Characterizing Task-Based Human–Robot Collaboration Safety in Manufacturing

Jeremy A. Marvel, *Member, IEEE*, Joe Falco, and Ilari Marstio

***Abstract*—A new methodology for describing the safety of human–robot collaborations is presented. Taking a task-based perspective, a risk assessment of a collaborative robot system safety can be evaluated offline during the initial design stages. This risk assessment factors in such elements as tooling, the nature and duration of expected contacts, and any amortized transfer of pressures and forces onto a human operator. Risk assessments of example tasks are provided for illustrative purposes.**

***Index Terms*—Collaborative work, human–machine interaction, manufacturing automation, risk analysis, safety.**

I. Introduction

PREVAILING visions of the future of manufacturing depict environments in which robots and humans work amicably to complete collaborative tasks (see [1]). Modern manufacturing practices, however, enforce a strict separation of man and machine due to a disproportionate distribution of power that may lead to significant workplace injuries. A majority of such injuries have historically been the result of an operator making physical contact with the robot when he was not supposed to (see [2]–[4]). In many of these incidents, safety protocols were absent, disabled, or temporarily bypassed. The integration of machines in a human-centric world requires proven safety and a better understanding of the nature and risks of human–robot collaboration.

Planning collaborative tasks requires a juxtaposition of detailed knowledge of the tasks at hand plus an understanding of the risks involved with the task-centric collaboration. Describing tasks at the planning level is an active field of ontological research, and numerous paradigms have been presented to enable automated task planning (see [5] for service robots), restructuring (see [6]), sustainability evaluation (see [7]), and optimization (see [8]). Such efforts rarely include the planning and acknowledgment of safety concerns resulting from the physical interactions between humans and robots. Safety systems should take as inputs the machines and their capabilities, the tooling and fixturing, and the workpieces.

Manuscript received November 18, 2013; revised April 17, 2014; accepted June 27, 2014. Date of publication July 28, 2014; date of current version January 13, 2015. This paper was recommended by Associate Editor K. Hirota.

J. A. Marvel and J. Falco are with the Department of Intelligent Systems Division, National Institute of Standards and Technology, Gaithersburg, MD 20899-8230 USA (e-mail: jeremy.marvel@nist.gov; joseph.falco@nist.gov).

I. Marstio is with VTT Technical Research Centre of Finland, Espoo FI-02044, Finland (e-mail: ilari.marstio@vtt.fi).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSMC.2014.2337275

This paper presents a methodology for the offline evaluation of human–robot collaborative tasks utilizing planning-stage risk assessments and performance characterizations prior to bringing the collaborative application online. The robots under consideration are industrial arms and manipulators. Candidate tasks of manufacturing applications are described in terms of the collaborative nature of the tasks, requisite hardware (including grippers, tooling, and fixtures), motion profiles, and potential hazards. This methodology differs from more traditional risk assessments in that we present an activity-based evaluation of risks rather than an environment-based hazard review. Here, the task decomposition and evaluation focuses on the manufacturing collaborative process. However, the same may also be applied to additional phases of operation, including robot programming, or as part of a risk assessment for accidental contact.

Section II provides a brief overview of human–robot interactions (HRI), including a discussion of the current guidelines for collaborative industrial robotics. Section III establishes a basis for task-based robot safety. Section IV introduces an ontology for collaborative tasks, while Section V details how that ontology is used as part of a risk assessment and abatement. An example case study is presented in Section VI to illustrate the application of the ontology to a collaborative process from task decomposition through risk abatement. Throughout the document we use an illustrative example to demonstrate the application of the ontology and risk minimization methodology.

II. Human–Robot Interaction

The history and applications of HRI are both vast and varied, and have been covered extensively in dedicated reviews and surveys (see [9], [10]). This section serves to introduce the broad concept of human–robot interaction in the context of manufacturing applications.

It has been argued that effective and meaningful HRI in collaborative tasks requires mutual understanding [10]. Some have interpreted this to include a theory of mind in which the robot attempts to model the intent of its human coworker through situational awareness and contextual clues. Such clues include natural language or dialogue (see [11]), gaze or attention inference (see [12]), and biomechanical and biochemical feedback (e.g., anxiety [13]). A substantial portion of collaboration research has focused on the human-centric cues that make effective communication in human-human collaborative tasks possible. It seems only natural that similar mechanisms

TABLE I
INJURY CRITERIA AND BODY MODELS FROM EARLY DRAFTS OF ISO TS 15066 [22]. CLAMPING/SQUEEZING FORCE (CSF), IMPACT FORCE (IMF), AND PRESSING CROSS SECTIONS (PRESSURE/SURFACE PRESSING, PSP) LIMITS ARE PROVIDED FOR SEVERAL REGIONS OF THE BODY (a), WITH THE DISTINCTIONS BETWEEN THE TWO BEING CHARACTERIZED BY DURATION AND MAGNITUDE (b)

Body model Main and individual regions with codification			**Maximum allowable Limit values of the injury severity criteria***		
Main body regions		**Individual body regions**	**CSF [N]**	**IMF [N]**	**PSP [N/cm²]**
Head with neck	1.1	Skull/Forehead	130	175	30
	1.2	Face	65	90	20
	1.3	Neck (sides/neck)	145	190	50
	1.4	Neck (front/larynx)	35	35	10
Trunk	2.1	Back/Shoulders	210	250	70
	2.2	Chest	140	210	45
	2.3	Belly	110	160	35
	2.4	Pelvis	180	250	75
	2.5	Buttocks	210	250	80
Upper extremities	3.1	Upper arm/Elbow joint	150	190	50
	3.2	Lower arm/Hand joint	160	220	50
	3.3	Hand/Finger	135	180	60
Lower extremities	4.1	Thigh/Knee	220	250	80
	4.2	Lower leg	140	170	45
	4.3	Feet/Toes/Joint	125	160	45

* CSF - Clamping/Squeezing force
IMF - Impact force
PSP - Pressure/Surface pressing

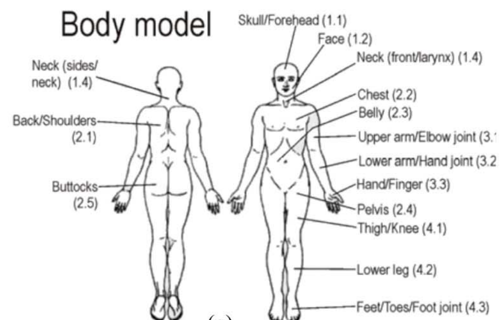

(a)

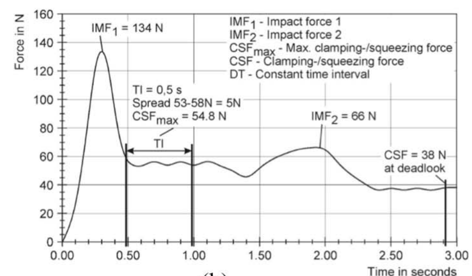

(b)

apply to human–robot collaborations, where the interactions themselves are the foci of the robot systems. An alternative to modeling the complex human psyche (see [14]) focuses instead on a deeper understanding of the nature of each collaborative task (see [15], [16]). Such approaches tend to be narrowly focused on single-purpose results pertaining to goals such as safety, ease of programming, and production throughput.

Four degrees of interaction between a human operator and an industrial robot have been identified for collaboration [17].

1) *Independent:* The human and the robot operate on separate workpieces without collaboration.
2) *Synchronous:* The human and the robot operate on sequential components of the same workpiece.
3) *Simultaneous:* The human and the robot operate on separate tasks on the same workpieces at the same time.
4) *Supportive:* The human and the robot work cooperatively in order to complete the processing of a single workpiece.

Of these four, simultaneous and supportive tasks are expected to have the highest potential for risk of injury resulting from collisions between the robot and the human operator. The actual evaluations of potential hazards are performed during the risk assessment (e.g., as described in [18]).

Pervez and Ryu [19] provide a good overview of the implementations and assessments of the safety of the robot's independent underlying technologies. In more general terms of manufacturing, robotic manipulators are expected to adhere to established international and national robot safety standards. In 2011, the International Organization for Standardization (ISO) revised the language of their robot safety standards [20], [21] to accommodate four new safe collaborative operational modes: safety-rated monitored stop; speed and separation monitoring (SSM); hand-guiding; and power and force limiting (PFL). These four operational modes are described in ISO technical specification (TS) 15066 [22].

The safety of collaborative systems has typically been characterized as a boolean metric: either the robot made contact with an obstacle or it did not. The PFL component of ISO TS 15066 addresses the physical impact between man and machine, and the factors directly relating to the transfer of pressure and force between the two. The current (as of Spring, 2013) metric for PFL is the onset of pain, though previously it was defined by the onset of injury (see Table I). In prior work [23], a generalized means of characterizing the safety of a system was provided that evaluated the robot in terms of mass, speed, and potential severity of impact. Other studies have focused on impact force (see [24], [25]), separation distance (see [26]), velocity and robot configuration (see [27]), and inertia (see [28], [29]). The following sections draw inspiration from all of these approaches, and present a methodology for describing and assessing the risks of human–robot collaborative tasks. The motivating factor is the output of the interaction rather than the interaction itself, and is therefore tuned to the industrial manufacturing problem.

Applications of HRI in manufacturing are ultimately limited by the mechanisms by which operator safety is ensured. Physical barriers separating humans and machines are the de facto means for ensuring operator safety by limiting the potential interactions between man and machine. For interactive HRI, operator safety is ultimately dependent on the technologies used to detect humans in the shared space. For ISO TS 15066, it is assumed that presence-sensing sensors for SSM meet the requirements of [21], which specifies that all electro-sensitive protective equipment is compliant with International Electrotechnical Commission (IEC) standards [30], [31]. Of particular importance is IEC/TS 62046 [31], which enumerates the technologies that are currently suitable as the sole means of protection as being laser detection and ranging devices, light curtains, and pressure-sensitive mats. Passive infrared devices are mentioned, but as they are not standardized are not considered reliable. Acceptable sensing technologies for robot safety are either 1-D or 2-D in terms of their detection zones. Many passive infrared and camera-based systems have 3-D detection zones, but, as mentioned previously, are unacceptable as sole protective devices according to the international standards. Moreover, another major limitation is that these sensing technologies are not specific to humans, but instead only detect intrusions of objects into a protected zone.

