Information Metrics for Assistive Human-In-The-Loop Cognitive Systems

Martin F. Stoelen, Alberto Jardon, Juan G. Victores, Carlos Balaguer, Fabio Bonsignorio

***Abstract*— The development of intelligent service robotic systems is currently an active field of research in the robotics community. For example assistive robots that can aid elderly and disabled people in daily life activities. One emerging requirement for this type of system is the inclusion of the user in the decision process through physical and cognitive collaboration. This human-in-the-loop (HIL) concept allows for the use of the human perception and cognitive abilities in order to safely achieve the tasks that would be too complex to perform in a purely autonomous way. However, the overall human-machine system is complex and may be difficult to analyze. The user and the robot are operating in a closed loop and both are potentially capable of adapting to the other. The users may have a disparate set of noisy channels available for communicating their intended commands to the robot. The robots are typically dexterous and are expected to operate in an unstructured environment. Metrics can help in the analysis, development, and benchmarking of this type of system, by quantifying performance and driving the mutual learning and adaptation process. However, there are currently few such metrics available. Information Theory and related information-based concepts have been applied in disparate fields such as communications, human factors, control theory and cognitive processes. The work presented here attempts to identify metrics based on these concepts for assistive human-in-the-loop cognitive systems.**

I. INTRODUCTION

Assistive robots are currently being developed to support disabled and elderly people inside their own homes and in other everyday environments. One example is the climbing assistive robot ASIBOT, developed at Universidad Carlos III de Madrid [1]. Several other assistive robot types exist, ranging from static systems like HANDY 1 [2] to mobile manipulators like KARES II [3].

Fig. 1 is a simplified model of the complete human-machine system for this type of robots. As can be seen in the figure, the model assumes that the user has some intentional commands for the robot, **h**, that are actuated through a set of input devices. The disabilities of the user are modeled as sources of noise, **z**, which can be independent for each input modality. The multimodal signals received by the cognitive part of the machine, here denoted the enabling interface, **d**, are thus noisy representations of the user's true intention. The goal of the enabling interface is to use these noisy signals

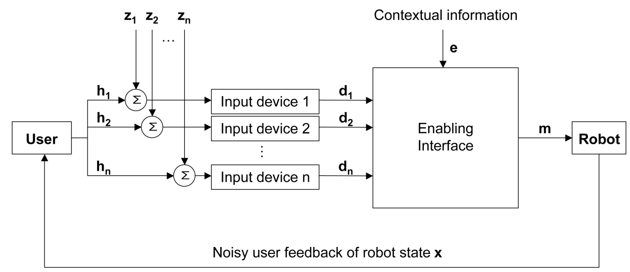

Fig. 1. Simplified representation of the complete human-machine system.

and time copies thereof, together with information from the context (**e**), to produce robot commands (**m**) that are as close as possible to the user's original intention. The user receives noisy feedback about the state of the robot, **x**, closing the loop. Feedback from the state of the input devices (visual and/or proprioceptive) is omitted for clarity in Fig. 1. Both the human and the machine are assumed to potentially be able to perform some form of adaptation and learning.

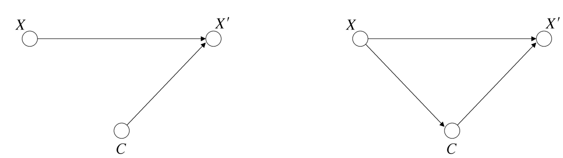

(a) Open-loop control. (b) Closed-loop control.

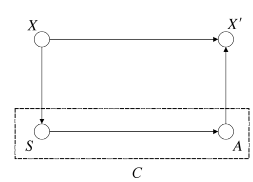

(c) Sensor (S) and actuator (A) constituting the controller (C).

Fig. 2. A control system as a directed acyclic graph.

One interesting approach to analyzing complex closed loop systems like the one shown in Fig. 1 can be found in [4]. This is based on representing a complete control system

Martin F. Stoelen, Alberto Jardon, Juan G. Victores and Carlos Balaguer are members of the RoboticsLab research group within the Department of Systems Engineering and Automation, Universidad Carlos III de Madrid, (mstoelen, ajardon, jcgvicto, balaguer)@ing.uc3m.es

Fabio Bonsignorio is a member of the RoboticsLab as Santander Chair of Excellence in Robotics at Universidad Carlos III de Madrid and CEO/founder of Heron Robots of Genova, Italy, bonsign@ing.uc3m.es

as a directed acyclic graph of random variables, see Fig. 2, and analyzing it using concepts from Information Theory [5]. The system includes the current state *X*, with values $x \in \mathcal{X}$, and the future state *X'*. The random variable representing the controller, *C*, then senses the current state (with sensor *S*) and actuates to achieve the future state (with actuator *A*) . This can be represented by conditional probabilities, $p(c|x)$ and $p(x'|x, c)$. These can be viewed as representing a sensor and actuation channel, respectively. The authors were further able to derive the conditions for observability, controllability, and optimality using this method. Fig. 3 depicts our extension of this method to the human-machine system, introduced in [6].

