

# Real-Time Balancing of Stability and Plasticity in Continual Learning Enables Adaptive Speed Estimation for Lower-Limb Prostheses

Cole B. Johnson<sup>✉</sup>, Jairo Maldonado-Contreras<sup>✉</sup>, Kinsey R. Herrin<sup>✉</sup>, and Aaron J. Young, *Senior Member, IEEE*

**Abstract**—A primary challenge in continual learning (CL) for wearable robotics, especially prosthetics, is balancing the need to retain learned knowledge (stability) with the necessity to adapt to new information (plasticity). This balance is crucial for online adaptation, enabling systems to transition between tasks without losing prior knowledge. In this paper, we introduce a novel online optimizer-based framework designed to manage the stability-plasticity balance through strategic datapoint replay and learning-rate adjustments of a deep neural network. We applied this framework to speed estimation systems for transfemoral prostheses (TFA users), conducting offline validation tests using data from 10 individuals with TFA, and online tests with three TFA and six able-bodied (AB) participants. Our results demonstrate statistically significant improvements: in offline settings, our method showed a 39.2% increase in stability and a 35.2% boost in plasticity over traditional CL approaches during leave-one-subject-out validation. Similarly, in real-time trials with AB participants, we observed statistically significant gains in handling both previously encountered and new walking speeds. Finally, trials with individuals with TFA showed that the system improved the plasticity of the baseline model by 67.45% and the stability of the traditional CL approach by 31.36%; reducing overall average walking speed estimation error by 19.47%.

**Index Terms**—Continual Learning, Machine Learning, Robotics, Lower-Limb Prosthetics.

## I. INTRODUCTION

CONTINUAL Learning (CL) emulates the adaptive learning behavior observed in animals, enabling systems to incrementally adapt and improve upon sequential tasks, a capability essential for generalized intelligent systems [1], [2]. The effectiveness of CL hinges on a delicate balance between preserving knowledge of past tasks (stability) and efficiently learning new information (plasticity). However, achieving this balance is challenging because of the inverse relationship between stability and plasticity in systems with limited resources. Too much stability can cause an

entrenchment effect, making the system inflexible and less capable of integrating new data [3], [4]. In contrast, excessive plasticity can lead to catastrophic interference, a weak form of catastrophic forgetting, where learning new information undermines the system’s proficiency in other tasks [5], [6]. This dynamic is referred to as the stability-plasticity dilemma [3] and is the focus of this study.

The focus in literature regarding the stability-plasticity dynamic in computer science has thus far been on mitigating its harmful effects on both sides of the spectrum – most notably that of catastrophic forgetting. Four classes of approaches have been developed: architecture-based methods introduce dynamic neural structures which add, prune, or modify network modules, thus allowing adjustment of the structure in response to new tasks [7]–[9]. Replay-based systems store representative historical data in a limited coresets and augment portions of it into real-time via fine-tuning cycles to capitalize on task recency bias and maintain competency on old tasks [10]–[14]. Knowledge distillation approaches involve training a student network using the outputs or logits (outputs before activation functions) of the parent network, thereby teaching competency on old tasks to the student network [9], [15]–[17]. Finally, weight regularization methods such as elastic weight consolidation and memory-aware synapses help maintain stability by adjusting the weights of the network to protect from changing ones essential for historical tasks [8], [10], [18], [19].

Although there are established methods to mitigate both entrenchment and catastrophic forgetting, these limitations persist in all systems. However, current research on CL lacks clarity in understanding and identifying the best method to balance stability and plasticity in models. Tuning this balance is vital as different learning scenarios demand varying degrees of stability and plasticity. For instance, environments with frequent and unpredictable changes require a model with higher plasticity to adapt swiftly, whereas more constant environments with long-range patterns benefit from greater stability to preserve accumulated knowledge. Furthermore, the ability to dynamically respond to data streams to maximize the benefit of model fine-tuning is essential: increasing stability in response to potentially harmful data while amplifying plasticity during rapid environmental changes can protect model performance and substantially reduce model reaction time, respectively [20].

Manuscript submitted . Asterisk indicates corresponding author. This work was supported in part by the Ford Foundation Fellowship, by a U.S. Department of Defense grant through the CDMRP Award Number W81XWH-21-1-0686, and National Institutes of Health Director’s New Innovator Award No. DP2-HD111709.

C. Johnson was with the School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332 USA.

J. Y. Maldonado-Contreras, K. R. Herrin, and A. J. Young were with the Mechanical Engineering Department, Georgia Institute of Technology, Atlanta, GA 30332 USA and with the Institute of Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30332 USA.

In recent years, robotics has emerged as a key application of CL due to the necessity of navigating and interacting with the dynamism inherent to various environments and tasks [21]. Here, we use CL in the domain of intelligent lower-limb prosthetic control. Frequent terrain alterations and walking gait pattern shifts of users create an environment that is varied and evolving, requiring the integration of multi-task systems for functionality. Furthermore, translation of raw sensor data into meaningful user context estimation (walking speed, terrain slope) and intent recognition (ambulatory mode) illustrates the necessary model complexity and constant fine-tuning processes that must occur to maintain an effective control system, especially for handling new, unseen tasks to the model without losing the ability to predict historical tasks. Finally, the unpredictability in the system, particularly due to manual donning/doffing of sensors, gait asymmetries/deviations, challenges with socket fittings, and sensor drops over long periods of use underscores the necessity of a robust CL-based approach to maintain consistent and accurate performance.

Numerical methods for intent recognition [22] and context estimation [23], [24] are commonly used in microprocessor-controlled prosthetic knees, such as the Ottobock C-Leg, to modulate resistance and improve gait efficiency. However, these methods often lack the accuracy needed to handle dynamic transitions, such as changes in walking speed or incline, where precise state estimation is critical. Recent advances in machine learning (ML) [22], [25]–[32] and deep learning (DL) [33], [34] provide adaptive pipelines that improve mode selection, walking speed estimation, and incline adaptation, addressing limitations of traditional numerical approaches. These learning-based methods enable more reliable handling of state transitions, which highlights the need for adaptive systems in prosthetic devices.

First investigated in 2009 by Sensinger, et al. [35], adaptation methods for wearable robotics have been applied to various fields, more recently with Spanias, et al. developing a method for intent recognition which uses PCA dimensionality reduction on a tuned feature space taken from EMG and mechanical sensor readings. This system used a collection of Dynamic Bayesian Networks (DBN), each corresponding to a phase and mode, totaling eight models, to then make modal classifications based on these readings. The system implemented an LDA classifier trained on mechanical sensor data as their backward estimator to fine-tune the DBNs [36]. Similar work from Woodward et al. introduced DL for intent recognition through using a Scaled Conjugate Gradient Artificial Neural Network as a forward predictor, taking as input feature-extracted data from load cell and IMU sensor readings. However, unlike Spanias, et al., the study relied on ground truth mode labels [37]. Other applications also benefit from this approach, such as gait phase estimation in exoskeletons [38].

Here, we build on the methods first described in [39], [40], using a model initially trained offline using participant-

independent (IND) data, which means that models were first trained on multi-participant data and tested using the novel participant data. During the course of a trial, the system fine-tunes using batched gait data specific to the user, thereby transitioning to a partially participant-dependent (DEP) model, meaning the model is now both trained and evaluated on novel participant data. The adaptive pipeline consists of a forward predictor which produces the real-time modal classifications and speed/slope estimates which are converted to assistance scaling parameters for prosthetic actuation [41], [42]. Concurrently, a backward estimator provides more accurate gait labels from completed stride data [41]. After a certain-sized batch of strides has been completed, these retrospective labels are used with the gait data to fine-tune the forward predictor in a pseudo self-supervised CL, ideally driving the accuracy of the real-time forward predictor to the asymptotic limit given by the accuracy of the slower, more accurate backward estimator.

In this paper, we implement novel stability-plasticity manipulation systems into DL-based speed estimation adaptation in lower-limb prosthetic control systems. Our contributions using this system are threefold, each system being demonstrated offline using a dataset of 21 trials from 10 distinct TFA participants and real-time efficacy shown online using 6 able-bodied participants wearing a prosthetic adapter and 3 TFA participants:

- 1) We propose a novel optimization-based method of simultaneously controlling stability and plasticity based on an analysis of their dynamics in the prosthetic adaptation pipeline. Specifically, we hypothesize that our system is capable of fine-tuning the levels of stability and plasticity within neural networks to precise levels.
- 2) Furthermore, we introduce the theory of Recency-Weighted Incremental Error which allows the introduction of a tunable parameter  $\lambda$  representing accuracy-optimal ratios of stability and plasticity along a continuum from very plastic to very stable. This theory allows us to test the hypothesis that controlling stability and plasticity levels in a CL system towards an optimal value is errorwise advantageous. Furthermore, we hypothesize that the system can attain higher plasticity than baseline models and higher stability than conventional CL methods.

## II. METHODS

### A. Prosthetic System

*Device:* In this study, we evaluate our control systems using the Open Source Leg (OSL) designed by the University of Michigan (detailed mechanical design and characterization available in [43], [44]) and assembled by the Exoskeleton and Prosthetic Intelligent Controls (EPIC) Lab at Georgia Institute of Technology. The OSL is a powered transfemoral prosthesis that utilizes DEPHY ActPack actuators at both the knee and ankle joints, a six degree-of-freedom (DOF) load cell located at the shank (Sunrise Instruments M3564F,

