Towards Instant Calibration in Myoelectric Prosthetic Hands: A Highly Data-Efficient Controller Based on the Wasserstein Distance

Digby Chappell[1, 2, 3], Zeyu Yang[1], Honn Wee Son[2], Fernando Bello[3], Petar Kormushev[2], and Nicolas Rojas[1]

***Abstract*— Prosthetic hand control research typically focuses on developing increasingly complex controllers to achieve diminishing returns in pattern recognition of muscle activity signals, making models less suitable for user calibration. Some works have investigated transfer learning to alleviate this, but such approaches increase model size dramatically—thus reducing their suitability for implementation on real prostheses. In this work, we propose a novel, non-parametric controller that uses the Wasserstein distance to compare the distribution of incoming signals to those of a set of reference distributions, with the intended action classified as the closest distribution. This controller requires only a single capture of data per reference distribution, making calibration almost instantaneous. Preliminary experiments building a reference library show that, in theory, users are able to produce up to 9 distinguishable muscle activity patterns. However, in practice, variation when repeating actions reduces this. Controller accuracy results show that 10 non-disabled and 1 disabled participant were able to use the controller with a maximum of two recalibrations to perform 6 actions at an average accuracy of 89.9% and 86.7% respectively. Practical experiments show that the controller allows users to complete all tasks of the Jebsen-Taylor Hand Function Test, although the task of picking and placing small common objects required on average more time than the protocol's maximum time.**

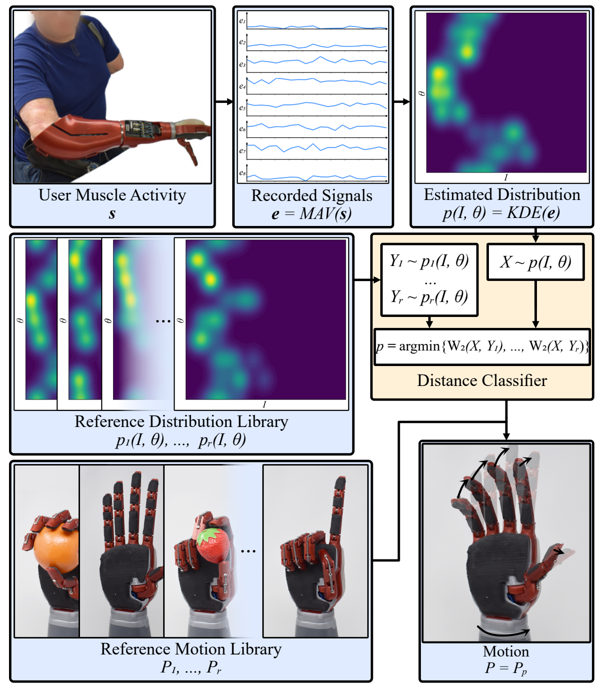

Fig. 1. Summary diagram of the controller presented in this work. EMG signals are collected from the forearm of the user via surface electrodes, then the underlying distribution producing the observed signals is estimated using kernel density estimation. The Wasserstein distance is computed between the estimated distribution and a set of reference distributions to determine the desired motion of the prosthetic hand.

I. Introduction

Myoelectric prosthetic hands are controlled by sensing the small electrical signals produced during muscle contraction in a process known as electromyography (EMG). In the vast majority of applications, these signals are measured at several points on the surface of the skin using a series of electrodes. After these signals are collected, a myoelectric hand controller must interpret the user's intention, and convert these to a desired movement or grasp for the prosthetic hand to perform [1]. The control of prosthetic hands using EMG has seen many developments in recent years, due to the rise of machine learning (ML) as an effective signal processing tool. However, ML-based controllers typically rely on a large amount of training data specific to the user, and a high number of trainable parameters, both of which are unsuitable for the embedded controllers used in prosthetic hands [2]. A further drawback of these is that calibration can be a difficult process in which new training data must be collected and used to update many model parameters. As such, there is progress to be made before next-generation prosthetic hand controllers are suitable for widespread use.

Controller performance and simple calibration are both important for new myoelectric controllers to become widely accepted by users [2], [3]. The vast majority of recent research has focused on the former; developing more and more complex models to attain diminishing returns in pattern recognition tasks [4], [5]. While improved performance is necessary to minimise the risk of prosthesis abandonment [6], it is clear that it must not come at the cost of ease of calibration.

