2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

Shared Control of Assistive Robots through User-intent Prediction and Hyperdimensional Recall of Reactive Behavior

Alisha Menon, Laura I. Galindez Olascoaga, Vamshi Balanaga, Anirudh Natarajan, Jennifer Ruffing, Ryan Ardalan and Jan M. Rabaey

***Abstract*— There is increasing interest in shared control for assistive robotics with adaptable levels of supervised autonomy. In this work, we present a user-adaptive multi-layer shared control scheme for control of assistive devices. The system leverages the advantages of brain-inspired hyperdimensional computing (HDC) for classification & recall of reactive robotic behavior including high performance, computational efficiency and intelligent sensor fusion, to execute actuation based on the user's goal while alleviating the burden of fine control. Using a multi-modal dataset of activities of daily living, we first recognize the user's most recent behaviors, then predict the user's next action based on their habitual action sequences, and finally, determine actuation through HDC recall-based shared control which intelligently deliberates between the predicted action and sensor feedback-based autonomy. In this work, we independently implement each layer to achieve >92% accuracy and then integrate the layers and discuss the combined performance and methods to reduce accumulated error.**

I. INTRODUCTION

Myoelectric hands aim to provide the user with a functional residual limb, allowing them to perform common activities of daily living that may include pinching, pushing, or grasping objects [1]. They improve finger movement and wrist rotation that passive prostheses cannot provide. However, the functionality on commercial hands is generally limited to EMG-based selection of a grip from a programmable set [2]. It also relies on consistent placement of the sensors on the forearm when donning the prosthetic device, involves long training sessions, and struggles with signal degredation over even one wear session [3].

Due to the unreliability and inconvenience of switching grips, users tend to use one grip 70% of the time [1], leading to poor voluntary usage of even the limited existing functionalities. Prosthetics impose other practical difficulties on the user as well including lack of feedback, limited grip gestures, [3] and a high expenditure of 90% of the user's visual attention on either their prosthetic or the area near the object [4]. EMG-based control also causes early fatigue since the user must continuously over-emphasize specific muscle patterns [5]. While these systems provide users with a functioning hand, the difficulties in addition to low dexterity and unintuitive assistance ultimately lead to an up to 75% rejection rate of myoelectric prostheses [6].

Robotic control approaches often eliminate the mental and physical prosthetic burden through optimal control schemes like reinforcement learning (RL). These use a cost-minimization policy to determine [7] an optimal trajectory. However, the models are limited to pre-determined scenarios or goals [8] and are data hungry and slow [9]. This is especially an issue for myoelectric prosthetics where training data is limited for each user and wear session, and where actuation latency impedes task completion. Also, these methods do not allow the user to naturally execute on their goals, make decisions or anticipate prosthetic behavior due to the lack of user involvement in trajectory selection.

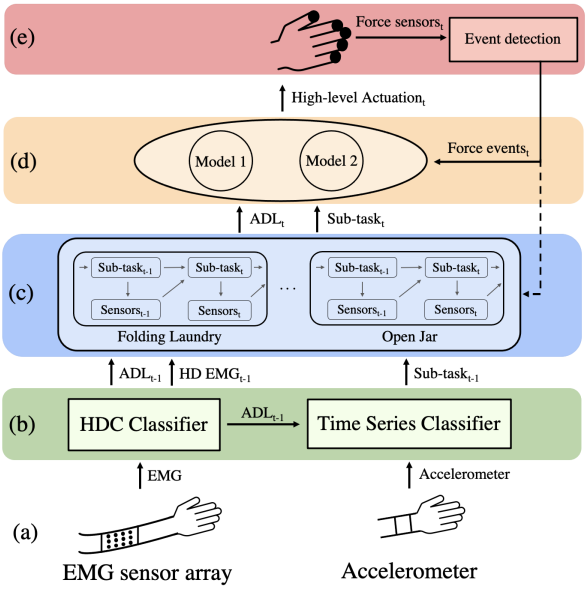

Fig. 1: Multi-layer shared control scheme with (a) sensor inputs (b) user behavior recognition (c) user intent prediction (d) shared control with HDC RORB (e) actuation & feedback.

"Shared control" is an emerging strategy that intelligently arbitrates between user control and prosthetic autonomy. This allows the system to execute on human intent while alleviating the fine control burden. Previous works using shared control schemes include either simple binary switching from human-driven control to fully-automated control where the user completely loses autonomy in the second mode [5], or a weighted sum of the human-driven trajectory with an inverse RL-based trajectory where the model is limited to very specific intents – even changing the object's location would necessitate a new intent and re-training [10].

In this work, we design & evaluate a shared control system that maintains user autonomy at all times, but displaces the burden of fine control. The system is user-adaptive, generalizes to new tasks without large amounts of training

All authors are with the Berkeley Wireless Research Center, EECS Department, University of California, Berkeley. Corresponding author: allymenon@berkeley.edu.

data, and aims for computational simplicity for in-device implementation to lower latency and avoid power-hungry data streaming. It has the potential to quickly update its model, avoiding performance degredation due to time or sensor placement inconsistency. Towards these goals, the system takes advantage of hyperdimensional computing (HDC), a brain-inspired computationally-efficient paradigm based on the high-dimensional circuits found in the brain [11].