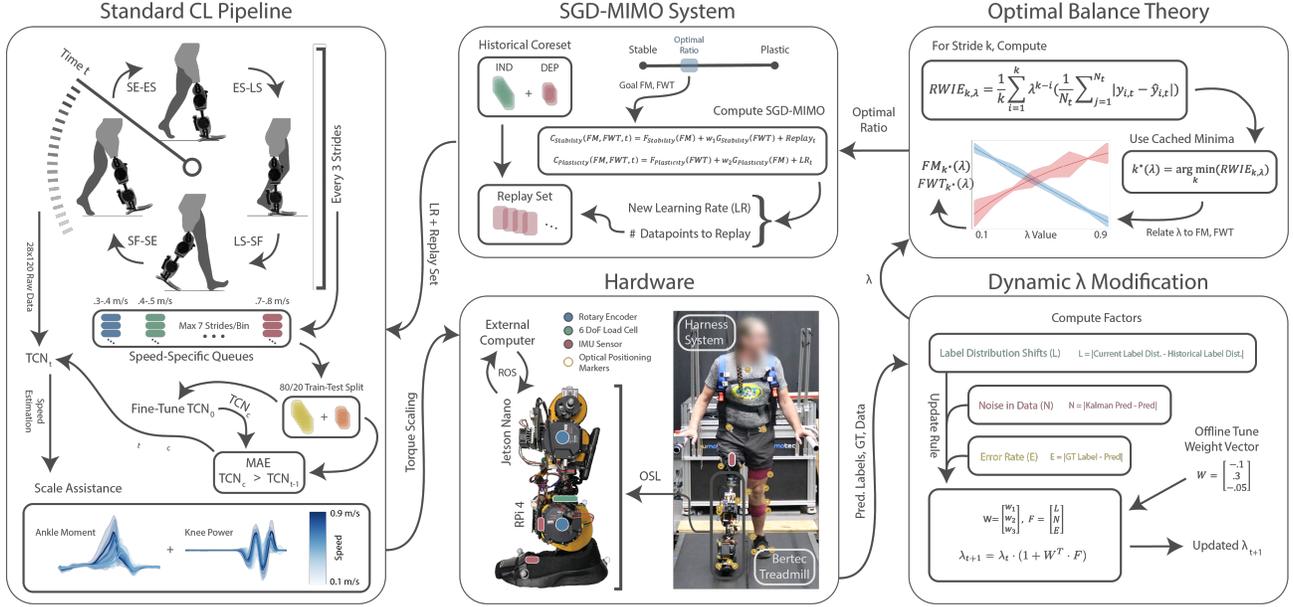


Fig. 1: Overview of hardware systems, the standard CL pipeline, the SGD-MIMO system, Optimal Balance Theory, and Dynamic  $\lambda$  Modification and the data that is transferred between each. The experiments involved a participant walking on a Bertec treadmill throughout a speed profile using the OSL, equipped with ankle and knee encoders, a loadcell, optical positioning markers, and IMU sensors, and communicating with an external computer using ROS. The standard CL pipeline operates using a forward predictor which estimates walking speed every 1200 ms and thus scales knee power and ankle moment used by the prosthetic. and a backwards estimator which outputs ground truth speed values. The backwards estimator output is iteratively binned into .1 m/s queues, and used to fine-tune the forward estimator. Every fine-tuning iteration, the SGD-MIMO system appends the bins with a replay set derived from a historical coreset. This subsystem determines the new learning rate and the number of data points to replay at each iteration using an optimal ratio given by the Optimal Balance system. The Optimal Balance system derives the goal FM and FWT values from the RWIE relation found using cached error minima found offline for each  $\lambda$  value in the range  $[.1, \dots, .9]$ . The Dynamic  $\lambda$  Modification system computes different factors using live data stream characteristics and an offline-tuned weight matrix is used to update  $\lambda$  for each fine-tuning timestep.

Nanning, China), an encoder at each joint (AS5047P & AK7452 – DEPHY ActPack, Maynard, MA), and three inertial measurement units (IMU) located on the shank (MPU-9250 InvenSense, San Jose, CA), the top of the foot, and attached to the participant’s thigh (2x 3DMCX5-25 IMU, LORD Microstrain, Williston, VT). A NVIDIA Jetson Nano was mounted in the knee housing and used to compute real-time context estimations (i.e., forward predictions) using the sensor readings, all of which were sampled at 100 Hz. Additionally, an external laptop assisted in signal visualization and impedance parameter adjustments, communicating with the Jetson Nano through Robotic Operating System (ROS) channel messages. Data communication between concurrently running nodes, such as the finite state machine (FSM), forward predictor, and backward estimator, were also managed by ROS, which broadcasts channels like sensor data, current phase, estimated speed, and ground truth labels. These are used by the adaptive pipeline to learn participant-specific gait patterns and progressively improve the models used for real-time context estimation.

**Control System:** The control architecture of the prosthesis encompasses three distinct levels: high-level, mid-level, and low-level control. The high-level controller makes context estimations for walking speed. These estimations are facilitated by onboard sensor data and ML models belonging to

the forward predictor, which we describe below. Estimates are filtered using a Kalman filter [45]. This technique functions by recursively estimating joint probability distributions over undetermined variables for every time frame. Namely, the Kalman Gain for the  $n$ th time frame:

$$K_n = \frac{P_{prior}}{P_{prior} + P_{meas}} \quad (1)$$

where  $P_{prior}$  is the prior variance and  $P_{meas}$  is the measurement variance. The gain is used to weight incoming data, estimating the  $n$ th time frame as  $X_{nn} = X_{prior} + K_n(M_{slope} - X_{prior})$ , where  $X_{prior}$  is the prior estimate and  $M_{slope}$  is the slope measurement. By using this filtered data, we decrease the noise from the speed estimates and thus facilitate a more predictable system. We also compute the estimate’s variance for the  $n$ th time frame as  $P_{nn} = P_{prior} \cdot (1 - K_n)$ . The mid-level controller is characterized by a FSM with four states corresponding to different gait phases: Early Stance (ES, stiffening of the knee and ankle joints), Late Stance (LS, powered plantar flexion during push-off), Swing Flexion (SF, powered knee swing flexion), and Swing Extension (SE, powered knee swing extension assistance). Each state manages transitions between gait phases, using impedance parameters of the knee and ankle joints, and generates a desired joint profile throughout the gait cycle, which have been described in detail in many prior works [46], [47]. The

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

Architecture	Value	Training	Value
kernel size	5	input size	28
dropout	0.2	output size	1
effective hist	120	LR	$1e^{-4}$
# channels	[10, 10, 10, 10]	epochs	2

controller at this level regulates torque ( $\tau$ ) at the joint level ( $i$ ) using a torque law that is a function of the current joint state (both angle  $\theta$  and angular velocity  $\dot{\theta}$ ), stiffness ( $k$ ), theta equilibrium ( $\theta_{eq}$ ), and damping parameters ( $b$ ), thereby defining the impedance law

$$\tau_i = -k_i(\theta_i - \theta_{eq,i}) - b_i\dot{\theta}_i. \quad (2)$$

Here, the peak ankle plantarflexion moment is found in Level Walking (LW) by fitting a linear regression model to a dataset collected from able-bodied participants. The moment was manually adjusted and it was found that it increased by a factor of 0.4898 Nm at a speed of 0.5 m/s [41]. Thus, we can accordingly find the baseline prosthetic ankle torque during plantarflexion and adjust the impedance control stiffness  $k_s$  according to current walking speed  $v$  (m/s) [48] through

$$k_s = k(1 + 0.422(v - 0.5)). \quad (3)$$

*Adaptation Pipeline:* We integrate the forward predictor and backward estimator into the adaptive pipeline to enable real-time speed estimations to be sent to the mid-level controller and thus scale torque parameters, while allowing for batched fine-tuning using retrospective labeling. The forward predictor in this embodiment consists of a Temporal Convolutional Neural Networks (TCN) [49] which outputs speed estimations based on 1200ms segments of data with 20ms overlap between successive windows which modulates torque at 50Hz, and uses an architecture described in Table I [50].

The backward estimator in previous work has been shown to be highly effective when utilizing numerical IMU distance calculations [39], [51]–[54]. However, in this paper, we will be using ground-truth labels to evaluate the efficacy of the stability-plasticity manipulation systems in isolation, so as to avoid an interaction effect with the accuracy of the backward estimator. Every batch of three strides, the backward estimator processes elapsed data and outputs speed labels for the associated data, both of which are used to fine-tune the forward predictor. Here, to simplify the analysis, we decided to use speeds collected in real-time from treadmill readings as our backward estimator labels. In this paper, we refer to the adaptation process without any stability-plasticity control as the Conventional CL method, presented in Algorithm 1. In this method, the labeled data were organized into bins by speed increments of 0.1 m/s from .3 m/s to .8 m/s. Each bin holds up to seven strides, using a first-in-first-out method to keep only the most recent strides. Adaptation occurred every third stride, using 80% of the strides in each bin for training and the remaining 20% for validation. At least two strides

per bin were required before starting the adaptation to ensure a balanced representation of walking speeds. The adaptation used a learning rate of 0.0001, a batch size of 32, Adam optimizer, mean squared error loss function, and ran for 2 epochs, the hyperparameters of which were tuned through a grid search. Successful adaptations, indicated by improved mean absolute error (MAE) on validation steps, led to the replacement of the current forward estimator with the newly adapted TCNs.

### B. Stability-Plasticity Control Theory Background

The system we propose establishes a feedback loop that continuously adjusts control strategies based on real-time metrics that measure stability and plasticity. This adjustment is informed by an understanding of how these metrics interact, allowing the system to dynamically optimize learning processes in response to evolving conditions.

*Metrics:* System performance was evaluated using Average Incremental Error (AIE) [12], [13], based on the Average Error (AE) [55]. AE assesses the model’s immediate performance on the current task,  $k$ ,

$$AE_k = \frac{1}{k} \sum_{j=1}^k e_{k,j}, \quad (4)$$

where  $e_{k,j}$  is the MAE for task  $j$  evaluated after training up to the  $k$ -th task:

$$e_{k,j} = MAE_{k,j} = \frac{1}{N} \sum_{i=1}^N |y_{i,j} - \hat{y}_{i,j,k}|, \quad (5)$$

where  $N$  is the number of data points,  $y_{i,j}$  is the true value, and  $\hat{y}_{i,j,k}$  is the predicted value for the  $i$ -th data point in the  $j$ -th task. AIE measures the model’s ability to maintain performance over time, defined for the  $k$ -th task as

$$AIE_k = \frac{1}{k} \sum_{i=1}^k AE_i. \quad (6)$$

This cumulative metric indicates whether the model’s performance changes with new tasks.

The Forgetting Measure (FM) assesses stability, specifically evaluating the model’s ability to avoid catastrophic interference (forgetting) in a regression context [18]. Forgetting manifests itself as an increase in error for previously learned tasks, similar to how catastrophic forgetting in classification leads to error rates near random chance. The increase in error on task  $j$  after training up to task  $k$ , compared to its prior best performance, is

$$f_{j,k} = e_{k,j} - \min_{i \in \{1, \dots, k-1\}} e_{i,j}, \quad \forall j < k \quad (7)$$

and the FM for the  $k$ -th task aggregates this increase across all previously learned tasks:

$$FM_k = \frac{1}{k-1} \sum_{j=1}^{k-1} f_{j,k}. \quad (8)$$

The Forward Transfer Error (FWT) measures plasticity, assessing how previous learning episodes affect the performance of a model on a new task [55]. A negative FWT indicates that prior learning has benefited the model in new tasks, while a positive FWT suggests interference from earlier tasks, highlighting adaptability challenges. This measure is inversely correlated with plasticity: it is positive when interference occurs and negative when past learning provides valuable insights for current data:

$$FWT_k = \frac{1}{k-1} \sum_{j=2}^{k-1} (e_{j,j} - \tilde{e}_j), \quad (9)$$

where  $\tilde{e}_j$  denotes the regression error when the model is exclusively trained on the  $j$ -th task, while  $e_{j,j}$  signifies the error on the  $j$ -th task after the model has incrementally been trained up to that point.