In order to address the issue of calibration, some works have investigated transfer learning in order to transfer a complex controller, trained on a large dataset, to individual users [7]–[9]. However, the models used in these works are an order of magnitude larger than controllers trained

[1]REDS Lab, Dyson School of Design Engineering, Faculty of Engineering, Imperial College London, SW7 1AL, United Kingdom.
[2]Robot Intelligence Lab, Dyson School of Design Engineering, Faculty of Engineering, Imperial College London, SW7 1AL, United Kingdom.
[3]SiMMS Research Group, Department of Surgery and Cancer, Faculty of Medicine, Imperial College London, SW10 9NH, United Kingdom.
Contact: d.chappell19@imperial.ac.uk
Video: https://www.youtube.com/watch?v=AWtHQU4buZI

specifically to single users, as they must also capture subtle differences between users. Further to this, presented controllers are rarely tested on actual prosthetic hand users, who exhibit a high degree of heterogeneity in muscle activity post-amputation [10], and moreover are often not validated on physical tasks, instead opting to test classification accuracy with non-disabled participants [8], [9]. It therefore follows that there is a need for a prosthetic hand control algorithm that is able to classify multiple patterns of muscle activity at a similar level to modern controllers, but is also able to be calibrated quickly at minimal effort to the user.

This work proposes a novel myoelectric hand controller that aims to provide both functionality and ease of calibration without requiring a large training dataset or a high number of trainable parameters. The controller, summarised in Fig. 1, is based on first estimating the underlying distribution that produced a pattern of observed muscle activity, and then computing the Wasserstein distance between the estimated distribution and a library of reference distributions pre-recorded by the user. Preliminary experiments are performed investigating the ability of the controller to distinguish between common motions for both disabled and non-disabled users. Classification accuracy of the controller with 6 references is computed for non-disabled and disabled participants before the controller is validated on the Jebsen-Taylor Hand Function Test (JTHFT) [11].

We find that the proposed controller can theoretically distinguish between 9 common actions. However, in practice a lesser number is required in order to account for variability between repeated actions. In practical experiments, 6 reference distributions congruent to target prosthesis actions are recorded from each non-disabled participant. Meanwhile, the disabled participant used non-congruent actions in order to produce 6 distinguishable references while accounting for severe heterogeneity in muscular structure post-amputation. Classification accuracy of the controller, achieved with a maximum of two re-calibrations, is comparable to similar ML-based controllers. Finally, all participants were able to complete the Jebsen-Taylor Hand Function Test using the controller, validating its use in a real-world setting, allowing the identification of important factors that influence task completion time. A video summarising this work is available at `https://www.youtube.com/watch?v=AWtHQU4buZI`.

II. Methods

A. EMG Data Acquisition

To record EMG signals from the forearm of the user, a Thalmic Labs Myo armband is used. The armband samples surface EMG at a rate of 200 Hz at 8 equally spaced points around the forearm. Pre-processing is performed to rectify and generate the envelope of the incoming samples. For recorded signal $\boldsymbol{s}(k)$ at sample k, the mean absolute value $\boldsymbol{e}(k)$ over $W = 10$ samples is computed:

$$\boldsymbol{e}(k) = \frac{1}{W}\sum_{w=1}^{W} |\boldsymbol{s}(k + w - W)|. \tag{1}$$

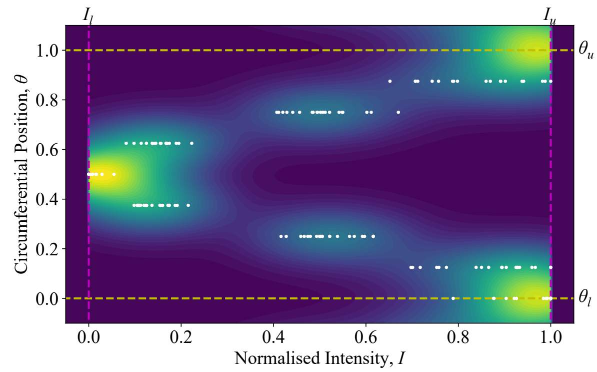

Fig. 2. An illustrative example of the generated kernel density estimate for 8 sensors equally spaced around the forearm sampled 20 times. I_l and I_u are the respective lower and upper bounds of measured intensity, beyond which the density estimate is truncated. θ_l and θ_u are the respective lower and upper bounds of circumferential position on forearm, beyond which the density estimate wraps.

