Robot Vitals and Robot Health: An Intuitive Approach to Quantifying and Communicating Predicted Robot Performance Degradation in Human-Robot Teams

Aniketh Ramesh
Axr1050@student.bham.ac.uk
Extreme Robotics Laboratory,
University of Birmingham
Birmingham, England

Manolis Chiou
M.Chiou@bham.ac.uk
Extreme Robotics Laboratory,
University of Birmingham
Birmingham, England

Rustam Stolkin
R.Stolkin@bham.ac.uk
Extreme Robotics Laboratory,
University of Birmingham
Birmingham, England

ABSTRACT

In this work we introduce the concept of Robot Vitals and propose a framework for systematically quantifying the performance degradation experienced by a robot. A performance indicator or parameter can be called a Robot Vital if it can be consistently correlated with a robot's failure, faulty behaviour or malfunction. Robot Health can be quantified as the entropy of observing a set of vitals. Robot vitals and Robot health are intuitive ways to quantify a robot's ability to function autonomously. Robots programmed with multiple levels of autonomy (LOA) do not scale well when a human is in charge of regulating the LOAs. Artificial agents can use robot vitals to assist operators with LOA switches that fix field-repairable non-terminal performance degradation in mobile robots. Robot health can also be used to aid a tele-operator's judgement and promote explainability (e.g. via visual cues), thereby reducing operator workload while promoting trust and engagement with the system. In multi-robot systems, agents can use robot health to prioritise robots most in need of tele-operator attention. The vitals proposed in this paper are: rate of change of signal strength; sliding window average of difference between expected robot velocity and actual velocity; robot acceleration; rate of increase in area coverage and localisation error.

CCS CONCEPTS

• **Human-centered computing** → *Interaction design*; • **Computer systems organization** → **Robotic autonomy**; Reliability.

KEYWORDS

variable autonomy, mixed initiative, robot health, predictive analytics

ACM Reference Format:
Aniketh Ramesh, Manolis Chiou, and Rustam Stolkin. 2021. Robot Vitals and Robot Health: An Intuitive Approach to Quantifying and Communicating Predicted Robot Performance Degradation in Human-Robot Teams. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21 Companion), March 8–11, 2021, Boulder, CO, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3434074.3447181

HRI '21 Companion, March 8–11, 2021, Boulder, CO, USA

ACM ISBN 978-1-4503-8290-8/21/03…$15.00
https://doi.org/10.1145/3434074.3447181

1 INTRODUCTION

Variable Autonomy (VA) robotic systems are capable of exhibiting multiple behaviours or autonomous capabilities, known as Levels of Autonomy (LOA) [21]. To carry out a diverse range of tasks, robots can be programmed with multiple LOAs. During run time, the LOA best suited for the task at hand can be chosen on-the-fly. Alternatively, programming tele-operation or manual control as one of the LOA is useful when the autonomous abilities of the robot are compromised. For example, if a robot gets stuck while navigating autonomously, the LOA can be switched to tele-operation. The robot can then be tele-operated and switched back to autonomous navigation when it is capable of functioning autonomously again. When a human tele-operator is in charge of carrying out such LOA switches during run time it is called a Human Initiative (HI) system. As the number of robots used increases or the task gets increasingly complex, tele-operators experience high cognitive workload [4, 14]. This may result in the tele-operators making sub-optimal LOA choices for robot(s), thereby reducing overall task performance[5, 10, 22]. Providing operators in HI systems with visual aids that suggest the optimal LOA choice or which robot needs manual operation or help, can reduce the overall cognitive load. Alternatively, allowing an artificial agent to assist the tele-operator with carrying out LOA switches will reduce the cognitive demand on the tele-operator, while improving overall task performance. A system where both the agent and the tele-operator collaborate to carry out LOA switches is called a Mixed Initiative (MI) [12] system. Our research deals with the design of artificial agents that can assist the tele-operator with carrying out LOA switches.

To initiate LOA switches, an artificial agent would require - 1) A set of parameters used to detect performance degradation experienced by robots in extreme environments, 2) A metric binding such parameters to quantify the severity of the performance degradation. In our study, we refer to the former as robot vitals and the latter as robot health. The use of robot vitals to carry out LOA switches will facilitate explainability of the agent's decisions and provide an intuitive way of communicating the predicted performance of the robot.

