Clodic, S. Fleury, M. Herrb, V. Montreuil, and E. A. Sisbot, "Towards human-aware cognitive robots", in *Proc. The Fifth International Cognitive Robotics Workshop (The AAAI-06 Workshop on Cognitive Robotics)*, 2006.

[23] S. Nikolaidis, P. Lasota, G. Rossano, C. Martinez, T. Fuhlbrigge, and J. Shah, "Human-robot collaboration in manufacturing: Quantitative evaluation of predictable, convergent joint action", in *IEEE ISR 2013*, Oct. 2013, pp. 1–6. DOI: 10.1109/ISR.2013.6695625.

[24] D. Henrich and S. Kuhn, "Modeling intuitive behavior for safe human/robot coexistence cooperation", in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, May 2006, pp. 3929–3934. DOI: 10.1109/ROBOT.2006.1642304.

[25] P. A. S. L.P. (2017). Pilz safetyeye, [Online]. Available: https://www.pilz.com/en-US/products-solutions/sensor-technology/safe-camera-systems (visited on 07/04/2017).

[26] A. de Luca and R. Mattone, "Sensorless robot collision detection and hybrid force/motion control", in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, Apr. 2005, pp. 999–1004. DOI: 10.1109/ROBOT.2005.1570247.

[27] A. D. Luca, A. Albu-Schaffer, S. Haddadin, and G. Hirzinger, "Collision detection and safe reaction with the dlr-iii lightweight manipulator arm", in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2006, pp. 1623–1630. DOI: 10.1109/IROS.2006.282053.

[28] S. Haddadin, A. Albu-Schaffer, A. D. Luca, and G. Hirzinger, "Collision detection and reaction: A contribution to safe physical human-robot interaction", in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sep. 2008, pp. 3356–3363. DOI: 10.1109/IROS.2008.4650764.

[29] M. Geravand, F. Flacco, and A. D. Luca, "Human-robot physical interaction and collaboration using an industrial robot with a closed control architecture", in *2013 IEEE International Conference on Robotics and Automation*, May 2013, pp. 4000–4007. DOI: 10.1109/ICRA.2013.6631141.

[30] B. Lacevic and P. Rocco, "Kinetostatic danger field - a novel safety assessment for human-robot interaction", in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2010, pp. 2169–2174. DOI: 10.1109/IROS.2010.5649124.

[31] D. Kulic and E. A. Croft, "Real-time safety for human - robot interaction", in *ICAR '05. Proceedings., 12th International Conference on Advanced Robotics, 2005.*, Jul. 2005, pp. 719–724. DOI: 10.1109/ICAR.2005.1507488.

[32] M. Indri, S. Trapani, and I. Lazzero, "A general procedure for collision detection between an industrial robot and the environment", in *2015 IEEE 20th Conference on Emerging Technologies Factory Automation (ETFA)*, Sep. 2015, pp. 1–8. DOI: 10.1109/ETFA.2015.7301539.

[33] A. Meguenani, V. Padois, and P. Bidaud, "Control of robots sharing their workspace with humans: An energetic approach to safety", in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 4678–4684. DOI: 10.1109/IROS.2015.7354043.

[34] G. Tonietti, R. Schiavi, and A. Bicchi, "Design and control of a variable stiffness actuator for safe and fast physical human/robot interaction", in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, Apr. 2005, pp. 526–531. DOI: 10.1109/ROBOT.2005.1570172.

[35] D. Ebert, T. Komuro, A. Namiki, and M. Ishikawa, "Safe human-robot-coexistence: Emergency-stop using a high-speed vision-chip", in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Aug. 2005, pp. 2923–2928. DOI: 10.1109/IROS.2005.1545242.

[36] H. Liu, X. Deng, and H. Zha, "A planning method for safe interaction between human arms and robot manipulators", in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Aug. 2005, pp. 2724–2730. DOI: 10.1109/IROS.2005.1545241.

[37] P. A. Lasota, G. F. Rossano, and J. A. Shah, "Toward safe close-proximity human-robot interaction with standard industrial robots", in *2014 IEEE International Conference on Automation Science and Engineering (CASE)*, Aug. 2014, pp. 339–344. DOI: 10.1109/CoASE.2014.6899348.

[38] D. Ebert and D. Henrich, "Safe human-robot-cooperation: Problem analysis, system concept and fast sensor fusion", in *Conference Documentation International Conference on Multisensor Fusion and Integration for Intelligent Systems. MFI 2001 (Cat. No.01TH8590)*, 2001, pp. 239–244. DOI: 10.1109/MFI.2001.1013541.