B. EMG Distribution Estimation

EMG signals are stochastic, meaning there is some underlying probability distribution governing the likelihood of observing sampled EMG data. For muscle activity data sampled from the forearm, three main factors play a part in this: the position of the electrode on the forearm, the intensity of the signal, and the time at which the signal was generated. The position on the forearm is the two-dimensional position along the axis of—and around the circumference of—the forearm. Muscles in the forearm are generally aligned with its axis, meaning EMG signal varies by a relatively small amount along the bulk muscles of the forearm compared to circumferential position on the forearm, which is directly related to which muscles' activities are being detected. The intensity of the EMG signal is limited by the detection range possible with the chosen electrodes, truncating the signal between a minimum I_l and maximum I_u value. Finally, over a suitable time window, EMG signals can be assumed to be stationary [12], and are therefore independent of time. Using this information, the probability of observing a sample at a particular intensity, position on the forearm, and time, can be reduced to a two-dimensional probability of observing the sample at a particular intensity I, and circumferential position θ on the forearm. In this sense, the received signal exists on a cylinder that wraps around the forearm, capped between a minimum and maximum intensity by the sensing limits of electrodes.

Kernel density estimation (KDE) can be used to estimate the underlying probability distribution from a window of recorded samples. In the intensity dimension, given circumferential position, the kernel function used is a truncated Gaussian:

$$p(I|\theta,\, \mu,\, \sigma,\, I_{\mathrm{l}},\, I_{\mathrm{u}}) \propto \frac{\exp\left(-\frac{(I-\mu)^2}{2\sigma^2}\right)}{\Phi(\frac{I_{\mathrm{u}}-\mu}{\sigma}) - \Phi(\frac{I_{\mathrm{l}}-\mu}{\sigma})}, \tag{2}$$

where $\Phi(x) = \frac{1}{2}(1 - \mathrm{erf}(-\frac{x}{\sqrt{2}}))$. A truncated Guassian is used to account for minimum and maximum sensing limits. In the circumferential direction, the kernel function used is a von Mises distribution, also known as a circular Gaussian:

$$p(\theta|\nu,\,\kappa) \propto \exp\Big(\kappa\cos\left(2\pi(\theta-\nu)\right)\Big), \tag{3}$$

where we have scaled the circumferential position $\theta-\nu$ such that a full rotation around the forearm is equal to 1. The von Mises distribution is chosen in order to preserve the periodicity of the recorded signal around the circumference of the arm. The two dimensional kernel formed from the outer product of the individual one dimensional kernels is used for kernel density estimation. After the estimate is produced, it is normalised to produce a valid probability distribution:

$$p(I,\,\theta) = \text{KDE}(\mathbf{e}). \tag{4}$$

An example KDE produced from a window of 20 samples (0.1 s) from 8 sensors recording EMG equally spaced around the forearm is shown in Fig. 2. As seen, the KDE is truncated at the upper and lower limits of intensity, and circumferential position is periodic in the range $[0,\,1)$.

The kernel density estimate can be rapidly updated by maintaining a rolling dataset of 20 samples from which the estimate is calculated. For each set of d incoming samples, the components of the d oldest samples in the dataset are subtracted from the estimate, and the d new samples' components added. This is considerably faster than recalculating the entire estimate, which has a complexity of $\mathcal{O}(DM)$, where D is total size of the rolling dataset and M is the number of points at which the new estimate is computed. Updating only d samples has a complexity of $\mathcal{O}(2dM)$, which, given $d << D$, results in a large saving in computation.

C. Inter-Distribution Distance

After the distribution responsible for producing the observed EMG signal is computed, it must be compared against reference distributions of known actions in order to classify the desired action of the user. A suitable method to compare two distributions is to calculate the distance between them, known as the Wasserstein distance. The Wasserstein distance can be calculated by first drawing n samples from two distributions: $X = \{\boldsymbol{x}_1,\,\ldots,\,\boldsymbol{x}_n\} \sim p_X(I,\theta)$, and $Y = \{\boldsymbol{y}_1,\,\ldots,\,\boldsymbol{y}_n\} \sim p_Y(I,\theta)$. Since each sample is comprised of circumferential position and intensity, $\boldsymbol{x}_i = [I_i^{(x)},\,\theta_i^{(x)}]^T$ and $\boldsymbol{y}_j = [I_j^{(y)},\,\theta_j^{(y)}]^T$, the distance $d_{i,\,j}$ between them is defined as the geodesic on the cylindrical space introduced previously:

$$d_{i,\,j} = \sqrt{(I_i^{(x)} - I_j^{(y)})^2 + a(\theta_i^{(x)},\,\theta_j^{(y)})^2}, \tag{5}$$

where $a(\theta_i^{(x)},\,\theta_j^{(y)})$ is the minimum unit circumferential distance between points $\theta_i^{(x)}$ and $\theta_j^{(y)}$, defined as $a(\theta_i^{(x)},\,\theta_j^{(y)}) = ((\theta_i^{(x)} - \theta_j^{(y)} + 0.5) \bmod 1) - 0.5$.

