

In a typical BMI (Shanechi, 2017) the decoder is initialized by a supervised calibration procedure, which creates a map based on labelled examples of desired actions and corresponding neural activities. However, a good calibration is not sufficient to achieve proficient BMI control, as the performance of the decoder often degrades and fluctuates when evaluated online, due to imperfect predictions, recording instabilities, changes in neuronal properties, attentional changes, and changes brought about by learning (Barrese et al., 2013; Downey et al., 2018). BoMI decoders suffer from similar issues. The decoder of a BoMI, also referred to as “forward map”, is initialized by unsupervised identification of the low-dimensional, latent manifold of unconstrained users’ movements recorded during an initial calibration (Casadio et al., 2010). This latent manifold is expected to change with subsequent practice. Therefore, a discrepancy is likely to develop in time between the evolving latent manifold of the user and the initial BoMI forward map. Indeed, it has been observed that extensive practice with a BoMI led to the consolidation of task-specific movement strategies (Pierella et al., 2017) and several studies in brain machine interfaces demonstrated the existence of a stable manifold of neural activity linked to BMI use and interpreted this as the result of neural adaptation following extensive practice with the interface (Gallego et al., 2020; Ganguly & Carmena, 2009; Oweiss & Badreldin, 2015; Shenoy & Carmena, 2014).

The consequence of adopting a fixed activity-intention map after its initial offline tuning, is that the user is left with the burden of learning how to use the interface for achieving new goals within new operating conditions that the decoder has not been optimized on.

Closed-loop supervised decoder adaptation in BMIs has been proposed as an effective way to increase decoder performance during use. This was the case when allowing the decoder parameters to smoothly change according to the inferred movement goal during within-session interface operation (Dangi et al., 2013; Orsborn et al., 2014). A very recent study considered instead the case of performance loss arising from decoder instabilities across sessions. They showed that by calibrating the decoder using features within the manifold of stable activity of the recorded neurons, performance of the BMI can be reliably recovered by manifold alignment across sessions (Degenhart et al., 2020). Despite contributing encouraging results, these current perspectives rely either on the knowledge of user intent or on the presence of an established manifold of neural activity linked to the use of the BMI. Hence, they cannot be easily extended to facilitate the use of the interface when the movement goal is unknown or when the formation of new neural strategies is still ongoing and the activity manifold has not yet consolidated (Oby et al., 2019).

Given the current limitations of closed-loop decoder adaptation, we propose a procedure for facilitating interface operation that does not rely on estimates of user intent and can be applied from the very initial stages of learning. The procedure is initially developed for application with body machine interfaces and exploits a non-linear autoencoder (AE) network (Kramer, 1991) trained iteratively to identify and track the evolution of the latent manifold of its inputs.

Previous work from our group suggested that individuals training with a linear BoMI that was adapted iteratively based on movement statistics increased movement efficiency compared to a fixed interface and were able to develop a more faithful internal representation of the BoMI forward map (De Santis et al., 2018;

explore their range of motion. This procedure is not task-related, as users do not receive visual feedback of their movements and are not yet connected to the device. Data collected during calibration are then used as training set for a dimensionality reduction (DR) algorithm that derives f by extracting the low-dimensional, m -D, manifold in which the highest amount of variance of body signals is available. Common iterations of BoMIs exploit linear DR methods to extract such low-dimensional manifold, e.g. Principal Component Analysis (PCA) (Wold et al., 1987), or Kalman Filter (Seáñez González et al., 2016). In this study, we explore the use of a non-linear DR algorithm, autoencoder networks.

2.2. Fixed autoencoder

AEs are unsupervised artificial neural networks capable of learning efficient lower dimensional representations of the input data without labelled data. An AE is a cascade of two components: an encoder E that converts the inputs to a lower dimensional




Fig. 3. Setup for the reaching task and training protocol. The participant was sitting in front of a computer monitor and was controlling a cursor using signals generated by IMUs (red boxes). Training (grey) and test (blue) targets were uniformly distributed on a circle. Four training targets were placed in four directions (45° + $k90^\circ$), while eight test targets were placed in eight directions (22.5° + $k45^\circ$). Each target was placed at the same distance L of 10.5 cm from HOME target.

all 4 targets had been reached. In the last repetition of the 4 targets visual feedback of the cursor was removed, and the participants were asked to stop moving when they believed to be in the target (blind trials). The goal of these blind trials was to establish if the participants were guided by error feedback or if, instead, they formed a feedforward command based on an internal representation of the cursor space.