HDC has shown high success for a variety of biosignal classification problems such as multi-modal emotion recognition [12], seizure detection [13] and EMG-based gesture recognition [14], and also in robotics through a 2-D navigation task [15]. It has major advantages in noise robustness [16], hardware efficiency [17] [15] and memory footprint [18]. HDC models for EMG systems perform well on limited training data and can be rapidly updated [19]. Embedded HDC achieves low latency, enabling in-device classification for fast & natural prosthetic actuation [14]. The paradigm is very versatile; it can address sensorimotor mappings [20], model superposition [21] and efficient symbolic reasoning [22]. These can be used in combination for a multi-layer system for prosthetic shared control as shown in Fig. 1.

Overall, this paper presents the following contributions:

Recognition of user behaviors. We recognize the user's behaviors, completing the first layer of the multi-layer scheme. We utilize an HDC classifier to recognize the user's long-term goal based on their EMG signals, and a time-series classifier to recognize their most recent short-term behavior based on accelerometer sensor information.

User intent prediction. We successfully predict the user's intended short-term behavior through probabilistic inference. This layer relies on the long-term goal recognition and time series classification from the first layer, and additional behavioral information derived from HDC-encoded EMG.

Shared control through HDC recall of reactive behavior. We implement the third layer to arbitrate between human-driven and automatic control to output the most appropriate actuation. An HDC-based recall algorithm is designed to store the mapping of rich heterogeneous inputs to reactive outputs. This mapping includes variations on prosthetic behaviors as it reacts to force feedback from the end-effector.

Integration of the layers and discussion of the outcome. Finally, we integrate the layers designed and evaluated in this work and observe the preliminary results. We discuss the outcome and propose methods to correct for error accumulated through the system.

II. Related work

When being fit for a myoelectric terminal device, patients are asked to describe their activities of daily living (ADLs) – this determines the order and range of grips pre-programmed into the hand. ADLs can include a wide range of activities, each with its own unique set of necessary grips. In our prior work, we collected a multi-sensor dataset for 6 ADLs (with 75 total sub-tasks or short-term behaviors) for 3 able-bodied, adult subjects [23], labeled using the pre-set time length allotted per sub-task. We also provided preliminary results for HDC-based ADL recognition using the EMG data. Lastly, a multi-layer control scheme was proposed (though not implemented) for assistive prosthetics. In this work, we expand the dataset to include another able-bodied, adult subject totalling 4 subjects and improve the HDC classifier over prior work. We then design and implement the full control scheme at each layer, modifying the prior proposed control scheme in the process to achieve higher performance.

Because HDC provides significant benefits for biosignal classification, prior work has explored its usage for other applications including robotics. Deep learning for robotics faces similar challenges as biosignal classification: limited training data, generalization, and in-device implementation [24]. Using hypervector representation and operations, prior work proposed a fully-HDC algorithm for robotic learning and recall of reactive behaviors (HDC RORB) to address these challenges [25], and further work exploring its design, utility and performance for navigation of a 2-D map showed that this algorithm can achieve better performance than a neural network with significantly better hardware efficiency [15]. The algorithm is even able to prioritize behavior of the agent by integrating prior knowledge about the task through intelligent sensor fusion of heterogenous input data streams. This property is what makes it attractive for the shared control aspect of prosthetic shared control.

Some examples of commonly used prosthetic grips in commercial devices include the power grasp, closed pinch, open pinch, key grip, and finger point with the pinch grips and power grasp being the most used overall [2]. For example, for the power grasp, all 5 digits begin to flex, surrounding the object of interest [26]. These grips are examples of the high-level prosthetic behavior outputted by the shared control layer. The full set of prosthetic behaviors used in this work was created based on the needs of the ADLs included in the dataset [23], commonly used prosthetic grips and gestures, and prosthetic responses during interactions with objects.

Prior work has successfully projected probability tables to HDC vectors [27]. While we use a traditional probabilistic model in this work, we make initial steps towards an HDC representation through taking input sensor data in the form of HDC vectors instead of raw data. Future work will fully replace this layer with an HDC representation making the full multi-layer control scheme majority HDC and able to take advantage of HDC properties like computational efficiency, low latency, and quick model updates.