As detailed in [13], the Wasserstein distance between X and Y aims to find a flow $\boldsymbol{F} = [f_{i,\,j}]$, between each sample $\boldsymbol{x}_i$ and $\boldsymbol{y}_j$. When n samples are drawn from each distribution, the flow is that which minimises the following optimal transport cost:

$$\text{WORK}(X, Y, \boldsymbol{F}) = \sum_{i=1}^{n}\sum_{j=1}^{n} d_{i,\,j} f_{i,\,j}, \tag{6}$$

subject to:

$$f_{i,\,j} \geq 0 \qquad 1 \leq i \leq n, \quad 1 \leq j \leq n \tag{7}$$

$$\sum_{j=1}^{n} f_{i,\,j} \leq 1 \qquad 1 \leq i \leq n \tag{8}$$

$$\sum_{i=1}^{n} f_{i,\,j} \leq 1 \qquad 1 \leq j \leq n \tag{9}$$

$$\sum_{i=1}^{n}\sum_{j=1}^{n} f_{i,\,j} = n. \tag{10}$$

When the optimal flow has been found by solving the optimal transport problem, the Wasserstein distance is the normalised optimal work:

$$W_2(X,\,Y) = \frac{\sum_{i=1}^{n}\sum_{j=1}^{n} d_{i,\,j} f_{i,\,j}}{\sum_{i=1}^{n}\sum_{j=1}^{n} f_{i,\,j}}. \tag{11}$$

For a multi-dimensional distribution, the number of calculations performed to solve the optimal transport problem rapidly grows with number of samples beyond computational feasibility. To reduce computational complexity, Sinkhorn divergences can be used to significantly improve calculation speed — this work is not intended to cover this approximation, and [14] provides a good overview of the practical methods used to calculate the Wasserstein distance.

In this work, the Wasserstein distance is calculated using $n = 250$ sampled points from each density estimate using the Python package 'GeomLoss' [14]. This value is chosen to ensure a representative sample is drawn from each distribution such that the calculated Wasserstein distance is consistent, while balancing speed of computation.

D. Prosthetic Hand Control

Once the Wasserstein distance between the current distribution and a set of r reference distributions has been computed, the desired action of the user is classified by selecting the action with the index p that corresponds to the closest reference distribution:

$$p = \text{argmin}\{W_2(X, Y_1),\,\cdots,\,W_2(X, Y_r)\}. \tag{12}$$

To improve robustness, the action is only accepted if its distance falls below a threshold (0.012 herein).

E. Experimental Setup

In this work, the OLYMPIC hand [15] is used as a test platform, with the addition of a modular powered forearm that is able to pronate and supinate the hand (shown in Fig. 3). The OLYMPIC hand is equipped with a small computer that is capable of utilising GPU acceleration [16], allowing full use of the software libraries used to compute the Wasserstein distances between references.

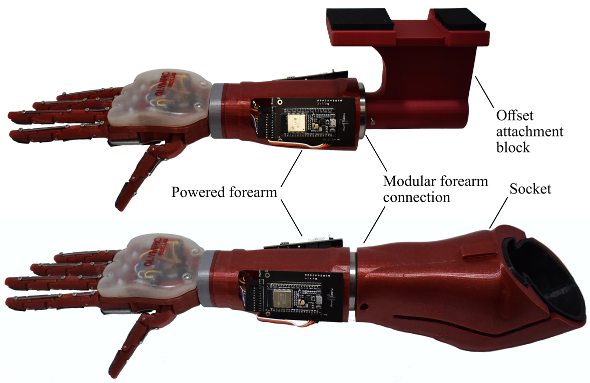

Fig. 3. The OLYMPIC hand with modular powered forearm. **Top:** forearm connected to an offset attachment block for non-disabled participants. **Bottom:** forearm connected to a custom socket for the disabled participant.