III. Task-Based Safety

A point of concern of human–robot interaction involves distinguishing collaborative tasks from noncollaborative tasks. Collaborative tasks such as moving the robot's tool center point (TCP) through direct physical contact (i.e., hand-guiding) necessitate co-location, simultaneous efforts, relaxed or absent physical barriers, and working together to achieve a common goal. The human's actions with the robot extend beyond conventional robot control mechanisms (e.g., jogging the TCP via a teach pendant), but rather serve to interact with the robot as it works on some task.

The risk assessment is a critical component of understanding the potential risks of designing robot cells. The need for a risk assessment applies to both developing new cells, and repurposing an existing cell for a new task where the hazards, required training, or safety controls have changed. The risk assessment identifies the underlying hazards inherent in the equipment or processes. The assessment process is a function of the implied dangers with some assignment of type or severity of injury. Haddadin *et al.* [32] proposed an integration of injury knowledge of impact events into the risk assessment of the physical design of robots, tools, and part components. Distinguishing between the safety of a robot and the safety of a robot system is a major shift in approaching robot safety. The semantic differences are subtle, but are important for limiting liability. The safety of a robot is limited to only the robot and its controller, where the safety of a robot system extends to controllable equipment outside the robot manufacturer, including external hardware such as tooling and stored energy sources.

A missing component of both the safety and new risk assessment models is an understanding of the risks involved with the task itself. All processes, tools, and environments involve an element of risk. Blunt force impacts, lacerations, slips, trips, falls, and exertion injuries from poor ergonomics (e.g., pushing, pulling, bending, and twisting) are inherent in most manufacturing tasks. Understanding the impacts of these risks allows the design of workcells and processes that minimize potential injuries. By decomposing a task to its atomic elements, discrete events and task components can be assessed independently. As an example, Tan *et al.* [33] present a method to determine the individual collaboration task roles of an entire manufacturing process. They propose a hierarchical approach for task decomposition to assign subtasks to humans or robots. In their approach, safety is presented as potential collision hazards, but the safety of the task itself is not considered.

IV. Ontology for Task Classification

We have defined an informal ontology to describe the nature of the interactions of task-oriented collaborations. This ontology does not follow the strict guidelines of, for instance, the ontology web language (OWL, [34]). The reasoning behind this is two-fold. First, integrators and manufacturers do not use formal descriptive languages [e.g., planning domain definition language (PDDL) [35]] to describe task functions when performing their risk assessments. Instead, representations are simple, and tend toward plain English language descriptors (see [36]). Second, because the risk assessment process typically occurs only once (i.e., prior to the initiation of a new activity within a workcell), this ontology needs only to capture the initial state of an anticipated interaction.

In contrast to more traditional implementations, this ontology does not require strict formality in its representations and relationships. This property allows the ontology to be more flexible in its implementation, and enables the integration of its structures and characterizations into a risk matrix (see [18], [37]) as part of a traditional risk assessment.

The risk assessment breaks each task into events (discrete motion steps and actions of a task) as described in Section V. Each event is assessed based on its subjects (the physical components of the task) and predicates (the properties and capabilities of the task subjects). Hazardous steps are identified, and may be isolated or replanned to minimize risk. Isolation focuses on removing the hazard by separating the human and the robot, and may consist either of installing additional safeguards or reconfiguring the task's collaboration such that simultaneous or supportive tasks are independent or synchronous, instead. Replanning intrinsically reduces the hazard through a redesign of the task steps, and may involve additional or reconfigured tooling, steps, or controls.

For our purposes, we define a robot workcell as the space within which a manufacturing task is performed by one or more robots. The workcell includes all of the tools, personnel, and materials needed to complete a given task. In traditional industrial robotics, the confines of a workcell are physical barriers that strictly delineate the regions in which the robots and human operators could work. As modern installations evolve, however, these barriers become increasingly ethereal.

We now revisit the workcell definition and describe it in terms of components summarized in Table II. A collaborative workcell consists of two or more agents. These agents may be either robot systems or humans. Robot systems are comprised of one or more robots, their attached tooling, a base, and any additional support equipment. Depending on the nature of the collaboration, there may exist zero or more humans in the collaborative robot workspace. From a system perspective, these humans have associated attributes such as names and roles,

TABLE II
SUBJECTS (PARTS, STRUCTURES, EQUIPMENT, AND PEOPLE) ASSOCIATED WITH A GIVEN INTERACTION WITHIN A WORKCELL

WORKCELL COMPONENT	**ELEMENTS**	***NOTES***
Robot system	Robot(s)	
	Tooling (attached)	
	Base	*(e.g., pedestals, linear rails, or mobile platforms)*
	Support equipment	*Hardware physically connected to the robot (e.g., hoses, holsters, or clamps)*
Human(s)	Name/title/role	*(e.g., operator, pedestrian, technician, or programmer)*
	Carried equipment	*(e.g., tooling, protective equipment, work components, trash, or food)*
Fixturing	Structures	*(e.g., scaffolding or tables)*
	Stations	*(e.g., tool changers, inspection stations, loading stations, or operator stations within workcell)*
Workpiece	Principal part	*Primary component to which tasks are applied or subcomponents are added*
	Subcomponents	*Parts that will be added to the primary component of the workpiece*
Machine tools	Tooling	*Non-portable tools such as lathes and belt sanders*
	Support equipment	*(e.g., cabling and hoses, vacuum systems, and permanently-mounted lighting)*
Safeguards	Physical barrier	*Physical guards to prevent access or entry to parts of the workcell*
	Sensing device	*Sensors that detect the occurrence of events such as entry into the workcell (e.g., laser gates) or the presence of hazardous conditions*

as well as any tools or equipment they bring with them into the workcell. Also included are static fixtures (e.g., scaffolds or tables) and task stations (i.e., locations within the workspace dedicated to specific tasks). Workpieces(s) are expected to be within the workcell while the robot is performing its task, and consist of the principal component and any subcomponents. Also included in the workcell are any tools associated with the task and dedicated safeguards (physical barriers and sensors) intended to maintain operator safety.

This definition of the collaborative workcell does not necessitate the presence of robots. This is due to the nature of evaluating only the task-based safety of a given process. Moreover, it reflects the expected evolution of human-scale automation [38] in which manufacturing features tasks in which a human or a robot may be employed without requiring changes in tooling, processes, or safeguards.

A. Degree of Collaboration

The risks associated with the collaborative task are evaluated according to the predicates and values given in Table III. The description of the collaborative task begins with the nature of the collaboration, as indicated by one of the four degrees of collaboration from [17], previously discussed in Section II: independent, synchronous, simultaneous, or supportive. These four degrees of collaboration establish the base nature of the collaboration, and set the stage for further factors that will need to be taken into consideration for evaluating the task-based safety. Collaborative tasks require some level of colocation, so even the independent scenarios have a degree of associated risk. The distinguishing characteristic between them the nature of that interaction.

Depending on the nature of the collaboration, minor process errors can impact the collaborative tasks in different ways. For instance, if the human operator works too slowly or makes mistakes, the timing and successful completion of the robot's efforts may be impacted. Similarly, if the robot's timing is off or performs outside of its tolerance specifications, a human operator or another robot may be required to compensate. With independent, synchronous, and simultaneous collaborations, these impacts can be measured in terms of time spent correcting or compensating for errors rather than productively working (which might also impact production quantity goals). The effects of timing asynchronies on supportive collaborations are more difficult to quantify. Error compensation occurs inline with the task actions and is thus difficult to separate. Moreover, timing and process errors may also introduce new hazards during supportive collaborations. This places an increased burden on process engineers and robot programmers to add additional safeguards and intelligent robustness against all foreseeable impacts on both safety and process quality. Enabling robot agility by increasing machine intelligence to automatically recognize and compensate for process errors is anticipated to play an increasing role as their integration into human-occupied environments grows. However, increased faith in such autonomous solutions also requires new evaluative test methods to simultaneously verify safe functionality.

B. Tooling

At the heart of any robotic manufacturing task are the tools used to complete work, which may be described in a myriad of ways. For instance, tools may be described either by their features (see [39]) or their functionality (see [40]). To maintain a safe working environment, both should be specified as part of the risk assessment. Tool features describe the physical characteristics of the tools, and include whether said tool is rigid, flexible, or articulated; blunt or edged; and powered or unpowered. Tool features are described in terms of their compliance (i.e., stiff, elastic, plastic, or articulated) and power (powered or unpowered). A tool's function is described in terms of the physical interaction and intended use. Interactions are classified as direct contact or noncontact, and are further described in terms of how they interact (i.e., push, pull, lift, grasp, or disperse). Tool uses describe the intended purpose of the interaction, and can be classified as adding (i.e., adding material to the surface of a component), connecting, moving, removing, inspecting, or heating.

