Accelerating Constrained Continual Learning with Dynamic Active Learning: A Study in Adaptive Speed Estimation for Lower-Limb Prostheses

C. Johnson[1], J. Maldonado-Contreras[2,3], and A. J. Young[2,3]

2024 International Symposium on Medical Robotics (ISMR) | 979-8-3503-7711-8/24/$31.00 ©2024 IEEE | DOI: 10.1109/ISMR63436.2024.10585934

***Abstract*—Continual Learning is quickly emerging as a fundamental technique in almost all technical domains. This study develops its application in robotics, with a specific focus on transfemoral prosthetics, where machine learning models are fine-tuned in real-time to better predict ambulatory speed. This process of model adaptation faces several challenges stemming from the necessity of learning fast enough to keep up with real-time gait, while also ensuring sufficient accuracy and plasticity when encountering changing speeds and modalities. To address these challenges, we introduce Dynamic Active Learning (DAL) and Intermittent DAL (IDAL), novel frameworks which employ uncertainty-based sampling, as potential precursor steps to learning in this adaptation pipeline. Our contributions not only provide a robust guarantee that adaptation will occur within the time constraints posed by gait cycles, but also increase the rate of accuracy convergence by 51%, IDAL has been shown to attain a 4% lower post-convergence error rate, and maintain 30% more reliable post-convergence predictions compared to non-AL based methods of adaptation. In developing this system, we assessed numerous uncertainty metrics, finding that the Query by Committee method performs the best, attaining a Spearman Correlation Coefficient of 0.81 with ground truth error. While showcased through transfemoral prosthetics, our results illustrate the wide reaching potential of our DAL systems across diverse robotics applications.**

***Index Terms*—Lower-limb prosthetics, robotics, machine learning, active learning, uncertainty sampling.**

I. Introduction

THE field of robotics is experiencing a shift in paradigms, epitomized by Google's RT-2-X [1] – a leap in generalized robotic intelligence – that parallels OpenAI's groundbreaking transitions in natural language AI from specialized to generalized systems [2], [3], [4]. Central to this shift is the imperative for robotic systems to autonomously adapt, recalibrate, and evolve without constant human intervention or exhaustive retraining. Continuous Learning (CL) is rapidly gaining traction across diverse technical sectors [5], as enables it systems to incrementally assimilate new information without forgetting previously acquired knowledge – a paradigm particularly relevant for robotics where adaptability to dynamic environments is imperative [6], [7]. The primary concerns with the integration of CL into robotic systems revolve around the constraints surrounding computational efficiency and real-time data processing, both of which are crucial for robotic applications. Additionally, maintaining a high degree of accuracy, especially amidst evolving and unanticipated data streams, can be particularly taxing [8], [9]. Thus, the onus rests on developing methodologies that meet these stringent criteria while leveraging the recent achievements in deep learning [10], [11].

Active Learning (AL), characterized by strategic data selection, addresses many of the challenges in integrating CL [12]. Though beneficial in Natural Language Processing and Big Data, particularly via Human-in-the-Loop systems [13], [14], its potential in robotics remains largely unexplored [15], [16], [17]. Uncertainty sampling, a notable method in AL, selects only the most uncertain data points for learning, thus prioritizing data with the greatest novel insights [18], [19], [20], [21]. This approach not only ensures efficient model adaptation but also stands as a pillar in the robotic context where minimizing computational overheads and speeding up the learning process are paramount. This positions it as a crucial tool to bridge the significant gap faced by robotic systems: navigating rapidly changing environments while ensuring swift, adaptive responses [22], [23].

Transfemoral prosthetics highlight the need for AL in robotics due to challenges like varying user activities, terrains, and inter-subject variability [24]. These prosthetics demand performance across a broad and evolving spectrum of scenarios, emphasizing the need for continually adaptive responses [25], [26]. While their embedded microprocessors are advanced, they face computational challenges during real-time gait cycles. In such scenarios, not only is a rapid rate of adaptation necessary to keep up with changing conditions, but the speed and reliability of these adaptations is paramount [27], [28]. In this study, we focus on walking speed estimation, a facet of prosthetic control where a misprediction, whether due to delay or divergence, can result in the user stumbling, with potential injuries as consequences

[1]College of Computing, Georgia Institute of Technology, Atlanta, GA, 30332, USA.

[2]Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA, 30332, USA.

[3]Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA, 30332, USA.

This work was supported by J. Maldonado-Contreras' Ford Foundation Fellowship, the Department of Defense Congressionally Directed Medical Research Programs (CDMRP) Award No. W81XWH-17-1-0031, and the National Institutes of Health (NIH) Award No. DP2-HD111709.

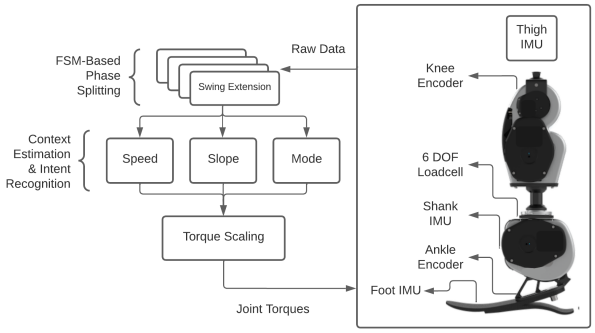

Fig. 1: OSL overview - sensors, three-tiered control structure

[29]. To estimate speed, we use an adaptive pipeline that progressively learns subject-dependent walking patterns after subject-independent initialization. AL's ability to rapidly and reliably select crucial data points makes its integration into this system a salient use case. In this study, we introduce novel AL paradigms, demonstrating their enhanced reliability and speed compared to traditional CL methods.