During the first training epoch, the encoder for the adaptive group (A) was initialized with $E^{0/n}$ and kept constant for 60 s (baseline). After that, the map was iteratively updated as described in Table 1. Update of the map was suspended during blind trials.

3.4.4. Test phase

The participants practised a centre-out reaching task to eight target locations uniformly distributed on a circle (Fig. 3, blue targets). After each successful reaching, they were asked to move the cursor back to the central HOME target (Fig. 3, green target). In each test the eight targets were presented once.

A total of three tests were presented: an initial test as baseline before starting the experiment, a midway test after four epochs of training, and a final test at the end of the experiment. For the whole duration of each test, the control group (F) practised with $E^{0/n}$, while the map update was suspended for the adaptive group (A). Hence both groups practised with the same map during the initial test epoch.


3.5. Outcome measures

3.5.1. Sensitivity tuning

To evaluate the convergence and the stability of the a-AE in the offline sensitivity tuning we computed three metrics.

First, we quantified the tracking performance of the a-AE based on the **reconstruction error (RE)**. RE is defined as the loss function described in Eq. (2), that was minimized every two seconds during the online retraining of the AE. A reduction of the

Reaching Task



reconstruction error implies that the mapping was providing a stable representation of the AE training data set.

To quantify the stability of the interface when using the a-AE, we considered the actual cursor vector $p^{n/n}$, and the cursor vector, $\hat{p}^{n/n}$, that would have been obtained without updating the encoder:

$$\begin{aligned} \hat{p}^{n/n} &\in E^{(n-1)}(q^{n/n}) \\ p^{n/n} &\in E^{n/n}(q^{n/n}) \end{aligned} \quad (7)$$

where $E^{(n-1)}$ and $E^{n/n}$ are the two time-consecutive encoders and $q^{n/n}$ is the movement set used in the training of the AE at the current update iteration. Note that, even though the a-AE at consecutive update iterations ($E^{(n-1)}$ and $E^{n/n}$) was trained with different movement set q , here we computed the cursor trajectory using the same set $q^{n/n}$. This allowed us to evaluate the **static jump**, defined as the L_2 norm between the final k value of $\hat{p}^{n/n}$ (end of an update iteration) and the initial value of $p^{n/n}$.

computed the number of trials completed during baseline and then during time intervals of two minutes, including the final two minutes of training. Blind trials were not considered in this metric.

Figure 5. Performance metrics for adaptive (red) and fixed (black) group. Panel A: Number of trials completed during baseline, after two, four, six minutes following baseline and during the final two minutes of training. The asterisk represents a significant difference between groups during the last 2 min of training. Panel B: Endpoint error during blind trials of each training epoch. Panel C-D: Linearity index and movement smoothness values during each training and test epoch. Mean values across participants are plotted with 95% confidence interval.




Figure 6. A: Bi-dimensionality index values for adaptive (red) and static (black) group during each training epoch. Mean values across participants are plotted with 95% confidence interval. Panel B: Variance accounted for (VAF) values for adaptive (red) and static (black) group during each batch of the training duration. Mean across subjects is shown as a bold line for both groups.




Fig. 7. Similarity matrix representing the similarity between final encoders among participants. The higher the value of a cell, the more similar the final encoders of the two participants of that cell. The matrix has been ranked to cluster participants with similar final encoders.



Fig. 8. Cursor trajectories of one participant (S8, black lines) during the final test epoch. We took the IMU data of S8 recorded during the final test epoch and applied the final encoder of a participant with whom S8 had a high similarity (S4, red lines, Panel A) and a low similarity (S7, red lines, Panel B).

substitute the final AE map that a participant used in the late test phase with that of another participant whose map had a comparable structure, the same participant would have been able to cover the target space much more consistently (Fig. 8A) than with a map whose similarity was not as pronounced (Fig. 8B).