*Basic Tools:* To enhance plasticity, adjusting the learning rate during fine-tuning in the CL cycle is effective [56], [57]. For stability, we use a simple replay method, specifically the reservoir technique proposed by Vitter in 1985 [58]. This technique selects a uniform subset  $S = |C|$  from the incoming data stream to include in the coreset  $C$ , constructed with partial replay (i.e., storing only part of the experienced data) to reduce storage costs [11], [55]. This strategy ensures that each sample has an equal chance of inclusion, while maintaining future training representativeness. Leveraging a temporally aware model, we store data as complete strides (exemplars) rather than individual points. Data selection for inclusion in the fine-tuning pool follows a uniform distribution. Furthermore, we employ experience replay (ER), a method that combines historical and current data during training [14], [59]. While training on a current batch  $B$ , the model’s objective includes both new and historical data, represented as

$$L(B; \theta) + \beta L(C; \theta), \quad (10)$$

where  $L$  denotes the cross-entropy loss,  $C$  is the coreset,  $\theta$  represents model parameters, and  $\beta$  balances the focus between new and historical tasks.

*Dynamics:* To accurately adjust stability and plasticity in learning models, we must understand each method’s impact on these metrics. For stability, represented by the FM metric, we analyze how the number of replayed datapoints influences this value. For plasticity, indicated by the FWT metric, we examine the effects of varying learning rates.

We conducted experiments on the offline dataset to establish these relationships before evaluation. In the Replay-FM trial, we varied the number of coreset exemplars added to the training set in each fine-tuning iteration, averaging FM changes across iterations and participants. The resulting formula for Replay-FM dynamics is  $F_{\text{Replay}}(x)$ . In the LR-FWT trial, we adjusted the learning rates between  $10^{-5}$  and  $10^{-2}$ , deriving the formula for  $F_{\text{LR}}(x)$ . Furthermore, given the inverse correlation between stability and plasticity, we explored the interactions between methods and metrics. This

yielded formulas for  $G_{\text{Replay}}(x)$ , representing the impact of the replay method on FWT, and  $G_{\text{LR}}(x)$ , for the effect of the changes in learning rate on FM.

### C. Stochastic Gradient Descent - Multi-Input Multi-Output

To derive the system for controlling stability and plasticity throughout a CL trial, we first consider controlling stability independently of its interaction with plasticity. This approach parallels optimizers in machine learning (ML) training: in iteration  $t$ , we aim to minimize a loss function that represents the difference between the current and target stability metric values. Accordingly, we can define an update rule for the number of datapoints to replay at iteration  $t + 1$ , analogous to Stochastic Gradient Descent (SGD):

$$\overbrace{\text{Replay}_{t+1}}^{\theta_{t+1}} = \overbrace{\text{Replay}_t}^{\theta_t} + \overbrace{\eta}_{\eta_{\text{Replay}}} \cdot \overbrace{L_{\text{Stability}} \cdot \frac{dF_{\text{Replay}}^{-1}}{d(\text{FM})}}^{\nabla L(\theta_t)} \quad (11)$$

$$\text{LR}_{t+1} = \text{LR}_t + \eta_{\text{LR}} \cdot L_{\text{Plasticity}} \cdot \frac{dF_{\text{LR}}^{-1}}{d(\text{FWT})} \quad (12)$$

In this method,  $\text{Replay}_t$  represents the number of replayed datapoints,  $\text{LR}_t$  is the learning rate at iteration  $t$ ,  $\eta_{\text{Replay}}$  is a tunable parameter determining the update step size, and  $L_{\text{Stability}}$  and  $L_{\text{Plasticity}}$  are loss functions. Although this approach can generalize to any optimizer, we use Stochastic Gradient Descent (SGD) due to several factors: the simplicity of the loss landscape suggested by functions  $F$  and  $G$ , guidance from their derivatives in shaping loss function behavior, and the high similarity between participants’ gait data, which implies robustness for future participants. We also use the mean absolute error (MAE) as a loss function, as it outperforms the mean squared error (MSE) in this context.

Given the strong interaction between plasticity and stability, each update rule must consider its impact on the other metric. To manage this, we implement a multiple-input multiple-output (MIMO) system with cross-interaction terms. These terms model the adverse effects of actions that target one metric on the opposing metric, represented by the previously derived functions  $G$ . Thus, we establish update rules for inputs  $FM$ ,  $FWT$ , and  $t$ :

$$C_{\text{Stability}} = F_{\text{Stability}}(FM) + w_1 G_{\text{Stability}}(FWT) + \text{Replay}_t \quad (13)$$

$$C_{\text{Plasticity}} = F_{\text{Plasticity}}(FWT) + w_2 G_{\text{Plasticity}}(FWT) + \text{LR}_t \quad (14)$$

Here, the term  $F$  influences a specific metric, while the term  $G$  dampens changes that could negatively impact the opposing metric. The effect of cross-interaction is adjusted

by the weight terms  $w_1$  and  $w_2$ . The identities of each  $F$  and  $G$  term are as follows:

$$F_{\text{Stability}}(FM) = \eta_{F_{\text{Replay}}} \cdot L_{FS} \cdot \left( \frac{dF_{\text{Replay}}^{-1}}{d(FM)} \right) \quad (15)$$

$$G_{\text{Stability}}(FWT) = \eta_{G_{\text{Replay}}} \cdot L_{GS} \cdot \left( \frac{dG_{\text{Replay}}^{-1}}{d(FWT)} \right) \quad (16)$$

$$F_{\text{Plasticity}}(FWT) = \eta_{F_{\text{LR}}} \cdot L_{FP} \cdot \left( \frac{dF_{\text{LR}}^{-1}}{d(FWT)} \right) \quad (17)$$

$$G_{\text{Plasticity}}(FM) = \eta_{G_{\text{LR}}} \cdot L_{GP} \cdot \left( \frac{dG_{\text{LR}}^{-1}}{d(FM)} \right) \quad (18)$$

#### D. Optimal Balance Theory

SGD-MIMO enables us to adjust stability and plasticity levels to target metric values, but finding the optimal values to aim for requires a new framework. Therefore, we develop a theory to balance metric values (FM and FWT) for various stability-plasticity levels, introducing the concept of Recency-Weighted Incremental Error (RWIE), derived from AIE.

$$RWIE_{k,\lambda} = \frac{1}{k} \sum_{i=1}^k \lambda^{k-i} \left( \frac{1}{N_t} \sum_{j=1}^{N_t} |y_{i,t} - \hat{y}_{i,t}| \right) \quad (19)$$

Here,  $\lambda$  serves as a tunable time-decay factor: higher  $\lambda$  weights historical data more heavily, while lower  $\lambda$  emphasizes recent data. This enables performance evaluation for specific stability and plasticity goals within the CL lifecycle (i.e., high  $\lambda$  to assess stability and low  $\lambda$  to assess plasticity). By sweeping  $RWIE_{k,\lambda}$  for all  $\lambda \in (0, 1)$ , we identified the optimal FM and FWT values throughout the continuum from highly plastic to highly stable.

$$k^*(\lambda) = \arg \min_k (RWIE_{k,\lambda}) \quad (20)$$

Then we computed the FM and FWT values at each of these minima (i.e., for each  $\lambda$  value swept):  $FM_{k^*(\lambda)}$  and  $FWT_{k^*(\lambda)}$ . Thus, we establish a precise relation between  $\lambda$  and the RWIE-optimal FM and FWT values, providing  $\lambda \in (0, 1)$  as a tunable parameter to set a target on the stability-plasticity continuum for SGD-MIMO throughout a CL model's lifecycle.

#### E. Offline Analysis and Evaluation

Offline, we first validated the functionality of SGD-MIMO. We then evaluated pipeline performance against Conventional CL and Baseline methods in two domains: stability (generalization from historical data after non-exposure) and plasticity (effective few-shot learning capability).

*Dataset Overview:* For offline analysis and system optimization, we used two datasets (mean  $\pm$  SD), both from individuals with TFA. The first dataset includes 11 participants (9 male, 2 female; age =  $50.36 \pm 12.09$  years; mass =  $80.13 \pm 15.64$  kg; height =  $1.77 \pm 0.10$  m) which we fully utilize. The second dataset comprises data from

10 participants (7 male, 3 female; age =  $42.40 \pm 12.70$  years; mass =  $71.44 \pm 14.46$  kg; height =  $1.69 \pm 0.10$  m), originally using the first two trials for benchmarking purposes and thus used here. The participants walked at discrete treadmill speeds [0.3, 0.5, 0.7, 0.9, 0.8, 0.6, 0.4] m/s, with each speed held for 20 seconds and  $0.1 \text{ m/s}^2$  acceleration between speeds. The studies received approval from the Georgia Institute of Technology Institutional Review Board, with participants providing their informed written consent and prosthetic alignment by a certified prosthetist.

*SGD-MIMO:* To validate the efficacy of the SGD-MIMO system, we examine its ability to drive FM and FWT metrics toward specified goal values. We evaluated this system through two analyses, each using a forward predictor trained on data from all participants except one, and evaluated through a standard adaptation trial (augmented with the SGD-MIMO system) using the data of the omitted participant. A standard adaptation trial involves testing data collected as the participant walks with the prosthetic, performing forward predictions in real-time, and adapting with ground-truth backward estimates every three strides. In the first analysis, we perform a grid search over goal values: Goal FM  $\in \{.001, .002, \dots, .009\}$  and Goal FWT  $\in \{-.001, -.002, \dots, -.009\}$ . The second analysis focuses on evaluating the efficacy of the SGD-MIMO system in driving stability-plasticity metrics to RWIE-optimal values specified by  $\lambda$ . For each  $\lambda \in \{0.1, \dots, 0.9\}$ , we compute optimal  $FM_{k^*(\lambda)}$  and  $FWT_{k^*(\lambda)}$  and conduct a similar test using these as goal values. These analyses were repeated for each participant in the offline TF dataset.