II. Methods

A. Prosthetic System

This study utilized the Univeristy of Michigan's Open Source Leg (OSL) [11], [30] and Georgia Tech's Prosthetic Intelligent Controls (EPIC) Lab. This advanced knee-ankle prosthesis uses a six-degree-of-freedom (DoF) loadcell, three six-DoF IMUs situated on the thigh, shank, and foot, and a pair of joint encoders and actuators.

For level walking (LW), the peak ankle plantarflexion moment was identified using a linear regression model from an able-bodied dataset. As walking speed increases, the moment escalates with a factor of $0.4898\,\frac{\text{Nm}}{\text{kg}\cdot\text{m/s}}$, equating to $\frac{42.2}{\text{m/s}}$ at a speed of $0.5\,\text{m/s}$, which was hand-tuned [31]. To adjust the baseline prosthetic ankle torque during plantarflexion, this factor was utilized to augment the impedance control stiffness k_s according to the current walking speed v (in $\frac{\text{m}}{\text{s}}$) [32]:

$$k_s = k(1 + 0.422(v - 0.5)) \tag{1}$$

The prosthesis operates on a three-tiered control structure: high, mid, and low, as illustrated in Fig. 1. Using onboard sensor data coupled with learning-based algorithms, the high-level control focuses on identifying the user's intent (mode classification) and making context estimations — walking pace and slope angle regression. The mid-level controller uses a finite state machine (FSM) to handle the transitions in different walking styles and the gait's specific phases, adjusting the ankle and knee joint torques as required. Finally, the low-level control translates the determined joint torques into commands for the actuators building off of (1), using the relation

$$\tau_i = -k_i(\theta_i - \theta_{ei}) - b\theta_i^{'} \tag{2}$$

where i, θ_i, and $\theta_i^{'}$, are the joint in question, angle, and angular velocity given by onboard encoders, respectively, k is the stiffness, b is the damping coefficient, and θ_{ei} is the target angle specified by the FSM.

In the process, the high-level controller's context estimations are processed using the Kalman filter. This filtering technique is well-known for providing enhanced approximations by recursively estimating the joint probability distribution over undetermined variables for every specific timeframe. Specifically, the Kalman Gain, given by

$$K_n = \frac{P_{\text{prior}}}{P_{\text{prior}} + P_{\text{meas}}} \tag{3}$$

where K_n is the Kalman Gain for the nth timeframe, P_{prior} is the prior variance, and P_{meas} is the measurement variance, adjusts the weight of the incoming measurements. Once the context has been filtered, impedance parameters are adjusted, translating to the desired torque values from (2). These are subsequently managed by low-level controllers and conveyed to the actuators. Incorporating PID controllers within these actuators ensures the leg receives appropriately scaled assistance [33], [34].

B. Dataset Overview

Nine individuals with transfemoral amputations (average $\pm$ SD; 7 males and 2 females, age 50.36 ± 12.09 years; weight = 80.13 ± 15.64 kg; height = 1.77 ± 0.10 m) took part in the research and underwent a bifurcated trial: treadmill walking at fixed speeds (ranging from $0.3\frac{\text{m}}{\text{s}}$ to $0.9\frac{\text{m}}{\text{s}}$ in $0.1\frac{\text{m}}{\text{s}}$ increments, which act as data labels) and dynamic speeds following triangular and step-wise patterns, identical to those presented in [25]. On average, participants took 171 steps (STD: 21), with approximately 80% at the fixed speed settings. We trained nine distinct groups of models, each excluding a different subject's data. Then, for each of these model groups, we evaluated it on the data from the participant on which it had not been trained. This validation technique ensured our system's performance when transitioning from subject-independent to subject-dependent scenarios.

C. Adaptation Pipeline

TABLE I: TCN model parameters - architecture parameters (left table), training parameters (right table).

Architecture	Value
kernel size	5
dropout	0.2
effective hist	120
# channels	$[10, 10, 10, 10]$

Training	Value
input size	28
output size	1
LR	$1e^{-4}$
epochs	2

The adaptive pipeline integrates both a forward predictor in order to enable real-time speed predictions, scaling respective torque parameters, while facilitating batched refinement of the predictive models using a backwards estimator which, in this paper, uses ground truth. This system is presented in [25], with modifications in the forward predictor and the insertion of the AL subprocess, which this study focuses on.