5. Discussion

In this study we proposed an adaptive platform based on the use of an iterative non-linear autoencoder to implement unsupervised tracking of user's manifold for improving the ease-of-use of a human machine interface. First and foremost, our results support the use of a non-linear AE as a proficient control map within the body machine interface scheme. Moreover, the adaptive approach led to an increased representational efficiency of the interface decoder while concurrently increasing users' task-related performance, both in terms of number of trials completed over time and accuracy during reaching of blind trials. This result suggests that the online co-adaptation process encourages the development of a more accurate internal model. Importantly, the proposed approach has three salient features that makes it appealing in many applications other than the one tested here: (i) it cancels the cost of interrupting the operation of the device to perform decoder recalibration, (ii) as no information about the state of the task and/or intended task goals is needed, the

manifold tracking algorithm can be applied to a great variety of contexts, iii) it does not rely on the existence of a stable neural or movement manifold to compensate for decoder instabilities, allowing it to be applied in the earliest stages of interface operation, when the formation of new neural strategies is still on-going.

5.1. Autoencoder networks can proficiently control low-dimensional devices

There is an increasing enthusiasm about using autoencoder networks in the field of human machine interfaces (HMIs). AEs give the freedom of choosing the type of network architecture and level of complexity (e.g., linear/non-linear, generative/convolutional/recurrent, number of hidden layers and neurons per layer), thus potentially allowing to better match the degree of complexity of the input signal and hence allow for a great variety of applications. For instance, AE-based approaches have recently been employed for adversarial and variational domain adaptation (Farshchian et al., 2018; Hsu et al., 2017) and for extracting precise estimates of neural dynamics (Pandarinath et al., 2018). However, applicability of non-linear AEs for the control of low dimensional devices have seen limited efforts.

The use of AEs, or more specifically of their encoder sub-network, as a forward map in a BoMI has been first proposed by 2018

5.4. Map adaptation is guided by the individual's learning trajectory

If on the one hand the a-AE was successful in tracking the movement manifold of each participant, on the other we noted that the changes in the encoder were not consistent across participants. Namely, we could divide the BoMI users into two main groups: those who converged towards a similar encoder structure at the end of the training (S1, S2, S4, S6, S8, S10 - Fig. 7) and those whose maps did not converge to a particular similarity value (S3, S5, S7, S9 - Fig. 7). All the participants were able to reach a satisfactory level of performance regardless of whether the map was converging towards a particular structure. Furthermore, as mentioned before, in addition to being adaptive, the redundant character of the interface allowed each participant to complete the task with a different strategy. We speculate that the final encoder was merely a result of the learning trajectory of each participant. Indeed, if two participants were to exhibit the same movement dynamics (*i.e.*, strategy) from the start to the end of the practice, it is reasonable to assume that they would have converged towards the same encoder. Since this was not always the case, we concluded that some participants effectively learned the task in a different way. It is important to remark that the a-AE was able to guide them to an efficient resolution of the task regardless of the inverse model employed by the participant. This is a remarkable characteristic that increases the generalization capability of the proposed interface.

5.5. Perspective on current adaptive interfaces

The problem of building adaptive interfaces is raising increasing interest in the realm of human machine interactions. Here we focused on the development of an adaptive interface that provides a seamless interaction with the user during online operation of the interface. In the field of BMIs, the closed-loop adaptation of the interface is driven by some policies associated with the user's known movement intention (Vidaurre et al., 2010). This results in supervised adaptation as, for example, it would require the user to perform pre-selected movements to guide the update of the interface parameters (Dangi et al., 2014;

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This work was supported by the Marie Curie Integration [Grant FP7- PEOPLE-2012-CIG-334201], the Ministry of Science and Technology, Israel (Joint Israel Italy lab in Biorobotics Artificial somatosensorial for humans and humanoids), the National Science Foundation [Grant 1632259], the NIDILRR [Grant 90REG0005-01], the NICHD [Grant 5R01HD072080], NIBIB, USA (Grant No. R01 EB024058-03) and the European Union Horizon 2020 research and innovation program under the Marie Skłodowska-Curie, project REBoT, [G.A. No 750464].