*Stability & Plasticity:* We evaluate the model's stability (ability to retain previously learned information) and plasticity (capacity to adapt to new labels) using a  $K$ -fold validation-like procedure across varying speeds and  $\lambda$  values. In both cases, we omit data for a specific participant during training and use these data for testing, creating a systematic participant exclusion for each iteration. For stability, forward predictors are trained on all data except for a specific participant and a particular speed, with adaptation cycles using only the data of the omitted participant. Performance is evaluated on the withheld speed, and this process is repeated for each speed in  $\{0.3, 0.4, \dots, 0.9\}$  m/s, each  $\lambda \in \{0.1, 0.2, \dots, 0.9\}$ , and each participant, assessing generalization from independent (IND) to dependent (DEP) settings.

To demonstrate the plasticity of the pipeline, we use a similar training and testing procedure as that used for stability. However, during training, we systematically withhold data for a specific speed. Adaptation then uses the data of the omitted participant, and performance at the withheld speeds is assessed. This is repeated similarly to the stability test and thus directly assesses the adaptability of the pipeline to new information.

## Speed Stability & Plasticity Speed Profile

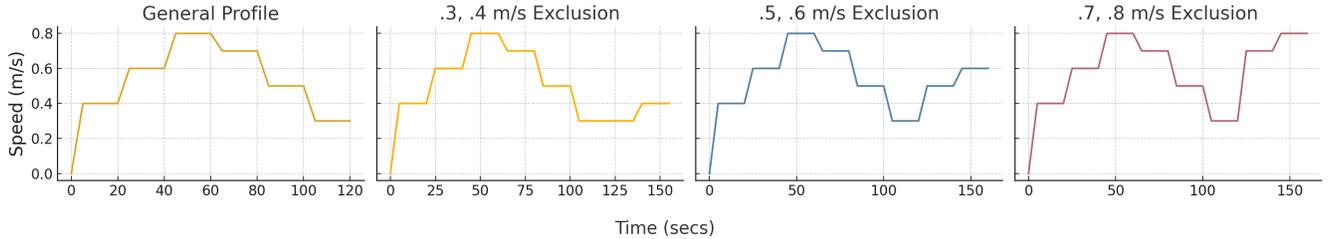


Fig. 2: Treadmill Speed Profile. General Profile terminates at 120 seconds after a staircase pattern. The exclusion trials consist of the General Profile appended with 15 seconds at each excluded speed.

### F. Online Validation

**Dataset Overview:** Six able-bodied individuals participated in the study (6 male; age =  $24.33 \pm 4.37$  years; mass =  $89.17 \pm 7.99$  kg; height =  $1.83 \pm 0.05$  m), as well as three participants with TFA (3 male; age =  $38.66 \pm 8.33$  years; mass =  $86.41 \pm 10.47$  kg; height =  $1.68 \pm 0.09$  m).

**Experimental Protocol:** Participants completed nine walking trials on a Bertec split-belt treadmill with force plates, with treadmill speeds set as in Fig. 2, varying by trial. Throughout,  $\lambda$  was set to 0.6. A certified prosthetist initially fitted the OSL to each participant with TFA. Before trials, impedance control parameters were tuned during treadmill walking using user preference. In addition, the scaling equations for ankle pushoff and knee extension stiffness were adjusted for walking speed. The trials were structured as follows:

**Trial 1:** The first trial involved walking at the speeds specified in the General Profile. During ambulation, forward estimates used ground-truth treadmill speed, ensuring precise prosthetic assistance scaling. The data from this trial served as a test set for real-time monitoring of the performance of the adaptation versus the baseline model.

**Trials 2-4:** These trials assessed the stability of the pipeline during adaptation, each consisting of two parts. In part (a), an IND forward predictor trained on the offline dataset was adapted in real-time using the S&P Controlled framework while walking the General Profile. During adaptation, treadmill speeds were maintained at specific values (trial 2: 0.3 & 0.4 m/s, trial 3: 0.5 & 0.6 m/s, trial 4: 0.7 & 0.8 m/s) without adaptation. In part (b), these speeds were walked at for 15 seconds each with adaptation frozen to record the accuracy of the speed estimation. Collected data was then used for offline simulation of Baseline and Conventional CL methods for comparison.

**Trials 5-7:** These trials were aimed to evaluate the plasticity of the pipeline, which also includes two parts. In part (a), adaptation occurred throughout the General Profile using a forward predictor trained on all offline TF data, excluding specific speeds, with the S&P Controlled pipeline. Afterward, adaptation was frozen, and participants walked for 15 seconds at the omitted speeds, recording the accuracy of the speed estimation, as in trials 2-4. The data gathered were then

used for offline simulation of Baseline and Conventional CL methods for comparison.

**Trials 8-9:** The last two trials involved participants walking the General Profile to enable error-based comparison between methods. A forward predictor trained on all the TF participant speeds in the offline dataset was used with the S&P Controlled pipeline. In trial 8, adaptation proceeded normally, while in trial 9, adaptation was frozen and the accuracy recorded. Offline simulations of Baseline and Conventional CL methods were also performed, allowing error comparison between methods after forward predictor adaptation.

**Statistical Analysis:** For Stability and Plasticity trials with able-bodied participants, we used Repeated Measures ANOVA to assess differences across the three methods: Conventional CL, S&P Controlled, and Baseline. This analysis tested two hypotheses: that the S&P Controlled method would exhibit greater stability than Conventional CL and higher plasticity than Baseline. Post hoc tests with Bonferroni corrections followed ANOVA to identify specific differences between methods. For trials with TFA participants, we calculated the effect sizes using Hedges'  $g$  to compare outcomes between methods.

## III. RESULTS

### A. Offline Results

**SGD-MIMO:** As shown in Fig. 3, averaged across participants, it is observed that the MAE of the Forgetting Measure (FM) is equal to  $3.61 \times 10^{-5}$  and the MAE of the Forwards Transfer Error (FWT) is equal to  $6.33 \times 10^{-5}$  m/s. These error rates are negligible, providing strong support that SGD-MIMO performs as expected.

**Stability:** We observed an approximately linear relationship between the value of  $\lambda$  and the MAE in Fig. 4, where the higher values of  $\lambda$  correspond to a lower error, enforcing the positive relationship between stability and the value of  $\lambda$ . We also noted that Conventional CL performed similarly to the S&P Controlled method when  $\lambda = 0.23$  and that Baseline performed similarly to the S&P Controlled method when  $\lambda = 0.55$ . Therefore, the S&P Controlled method demonstrated superior stability when  $\lambda \geq 0.55$ , further supporting the use of  $\lambda = 0.6$  in static  $\lambda$  use cases. Moreover, it was clear that extremely high  $\lambda$  values approach accuracy levels expected of

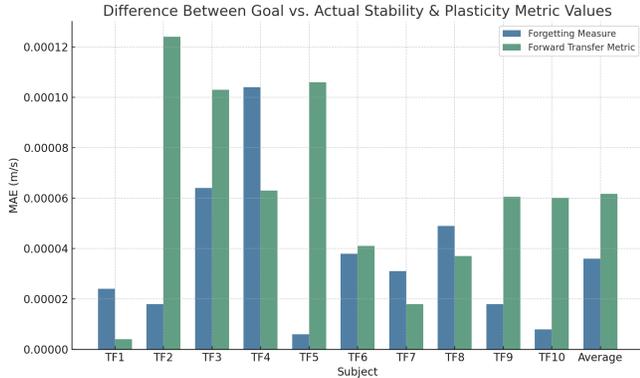


Fig. 3: Results from offline test of SGD-MIMO efficacy. The evaluation shows negligible discrepancy between actual and desired metric values for both Forgetting Measure and Forward Transfer Metric.

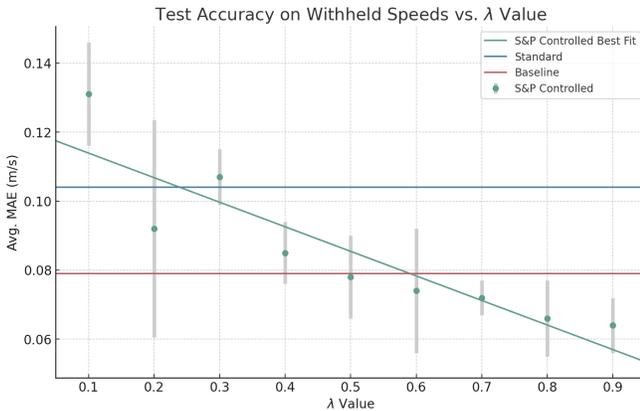


Fig. 4: Results from offline stability trials. S&P Controlled method shows direct correlation between the value of  $\lambda$  and the performance of the task, outperforming Baseline when  $\lambda \geq .55$  and Conventional CL (Standard) when  $\lambda > .23$ .

speeds recently seen in the Conventional CL method, which should be viewed as the asymptotic limit here.

**Plasticity:** In Fig. 5, we observed a very high error rate for the Baseline model, which should be expected as it has no ability to learn the speeds tested on since they are omitted from the initial training set and it does not adapt. Thus, the high MAE of 0.219 m/s was due to the model predicting some of the closest speeds that it had previously seen. As seen in the Stability test, the S&P Controlled method displayed a near-linear relationship between  $\lambda$  values and MAE, reflecting the expectation that low  $\lambda$  values are more plastic. The S&P Controlled method seems to perform better (i.e., is more plastic) than Conventional CL at  $\lambda$  values less than 0.7. This, again, supports the use of the static  $\lambda = 0.6$  value. Similarly to the stability test, low  $\lambda$  values also approached error values expected from models fully trained on withheld speeds.

### B. Online Results

**Stability Evaluation (trials 2-4):** In the able-bodied trials, the statistical analysis demonstrated that the S&P Controlled

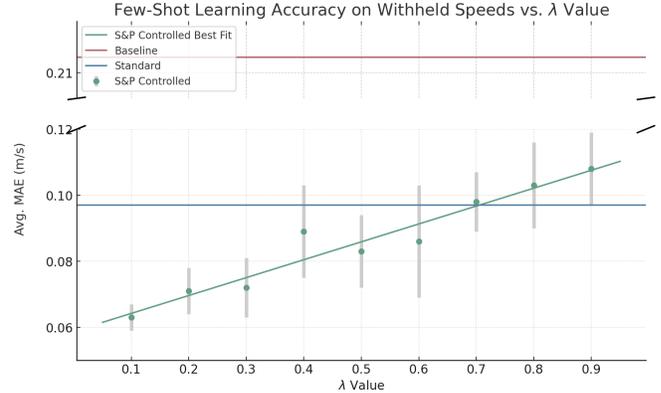


Fig. 5: Results from offline plasticity trials. S&P Controlled method shows inverse correlation between  $\lambda$  value and performance on task, outperforming Conventional CL when  $\lambda < .7$  and Baseline universally.

method is more stable than Conventional CL with statistically significant results ( $p_{val\_corr} < 0.05$ ) for all exclusion profiles, as shown in Fig. 6. On average, the S&P Controlled method improved the Baseline method by 10.25% and Conventional CL by 31.36% (avg.  $p_{val\_corr} = 0.01$ ).