Features and functions are combined to describe a tool during a collaborative task. For instance, a metal inert gas welding tool would be described as stiff, powered electric, contact

TABLE III
PREDICATES DESCRIBING THE CAPABILITIES AND COMMON HAZARDS OF THE TASK SUBJECTS

	PREDICATE	**VALUE**	**SUB-VALUE**	**UNITS**	**COMMON HAZARDS (*NOTES*)**
	Collaboration	Independent			*(Used to profile hazard)*
		Synchronous			
		Simultaneous			
		Supportive			
Tooling	Tool Compliance	Stiff			*(Used to profile hazard)*
		Elastic			
		Plastic			
		Articulated	Active/Passive		
	Tool Power	Powered	Electric	V & A	Unexpected startup, electric shock, burns
			Hydraulic		High pressure
			Pneumatic	PSI	High pressure, unexpected startup
			Mechanical		Crushing, pinching, stabbing, impact *(E.g., springs)*
			Thermal	°C, °F	Fire, burns, scalds, damage by hot or cold environment
		Unpowered			
	Tool Interaction	Contact	Grasping		Crushing, pinching
			Carrying		Crushing, pinching, falling load
			Pushing		Crushing (load instability)
			Pulling		Crushing (load instability)
			Lifting		Crushing, falling load
			Dispersing		Contact with or inhalation of harmful fluids, gases, mists, fumes, and dusts
			None		
		Non-contact	Dispersing		Contact with or inhalation of harmful fluids, gases, mists, fumes, and dusts
			None		
	Tool Use	Adding			Burns, scalds, contact/inhalation of harmful fluids, gases, mists, fumes, and dusts
		Connecting			Crushing, pinching, stabbing *(Parts are joined together)*
		Moving			Pinching, impact, crushing *(Tool used for material handling and assembly functions)*
		Removing			Cutting, severing, friction, abrasion, stabbing, or inhalation of dust
		Inspecting			Lasers, UV radiation, radio frequency radiation, microwaves
		Heating			Fire, burns, inhalation of fumes
		None			
Physical Interaction	Force Transfer	Applied			Crushing, pinching, impact *(Robot/human hits/pushes other robot/human)*
		Incurred			*(The other robot/human hits/pushes robot/human; may have ergonomic impact)*
		None			*(No physical interaction possible)*
	Force Magnitude	Force/Torque/Pressure		N, Nm, N/mm^2	*(Used to profile hazard)*
	Force Type	Compression			Crushing, pinching *(Force application without retraction)*
		Impact			Impact, crushing
	Force Application	Constant			Crushing
		Ramping			Crushing
		Pulsed	Regular		Impact, whole body vibration
			Irregular		Impact, whole body vibration
	Process Speed	Speed		mm/s	Impact, excessive speed, *(The value of the highest-speed component of the task)*
	Contact State	Constrained			Crushing, pinching
		Free body			Impact
	Contact Edge	Sharp			Cutting, severing, stabbing
		Blunt			Blunt impact, crushing, pinching
Ergonomics	Human Posture	Standing			Neck/back strain
		Kneeling/Squatting			Knee strain, excessive flexion
		Sitting			Back strain
	Manual Material Handling	Carrying			Neck/back/arm strain, excessive flexion, muscle strain, excessive carry weight
		Lifting			Neck/back/arm strain, excessive carry weight
		Pushing			Neck/back/arm strain, excessive flexion
		Pulling			Neck/back/arm strain, excessive flexion
	Manual Action Duration	Time		s	*(Used to profile hazard)*
	Manual Action Load	Mass		kg	*(Used to profile hazard)*

dispersing, connecting, a pneumatic paint sprayer would be stiff, powered pneumatic, noncontact dispersing, adding, a vacuum gripper for palletizing would be elastic, powered pneumatic, contact lifting, moving, and a ruby-tipped stylus probe would be stiff, unpowered, contact pushing, and inspecting. It is generally assumed that tooling for intentional interactions during human–robot collaboration is limited for specific applications such as training, assembly, and assisted lifting. Tools that potentially result in physical harm (e.g., blades, grinding wheels, and welding tips) should be avoided during close-quarters work, and will likely require independent or synchronous collaborations. An alternative to restructuring the collaboration separation involves changing tool designs or types, task processes (e.g., using nontoxic adhesives rather than welding), or even having the human perform certain task components to reduce the risk of hazards.

C. Physical Interaction

Beyond the hazards of the tooling, the physical interactions between the robot and human have historically been the leading factors of robot-related workspace injury [3]. The transfer of forces and pressures from machine to object (applied), and vice versa (incurred), may lead to direct and tertiary injury as part of the normal task process. When such transfers occur, they can be assessed in terms of their expected highest magnitude of the forces, torques, and pressures applicable to both the task and the potential hazards. Moreover, the types of transfer of forces and pressures are describable as being impacts (force applications with subsequent retractions or reductions in force) or compressions (extended force applications with delayed retractions), and the application of forces and pressures (i.e., constant, ramping, or pulsed; pulsed applications, which can be further characterized as being regular or irregular), may also lead to additional hazards.

In this analysis, the point of contact between the robot and human is generally considered to be the most likely site of localized injury. Resistance to injury and tolerance to pain vary between individuals and by regions of the body [22]. Moreover, depending on the contact state (constrained or free body) and contact edge (sharp or blunt), the limits on force, torque, and pressure will go up or down accordingly, and many localized injuries can be avoided by rounding sharp edges or increasing the area of possible contact points.

These physical interaction descriptors are then chained together much the same way as the tool features and functions. For instance, a steady tapping motion of a 1 mm probe against the surface of a table can be classified as, applied, impact, pulsed regular, sharp, constrained, at 50 N. In contrast, a 15 mm ball bearing dropped onto a table could be classified as, “applied, impact, pulsed irregular, blunt, constrained, at 30 N, 300 mm/s, for 0.08 s duration” followed by “applied, compression, constant, blunt, constrained, at 0.6 N.”

D. Task Ergonomics

When a human is directly involved with a task, ergonomics has a direct impact on not only the operator’s comfort, but on the operator’s health and wellbeing. Repetitive strain or deep-tissue injury [41] may result in incorrect ergonomics, and should therefore also be included in the application assessment. Several quantitative (albeit subjective) metrics for ergonomics that map physical exertion to the perception of pain (e.g., Borg’s ratings of perceived exertion (RPE) [42], category of scale with ratio properties, CR10 [43], and CR100 [44]) have been proposed, and such functions are often included in addition to other risk assessments (e.g., the automotive assembly worksheet [45]) for workspace and process evaluation (see [46]). For tasks or events that do not involve humans, the topic of ergonomics is not applicable.

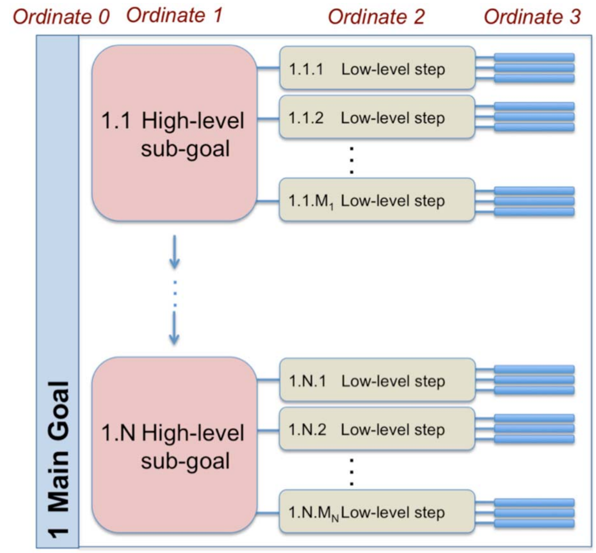

Fig. 1. Task decomposition ordinate hierarchy in which the ordinate-0 main goal is completed by following the ordinate-1 sub-goals that represent discrete stages of the manufacturing process. Ordinate-2 and -3 steps compose the task-level instructions and low-level motion primitives, respectively, to complete the ordinate-1 sub-goals.

To capture the potential risks associated with ergonomics, a task may be characterized by any number of stress or strain-related factors. When assessing the potential impact on the knees, legs, back, or feet, we consider the required or expected human posture of the workforce (standing or kneeling/squatting). Moreover, similar to the functionality of the tooling involved with a task, when considering the human body as a tool, the manual material handling (e.g., carrying, lifting, pushing, or pulling), the manual action duration, and the manual action load should also be assessed.

V. Risk Assessment and Abatement

For a given task, the various atomic subtasks and actions should be segmented and assessed independently to determine each step’s potential risk severity. In this section, we describe the process for task decompositions and, subsequently, assessing and abating the risks associated with that task.

A. Task Decomposition

A key step to assessing task-based safety is to identify all subtasks necessary to complete a given process. Using an extension of the Hierarchical Task Analysis (HTA) [47] similar to those proposed by Tan *et al.* [33] and Woodman *et al.* [48], a task can be decomposed into the ordinate steps that compose the task process plan (Fig. 1).

0: The main goal of the task, which consists of a plan of sub-ordinate steps that describe high-level sub-goals used to complete a product or finished sub-component.

1: High-level sub-goals that represent the steps necessary to complete important milestone stages of a process, but do not constitute finished products, themselves.

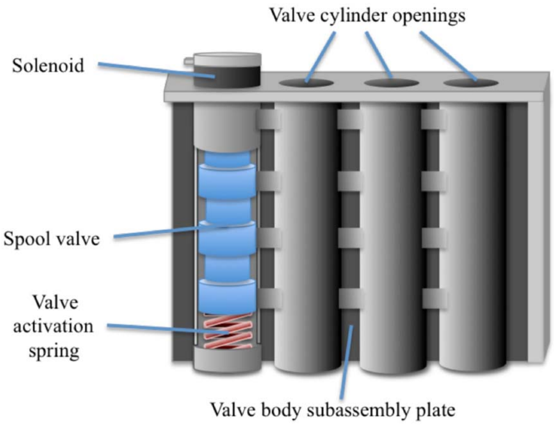

Fig. 2. Simplified illustration of a valve body subassembly. The cutaway (left) shows the component parts of the subassembly in order of insertion into the valve cylinder opening (top): the valve activation spring, the spool valve, and then the solenoid.

2: Low-level steps required to complete the high-level subgoals, and consist largely of motion primitives and base actions.
3: Joint-, tool-, and sensor-level steps that are required to complete the low-level motion primitives described in ordinate-2.

