Decision Making for Contextual Neuroprosthetic Control with Dendritic Gated Networks

Laura C. Petrich
Department of Computing Science
University of Alberta
Alberta Machine Intelligence Institute
Edmonton, AB, Canada
laurapetrich@ualberta.ca

Heather E. Williams
Institute of Biomedical Engineering
University of New Brunswick
Fredericton, NB, Canada
heather.williams@unb.ca

Matthew E. Taylor
Department of Computing Science
University of Alberta
Alberta Machine Intelligence Institute
Edmonton, AB, Canada
mtaylor3@ualberta.ca

Patrick M. Pilarski
Department of Medicine
Department of Computing Science
University of Alberta
Alberta Machine Intelligence Institute
Edmonton, AB, Canada
pilarski@ualberta.ca

Abstract

In this work, we examine how an emerging computational neural network architecture—Dendritic Gated Networks (DGNs)—may allow machine agents to detect and accommodate changing contexts while learning. Specifically, we study how DGNs impact real-time, sequential decision making by a machine agent as it supports a human user in their control of a neuroprosthesis. Machine learning algorithms for upper-limb prosthetic control are required to make moment-by-moment decisions on how to map biosignals flowing from the human body to movements of the person's body-affixed robotic device. When the context changes, as it does during activities in a person's daily home and working life, such mappings can become inappropriate for a new setting and device performance can be compromised. Dendritic gating enables the acquisition of new learned patterns while retaining previously learned patterns, even during continual or ongoing learning. DGN performance was compared against the current clinical and commercial standard, linear discriminant analysis (LDA), on biosignal datasets across four contextual settings known to impact model performance over time: differing arm postures, electrode displacement, varying days of use, and minimal training data. We found that DGNs are able to accommodate for changing contexts and exhibit less inter-subject variability than LDA. Interestingly, we discovered that DGNs are *less* confident in their incorrect predictions than LDA. When combined with a rejection threshold, this enables a substantial reduction in the number of incorrect predictions made by the model that would be mapped to robotic movements. DGNs appear to combine the benefits of both linear and deep learning to achieve sample efficient training on complex, non-linear data. This work represents the first investigation of DGNs trained on biosignal data in a real-world setting, and our findings suggest that DGNs hold promise for deployment in human-machine shared decision making scenarios.

Keywords: real-world decision making, neuroprostheses, machine learning, neural networks, human-machine interaction

Acknowledgements

The authors thank Johannes Günther for insight on this manuscript, Ethan Eddy for coding support, and Shaylee Lorrain for graphical design support. This work took place in the Bionic Limbs for Improved Natural Control (BLINC) and the Intelligent Robot Learning (IRL) labs at the University of Alberta, as this work was supported in part by research grants from Alberta Innovates, Alberta Machine Intelligence Institute (Amii), the Canada CIFAR AI Chairs Program, DRAC (Digital Research Alliance of Canada), the National Science and Engineering Research Council (NSERC), and the Department of Defense Peer Reviewed Orthopaedic Research Program (PRORP) Award HT9425-23-1-0398.

1 Introduction and Background

Machine learning models for neuroprosthetic control make decisions about what action the robotic arm should take with the goal of predicting the user's intent. These decisions happen continuously over time, with the system acquiring biosignals (typically electromyographic (EMG) signals generated by residual muscles), making predictions of what action the user intended to do, deciding whether to use the new prediction, and then translating the predictions into control signals to move the robotic arm or hand. Context is important when making these predictions because even though there are a finite number of actions (e.g., elbow flexion or closing the hand) the prosthetic arm is capable of making, the space of intended daily tasks and different physical environments, social settings, homeostatic or physiological states is infinite. Context provides us with the means to make informed decisions and brings understanding of the world around us and within us. It helps us see how different pieces of information are related to each other and informs how we behave in different situations. Without context, it is easy to see how information could be misinterpreted leading to the wrong decisions being made by a machine agent. For example, if a prosthetic control model was trained from data gathered in the user's home, when they try and use it after going to the gym their muscle signals would be different due to muscle fatigue and electrodes shifting; the machine agent may predict the user would like to open the hand instead of closing it when picking up their phone, resulting in dropping the phone to the ground. Currently, neuroprosthetic control models are unable to detect and accommodate for changing contexts such as these. When the context changes (as it continuously does in the real world) their performance plummets and it's easy to understand why the rate of device abandonment for upper limb prostheses remains high ($\sim 44\%$) in the real world, despite significant technological advances (Salminger et al., 2022). Understanding context would help give meaning to the acquired biosignals as the model would be able to adapt to the EMG signal non-stationary in real life environments. The model would be able to understand that regardless of where the user is or how they are feeling the intended action remains the same. Dendritic gated networks (DGNs) could be the answer for detecting and accommodating to changing contexts when making decisions during myoelectric control. In this work, we carry out the first investigation of DGNs in a real world setting where decisions need to be made continuously and incorrect decisions can have a significant impact on a person's quality of life.