III. Implementation

A. Overview

The multi-layer shared control scheme designed and built in this work is shown in Fig. 1. The lowest layer acquires feed forward sensor data: 59 EMG channels collected with a flexible sensor array and accelerometer data on the x, y and z axes, both sampled at 1kHz and pre-processed with a mean absolute value over 50ms windows. This information is used to determine the user's long-term goal (ADL) using an HDC classifier. A time series classifier recognizes the most recent sub-task completed by the user within the selected

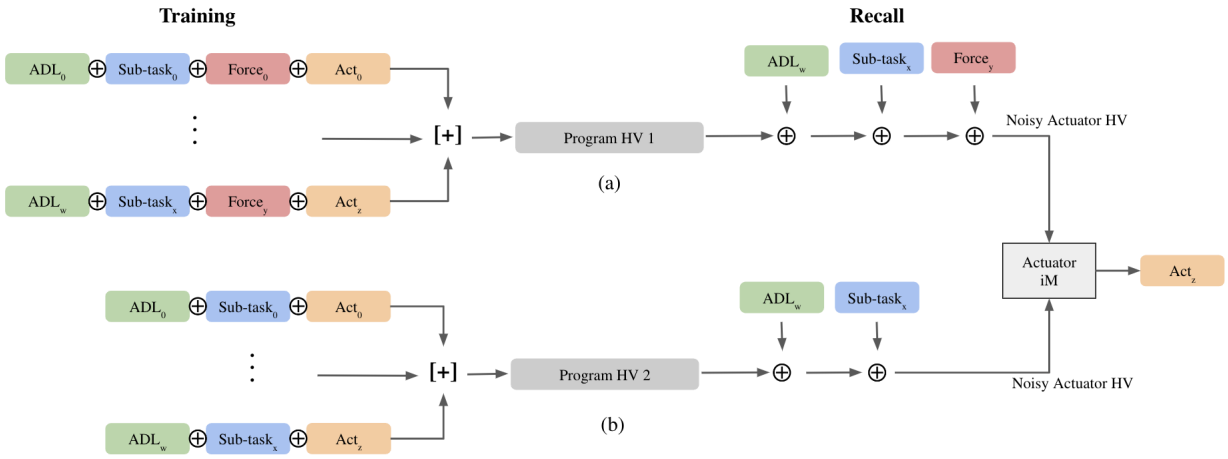

Fig. 2: The adopted HDC RORB algorithm with (a) the originally proposed algorithm combining all training data into a single PV, representing both the naive method and the version in which force-independent cases are pruned to prevent redundancy and (b) the added second PV for the force-independent subset of the data which results in higher recall performance.

ADL using the accelerometer data. Because the classification layer classifies on prior sensor data, it determines what has already happened and therefore can not direct actuation that corresponds to the user's actions at the present. To solve this, the probabilistic layer stores information on the user's habitual sequences of sub-tasks used to achieve their ADLs.

This information is used to predict the sub-task that the user intends to do at the present based on their most recent sub-task recognized by the time series classifier. This output is provided to the shared control layer as the human-driven control. It also uses feedback from force sensors on the fingers of the end-effector during actuation. The shared control is achieved through HDC RORB which is designed to intelligently use both the human-driven control and the feedback-based automatic control to determine the appropriate reactive high-level prosthetic actuation.

B. User Behavior Recognition

The lowest layer of the system recognizes the user's ADL. The HDC classifier used includes a spatial encoder, temporal encoder and associative memory as outlined in [23]. The overall classification window is 5 seconds long. The ADL recognition occurs primarily through the impact of arm position and gestures on the EMG signals. In this work, we shift the window over by 2 seconds, allowing for inclusion of EMG data during transitions into the sub-task at the beginning of the classification window.

ADL recognition provides long-term information to the system, but does not provide information on the short-term actuation the user is doing. While EMG primarily reflects arm position and hand gestures, accelerometer signals reflect arm movements and hence sub-tasks. To recognize the temporal accelerometer sub-task patterns, we used ROCKET, a method that transforms time series through random convolutional kernels and uses these transformations to train a linear classifier [28], enabling state-of-the-art performance with far less computational resources [29]. We train ROCKET on the full window (5s, 100 samples) of each training example and test on the remainder. The repetitive sub-tasks in ADLs like *opening jar* and *combing hair* such as *untwist* or *comb through hair* are given the same label. Through iteration, we included additional features for each channel on top of the raw data including standard deviation over a moving 10-sample window and jerk (change in acceleration).

C. User Intent Prediction

The second layer comprises a collection of Dynamic Bayesian Networks (DBNs), one for each ADL. The goal is to execute short-term intent prediction by taking into consideration the prior sub-task as well as additional periodic information extracted from the sensors, as shown in Fig. 1b which depicts a singe slice of the DBNs. Each DBN relates the sub-tasks probabilistically over time, for example in the *open jar* ADL, *untwist* (cap) follows (hand) *realign* most frequently, except when the user has finished unscrewing in which case *place cap* follows.

To address the challenge of correctly predicting future sub-tasks in ADLs with recurring behaviors, we introduce a binary indicator variable, denoted by $Sensors_{t-1}$ in Fig. 1b, that relies on HDC classification of EMG data to identify whether the user's repetitive sequence is likely to change. For example, for *open jar* in Fig. 1b, we identify through accelerometer time series classification that $subtask_{t-1}$ is *realign* and through HDC-encoded EMG data we find that this particular instance of *realign* is more likely to lead to *place cap* than to *untwist* as the EMG signals indicate that the user is now preparing to lift the cap rather than realign their hand for another untwist. Thus, the binary indicator enables selection of the appropriate conditional distribution, and inference is made. In general we infer future sub-tasks by evaluating the query $P(subtask_t|subtask_{t-1}, sensors_{t-1})$ on the relevant DBN at each time step.