Scoping these ordinates is not a trivial task, and proper care should be taken when developing the HTA. It is important to note that ordinate-0 goals result in a finished component, which may or may not necessarily be a finished product. Therefore, ordinate-0 goals should be considered discrete processes performed at a manufacturing station that result in some output for subsequent input to a manufacturing station. As such, an entire manufacturing process may be composed of several ordinate-0 goals chained together. For example, the construction of an automobile would most likely consist of several thousand ordinate-0 goals such as attach front windshield, or insert rear seat assembly. The creation of these subcomponents, may likewise be composed of several ordinate-0 goals, but are not directly linked to the process of final vehicle assembly.

In many cases, knowledge of parts and tools can be used to automatically generate and populate hazard fields. For example, when pairing a task decomposition with an existing knowledge database, known hazards associated with the workpieces should be linked with ordinate-1 tasks. To illustrate, if it is known that a workpiece has sharp edges, this information can be used to automatically generate fields for cutting and stabbing hazards. Similarly, tools and processes, if known, should be specified for each ordinate-2 subtask. These tools and processes may be cross-referenced for repetitive subtasks.

While scoping a task to ordinate-3 (or beyond) may be necessary for intelligent, agile robot planning and optimization, it does not provide useful insight to the task-based safety of an operation. For example, defining the insertion depth of a rod or the target torque on a bolt are necessary for assessing the termination conditions of an assembly action, but are not directly related to the safety of the process. We suggest, for the purpose of task decomposition for collaborative safety, that a given goal be reduced no further than ordinate-2.

As an example, we consider the simple case of assisted valve body subassembly, in which a series of valve activation springs and spool valves are inserted into the valve subassembly plate (Fig. 2). The assembly process consists of the subassembly plate being fitted with springs and spool valves into four spool channels, and then handed off to the next station for the final assembly of the transmission valve body. There are four spring-spool pairs that must be inserted. For this task, the human is responsible for inserting the springs into the valve cylinder, while the robot inserts the spool valves. The robot was chosen for the spool valve insertion to reduce the likelihood and severity of potential parts damage resulting from the binding of the metal components. For this process, the task decomposition is as follows.

1. Valve body subassembly:
   1.1. Fixture subassembly plate:
      1.1.1. Hold channel plate such that the inner bolt holes are facing upward and are on the left edge.
      1.1.2. Align plate outer bolt holes with fixture pegs.
      1.1.3. Gently lower plate onto pegs to secure it in place.
   1.2. Insert spring #1:
      1.2.1. Align spring with valve cylinder opening.
      1.2.2. Drop spring into cylinder.
   1.3. Insert spring #2:
      1.3.1. Align spring with valve cylinder opening.
      1.3.2. Drop spring into cylinder.
   1.4. Insert spool valve #1:
      1.4.1. Align spool with valve cylinder opening.
      1.4.2. Insert spool into valve cylinder opening.
   1.5. Insert spring #3:
      1.5.1. Align spring with valve cylinder opening.
      1.5.2. Drop spring into cylinder.
   1.6. Insert spool valve #2:
      1.6.1. Align spool with valve cylinder opening.
      1.6.2. Insert spool into valve cylinder opening.
   1.7. Insert spring #4:
      1.7.1. Align spring with valve cylinder opening.
      1.7.2. Drop spring into cylinder.
   1.8. Insert spool valve #3:
      1.8.1. Align spool with valve cylinder opening.
      1.8.1. Insert spool into valve cylinder opening.
   1.9. Insert spool valve #4:
      1.9.1. Align spool with valve cylinder opening.
      1.9.2. Insert spool into valve cylinder opening.
2. Evaluate subassembly:
   2.1. Evaluate free motion:
      2.1.1. Push down on each spool valve to verify free motion of all valves.
   2.2. Evaluate sitting/insertion depth:

TABLE IV
RISK MATRIX USED BY THE NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY'S ENGINEERING LABORATORY FOR ASSESSING HAZARDS

		POTENTIAL SEVERITY			
		CATASTROPHIC	SEVERE	MODERATE	MINOR
LIKELIHOOD OF OCCURRENCE (OVER 1-YEAR PERIOD)	FREQUENT (Likely to occur repeatedly)	HIGH (4)	HIGH (4)	SERIOUS (3)	MEDIUM (2)
	PROBABLE (Likely to occur multiple but infrequent times)	HIGH (4)	HIGH (4)	SERIOUS (3)	MEDIUM (2)
	OCCASIONAL (Likely to occur at some time)	HIGH (4)	SERIOUS (3)	MEDIUM (2)	LOW (1)
	REMOTE (Possible, but not likely to occur)	SERIOUS (3)	MEDIUM (2)	MEDIUM (2)	LOW (1)
	IMPROBABLE (Very unlikely, can reasonably assume it will not occur)	MEDIUM (2)	LOW (1)	LOW (1)	LOW (1)

2.2.1. Measure distance from the lip of each valve cylinder opening to the top of the respective spool valve to verify insertion depth.

Once completed, the transmission valve body subassembly is passed to the next workstation, where it is integrated onto the transmission's valve body channel plate. At the next station, a solenoid cable assembly is connected to the transmission valve body assembly, and the entire assembly is attached to the automatic transmission prior to being coupled to the engine block.

Here we have provided the step-by-step instructions for making a valve-body subassembly. Notice, however, that there are no roles or indications of timing assigned to the various subtasks. Evaluating the task decomposition, it is not immediately clear where and how the collaborative elements manifest. Such information can be added to the task decomposition to enable more accurate risk assessments.

Recall that the human is responsible for inserting the springs while the robot is expected to insert the valve spools. As such, for subtasks 1.2, 1.3, 1.5, and 1.7, we can add the tag (human), and add (robot) to subtasks 1.4, 1.6, 1.8, and 1.9. These tags apply to all subordinates unless otherwise noted.

Also, note that the task scheduling calls for springs #1 and #2 to be inserted prior to the insertion of the first spool. Based on Fig. 2 and what is known of the assembly process, we know that spring #1 must be inserted before spool #1. We also see that the insertion of spring #2 does not impact the ability to insert spool #1. We may thus assume that spool #1 may be inserted simultaneous to the insertion of spring #2. The ordering of tasks generally implies a temporal order of operations. For example, subtask 1.3 happens after subtask 1.2. For cases of simultaneity, we may add notes denoting such to other subtasks or substeps. For example, subtask 1.4 may be noted that it occurs simultaneous to subtask 1.3.

1.3. Insert spring #2 (human):
  1.3.1. Align spring with valve cylinder opening.
  1.3.2. Drop spring into cylinder.
1.4. Insert spool valve #1 (robot, simultaneous 1.3).

Similar role and temporal order applications can thus be applied throughout the entire task.

Based on this additional metadata, we are now better prepared to perform the risk assessment.

B. *Risk Assessment*

As part of the risk assessment process, risk matrices help end users and integrators identify highest priority hazards that may result in injury. Risk matrices assign a hazard priority based on the combination of the worst-case expected severity of an injury and the likelihood that said injury will occur. For example, if a potential injury is expected to be minor, at worst, and unlikely to actually occur, then there is a low risk of injury. The risk matrix given in Table IV, from the American National Standards Institute's (ANSI) Z10 standard [49], is used by the National Institute of Standards and Technology (NIST) Engineering Laboratory as part of a first-level hazard review process for new laboratories, testbeds, and research activities. Different organizations may use alternative risk matrices per their organizations' policies. For instance, this methodology does not factor in the likelihood of detectability of hazards (see [50]). In this report, we will use the risk matrix given in Table IV in an illustrative case study in Section VI.

Risk severity is assessed on a scale of 1 to 4, where 1 is low risk, and 4 is high risk. A value of 0 may be assigned only when the potential for injury is impossible. The risk severity for the combined task is the maximum value of all associated subtasks. For instance, if steps 1 through 5 of a given 7-step task have risk level 1, while steps 6 and 7 have severity 3 and 2, respectively, then the task's total risk severity is 3. We believe that the total risk for a collaborative task should be at most 1. A risk severity of 2 may be acceptable in some organizations and settings, but would ideally be reduced to 1 by means of risk abatement. Risk severities greater than 2 are unacceptable for any collaborative task. Note that this methodology is used to identify and assess the risks of a given task, only, and does not assess the total risk of the system or take into account an organization's risk tolerance (see [51]). Such considerations are important, but are beyond the scope of this process.

The process for identifying risks is largely guided by the task decomposition paired with the ontology discussed in Section IV. Each ordinate-2 subtask is evaluated using the descriptive predicate table (Table III), with the applicable fields identified and the known magnitudes specified. These then draw attention to potential hazards, which are then assessed using Table IV. Note that the hazards listed in Table III are common to the tasks covered in this methodology. Additional

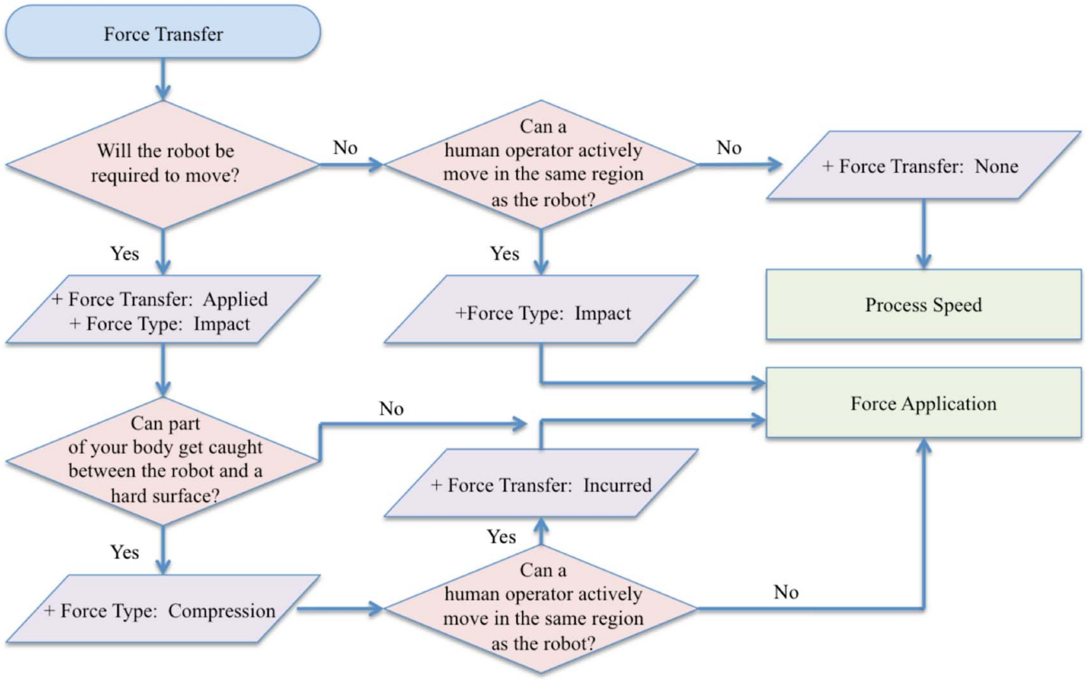

Fig. 3. Simplified decision flowchart for prompting the user for risk assessment information regarding the force transfer and force type data fields. Following the population of these fields, the process tree then forwards the user to evaluate either process speed or force application fields.

