2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)
Beijing, China, November 6-9, 2018

A Risk Informed Task Planning Framework for Humanoid Robots in Hazardous Environments*

Philip Long, Murphy Wonsick and Taskin Padir [1]

Abstract— **This paper presents a generalized method to evaluate risks associated with humanoid robots executing manipulation tasks. Risks are defined as the product of probability, the likelihood of an event occurring, and severity, the resulting magnitude of harm should it occur. Rather than try to reduce the probability of failure events to zero, the objective of this work is to allow an experienced operator/supervisor to define if some failures are worse than others. In doing so, this allows the operator to judge whether high risk motions are necessary for the task at hand. Utilizing NASA's humanoid robot Valkyrie, our framework is demonstrated in both simulation and on the physical robot, with a pick and place task. We show that our method is capable of predicting failures for given motions based on their calculated risk.**

I. INTRODUCTION

Nuclear facilities worldwide are approaching the end of their lifespan and must be decontaminated and decommissioned in a safe and secure way [1]. These tasks are dull, dangerous and dirty and therefore are ideal candidates for robots. However, due to their hazardous nature, it is unlikely that such tasks can or should be fully automated; in fact, it is desirable to exploit the operator's experience. Thus, supervised autonomy is envisaged where the operator selects a task level objective and the robot obtains a complete planning solution. With this in mind, it is extremely important that the risks associated with different solutions are calculated before execution and that these risks consider sensors, actuator and robot state uncertainties along with severity associated with failure.

Decommissioning tasks involve a wide range of locomotion and manipulation actions meaning that no one robotic solution is optimal. Recently humanoid robots have been proposed as a general purpose solution [2]. However, even with recent efforts [3], [4], deploying humanoid robots in unstructured environments with reliable autonomy remains a significant challenge. The core problem is the inability to fully model and react to unstructured and dynamic environments. In spite of sensor improvement, certain errors are unavoidable and frequently can only be mitigated by repeated calibration. Within an industrial process slight degradations in performance can be measured and thus corrected before failures occur. In safety critical applications, for instance, handling high-consequence materials, it is imperative that the system achieves its goal or at the very least fails in a safe and/or operational manner.

Risk associated with an event is defined from two basic concepts: the likelihood of the event occurring (probability) and the resulting magnitude of harm (severity), also known as the expected loss. Humans intuitively integrate risk into the decision making process [5], [6], [7], even when planning motor behaviors [8]. Decisions are based on confidence in expected results, either successes or failures. While the probability of an event can be calculated within defined bounds, the expected loss is a subjective measure. One definition is the inverse of payoff [9] in the case of success, but this metric can lead to incorrect and counter-intuitive evaluations [8]. Risk propensity [10] can vary with time, problem definition and prior experience [5] thus evaluation of risk cannot be separated from the decision making agent [11].

Risk is analogous to probability if the severity is constant, i.e. all failures are equally bad. In robotics this paradigm is often applied to autonomous vehicles, for example [12], [13] where risk is directly defined as the probability of collision. In this case Markov Decision Processes [14], [15] can be efficiently used to solve the selection process. However, most human drivers recognize a head-on collision is more dangerous than two vehicles scraping against each other, thus making the concept of severity an integral part of risk planning. Indeed, recent work [16], [17] has recognized the limitations of considering risk and probability in robotics to be equivalent. To embed the idea of severity into risk management, a known safety criteria can be used, for instance international norms (i.e. *ISO 13482:2014*). In [18], the authors propose a robot control scheme that evaluates the severity of a potential collision to change its controller. Likewise in [19], [20], the objective is not to reduce the probability of failure but rather ensure that the magnitude of harm incurred during this failure is reduced. This is applicable where severity is quantifiable, for example robot/human collision [21], [22]. In other cases, expert or empirical knowledge is used, for instance in [23] where a global measure of severity is empirically defined and fused with probability of failure to allow for risk minimization.

In this paper, we extend upon our work first presented in [24] by two contributions (1) a generalized method to evaluate risk for a failure event and (2) the introduction of the concept of severity based on robot state and operator experience. In our target scenario, displacing nuclear material in an unknown dynamic environment, it is unreasonable to

* This material is based upon work supported by the Department of Energy under Award Number DE-EM0004482, by the National Aeronautics and Space Administration under Grant No. NNX16AC48A issued through the Science and Technology Mission Directorate and by the National Science Foundation under Award No. 1451427.