In trials involving participants with transfemoral amputations (TFA), we observed that S&P Controlled still exhibited the lowest error (shown in Fig. 6). On average, the S&P Controlled method improved Baseline by 21.34% (1.91 avg. Hedges'  $g$  over exclusion profiles) and Conventional CL by 37.65% (3.5 avg. Hedges'  $g$ ), improving consistently upon the two alternative methods in every exclusion profile.

**Plasticity Evaluation (trials 5-7):** In the able-bodied trials, we similarly noted that S&P Controlled outperformed the other methods, as seen in Fig. 6. The statistical analysis confirmed the hypothesis that the S&P Controlled method is more plastic than Baseline. On average, the S&P Controlled method improved Baseline by 64.4% (avg.  $p_{val\_corr} = 0.0001$ ) and Conventional CL by 21.06%.

In trials with TFA participants, we observed that S&P Controlled outperforms the other methods, as seen in Fig. 6. On average, the S&P Controlled method improved Baseline by 67.45% (1.5 avg. Hedges'  $g$ ) and Conventional CL by 15.03% (1.23 avg. Hedges'  $g$ ).

**Error Comparison (trials 8-9):** In the able-bodied trials, we compared the Conventional CL and S&P Controlled methods against Baseline, focusing on their final and average Mean Absolute Error (MAE) improvements throughout trials. As shown in Fig. 6, the S&P Controlled method reduced final and average errors by 9.8% and 4.3% over Baseline, respectively, marking a 7.8% improvement in final error and a 38.1% greater improvement in average error than Conventional CL. Similarly, in trials involving participants with TFA, the S&P Controlled method yielded final and average error reductions of 19.47% and 9.58% over Baseline. Thus, compared to Conventional CL, S&P Controlled demonstrated a 10.73% greater reduction in final error 144.51% in average error.

Results from Able-Bodied and Amputee Stability and Plasticity Trials

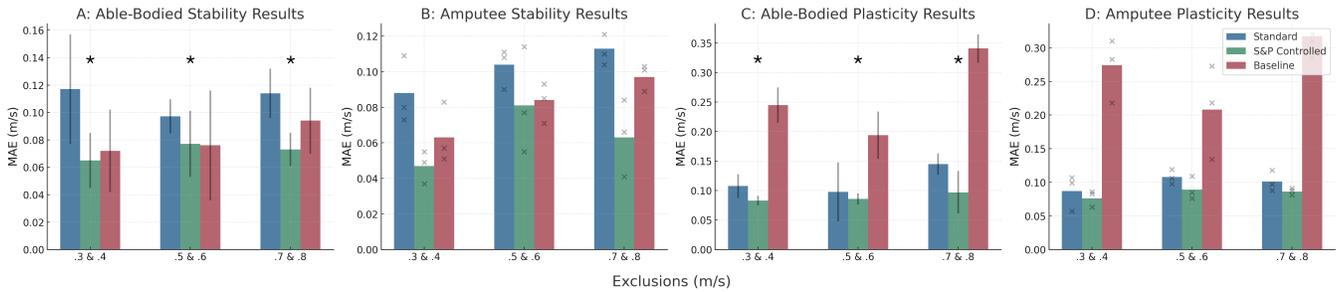


Fig. 6: Results from stability and plasticity trials with able-bodied ( $n = 6$ ) and TFA ( $n = 3$ ) participants. Plot A: Able-Bodied Stability Trials; Plot B: TFA Stability Trials; Plot C: Able-Bodied Plasticity Trials; Plot D: TFA Plasticity Trials. \* denotes statistically significant differences supporting hypotheses

Real-Time Improvement Over Baseline

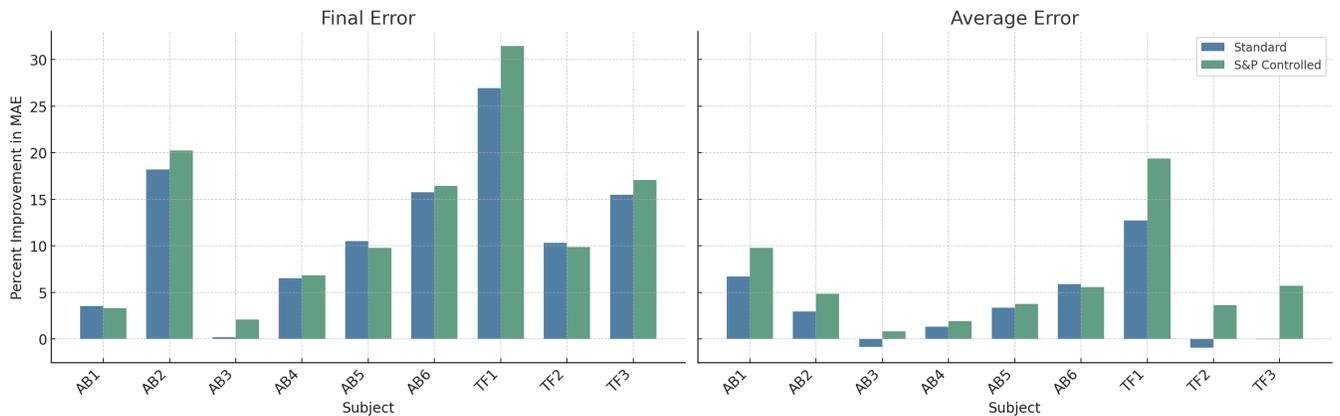


Fig. 7: Percent Improvement of S&P Controlled and Conventional CL methods in able-bodied and TFA participant trials over baseline. Averaged over all able-bodied participants, S&P Controlled has marginal improvement in final error (7.8%) and significant improvement in average error (38.1%) over Conventional CL. Averaged over all TFA participant trials, S&P Controlled has a strong improvement in both final error (10.73%) and average error (144.51%) over Conventional CL.

#### IV. DISCUSSION

The first hypothesis of this study was that the SGD-MIMO system can be used to control the levels of stability and plasticity precisely and simultaneously. Our results support this hypothesis by demonstrating that the SGD-MIMO system effectively drives stability and plasticity metrics toward the desired levels during continual learning (CL) trials. Offline tests indicated that the mean MAE between the actual and desired metrics for Forgetting Measure (FM) was  $3.61 \cdot 10^{-5}$  and for Forwards Transfer Error (FWT) was  $6.33 \cdot 10^{-5}$  after sweeping different target values. This confirms that the system consistently controlled these values with negligible error. Furthermore, when evaluated in different trials, the oscillations around the goal values for both FM and FWT were minimal, indicating that the SGD-MIMO system maintained the pre-specified targeted balance between stability and plasticity throughout training. This precise control reinforces our claim that the system can manipulate these two metrics to pre-specified levels, making it suitable for diverse learning environments that require varying levels of retention and adaptability.

Next, we hypothesized that we can control the stability and plasticity levels of a network in real-time to error-optimal levels, thereby yielding higher plasticity than baseline models and higher stability than conventional CL methods. The results of both offline and online trials provide strong evidence for the real-time efficacy of the SGD-MIMO system in achieving error-optimal control over stability and plasticity. In the online evaluation tests, the S&P Controlled method outperformed the Conventional CL models by improving stability by 31.36% and 37.65% for the able-bodied and TFA trials, respectively. Similarly, the S&P Controlled method improved upon the static Baseline models by increasing plasticity levels by 64.4% and 67.45% for able-bodied and TFA participant trials, respectively. Furthermore, statistical analysis shows significant differences between the S&P Controlled and Conventional CL method in stability trials and between the S&P Controlled and Baseline method in the plasticity trials, with p-values less than 0.05 in all cases. These results confirm that real-time control of stability and plasticity, achieved through responsive fine-tuning, leads to performance improvements that exceed conventional CL systems, even with a static balance (e.g.  $\lambda = 0.6$ ). We note

that in the evaluation of all hypotheses, we observe exactly the same trends between the trials involving AB participants and those involving individuals with TFA.

In our able-bodied experiments on Stability, the S&P Controlled model demonstrated enhanced generalization capabilities compared to the Baseline model in two of the three exclusion profiles (one of which has a statistically significant difference). This superior performance is attributed to its strategy of replaying data points along with real-time learning. By constantly reintegrating historical data, the S&P model retains previously acquired knowledge also belonging to the Baseline model but also integrates new related information. This process prevents overwriting of existing memories and supports the formation of abstract representations that apply across similar speeds. Similarly, in our able-bodied Plasticity tests, the S&P Controlled method was more plastic on average than Conventional CL over all three trials (two of which were statistically significant). This is simply explained by looking at the average learning rate ( $\eta$ ) of both methods throughout the trials. The Conventional CL method used a constant  $\eta = 1 \times 10^{-4}$  and the S&P Controlled method used  $\eta = 3.7 \times 10^{-3}$  on average. Thus, the higher  $\eta$  used by the S&P Controlled method allows it to exhibit greater plasticity.

Direct comparisons to explicit stability and plasticity measurements cannot be made with other prosthetic or wearable robotic studies as other studies in the field do not yet formally test this important concept. However, there are error-wise comparisons with other studies of transfemoral prosthetics for walking speed estimation. For example, IND models have been trained to achieve a root mean square error (RMSE) of 0.07 m/s in speed estimation tasks, as reported by Bhakta, et al [60]. Kinematic methods have achieved comparable RMSE values, with 0.09 m/s using a shank IMU [61] and 0.08 m/s through double integration techniques [62]. A more recent study by Maldonado-Contreras, et al. [63], which focused on the CL pipeline and is comparable with this study, achieved an average MAE of 0.074 m/s after the adaptation was complete. As shown in the results, our approach improves the final error of this method by 10.73% and the average error by 144.51% in the TFA trials. Compared to related studies that achieved improvements of 7% [64] and 45% [37] in classifying walking modes through adaptation techniques, our work demonstrates a 19.47% improvement—an especially notable result given that our task focuses on the regression challenge of speed estimation rather than classification. Our study not only achieves superior accuracy, but also addresses the critical challenge of system robustness by demonstrating adaptability to previously unseen speeds while maintaining stable accuracy in real-time. This highlights a significant advancement over existing systems, which often rely on frequent online fine-tuning.

