Continual Learning for EEG based Brain Computer Interfaces

Abstract

A healthy human brain manifests a high variance in signals captured from different modalities like EEG, MEG, and fMRI during ongoing activity. Further, there is brain signal variability exhibited at the user level. The study of within-subject and cross-subject variance is essential to design general Brain-Computer Interfaces (BCI) that help interpret these signals into intended outcomes. We propose that these variations can be studied under the umbrella of Continual Learning (CL), considering Domain Incremental and Task Incremental Continual Learning. We performed an empirical evaluation to understand the impact of CL strategies on the benchmark dataset. Our findings in within-subject and cross-subject scenarios suggest that CL strategies can outperform offline learning and build robust models for BCI applications. In the cross-subject scenario, CL can lead to learning invariant subject representation when transferring knowledge from one subject to another. In the within-subject scenario, CL can enhance performance when transferring knowledge from one session to another.

Keywords: Continual Learning; Brain Computer Interface; EEG;

1. Introduction

Noninvasive techniques that measure brain signals at the ensemble levels are sensitive to dynamic changes in neural activity arising as a response to various stimuli or task conditions(Faisal et al. (2008)). Study in individuals responding to stimuli suggests that specific individuals maintain high baseline activity or lower baseline activity while performing different tasks(Arazi et al. (2017a,b); Sheehan et al. (2018)). Hence, there is evidence of variability in brain signal across task conditions and subjects. Furthermore, components of brain function that vary across time lead to unreliable EEG signatures and show a lack of repeatability within-subject across trials. (Meyer et al. (2013)). In order to build robust BCI systems, the earliest attempts were to train users to modulate brain signals to overcome within-subject variations. Later, machine learning models were calibrated for each session and user, respectively. Recent studies have focused on transfer learning to provide a mechanism to indirectly transfer knowledge pertaining to the sources of within-subject and cross-subject variability (Saha and Baumert (2020)). Bakas et al. (2022) won the recent EEG Transfer Learning Competition at NeurIPS 2021. They address this difficulty by explicitly aligning feature distributions at various layers of the deep learning model, using both simple statistical techniques and trainable methods with more representational capacity. The feature alignment process is either subject or session-specific. Gibson et al. (2022) demonstrate that the variation in EEG signal strength and variability was found across subjects rather than across sessions. This was true across sessions that differed considerably in their behavioral and cognitive demands on the same task—further suggesting that EEG variability can be a sensitive subject-driven signal of interest.

© 2022 .

In this work, we propose that continual learning strategies must form an integral part of BCI systems to further the progress in the field because its crucial to determine how adaptation needs to be performed with a user in the loop. Especially considering cases of mental imagery in which human learning is involved (Lotte et al. (2013); Pfurtscheller and Neuper (2010)). Adaptation may not always assist; it may even be an impediment. A two-learner problem was proposed in Müller et al. (2017), indicating that either too fast or too slow adaptation can be detrimental to user learning. Thus, adaptive classifiers must be designed to ensure and favor human learning (Lotte et al. (2018)). The combination of transfer learning and online adaptive learning for BCI can be studied under the umbrella of CL as discussed in Mundt et al. (2021). The unique challenges in the EEG modality can further lead to robust strategies in CL. In this paper, we empirically evaluate CL strategies on the motor imagery dataset as a benchmark and try to assess the following:

- The impact of catastrophic forgetting in transfer learning for cross-subject and within-subject scenarios.
- The impact of existing CL strategies on classification accuracy for cross-subject and within-subject scenarios.

2. Methodology

2.1. Benchmark Dataset

We focus on EEG brain signals captured during motor movements or imagery for the study since they are considered fundamental applications for BCI. Our benchmark is based on the existing dataset: EEGMMID created and contributed to PhysioNet (Goldberger et al. (2000)) by the developers of the BCI2000 (Schalk et al. (2004)) instrumentation system, which they used in making these recordings. EEGMMID provides 64-channel EEG recordings of 109 volunteers at the sampling frequency 160Hz performing different motor movement and imagery tasks. Each subject performed 14 experimental runs: two one-minute baseline runs (one with eyes open, one with eyes closed) and three two-minute runs of each of the four following tasks:

- Task 1 (open and close left or right fist)
- Task 2 (imagine opening and closing left or right fist)
- Task 3 (open and close both fists or both feet)
- Task 4 (imagine opening and closing both fists or both feet)