[39] D. Ebert and D. Henrich, "Safe human-robot-cooperation: Problem analysis, system concept and fast sensor fusion", in *Conference Documentation International Conference on Multisensor Fusion and Integration for Intelligent Systems. MFI 2001 (Cat. No.01TH8590)*, 2001, pp. 239–244. DOI: 10.1109/MFI.2001.1013541.

[40] D. M. Ebert and D. D. Henrich, "Safe human-robot-cooperation: Image-based collision detection for industrial robots", in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, 2002, 1826–1831 vol.2. DOI: 10.1109/IRDS.2002.1044021.

[41] S. Kuhn and D. Henrich, "Fast vision-based minimum distance determination between known and unkown objects", in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2007, pp. 2186–2191. DOI: 10.1109/IROS.2007.4399208.

[42] T. Gecks and D. Henrich, "Human-robot cooperation: Safe pick-and-place operations", in *ROMAN 2005. IEEE International Workshop on Robot and Human Interactive Communication, 2005.*, Aug. 2005, pp. 549–554. DOI: 10.1109/ROMAN.2005.1513837.

[43] M. Fischer and D. Henrich, "Surveillance of robots using multiple colour or depth cameras with distributed processing", in *2009 Third ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC)*, Aug. 2009, pp. 1–8. DOI: 10.1109/ICDSC.2009.5289381.

[44] D. Henrich and T. Gecks, "Multi-camera collision detection between known and unknown objects", in *2008 Second ACM/IEEE International Conference on Distributed Smart Cameras*, Sep. 2008, pp. 1–10. DOI: 10.1109/ICDSC.2008.4635717.

[45] J. Bernshausen, J. Wahrburg, and H. Roth, "Obstacle recognition in the workspace of an industrial robot by a 3d-pmd-camera", in *2009 European Control Conference (ECC)*, Aug. 2009, pp. 2787–2791.

[46] S. Kuhn, T. Gecks, and D. Henrich, "Velocity control for safe robot guidance based on fused vision and force/torque data", in *2006 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, Sep. 2006, pp. 485–492. DOI: 10.1109/MFI.2006.265623.

[47] O. Ogorodnikova, "Creating an active awareness system for humans in robotic workcell", *Acta Polytechnica Hunarica*, vol. 5, no. 2, pp. 11–20, 2008.

[48] O. O. Ogorodnikova, "Safe and reliable human-robot interaction in manufactory, within and beyond the workcell", in *19th International Workshop on Robotics in Alpe-Adria-Danube Region (RAAD 2010)*, Jun. 2010, pp. 65–70. DOI: 10.1109/RAAD.2010.5524607.

[49] T. Fong, J. Scholtz, J. A. Shah, L. Fluckiger, C. Kunz, D. Lees, J. Schreiner, M. Siegel, L. M. Hiatt, I. Nourbakhsh, R. Simmons, B. Antonishek, M. Bugajska, R. Ambrose, R. Burridge, A. Schultz, and J. G. Trafton, "A preliminary study of peer-to-peer human-robot interaction", in *2006 IEEE International Conference on Systems, Man and Cybernetics*, vol. 4, Oct. 2006, pp. 3198–3203. DOI: 10.1109/ICSMC.2006.384609.

[50] J. Shah, J. Wiken, B. Williams, and C. Breazeal, "Improved human-robot team performance using chaski, a human-inspired plan execution system", in *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Mar. 2011, pp. 29–36. DOI: 10.1145/1957656.1957668.

[51] R. Meziane, P. Li, M. J. D. Otis, H. Ezzaidi, and P. Cardou, "Safer hybrid workspace using human-robot interaction while sharing production activities", in *2014 IEEE International Symposium on Robotic and Sensors Environments (ROSE) Proceedings*, Oct. 2014, pp. 37–42. DOI: 10.1109/ROSE.2014.6952980.

[52] G. J. Garcia, J. A. Corrales, J. Pomares, F. A. Candelas, and F. Torres, "Visual servoing path tracking for safe human-robot interaction", in *2009 IEEE International Conference on Mechatronics*, Apr. 2009, pp. 1–6. DOI: 10.1109/ICMECH.2009.4957114.

[53] ABB. (2017). IRB 2400 - Industrial Robots From ABB Robotics, [Online]. Available: https://www.new.abb.com/products/robotics/industrial-robots/irb-2400 (visited on 07/04/2017).

[54] R. Diankov, "Automated construction of robotic manipulation programs", PhD thesis, Carnegie Mellon University, Robotics Institute, Aug. 2010. [Online]. Available: http://www.programmingvision.com/rosen_diankov_thesis.pdf.

[55] "Back Disorders and Injuries", Occupational Safety and Health Administration, OSHA Technical Manual Section VII, Chapter 1.