10 non-disabled participants and 1 (body powered) prosthetic hand user were recruited to take part in this study. A pilot study was conducted in which participants were asked to record reference distributions corresponding to the six prehensile grips identified in the Southampton Hand Assessment Procedure [17]: *lateral*, *power*, *tripod*, *tip*, *extension*, *spherical*, as well as 5 common actions used in prosthetic hand control: *wrist flexion* (close hand), *wrist extension* (open hand), *rest*, *pronate hand*, *supinate hand*. To build a reference library, reference distributions were iteratively appended to the library in order of minimum distance to the existing library, after starting with the reference distribution with the greatest average distance from all other references. As seen in Fig. 4, as more references are added to the library, the minimum distance between them decreases, until it falls below the minimum detection threshold. The detection threshold is calculated based on the mean distance between two identical distributions (which is non-zero due to random sampling when computing the Wasserstein distance) plus two standard deviations. If the distance between two references is below this threshold it is not possible to discriminate between them with statistical confidence. The minimum distance between reference distributions for the disabled participant was found to be consistently lower than for non-disabled participants. After normalising according to the mean and standard deviation of each set of results for non-disabled participants, an uneven t-test found that this difference is statistically significant ($p < 0.01$). This highlights that references based on common actions may not always be suitable for control by actual prosthetic hand users, and references should be carefully chosen according to each individual user.

In practical experiments used in this work, 6 references are used in order for recorded references to be far above the minimum detection threshold. This is because a user is rarely able to perfectly recreate a pattern of muscle activity, and physiological factors such as fatigue and quality of the electrode-skin contact can degrade recorded signals over time; a myoelectric controller should be robust to this. The 6 reference actions used in this work are: *1: rest*, *2: wrist flexion*, *3: wrist extension*, *4: tripod*, *5: supinate hand*, and *6: pronate hand*. Note that these references do not necessarily correspond to the optimal 6 references identified from the 11 actions trialled in the pilot study, rather these are the references that correspond to the desired motions of the prosthesis. For the disabled participant, it was found that actions congruent to the desired motion of the prosthesis were not distinguishable by the controller, corroborating the results of the pilot study. To handle this, the controller was calibrated with the following actions, which were found to produce distinct patterns of muscle activity; *rest*, *pronate hand*, *supinate hand*, *power grip*, *thumb opposition*, and *pinky finger flexion*.

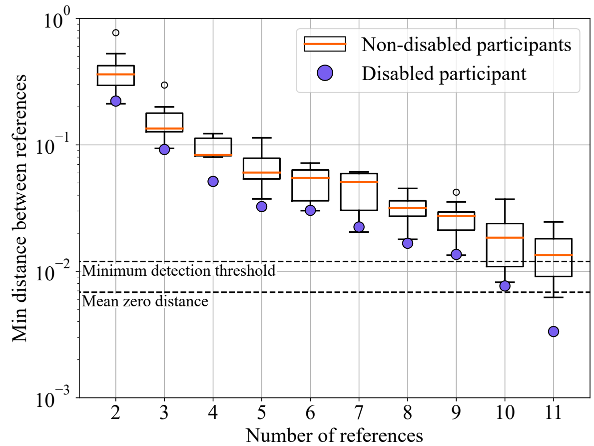

Fig. 4. The minimum distance between reference distributions against number of reference distributions. References are added in descending order of minimum distance to references already in the library, starting with the reference with the greatest average distance to other references. Minimum detection threshold and mean zero distance are shown as dashed lines. Note that mean zero distance is non-zero due to random sampling.

After initially recording the 6 reference actions to estimate their starting distributions, participants were asked to perform 10 repetitions of each action within a 5 second window in a dictated random order, during which time the classification rate of the controller was monitored. Participants were allowed a maximum of three attempts at this task, after which the performance of their final attempt was recorded. During each attempt no feedback was available to the participant, only the notification of the next action to perform. After a complete attempt, participants were able to view their performance in order to make a decision whether to recalibrate the controller and repeat the task. Participants were then given 1 hour to practice controlling the hand, during which time all items used in subsequent dexterity tests were available for practice. After the practice period, participants completed the Jebsen-Taylor Hand Function Test [11], chosen because it requires a mixture of fine and coarse manipulation skills at varying wrist pronation and supination. Participants were able to recalibrate the controller at any point during the training period and during the JTHFT.