The study has several limitations, including the selection of FM and FWT for stability and plasticity metrics without evaluating other options like Backward Transfer

Error (BWT) and Intransigence Measure (IM), which could offer different insights. Alternative sampling methods, such as prioritized replay, cluster-based sampling, or temporal coherence sampling, could also enhance model performance but were not explored. Moreover, methods such as adjusting momentum, dropout rates, or regularization strength could serve as alternatives to replay for stability, though these were not considered in depth. In the results, performance differences were more pronounced when extrapolating speeds (e.g., 0.3–0.4 m/s, 0.7–0.8 m/s) rather than interpolating (0.5–0.6 m/s), likely due to the greater distance of these edge speeds from the core training data, which limits the model’s ability to generalize effectively across the full speed range. We also simplified our pipeline by using ground truth (GT) labels from treadmill speeds, limiting the approach to overall applicability, and future work should explore non-GT estimators, perhaps by using IMUs on foot insoles. The statistical analysis of the study relied on able-bodied participants with limited data from participants with TFA, indicating a need for more comprehensive trials with a larger and more diverse participant pool. Overground trials, once an effective backward estimation method is established, are necessary to assess real-world pipeline performance.

## V. CONCLUSION

This paper has demonstrated the potential of the SGD-MIMO framework and Optimal Balance Theory to significantly mitigate the stability-plasticity dilemma in continual learning systems. Our empirical studies in the domain of speed estimation for lower-limb prostheses underscore the ability of these methodologies to effectively enhance both stability and plasticity. The offline and online experiments conducted using data from individuals with TFA and able-bodied individuals reveal that our framework not only outperforms traditional continual learning systems in this domain, but also shows statistically significant improvements in real-time applications. This suggests that precision in tuning stability and plasticity can be critical to the success of adaptive systems, especially in dynamic and physically interactive environments, such as robotic control. The principles and methods in this study could extend to domains where continual learning is essential, enabling adaptive systems to adapt to changing conditions while preserving knowledge without extensive retraining. This work advances the development of robust and adaptable systems that thrive in dynamic environments.

## VI. ACKNOWLEDGEMENT

We appreciate the contributions of all participants in this study, as well as the expertise and resources provided by the EPIC Lab at the Georgia Institute of Technology and the assistance from Sixu Zhou and Hanjun Kim.

## REFERENCES

- [1] L. Wang, X. Zhang, H. Su, and J. Zhu, "A Comprehensive Survey of Continual Learning: Theory, Method and Application," Feb. 2024, arXiv:2302.00487 [cs]. [Online]. Available: <http://arxiv.org/abs/2302.00487>
- [2] G. M. van de Ven and A. S. Tolias, "Three scenarios for continual learning," Apr. 2019, arXiv:1904.07734 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1904.07734>
- [3] M. Mermillod, A. Bugajska, and P. Bonin, "The stability-plasticity dilemma: investigating the continuum from catastrophic forgetting to age-limited learning effects," *Frontiers in Psychology*, vol. 4, p. 504, Aug. 2013. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3732997/>
- [4] H.-J. Schmid, "Understanding the process of entrenchment," in *The Dynamics of the Linguistic System: Usage, Conventionalization, and Entrenchment*, H.-J. Schmid, Ed. Oxford University Press, Jan. 2020, p. 0. [Online]. Available: <https://doi.org/10.1093/oso/9780198814771.003.0011>
- [5] I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, and Y. Bengio, "An Empirical Investigation of Catastrophic Forgetting in Gradient-Based Neural Networks," Mar. 2015, arXiv:1312.6211 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1312.6211>
- [6] M. McCloskey and N. J. Cohen, "Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem," in *Psychology of Learning and Motivation*, G. H. Bower, Ed. Academic Press, Jan. 1989, vol. 24, pp. 109–165. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0079742108605368>
- [7] G. Lin, H. Chu, and H. Lai, "Towards Better Plasticity-Stability Trade-off in Incremental Learning: A Simple Linear Connector," Mar. 2022, arXiv:2110.07905 [cs]. [Online]. Available: <http://arxiv.org/abs/2110.07905>
- [8] L. Wang, M. Zhang, Z. Jia, Q. Li, C. Bao, K. Ma, J. Zhu, and Y. Zhong, "AFEC: Active Forgetting of Negative Transfer in Continual Learning," Nov. 2021, arXiv:2110.12187 [cs]. [Online]. Available: <http://arxiv.org/abs/2110.12187>
- [9] J. Zhang, J. Zhang, S. Ghosh, D. Li, S. Tasci, L. Heck, H. Zhang, and C.-C. J. Kuo, "Class-incremental Learning via Deep Model Consolidation," Jan. 2020, arXiv:1903.07864 [cs]. [Online]. Available: <http://arxiv.org/abs/1903.07864>
- [10] "[1711.09601] Memory Aware Synapses: Learning what (not) to forget." [Online]. Available: <https://arxiv.org/abs/1711.09601>
- [11] S.-A. Rebuffi, A. Kolesnikov, G. Sperl, and C. H. Lampert, "iCaRL: Incremental Classifier and Representation Learning," Apr. 2017, arXiv:1611.07725 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1611.07725>
- [12] A. Douillard, M. Cord, C. Ollion, T. Robert, and E. Valle, "PODNet: Pooled Outputs Distillation for Small-Tasks Incremental Learning," Oct. 2020, arXiv:2004.13513 [cs]. [Online]. Available: <http://arxiv.org/abs/2004.13513>
- [13] S. Hou, X. Pan, C. C. Loy, Z. Wang, and D. Lin, "Learning a Unified Classifier Incrementally via Rebalancing," 2019, pp. 831–839. [Online]. Available: [https://openaccess.thecvf.com/content\\_CVPR\\_2019/html/Hou\\_Learning\\_a\\_Unified\\_Classifier\\_Incrementally\\_via\\_Rebalancing\\_CVPR\\_2019\\_paper.html](https://openaccess.thecvf.com/content_CVPR_2019/html/Hou_Learning_a_Unified_Classifier_Incrementally_via_Rebalancing_CVPR_2019_paper.html)
- [14] Y. Wu, Y. Chen, L. Wang, Y. Ye, Z. Liu, Y. Guo, and Y. Fu, "Large Scale Incremental Learning," 2019, pp. 374–382. [Online]. Available: [https://openaccess.thecvf.com/content\\_CVPR\\_2019/html/Wu\\_Large\\_Scale\\_Incremental\\_Learning\\_CVPR\\_2019\\_paper.html](https://openaccess.thecvf.com/content_CVPR_2019/html/Wu_Large_Scale_Incremental_Learning_CVPR_2019_paper.html)
- [15] P. Dhar, R. V. Singh, K.-C. Peng, Z. Wu, and R. Chellappa, "Learning without Memorizing," Apr. 2019, arXiv:1811.08051 [cs]. [Online]. Available: <http://arxiv.org/abs/1811.08051>
- [16] H. Jung, J. Ju, M. Jung, and J. Kim, "Less-forgetting Learning in Deep Neural Networks," Jul. 2016, arXiv:1607.00122 [cs]. [Online]. Available: <http://arxiv.org/abs/1607.00122>
- [17] Z. Li and D. Hoiem, "Learning without Forgetting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 12, pp. 2935–2947, Dec. 2018, conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. [Online]. Available: <https://ieeexplore.ieee.org/document/8107520>
- [18] A. Chaudhry, P. K. Dokania, T. Ajanthan, and P. H. S. Torr, "Riemannian Walk for Incremental Learning: Understanding Forgetting and Intransigence," 2018, vol. 11215, pp. 556–572, arXiv:1801.10112 [cs]. [Online]. Available: <http://arxiv.org/abs/1801.10112>
- [19] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3521–3526, Mar. 2017, arXiv:1612.00796 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1612.00796>
- [20] Z. Chen and B. Liu, "Continual Learning and Catastrophic Forgetting," in *Lifelong Machine Learning*. Cham: Springer International Publishing, 2018, pp. 55–75, series Title: Synthesis Lectures on Artificial Intelligence and Machine Learning. [Online]. Available: [https://link.springer.com/10.1007/978-3-031-01581-6\\_4](https://link.springer.com/10.1007/978-3-031-01581-6_4)
- [21] T. Lesort, V. Lomonaco, A. Stoian, D. Maltoni, D. Filliat, and N. Díaz-Rodríguez, "Continual Learning for Robotics: Definition, Framework, Learning Strategies, Opportunities and Challenges," Nov. 2019, arXiv:1907.00182 [cs]. [Online]. Available: <http://arxiv.org/abs/1907.00182>
- [22] H. A. Varol, F. Sup, and M. Goldfarb, "Multiclass Real-Time Intent Recognition of a Powered Lower Limb Prosthesis," *IEEE transactions on bio-medical engineering*, vol. 57, no. 3, pp. 542–551, Mar. 2010. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2829115/>
- [23] "Real-Time continuous gait phase and speed estimation from a single sensor | IEEE Conference Publication | IEEE Xplore." [Online]. Available: <https://ieeexplore.ieee.org/document/8062565>
- [24] Y. Liu, H. An, H. Ma, and Q. Wei, "Online Walking Speed Estimation Based on Gait Phase and Kinematic Model for Intelligent Lower-Limb Prosthesis," *Applied Sciences*, vol. 13, no. 3, p. 1893, Jan. 2023, number: 3 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/2076-3417/13/3/1893>
- [25] B.-Y. Su, J. Wang, S.-Q. Liu, M. Sheng, J. Jiang, and K. Xiang, "A CNN-Based Method for Intent Recognition Using Inertial Measurement Units and Intelligent Lower Limb Prosthesis," *IEEE transactions on neural systems and rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 27, no. 5, pp. 1032–1042, May 2019.
- [26] A. J. Young, A. M. Simon, and L. J. Hargrove, "A Training Method for Locomotion Mode Prediction Using Powered Lower Limb Prostheses," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 3, pp. 671–677, May 2014, conference Name: IEEE Transactions on Neural Systems and Rehabilitation Engineering. [Online]. Available: <https://ieeexplore.ieee.org/document/6650103>
- [27] A. M. Simon, L. J. Hargrove, B. A. Lock, and T. A. Kuiken, "A decision-based velocity ramp for minimizing the effect of misclassifications during real-time pattern recognition control," *IEEE Transactions on Biomedical Engineering*, vol. 58, no. 8, pp. 2360–2368, Aug. 2011. [Online]. Available: <http://www.scopus.com/inward/record.url?scp=79960735491&partnerID=8YFLogxK>
- [28] L. J. Hargrove, A. J. Young, A. M. Simon, N. P. Fey, R. D. Lipschutz, S. B. Finucane, E. G. Halsne, K. A. Ingraham, and T. A. Kuiken, "Intuitive control of a powered prosthetic leg during ambulation: a randomized clinical trial," *JAMA*, vol. 313, no. 22, pp. 2244–2252, Jun. 2015.
- [29] H. Huang, F. Zhang, L. J. Hargrove, Z. Dou, D. R. Rogers, and K. B. Englehart, "Continuous Locomotion-Mode Identification for Prosthetic Legs Based on Neuromuscular-Mechanical Fusion," *IEEE transactions on bio-medical engineering*, vol. 58, no. 10, pp. 2867–2875, Oct. 2011. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3235670/>
- [30] Q. M. Billah, L. Rahman, J. Adan, A. M. Kamal, M. K. Islam, C. Shahnaz, and A. Subhana, "Design of Intent Recognition System in a Prosthetic Leg for Automatic Switching of Locomotion Modes," in *TENCON 2019 - 2019 IEEE Region 10 Conference (TENCON)*, Oct. 2019, pp. 1638–1642, iSSN: 2159-3450. [Online]. Available: <https://ieeexplore.ieee.org/document/8929624>
- [31] "Evolutionary optimization of user intent recognition for transfemoral amputees | IEEE Conference Publication | IEEE Xplore." [Online]. Available: <https://ieeexplore.ieee.org/document/7348280>
- [32] K. Zhang, C. Xiong, W. Zhang, H. Liu, D. Lai, Y. Rong, and C. Fu, "Environmental Features Recognition for Lower Limb Prostheses Toward Predictive Walking," *IEEE transactions on neural systems and rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 27, no. 3, pp. 465–476, Mar. 2019.