TABLE V
TRANSMISSION VALVE BODY SUBASSEMBLY SUBTASK 1.2.2 (DROP SPRING INTO CYLINDER) RISK ASSESSMENT

PREDICATE	VALUE	MEASURE	HAZARD (RISK)	ABATEMENT STRATEGY	POST-ABATEMENT RISK
Collaboration	Simultaneous				
Tool Compliance	Elastic (*Hands*)				
Tool Power	Unpowered		N/A (0)		N/A (0)
Tool Interaction	Contact – Grasping		N/A (0)		N/A (0)
Tool Use	Moving		N/A (0)		N/A (0)
Force Transfer	None		N/A (0)		N/A (0)
Force Magnitude	Force	0 N			
Force Type	Impact		Impact (1)		Impact (1)
Force Application	Pulsed – Irregular		Impact (1)		Impact (1)
Process Speed	Speed	500 mm/s	Impact (1)		Impact (1)
Contact State	Free body		Impact (1)		Impact (1)
Contact Edge	Sharp		Cutting (2), Stabbing (2)	Cut-resistant gloves	Cutting (1), Stabbing (1)
Human Posture	Sitting		Back strain (1)		Back strain (1)
Manual Material Handling	Carrying		Excessive flexion (1)		Excessive flexion (1)
Manual Action Duration	Time	4 s			
Manual Action Load	Mass	0.057 kg			
		Max Risk:	2	Max Risk:	1

hazards may exist. Refer to relevant documentation such as a material data safety sheet (MSDS) or tool instruction manual to identify these hazards.

Data entry for the risk assessment is historically a manual process carried out by a safety officer and/or a process engineer. Due to the formalization of an ontology for task-based risk assessments, however, the population of forms can be guided by software tools that prompt the user for specific inputs in plain English and simultaneously validate the inputs as they are entered. For instance, when evaluating force transfer, the decision tree for specifying and entering inputs can be implemented as shown in Fig. 3. In this flowchart, the decision rules are fairly simple, but can capture the nature of some of the possible physical hazards associated with a given subtask. Additional inputs or database lookups can be used to automatically populate the Force Magnitude field. For instance, the worst-case magnitude of the force transfer can be approximated using known properties of the robot, workpiece, and process parameters. Otherwise the safety officer may enter magnitudes based on estimated injury criterion (see [52]).

For subtasks 1.4, 1.6, 1.8, and 1.9 in our valve body subassembly example, let us assume that the robot is holding the spool valve by a single point-of-contact steel vacuum gripper with a bellowed rubber tip. In contrast, the human operator is only using his hands to pick up, carry, and insert the springs into the valve cylinder opening. The steel springs have sharp

TABLE VI
TRANSMISSION VALVE BODY SUBASSEMBLY SUBTASK 1.4.2 (INSERT SPOOL INTO VALVE CYLINDER OPENING) RISK ASSESSMENT

PREDICATE	VALUE	MEASURE	HAZARD (RISK)	ABATEMENT STRATEGY	POST-ABATEMENT RISK
Collaboration	Simultaneous				
Tool Compliance	Stiff				
Tool Power	Powered - Pneumatic	100 PSI	High pressure (1)		High pressure (1)
Tool Interaction	Contact – Carrying		Pinching (1)		Pinching (1)
Tool Use	Moving		Pinching (1), Impact (1)		Pinching (1), Impact (1)
Force Transfer	Applied		Pinching (2)		Pinching (1)
Force Magnitude	Force/Torque	100 N / 120 Nm		Reduce force to 40 N	
Force Type	Compression		Pinching (2)		Pinching (1)
Force Application	Constant		Crushing (1)		Crushing (1)
Process Speed	Speed	1000 mm/s	Impact (1)		Impact (1)
Contact State	Constrained		Crushing (1), Pinching (2)		Crushing (1), Pinching (1)
Contact Edge	Blunt		Crushing (1), Pinching (2)		Crushing (1), Pinching (1)
Human Posture	N/A		N/A (0)		N/A (0)
Manual Material Handling	N/A		N/A (0)		N/A (0)
Manual Action Duration	N/A				
Manual Action Load	N/A				
		Max Risk:	2	Max Risk:	1

points at both ends of the wide coil, and the spool valves are smooth, cast aluminum with rounded edges.

Despite its brevity, there are a considerable number of ordinate-3 subtasks to evaluate. To illustrate the risk assessment and abatement procedures, we will provide as exemplars the details of subtasks 1.2.2 and 1.4.2. Tables V and VI illustrate the predicates and risk assessments of subtasks 1.2.2 and 1.4.2, respectively. Each of these tables is derived from Table III, and the potential hazards are assessed using the risk matrix given in Table IV.

Because the human's subtask of inserting the valve activation spring does not involve the use of tools, there is no risk of tool-related injury, while there exists the possibility of injury from the robot's pneumatic gripper. Moreover, while the human is performing his subtasks, he is just as likely to hit the robot as the robot is to hit him. In either event, there is the potential for occasional impact, although the risk of injury is minor. Using our risk matrix in Table IV, the combination of occasional occurrence and minor severity, the potential hazard is ranked as low for both Tables V and VI.

Instead, the greatest risk of injury stems from the parts and processes of the assembly. For subtask 1.2.2, there is the potential for cutting or stabbing of the human operator's hands by the sharp ends (contact edge—sharp) of the spring. Even though such injuries are expected to occur, at most, only occasionally, there is the potential for moderate severity given the sensitivity of the hands. Similarly, there is a risk of pinching injury during subtask 1.4.2 based on several factors.

1) The robot is applying force to perform the assembly (force transfer—applied).
2) The applied forces are of a compressive nature (force type—compression).
3) The subassembly plate is fixtured on a table (contact state—constrained).
4) The parts being assembled are rounded (contact edge—blunt).

For each, there is a remote possibility that the human's hands could be pinched between the spool valve and the subassembly plate while the human is working near the robot. This risk of injury is expected to have, at most, moderate severity. Using the risk matrix in Table IV, both the cutting/stabbing hazard (Table V) and the pinching hazard (Table VI) are given medium risk severities.

Based on the task decomposition for the valve body subassembly, we see that a number of subtasks are either repeated or directly related to other subtasks. While this may not necessarily be the case for all collaborative tasks, it does occur often enough that the risk assessment process is simplified through repetition and derivation. Moreover, even in cases where the products and processes of a given manufacturing plant are prone to change between model years, the actual subtasks are not expected to change significantly. This allows for the risk assessments to be recorded and, in many instances, reused, further simplifying the risk assessment process.

C. Risk Abatement

This severity assessment enables two critical system capabilities. First it identifies the individual steps and processes that pose risks to human operators. This information can be used for the automatic reassignment of roles and responsibilities provided sufficient system capacity, and for the automatic restructuring of task processes when known alternatives exist. Second, the severity assessment identifies the risk for the task as a whole. This information can be used as a larger input when evaluating higher-level processes for safety and system prognostics.

The goal of the risk abatement is to improve the safety of a task by minimizing either the risk or the impact of the hazard. The ideal abatement solution reduces both the likelihood and potential severity of the hazards such that the risk severity for the entire task is at most 1. This is accomplished by two different mechanisms. First, limit or modify the exposure to the risks by choosing alternative collaboration types, changing tools, changing task processes, or adding or changing safeguards. Second, limit or modify the effects of the exposure by changing process settings (e.g., lowering the applied forces or pressures, or slowing down the robot), adding or modifying personal protective equipment (PPE), or changing design characteristics of the workcell (e.g., adding padding to the robot or rounding sharp edges).

Proposed changes should be weighed against their potential impact on the task performance and any prior or ensuing tasks. For example, adding safeguards to prevent access to the workcell while a robot is drilling parts may also slow or inhibit collaborative subtasks in which the robot and a human worker are inserting retaining bolts into the holes the robot drilled. In a previous report [23], we presented a mechanism to measure the impact that implementing new safety protocols on an existing process will have on the process' productivity. The productivity metric is simple, but provides a convenient input when designing and assessing hazard abatement strategies. The task impact is measured as

$$p_r = \frac{\hat{t}}{t_r}. \quad (1)$$

Here, p_r is the impact for implementing the rth abatement strategy, $\hat{t}$ is the nominal time to accomplish a task without hazard abatement, and t_r is the time to complete the same task with the hazard abatement strategy in effect.

In our example, subtasks 1.2.2 and 1.4.2 both have a maximum hazard risk of 2. The highest risk of injury in subtask 1.2.2 pertains to the risk of cutting/stabbing from the sharp ends of the springs, while several aspects of subtask 1.4.2 contribute to the potential for a pinching injury.