We merge the movement and imagery tasks: Task 3 and Task 4, to augment the dataset. Each subject, on average, did 93 trials that ran for approximately 4 seconds where they opened or closed their fists or feet multiple times. We divided those 4-second trials into 2 seconds to augment the data and reduce class imbalance, resulting in 320 sample recordings in each sequence and an average of 185 sequences for each subject. Each sequence was labeled based on the activity: 0 for opening/closing of both fists and 1 for opening/closing of both feet. Data from each subject was first split 75%-25% into training and test sets while

maintaining class balance. For the within-subject scenario, the training and test set were from the same subject, while for the cross-subject scenario, the training set represented a different subject than the one from the test set. In both scenarios, the training set was then divided into five experiences used to train the model employed with CL strategy iteratively, each iteration followed by an evaluation on the test set.

2.2. Experimental Setup

We conduct the empirical evaluation to understand CL strategies' behavior in within-subject and cross-subject scenarios. We follow the experimental setup and implementation of Matteoni et al. (2022) for the empirical evaluation. We test CL strategies, namely:

- Replay (Hayes et al. (2021)): retains the previous dataset partially for the next experience.
- Elastic Weighted Consolidation (EWC) (Kirkpatrick et al. (2017)): restricts change in model weights based on its importance value.
- Learning without Forgetting (LwF) (Li and Hoiem (2017)): uses a combination of knowledge distillation and transfer learning.
- Gradient of Episodic Memory (GEM) (Lopez-Paz and Ranzato (2017)): projects the gradient on the current minibatch by using an external episodic memory of patterns from previous experiences.

Apart from these CL strategies, naive transfer learning and cumulative strategy (including current and all previous experiences as the training set) are also considered. We also compare the CL strategies with offline training, where the complete training dataset is trained in a single phase. Offline training acts as a higher baseline for the given task. We use an LSTM network with two layers and 64 hidden units. For Replay, the memory is set to 25%. Every training phase and strategy uses the same hyperparameters for Adam optimiser and cross-entropy loss as the loss criterion. We measure the accuracy on the test set after every experience for each strategy and scenario. All the experiments use the CL framework provided by Avalanche library(Lomonaco et al. (2021)) to manage dataset benchmarks, training phases, and evaluations.

3. Results and Discussion

1 and 2 report the accuracy achieved at the end of the training on all five experiences for cross-subject and within-subject scenarios, respectively. 1 and 2 report the accuracy achieved after training on each experience for cross-subject and within-subject scenarios, respectively.

3.1. CL strategies perform better than offline

Results reported in 1 and 2 show that in most cases except three cases, CL strategies outperform offline learning by a considerable margin. While only in one case for Subject ID 8 in the within-subject scenario, we find offline learning outperforming all the learning methods, suggesting considerable scope for the CL in BCI applications.

3.2. Memory driven vs. Regularisation Strategy

For both cross-subject and within-subject scenarios, memory or replay-driven strategies, namely Replay and Gradient of Episodic Memory(GEM), outperform other strategies. The reason may be correlated with the finding that there is no catastrophic forgetting in this case and other methods focus on avoiding forgetting by regularisation.

3.3. Better performance in Cross-Subject

It is surprising to observe that cross-subject scenarios have greater accuracy than within-subject scenarios for their respective subject ids. This finding needs further investigation to understand the cause.

3.4. Stopping or control criteria for CL

Referring to 1 and 2, one can note a significant drop in accuracy after training on specific experiences. It is crucial to define a control strategy or stopping criteria for CL strategies in order to protect the model from deterioration by learning on new data. We observe that further experimentation and theoretical framework is required to understand this space.

Subject ID	**Offline**	**Naive**	**Cumulative**	**Replay**	**Episodic**	**EWC**	**LwF**
3	65.96	55.32	48.94	**68.09**	59.57	55.32	63.83
4	36.96	50.00	45.65	**52.17**	**52.17**	36.96	39.13
5	**65.22**	**65.22**	45.65	47.83	58.70	60.87	58.70
6	60.87	60.87	56.52	**65.22**	63.04	54.35	63.04
7	63.83	61.70	63.83	48.94	65.96	57.45	**70.21**
8	50.00	**63.04**	45.65	56.52	50.00	60.87	60.87

Table 1: Final accuracy evaluated on the test set of given subject id after training on five experiences for the cross-subject scenario.