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., ... Zheng, X. (2016). TensorFlow: A system for large-scale machine learning. In *12th USENIX symposium on operating systems design and implemen-*

(2716).350(Body-1(machine35147(intfacine3514enablesne3514peopline3514wi12th3514ce(lr-sc.))TJO 0 0 rg 0 0 0 RGO 0 0 rg 0 0 0 RG 0 -8.568 Td sp

- Miehlbradt, J., Cherpillod, A., Mintchev, S., Coscia, M., Artoni, F., Floreano, D., & Micera, S. (2018). Data-driven body-machine interface for the accurate control of drones. *Proceedings of the National Academy of Sciences of the United States of America*, 115(31), 7913–7918. <http://dx.doi.org/10.1073/pnas.1718648115>.
- Morcos, A. S., Raghu, M., & Bengio, S. (2018). Insights on representational similarity in neural networks with canonical correlation. In *Advances in neural information processing systems* (pp. 5727–5736).
- Mosier, K. M., Scheidt, R. A., Acosta, S., & Mussa-Ivaldi, F. A. (2005). Remapping hand movements in a novel geometrical environment. *Journal of Neurophysiology*, 94(6), 4362–4372. <http://dx.doi.org/10.1152/jn.00380.2005>.
- Müller, J. S., Vidaurre, C., Schreuder, M., Meinecke, F. C., Von Büna, P., & Müller, K.-R. (2017). A mathematical model for the two-learners problem. *Journal of Neural Engineering*, 14(3), 36005. <http://dx.doi.org/10.1088/1741-2552/aa620b>.
- Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., Chase, S. M., & Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences of the United States of America*, 116(30), 15210–15215. <http://dx.doi.org/10.1073/pnas.1820296116>.
- Orban De Xivry, J. J., & Lefèvre, P. (2015). Formation of model-free motor memories during motor adaptation depends on perturbation schedule. *Journal of Neurophysiology*, 113(7), 2733–2741. <http://dx.doi.org/10.1152/jn.00673.2014>.
- Orsborn, A. L., Dang, S., Moorman, H. G., & Carmena, J. M. (2011). Exploring time-scales of closed-loop decoder adaptation in brain-machine interfaces. In *2011 annual international conference of the IEEE engineering in medicine and biology society* (pp. 5436–5439).
- Orsborn, A. L., Dang, S., Moorman, H. G., Carmena, J. M., & Member, S. (2012). Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(4), 468–477.
- Orsborn, A. L., Moorman, H. G., Overduin, S. A., Shanechi, M. M., Dimitrov, D. F., & Carmena, J. M. (2014). Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron*, 82(6), 1380–1393. <http://dx.doi.org/10.1016/j.neuron.2014.04.048>.
- Oweiss, K. G., & Badreldin, I. S. (2015). Neuroplasticity subserving the operation of brain-machine interfaces. *Neurobiology of Disease*, 83, 161–171. <http://dx.doi.org/10.1016/j.nbd.2015.05.001>.
- Pandarinath, C., O'Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., Henderson, J. M., Shenoy, K. V., Abbott, L. F., & Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10), 805–815. <http://dx.doi.org/10.1038/s41592-018-0109-9>.
- Pierella, C., Abdollahi, F., Thorp, E., Farshchiansadegh, A., Pedersen, J., Seáñez González, I., Mussa-Ivaldi, F. A., & Casadio, M. (2017). Learning new movements after paralysis: Results from a home-based study. *Scientific Reports*, 7(1), 1–11. <http://dx.doi.org/10.1038/s41598-017-04930-z>.
- Pierella, C., Sciacchitano, A., Farshchiansadegh, A., Casadio, M., & Mussa-Ivaldi, S. A. (2018). Linear vs non-linear mapping in a body machine interface based on electromyographic signals. In *Proceedings of the IEEE RAS and EMBS international conference on biomedical robotics and biomechanics* (pp. 162–166). <http://dx.doi.org/10.1109/BIROB.2018.8487185>.