[1]Philip Long, Murphy Wonsick and Taskin Padir are with Department of Electrical and Computer Engineering, Northeastern University, Boston, MA 02115, USA p.long(t.padir)@northeastern.edu, wonsick.m@husky.neu.edu,

978-1-5386-7282-2/18/$31.00 ©2018 IEEE

believe the robot can/should compute severity indexes for everything that can go wrong. Instead, the likelihood of this occurrence should be communicated to the teleoperator who then can make an informed decision, based on their past experiences dealing with hazardous material. The first step, which we do not treat in this paper, is to obtain the possible failure events and their associated severity, either using safety guidelines or empirical evidence. The selected severity for the failure events is then combined with the probability to obtain the risk for a sequence of tasks. The paper is organized as follows. Section II defines the problem and a new risk function for a generic action is introduced. The experimental setup is explained in Section III followed by the results in simulation and in experimentation in Section IV. Finally in Section V conclusions are drawn and future work discussed.

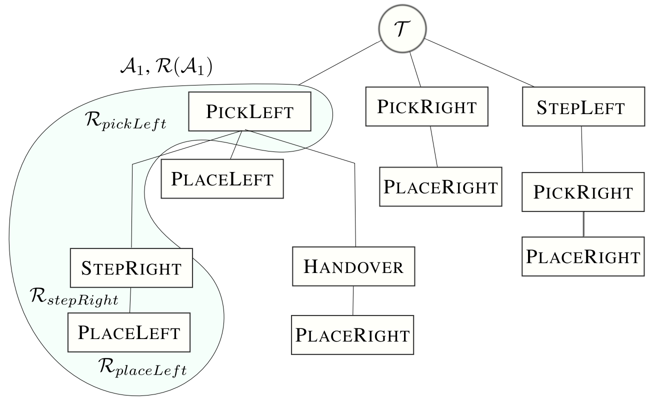

Fig. 1: Compositions for the humanoid pick and place task, with an example composition highlighted.

II. Methodology

A. Problem Statement

In the following, the notation defined in [24] is recalled. A high level task, denoted $\mathcal{T}$ can be achieved by composing a sequence of robot behaviors. This is known as compositional robot autonomy. The i^{th} composition, denoted $\mathcal{A}_i$, comprises a set of actions, i.e. $\mathcal{A}_i = \{\mathcal{M}_1, \mathcal{M}_2 \ldots \mathcal{M}_j\}$. An action, $\mathcal{M}_j$, drawn from a feasible action set, is typically represented by a motion trajectory, $Traj(\mathcal{M}_j) = \mathbf{q}(t),\ t \in [t_0,\ t_f]$, where $\mathbf{q}(t)$ is the vector of joint position variables that define the robot configuration. Each composition that would successfully complete the task is contained in the set of feasible compositions, $\mathcal{A}_i \in \mathbb{A}$. Thus in the absence of errors and in a known static environment, the decision maker selects between compositions using a criterion, such as time or energy minimization, or intuition. In reality, sensor, actuation, and controller errors are unavoidable. Therefore, our objective in this work, is to obtain a risk measure that permits the operator to make an informed composition selection, based on probabilistic bounds and severity costs. In the following section, our proposed method for evaluating risks is outlined. In summary,

- Actions are comprised of motions and are defined with respect to robot capability. Examples for a humanoid robot can be: PICKLEFT: *Grasp an object with the left hand*, PLACELEFT: *Move left hand to target location and release object*, SCANSCENE: *Look around for target object*, SIDESTEPLEFT: - *Take one side-step left*, STEPMULTIPLE - *Walk to a desired base pose with multiple steps*, etc.
- A risk assessment listing failure events is carried out for members of the feasible action set. An example event for PLACELEFT: is COLLISIONLEFTARMWALL: *the left arm colliding with a wall*. To calculate the risk associated with an event two function must be defined. Firstly, a function that returns the event's probability during PLACELEFT based on robot state and the perceived environment. Secondly, a function that evaluates the cost, based on robot state and operator experience. For instance, an operator may consider a collision with a wall more serious than a collision with a movable object.
- The compositions are formed by stitching feasible actions together. For example LOCATEOBJECT $\rightarrow$ PICKLEFT $\rightarrow$ PLACELEFT or alternatively STEPLEFT $\rightarrow$ PICKLEFT $\rightarrow$ HANDOVER $\rightarrow$ PLACERIGHT can be used to execute a pick and place task. The risk for the action composition is obtained by the sum of the action's risks. The high level decision maker can select between compositions using the risk metric. Furthermore, the risk can be monitored on-line in case of large variations due to controller errors and/or environmental changes.