D. Shared Control with Recall of Reactive Behavior

The previous layer predicts the current ADL as well as the sub-task. We combine these with feedback force sensors on

the end-effectors fingers to determine the desired actuation of the prosthetic device. This occurs via an HDC RORB scheme in which the ADL, sub-task and force information are all equally weighted in determining the actuation. The primary benefit of using HDC for this layer is the ability to fuse the various sensor input streams efficiently [12], [30].

The force data is processed to detect one of four states: no contact, touch, grip with slipping, or gripping. Each of these is mapped to a pseudo-orthogonal hypervector (HV) that gets bound (via the XOR operator $\oplus$) to the output of the ADL–sub-task binding and then to the associated actuation HV for that state. The fused vectors for the various ADL, sub-task, and force combinations get bundled together with equal weighting into a hypervector we call the program memory or program vector. This process is explained in more detail with several variations in [15]. However, there are several sub-tasks for which all 4 force states will have the same actuation. This occurs in force-independent sub-tasks, such as the L grip actuation which has a prescribed final position.

Given 75 unique sub-tasks across all ADLs and 4 force states, there are 300 sub-task-force-actuation combinations that get bundled together with equal weight into the Program Vector (PV) as in Eq. 1. The appropriate actuation can be extracted in real time by sequentially binding each of the HVs associated with input values to the PV. The resulting HV is a noisy approximation to one of the actuation HVs and the appropriate actuation HV can be extracted by comparison with an item memory. See Fig 2.

$$PV = \sum_{w,x,y,z} ADL_w \oplus S_x \oplus F_y \oplus Act_z \qquad (1)$$

However, this naive implementation doesn't take into consideration the 22 sub-tasks that are force-independent. Each of these sub-tasks will have 4 training samples in the current training scheme, 3 of which are redundant. Exploiting this, we store only the training sample for the no contact state for force-independent sub-tasks. Similarly, at test time we query only the no contact state for the force HV binding. This reduces the number of training samples from 300 to 234 which decreases the magnitude of noise in the recovered noisy actuation HV, increasing the probability that we can correctly identify the right entry in the item memory.

Given that each training sample in the bundle contributes equally to the noise during extraction, we can further increase our accuracy by adding a second PV. We split our model into two parallel channels for the force-dependent and independent sub-tasks. Our primary channel is the same as our original naive model in which we bind ADL, sub-task, force, and actuator HV together and bundle those together into a program vector (now PV1). However, for our force-independent sub-tasks, we bind only the ADL, sub-task, and actuation HV together – without the force HV – and bundle all of these into PV2. This further decreases the number of training samples stored in PV1 to 212, with only 22 samples stored in PV2. This third model alleviates noise during extraction at the expense of slightly more processing. We implement and evaluate all three models in Section IV.

There are 19 unique high-level actuation outputs that are passed to the end-effector ranging from *Wrist flex* to *2 finger closed pinch*. These are the simplest repeatable actuations we were able to define based on our ADLs and sub-tasks.

E. User-adaptivity

This system is designed to be able to adapt to users through personalization based on the user's activities of daily living, similar to commercial prosthetics. Though data was collected for and performance evaluated on a set of 6 ADLs, the control scheme can adjust to a new user's ADLs and sub-tasks. The probabilistic layer stores information on these sequences, along with the actuation mapping for the HDC RORB layer, both of which would require a one-time setup when the user is fitted for their prosthetic. The user would then need to provide demonstrations of each ADL/sub-task with corresponding EMG and accelerometer data in order to train the classifiers to recognize them. The dataset used in this work included 5 examples of each, though the classifiers can theoretically be trained on just one example instead of four. Future studies are needed to validate the limits of this.

IV. Experimental Results

A. EMG & Accelerometer Classification

The updated HDC EMG classifier resulted in improved performance over [23] including results for an additional subject. As shown in Fig. 3a, using leave-one-out cross-validation over 5 examples, every single subject achieves classification accuracy greater than 91% with an average accuracy of 92.5%. Generally, the performance of each subjects varies for each ADL. This is hypothesized to be due to different levels of consistency of overall arm position and gestures between examples of the same task. The time series classification of the accelerometer sensor data achieves an average sub-task recognition accuracy of 91.7% across all the subjects with leave-one-out cross-validation over 5 examples, shown in Fig. 3b. This algorithm also sees the same varied performance across tasks between different subjects which, in this case, is attributed to inconsistencies in the exact method of accomplishing each ADL between examples. Overall, the performance of both classifiers is very high and showcases the utility of both the EMG and accelerometer information in recognizing the user's behaviors.