The forward predictor in this paper adopts a Temporal Convolutional Network (TCN) [35] architecture with distinct

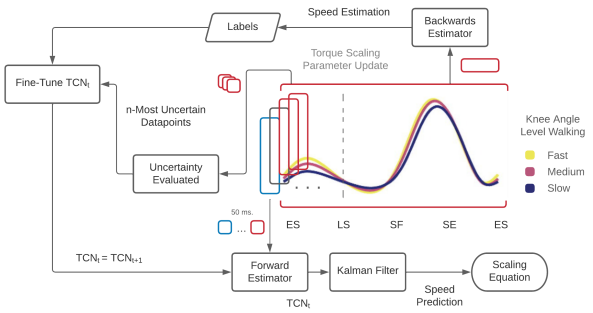

Fig. 2: Adaptation pipeline overview - forward predictor, backwards estimator, fine-tuning system, uncertainty sampling

models being trained on and used to predict speeds in each phase: Early Stance (ES), Late Stance (LS), Swing Flexion (SF), and Swing Extension (SE). These models make predictions using 120 segments of data, with 20ms overlap between successive windows, optimizing prediction accuracy. The TCN utilizes the architecture and hyperparameters found in Table 1 [36]. Finally, this study uses ground truth as a stand-in for the backwards estimator subprocess for simplicity.

Fig. 2 shows the adaptive pipeline managing the forward predictor and adaptive components concurrently. Sensor data is segmented in 120ms intervals, driving intent recognition and context estimation to modulate torque at 50Hz. Following batches of three strides, the backwards estimator labels the accumulated gait data, thereby providing more reliable speed labels used to fine-tune the forward predictor models.

D. Active Learning System

TABLE II: Speed ($\frac{\text{m}}{\text{s}}$) vs. d_{adapt} (sec), used to compute time budget at each walking speed

β	0.3	0.4	0.5	0.6	0.7	0.8	0.9
d_{adapt}	1.69	1.56	1.52	1.42	1.41	1.36	1.33

1) Time Budget Computation: AL is an ML paradigm where the model selectively queries the most informative data points from an unlabeled dataset to be labeled, optimizing the learning process. When a predetermined number of queries is set, this is referred to as a "query budget," constraining the total number of data points the model can request for labeling [37], [38]. In the use case presented here, we use a time budget for each adaptation window, demonstrated in Fig. 3. We compute this at every adaptation cycle (3 strides) using the highly accurate estimation given by the backwards estimator. We compute the maximum allowable adaptation duration ($d_{adapt}(t)$, at time t) per speed in $\frac{\text{m}}{\text{s}}$ through taking the average 3-stride completion time by speed, shown in Table 2. We will refer to the function mapping speeds (β) to d_{adapt} as $T(\beta)$.

2) Uncertainty Sampling: In AL, one prominent strategy for selecting the most informative data points is called Uncertainty Sampling, in which, the model queries the instances about which it is most uncertain, often those closest to the decision boundary, see Algorithm 1 [39]. In accordance with

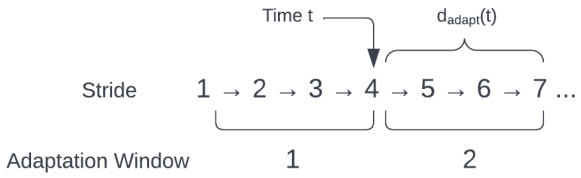

Fig. 3: Time budget example - at time $t = 4$, time budget $d_{adapt}(t)$ is the duration of strides 4-7

the theory of Supervised Learning, the underlying assumption is that by labeling these uncertain (which is a surrogate for error) instances, the model can achieve more significant improvements in performance with fewer labeled samples [40], [41].

Algorithm 1 Uncertainty Sampling

Require: U (initial pool), D (training data), θ (model), B (budget), s (uncertainty metric)

$b \leftarrow 0$
while $b < B$ **and not** $U = \emptyset$ **do**
  $i^* \leftarrow \text{argmax}_{i \in U} s(\theta, i)$
  Query label y^* for datapoint i^*
  $D \leftarrow D \cup (i^*, y^*)$, $U \leftarrow U \setminus i^*$
  Update θ with D
  $b \leftarrow b + 1$
end while
return U, D, θ

3) Uncertainty Metrics: This prompts the question of what the optimal uncertainty metric $s(\theta, x)$ is. To determine this, we assessed four candidate methods, using the Spearman Correlation Coefficient as the evaluation criterion to measure the relationship between the uncertainty trends and the ground truth errors. This choice of criterion is grounded in the fundamental theories of supervised learning, as a strong correlation indicates that as our model's uncertainty increases, so does its likelihood of making errors, reflecting the intuitive notion that areas of higher uncertainty are typically those where the model lacks sufficient knowledge or information [42].

First, we evaluated the traditional Query by Committee method [43], [44]. We trained n=3 forward predictor models with restricted datasets. This restriction of $k\%$ causes differences between the models and thus hopefully meaningful disagreement indicative of uncertainty. We first tuned the restriction parameter k and furthermore evaluated seven different disagreement measures.

The disagreement measures used in the Query by Committee Approach serve to quantify the level of disagreement among the n models. First, the coefficient of variation is calculated as the ratio of the standard deviation to the mean of the predictions given by the models, defined as $\text{CV} = \frac{\sigma}{\mu}$. The mean absolute deviation is the average of absolute differences from the mean inference, represented as $\text{MAD} = \frac{1}{n}\sum_{i=1}^{n} |x_i - \bar{x}|$. Variance measures the dispersion

of the predictions around their mean inference, calculated as $\sigma^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^2$, offering a normalized measure of dispersion. The range is simply the difference between the maximum and minimum predictions of the models.