B. Action Risk Assessment

The risk assessment tool is used to define events, denoted as $(\mathcal{E}_{ij})_{i=1}^{k}$, that may lead to failure during the execution of this action, where i denotes a failure event, j denotes the action, and k is the total number of failure events. Hence $\mathcal{E}_{ij}$ denotes one of k possible failure events associated with action j.

In order to evaluate the risk $\mathcal{R}_{ij}$, associated with event $\mathcal{E}_{ij}$, functions are defined that calculate the probability of $\mathcal{E}_{ij}$ occurring and the magnitude of harm (also known as cost and called throughout this paper as severity) should it occur. The probability and severity, denoted $\mathcal{P}_{ij}(t)$ and $\mathcal{S}_{ij}(t)$, are evaluated for the duration of the action. A time varying risk metric, for event $\mathcal{E}_{ij}$ is given as

$$\mathcal{R}_{ij}(t) = \left(\mathcal{P}_{ij}(t) \times \mathcal{S}_{ij}(t)\right). \tag{1}$$

Hence, the unit of risk is defined by the measure of severity, the most logical being monetary [16], i.e. the expected cost of failure. Indeed, the risk is equal to the cost if the action is inevitable. It should be noted that events themselves are not necessarily *elementary* (atomic) and the system may be capable of reacting in order to reduce the cost. In this case (1) becomes

$$\mathcal{R}(t)_{ij} = \left(\mathcal{P}_{ij}(t) \times \mathcal{R}^{r}_{ij}(t)\right), \tag{2}$$

where the risk associated with the recovery action is $\mathcal{R}^r_{ij}(t) \leq \mathcal{S}_{ij}(t)$ and is expressed as

$$\mathcal{R}^r_{ij}(t) = \mathcal{P}^r_{ij} \times \mathcal{S}^r_{ij} + (1 - \mathcal{P}^r_{ij}) \times \mathcal{S}_{ij}(t). \tag{3}$$

$\mathcal{P}^r_{ij}$ is the probability the reactive action succeeds and $\mathcal{S}^r_{ij}$ is the reduced cost incurred thanks to the reactive behavior. This type of recursive analysis in risk assessment is known as event trees, [25]. In this paper, we select the maximum risk, i.e. worse case scenario for each event. Additionally, we treat the events as independent, hence the risk $\mathcal{R}_j$ associated with action $\mathcal{M}_j$ can be obtained as

$$\mathcal{R}_j = \sum_{i=1}^{k} \left(\max(\mathcal{R}_{ij}(t)\right). \tag{4}$$

C. Risk Aware Planning Framework

For a desired task, $\mathcal{T}$ and a robot configuration $\mathbf{q}$, we assume the existence of $\mathbb{A}$, which is the set of feasible action compositions to accomplish the task. The motion generation for each action is calculated by an optimization based motion planner described in [3], [4], resulting in a set of robot states that satisfy the task constraints. A cubic spline interpolation is implemented to obtain the desired trajectory, $Traj(\mathcal{M}_j) = \mathbf{q}(t),\ t \in [t_0,\ t_f]$. The procedure described in Section II-B then is used to obtain the risk for action $Traj(\mathcal{M}_j)$, assuming the information used by the motion planner is perfectly representative of the real world. The risk for composition $\mathcal{A}_i$ is given as

$$\mathcal{R}(\mathcal{A}_i) = \sum_{j=1}^{M} \mathcal{R}_j, \tag{5}$$

where M denotes the total number of motions in the given action. The composition with the minimum risk from the set of feasible action compositions may then be chosen. Alternatively, the risk associated with each composition can be compared to a reward function such as total time.

III. Experimental Setup

In order to demonstrate the concept, a simple pick and place task is executed by NASA's R5 humanoid robot, Valkyrie. Several different compositions can complete this task, a subset can be seen in Fig. 1. In this section, three different failure events are introduced.