B. Probabilistic Model

Fig. 4a shows the prediction accuracy of the DBN, calculated over the set of all sub-tasks for each ADL starting at $t = 1$. Note that, unlike time series sub-task classification, we consider repeated sub-tasks because one goal of this layer is to correctly predict future sub-tasks in ADLs with recurrent behavior despite the number of times this behavior is repeated. This allows the user to repeat the sub-task as needed (e.g. prediction must be successful whether the user opens a jar by twisting the cap 3, 4, 5, etc. times). Note also that the results in Fig. 4a are evaluated on the stand-alone DBN which assumes that the layer below correctly recognizes $subtask_{t-1}$ (integration of the two is discussed

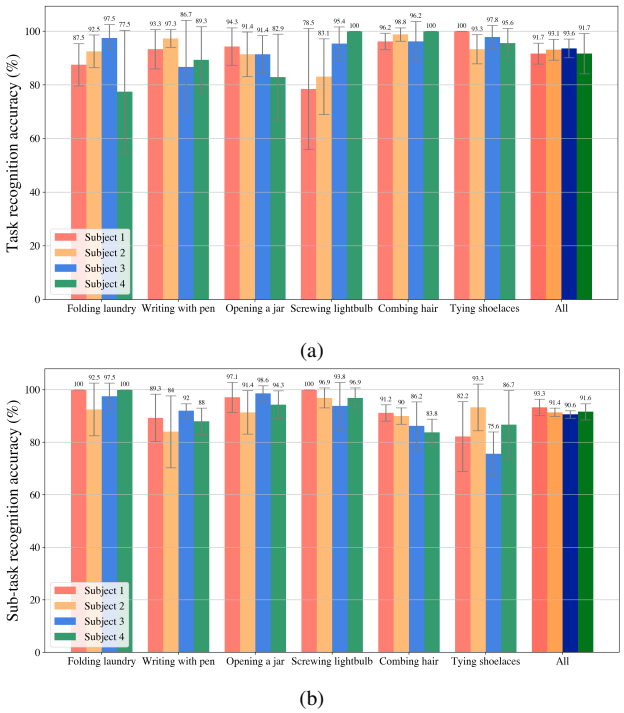

Fig. 3: (a) EMG-based task recognition accuracy with HDC per ADL and overall for the 4 subjects (b) sub-task recognition accuracy with a time-series classifier, given an accurate ADL from the HDC classifier. The error bars indicate standard deviation.

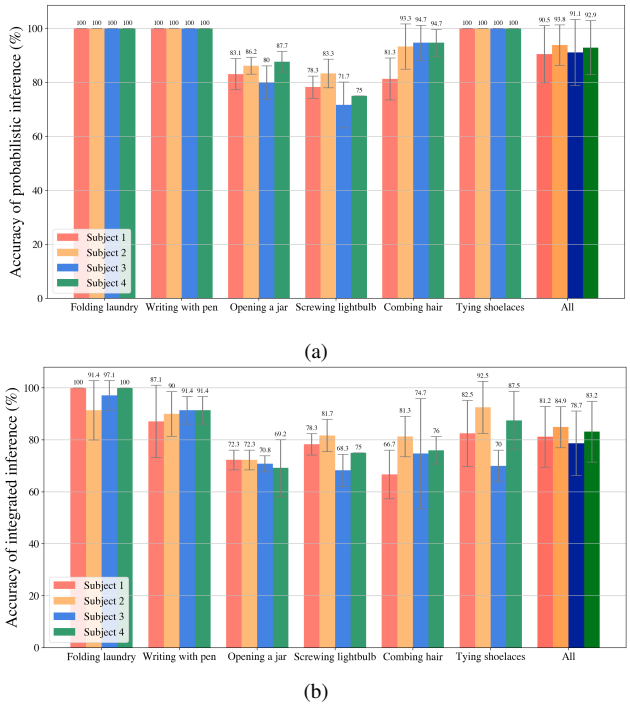

Fig. 4: Accuracy of sub-task prediction using the probabilistic layer (a) with ideal sub-task classifications (b) after integration of layers

in Section IV-D). The average prediction accuracy for the 4 subjects in a leave-one-out cross validation scheme over 5 trials is 92%. ADLs that present recurrent behaviour attain lower accuracy as the repetitions of this recurrence is non-deterministic, relying on the precision of the EMG-based indicator variable generator. Yet, accuracy remains higher than 80% on all ADLs but *screwing lightbulb*, where the quality of EMG data tends to be the least consistent across users due to the range of arm motions executed as seen in Fig. 3a. Conversely, the arm motions are beneficial in accelerometer time-series classification of sub-tasks, as seen in Fig. 3b, making for an interesting tradeoff that the probabilistic model aims to balance.

C. HDC Recall of Reactive Behavior

The HDC RORB control system first determines which PV to query based on a flag passed by the probabilistic layer which contains a map of all ADLs and their sub-tasks including their force dependency - manually labelled for the user's activities of daily living. We then bind the ADL and sub-task HVs to the appropriate PV. In the case where force feedback is not relevant, we are left with a noisy approximation of the actuation HV which we clean up via an item memory. When force feedback is necessary, we identify the appropriate force state: no contact, touch, grip with slipping, or grip, and bind that HV to the earlier HV. This produces a noisy actuation HV for further clean-up.