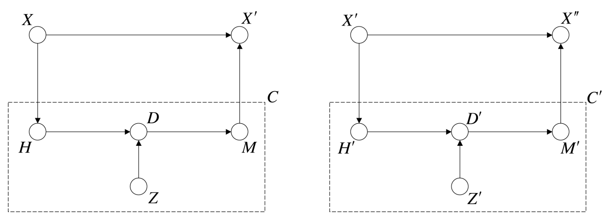

(a) System at t. (b) System at $t + 1$.

Fig. 3. The human-machine system as a directed acyclic graph, shown for two consecutive instances in time.

The controller C here includes both the user (more generally the Human, *H*) and the assistive robot (more generally the Machine, *M*). The goal of the human-machine system in the most general sense is then to maximize the flow of useful information between the human and the machine over a noisy medium. Thus, we are interested in the communication channel existing between a source H and a receiver M, which will be denoted the "human-machine channel" in the following discussion and which has channel capacity C_{HM}. The information available in the source can be represented by the Shannon entropy of the random variable representing the human, here denoted as $S(H)$. The definition for entropy used here is shown in (1).

$$S(X) = -\sum_{x \in \mathcal{X}} p(x) \log p(x). \tag{1}$$

In Information Theory terminology the stated goal is then equivalent to transmitting this information over the human-machine channel with a minimum of errors. Chan and Childress [7] also applied information theory principles to analyze the information transmission in the human-machine system for tracking tasks. The analysis here differs in that it entails multimodality as well as learning, and is applied on the directed acyclic graph representing the system.

More specifically, the goal of the human-machine system can be defined as maximizing the flow of useful information between the user and the assistive robot, given the user's physical disability. As can be seen in Fig. 3, the disability is also here modeled as a source of noise, *Z*. The random variable representing a given input device, *D*, will then depend probabilistically on both the user's true intentions, *H*, and the noise *Z*.

A model with a user both mentally and physically healthy will not include this noise. Assuming that input devices with sufficient performance are available to the user, we would then have $S(H)_h \leq C_{HMh}$. The subscript *h* is used to denote physically healthy here. As stated in the channel coding theorem [5], there exists a coding system for this situation such that the information from the source, the user's intended commands, can be transmitted with an arbitrarily low error.

The interpretation of a mentally healthy, but physically disabled user attempting to control a complex system like an assistive robot is then that of a source rich in information, but acting over a human-machine channel with limited channel capacity. We are assuming this user has no limitation in his/her ability to imagine commands, thus $S(H)_d = S(H)_h$. The subscript *d* is used to denote physically disabled here. The difference from the physically healthy user is then the noise added in the human-machine channel, leading to $S(H)_d > C_{HMd}$.

This analysis is of little practical use however, if these quantities cannot be measured. The purpose of this article is to identify information-based metrics for this purpose, and attempt to apply these metrics to representative data. It is hoped that these metrics may help drive the development of assistive human-in-the-loop cognitive systems, by quantifying performance and potentially motivating learning.

II. METRICS

A. *"Empowerment"*

One of the central concerns when designing a human-machine interface is to ensure the user feels in control of the machine. This is among other expressed in one of the "eight golden rules" of user interface design for computers [8], promoting that the interface should "support internal locus of control." In robotics, the term "mode confusion" is sometimes applied, referring to the undesirable situation that occurs when the system's true state differs from what the user predicts based on his/her mental model of the system. Central to both these concepts is the relationship between what the user inputs into the system, over his/her actuators, and what the user perceives from the system, over his/her sensors.

Klyubin, Polani and Nehaniv [9] proposed "empowerment" as a task-independent driving principle for sensorimotor systems: "Empowerment quantifies the agent's *potential* ability to influence the environment as measured by the capacity to "imprint" information onto the environment and later perceive the information via the sensors". Formally the measure is defined as the mutual information across a finite number of past actuations and the current sensor value of an agent.