For n models making predictions, each method provides a unique way to quantify the level of disagreement between them. Specifically, for each unique pair (i, j) among the n models, the disagreement is measured by a function F such that:

$$\text{Pairwise Disagreement} = \frac{1}{\binom{n}{2}} \sum_{i=1}^{n} \sum_{j=i+1}^{n} F(i,j). \quad (4)$$

For the Kullback-Leibler (KL) Divergence, denoted as P-KL: $F(i,j) = \text{KL}(P_i||P_j) = \sum_x P_i(x) \ln\left(\frac{P_i(x)}{P_j(x)}\right)$, where the sum runs over all speed $0.1\frac{\text{m}}{\text{s}}$ bins x. For the Joint Entropy, denoted as P-H: $F(i,j) = H(V_i, V_j) = -\sum p(v_i, v_j) \ln p(v_i, v_j)$, where the sum runs over all vote outcomes in $.1\frac{\text{m}}{\text{s}}$ intervals. For the Euclidean Distance, denoted as P-D: $F(i,j) = D(X_i, X_j) = \sqrt{\sum_{k=1}^{m}(x_{ik} - x_{jk})^2}$, where the sum runs over all dimensions of the vectors [19], [45], [46].

After QBC, we explored distance metrics, specifically the Euclidean and Mahalanobis Distances [47], [48]. In the interest of enhancing computational efficiency, to enable real-time use, instead of calculating the distance from a datapoint to its nearest neighbor, we determined the distance to the nearest centroid. This centroid represents the pre-computed average center of all data points associated with a specific speed. For the Euclidean Distance, this involved measuring the straight-line distance between the datapoint and the centroid within the feature space, mathematically defined as

$$d_{\text{euclidean}}(x, c) = \sqrt{\sum_{i=1}^{n}(x_i - c_i)^2} = \sqrt{(x-c)^T(x-c)} \quad (5)$$

where n denotes the number of data dimensions or features. On the other hand, the Mahalanobis Distance provided a more context-aware measure, taking into account both the variability and correlation of the data. It's represented by the formula

$$d_{\text{mahalanobis}}(x, c) = \sqrt{\sum_{i=1}^{n}\sum_{j=1}^{n}(x_i - c_i)\Sigma_{ij}^{-1}(x_j - c_j)}$$
$$= \sqrt{(x-c)^T\Sigma^{-1}(x-c)} \quad (6)$$

wherein Σ is the covariance matrix of the data. The value of the Mahalanobis Distance lies in its consideration of the covariance structure, making it particularly potent when dealing with features that aren't independent or vary in scales.

Finally, we assessed the effectiveness of using the difference between backwards estimators and forward predictors as an uncertainty metric. Backwards estimators are computed using $v = \frac{d_{hc}}{t}$, with d_{hc} representing the Euclidean distance between heel contacts and t representing the stride duration, measured by a foot-mounted IMU from Navigation Solutions LLC [49], [50], [51], [52]. This backwards estimation method, computed post-stride, demonstrated significantly higher accuracy compared to the forward predictor, as indicated in [25], and thus their discrepancy is reasoned to be a performant uncertainty metric.

4) Dynamic Active Learning System: In our Dynamic Active Learning (DAL) algorithm, the primary challenge was determining the appropriate query budget in terms of both datapoint number and epoch number during model retraining. We begin with a foundational setting of the epoch number to 2. Using the relationship between datapoint count (d), epoch count (ϵ), and the resulting adaptation time, $R(d, \epsilon) = 1.113e^{-4}d + .108\epsilon - 1.1673$, which was derived through a grid search, we can fix epoch count and directly compute the maximum number of datapoints we can process without exceeding the maximum allowable time. Finally, if processing all data points within the 3-stride adaptation window stays within the time limit, we gradually raise the epoch count to improve model accuracy, effective up to about 10 epochs.

Because we utilize a time-dependent model – a TCN – we need to ensure temporal coherence in the dataset, and thus cannot simply select the most uncertain data points, since there is no guarantee of temporal coherence. Therefore, we use a hybrid sampling method which integrates both the uncertainty of a data point and its temporal position to guide the AL process for the training of the TCN. For a given data point, its uncertainty is denoted by $s(x)$. To maintain temporal coherence in the selected samples, a temporal penalty function is introduced, defined as

$$p(x) = \min_{y \in \text{selected}} |x - y|, \quad (7)$$

where y represents the indices of previously selected points. The combined metric for sampling is then given by:

$$H(\theta, x, \alpha) = s(\theta, x) - \alpha p(x), \quad (8)$$

where α is a weighting factor to balance uncertainty and temporal coherence, here set to 5. Data points are ranked based on this combined metric, and those with the highest ranks are used for fine-tuning, ensuring that both informativeness and temporal structure are considered during the selection process.

This dynamic adjustment mechanism ensures that we leverage the most out of our available time budget, either by accommodating more data points or allowing for more training epochs, ensuring an optimal balance between data ingestion and model refinement. This novel system is presented in general form in Algorithm 2.