Subject ID	**Offline**	**Naive**	**Cumulative**	**Replay**	**Episodic**	**EWC**	**LwF**
3	59.57	48.94	51.06	59.57	57.45	51.06	**63.83**
4	47.83	**56.52**	47.83	**56.52**	52.17	**56.52**	41.30
5	56.52	54.35	58.70	54.35	**65.22**	52.17	43.48
6	**56.52**	50.00	52.17	45.65	**56.52**	52.17	54.35
7	53.19	65.96	**70.21**	55.32	55.32	51.06	65.96
8	**69.57**	58.70	54.35	41.30	54.35	58.70	58.70

Table 2: Final accuracy evaluated on the test set of given subject id after training on five experiences for the within-subject scenario.

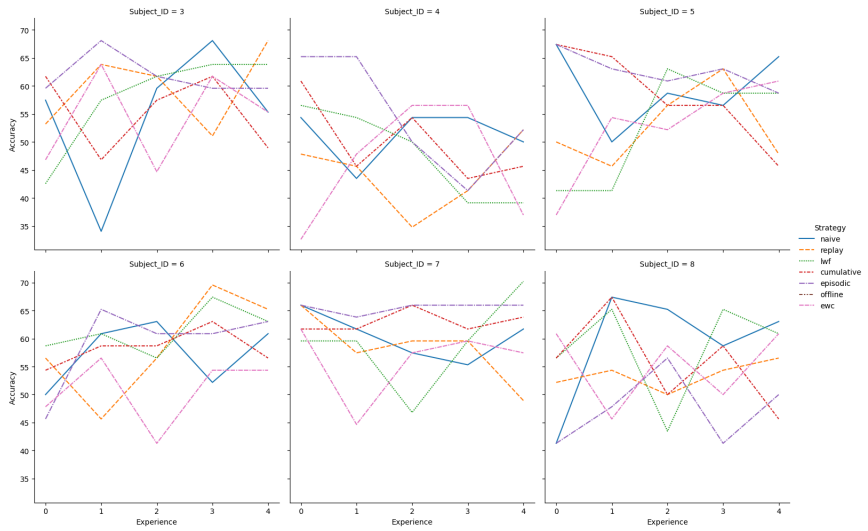

Figure 1: Accuracy on the test dataset of given subject id after training on each experience for cross subject scenario.

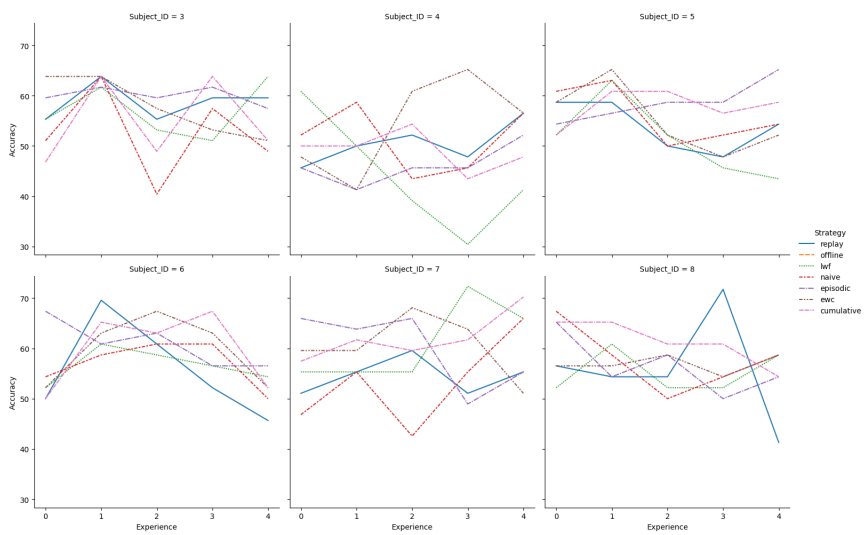

Figure 2: Accuracy on the test dataset of given subject id after training on each experience for within subject scenario.

4. Conclusion and Future Work

We can conclude that there is a necessity and an opportunity for inducting continual learning in Brain-computer interface(BCI) systems, having observed better performance than offline learning. These findings certainly need to be corroborated by performing experiments on more subjects and other datasets. We could run experiments for six subjects; however, the work is in progress to include all 109 subjects for the analysis. Further studies can also be centered on identifying distributions to define control strategies or stopping criteria for continual learning. The work can be extended by leveraging CL strategies for domain incremental and task incremental scenarios on a large corpus of EEG data with multiple tasks and sessions.