To investigate this and similar measures the experimental setting seen in Fig. 4 was constructed. Here the system may allow for multimodal input and the user is being given an explicit desired state, X_d. Noise can be added to the input devices if required. This could be seen as a crude disability simulation, and is of interest for testing this type of systems outside of clinical conditions during the development phase. It could also facilitate more uniform subject pools for larger experiments, although only as an approximation to the real disabled users. For the purpose of this experiment the metric "empowerment" was then measured across the desired state (the position of a virtual object on a screen), and the actual state of the system (the position of a user-controlled virtual object). It was assumed that the user's intention was to follow the desired state as closely as possible. The one-step mutual information was used, with "empowerment" being defined as $I(X_d; X)$.

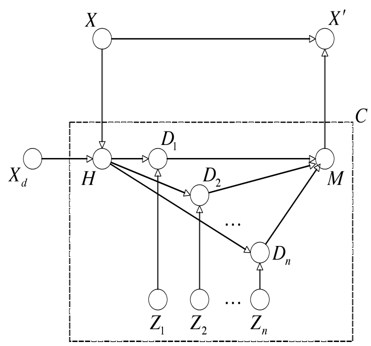

Fig. 4. The experimental setting considered.

B. Predictive Information

Bialek, Nemenman and Tishby [10] proposed predictive information, in the form of the mutual information between the past and the future, as a general measure of complexity of a time series. The measure can be said to quantify the total information of past experience that can be used for predicting future events, and has among other been applied to the behavior of mobile robots in an unknown environment, see [11]. Here the measure was found empirically to have a maximum for a behavior which is both explorative and sensitive to the environment. On the basis of this result, it was proposed that predictive information could be a prospective candidate as an objective function for the autonomous development of behaviors. This property may be of interest also for the system considered here, with two agents capable of adapting to each other. The predictive information of an agent's actuations should at least have some relation with how random, or explorative, the agent's behavior is. This may be useful information for whoever is trying to cooperate with this agent.

In experiments on human-human collaboration an increase in the speed and accuracy of movements with respect to when performing a task alone has been observed. This may perhaps be interpreted as an unconscious strategy to make the movements easily predictable by the other part [12], and might also be a suitable strategy for a machine agent cooperating with a human.

It is important to note that the measure uses only the time series of actuations performed, and does not for example require knowledge of the true intention of an agent's behavior. This knowledge can be hard to obtain outside the experimental setting. The one-step mutual information was used, with predictive information being defined as $I(X; X')$.

C. Quantifying Coordination

Another quantity of interest in the system shown in Fig. 4 is the amount of coordination across input modalities. A very general definition of modality is here used, including different DOF for the same input device. Zhai and Milgram [13] proposed "efficiency" as a measure for the coordination involved in movements using input devices with N rotational or translational DOF, based on comparing the actual trajectory in N-dimensional space with the shortest possible. Here we propose calculating the mutual information across two modalities, for example two DOF, as an alternative metric for quantifying coordination. The one-step mutual information was used, with $I(D_i; D_j)$ as the metric for quantifying coordination.

III. MATERIALS AND METHODS

A. Introduction

A pilot study was performed to explore the application of the metrics defined on representative data. The metrics are based on mutual information, which is calculated from probabilities. This is typically estimated from the relative frequency of events occurring, and thus requires a large number of samples to be accurate. However, if the metrics are to be applied to human-in-the-loop systems there are limitations on the amount of data available, typically tens to hundreds of repetitions, and minutes to hours of data recorded at tens of Hz. This contrasts the typical applications of information-based metrics so far, for example in motivating self-organizing behavior, where simulated environments give the possibility of an nearly unlimited amount of samples. A one degree-of-freedom (DOF) pursuit tracking task was chosen, with simple constant-amplitude sinusoidal movements of the target. The dependent variables were thus the trajectories of the target and the users-controlled cursor. The first independent variable was the frequency of oscillation of the target, with four levels: 0.05, 0.2, 0.8 and 1.6 Hz. The second independent variable was a noise added to the user's input, with two levels: with and without noise.

B. Subjects

The subjects were 3 male students of Universidad Carlos III de Madrid. None of the subjects were involved in the project related to this study. All subjects were right-handed and were between 25 and 27 years old with a mean of 26 years. There was 1 subject with corrected vision, and 2 had previous experience with 3D input devices. All gave their informed consent to participate in the study.