5) Intermittent Dynamic Active Learning System: Using DAL offers data efficiency in model training, but recent research points to accuracy loss in deep learning due to AL-induced localized overfitting. This anomaly arises from the hyper-constriction of neuron localization during training, causing the model to over-specialize [41]. Recognizing this

Algorithm 2 Temporally-Coherent Dynamic Active Learning

Require: TrainingData D, RegressionModel θ
Require: MaxTimeRelation T, RuntimeRelation R, UncertaintyModel H
Require: speed β, UncertaintyPool U
Maximum Allowable Time for Adaptation $T_{\max} \leftarrow T(\beta)$
Epochs $\epsilon \leftarrow 2$
Derive $D_{\max}$ solving $T_{\max} = R(D_{\max}, 2)$
$\alpha \leftarrow 5$
For each $i \in U$:
    $L \leftarrow H(\theta, i, \alpha)$
For $i = 1$ to $D_{\max}$ **and** while $U \neq \emptyset$:
    $i^* \leftarrow \text{argmax}_{i \in U} L$
    $D \leftarrow D \cup \{i^*\}$
    $U \leftarrow U \setminus \{i^*\}$
    TimeEstimate $\leftarrow R(|D|, \epsilon)$
    If TimeEstimate $> T_{\max}$: **Break**
If $U = \emptyset$:
    Do
        $\epsilon \leftarrow \epsilon + 1$
        TimeEstimate $\leftarrow R(|D|, \epsilon)$
    While TimeEstimate $\leq T_{\max}$
Train θ with D for ϵ epochs
return Updated U, D, and Trained θ

challenge, our initial AL methodology capitalizes on the inherent data variance present at the inception of a trial, selectively refining the most pertinent neural segments and thereby theoretically suppressing error surges through reduced network noise. As trials continue, we anticipated increased post-convergence errors compared to non-AL methods. To rectify this, we introduce a derivative approach: Intermittent DAL (IDAL). Herein, DAL is exclusively employed post-speed transitions – a period characterized by a moderate spike in dataset variance akin to trial commencements. By harnessing the strengths of both localized neural training and the holistic benefits of standard adaptation, IDAL strikes a balance, effectively mitigating the pitfalls of over-specialization while maximizing model robustness and adaptability across various speeds.

E. Offline Analysis

1) Uncertainty Metric & Disagreement Measure Optimization: We assessed the proposed uncertainty metrics and disagreement measures for robustness by correlating them with actual errors using the Spearman Correlation Coefficient. In the Query by Committee approach, we examined training set overlaps from 0% to 100% in 10% increments to identify the most effective overlap for detailed analysis of seven disagreement measures, aiming to find the metric of uncertainty that acted as the best proxy for prediction errors.

2) Convergence Rate, Residuals, Post-Convergence Accuracy Analysis: We assessed our AL-based methods' performance by fitting their errors with quadratic curves and calculating residuals

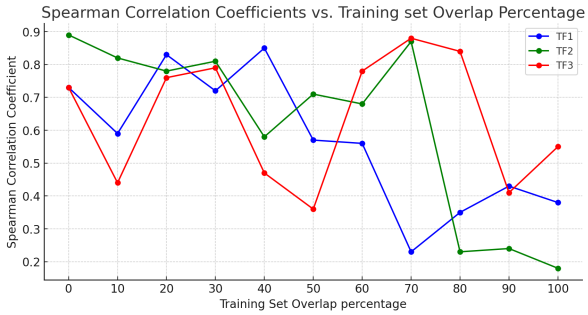

Fig. 4: QBC model training set overlap analysis - Spearman Correlation Coefficient vs. Overlap Percentage

$$r_i = |y_{actual}(x_t) - y_{fitted}(x_t)| \tag{9}$$

to measure reliability. The convergence rate, reflecting adaptation speed, was noted when the fit line came within a 10% tolerance of its asymptote. Post-convergence accuracy, indicating potential for optimal steady-state accuracy, was determined by averaging errors after convergence occured.

3) Time Constraint Satisfaction: We evaluated the efficiency of AL-based methods (DAL, IDAL) against adaptation time constraints by introducing a synthetic triangular time constraint, ranging from 0.1 to 1.5 seconds, in a simulated trial. This dynamic constraint tested the adaptability of our methods. We compared the AL methods' performance with a standard adaptation [25] system to gauge their relative proficiency in meeting dynamic time constraints.

III. Results

A. Uncertainty Metric Analysis

1) Uncertainty Metrics: We obtain Spearman Correlation Coefficients with Ground Truth Error as 0.22, 0.48, 0.57, and 0.75 for Mahalanobis Distance, Backwards Estimator Discrepancy, Euclidean Distance, and QBC, respectively. Both Euclidean and Mahalanobis Distance metrics compute approximately 2.6 times too slowly on average for efficacious real-time implementation. Therefore, QBC is the most effective metric, warranting further optimizations in its design.

TABLE III: Disagreement Measure vs. Spearman Correlation Coefficient (r)

Measure	CV	P-KL	Range	P-H	P-D	MAD	σ^2
r	-0.471	-0.185	0.173	0.330	0.571	0.609	0.807

2) QBC Optimization: Optimal training data overlap for the $n = 3$ ensemble models is determined to be 20%, as evidenced by evaluations on data from three transfemoral amputee subjects (see Fig. 4). Moreover, the variance disagreement measure emerges as the most effective, achieving an average Spearman Correlation Coefficient of 0.773 (refer to Table III).