2 RELATED WORK

Various studies [3, 8, 18] have catalogued the different ways in which robots fail or malfunction in real and experimental settings.

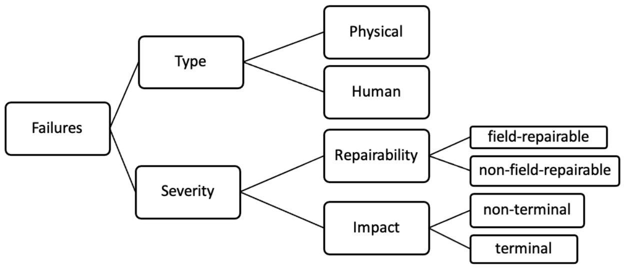

Figure 1: Taxonomy of mobile robot failures used in analysis adapted from [3]

Carlson et al. proposed the taxonomy[3] adapted in figure 1 for classifying how UGVs physically failed in the field. Based on the repairability of a UGV's physical failures they were classified as field-repairable and non field-repairable failures. Similarly based on the impact of their failures they are classified as terminal and non-terminal failures. Some examples of non-terminal field-repairable issues include sub-optimal path planning, immobility due to camera or laser scanner noise, camera occlusion etc. In our study, performance degradation refers to field-repairable non-terminal failures of mobile UGVs. In MI systems, an artificial agent can assist the tele-operator by detecting when a robot is experiencing performance degradation and switching to the LOA best suited to mitigate it.

Designing LOA regulating agents has received considerable attention in existing literature[15, 16, 19]. However, there is little agreement on parameters such agents would use, and a metric that would quantify the extent of performance degradation experienced by the robot during run time. One of the earliest studies on quantifying robot performance degradation for LOA switching was carried out by Valero et al. [24]. The authors proposed an agent based LOA switching system that used four performance indicators - the robot's mean velocity after 4 and 15 seconds, displacement in the last 2 seconds and area explored. The trigger values of these indicators were tuned during experimentation. Similarly, a pilot study of Mixed Initiative LOA switching for a single robot[7] used the difference between the actual velocity of the robot and an expert planner. A fuzzy rule base then uses this value along with the robot's actual velocity to trigger LOA switches. Mendoza et al. [17] estimated faulty behaviour in the 'CoBots' platform by calculating the robot's location independently using different sensors. The robot compares localisation information generated individually from wheel encoders, laser range finders and infrared scanners to check for mismatches. The presence and magnitude of mismatches are used to detect faults. Additionally, the robot's knowledge base generates an estimated time to task completion for each task allocated to it. When the task takes longer than the estimated time to complete, the robot flags a possible fault. In the absence of multiple sensors that can give location information, the robot's Lidar scan data can be compared with the occupancy grid map to calculate a reliability estimate[1].

3 ROBOT VITALS

In our work, we define robot vitals as a set of parameters that indicate the status of the robot and it's ability to continue operating autonomously during performance degradation. When a robot is experiencing little to no performance degradation, the robot vitals should be within a fixed window of values. The window of values may vary based on the robotic hardware or environmental conditions. A change in one or more vitals can give clues about the exact nature of performance degradation the robot is facing. Each vital is ideally symptomatic of at least one performance degrading factor. For the purpose of this study, we assume that any degradation in performance is non terminal and can be repaired by timely switching of LOAs. Therefore, by monitoring a robot's vitals it is possible to monitor the robot's performance and predict whether it needs tele-operator assistance.

The existing literature usually does not differentiate between metrics for a robot's performance and metrics for overall performance of the robot(s) in the task. Task performance metrics assess how effective robots are in carrying out the assigned task. For example in an urban search and rescue mission undertaken by a multi-robot system (MRS), some of the task performance metrics are [15] - targets/victims found, time taken to cover the area, workload distribution, rate of area coverage, number of repeated visits to any given point on the map, tele-operator cognitive workload, etc. While these metrics are useful to evaluate how well the MRS carried out the task, they give no information about how each robot performed in the task. However, metrics like robot velocity, robot information entropy [27], sensor noise, number of collisions are useful to understand how well each robot performed during the task. Consider the example of a game of soccer between two teams. A team's performance can be determined using metrics like - number of goals scored, ball possession, number of yellow and red cards issued, shots on target, fouls etc. However, some players may feel unhealthy during the game and require medical aid or substitution. This can be monitored by measuring each player's pulse, respiration rate, body temperature and gait speed [23] during the game. Analogously, robot vitals along with overall task performance indicators can communicate the status of each robot to the tele-operator in an easy and intuitive way. Artificial agents can use this information to trigger LOA switches in a manner that promotes explainability.