C. Apparatus

The pilot study was conducted on a PC workstation in an office environment, see Fig. 5. The participants worked on a 19 inch (482.6 mm) widescreen external liquid crystal display monitor (Asus VW195S) with a 60 Hz refresh rate. The input device was a 3Dconexion SpaceNavigator joystick with 6 DOF. The sensor of this system was held in the participants dominant hand and measured one rotational DOF, nominally the roll angle. The remaining DOF were not physically obstructed, but did not impact the movement of the cursor on the display. The subjects were performing movements in a virtual environment based on the OpenRAVE simulator [14], updating at approximately 50 Hz.

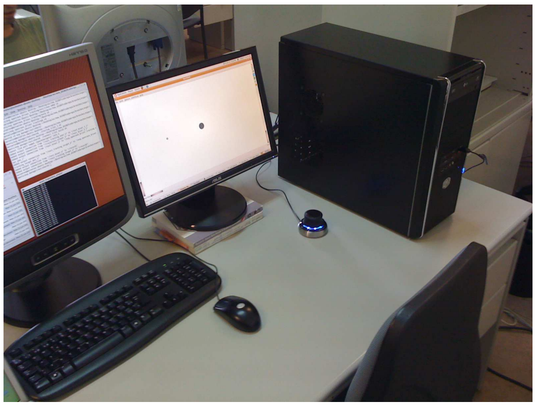

Fig. 5. The experiment setup.

D. Stimuli

The stimuli presented on the display can be seen in Fig. 6. The stimuli included a small, light-grey disc representing the target and a larger dark-grey disc representing the user's cursor. The latter was allowed to be obscured by the target disc, but was always visible due to the larger size. For half of the trials noise was added to the user's input. This was an approximate white noise, low-pass filtered with a cut-off frequency of 5 Hz.

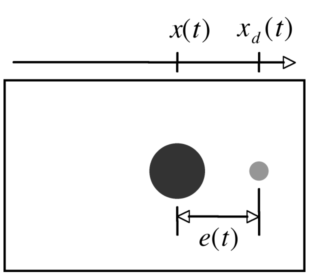

Fig. 6. Representation of the experiment stimuli.

E. Procedure

Each subjects session lasted approximately 40 minutes, starting with a thorough brief of the task procedures. Each experiment included 8 different combinations of tracking frequencies and noise added, each repeated 2 times for non-training trials. All subjects were first given 2 minutes for training to get acquainted with the task and the experimental setup (0.2 Hz and no noise added to input). The participants then performed the 16 attempts of one minute each, with breaks in-between as needed.

F. Data Collection and Conditioning

The data recorded was the position of the target and the user controlled cursor for each point in time, at 120 Hz. Thus approximately 14400 data points were available from each subject for each experiment condition. The data was then normalized to between 0 and 100 and generally dicretized to 50 states. The data discretization heuristic used was to keep the number of bins for the state less than the number of samples divided by 3. For these operations the Matlab toolbox of Lungarella, Pegors, Bulwinkle and Sporns [15] was used.

IV. RESULTS AND DISCUSSION

The metrics "empowerment" and predictive information was applied to the data gathered from the pilot study, and the results can be seen in Fig. 7 and Fig. 8, respectively. All subjects seem to behave in a reasonably similar manner, suggesting that the experiment design and protocol was sufficient for the purposes of this study. As can be seen from the figures, there was a reduction with higher frequencies of movement for both metrics. This makes sense qualitatively, as the subject struggles to keep up with the faster moving target. It is also consistent with results form the literature that shows an increase in mean square error with frequency, see for example [16]. Furthermore, the results show a drop in both metrics with noise added to the user's input. This is also to be expected, as the increased noise will reduce the statistical relationship between the desired and actual state as well as across past and current user input. As the frequency of the target increases, the difference between the condition with noise and the one without is reduced.

It is interesting to note that predictive information provides similar results as "empowerment", without knowledge of the desired trajectory. As mentioned before, this makes the metric easier to apply in non-experimental settings, for example in assessing performance during usage by the end-user. For both metrics similar developments were found for different state discretization, although the actual values increase with number of states.