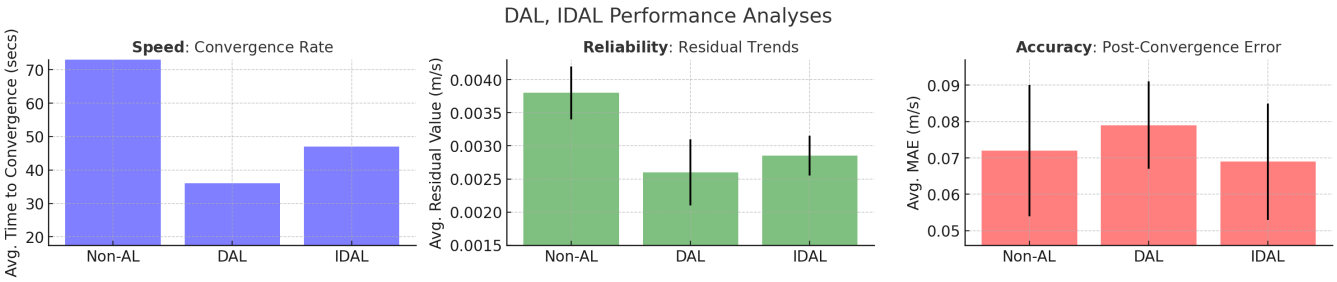

Fig. 5: DAL, IDAL performance analyses - cross subject average convergence rate, residual trends, post-convergence error. Error bars represent interquartile range, and are not shown for convergence rate due to inconsistency in trial duration.

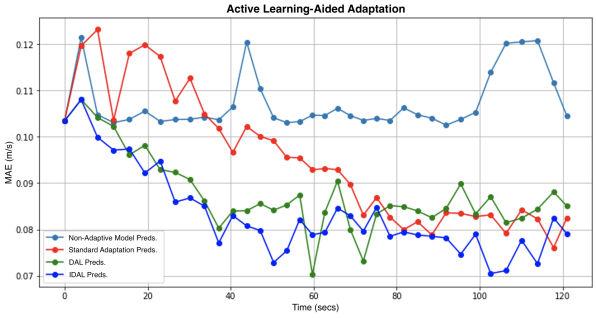

Fig. 6: Convergence trends - Data from TF1 used to demonstrate sweep used to determine convergence rate, residuals, and post-convergence accuracy.

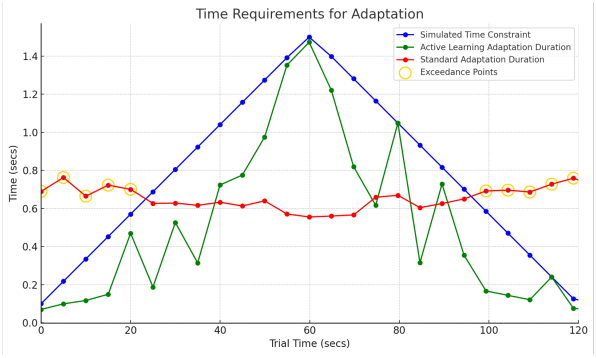

Fig. 7: Constraint satisfaction analysis - the DAL and non-AL method performance under synthetic time constraint

B. Convergence Rate, Residuals, Post-Convergence Accuracy Analysis

The results in Fig. 5 are the averages of the trial shown in Fig. 6 over all subjects considered in the trial. We compute time to convergence, residuals, and post-convergence accuracies using fitted quadratic lines. On average, DAL converges 51.35% and IDAL 49.8% faster than the non-AL method. Residual analysis shows DAL and IDAL achieve 29.53% and 28.17% lower residuals, respectively, than non-AL. Finally, DAL has 7.4% lower while IDAL has 3.95% higher post-convergence accuracy than the non-AL method.

C. Constraint Satisfaction Analysis

In the synthetic constraint trial, increasing epochs proves beneficial when time permits. Upon sweeping all subjects, similar to as shown in Fig. 7 for TF1, we find that the non-AL method averages eight exceedance points compared to zero for DAL.

IV. Discussion

AL methodologies facilitate data efficiency but can cause localized overfitting due to the hyper-constriction of neuron localization during training [41]. IDAL mitigates this over-specialization by maintaining model flexibility and adaptability, as evidenced by the results in Fig. 5. These results demonstrate IDAL's effectiveness, where it not only converges rapidly, almost paralleling DAL at 49.8% faster than non-AL methods, but also achieves lower residuals and enhances post-convergence accuracy by 3.95% compared to non-AL.

The study is limited by potential discrepancies between offline and real-time performance from the offline use of 2GHz Quad-Core Intel Core i5 and Core ARM Cortex-A57 MPCore in real-time, which may influence DAL's time budget calculations. Additionally, real-time adaptation will be impacted by imperfect backwards estimates.

Traditional AL methods risk introducing biases by focusing on high-uncertainty data points, potentially reducing model generality and robustness [41], [53]. A potential novel solution to this uses the Beta distribution to probabilistically modulate the selection of datapoints during uncertainty sampling [54]. Under this paradigm, high-uncertainty points have a higher selection likelihood, but low-uncertainty points are not excluded.

V. Conclusion

Our DAL framework, as demonstrated in transfemoral prosthetics, provides a robust solution to the challenges of CL in robotics, boosting the rate of accuracy convergence by 51% and maintaining 30% more reliable post-convergence predictions. We attained a Spearman Correlation Coefficient of 0.81 between our tuned QBC uncertainty metric and error, illustrating the efficacy of our approach when coupled with Uncertainty Sampling. Moreover, IDAL further strengthens the system, delivering a 4% lower post-convergence error rate than non-AL methods. Collectively, this research demonstrates the tangible contribution that DAL has to CL in robotics.