3.1 Properties

We propose that robot vitals used for mobile UGVs should have the following properties:

(1) **Task Agnostic**: Any chosen robot vital should work irrespective of the task given to the robot(s). It should be robust to on-the-fly changes in task or mission goals.
(2) **Correlation**: Every robot vital should be demonstrably correlated to one or many forms of performance degradation that the robot can experience. To establish this correlation, the vitals may or may not require pre-processing like noise removal, filtering or smoothing etc.
(3) **Performance Metric**: It should be possible to derive a scalar measure of any robot's health using the robot vitals so as to quantitatively show that the health of one robot is more 'critical' than the other. That is, at any given time the robot health should be proportional to the performance degradation experienced by the robot.

3.2 List of robot vitals

Data from experiments carried out by [6, 7] was analysed for identifying an initial set of meaningful robot vitals. In [6] the authors study HI LOA switching for a simulated navigation task using a robot with 2 levels of autonomy - Joystick Controlled Tele-operation and Autonomous Control (robot navigates autonomously towards way points given by the tele-operator). The robot's performance is degraded by introducing Gaussian white noise briefly during the simulation. Similarly, the cognitive workload of the tele-operator is increased by giving them a secondary task (mental rotation of 3D objects). In [7] the authors extend this work, by comparing the performance of HI LOA switching to that of Robot Initiative LOA switching (here the artificial agent is a software layer on the robot's on board computing) and MI LOA switching. These systems are also then tested on a physical robot in field experiments and the challenges of implementing such systems practically are identified. The set of robot vitals for our preliminary work are listed below

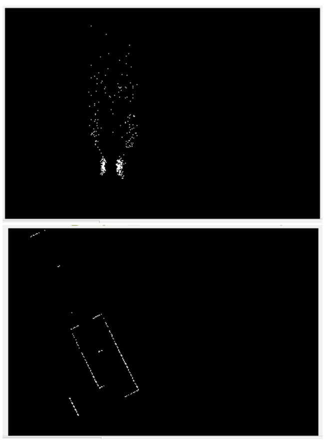

Figure 2: Laser Scanner Reading with (above) and without (below) Gaussian noise

3.2.1 ***Change in Signal Strength:*** One of the factors that cause robots to fail in the field is laser scanner noise. The amount of noise in a signal is usually measured using the signal-to-noise ratio (SNR)[13]. SNR is the ratio of the power of the signal P_s to the power of the noise P_N. Most laser scans can be visualised as a black and white image. An example of laser scans by robots in [7] with and without Gaussian noise is given in figure 2. The Peak SNR (PSNR) is calculated using equations 1 and 2 for any given signal.

$$PSNR = 10 \cdot log_{10} \frac{MAX_I^2}{MSE} \tag{1}$$

$$MSE = \frac{1}{c \cdot i \cdot j} \sum (I_1 - I_2)^2 \tag{2}$$

The maximum valid value per pixel MAX_I is 255 in the case of simple 8-bit black and white images. To calculate the $PSNR$, this value is divided by the mean squared error. Here, the MSE or the mean squared error[25] is calculated for the noise free image I_1 and a noisy image I_2 with two dimensional size ixj and c channels. The noise free image is available before experimentation. Instead, we propose to measure the rate of change in $PSNR$ between laser scans made a few time steps apart to check the relative change in signal strength. This value is denoted by $\dot{\delta}_{PSNR}$. To remove jitter and artefacts, the rate of change of this value over a minimum time of 0.5 - 1.0 seconds were calculated.

3.2.2 ***Velocity Error Average:*** Similar to the work carried out in [7], we assume the existence of a task expert (e.g. an expert navigation planner) that can provide the ideal task performance for the human-robot system in the absence of performance-degrading or other unexpected factors. The velocity error is calculated as the difference between the robot's actual velocity to the ideal velocity. The exponential moving average of this value is used as the next vital, denoted by v_{error}.