The cutting/stabbing hazard of subtask 1.2.2 can be mitigated by three possible means. First, the sharp ends can be rounded to eliminate the cutting and stabbing hazards entirely. Second, the human operator can be required to wear cut-resistant gloves to reduce the likelihood of hazard occurrence from occasional to improbable. Third, a second robot process could be inserted into the assembly task to perform the human's subtasks. Depending on the sourcing of the springs, the first option may not be a viable solution without significantly impacting the throughput or efficiency of the manufacturing process. Specifically, grinding springs is both difficult and labor-intensive, and often introduces new hazards. The third option (adding a new robotic process) will require specialized tooling (i.e., a robot gripper capable of grasping springs) and additional time and effort to write and maintain the robot's programming. The risk abatement strategy that has the least impact on productivity is thus the second option in which the human operator is required to wear PPE to effectively remove the likelihood of injury. Reassessing the risk severity, this reduction in severity is reflected in Table V in the abatement strategy column, and the post-abatement hazard is recorded.

As with the cutting/stabbing hazard of subtask 1.2.2, the pinching hazard of subtask 1.4.2 has a number of potential abatement strategies available. First, the collaboration mode could be changed from simultaneous to synchronous by reordering the subtasks such that the human must insert all springs prior to the spool valve being inserted. This effectively removes the pinching hazard by ensuring that the human's hands will never be in the working envelope of the active robot. However, it will remove the possibility of parallelizing the assembly process, and will impact the time necessary to complete the subassembly. A second option stems from the observation that the severity of the pinching hazard is based on the force with which the robot pushes the spool valve into place. Minimizing the applied force from 100 to 40 N also reduces the potential severity of the pinching hazard from moderate to minor. However, this will also extend the assembly time by subsequently slowing the motions of the robot as it inserts the spool valve. A third option is to remove the robot altogether and have the human operator insert the spool valve. While this completely eliminates the pinching hazard, it has the same effect on the assembly time as the first option, and may incur additional time penalties due to the binding of parts that the robots were initially employed to avoid. Of the three, the second option (to reduce the applied force from 100 to 40 N) has the lowest potential for negatively impacting the assembly process. This abatement strategy and the post-abatement risk are recorded in Table VI.

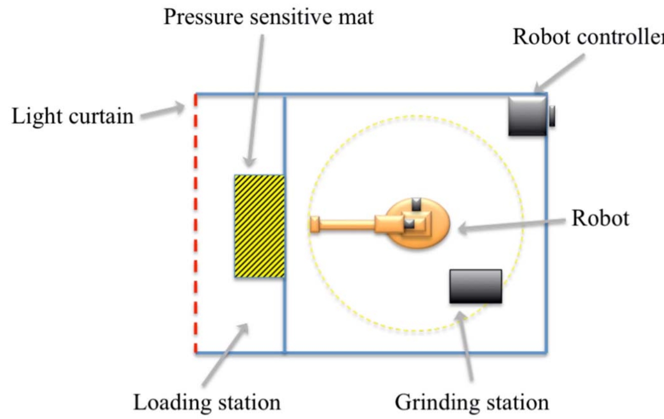

Fig. 4. Example robot cell with a grinding station and a loading station. A protective guard fence surrounds the robot and the loading station is separated from the robot's work volume by a low barrier that allows the operator to reach inside to hand a part to the robot.

VI. Case Study: Part Handoff

In this section we present a simple, illustrative example of a collaborative task in which an operator must hand a workpiece to a robot for further machining.

A. Task Description

An assembly line is being designed that integrates both human workers and industrial, collaborative. The manufacturing process involves the assembly of parts and post-assembly surface finishing. The part has a mass of 2.5 kg, and its surfaces are rough, but not sharp. During the initial design phase, it was decided that the robots would handle the surface finishing in a closed workcell into which entry is impossible while the robot is active. The surface finishing is accomplished using a stationary grinder, against which the robot holds the workpiece using a custom pneumatic gripper that fixtures the part firmly in its grasp with a known orientation. The human operators, in turn, would perform the more difficult task of assembling the parts prior to finishing. To save time and integration costs, the transition between the two processes would be a direct transfer of parts from man to machine. The collaborative task, therefore, would be a part handoff, where the robot would take the assembled parts out of the hands of a human operator. Prior to bringing this new assembly process online, a risk assessment must be completed to ensure operator safety. The hand assembly process is legacy, and the risks and abatement strategies are well known. The robotic finishing process is new, but was installed by an integrator who also provided a full risk assessment, and installed safeguards to ensure operator safety. The parts handoff, however, has not yet had a hazard review.

The handoff process consists of the following steps. Once a part has been assembled, the operator moves to a loading station (Fig. 4). The robot senses the operator's presence via a pressure sensitive mat, and moves its end effector to a predefined location within reach of the operator at 500 mm/s while the operator stands at the ready. Prior to this point the robot is stationary within its workcell. The operator then presents the part to the robot by placing it inside the gripper's jaws, which clamp down on the part when it is detected to be within acceptable pose tolerances. The gripper has been designed to automatically correct for minor part orientation and position errors, and holds the part firmly by applying 200 N/mm^2 of force at all contact points, and the robot's program logic will close the gripper only when the robot itself is stationary. The robot will not move again until the operator moves off of the pressure sensitive mat and away from the loading station, after which it performs the surface finishing process inside the confines of its workcell. Three distinguishable phases can be easily identified in this task: 1) bring the part to the robot; 2) hand the part to the robot; and 3) then depart from the robot station.

B. Task Decomposition

From the task description above, one possible task decomposition for the part handoff process is as follows.

1. Part handoff:
   - 1.1. Bring part to loading station:
     - 1.1.1. Approach loading station with part, stepping onto pressure-sensitive mat.
     - 1.1.2. Move robot to loading configuration.
   - 1.2. Hand off part to robot:
     - 1.2.1. Align part with robot's gripper fingers in proper orientation.
     - 1.2.2. Close robot gripper firmly onto part.
   - 1.3. Leave loading station:
     - 1.3.1. Step off of pressure-sensitive mat and away from loading station.
     - 1.3.2. Move robot into surface finishing workcell.

Subtasks 1.1 and 1.2 involve the operator's contact with the part, so known part-related hazards can be automatically associated with these subtasks. Of the three ordinate-1 subtasks, only 1.2 involves physical interaction with the robot, though subtask 1.1.2 involves a potential collision hazard as the robot moves into the loading configuration. Subtask 1.3 involves no possible interaction with either the part or the robot, so only normal environmental risks need to be considered.

C. Risk Assessment

In the task decomposition, three ordinate-2 subtasks are identified that have possible physical interaction between the operator and the robot: 1) 1.1.2 "Move robot to loading configuration;" 2) 1.2.1 "Align part with robot's gripper fingers in proper orientation;" and 3) "Close robot gripper firmly onto part." We will evaluate these in the order presented.

In subtask 1.1.2, the robot moves at 500 mm/s toward the human operator, and the activities of both the operator and robot are simultaneous as they near the loading position. Beyond what was presented in the task description, not much is currently known about the robot. Assume that as we look at the robot's specifications, we learn that both the robot and its gripper are stiff, the gripper is actuated by 100 PSI to achieve the 200 N/mm^2 of applied pressure at each contact point (though in this subtask the gripper is not being used), and that the robot has a maximum acceleration of 1000 mm/s^2 and a total mass of 300 kg. The loading station was designed such that all immovable surfaces are at least 500 mm beyond the robot's work volume, and the process engineers have assessed nominal completion times for manual operations based on average walking speeds and empirical evaluations. Based on this information, we can identify the following hazards for subtask 1.1.2: high pressure (pneumatic gripper), impact (robot, pneumatic gripper), neck/back strain (human standing), neck/back/arm strain (human carrying), excessive flexion (human carrying), muscle strain (human carrying), excessive carry weight (human carrying). During a detailed evaluation of the process, it was determined that the ergonomic hazards may occasionally occur, but their severity is expected to be minor (low risk). Moreover, the high-pressure hazard is determined to be both remote and minor (low risk) due to engineering constraints on the hose connectors. The impact hazard is deemed to be remote, but could result in moderate to severe injury (medium risk) should the human move too close to the loading position while the robot is still moving. The risk assessment for subtask 1.1.2 is thus reflected in Table VII.

In subtask 1.2.1, the operator must lift and position the part into the robot's waiting gripper. Here, it is assumed that the robot is stationary and that actions occur sequentially, but it is also possible that the operator begins moving the part toward the robot while the robot is still moving. We can reuse much of the information that has already been identified, and thus assess the following risks: high pressure (pneumatic gripper), crushing (robot), pinching (robot), impact (robot, pneumatic gripper), neck/back strain (human standing), neck/back/arm strain (human carrying, human lifting), excessive flexion (human carrying, human lifting), muscle

TABLE VII
PART HANDOFF SUBTASK 1.1.2 (MOVE ROBOT TO LOADING CONFIGURATION) RISK ASSESSMENT

PREDICATE	VALUE	MEASURE	HAZARD (RISK)
Collaboration	Simultaneous		
Tool Compliance	Stiff		
Tool Power	Powered - Pneumatic	100 PSI	High pressure (1)
Tool Interaction	Contact – None		
Tool Use	None		
Force Transfer	Applied		Impact (2)
Force Magnitude	Force/Torque	300 N	
Force Type	Impact		Impact (2)
Force Application	Pulsed – Regular		Impact (2)
Process Speed	Speed	500 mm/s	Impact (2)
Contact State	Free body		Impact (2)
Contact Edge	Blunt		Blunt impact (2)
Human Posture	Standing		Neck/back strain (1)
Manual Material Handling	Carrying		Neck/back strain (1) Excessive carry weight (1)
Manual Action Duration	Time	15 s (nominal)	
Manual Action Load	Mass	2.5 kg	
		Max Risk:	2

TABLE VIII
PART HANDOFF SUBTASK 1.2.1 (ALIGN PART WITH ROBOT'S GRIPPER FINGERS IN PROPER ORIENTATION) RISK ASSESSMENT

PREDICATE	VALUE	MEASURE	HAZARD (RISK)
Collaboration	Simultaneous		
Tool Compliance	Stiff		
Tool Power	Powered - Pneumatic	100 PSI	High pressure (1)
Tool Interaction	Contact – None		
Tool Use	None		
Force Transfer	Applied		Impact (2)
Force Magnitude	Force/Torque	300 N	
Force Type	Impact		Impact (2)
Force Application	Pulsed – Regular		Impact (2)
Process Speed	Speed	500 mm/s	Impact (2)
Contact State	Free body		Impact (2)
Contact Edge	Blunt		Blunt impact (2)
Human Posture	Standing		Neck/back strain (1)
Manual Material Handling	Carrying Lifting		Neck/back strain (1) Excessive flexion (1) Muscle strain (1) Excessive carry weight (1)
Manual Action Duration	Time	5 s (nominal)	
Manual Action Load	Mass	2.5 kg	
		Max Risk:	2

strain (human carrying, human lifting), excessive carry weight (human carrying, human lifting). As before, only the impact hazard has a risk of medium or greater due to the potential for simultaneous occupancy and motion within the shared workspace. These risks are reflected in Table VIII.