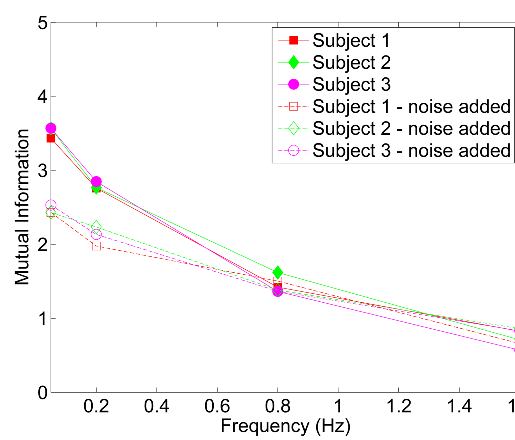

Fig. 7. "Empowerment" with varying target frequency and noise added to user input.

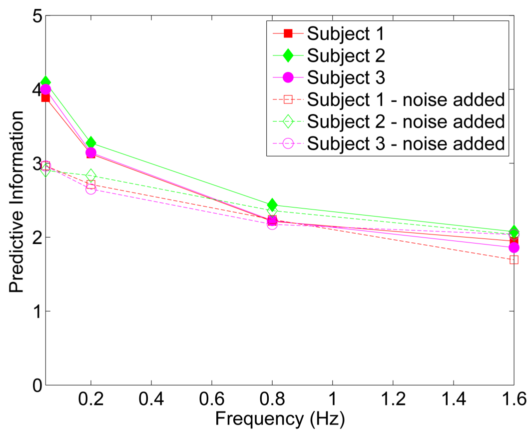

Fig. 8. Predictive information with varying target frequency, as well as with/without noise added to the user input.

To explore the utility of predictive information in online applications, two subjects were asked to perform a simple two-minute session after having completed the original pilot study. The target was moving at 0.8 Hz and no noise was added to the user's input. For this test only, the mapping of the deflections of the joystick to the movements of the cursor on the display was altered. Where before pitch movements were consistent with both direction and magnitude of the movement of the cursor, the new mapping involved negative yaw and a non-linear magnitude relationship. This caused the device to became more sensitive with higher deflections. The subjects were not informed of the exact mapping before performing the test, only that the mapping had changed from the previous study. The resulting trajectories can be seen in Fig. 9. The subject numbers do not necessarily correspond to that in the original pilot study.

As can be seen from the trajectories, subject 1 was able to achieve reasonable tracking after about 40 seconds, while subject 2 took less than 15 seconds. For the analysis of this data the time was binned at approximately 10 Hz. The predictive information was then calculated over 60 overlapping periods of 120 bins each, corresponding to about 12 seconds of data per period. The state was here discretized to 25 bins. As can be seen from the results in Fig. 10, the development of the metric confirms the qualitative observation from the trajectories made above. After discovering the approximate mapping, each subject's performance converges to more or less the same level. Applied in this fashion the metric shows promise as a candidate for motivating online adaptation to the user, with values available on the order of tens of seconds.

The coordination across DOF was also investigated. For this analysis, data from a previous study on applying Fitts' law [17] to combined rotational and translational movements was used. In this study subjects were asked to perform movements that had both a 1 DOF rotational and a 1 DOF translational magnitude and accuracy requirement, see Stoelen and Akin [18] for details. Each condition had a different magnitude and accuracy requirement. The magnitudes ranged from 4.8 to 12.7 cm for translations and 50 to 130 degrees for

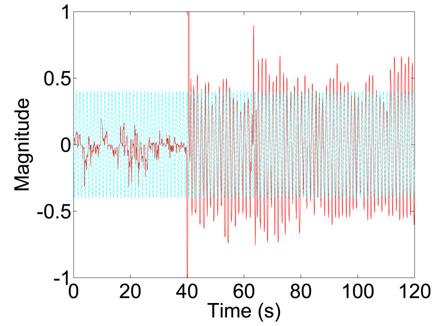

(a) Subject 1.

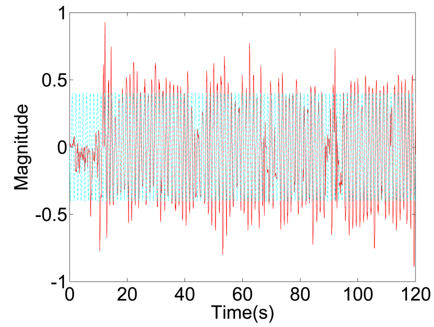

(b) Subject 2.

Fig. 9. Trajectories of target (dotted cyan line) and user controlled cursor (solid red line) when adapting to a new mapping from the input device to the cursor movements on the display.