REFERENCES

[1] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Lisa Lee, Tsang-Wei Edward Lee, Sergey Levine, Yao Lu, Henryk Michalewski, Igor Mordatch, Karl Pertsch, Kanishka Rao, Krista Reymann, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Pierre Sermanet, Jaspiar Singh, Anikait Singh, Radu Soricut, Huong Tran, Vincent Vanhoucke, Quan Vuong, Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Jialin Wu, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-2: Vision-language-action models transfer web knowledge to robotic control, 2023.

[2] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with gpt-4, 2023.

[3] Sascha Freyberg and Helmut Hauser. The morphological paradigm in robotics. *Studies in History and Philosophy of Science*, 100:1–11, 2023.

[4] Víctor Mayoral, Alejandro Hernández, Risto Kojcev, Iñigo Muguruza, Irati Zamalloa, Asier Bilbao, and Lander Usategui San Juan. The shift in the robotics paradigm — the hardware robot operating system (h-ros); an infrastructure to create interoperable robot components. *2017 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*, pages 229–236, 2017.

[5] German Ignacio Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *CoRR*, abs/1802.07569, 2018.

[6] Haoxuan Qu, Hossein Rahmani, Li Xu, Bryan Williams, and Jun Liu. Recent advances of continual learning in computer vision: An overview, 2021.

[7] Sayantan Auddy, Jakob Hollenstein, Matteo Saveriano, Antonio Rodríguez-Sánchez, and Justus Piater. Continual learning from demonstration of robotics skills, 2023.

[8] Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz-Rodríguez. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges, 2019.

[9] Rachel Gehlhar, Maegan Tucker, Aaron J. Young, and Aaron D. Ames. A review of current state-of-the-art control methods for lower-limb powered prostheses. *Annual Reviews in Control*, 55:142–164, 2023.

[10] Xiangli Nie, Zhiguang Deng, Mingdong He, Mingyu Fan, and Zheng Tang. Online active continual learning for robotic lifelong object recognition. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–15, 2023.

[11] A. F. Azocar, L. M. Mooney, J. F. Duval, et al. Design and clinical implementation of an open-source bionic leg. *Nature Biomedical Engineering*, 4:941–953, 2020.

[12] Burr Settles. Active learning literature survey. Technical Report TR1648, University of Wisconsin-Madison Department of Computer Sciences, 2009.

[13] Samuel Budd, Emma C. Robinson, and Bernhard Kainz. A survey on active learning and human-in-the-loop deep learning for medical image analysis. *Medical Image Analysis*, 71:102062, 2021.

[14] Zimo Liu, Jingya Wang, Shaogang Gong, Huchuan Lu, and Dacheng Tao. Deep reinforcement active learning for human-in-the-loop person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.

[15] Marcos Maroto-Gómez, Sara Marqués-Villaroya, José Carlos Castillo, Álvaro Castro-González, and María Malfaz. Active learning based on computer vision and human–robot interaction for the user profiling and behavior personalization of an autonomous social robot. *Engineering Applications of Artificial Intelligence*, 117:105631, 2023.

[16] Ian Abraham and Todd D. Murphey. Active learning of dynamics for data-driven control using koopman operators, 2019.

[17] Ali Ayub and Carter Fendley. Few-shot continual active learning by a robot. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.

[18] David D. Lewis and Jason Catlett. Heterogeneous uncertainty sampling for supervised learning. In William W. Cohen and Haym Hirsh, editors, *Machine Learning Proceedings 1994*, pages 148–156. Morgan Kaufmann, San Francisco (CA), 1994.

[19] V. L. Nguyen, M. H. Shaker, and E. Hüllermeier. How to measure uncertainty in uncertainty sampling for active learning. *Machine Learning*, 111:89–122, 2022.

[20] Shang Liu and Xiaocheng Li. Understanding uncertainty sampling, 2023.

[21] Jingbo Zhu, Huizhen Wang, Benjamin K. Tsou, and Matthew Ma. Active learning with sampling by uncertainty and density for data annotations. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(6):1323–1331, 2010.

[22] Annalisa T. Taylor, Thomas A. Berrueta, and Todd D. Murphey. Active learning in robotics: A review of control principles. *Mechatronics*, 77:102576, 2021.

[23] Nutan Chen, Alexej Klushyn, Alexandros Paraschos, Djalel Benbouzid, and Patrick van der Smagt. Active learning based on data uncertainty and model sensitivity, 2018.

[24] Krishan Bhakta, Jonathan Camargo, William Compton, Kinsey Herrin, and Aaron Young. Evaluation of continuous walking speed determination algorithms and embedded sensors for a powered knee ankle prosthesis. *IEEE Robotics and Automation Letters*, 6(3):4820–4826, 2021.

[25] C. Johnson, J. Cho, J. Maldonado-Contreras, S. Chaluvadi, and A. J. Young. Adaptive lower-limb prosthetic control: Towards personalized intent recognition context estimation. In *2023 International Symposium on Medical Robotics (ISMR)*, pages 1–7, 2023.

[26] Cesar H. Guzmán-Valdivia Carlos M. Lara-Barrios, Andrés Blanco-Ortega and Karla D. Bustamante Valles. Literature review and current trends on transfemoral powered prosthetics. *Advanced Robotics*, 32(2):51–62, 2018.