The use of non-disabled participants and disabled participants in this work was granted ethical approval by the Imperial College London Research Governance and Integrity Team (RGIT), Science Engineering and Technology Research Ethics Committee (SETREC) number 21IC6716 and

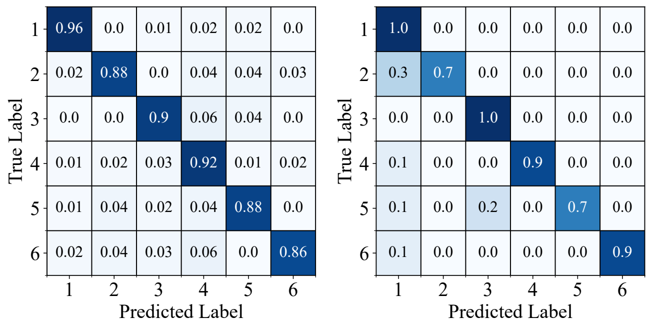

Fig. 5. **Left:** Average confusion matrix across 10 non-disabled participants normalised across each true label (1: *rest*, 2: *wrist flexion*, 3: *wrist extension*, 4: *tripod*, 5: *supinate hand*, and 6: *pronate hand*). **Right:** Confusion matrix for disabled participant also normalised (1: *rest*, 2: *supinate hand*, 3: *pronate hand*, 4: *power grip*, 5: *thumb opposition*, and 6: *pinky finger flexion*).

Imperial College Research Ethics Committee (ICREC) number 21IC7227, respectively. Written consent was obtained from all participants prior to taking part in the study.

III. Results & Discussion

To measure the ability of the controller to classify user muscle activity, 10 non-disabled users were asked to perform 10 repetitions of the 6 reference actions. Given a maximum of only two calibrations, users were able to achieve an average accuracy of 89.9%, comparable with existing deep learning-based controllers [8], [16]. The average confusion matrix corresponding to this task is shown in Fig. 5 Left. As seen, reference 4, tripod grip, elicits the most type I (false positive) errors, which are most prevalent when users attempt to perform reference 3, wrist extension, and 6, pronate hand. This is likely due to variance when users repeatedly perform actions; users may not perform actions identically every time, meaning a reference distribution generated from a single sample of data may not fully represent that action. It was observed that there may be a relationship between the distance between reference distributions and classification rate; for example, if two references are close together then it is more likely that the controller will misclassify these two references. However, this is subject to further investigation. Fig. 5 Right shows the confusion matrix produced by the disabled participant, who achieved an average accuracy of 86.7%. In this case, reference 1, rest, elicits the most false positive errors. This could be due to weaker muscles present in the disabled user's residual limb, which are not used on a regular basis when operating a body-powered prosthesis.

The results of the Jebsen-Taylor Hand Function Test can be seen in Fig. 6. All participants were able to complete each task, although picking and placing small objects (J3) required on average more time than the JTHFT protocol's 120 s maximum time limit. In general, it can be observed that task completion requires greater time using the OLYMPIC hand than the disabled participant's current body powered prosthetic hand. This time increase can be resolved into four primary sources; user expertise, classifier latency, motion latency, and mechanical features.

User expertise refers to the level of practice a user has experienced with each prosthesis. The disabled participant currently uses a body powered prosthesis for day to day tasks, so can be considered an expert user. With the OLYMPIC hand, all users can be considered novice users, because they had only been given 1 hour to practice. The effect of user expertise can be seen in the writing (J1) and feeding (J4) tasks; the disabled participant outperformed non-disabled participants when using the OLYMPIC hand. These tasks were accomplished with a static grip in the same way that a body-powered prosthesis grasps an object.

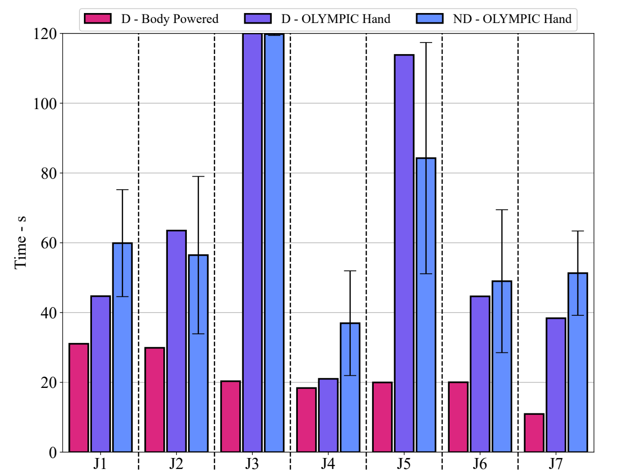

Fig. 6. Results of the Jebsen-Taylor Hand Function Test for 10 non-disabled participants and 1 disabled participant. D: disabled participant, ND: Non-disabled participant. J1: Writing, J2: Simulated page turning, J3: Small objects, J4: Simulated feeding, J5: Stacking checkers, J6: Large, light objects, J7: Large, heavy objects.