3.2.3 ***Rate of Change of Area Coverage:*** The Manhattan distance [2] of the robot from the start position is used for the next vital. The rate of increase of this distance from the start position, or the robot's previous way point gives the rate of increase of area coverage $\dot{d}_{waypoint}$.

3.2.4 ***Robot Acceleration:*** The rate of change of the robot's actual velocity $\dot{v}$ is used as the next vital. Usually a sudden drop in acceleration is observed only when the robot is making a turn. If such a drop is observed when in straight passage, it may be due to some debris blocking the robot's path.

3.2.5 ***Localisation Error:*** As described in [17], the position of the robot can be calculated independently by using different sensors like LIDAR sensors of the robot, its camera, wheel encoders and GPS etc. While SLAM techniques combine independent measures of the robot's location to improve the reliability of the estimate [9], measuring the difference between the location estimates by different sensors gives an estimate of the noise in each sensors value. In some cases, odometry noise could be caused by debris lodged in the wheel. To accommodate for this source of performance degradation, the localisation error is used as the next vital $\delta_{position}$. Let X_1 be the position of the robot calculated using its odometry, and X_2 be the robot's position calculated according to SLAM using LIDAR [11]. The position error is given by $\delta_{position} = X_1 - X_2$. This error value has an window of acceptable values. A high error outside this window indicates that either the odometry calculation or SLAM is not functioning properly.

4 ROBOT HEALTH

Let the vector of robot vitals for any robot i at time t be given by V_i^t. This vector is given by

$$V_i^t = \{\dot{\delta}_{PSNR}, v_{error}^t, \dot{d}_{waypoint}^t, \dot{v}, \delta_{position}^t\} \tag{3}$$

This vector of vitals characterise a robot's health. Each of these vitals are indicative of the degree of performance degradation experienced by the robot at any given time. A very high amount of performance degradation may or may not directly result in the robots failure. However when a robot experiences performance degradation, it is highly probable that the robot will fail. Hence, it is assumed that the probability of a robot's failure is conditionally dependent on each of the robot vitals.

The probability of failure of robot i at time t is calculated using the law of total probability[28] as given below

$$P(f_i) = \sum_{v \in V_i^t} P(f_i|v)P(v) \tag{4}$$

As the robot is equally likely to fail due to a change in of the above mentioned vitals, $P(v)$ is assumed the same for each of the vitals (in this specific case $P(v) = 0.20$).

The probability of robot failure given vital v is close to 0, as long as the value of v is within a standard range. This probability gradually starts increasing as the value of v deviates from the standard operating window. Figure 3 shows a density plot of different values of v_{error} observed during experiments on MI LOA switching[7]. It was observed during experimentation that in many cases where the robot was experiencing high performance degradation the value of v_{error} was close to 0.07. When the robot's performance was satisfactory, this value was closer to $0 - 0.01$. This essentially indicates that the robot's performance deviated from that of an expert planner the most when it was experiencing performance degradation.

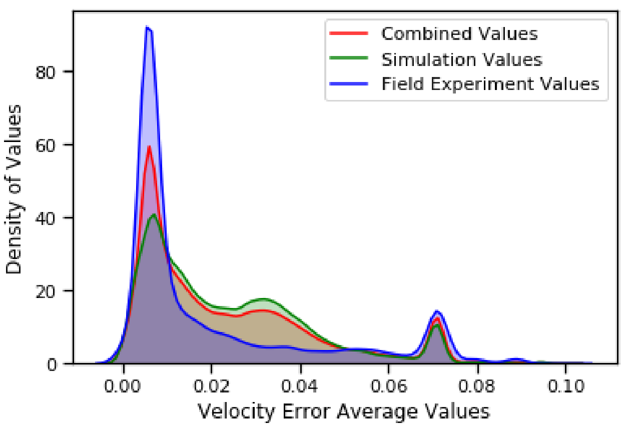

Figure 3: Density plot of velocity error average plotted from previous experiment data

Modelling the probability of robot failure given v_{error} as a sigmoid function[26] captures this probability distribution the best. Thus, the probability distribution for robot failure given the velocity error $P(f|v_{error})$ is given by:

$$P(f|v_{error}) = \frac{1}{1 + e^{-c_1(x-c_2)}} \tag{5}$$

The values of $c1$ and $c2$ are determined empirically to adapt to the specific robotic hardware used for experimentation. For the case given in Figure 3, the value of $c_1 = 1000$ and $c_2 = 0.0685$ to get sigmoid function where $P(f|v_{error} = 0.07) = 0.8176$. For $v_{error} < 0.06$ the probability of failure is trivial, and for values $v_{error} < 0.07$ the probability sharply rises to 1. A similar model for the probability of robot failure given each of the robot vitals can be modelled as a sigmoid function. The exact trigger values can be determined empirically.