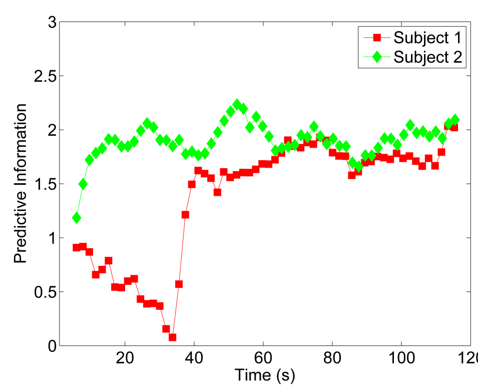

Fig. 10. Predictive information (calculated over 12 second periods) after a change in the mapping from the input device to the cursor movements on the display.

rotations. The resulting 12 attempts for 1 out of 16 conditions can be seen in Fig. 11, for three different subjects.

From the trajectories it can be seen that subject 1 and 2 performed the task in a relatively parallel fashion, timing the translational and rotational movement to start and finish at approximately the same time. Subject 3 on the other hand, did not begin the rotational movement until the translational movement was more or less completed. This subject was the only one among the 12 subjects used that did not perform the task in paralell. It can also be seen that subject 1 performed the task in a more coordinated fashion than subject 2, with the trajectories varying little across attempts. The trajectories shown are representative of the performance of the subjects on the remaining 15 conditions.

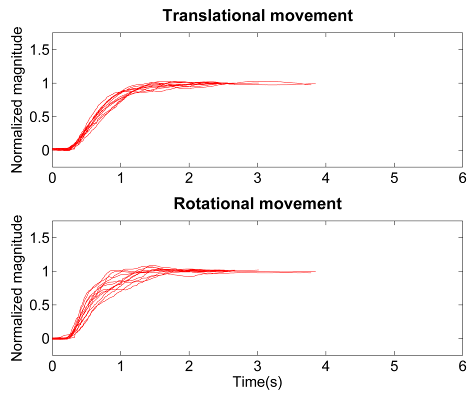

(a) Subject 1.

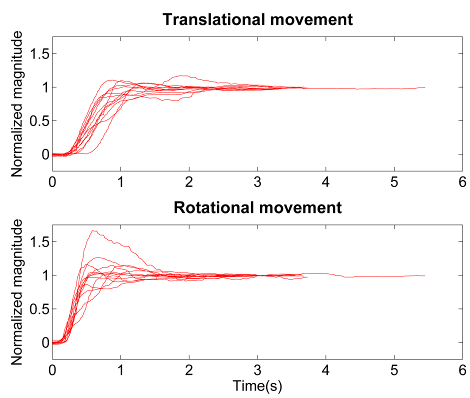

(b) Subject 2.

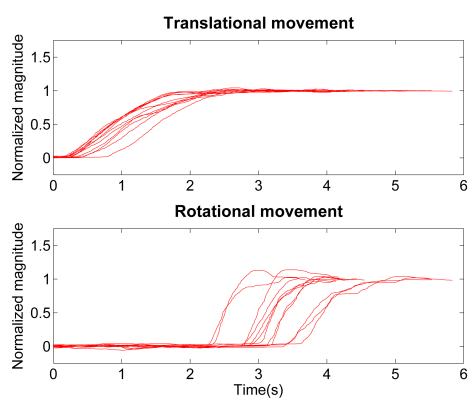

(c) Subject 3.

Fig. 11. Example trajectories for 12 attempts on a given discrete task with combined rotational and translational movement magnitude and accuracy requirements. Data from [18].

For each subject 192 sample trajectories were used, with the magnitude of rotations and translations normalized based on the requirements on magnitude for each task. The data was discretized to 8 states. The corresponding mutual information across rotational angle and translational position was 1.11, 0.56 and 0.34, respectively, for subject 1, 2 and 3. This is an interesting result, as not only does the metric quantify a difference in coordination for subjects that performed the task in parallel and in a serial fashion, but it is also twice as large for the more coordinated subject 1, as compared to subject 2. Although the application here was on two DOF from the same input device, it is not unlikely that the metric could be used in a similar fashion for multimodal interfaces. For example to quantify the ability of a user to effectively coordinate two input devices to achieve a given goal.