Fig. 5 shows the accuracy of our adopted HDC RORB method, as well as the two simpler models implemented, as a function of the dimensionality of our HD space. We highlight performance at a dimensionality of 10,000 due to its common usage in the HDC community [31]. The accuracy of 98.4% is the result of combining a 100% accurate force-independent model (PV2) and 97.9% accurate force-dependent model (PV1). The difference between PV1 and PV2 accuracies can be entirely attributed to the difference in total training sample vectors stored in each; 22 in PV1, 212 in PV2. 212 vectors evidently exceeds the storage capacity of 10,000-bit hypervector, demonstrated further by the two simpler models which store even more samples into their PVs. In terms of computational simplicity, though the 2-model method uses two PVs, the simpler methods would require a larger HD dimension to achieve the same accuracy which would cumulatively be much more expensive.

D. Integration

Fig. 4b shows the performance of the probabilistic model when considering the output of the accelerometer-based time series classifier from Section IV-A. Overall, the accuracy suffers a 10% average accuracy degradation, as this end-to-end implementation is affected by the error accumulating throughout the layers. Accuracy degradation is most severe for subject 4 in *tying shoelaces*, which indeed corresponds to the lowest time series classification accuracy shown in Fig. 3b. Overall, reliably predicting future sub-tasks with a limited amount of data (only 5 trials per ADL) remains the main challenge, which the hybrid probabilistic model has shown promising results towards. Closed-loop feedback as well as incorporating additional user-specific information may improve the performance of this model. Regarding the former, feedback from the force sensors will be used in a

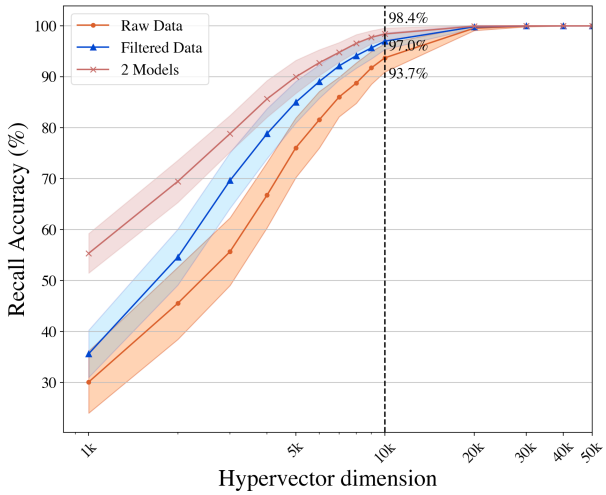

Fig. 5: Shared control prosthetic actuation recall accuracy for HDC RORB implemented naively, after eliminating redundancy, and with two shared control models. Shaded area shows standard deviation.

future iteration of this model to improve robustness of the EMG and accelerometer-dependent variables. Furthermore, once we move to the following time slice in the DBN, we can verify our prior prediction by comparing it with the accelerometer-based time series sub-task classification and correcting if necessary through, for example, an ensemble scheme. Towards adding user-specific information, an interesting future work direction would be to add a variable based on theories of mind commonly used in robotics [32] (such as the Boltzmann noisily-rational decision model [10], which assumes people tend to choose trajectories in proportion to some exponentiated cost), for recurrent behaviors. For example, the longer a person has combed their hair the less likely they are to continue. Finally, transfer learning could help address some of the intra-user data quality inconsistencies that can plague the overall performance of particular users.

E. Actuation

Given the appropriate high-level actuation from the HDC RORB layer, we are able to implement these actuations in simulation as well as with a physical robot hand. We are currently primarily using the simulation to model the end-to-end system including feedback. The simulation environment uses the PyBullet physics simulation along with the Shadow Dextrous Hand (Fig. 6a) to model a human hand and its movements [33]. This model has 20 actuated degrees of freedom and is able to replicate the movements of a human hand almost entirely (the human hand has 27 degrees of freedom [34]). Objects used in ADLs are instantiated from OpenAI's Roboschool Objects [35]. This allows for testing of shared control with the shadow hand and its interaction with a wide range of objects. The simulation environment provides force feedback at the contact points between objects mirroring the force sensors on the physical setup. The hardware being used for system prototyping is a DFRobot Rob0142 Robot hand (Fig. 6b) [36]. The Rob0142 has only 1 degree of freedom for each of the 5 fingers which is far fewer than both an

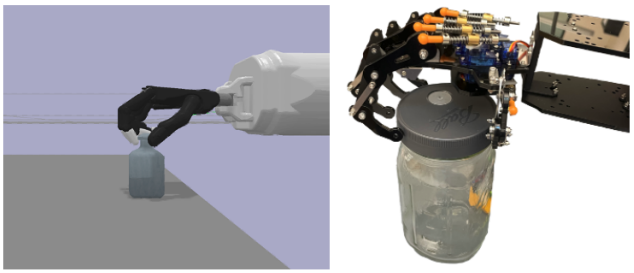

(a) (b)

Fig. 6: Actuation grips with (a) the Shadow Dexterous hand in a Pybullet simulation environment (b) the DFRobot Rob0142 hand.

actual hand as well as the hand in simulation. It's inclusion is in anticipation of data collection with amputees during which the hand can provide real-time rapid feedback to the patient during the training process.