Dendritic gated networks are feed-forward artificial neural networks proposed by Sezener et al. (2021) that combine dendritic gating with local learning rules to solve both classification and regression tasks. Each neuron in the network has a number of dendritic branches that selectively updates the connecting weights depending on the context. DGNs function as *one-vs-all* classifiers, with a DGN trained for each class and then ensembled to output the final prediction. Petrich et al. (2024a) recently showed that DGNs are capable of making predictions suitable for prosthetic control in different limb positions (i.e., different contexts). DGNs are made up of multiple layers of neurons where the input to each layer is made up of a linear combination of the activity in the previous layer. Each neuron has a number of dendritic branches that receives input from the previous layer and side information via the branches inhibitory interneuron. The last layer consists of one neuron that outputs the final prediction. A gating function decides whether a branch should be

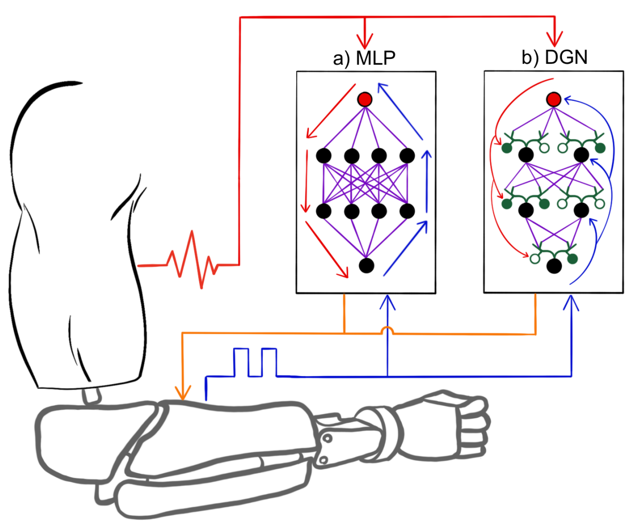

Figure 1: **Comparison of MLPs and DGNs**. Biological signals (red) are used as input for each network with the intended motor control action (orange) as output. Within the networks arrows show feedforward computations (red) and feedback information (blue). DGNs (b) differ from MLPs (a) such that each inhibitory interneuron (green circles) directly receives the input signal, each neuron receives the target (blue) directly instead of backpropagating error, and each dendritic branch (green) receives input from all prior-layer neurons. *Figure originally published in Petrich et al. (2024a).*

gated on or off, resulting in input-dependent weight updating. Figure 1 shows how the DGN architecture differs from a multilayer perceptron (MLP), a foundational learning block of deep learning methods.

Due to their feedforward architecture and dendritic branching, DGNs exhibit the benefits of both deep and linear learning methods, without many of the disadvantages of either. Linear learning methods are easy to understand and interpret, less prone to overfitting (a common problem with deep learning models), and can learn quickly from a small amount of data. However, they tend to perform poorly if the relationship between the input features and target variable is not linear, require careful feature engineering, and do not generalized well to data differing from those seen during training (e.g., when there is a dramatic shift in input space). Conversely, deep learning methods are able to learn complex, non-linear relationships and do not require feature engineering, however they are challenging to interpret, computationally expensive, and require a lot of training data. DGNs display many of the benefits of deep learning methods with the added benefits from linear learning methods such as sample efficient learning, being less prone to overfitting, and delivering increased interpretability. Contrary to the machine learning models typically (clinically and commercially) used in upper-limb prosthetic control, DGNs have been shown to be able to represent essentially arbitrary nonlinear functions due to the gating of their dendritic branches. Furthermore, DGNs have been shown to be highly resilient to forgetting old tasks when learning new ones, a common problem in deep learning known as *catastrophic forgetting* (Sezener et al., 2021). In terms of computational complexity, DGNs take longer to train and are more costly than linear learning methods, but less so than deep learning methods. Despite promising results from previous works there remains a paucity of research investigating DGNs on real world data (Petrich et al., 2024a). This work presents the first comparison of dendritic gated networks and a state-of-the-art linear learning method in prosthetic control, showcasing their performance across diverse EMG datasets and changing contexts.