Classifier latency is the time period between a user attempting to produce muscle activity required to perform a desired motion and the point in time at which the classifier successfully classifies the muscle activity produced. In the controller presented in this work, references are recorded of muscles at steady state, so users must match this. Identifying exactly when a user begins to attempt a motion is difficult, but from observations made during these experiments can range from 0.25 s to 2 s depending on the participant and the motion which they are attempting. Classifier latency dominates tasks involving large objects (J6 and J7), and is not present in body-powered prostheses.

Motion latency is the time period between muscle activity being successfully classified and the prosthetic hand completing the desired motion. This depends on the motion and any object the hand interacts with. Since the controller used in this work is discrete, i.e., it moves between fixed discrete poses, motion latency is relatively high. Motion latency makes up a significant proportion of the time increase observed between the body powered prosthesis and the myoelectric prosthesis, particularly when grasping small objects (J2, J3, and J5). The body powered prosthesis is able to remain in a semi-closed state making grasping much quicker than fully opening and closing the OLYMPIC hand.

Certain mechanical features of the OLYMPIC hand lead to observed time differences in some tasks. The tasks which exhibit the largest increase in time taken to complete are picking and placing small objects (J3) and stacking checkers (J5). As described in [18], each finger of the OLYMPIC hand is underactuated. The equilibrium between the extension springs mounted on the dorsal side and the flexion tendon that runs within them governing the path taken during closing if motion is unrestricted. If motion is hindered in any way, this equilibrium is disrupted, and the finger will deviate from its path. It should be noted that this makes grasping small objects challenging in practice, since all fingers used in the grasp must move unhindered to the grasping posture for a repeatable, precise grip. So, although users were able to perform the correct patterns of muscle activity required to close the hand in tripod grip and open the hand again, many more attempts were required to successfully grasp each object. A time difference on the writing (J1) and simulated feeding (J4) tasks is observed between the results of the disabled participant using their body powered prosthesis compared to results of both the disabled and non-disabled participants performing this task with the OLYMPIC hand. This increase is due to the OLYMPIC hand having soft pads located on the palmar side of the hand, which means that a grasped object is able to move slightly, hindering static tasks.

IV. Conclusions & Future Work

In this work, we have presented a new prosthetic hand controller that uses minimal data. By using a non-parametric model to estimate the underlying distribution of observed EMG signals, and computing the Wasserstein distance between the observed distribution and a library of reference distributions, real-time control of a prosthesis was achieved. Near-instant calibration is achieved by rerecording the library of reference distributions; with each reference distribution generated from only 0.1 s of user data, full calibration of multiple references can be achieved in under 1 minute.

Results of the pilot study showed that participants could theoretically produce up to 9 distinguishable references based on common grasps and actions, however references were significantly closer for a disabled participant. In practical experiments, 10 non-disabled participants were able to achieve an average accuracy of 89.9% on 6 references congruent to desired actions of a prosthetic hand, while a disabled participant was able to achieve an average accuracy of 86.7% on 6 non-congruent references. Participants were able to complete all tasks of the JTHFT, although results show that task completion times were slower than the disabled participant using their own body-powered prosthesis, with the majority of participants unable to complete the task of picking and placing small common objects within the maximum time limit. This difference in completion time can be resolved into a combination of user expertise, controller latency, motion latency, and mechanical features of the prosthesis, and future work will investigate reducing these.

Acknowledgement

This work was supported by the UKRI CDT in AI for Healthcare `http://ai4health.io` (Grant No. EP/S023283/1). The authors would like to thank the Alex Lewis Trust and Koalaa for their support.

References

[1] M. B. I. Reaz, M. S. Hussain, and F. Mohd-Yasin, "Techniques of EMG signal analysis: detection, processing, classification and applications," *Biological procedures online*, vol. 8, no. 1, pp. 11–35, 2006.

[2] D. Farina, N. Jiang, H. Rehbaum, A. Holobar, B. Graimann, H. Dietl, and O. C. Aszmann, "The Extraction of Neural Information from the Surface EMG for the Control of Upper-Limb Prostheses: Emerging Avenues and Challenges," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 4, pp. 797–809, 2014.