- [33] H. T. T. Vu, F. Gomez, P. Chelle, D. Lefeber, A. Nowé, and B. Vanderborcht, "ED-FNN: A New Deep Learning Algorithm to Detect Percentage of the Gait Cycle for Powered Prostheses," *Sensors*, vol. 18, no. 7, p. 2389, Jul. 2018, number: 7 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/1424-8220/18/7/2389>
- [34] I. Kang, D. Molinaro, S. Duggal, Y. Chen, P. Kunapuli, and A. Young, "Real-Time Gait Phase Estimation for Robotic Hip Exoskeleton Control During Multimodal Locomotion," *IEEE Robotics and Automation Letters*, vol. PP, pp. 1–1, Feb. 2021.
- [35] J. W. Sensinger, B. A. Lock, and T. A. Kuiken, "Adaptive pattern recognition of myoelectric signals: Exploration of conceptual framework and practical algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 17, no. 3, pp. 270–278, 2009.
- [36] J. A. Spanias, A. M. Simon, and L. J. Hargrove, "Across-user adaptation for a powered lower limb prosthesis," *IEEE ... International Conference on Rehabilitation Robotics: [proceedings]*, vol. 2017, pp. 1580–1583, Jul. 2017.
- [37] R. B. Woodward, A. M. Simon, E. A. Seyforth, and L. J. Hargrove, "Real-Time Adaptation of an Artificial Neural Network for Transfemoral Amputees Using a Powered Prosthesis," *IEEE transactions on bio-medical engineering*, vol. 69, no. 3, pp. 1202–1211, Mar. 2022. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8988236/>
- [38] I. Kang, P. Kunapuli, and A. J. Young, "Real-time neural network-based gait phase estimation using a robotic hip exoskeleton," *IEEE Transactions on Medical Robotics and Bionics*, vol. 2, no. 1, pp. 28–37, 2020.
- [39] 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)*. Atlanta, GA, USA: IEEE, Apr. 2023, pp. 1–7. [Online]. Available: <https://ieeexplore.ieee.org/document/10130251/>
- [40] jairo, "replace me," *replace me*, vol. 2, no. 1, pp. 28–37, 2020.
- [41] A. J. Young, A. M. Simon, N. P. Fey, and L. J. Hargrove, "Intent recognition in a powered lower limb prosthesis using time history information," *Annals of Biomedical Engineering*, vol. 42, no. 3, pp. 631–641, Mar. 2014.
- [42] A. J. Young and L. J. Hargrove, "A Classification Method for User-Independent Intent Recognition for Transfemoral Amputees Using Powered Lower Limb Prostheses," *IEEE transactions on neural systems and rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 24, no. 2, pp. 217–225, Feb. 2016.
- [43] A. F. Azocar, L. M. Mooney, J.-F. Duval, A. M. Simon, L. J. Hargrove, and E. J. Rouse, "Design and clinical implementation of an open-source bionic leg," *Nature Biomedical Engineering*, vol. 4, no. 10, pp. 941–953, Oct. 2020, publisher: Nature Publishing Group. [Online]. Available: <https://www.nature.com/articles/s41551-020-00619-3>
- [44] A. F. Azocar, L. M. Mooney, L. J. Hargrove, and E. J. Rouse, "Design and Characterization of an Open-Source Robotic Leg Prosthesis," *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, pp. 111–118, Aug. 2018, conference Name: 2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob) ISBN: 9781538681831 Place: Enschede Publisher: IEEE. [Online]. Available: <https://ieeexplore.ieee.org/document/8488057/>
- [45] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, Mar. 1960. [Online]. Available: <https://doi.org/10.1115/1.3662552>
- [46] K. Bhakta, J. Camargo, P. Kunapuli, L. Childers, and A. Young, "Impedance control strategies for enhancing sloped and level walking capabilities for individuals with transfemoral amputation using a powered multi-joint prosthesis," *Military Medicine*, vol. 185, no. Supplement\_1, pp. 490–499, January-February 2020. [Online]. Available: <https://doi.org/10.1093/milmed/usz229>
- [47] A. M. Simon, K. A. Ingraham, N. P. Fey, S. B. Finucane, R. D. Lipschutz, A. J. Young, and L. J. Hargrove, "Configuring a powered knee and ankle prosthesis for transfemoral amputees within five specific ambulation modes," *PLoS One*, vol. 9, no. 6, p. e99387, Jun 2014.
- [48] S. Heins, L. Flynn, J. Geeroms, D. Lefeber, and R. Ronsse, "Torque control of an active elastic transfemoral prosthesis via quasi-static modelling," *Robotics and Autonomous Systems*, vol. 107, pp. 100–115, Sep. 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889018300356>
- [49] Y. Chen, Y. Kang, Y. Chen, and Z. Wang, "Probabilistic Forecasting with Temporal Convolutional Neural Network," Mar. 2020, arXiv:1906.04397 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1906.04397>
- [50] "Subject-Independent, Biological Hip Moment Estimation During Multimodal Overground Ambulation Using Deep Learning | IEEE Journals & Magazine | IEEE Xplore." [Online]. Available: <https://ieeexplore.ieee.org/document/9687847>
- [51] A. Kose, A. Cereatti, and U. Della Croce, "Estimation of traversed distance in level walking using a single inertial measurement unit attached to the waist," *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference*, vol. 2011, pp. 1125–1128, 2011.
- [52] Q. Li, M. Young, V. Naing, and J. M. Donelan, "Walking speed estimation using a shank-mounted inertial measurement unit," *Journal of Biomechanics*, vol. 43, no. 8, pp. 1640–1643, May 2010.
- [53] K. 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)*, May 2023, pp. 12665–12671. [Online]. Available: <https://ieeexplore.ieee.org/document/10160770>
- [54] "Sensors | Free Full-Text | Estimation of Walking Speed and Its Spatiotemporal Determinants Using a Single Inertial Sensor Worn on the Thigh: From Healthy to Hemiparetic Walking." [Online]. Available: <https://www.mdpi.com/1424-8220/21/21/6976>
- [55] D. Lopez-Paz and M. Ranzato, "Gradient Episodic Memory for Continual Learning," Sep. 2022, arXiv:1706.08840 [cs]. [Online]. Available: <http://arxiv.org/abs/1706.08840>
- [56] L. N. Smith, "Cyclical Learning Rates for Training Neural Networks," Apr. 2017, arXiv:1506.01186 [cs]. [Online]. Available: <http://arxiv.org/abs/1506.01186>
- [57] M. D. Zeiler, "ADADELTA: An Adaptive Learning Rate Method," Dec. 2012, arXiv:1212.5701 [cs]. [Online]. Available: <http://arxiv.org/abs/1212.5701>
- [58] J. S. Vitter, "Random sampling with a reservoir," *ACM Transactions on Mathematical Software*, vol. 11, no. 1, pp. 37–57, Mar. 1985. [Online]. Available: <https://dl.acm.org/doi/10.1145/3147.3165>
- [59] T. L. Hayes, K. Kafle, R. Shrestha, M. Acharya, and C. Kanan, "REMIND Your Neural Network to Prevent Catastrophic Forgetting," in *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII*. Berlin, Heidelberg: Springer-Verlag, Aug. 2020, pp. 466–483. [Online]. Available: [https://doi.org/10.1007/978-3-030-58598-3\\_28](https://doi.org/10.1007/978-3-030-58598-3_28)
- [60] K. Bhakta, J. Camargo, W. Compton, K. Herrin, and A. Young, "Evaluation of continuous walking speed determination algorithms and embedded sensors for a powered knee ankle prosthesis," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4820–4826, 2021.
- [61] B. Dauriac, X. Bonnet, H. Pillet, and F. Lavaste, "Estimation of the walking speed of individuals with transfemoral amputation from a single prosthetic shank-mounted imu," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 233, no. 9, pp. 931–937, Sep 2019.
- [62] S. Miyazaki, "Long-term unrestrained measurement of stride length and walking velocity utilizing a piezoelectric gyroscope," *IEEE Transactions on Biomedical Engineering*, vol. 44, no. 8, pp. 753–759, 1997.
- [63] J. Maldonado-Contreras, C. Johnson, S. Zhou, H. Kim, I. Knight, K. R. Herrin, and A. J. Young, "Real-time adaptation of deep learning walking speed estimators enables biomimetic assistance modulation in an open-source bionic leg," *Transactions on Biomedical Engineering*, 2024, in review.
- [64] J. A. Spanias, A. M. Simon, S. B. Finucane, E. J. Perreault, and L. J. Hargrove, "Online adaptive neural control of a robotic lower limb prosthesis," *Journal of Neural Engineering*, vol. 15, no. 1, p. 016015, Feb 2018.