References

Ayelet Arazi, Nitzan Censor, and Ilan Dinstein. Neural variability quenching predicts individual perceptual abilities. *Journal of Neuroscience*, 37(1):97–109, 2017a.

Ayelet Arazi, Gil Gonen-Yaacovi, and Ilan Dinstein. The magnitude of trial-by-trial neural variability is reproducible over time and across tasks in humans. *eneuro*, 4(6), 2017b.

Stylianos Bakas, Siegfried Ludwig, Konstantinos Barmpas, Mehdi Bahri, Yannis Panagakis, Nikolaos Laskaris, Dimitrios A Adamos, and Stefanos Zafeiriou. Team cogitat at neurips 2021: Benchmarks for eeg transfer learning competition. *arXiv preprint arXiv:2202.03267*, 2022.

A Aldo Faisal, Luc PJ Selen, and Daniel M Wolpert. Noise in the nervous system. *Nature reviews neuroscience*, 9(4):292–303, 2008.

Erin Gibson, Nancy J Lobaugh, Steve Joordens, and Anthony R McIntosh. Eeg variability: Task-driven or subject-driven signal of interest? *NeuroImage*, 252:119034, 2022.

Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *circulation*, 101(23):e215–e220, 2000.

Tyler L Hayes, Giri P Krishnan, Maxim Bazhenov, Hava T Siegelmann, Terrence J Sejnowski, and Christopher Kanan. Replay in deep learning: Current approaches and missing biological elements. *Neural Computation*, 33(11):2908–2950, 2021.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.

Vincenzo Lomonaco, Lorenzo Pellegrini, Andrea Cossu, Antonio Carta, Gabriele Graffieti, Tyler L Hayes, Matthias De Lange, Marc Masana, Jary Pomponi, Gido M Van de Ven, et al. Avalanche: an end-to-end library for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3600–3610, 2021.

David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017.

Fabien Lotte, Florian Larrue, and Christian Mühl. Flaws in current human training protocols for spontaneous brain-computer interfaces: lessons learned from instructional design. *Frontiers in human neuroscience*, 7:568, 2013.

Fabien Lotte, Laurent Bougrain, Andrzej Cichocki, Maureen Clerc, Marco Congedo, Alain Rakotomamonjy, and Florian Yger. A review of classification algorithms for eeg-based brain–computer interfaces: a 10 year update. *Journal of neural engineering*, 15(3):031005, 2018.

Federico Matteoni, Andrea Cossu, Claudio Gallicchio, Vincenzo Lomonaco, and Davide Bacciu. Continual learning for human state monitoring. *arXiv preprint arXiv:2207.00010*, 2022.

Matthias C Meyer, Erik SB van Oort, and Markus Barth. Electrophysiological correlation patterns of resting state networks in single subjects: a combined eeg–fmri study. *Brain topography*, 26(1):98–109, 2013.

Jan Saputra Müller, Carmen Vidaurre, Martijn Schreuder, Frank C Meinecke, Paul Von Bünau, and Klaus-Robert Müller. A mathematical model for the two-learners problem. *Journal of neural engineering*, 14(3):036005, 2017.

Martin Mundt, Steven Lang, Quentin Delfosse, and Kristian Kersting. Cleva-compass: A continual learning evaluation assessment compass to promote research transparency and comparability. *arXiv preprint arXiv:2110.03331*, 2021.

G Pfurtscheller and C Neuper. Neurofeedback training for bci control, brain–computer interfaces. *The Frontiers Collection*, page 65, 2010.

Simanto Saha and Mathias Baumert. Intra-and inter-subject variability in eeg-based sensorimotor brain computer interface: a review. *Frontiers in computational neuroscience*, 13:87, 2020.

Gerwin Schalk, Dennis J McFarland, Thilo Hinterberger, Niels Birbaumer, and Jonathan R Wolpaw. Bci2000: a general-purpose brain-computer interface (bci) system. *IEEE Transactions on biomedical engineering*, 51(6):1034–1043, 2004.

Timothy C Sheehan, Vishnu Sreekumar, Sara K Inati, and Kareem A Zaghloul. Signal complexity of human intracranial eeg tracks successful associative-memory formation across individuals. *Journal of Neuroscience*, 38(7):1744–1755, 2018.