A. Collision Event

A collision event is defined as an undesired collision between a robot's link and the environment or another of the links (self-collision). The motion planner ensures that discrete robot state's are collision free. The intermediate volume is verified using the swept volume method [26]. Nevertheless, unforeseen collisions can occur due to errors in the environmental model, mesh volume estimations, trajectory interpolation/smoothing approximations and controller errors due to imprecise robot state. To model the probability of collision between a link $\mathcal{L}$ and an object $\mathcal{O}$, the minimum distance $d_{lo}(t)$ is defined for the duration of $Traj(\mathcal{M}_j)$. This is defined as the minimum distance that $\mathcal{L}$ can move in any direction before its mesh intersects with the mesh representing object $\mathcal{O}$. A normal distribution is used with the probability density function (PDF)

$$f(x \mid d_{lo}(t), \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}}\, e^{-\dfrac{(x - d_{lo}(t))^2}{2\sigma^2}}, \tag{6}$$

Then the probability of collision, i.e. $d_{lo}(t) \leq 0$, at an instant, is denoted t $\mathcal{P}_{cj}(t)$ and is defined by the cumulative distribution function (CDF)

$$\mathcal{P}cj(t) = \int_{-\infty}^{0} f(x \mid d_{lo}(t), \sigma^2)dx, \tag{7}$$

where σ denotes the standard deviation of the PDF. In this work, we consider self collisions ($\sigma = 0.025$) and possible collision between all objects and the robot's hands ($\sigma = 0.4$).

The severity of a collision depends on the mass and the velocity of the link in question. Therefore, severity of a potential collision between $\mathcal{L}$ and an object $\mathcal{O}$ at any instant during the trajectory is defined as using the link's kinetic energy:

$$\mathcal{S}cj(t) = k_{lo}\left(m_l\ \mathbf{v}_l^T \mathbf{v}_l + \omega_\mathbf{l}{}^T \mathbf{I}_l \omega_\mathbf{l}\right), \tag{8}$$

where m_l, $\mathbf{I}_l$ denote the mass and inertia tensor of $\mathcal{L}$ and $\mathbf{v}_l$ and $\omega_\mathbf{l}$ denotes its translational and angular velocity. k_{lo} is a user defined variable that is chosen based on $\mathcal{O}$, with units J^{-1}, where J denotes a *Joule*.

B. Falling Event

Falling over is a critical event, which, for humanoid robots, typically signals the end of any operation. To prevent falling, for statically stable humanoid such as Valkyrie, it is necessary to maintain the projection of the center of mass within the support polygon. Motions are planned such that this constraint is respected. However, errors may occur due to incorrect robot inertial parameters and/or object mass properties. In addition, the exact bounds of the support polygon may vary with terrain, for instance on uneven or soft terrain. The minimum distance from the projected center of mass to an edge of the support polygon formed by the robot's feet, denoted $d_z(t)$, is calculated for the duration of $Traj(\mathcal{M}_j)$. The probability of falling, i.e. $d_z(t) \leq 0$, denoted $\mathcal{P}zj(t)$ is given as

$$\mathcal{P}zj(t) = \int_{-\infty}^{0} f(x \mid d_z(t), \sigma^2)dx, \tag{9}$$

with $\sigma = 0.08$ in general operation and $\sigma = 0.16$ when carrying an object. The severity of a fall depends on robot configuration, surroundings and durability of respective link's/components. In this paper, severity of a fall event is defined as

$$\mathcal{S}zj(t) = k_z\left(p_z^{com}\right), \tag{10}$$

where p^{com} denotes the center of mass and the subscript $_z$ denotes its z component. The user selected gain term is defined in this case as $k_z = 10$.

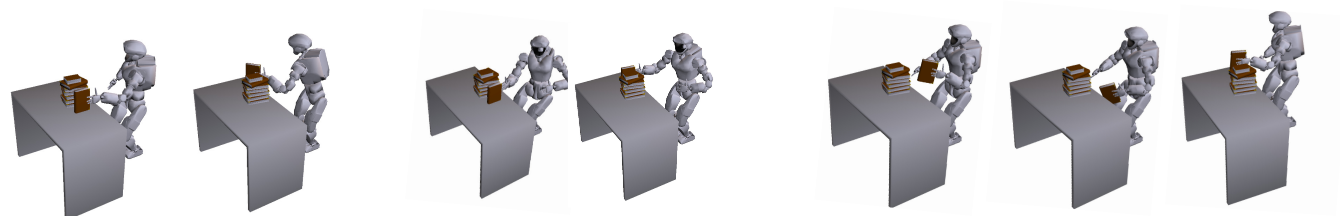

(a) $\mathcal{A}_1$: PICKLEFT → PLACELEFT (b) $\mathcal{A}_2$ PICKRIGHT → PLACERIGHT (c) $\mathcal{A}_3$ PICKLEFT → STEPRIGHT → PLACELEFT