[3] L. C. Smail, C. Neal, C. Wilkins, and T. L. Packham, "Comfort and function remain key factors in upper limb prosthetic abandonment: findings of a scoping review," *Disability and Rehabilitation: Assistive Technology*, vol. 16, no. 8, pp. 821–830, 2020.

[4] P. Tsinganos, B. Cornelis, J. Cornelis, B. Jansen, and A. Skodras, "Deep Learning in EMG-based Gesture Recognition," in *5th International Conference on Physiological Computing Systems*, pp. 107–114, 6 2018.

[5] M. Jafarzadeh, D. C. Hussey, and Y. Tadesse, "Deep learning approach to control of prosthetic hands with electromyography signals," in *2019 IEEE International Symposium on Measurement and Control in Robotics (ISMCR)*, pp. 1–4, 9 2019.

[6] D. K. Kumar, B. Jelfs, X. Sui, and S. P. Arjunan, "Prosthetic hand control: A multidisciplinary review to identify strengths, shortcomings, and the future," *Biomedical Signal Processing and Control*, vol. 53, p. 101588, 2019.

[7] F. Demir, V. Bajaj, M. Ince, S. Taran, and A. Sengur, "Surface EMG signals and deep transfer learning-based physical action classification," *Neural Computing and Applications*, vol. 31, 6 2019.

[8] U. Côté-Allard, C. L. Fall, A. Drouin, A. Campeau-Lecours, C. Gosselin, K. Glette, F. Laviolette, and B. Gosselin, "Deep Learning for Electromyographic Hand Gesture Signal Classification Using Transfer Learning," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 27, no. 4, pp. 760–771, 2019.

[9] E. Rahimian, S. Zabihi, A. Asif, D. Farina, S. F. Atashzar, and A. Mohammadi, "FS-HGR: Few-shot Learning for Hand Gesture Recognition via ElectroMyography," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, pp. 1–1, 2021.

[10] A. Waris, I. K. Niazi, M. Jamil, O. Gilani, K. Englehart, W. Jensen, M. Shafique, and E. N. Kamavuako, "The effect of time on EMG classification of hand motions in able-bodied and transradial amputees," *Journal of Electromyography and Kinesiology*, vol. 40, 2018.

[11] R. H. Jebsen, N. Taylor, R. B. Trieschmann, M. J. Trotter, and L. A. Howard, "An objective and standardized test of hand function," *Arch. Phys. Med. Rehabil.*, vol. 50, no. 6, pp. 311–319, 1969.

[12] N. Nazmi, M. A. A. Rahman, S.-i. Yamamoto, S. A. Ahmad, M. Malarvili, S. A. Mazlan, and H. Zamzuri, "Assessment on Stationarity of EMG Signals with Different Windows Size During Isotonic Contractions," *Applied Sciences 2017, Vol. 7, Page 1050*, vol. 7, p. 1050, 10 2017.

[13] Y. Rubner, C. Tomasi, and L. J. Guibas, "The earth mover's distance as a metric for image retrieval," *International journal of computer vision*, vol. 40, no. 2, pp. 99–121, 2000.

[14] J. Feydy, T. Séjourné, F.-X. Vialard, S.-I. Amari, A. Trouvé, and G. Peyré, "Interpolating between Optimal Transport and MMD using Sinkhorn Divergences," in *AISTATS*, pp. 2681–2690, 2019.

[15] L. Liow, A. B. Clark, and N. Rojas, "OLYMPIC: A modular, tendon-driven prosthetic hand with novel finger and wrist coupling mechanisms," *IEEE Robotics and Automation Letters*, vol. 5, pp. 299–306, 4 2020.

[16] Z. Yang, A. B. Clark, D. Chappell, and N. Rojas, "Instinctive Real-time sEMG-Based Control of Prosthetic Hand with Reduced Data Acquisition and Embedded Deep Learning Training," in *IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2022.

[17] C. M. Light, P. H. Chappell, and P. J. Kyberd, "Establishing a standardized clinical assessment tool of pathologic and prosthetic hand function: Normative data, reliability, and validity," *Archives of Physical Medicine and Rehabilitation*, vol. 83, pp. 776–783, 6 2002.

[18] D. Chappell, H. W. Son, A. B. Clark, Z. Yang, F. Bello, P. Kormushev, and N. Rojas, "Virtual Reality Pre-Prosthetic Hand Training with Physics Simulation and Robotic Force Interaction," *IEEE Robotics and Automation Letters*, 2022.