V. FUTURE WORK

The study performed for this article was very limited in its scope, both in terms of the number of subjects and in the simplicity of the task performed. The future work will therefore include more extensive controlled experiments, where these weaknesses can be adressed. These experiments should include a larger pool of subjects, should address the stationarity of the data collected and include comparisons with existing metrics for the tasks performed. In addition it would be beneficial to begin to explore the application of the metrics to more realistic situations. This could include attempting to quantify the performance of users with real disabilities, coordination of disparate input modalities, and exploring the use in a system with mutual adaptation by the human and the machine agent.

REFERENCES

[1] A. Jardón, A. Giménez, R. Correal, R. Cabas, S. Martínez, and C. Balaguer, “A portable light-weight climbing robot for personal assistance applications,” *Industrial Robot: An International Journal*, vol. 33. no. 4. pp. 303-307, 2006.

[2] M. Topping, “Early Experiences Encountered in the Placement of 60 Handy 1 Robotic Aids to Eating,” in *Proc. of the IEEE 93 Conference on Cybernetics*, Le Touguet, France, 1993.

[3] Z. Bien, M. Chung, P. Chang, D. Kwon, D. Kim, J. Han, J. Kim, D. Kim, H. Park, S. Kang, K. Lee, and S. Lim, “Integration of a Rehabilitation Robotic System (KARES II) with Human-Friendly Man-Machine Interaction Units,” *Autonomous Robots*, vol. 16, pp. 165-191, 2004.

[4] H. Touchette and S. Lloyd, “Information-theoretic approach to the study of control systems,” *Physica A: Statistical Mechanics and its Applications*, vol. 331, pp. 140-172, 2004.

[5] C. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal*, vol. 27, pp. 379-423, 1948.

[6] M.F. Stoelen, A. Jardon, F. Bonsignorio, J.G. Victores, C. Monje and C. Balaguer, “Towards an Enabling Multimodal Interface for an Assistive Robot,” *Workshop on Mutimodal Human-Robot Interfaces, IEEE International Conference on Robotics and Automation*, Anchorage, Alaska, 2010.

[7] R. Chan and D. Childress, “On information transmission in human-machine systems: channel capacity and optimal filtering,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 20, pp. 1136-1145, 1990.

[8] B. Shneiderman and C. Plaisant, “Designing the User Interface: Strategies for Effective Human-Computer Interaction (4th Edition)”, Pearson Addison Wesley, 2004.

[9] A.S. Klyubin, D. Polani, and C.L. Nehaniv, “Keep your options open: an information-based driving principle for sensorimotor systems.,” *PloS one*, vol. 3, p. e4018, 2008.

[10] W. Bialek, I. Nemenman, and N. Tishby, “Predictability, complexity, and learning.,” *Neural computation*, vol. 13, pp. 2409-63, 2001.

[11] N. Ay, N. Bertschinger, R. Der, F. Guttler, and E. Olbrich, “Predictive information and explorative behavior of autonomous robots,” *The European Physical Journal B*, vol. 63, pp. 329-339, 2008.

[12] C. Vesper, S. Stork, and A. Schubo, “Movement Times in Inter-and Intrapersonal Human Coordination,” *Proceedings of the 2008 ECSIS Symposium on Learning and Adaptive Behaviors for Robotic Systems*, IEEE Computer Society, pp. 17-22, 2008.

[13] S. Zhai, and P. Milgram, “Quantifying coordination in multiple DOF movement and its application to evaluating 6 DOF input devices,” *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Los Angeles, California, United States, April 18 - 23, 1998.

[14] R. Diankov and J. Kuffner, “OpenRAVE: A Planning Architecture for Autonomous Robotics,” tech. report CMU-RI-TR-08-34, Robotics Institute, Carnegie Mellon University, July, 2008.

[15] M. Lungarella, T. Pegors, D. Bulwinkle, and O. Sporns, “Methods for Quantifying the Informational Structure of Sensory and Motor Data,” *Neuroinformatics*, pp. 243-262, 2005.

[16] J.R. Ware, “An Input Adaptive, Pursuit Tracking Model of the Human Operator,” *Seventh Annual Conference on Manual Control*, Los Angeles, California, 1972.

[17] P.M. Fitts, “The Information Capacity of the human Motor System in Controlling the Amplitude of Movement,” *Journal of Experimental Psychology*, vol. 47, pp. 381-391, 1954.

[18] M.F. Stoelen and D.L. Akin, “Assessment of Fitts’ Law for Quantifying Combined Rotational and Translational Movements,” *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 52, no. 1, pp. 63-77, 2010.