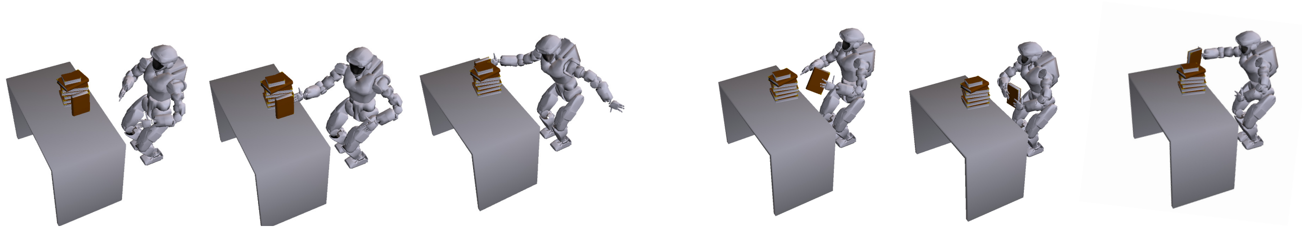

(d) $\mathcal{A}_4$ STEPLEFT → PICKRIGHT → PLACERIGHT (e) $\mathcal{A}_5$ PICKLEFT → HANDOVER → PLACERIGHT

Fig. 2: Task compositions broken into individual actions.

C. Torque Limits Violation Event

A torque limit violation occurs when the controller set-point exceeds torque limits. In this case, the system sends the maximum torque yet will be unable to achieve the desired action. Consequently, the trajectory may not be collision free and may lead to instabilities. This may occur due to poor estimation of the grasped object and unmodeled dynamic disturbances. The probability of the i^{th} joint exceeding a torque limit, τ_i^{max}, is given as

$$\mathcal{P}^i tj(t) = \int_{-\infty}^{0} f(x \mid \|\tau_i^{max} - \tau_i(t)\|, \sigma^2) dx, \tag{11}$$

where τ_i is the estimated joint torque at time t. We consider the severity of a violated torque to be proportional to the magnitude of the violation.

D. Calculation of Total Risk

For each event the risk value is obtained from (2). The maximum risk value across the whole trajectory is taken as the risk value for the action $Traj(\mathcal{M}_j)$. For example, composition $\mathcal{A}_1$ that consists of actions {PICKLEFT, PLACELEFT}, the total risk is defined by (5). In our pick and place example, the risk is a scalar value composed of collision risk, joint limit violation and falling risk for each action, bounded by $0 \leq \mathcal{R}(\mathcal{A}_i) \leq \mathcal{S}_{failure}$, where $\mathcal{S}_{failure}$ denotes the cost of a critical event defined by the user, for instance a high velocity collision or falling event.

IV. EXPERIMENTS

The objective is to transfer an object on a table from a pick position to a place position with an obstacle between the two points. In order to calculate the motion trajectories, the mesh and mass properties of the transported object are given. For each composition, the motion planner generates a priori safe plans. The initial configuration for a subsequent action is assumed to be the desired final configuration of the previous action.

A. Monte Carlo Simulations

The planner environment is shown in Fig. 2. For each action the motion planner generates a feasible trajectory. The risk associated for each action from the planner data is computed and presented in Fig. 3a. The risk associated with each composition as calculated by (5) is shown in Fig. 3b. The highest risks are generally those associated with falling, this is due to the high cost of failure, in spite of a low probability. The highest predicted risks are for $\mathcal{A}_3$ and $\mathcal{A}_5$, in the former case due to the possibility of falling while transporting the object, in the latter case due to the possibility of a high speed collision with the table[1].

The planned motions are executed by the dynamic simulator 100 times during which Gaussian noise is added to the object's position ($\sigma = 0.05[m]$), object's orientation normal to the ground plane ($\sigma = 0.01[rad]$), the book's mass ($\sigma = 1.5[kg]$) and the robot's center of mass ($\sigma = 0.02[m]$). In addition, the planned joint trajectory is not followed exactly by the dynamic simulator due to controller errors and interpolations. The failure events during each execution are recorded and presented for each composition in Fig. 4. It shows that while $\mathcal{A}_1$, $\mathcal{A}_2$ and $\mathcal{A}_4$ suffer from many failures these failures are of a very low severity, typically collisions at low speeds or contact before grasping. In contrast, $\mathcal{A}_3$ suffers from high cost failure events and thus should be avoided.