2 Methods and Results

We carried out offline analyses of five electromyographic (EMG) datasets, each chosen specifically because it covered a context of use relevant to the instability of myoelectric prosthetic control as described by Kyranou et al. (2018) and Campbell et al. (2024): *arm posture variability*; *electrode displacement*; *cross-day use*; and a *lack of training data*. Linear discriminant analysis (LDA) is the current clinical and commercial standard for neuroprosthetic control and serves as the baseline to compare DGNs with. We report the classification accuracy for each dataset variation both before and after applying a rejection threshold of 0.8 (i.e., any prediction with a confidence of less than 0.8 was rejected). An individualized model was trained for each participant and results reported are averaged across participants and five random seeds. We further looked into whether the variability was due to the different random seeds or variability between participants and concluded that the main source was due to differences between participants. This is consistent with the findings of Kyranou et al. (2018) that inter-subject variability is one of the five key causes of surface EMG signal changes that affect model performance.

We followed similar hyperparameter settings as Simon et al. (2023). Each classifier was trained with 200-ms analysis windows and a 25-ms update increment on the Hudgins feature set (Hudgins et al., 1993) consisting of the following time domain features: mean absolute value (MAV), waveform length (WL), zero crossings (ZC), and slope sign change (SSC). Following Simon et al., we also included the autoregressive coefficients feature with an order of $o = 7$, which showed a slight inflection point after sweeping over $o \in \{2, ..., 12\}$. The same feature set was used for training all control algorithms in this study. To tune the DGNs we performed parameter sweeps over the worst, median, and best participants from the 3DC dataset. The best performing parameters were then frozen for all experiments included in this study. We swept over number of hidden layers $K \in \{1, 2, 3\}$ with number of dendritic branches $b \in \{32, 64, 128, 256, 512, 1024, 2048\}$, and learning rate $\eta \in \{1e^{-4}, 1e^{-5}, 1e^{-6}\}$. The DGN parameters were frozen as a one neuron one hidden layer network with 128 dendritic branches. Each DGN was trained for 30 epochs with a learning rate of $1e^{-5}$.

No context change between the training and testing sets. To serve as a baseline, we chose the *Full Reachable Workspace (FRW)* dataset (Petrich et al., 2024b). The *FRW* dataset was split into one trial for training the models and two trials for testing. Each trial contains data from all the actions carried out in all nine limb positions, therefore there is no context change between how the model was trained and the context in which it was evaluated. It features sixteen able-bodied participants performing eleven upper limb actions (including elbow flexion/extension) in nine different limb positions. The data was collected with one arm braced which elicits isometric muscle contractions to simulate the muscle behaviour of a person with amputation. LDA has a base accuracy of 0.6820 (SD 0.1012) with an accuracy of 0.7762 (SD 0.0961) after rejection was applied with rejection rate 0.2752 (SD 0.0892). DGN has an accuracy of 0.7049 (SD 0.0699) with 0.8514 (SD 0.0504) after rejection with rejection rate 0.4145 (SD 0.0610) (see Fig. 2a).

Differing arm postures. To cover the context shift of differing arm postures (i.e., the testing set includes limb positions that are not seen during training), we split the *FRW* dataset so that the training set consists of data collected in one limb position (with the participants arms down by their side) and the testing set consisting of the data from the remaining eight limb positions (we will refer to this variation as *FRW-LP*). LDA has a base accuracy of 0.4925 (SD 0.0920) with an

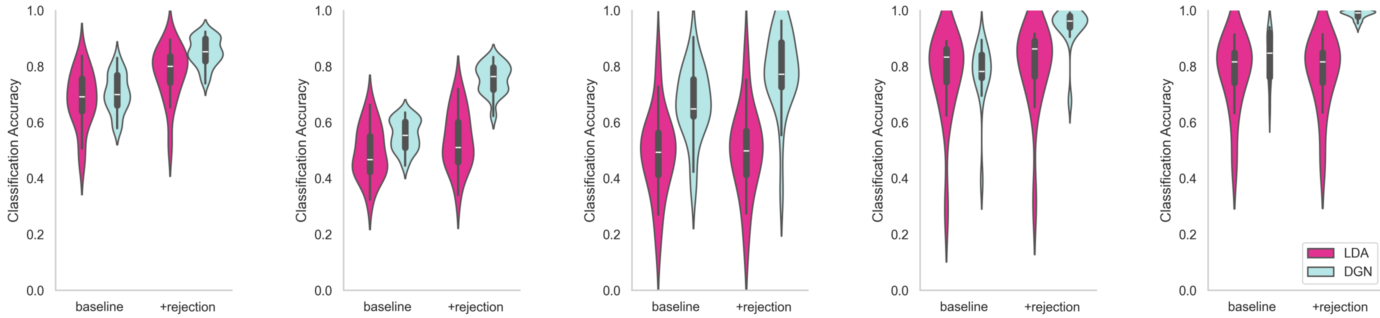

(a) No change in context (b) Arm posture (c) Electrode shift (d) Days of use (e) Minimal data