APPENDIX A - EXPERIMENTAL PROTOCOL

Trial Number	Description	Output
<b>Baseline Data (1 trial, 2 mins)</b>		
1	Use baseline model for forward prediction; repeat walking profile	Baseline accuracy + adaptation test data
<b>Stability Demonstration (3 trials, 8 mins)</b>		
2a	Adapt using S&P Controlled Method; freeze at 0.3 and 0.4 m/s	Trained Model 1
2b	Test Model 1; walk at 0.3 and 0.4 m/s for 15 seconds, no adaptation	Accuracy at 0.3 and 0.4 m/s
3a	Adapt using S&P Controlled Method; freeze at 0.5 and 0.6 m/s	Trained Model 2
3b	Test Model 2; walk at 0.5 and 0.6 m/s for 15 seconds, no adaptation	Accuracy at 0.5 and 0.6 m/s
4a	Adapt using S&P Controlled Method; freeze at 0.7 and 0.8 m/s	Trained Model 3
4b	Test Model 3; walk at 0.7 and 0.8 m/s for 15 seconds, no adaptation	Accuracy at 0.7 and 0.8 m/s
<b>Plasticity Demonstration (3 trials, 8 mins)</b>		
5a	Adapt using S&P Controlled Method; no training at 0.3 and 0.4 m/s	Trained Model 1
5b	Test Model 1; walk at 0.3 and 0.4 m/s for 15 seconds, no adaptation	Accuracy at 0.3 and 0.4 m/s
6a	Adapt using S&P Controlled Method; no training at 0.5 and 0.6 m/s	Trained Model 2
6b	Test Model 2; walk at 0.5 and 0.6 m/s for 15 seconds, no adaptation	Accuracy at 0.5 and 0.6 m/s
7a	Adapt using S&P Controlled Method; no training at 0.7 and 0.8 m/s	Trained Model 3
7b	Test Model 3; walk at 0.7 and 0.8 m/s for 15 seconds, no adaptation	Accuracy at 0.7 and 0.8 m/s
<b>Method Error (2 trials, 4 mins)</b>		
8	Adapt using S&P Controlled Method; trained on all speeds	Trained Model 1
9	Test Model 1; repeat walking profile	Accuracy with adapted predictor

TABLE II: Experiment Trials and Results

---

**Algorithm 1** One-Shot Conditional Continuous Learning for Adaptive Speed Estimation

---

**Require:**  $\Theta$ : Forward predictor model,  $\mathcal{D}_{\text{init}}$ : Initial data,  $\mathcal{B}$ : Backwards estimator

**Require:**  $\Delta s$ : Strides per tuning,  $k$ : Stiffness,  $b$ : Damping

```

0:  $\mathcal{F} \leftarrow \text{Initialize}(\Theta, \mathcal{D}_{\text{init}})$  {Initial forward predictor}
0:  $\mathcal{T} \leftarrow \{\}$ ;  $\mathcal{S} \leftarrow [\ ]$ ;  $\hat{\mathcal{D}} \leftarrow [\ ]$ ;  $s \leftarrow 0$  {Initialize bins, test data, gait data, stride count}
0: while trial ongoing do
0:    $\mathcal{D}_{\text{filtered}} \leftarrow \text{KalmanFilter}(\text{ReceiveData}(), \hat{\mathcal{D}}, \mathcal{F})$ 
0:    $\hat{\mathcal{D}} \leftarrow \mathcal{D}_{\text{filtered}} + \hat{\mathcal{D}}$ ;  $v \leftarrow \text{PredictSpeed}(\mathcal{D}_{\text{filtered}})$ 
0:    $k_s \leftarrow k(1 + 0.422(v - 0.5))$ 
0:   for each joint  $i$  do
0:      $\tau_i \leftarrow -k_i(\theta_i - \theta_{ei}) - b\dot{\theta}_i$ 
0:   end for
0:   if StrideComplete( $\hat{\mathcal{D}}$ ) then
0:      $s \leftarrow s + 1$ 
0:     if  $s \bmod \Delta s = 0$  then
0:        $\mathcal{D}_{\text{odd}}, \mathcal{D}_{\text{even}} \leftarrow \text{Split}(\hat{\mathcal{D}})$ 
0:        $\bar{\mathcal{L}}_{\text{even}}, \bar{\mathcal{L}}_{\text{odd}} \leftarrow \text{MAV}(\mathcal{B}(\mathcal{D}_{\text{even}})), \text{MAV}(\mathcal{B}(\mathcal{D}_{\text{odd}}))$ 
0:        $\mathcal{S} \leftarrow \mathcal{S} + (\mathcal{D}_{\text{even}}, \bar{\mathcal{L}}_{\text{even}})$ 
0:        $\mathcal{T} \leftarrow \text{Enqueue}(\mathcal{D}_{\text{odd}}, \bar{\mathcal{L}}_{\text{odd}})$ 
0:        $\mathcal{F}' \leftarrow \text{FineTune}(\mathcal{F}, \mathcal{T})$ 
0:       if EvaluateAccuracy( $\mathcal{F}', \mathcal{S}$ ) > EvaluateAccuracy( $\mathcal{F}, \mathcal{S}$ ) then
0:          $\mathcal{F} \leftarrow \mathcal{F}'$  {Update predictor if accuracy improves}
0:       end if
0:        $s \leftarrow 0$ ;  $\hat{\mathcal{D}} \leftarrow \{\}$ 
0:     end if
0:   end if
0: end while=0

```

---

## APPENDIX C – DYNAMIC $\lambda$ CONTROL

We designed an additional system to dynamically modify  $\lambda$  in real-time based on input data stream characteristics to increase plasticity when appropriate (e.g. changes in label distributions) and increase stability when appropriate (e.g. noisy data). In the following, we test our hypothesis that this dynamic control is more error-wise advantageous than using static  $\lambda$  values in CL systems.

### A. Methods

From Optimal Balance Theory, we have a method to modify the stability and plasticity characteristics of the network in the CL pipeline to achieve an optimal balance, enabling the tuning of  $\lambda$  for tasks requiring extreme stability (e.g., repetitive, long-range patterns) or high plasticity (e.g., constantly varying environments). Building on this, we can create a dynamic system that adjusts  $\lambda$  in real time based on incoming data, incorporating the entire stability-plasticity control system into an adaptive feedback loop, as shown in Fig. 1.

We define this update rule generally using the relation:

$$\lambda_{t+1} = \lambda_t \cdot (1 + W^T \cdot F) \quad (21)$$

where  $\lambda_t$  is  $\lambda$  at time  $t$  and  $F$  is a factor vector consisting of metrics computed from the input data stream at every iteration which is scaled in effect by  $W$ , a tunable weight vector. For our purposes, we construct  $F$  using three illustrative factors: label distribution shifts (L), data noise (N) and error rate (E). The greater the frequency of the label shifts, the more plastic the model should behave to adapt to new environmental conditions, thus we define:

$$L(l_i) = |P_{\text{current}}(l_i) - P_{\text{historical}}(l_i)| \quad (22)$$

$$P(l_i) = \frac{F(l_i)}{\sum_{j=1}^n F(l_j)} \quad (23)$$

where  $F(l_i)$  is the historical count of label  $i$ . We scale stability proportionate to  $N$  to protect the model from fine-tuning on unreliable data, and define  $N$  as the difference between Kalman-filtered and unfiltered data. Finally, we define  $E$  as MAE(Ground Truth, Prediction) or MAE(Backward Estimate, Prediction), depending on the availability of ground-truth labels.

We tune the weight matrix using a grid search in the range  $[-1, 1]$  and step size 0.05 for each individual weight, according to the accuracy of the adaptation trial of all participants with TFA. We obtain the weights:

$$W = \begin{bmatrix} L \\ N \\ E \end{bmatrix} = \begin{bmatrix} -0.1 \\ 0.3 \\ -0.05 \end{bmatrix} \quad (24)$$

To demonstrate the functionality of dynamically modifying  $\lambda$  in conjunction with the other two subsystems, we conduct a simple comparison. We train the forward predictor and sweep

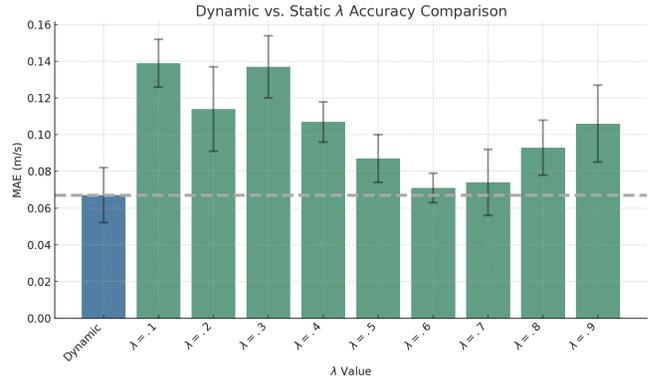


Fig. 8: Results from offline comparison between Dynamic  $\lambda$  Modification and static  $\lambda$  values. Dynamic balancing improves the best  $\lambda$  value of .6 by 5.6%.

participants identically to the method used to assess the SGD-MIMO subsystem. In one trial,  $\lambda$  is dynamically modified using the relation and weight matrix described previously throughout each adaptation trial. In the other nine trials, we use static  $\lambda$  values in  $\{0.1, \dots, 0.9\}$ . This procedure explicitly tests the superiority of dynamically modifying  $\lambda$  versus setting it statically.

### B. Results

As shown in Fig. 8, we found that the best performance with a static  $\lambda$  value is 0.6 with an MAE of 0.071 m/s. Dynamic  $\lambda$  Modification achieved a slightly lower error rate of 0.067 (5.6% improvement).

### C. Discussion

We hypothesized that using Dynamic  $\lambda$  Control, we attain more error-wise optimal stability and plasticity values to drive the models towards achieving than static  $\lambda$  values in CL systems. The average results from the offline comparison between dynamic  $\lambda$  Control and static  $\lambda$  trials indicate that dynamic adjustment of the stability-plasticity balance leads to more optimal outcomes, but the results lacked statistical significance, and thus the hypothesis is rejected. Dynamic  $\lambda$  Control achieved an MAE of 0.067 m/s, a 5.6% improvement over the best static  $\lambda$  value (0.071 m/s). Although this improvement is small, it highlights the adaptability of the dynamic system to respond to varying environmental conditions and learning requirements. Furthermore, while the statistical difference between Dynamic  $\lambda$  and the best static  $\lambda$  value was not significant enough to justify further online evaluation, the trend toward lower error rates with dynamic control suggests that it can be particularly beneficial in more complex or rapidly changing scenarios than simple treadmill speed-varying trials.