V. Conclusion & Future Work

In conclusion, this work demonstrates the successful design of a multi-layer shared control scheme, primarily through HDC. Each layer was shown to achieve individual performance greater than 92%, however preliminary work on integration indicates error accumulation through the system. Future work will address mitigation of this through error correction in order to improve accuracy of the final actuation selection. Live experiments of the full control stack in simulation and then later with the robotic hand, along with on-device implementation to observe the overall hardware efficiency and latency are left as future work.

Given the prosthetic application, the dataset must also be expanded to amputee studies. By fitting a population of right radioulnar amputees with a socket containing the 64 electrode array, we can compare our control scheme to the amputees' existing device. To standardize the study, we would fit each subject with the same type of socket and myoelectric terminal device. We can also track endurance levels on a subject basis, as our shared control system aims to minimize the fatigue amputees undergo when controlling their device for extended periods of time. By expanding the research to the prosthetic wearing community, future devices can be worn for longer periods of time and provide the user with more accurate and intentional movements.

ACKNOWLEDGMENT

The authors thank Aayush Shah for his support. This work was supported by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program, and the VIP and HYDDENN programs sponsored by DARPA. This work was also supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1752814. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, DARPA or the U.S. government. Support was also received from Berkeley Wireless Research Center sponsors.

REFERENCES

[1] E. A. Biddiss and T. T. Chau, "Upper limb prosthesis use and abandonment: A survey of the last 25 years," *Prosthetics and orthotics international*, vol. 31, no. 3, pp. 236–257, 2007.

[2] A. Mohammadi, J. Lavranos, P. Choong, and D. Oetomo, "X-limb: A soft prosthetic hand with user-friendly interface," in *International Conference on NeuroRehabilitation*, Springer, 2018, pp. 82–86.

[3] F. Cordella, A. L. Ciancio, R. Sacchetti, *et al.*, "Literature review on needs of upper limb prosthesis users," *Frontiers in neuroscience*, vol. 10, p. 209, 2016.

[4] M. Espinosa and D. Nathan-Roberts, "Understanding prosthetic abandonment," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, SAGE Publications Sage CA: Los Angeles, CA, vol. 63, 2019, pp. 1644–1648.

[5] K. Z. Zhuang, N. Sommer, V. Mendez, *et al.*, "Shared human–robot proportional control of a dexterous myoelectric prosthesis," *Nature Machine Intelligence*, vol. 1, no. 9, pp. 400–411, 2019.

[6] C. H. Jang, H. S. Yang, H. E. Yang, *et al.*, "A survey on activities of daily living and occupations of upper extremity amputees," *Annals of rehabilitation medicine*, vol. 35, no. 6, pp. 907–921, 2011.

[7] M. Sharif, D. Erdogmus, C. Amato, and T. Padir, "Towards end-to-end control of a robot prosthetic hand via reinforcement learning," in *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*, IEEE, 2020, pp. 641–647.

[8] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: A survey," in *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, IEEE, 2020, pp. 737–744.

[9] D. Salwan, S. Kant, H. Pareek, and R. Sharma, "Challenges with reinforcement learning in prosthesis," *Materials Today: Proceedings*, vol. 49, pp. 3133–3136, 2022.

[10] M. Zurek, A. Bobu, D. S. Brown, and A. D. Dragan, "Situational confidence assistance for lifelong shared autonomy," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021, pp. 2783–2789.

[11] M. M. Churchland, J. P. Cunningham, M. T. Kaufman, *et al.*, "Neural population dynamics during reaching," *Nature*, vol. 487, no. 7405, pp. 51–56, 2012.

[12] A. Menon, A. Natarajan, R. Agashe, *et al.*, "Efficient emotion recognition using hyperdimensional computing with combinatorial channel encoding and cellular automata," *Brain informatics*, vol. 9, no. 1, pp. 1–13, 2022.

[13] A. Burrello, K. Schindler, L. Benini, and A. Rahimi, "One-shot learning for ieeg seizure detection using end-to-end binary operations: Local binary patterns with hyperdimensional computing," in *2018 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, IEEE, 2018, pp. 1–4.

[14] A. Moin, A. Zhou, A. Rahimi, *et al.*, "A wearable biosensing system with in-sensor adaptive machine learning for hand gesture recognition," *Nature Electronics*, vol. 4, no. 1, pp. 54–63, 2021.