B. Experimental Validation

To experimentally validate our framework on physical hardware, the composition with the lowest predicted risk, $\mathcal{A}_1$, is compared with the composition of highest predicted risk, $\mathcal{A}_3$. The experimental setup is shown in Fig. 5. Each composition is repeated five times during which the position of the table and obstacle is slightly altered. Moreover, for tests 3, 4 and 5, 0.372 kg, 0.720 kg and 1.2 kg are added to the book to perturb the system.

[1]The higher velocities attained during $\mathcal{A}_5$ are generated by the motion planner to ensure balance and synchronized motion

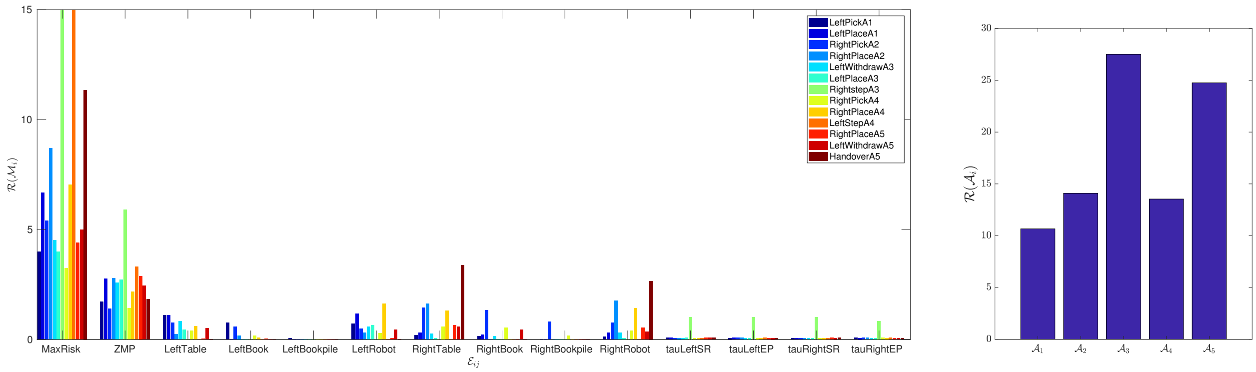

(a) Predicted risks using (4) versus failure event associated for each action, where *LeftBook* denotes a collision between the *left arm* and the *book*, *tauLeft* denotes a torque limit for the left arm and *SR* and *EP* denotes the shoulder roll and elbow pitch respectively.

(b) Evaluation and comparison the total risk i.e., (5), for each composition for task $\mathcal{T}$.

Fig. 3: Predicted risks based on planner output.