Figure 2: Results from the biosignal datasets investigated in this work. Each dataset targets a specific change in context known to affect real world model performance. The (a) FRW dataset results serves as the baseline as there is no context change between how the model was trained and how it is used; the (b) FRW-LP dataset shows performance under the context of differing arm postures (i.e., limb position); the (c) CIIL electrode shift dataset results shows the context of displaced electrodes (i.e., electrodes shifting within the prosthetic socket); the (d) MyoDisCo cross day dataset mimics the doffing and donning of the prosthesis on different days; and the (e) CIIL minimal data shows performance when a small amount of data is available for training.

accuracy of 0.5376 (SD 0.1036) after rejection was applied with rejection rate 0.1782 (SD 0.0442). DGN has an accuracy of 0.5558 (SD 0.0552) with 0.7547 (SD 0.0543) after rejection with rejection rate 0.5594 (SD 0.0401) (see Fig. 2b).

Electrode displacement. The *CIIL Electrode Shift* dataset (Campbell et al., 2024) captures when the input space has shifted substantially from the training data (e.g., the electrodes have shifted in the prosthetic socket), this was done by shifting each electrode by 45 degrees for collecting the testing data. LDA has a base accuracy of 0.4883 (SD 0.1519) with an accuracy of 0.4976 (SD 0.1592) after rejection was applied with rejection rate 0.0785 (SD 0.0510). DGN has an accuracy of 0.6653 (SD 0.1288) with 0.7820 (SD 0.1426) after rejection with rejection rate 0.4409 (SD 0.0818) (see Fig. 2c).

Varying days of use. The *MyoDisCo Cross Day* dataset (Eddy et al., 2024) mimics the real world scenario of prosthetic users doffing and donning their device between different days of use. Doffing and donning results in shifts in the electrode placement as well as other factors that change between days such as varying levels of muscle fatigue. LDA has a base accuracy of 0.7644 (SD 0.1598) with an accuracy of 0.7951 (SD 0.1588) after rejection was applied with rejection rate 0.0828 (SD 0.0420). DGN has an accuracy of 0.7727 (SD 0.1182) with 0.9409 (SD 0.0769) after rejection with rejection rate 0.4809 (SD 0.0768) (see Fig. 2d).

Minimal training data. The *CIIL Minimal Data* dataset (Campbell et al., 2024) captures the setting where only a small amount of data is available during model training, therefore there will be substantially more variability in the biosignals during testing than was seen during training. LDA has a base accuracy of 0.7751 (SD 0.1333) with an accuracy of 0.7759 (SD 0.1330) after rejection was applied with rejection rate 0.0024 (SD 0.0017). DGN has an accuracy of 0.8364 (SD 0.0873) with 0.9843 (SD 0.0167) after rejection with rejection rate 0.7734 (SD 0.0920) (see Fig. 2e). Note that the *FRW-LP* dataset also captures the minimal training data context as the testing set contains eight times more data than the training set.

3 Discussion and Conclusion

Overall, our results show that not only do DGNs perform no worse than LDA, the current state-of-the-art algorithm commonly (clinically and commercially) used in neuroprosthetic control, but they exhibit lower variability between participants in all cases except for the electrode displacement experiment. In all experiments, the worst performing participant was higher for DGNs than LDA and applying a rejection threshold had more of an impact on DGN performance. In both the change in arm posture (Fig. 2b) and electrode displacement (Fig. 2c) experiments, DGNs performed notably better than LDA even before rejection was applied. This suggests that DGNs can detect and accommodate for the changing context in the biosignals. When faced with changing contexts, Figs. 2b, 2c, 2d, and 2e, the DGN results clearly show a larger increase in classification accuracy after applying a rejection threshold than LDA, highlighting how DGNs are able to more effectively apply rejection to tease out incorrect predictions. A closer look at the predictions of each model revealed that DGNs become less confident in their incorrect predictions, whereas LDAs are highly confident that their incorrect predictions are correct. This allows us to make a more informed decision as to which predictions should be rejected and apply an appropriate rejection threshold. Future work will look into optimizing the rejection threshold when we carry out experiments with real world prosthesis users. Even in the *FRW* results (Fig. 2a) with no context change between the training and testing environments, the reduction in inter-participant variability suggests that DGNs are able to better pick up on the patterns common between participants that are lost when using linear learning methods such as LDA. The minimal training data results (Figs. 2e and 2b) highlights how DGNs are not only able to accommodate for changing context, as there was substantially more variability in the biosignals in the testing set than in the training set, but also DGNs ability for sample efficient training.