Again, for subtask 1.2.2 we can reapply many of the previous ergonomic assessments. Because the robot will only close the gripper when the robot is stationary, any possible impact with the robot will be the result of the operator's actions. The question now is, under what circumstances could injury possibly occur? For this subtask, the biggest risk is of crushing and pinching the operator by the pneumatic gripper. Such injuries can occur if the operator's hands are between the gripper fingers and the workpiece, or if a sensing error mistakes the operator's arm or hand as the work piece and clamps down prematurely. Moreover, if the part is not securely gripped, it can be ejected from the gripper and impact the operator. Thus, the following hazards are identified: high pressure (pneumatic gripper), crushing (pneumatic gripper), pinching (pneumatic gripper), impact (pneumatic gripper, workpiece), falling load (workpiece), neck/back strain (human standing), neck/back/arm strain (human carrying, human lifting), excessive flexion (human carrying, human lifting), muscle strain (human carrying, human lifting), and excessive carry weight (human carrying, human lifting). Without proper safeguards, it is determined that the crushing, pinching, impact, and falling load hazards can occasionally happen, and can result in injuries ranging from minor to severe (low to serious risk). These risks are reflected in Table IX.

D. Risk Abatement

From the risk assessment, it is clear that the largest hazards stem from the potential for the robot and operator to simultaneously move in the shared workspace and from crushing and pinching hazards caused by the robot's gripper. A number of preventative measures can be taken to abate these risks, though

TABLE IX
PART HANDOFF SUBTASK 1.2.2 (CLOSE ROBOT GRIPPER FIRMLY ONTO PART) RISK ASSESSMENT

PREDICATE	VALUE	MEASURE	HAZARD (RISK)
Collaboration	Simultaneous		
Tool Compliance	Stiff		
Tool Power	Powered - Pneumatic	100 PSI	High pressure (1)
Tool Interaction	Contact – Grasping Contact – Carrying		Crushing (3) Pinching (2) Falling load (1)
Tool Use	Moving		Crushing (3) Pinching (2)
Force Transfer	Applied		Impact (2)
Force Magnitude	Pressure	200 N/mm^2	
Force Type	Impact Compression		Impact (2) Crushing (3)
Force Application	Constant		Crushing (3)
Process Speed	Speed	0 mm/s	
Contact State	Free body Constrained		Impact (2) Crushing (3) Pinching (2)
Contact Edge	Blunt		Crushing (3) Pinching (2)
Human Posture	Standing		Neck/back strain (1)
Manual Material Handling	Carrying Lifting		Neck/back strain (1) Excessive flexion (1) Muscle strain (1) Excessive carry weight (1)
Manual Action Duration	Time	<3 s (nominal)	
Manual Action Load	Mass	2.5 kg	
		Max Risk:	3

many of these steps are not without drawbacks. For instance, let us first look at the impact risk with a moving robot. The handoff process can be altered such that the robot must already be in the waiting configuration before the operator is allowed to enter the loading station. However, if the robot is not already in this position before the operator arrives at the loading station, the operator will be forced to wait for the robot to be ready. This ultimately impacts the process time, and can result in delays elsewhere on the assembly line. Similarly, reducing the robot's speed and acceleration mitigates the risk, but also impacts the process time.

However, neither of these abatement strategies address the severe risk of crushing caused by the robot's gripper. The operator could be required to wear PPE that reduces the crushing hazard, but it is likely that such PPE would also get in the way during the assembly process preceding the handoff. This would require the operator to put on and remove the PPE multiple times throughout the production process. Reducing the gripping force would also reduce the crushing and pinching hazards, but may result in the part being ejected during the surface finishing stage. Similarly, a new gripper could be designed such that the part could be held just as securely with less pressure. But this would require a new, expensive process of design and fabrication, and would impact the assembly line until the new gripper is ready for deployment.

One possible abatement strategy involves a redesign of the handoff process. Rather than the operator handing the part directly to the robot, he will instead insert the part into a fixture that will hold the part in a known position and orientation. The same sensors that were used to verify the part was between the robot's gripper fingers can now be used, instead, to verify the presence of the part in the fixture. Meanwhile, the robot remains in its workcell until it verifies that: 1) a part has been placed in the fixture and 2) the operator has left the loading station. Once both conditions are met, the robot will then acquire the part from the fixture and proceed with the surface finishing process. Standard safeguards can be put in place at the workstation such that, if an operator enters while the robot is acquiring the part from the fixture, the robot will immediately stop until the operator has left again. This abatement strategy eliminates all impact, crushing, and pinching hazards at the cost of requiring a new fixture to be designed and turning a collaborative task into a noncollaborative task.

VII. DISCUSSION

This paper introduced a strategy for characterizing and assessing the task-based safety of collaborative manufacturing tasks. By utilizing a flexible ontology for task decomposition, the safety of a given task can be assessed quickly by evaluating the base elements of its subtask components. These elements then provide the bases for mitigating hazard risks.

When coupled with the construction of a database of similar tasks and subtasks, this methodology lends itself to the partial or full automation of risk assessment and abatement for collaborative tasks. Through the breakdown of a task's subcomponents, the process of identifying and assessing risks, assigning roles and responsibilities to humans and robots, and suggesting risk-reducing steps can be executed efficiently through software. Ongoing efforts at NIST include extending, generalizing, and refining this methodology to describe both the processes and the roles of task contributors. Across many different fields, activity-based risk assessments are being defined where the hazards are separate from the tools and environments utilized. As more flexible and intelligent robot systems are integrated into human-occupied environments, the number and requirements for these task-based risk assessments will increase. NIST is actively working with industry partners and standards organizations to identify and refine test methods and metrics for the verification and validation of the safety of these flexible robotic systems.

REFERENCES

[1] H. I. Christensen *et al.*, *A Roadmap for U.S. Robotics: From Internet to Robotics*. Computing Community Consortium, 2009.

[2] N. Sugimoto, "Safety engineering on industrial robots and their draft standard safety requirements," in *Proc. 7th Int. Symp. Ind. Robots*, 1977, pp. 461–470.

[3] B. C. Jiang and C. A. Gainer, "A cause-and-effect analysis of robot accidents," *J. Occup. Accid.*, vol. 9, no. 1, pp. 27–45, 1987.

[4] T. Malm *et al.*, "Safety of interactive robotics—Learning from accidents," *Int. J. Soc. Robot*, vol. 2, no. 3, pp. 221–227, 2010.

[5] Z. Ji, R. Qiu, D. Li, and S. Xu, "Towards automated task planning for service robots using semantic knowledge representation," in *Proc. IEEE Int. Conf. Ind. Inf.*, Beijing, China, 2012, pp. 1194–1201.

[6] S. Minhas, C. Juzek, and U. Berger, "Ontology based intelligent assistance system to support manufacturing activities in a distributed manufacturing environment," in *Proc. 45th CIRP Conf. Manuf. Syst.*, vol. 3. 2012, pp. 215–220.

[7] A. Giovannini *et al.*, "Ontology-based system for supporting manufacturing sustainability," *Annu. Rev. Control*, vol. 36, no. 2, pp. 309–317, 2012.

[8] Q. Guo and M. Zhang, "An agent-oriented approach to resolve scheduling optimization in intelligent manufacturing," *Robot. Comput. Integr. Manuf.*, vol. 26, no. 1, pp. 39–45, 2010.

[9] M. Rahimi and W. Karwowski, Eds., *Human-Robot Interaction*. Bristol, PA, USA: Taylor & Francis, 1992.

[10] S. A. Green, M. Billinghurst, X. Chen, and J. G. Chase, "Human-robot collaboration: A literature review and augmented reality approach in design," *Int. J. Adv. Robot. Syst.*, vol. 5, no. 1, pp. 1–18, 2008.

[11] T. Fong, C. Thorpe, and C. Baur, "Collaboration, dialogue, and human-robot interaction," *Springer Tracts Adv. Robot.*, vol. 6, pp. 255–266, 2003.

[12] H. Knight and R. Simmons, "Tracking aggregate vs. individual gaze behaviors during a robot-led tour simplifies overall engagement estimates," in *Proc. 7th Annu. ACM/IEEE Int. Conf. Human-Rob. Interact.*, Boston, MA, USA, 2012, pp. 175–176.

[13] P. Rani, N. Sarkar, C. A. Smith, and L. D. Kirby, "Anxiety detecting robotic system—Towards implicit human-robot collaboration," *Robotica*, vol. 22, no. 1, pp. 85–95, 2004.

[14] M. A. Goodrich and E. R. Boer, "Model-based human-centered task automation: A case study in ACC system design," *IEEE Trans. Syst., Man, Cybern. A, Syst. Humans*, vol. 33, no. 3, pp. 325–336, May 2003.