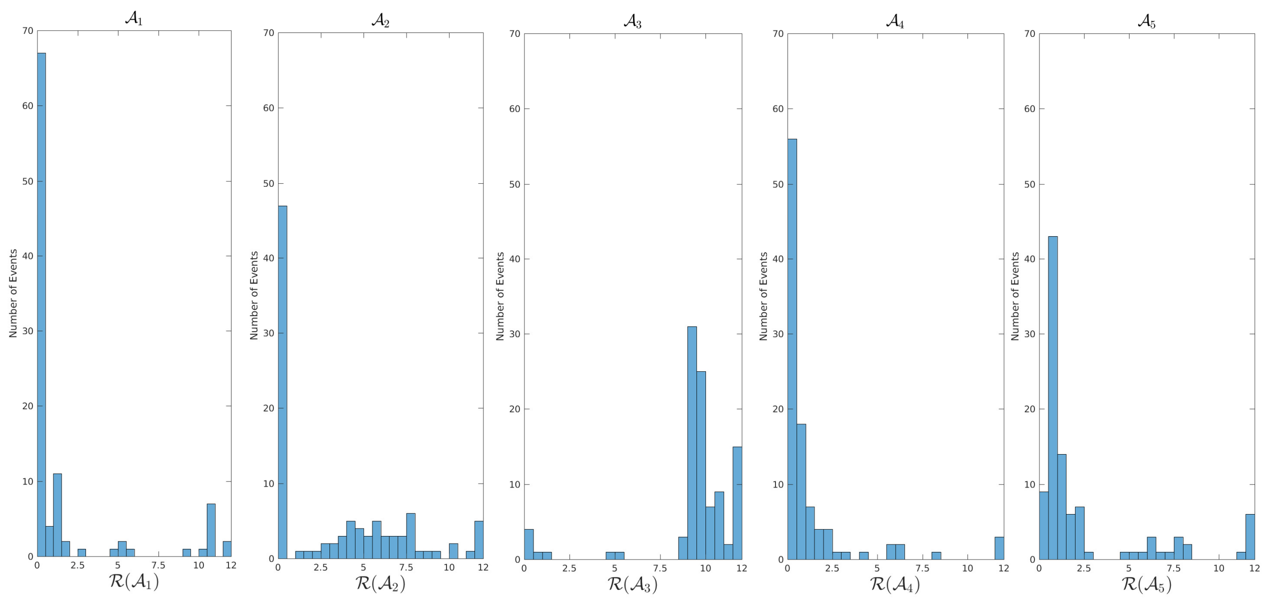

Fig. 4: Histogram of recorded events over 100 trials, the $x-$ axis shows the severity of the event and the $y-$ axis shows the number of occurrences. While $\mathcal{A}_3$ has less overall failure events, the resulting severity is much higher leading to a *riskier* overall composition. In contrast, $\mathcal{A}_1$, has a high number of low severity failure events such as low speed collisions.

This study presents a relatively small sample size comprising five experiments for $\mathcal{A}_1$ and $\mathcal{A}_3$, nevertheless there is a clear performance disparity between the two compositions. In all five experiments, $\mathcal{A}_1$ is successful in transferring the book from one side of the table to another, in two cases without any failure event. In the other three cases slight collisions were observed between the obstacle and the grasped object during the transfer and also during the placing action with the back of the left palm. These collisions occurred at low velocities and had no significant impact on the task. In contrast, $\mathcal{A}_3$ experienced no collision errors when placing the object, though occasionally minor collision between the book and box occurred during the withdrawal phase. However a critical error occurred as the robot fell during test 4 with the addition of 0.72 kg on the grasping hand. Additionally, the robot exceeded a torque limit towards the end of the composition in test 2. Furthermore, although not quantitatively evaluated the operators noted stark differences between the compositions including increased oscillations due to corrective controller action and a above average pelvis height during $\mathcal{A}_3$.

V. CONCLUSIONS

In this paper, a risk informed task planning framework is presented. that allows an operator to associate risks with actions by defining generic failure events. Each event is defined with both a probability and severity function. The resulting risk metric is a single scalar value of monetary unit

Fig. 5: Experimental setup for compositions $\mathcal{A}_1$ and $\mathcal{A}_3$. Top row, $\mathcal{A}_1$: PICKLEFT $\rightarrow$ PLACELEFT, Bottom row $\mathcal{A}_3$: STEPRIGHT $\rightarrow$ PLACELEFT

which should be associated with a real world cost of failure. We have shown by simulation how this risk metric is a valid indicator for predicting high severity failures. Moreover, the experimental work, albeit limited, appears to validate these metrics.

There are several avenues of future work. The enumeration of failure events is a time consuming process and in this work we have focused on three prevalent failure events, collision, falling and torque limit avoidance. In future work the automatic identification of a failure event's probability and severity will be explored. To do so, we aim to analyze the system's behavior when directly controlled by an experienced operator. Finally, the possibility of reducing total cost by allowing reactive control will be investigated.

REFERENCES

[1] D. W. Seward and M. J. Bakari, "The use of robotics and automation in nuclear decommissioning," in *22nd Int. Symposium on Automation and Robotics in Construction*, pp. 11–14, 2005.

[2] T. Padir, H. Yanco, and R. Platt, "Towards cooperative control of humanoid robots for handling high-consequence materials in gloveboxes," in *WM'17*, WM Symposia, Inc, 2017.

[3] M. DeDonato, F. Polido, K. Knoedler, B. P. Babu, N. Banerjee, C. P. Bove, X. Cui, R. Du, P. Franklin, J. P. Graff, *et al.*, "Team wpi-cmu: Achieving reliable humanoid behavior in the darpa robotics challenge," *Journal of Field Robotics*, vol. 34, no. 2, pp. 381–399, 2017.

[4] N. Banerjee, X. Long, R. Du, F. Polido, S. Feng, C. G. Atkeson, M. Gennert, and T. Padir, "Human-supervised control of the atlas humanoid robot for traversing doors," in *IEEE-RAS 15th Int. Conf. on Humanoid Robots (Humanoids)*, pp. 722–729, IEEE, 2015.

[5] L. L. Lopes, "Between hope and fear: The psychology of risk," in *Advances in experimental social psychology*, vol. 20, pp. 255–295, Elsevier, 1987.