Dendritic gated networks offer powerful new representations that show promise for contextualizing machine agent learning and actions during neuroprosthetic control, where continuous decision making from real-world data streams is required and incorrect decisions are detrimental to a persons quality of life. Incorrect decisions result in the prosthetic device carrying out actions the user never intended to do, this can be frustrating and contributes to the high rates of prosthesis abandonment (Salminger et al., 2022). In this work, we examined how DGNs translate to real-world biosignal data and believe they will be instrumental to continual or ongoing learning after deployment in the real world. Furthermore, DGNs provide the benefits of linear learning methods, such as sample efficient training and interpretability, combined with the ability of deep learning to learn complex, non-linear patterns and representations as well as removing the need for manual feature engineering. Next steps include taking a closer look at the activation patterns of the dendritic branches so we can begin to understand how real world data streams influence how decisions are made depending on context in the flow of time. DGNs similarity to how the cerebellum works could also help us learn about how biological systems make decisions (Sezener et al., 2021). Machine agents need context to be able to make informed decisions in the real world and we believe dendritic gated networks are the key for successful deployment of contextual neuroprosthetic control that learns and adapts with the user.

References

Campbell, E., Eddy, E., Bateman, S., Côté-Allard, U., and Scheme, E. (2024). Context-informed incremental learning improves both the performance and resilience of myoelectric control. *Journal of NeuroEngineering and Rehabilitation*, 21(1):70.

Eddy, E., Campbell, E., Bateman, S., and Scheme, E. (2024). Understanding the influence of confounding factors in myoelectric control for discrete gesture recognition. *Journal of Neural Engineering*, 21(3):036015.

Hudgins, B., Parker, P., and Scott, R. (1993). A new strategy for multifunction myoelectric control. *IEEE Transactions on Biomedical Engineering*, 40(1):82–94.

Kyranou, I., Vijayakumar, S., and Erden, M. S. (2018). Causes of performance degradation in non-invasive electromyographic pattern recognition in upper limb prostheses. *Frontiers in Neurorobotics*, 12:58:1–58:22.

Petrich, L. C., Williams, H. E., Günther, J., Taylor, M. E., and Pilarski, P. M. (2024a). Feasibility study of dendritic gated networks for upper limb prosthetic control. In: Pons, J. L., Tornero, J., and Akay, M., editors, *Converging Clinical and Engineering Research on Neurorehabilitation V. ICNR 2024*, pages 496–500, Cham. Springer Nature Switzerland.

Petrich, L. C., Williams, H. E., Taylor, M. E., Hebert, J. S., Lemelin, P., Shehata, A. W., and Pilarski, P. M. (2024b). A preliminary investigation into bio-inspired data collection for transhumeral targeted muscle reinnervation prosthetic control. *Myoelectric Controls Symposium*, 12-–15 August, Fredericton, NB, Canada.

Salminger, S., Stino, H., Pichler, L. H., Gstoettner, C., Sturma, A., Mayer, J. A., Szivak, M., and Aszmann, O. C. (2022). Current rates of prosthetic usage in upper-limb amputees—have innovations had an impact on device acceptance? *Disability and Rehabilitation*, 44(14):3708–3713.

Sezener, E., Grabska-Barwińska, A., Kostadinov, D., Beau, M., Krishnagopal, S., Budden, D., Hutter, M., Veness, J., Botvinick, M., Clopath, C., Häusser, M., and Latham, P. E. (2021). A rapid and efficient learning rule for biological neural circuits. *bioRxiv* 2021.03.10.434756.

Simon, A. M., Turner, K. L., Miller, L. A., Potter, B. K., Beachler, M. D., Dumanian, G. A., Hargrove, L. J., and Kuiken, T. A. (2023). User performance with a transradial multi-articulating hand prosthesis during pattern recognition and direct control home use. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:271–281.