[27] A. M. El-Sayed, N. A. Hamzaid, and N. A. Abu Osman. Technology efficacy in active prosthetic knees for transfemoral amputees: a quantitative evaluation. *The Scientific World Journal*, 2014:297431, Jul 2014.

[28] Frank Sup, Amit Bohara, and Michael Goldfarb. Design and control of a powered transfemoral prosthesis. *The International Journal of Robotics Research*, 27(2):263–273, 2008. PMID: 19898683.

[29] Krishan Bhakta, Jonathan Camargo, Luke Donovan, Kinsey Herrin, and Aaron Young. Machine learning model comparisons of user independent dependent intent recognition systems for powered prostheses. *IEEE Robotics and Automation Letters*, 5(4):5393–5400, 2020.

[30] Alejandro F. Azocar, Luke M. Mooney, Levi J. Hargrove, and Elliott J. Rouse. Design and characterization of an open-source robotic leg prosthesis. In *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, pages 111–118, 2018.

[31] AJ Young, AM Simon, NP Fey, and LJ Hargrove. Intent recognition in a powered lower limb prosthesis using time history information. *Annals of Biomedical Engineering*, 42(3):631–641, Mar 2014. Epub 2013 Sep 20.

[32] Sophie Heins, Louis Flynn, Joost Geeroms, Dirk Lefeber, and Renaud Ronsse. Torque control of an active elastic transfemoral prosthesis via quasi-static modelling. *Robotics and Autonomous Systems*, 107:100–115, 2018.

[33] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, March 1960.

[34] Greg Welch and Gary Bishop. An introduction to the kalman filter. Technical report, Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3175, 2006.

[35] Yitian Chen, Yanfei Kang, Yixiong Chen, and Zizhuo Wang. Probabilistic forecasting with temporal convolutional neural network. *Neurocomputing*, 399:491–501, 2020.

[36] Dean D. Molinaro, Inseung Kang, Jonathan Camargo, Matthew C. Gombolay, and Aaron J. Young. Subject-independent, biological hip moment estimation during multimodal overground ambulation using deep learning. *IEEE Transactions on Medical Robotics and Bionics*, 4(1):219–229, 2022.

[37] Guy Hacohen, Avihu Dekel, and Daphna Weinshall. Active learning on a budget: Opposite strategies suit high and low budgets, 2022.

[38] Patrick K Gikunda and Nicolas Jouandeau. Cost-based budget active learning for deep learning. In *9th European Starting AI Re-*

searchers' Symposium (STAIRS), Saint-Jacques-de-Compostelle, Spain, Aug 2020.
[39] Burr Settles and Mark Craven. An analysis of active learning strategies for sequence labeling tasks. 2008.
[40] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2 edition, 2009.
[41] Sebastian Farquhar, Yarin Gal, and Tom Rainforth. On statistical bias in active learning: How and when to fix it, 2021.
[42] Xinlei Zhou, Han Liu, Farhad Pourpanah, Tieyong Zeng, and Xizhao Wang. A survey on epistemic (model) uncertainty in supervised learning: Recent advances and applications. *Neurocomputing*, 489:449–465, June 2022.
[43] Y. Freund, H. S. Seung, E. Shamir, et al. Selective sampling using the query by committee algorithm. *Machine Learning*, 28:133–168, 1997.
[44] Andrew McCallum and Kamal Nigam. Employing em and pool-based active learning for text classification. In *Proceedings of the Fifteenth International Conference on Machine Learning*, ICML '98, page 350–358, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.
[45] Jerzy Stefanowski and Mateusz Pachocki. Comparing performance of committee based approaches to active learning. In *Recent Advances in Intelligent Information Systems*, pages 457–470, Poznan, Poland, 2009. Poznan University of Technology.
[46] Hideitsu Hino and Shinto Eguchi. Active learning by query by committee with robust divergences, 2022.
[47] Adam Conkey and Tucker Hermans. Active learning of probabilistic movement primitives, 2022.
[48] P. C. Mahalanobis. On the generalized distance in statistics. Tech. rep., National Institute of Science of India, 1936.
[49] Alper Köse, Andrea Cereatti, and Ugo Della Croce. Estimation of traversed distance in level walking using a single inertial measurement unit attached to the waist. *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 1125–1128, 2011.
[50] Q. Li, M. Young, V. Naing, and J.M. Donelan. Walking speed estimation using a shank-mounted inertial measurement unit. *Journal of Biomechanics*, 43(8):1640–1643, 2010.
[51] Keehong Seo. Real-time estimation of walking speed and stride length using an imu embedded in a robotic hip exoskeleton. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12665–12671, 2023.
[52] D. Arumukhom Revi, S. M. M. De Rossi, C. J. Walsh, and L. N. Awad. Estimation of walking speed and its spatiotemporal determinants using a single inertial sensor worn on the thigh: From healthy to hemiparetic walking. *Sensors (Basel)*, 21(21):6976, Oct 2021.
[53] Deepesh Agarwal and Balasubramaniam Natarajan. Tracking and handling behavioral biases in active learning frameworks. *Information Sciences*, 641:119117, 2023.
[54] J. Elmer, B. L. Jones, and D. S. Nagin. Using the beta distribution in group-based trajectory models. *BMC Medical Research Methodology*, 18:152, 2018.