- Portnova-Fahreva, A. A., Rizzoglio, F., Nisky, I., & Casadio, M. (2020). Linear and non-linear techniques on full hand kinematics. *Frontiers in Bioengineering and Biotechnology*, 8(May), 1–18. <http://dx.doi.org/10.3389/fbioe.2020.00429>.
- Raghu, M., Gilmer, J., Yosinski, J., & Sohl-Dickstein, J. (2017). SVCCA: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *Advances in neural information processing systems* (pp. 6077–6086).
- Ranganathan, R., Wieser, J., Mosier, K. M., Mussa-Ivaldi, F. A., & Scheidt, R. A. (2014). Learning redundant motor tasks with and without overlapping dimensions: facilitation and interference effects. *Journal of Neuroscience*, 34(24), 8289–8299.
- Rizzoglio, F., Pierella, C., Santis, D., De, Mussa-Ivaldi, F. A., De Santis, D., Mussa-Ivaldi, F. A., & Casadio, M. (2020). A hybrid body-machine interface integrating signals from muscles and motions. *Journal of Neural Engineering*. <http://dx.doi.org/10.1088/1741-2552/ab9b6c>.
- Ruder, S. (2016). An overview of gradient descent optimization algorithms. (pp. 1–14). <http://arxiv.org/abs/1609.04747>.
- Sanchez, J. C., Mahmoudi, B., DiGiovanna, J., & Principe, J. C. (2009). Exploiting co-adaptation for the design of symbiotic neuroprosthetic assistants. *Neural Networks*, 22(3), 305–315. <http://dx.doi.org/10.1016/j.neunet.2009.03.015>.
- Scholz, M., Fraunholz, M., & Selbig, J. (2008). Nonlinear principal component analysis: Neural network models and applications. *Lecture Notes in Computational Science and Engineering*, 58, 44–67. http://dx.doi.org/10.1007/978-3-540-73750-6_2.
- Shadmehr, R., & Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14(5), 3208–3224.
- Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, 33(1), 89–108. <http://dx.doi.org/10.1146/annurev-neuro-060909-153135>.
- Shanechi, M. M. (2017). Brain-machine interface control algorithms. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(10), 1725–1734. <http://dx.doi.org/10.1109/TNSRE.2016.2639501>.
- Shenoy, K. V., & Carmena, J. M. (2014). Combining decoder design and neural adaptation in brain-machine interfaces. *Neuron*, 84(4), 665–680. <http://dx.doi.org/10.1016/j.neuron.2014.08.038>.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT Press.
- Tenenbaum, J. B., De Silva, V., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500), 2319–2323. <http://dx.doi.org/10.1126/science.290.5500.2319>.
- Thompson, B. (2005). Canonical correlation analysis. In *Encyclopedia of statistics in behavioral science*. American Cancer Society, <http://dx.doi.org/10.1002/0470013192.bsa068>.
- Vidaurre, C., Sannelli, C., Müller, K.-R., & Blankertz, B. (2010). Machine-learning-based coadaptive calibration for brain-computer interfaces. *Neural Computation*, 23(3), 791–816. http://dx.doi.org/10.1162/NECO_a_00089.
- Vujaklija, I., Shalchyan, V., Kamavuako, E. N., Jiang, N., Marateb, H. R., & Farina, D. (2018). Online mapping of EMG signals into kinematics by autoencoding. *Journal of NeuroEngineering and Rehabilitation*, 15(1), <http://dx.doi.org/10.1186/s12984-018-0363-1>.
- Wei, Y., Bajaj, P., Scheldt, R., & Patton, J. (2005). Visual error augmentation for enhancing motor learning and rehabilitative relearning. 2005. In *Proceedings of the 2005 IEEE 9th international conference on rehabilitation robotics* (pp. 505–510). <http://dx.doi.org/10.1109/ICORR.2005.1501152>.
- Wold, S., Esbensen, K., & Geladi, P. (1987). Principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, [http://dx.doi.org/10.1016/0169-7439\(87\)80084-9](http://dx.doi.org/10.1016/0169-7439(87)80084-9).