In this study, we define the robot health as a scalar value that captures the degree of uncertainty or 'surprise' around the possible outcomes of V_i^t. Uncertainty about the robot's performance degradation is directly proportional to the health of the robot. When a robot is experiencing high performance degradation for a sustained period of time, unless remedial actions are taken to address the performance degradation the probability of failure is going to remain high. Hence there is little uncertainty in the values of $P(f)$ observed. In such cases, the robot health should be low. Alternatively if the robot is experiencing little to no performance degradation, there is high uncertainty. This is because the robot may or may not immediately encounter a situation that causes its performance to degrade. During high uncertainty periods the robot health should be high. To measure robot health as the degree of uncertainty surrounding the vitals, we propose the use of information entropy[20]. The health of robot i calculated from time steps t_1 to t_2 is given by

$$H_i^{t_1:t_2} = \sum_{t=t_1}^{t=t_2} -P(f_i^t) log(P(f_i^t)) \tag{6}$$

We conjecture that using threshold values of information entropy, it is possible to assess whether the robot requires an LOA switch. In MRS, tele-operator attention for two or more robots simultaneously experiencing performance degradation can potentially be prioritised based on the entropy value.

5 CONCLUSION AND FUTURE WORK

This paper proposes a set of five robot vitals and the use of information entropy as a metric for calculating the robot health. Using the concept of robot vitals and robot health, we aim to provide the tele-operator with information on the status of each robot in an easy and intuitive way through visual cues. Robot vitals can show which robots are performing poorly and need tele-operator attention. Our first set of experiments will focus on validating the use of robot vitals and robot health for quantifying performance degradation in mobile UGVs. In these experiments, UGVs shall be subjected to a set of performance degrading factors that increasingly impair their ability to function autonomously. We aim to demonstrate that changes in robot vitals and robot health during run time are statistically correlated to the extent of performance degradation faced by the robot in each experiment.

The current set of robot vitals work only for a limited set of use cases. For example, the vital $\dot{\delta}_{PSNR}$ can only be used for types of noise that affect the signal strength (e.g. Gaussian White Noise). Similarly, the vital v_{error}^t is contingent on the availability of an an expert planner during run time. We hope that our preliminary experiments inform the creation of a robust set of robot vitals that characterise more forms of non-terminal field-repairable performance degradation observed in extreme environments.

Subsequent work shall extend the work carried out by Chiou et al. [6, 7] and study how humans interact with variable autonomy multi-robot systems. Using the robot vitals and robot health, we aim to build artificial agents that can assist tele-operators with LOA regulation in MI multi-robot systems. By measuring how overall task performance and operator cognitive workload are affected by the use of HI and MI systems, we hope to understand how better interaction design can improve human multi-robot teaming.

REFERENCES

[1] Naoki Akail, Luis Yoichi Moralesl, and Hiroshi Murase. 2018. Reliability estimation of vehicle localization result. In *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 740–747.

[2] Paul E Black. 2006. Manhattan distance"" Dictionary of algorithms and data structures. *http://xlinux. nist. gov/dads//* (2006).

[3] Jennifer Carlson and Robin R Murphy. 2005. How UGVs physically fail in the field. *IEEE Transactions on robotics* 21, 3 (2005), 423–437.

[4] Jennifer Casper and Robin R. Murphy. 2003. Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 33, 3 (2003), 367–385.

[5] JYC Chen and Peter I Terrence. 2009. Effects of imperfect automation and individual differences on concurrent performance of military and robotics tasks in a simulated multitasking environment. *Ergonomics* 52, 8 (2009), 907–920.

[6] Manolis Chiou, Goda Bieksaite, Nick Hawes, and Rustam Stolkin. 2016. Human-Initiative Variable Autonomy: An Experimental Analysis of the Interactions Between a Human Operator and a Remotely Operated Mobile Robot which also Possesses Autonomous Capabilities.. In *AAAI Fall Symposia*.