[15] B. S. Medikonda and S. R. Panchumarthy, "An approach to modeling software safety in safety-critical systems," *J. Comput. Sci.*, vol. 5, no. 4, pp. 311–322, 2009.

[16] J. Fryman and B. Matthias, "Safety of industrial robots: From conventional to collaborative applications," in *Proc. 7th German Conf. Robot.*, Munich, Germany, 2012, pp. 1–5.

[17] E. Helms, R. D. Schraft, and M. Hagele, "rob@work: Robot assistant in industrial environments," in *Proc. IEEE Int. Workshop Robot. Humun Interact. Commun.*, 2002, pp. 399–404.

[18] International Organization for Standardization (ISO). *Safety of Machinery—Risk Assessment—Part 1: Principles*, ISO 14121-1:2007, 2007.

[19] A. Pervez and J. Ryu, "Safe physical human robot interaction—Past, present and future," *J. Mech. Sci. Tech.*, vol. 22, no. 3, pp. 469–483, 2008.

[20] *Robots and Robotic Devices—Safety Requirements—Part 1: Robots*, ISO 10218-1, 2011.

[21] *Robots and Robotic Devices—Safety Requirements—Part 2: Industrial Robot Systems and Integration*, ISO 10218-2, 2011.

[22] *Robots and Robotic Devices—Industrial Safety Requirements—Collaborative Industrial Robots*, ISO. ISO/TS 15066.

[23] J. Marvel, "Performance metrics of speed and separation monitoring in shared workspaces," *IEEE Trans. Autom. Sci. Eng.*, vol. 10, no. 2, pp. 405–414, Apr. 2013.

[24] A. Bicchi, M. A. Peshkin, and J. E. Colgate, "Safety for physical human-robot interaction," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Berlin, Germany: Springer, 2008, pp. 1335–1348.

[25] S. Haddadin, A. Albu-Schäffer, and G. Hirzinger, "Requirements for safe robots: Measurements, analysis and new insights," *Int. J. Robot. Res.*, vol. 28, nos. 11–12, pp. 1507–1527, 2009.

[26] P. Trautman and A. Krause, "Unfreezing the robot: Navigating in dense, interacting crowds," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, Taipei, Taiwan, 2010, pp. 797–803.

[27] B. Lacevic and P. Rocco, "Kinetostatic danger field—A novel safety assessment for human-robot interaction," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, Taipei, Taiwan, 2010, pp. 2169–2174.

[28] D. Kulić and E. Croft, "Safe planning for human-robot interaction," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2004, pp. 1882–1887.

[29] M. Zinn, "A new actuation approach for human friendly robotic manipulation," Ph.D. thesis, Stanford University, Stanford, CA, USA, 2005.

[30] International Electrotechnical Commission (IEC). *Safety of Machinery—Electro-Sensitive Protective Equipment—Part 1: General Requirements and Tests*, IEC 61496-1, 2012.

[31] *Safety of Machinery—Application of Protective Equipment to Detect the Presence of Persons*, IEC/TS 62046, 2008.

[32] S. Haddadin *et al.*, "A truly safely moving robot has to know what injury it may cause," in *Proc. IEEE Int. Conf. Intell. Robot. Syst.*, Vilamoura, Portugal, 2012, pp. 5406–5413.

[33] J. T. C. Tan, F. Duan, R. Kato, and T. Arai, "Collaboration planning by task analysis in human-robot collaborative manufacturing system," in *Advances in Robot Manipulators*, E. Hall, Ed. Shanghai, China: InTech China, 2010, pp. 113–132.

[34] D. McDermott *et al.*, "PDDL—The planning domain definition language, version 1.2," Yale Center Comput. Vis. Control, Yale Univ., New Haven, CT, USA, Tech. Rep. CVC TR-98-003/DCS TR-1165, 1998.

[35] G. Antoniou and F. van Harmelen, "Web ontology language: OWL," in *Handbook on Ontology*, S. Staab and R. Studer R, Eds. Berlin, Germany: Springer, 2004, pp. 67–92.

[36] *Safety of Machinery—Risk Assessment—Part 1: Principles*, ISO 14121-1, 2007.

[37] P. R. Garvey and Z. F. Lansdowne, "Risk matrix: An approach for identifying, assessing, and ranking program risks," *Air Force J. Logist.*, vol. 22, no. 1, pp. 18–21, 1998.

[38] T. Anandan. (2013, Jun. 10). The end of separation: Man and robot as collaborative coworkers on the factory floor. *Robotics* [Online]. Available: http://www.robotics.org/content-detail.cfm/Industrial-Robotics-Featured-Article/The-End-of-Separation:-Man-and-Robot-as-Collaborative-Coworkers-on-the-Factory-Floor/content_id/4140

[39] J. Connell and M. Brady, "Generating and generalizing models of visual objects," *Artif. Intell.*, vol. 31, pp. 159–183, 1987.

[40] D. L. Wu, H. M. Zhu, X. J. Zhen, and X. M. Fan, "Tools and equipment modeling for interactive assembling operations in a virtual environment," *Int. J. Prod. Res.*, vol. 49, no. 7, pp. 1851–1876, 2011.

[41] J. Falco, J. Marvel, and R. Norcross, "Collaborative robotics: Measuring blunt force impacts on humans," in *Proc. 7th Int. Conf. Safety Ind. Autom. Syst.*, 2012, Montreal, QC, Canada, pp. 186–191.

[42] G. Borg, "Perceived exertion as an indicator of somatic stress," *Scand. J. Rehabil. Med.*, vol. 2, no. 2, pp. 92–98, 1970.

[43] G. Borg, "A category scale with ratio properties for intermodal and interindividual comparisons," in *Psychophysical Judgment and the Process of Perception*, H.-G. Geissler and P. Petzold, Eds. Berlin, Germany: VEB Deutscher Verlag der Wissenschaften, 1982, pp. 25–34.

[44] G. Borg and E. Borg, "Principles and experiments in category-ratio scaling," Dept. Psychol., Stockholm Univ., Stockholm, Sweden, Tech. Rep. 789, 1994.

[45] G. Winter, K. Schaub, and K. Landau, "Stress screening procedure for the automotive industry: Development and application of screening procedures in assembly and quality control," *Occup. Ergon.*, vol. 6, no. 2, pp. 107–120, 2006.

[46] L. Fritzsche, "Ergonomics risk assessment with digital human models in car assembly: Simulation versus real life," *Human Fact. Ergon. Manuf. Serv. Ind.*, vol. 20, no. 4, pp. 287–299, 2010.

[47] N. A. Stanton, "Hierarchical task analysis: Developments, applications and extensions," *Appl. Ergon.*, vol. 37, no. 1, pp. 55–79, 2006.

[48] R. Woodman, A. F. T. Winfield, C. Harper, and M. Fraser, "Building safer robots: Safety driven control," *Int. J. Robot. Res.*, vol. 31, no. 13, pp. 1603–1626, 2012.

[49] American National Safety Institute. ANSI Z10. Occupational health and safety management systems. 2012.

[50] P. Clemens and T. Pfitzer, “Risk assessment and control,” *Prof. Safety*, vol. 51, no. 1, pp. 41–44, 2006.
[51] M. Braglia, M. Frosolini, and R. Montanari, “Fuzzy criticality assessment model for failure modes and effects analysis,” *Int. J. Qual. Reliab. Manag*., vol. 20, no. 4, pp. 503–524, 2003.
[52] D. Gao and C. W. Wampler, “Head injury criterion,” *IEEE Robot. Autom. Mag.*, vol. 16, no. 4, pp. 71–74, Dec. 2009.

Jeremy A. Marvel (M’10) received the bachelor’s degree in computer science from Boston University, Boston, MA, USA, the master’s degree in computer science from Brandeis University, Waltham, MA, USA, and the Ph.D. degree in computer engineering from Case Western Reserve University, Cleveland, OH, USA.

He is currently a Research Scientist with the Intelligent Systems Division, National Institute of Standards and Technology (NIST), Gaithersburg, MD, USA. Following a tour at Akron University, Akron, OH, USA, as an Adjunct Professor, he went to the Institute for Research in Engineering and Applied Physics, University of Maryland, College Park, MD, USA. He then moved to NIST as a Project Leader for multiple industrial robot safety projects. His current research interests include intelligent and adaptive solutions for robot applications, with particular attention paid to human–robot collaborations, multirobot coordination, safety, perception, and automated parameter optimization. He currently leads a team of researchers and engineers in metrology efforts at NIST for collaborative robot performance.

Joe Falco received the B.S. degree in mechanical engineering from the University of Massachusetts, Amherst, MA, USA and the M.S. degree in computer science from the Johns Hopkins University, Baltimore, MD, USA.

He is currently an Engineer with the Intelligent Systems Division, National Institute of Standards and Technology (NIST), Gaithersburg, MD, USA, within the U.S. Department of Commerce. His most recent work has been performed within the NIST Next-Generation Robotics and Automation Program and is focused on robot standards and performance metrics in the areas of safety, manipulation, grasping, and human–robot interaction.

Mr. Falco is a member of the BSR/RIA R15 Robot Standards Approval Committee and ISO TC184/SC2 WG 3 working group for Industrial Robot Safety. He has received one gold and three bronze U.S. Department of Commerce medals for superior performance and technical leadership.

Ilari Marstio received the M.Sc. degree in automation technology from the Helsinki University of Technology, Espoo, Finland.

He is currently a Senior Scientist at VTT, Espoo, Finland. He gained experience on robotic systems while researching for ABB Robotics from 2004 to 2006 as a User Interface Designer and Programmer. His master’s thesis was made for the Robotic Association in Finland in 2006. In his master’s thesis, he researched machine vision and sensor based natural interfaces with robot cells. He joined VTT in 2007. Since then he has been Project Manager in several research projects. His current research interests include robot-based production systems development and human–robot interaction.