[15] A. Menon, A. Natarajan, L. Galindez Olascoaga, Y. Kim, B. Braeden, and J. M. Rabaey, "On the role of hyperdimensional computing for behavioral prioritization in reactive robot navigation tasks," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2022.

[16] A. Rahimi, P. Kanerva, J. d. R. Millán, and J. M. Rabaey, "Hyperdimensional computing for noninvasive brain-computer interfaces: Blind and one-shot classification of eeg error-related potentials," in *10th EAI Int. Conf. on Bio-inspired Information and Communications Technologies*, 2017.

[17] A. Menon, D. Sun, S. Sabouri, *et al.*, "A highly energy-efficient hyperdimensional computing processor for biosignal classification," *IEEE Transactions on Biomedical Circuits and Systems*, 2022.

[18] A. Zhou, R. Muller, and J. Rabaey, "Memory-efficient, limb position-aware hand gesture recognition using hyperdimensional computing," *arXiv preprint arXiv:2103.05267*, 2021.

[19] A. Moin, A. Zhou, *et al.*, "An emg gesture recognition system with flexible high-density sensors and brain-inspired high-dimensional classifier," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, IEEE, 2018, pp. 1–5.

[20] A. Mitrokhin, P. Sutor, C. Fermüller, and Y. Aloimonos, "Learning sensorimotor control with neuromorphic sensors: Toward hyperdimensional active perception," *Science Robotics*, vol. 4, no. 30, eaaw6736, 2019.

[21] B. Cheung, A. Terekhov, Y. Chen, P. Agrawal, and B. Olshausen, "Superposition of many models into one," *Advances in Neural Information Processing Systems*, vol. 32, pp. 10 868–10 877, 2019.

[22] A. Mitrokhin, P. Sutor, D. Summers-Stay, C. Fermüller, and Y. Aloimonos, "Symbolic representation and learning with hyperdimensional computing," *Frontiers in Robotics and AI*, vol. 7, p. 63, 2020.

[23] A. Menon, L. I. G. Olascoaga, N. Shakouri, J. Ruffing, V. Balanaga, and J. M. Rabaey, "Brain-inspired multi-level control of an assistive prosthetic hand through emg task recognition," in *2022 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, IEEE, 2022, pp. 1–4.

[24] N. Sünderhauf, O. Brock, W. Scheirer, *et al.*, "The limits and potentials of deep learning for robotics," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 405–420, 2018.

[25] P. Neubert, S. Schubert, and P. Protzel, *Learning vector symbolic architectures for reactive robot behaviours*. Universitätsbibliothek Chemnitz, 2017.

[26] P. Slade, A. Akhtar, M. Nguyen, and T. Bretl, "Tact: Design and performance of an open-source, affordable, myoelectric prosthetic hand," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2015, pp. 6451–6456.

[27] L. I. Galindez Olascoaga, A. Menon, M. Ibrahim, and J. Rabaey, "A brain-inspired hierarchical reasoning framework for cognition-augmented prosthetic grasping," in *Combining Learning and Reasoning: Programming Languages, Formalisms, and Representations*, 2021.

[28] A. Dempster, F. Petitjean, and G. I. Webb, "Rocket: Exceptionally fast and accurate time series classification using random convolutional kernels," *Data Mining and Knowledge Discovery*, vol. 34, no. 5, pp. 1454–1495, 2020.

[29] B. Dhariyal, T. Le Nguyen, S. Gsponer, and G. Ifrim, "An examination of the state-of-the-art for multivariate time series classification," in *2020 International Conference on Data Mining Workshops (ICDMW)*, IEEE, 2020, pp. 243–250.

[30] E.-J. Chang, A. Rahimi, *et al.*, "Hyperdimensional computing-based multimodality emotion recognition with physiological signals," in *2019 IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, IEEE, 2019, pp. 137–141.

[31] P. Kanerva, "Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors," *Cognitive computation*, vol. 1, no. 2, pp. 139–159, 2009.

[32] C. Frith and U. Frith, "Theory of mind," *Current biology*, vol. 15, no. 17, R644–R645, 2005.

[33] M. T. Saeed, S. Khaydarov, B. L. Ashagre, and M. S. Zafar, "Comprehensive bond graph modeling and optimal control of an anthropomorphic mechatronic prosthetic hand," in *2019 IEEE International Conference on Mechatronics and Automation (ICMA)*, 2019, pp. 2006–2011. DOI: 10.1109/ICMA.2019.8816325.

[34] A. Rahman and A. Al-Jumaily, "Design and development of a bilateral therapeutic hand device for stroke rehabilitation," *International Journal of Advanced Robotic Systems*, vol. 10, no. 12, p. 405, 2013.

[35] G. Brockman, V. Cheung, L. Pettersson, *et al.*, "Openai gym," *CoRR*, vol. abs/1606.01540, 2016. arXiv: 1606.01540. [Online]. Available: http://arxiv.org/abs/1606.01540.

[36] *Bionic robot hand (right)*. [Online]. Available: https://www.dfrobot.com/product-1623.html.