[6] S. M. Tom, C. R. Fox, C. Trepel, and R. A. Poldrack, "The neural basis of loss aversion in decision-making under risk," *Science*, vol. 315, no. 5811, pp. 515–518, 2007.

[7] M. Hsu, M. Bhatt, R. Adolphs, D. Tranel, and C. F. Camerer, "Neural systems responding to degrees of uncertainty in human decision-making," *Science*, vol. 310, no. 5754, pp. 1680–1683, 2005.

[8] D. A. Braun, A. J. Nagengast, and D. Wolpert, "Risk-sensitivity in sensorimotor control," *Frontiers in human neuroscience*, vol. 5, p. 1, 2011.

[9] T. D. Sanger, "Risk-aware control," *Neural computation*, vol. 26, no. 12, pp. 2669–2691, 2014.

[10] S. B. Sitkin and L. R. Weingart, "Determinants of risky decision-making behavior: A test of the mediating role of risk perceptions and propensity," *Academy of management Journal*, vol. 38, no. 6, pp. 1573–1592, 1995.

[11] D. Bernoulli, "Exposition of a new theory on the measurement of risk," *Econometrica*, vol. 22, no. 1, pp. 23–36, 1954.

[12] S. Noh and K. An, "Risk assessment for automatic lane change maneuvers on highways," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 247–254, IEEE, 2017.

[13] D. Althoff, J. J. Kuffner, D. Wollherr, and M. Buss, "Safety assessment of robot trajectories for navigation in uncertain and dynamic environments," *Autonomous Robots*, vol. 32, no. 3, pp. 285–302, 2012.

[14] S. Feyzabadi and S. Carpin, "Risk-aware path planning using hirerachical constrained markov decision processes," in *2014 IEEE Int. Conf. on Automation Science and Engineering (CASE)*, pp. 297–303, IEEE, 2014.

[15] D. Mehta, G. Ferrer, and E. Olson, "Fast discovery of influential outcomes for risk-aware mpdm," in *2017 IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 6210–6216, IEEE, 2017.

[16] A. Majumdar and M. Pavone, "How should a robot assess risk? towards an axiomatic theory of risk in robotics," *arXiv preprint arXiv:1710.11040*, 2017.

[17] S. Singh, J. Lacotte, A. Majumdar, and M. Pavone, "Risk-sensitive inverse reinforcement learning via semi-and non-parametric methods," *arXiv preprint arXiv:1711.10055*, 2017.

[18] M. Ragaglia, L. Bascetta, P. Rocco, and A. M. Zanchettin, "Integration of perception, control and injury knowledge for safe human-robot interaction," in *2014 IEEE Int. Conf. on Robotics and Automation (ICRA),*, pp. 1196–1202, IEEE, 2014.

[19] D. Kulić and E. Croft, "Pre-collision safety strategies for human-robot interaction," *Autonomous Robots*, vol. 22, no. 2, pp. 149–164, 2007.

[20] P. Long, C. Chevallereau, D. Chablat, and A. Girin, "An industrial security system for human-robot coexistence," *Industrial Robot: An International Journal*, 2017.

[21] S. Haddadin, A. Albu-Schäffer, and G. Hirzinger, "Safety evaluation of physical human-robot interaction via crash-testing.," in *Robotics: Science and Systems*, vol. 3, pp. 217–224, 2007.

[22] S. Oberer-Treitz, A. Puzik, and A. Verl, "Measuring the collision potential of industrial robots," in *Robotics (ISR), 2010 41st Int. Symposium on and 2010 6th German Conference on Robotics (ROBOTIK)*, pp. 1–7, VDE, 2010.

[23] M. Nagel, G. Schmidt, G. Schnuetgen, and W. A. Kalender, "Risk management for a robot-assisted needle positioning system for interventional radiology," in *International Congress Series*, vol. 1268, pp. 549–554, Elsevier, 2004.

[24] X. Long, P. Long, and T. Padır, "Compositional autonomy for humanoid robots with risk-aware decision-making," in *IEEE-RAS 17th Int. Conf. on Humanoid Robots (Humanoids)*, IEEE, 2017.

[25] B. S. Dhillon and A. Fashandi, "Safety and reliability assessment techniques in robotics," *Robotica*, vol. 15, no. 6, pp. 701–708, 1997.

[26] G. v. d. Bergen, "A fast and robust gjk implementation for collision detection of convex objects," *Journal of graphics tools*, vol. 4, no. 2, pp. 7–25, 1999.