[7] Manolis Chiou, Nick Hawes, and Rustam Stolkin. 2019. Mixed-Initiative variable autonomy for remotely operated mobile robots. *arXiv preprint arXiv:1911.04848* (2019).

[8] Didier Crestani, Karen Godary-Dejean, and Lionel Lapierre. 2015. Enhancing fault tolerance of autonomous mobile robots. *Robotics and Autonomous Systems* 68 (2015), 140–155.

[9] Hugh Durrant-Whyte and Tim Bailey. 2006. Simultaneous localization and mapping: part I. *IEEE robotics & automation magazine* 13, 2 (2006), 99–110.

[10] Michael A Goodrich, Morgan Quigley, and Keryl Cosenzo. 2005. Task switching and multi-robot teams. In *Multi-Robot Systems. From Swarms to Intelligent Automata Volume III*. Springer, 185–195.

[11] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. 2016. Real-time loop closure in 2D LIDAR SLAM. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1271–1278.

[12] Shu Jiang and Ronald C Arkin. 2015. Mixed-initiative human-robot interaction: definition, taxonomy, and survey. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, 954–961.

[13] Don H Johnson. 2006. Signal-to-noise ratio. *Scholarpedia* 1, 12 (2006), 2088.

[14] Andreas Kolling, Katia Sycara, Steve Nunnally, and Michael Lewis. 2013. Human swarm interaction: An experimental study of two types of interaction with foraging swarms. *Journal of Human-Robot Interaction* 2, 2 (2013).

[15] Yugang Liu and Goldie Nejat. 2013. Robotic urban search and rescue: A survey from the control perspective. *Journal of Intelligent & Robotic Systems* 72, 2 (2013), 147–165.

[16] Yugang Liu and Goldie Nejat. 2016. Multirobot cooperative learning for semi-autonomous control in urban search and rescue applications. *Journal of Field Robotics* 33, 4 (2016), 512–536.

[17] Juan Pablo Mendoza, Manuela M Veloso, and Reid Simmons. 2012. Mobile robot fault detection based on redundant information statistics. (2012).

[18] Robin R Murphy. 2010. Navigational and mission usability in rescue robots. *Journal of the Robotics Society of Japan* 28, 2 (2010), 142–146.

[19] Ariel Rosenfeld, Noa Agmon, Oleg Maksimov, and Sarit Kraus. 2017. Intelligent agent supporting human–multi-robot team collaboration. *Artificial Intelligence* 252 (2017), 211–231.

[20] Claude E Shannon. 1948. A mathematical theory of communication. *The Bell system technical journal* 27, 3 (1948), 379–423.

[21] Thomas B Sheridan and William L Verplank. 1978. *Human and computer control of undersea teleoperators*. Technical Report. Massachusetts Inst of Tech Cambridge Man-Machine Systems Lab.

[22] PN Squire and R Parasuraman. 2010. Effects of automation and task load on task switching during human supervision of multiple semi-autonomous robots in a dynamic environment. *Ergonomics* 53, 8 (2010), 951–961.

[23] Stephanie Studenski, Subashan Perera, Dennis Wallace, Julie M Chandler, Pamela W Duncan, Earl Rooney, Michael Fox, and Jack M Guralnik. 2003. Physical performance measures in the clinical setting. *Journal of the American Geriatrics Society* 51, 3 (2003), 314–322.

[24] Alberto Valero, Massimo Mecella, Fernando Matia, and Daniele Nardi. 2008. Adaptative human-robot interaction for mobile robots. In *RO-MAN 2008-The 17th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 243–248.

[25] Z. Wang and A. C. Bovik. 2009. Mean squared error: Love it or leave it? A new look at Signal Fidelity Measures. *IEEE Signal Processing Magazine* 26, 1 (2009), 98–117. https://doi.org/10.1109/MSP.2008.930649

[26] Eric W Weisstein. 2002. Sigmoid function. *https://mathworld. wolfram. com/* (2002).

[27] Michael Young, Mahdieh Nejati Javaremi, and Brenna D Argall. 2019. Discrete N-Dimensional Entropy of Behavior: DNDEB. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2227–2233.

[28] Daniel Zwillinger and Stephen Kokoska. 1999. *CRC standard probability and statistics tables and formulae*. Crc Press.