

Transfer learning promotes acquisition of individual BCI skills

Satyam Kumar ^{a,*}, Hussein Alawieh ^a, Frigyes Samuel Racz ^{b,c}, Rawan Fakhreddine ^a and José del R. Millán ^{a,b,c,d,*}

^aChandra Family Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX 78712, USA

^bDepartment of Neurology, The University of Texas at Austin, Austin, TX 78712, USA

^cMulva Clinic for the Neurosciences, The University of Texas at Austin, Austin, TX 78712, USA

^dDepartement of Biomedical Engineering, The University of Texas at Austin, Austin, TX 78712, USA

*To whom correspondence should be addressed: Email: satyam.kumar@utexas.edu (S.K.); jose.millan@austin.utexas.edu (J.d.R.M.)

Edited By: Lydia Kavraki

Abstract

Subject training is crucial for acquiring brain–computer interface (BCI) control. Typically, this requires collecting user-specific calibration data due to high inter-subject neural variability that limits the usability of generic decoders. However, calibration is cumbersome and may produce inadequate data for building decoders, especially with naïve subjects. Here, we show that a decoder trained on the data of a single expert is readily transferrable to inexperienced users via domain adaptation techniques allowing calibration-free BCI training. We introduce two real-time frameworks, (i) Generic Recentering (GR) through unsupervised adaptation and (ii) Personally Assisted Recentering (PAR) that extends GR by employing supervised recalibration of the decoder parameters. We evaluated our frameworks on 18 healthy naïve subjects over five online sessions, who operated a customary synchronous bar task with continuous feedback and a more challenging car racing game with asynchronous control and discrete feedback. We show that along with improved task-oriented BCI performance in both tasks, our frameworks promoted subjects' ability to acquire individual BCI skills, as the initial neurophysiological control features of an expert subject evolved and became subject specific. Furthermore, those features were task-specific and were learned in parallel as participants practiced the two tasks in every session. Contrary to previous findings implying that supervised methods lead to improved online BCI control, we observed that longitudinal training coupled with unsupervised domain matching (GR) achieved similar performance to supervised recalibration (PAR). Therefore, our presented frameworks facilitate calibration-free BCIs and have immediate implications for broader populations—such as patients with neurological pathologies—who might struggle to provide suitable initial calibration data.

Significance Statement

In this manuscript, we show for the first time how inter-subject transfer learning approaches enable inexperienced users to immediately operate a noninvasive brain–computer interface (BCI), thus avoiding the need for the standard individual calibration session. Acquisition of calibration data to build the BCI decoder is expensive and error-prone since naïve subjects may likely generate brain signals with inadequate discriminant power due to the lack of feedback. Consequently, the resulting decoder will perform poorly, thus hampering BCI training. Our approach provides a solution to this issue as it relies on a subject-independent decoder built on data of one single expert subject that renders calibration unnecessary. The efficacy of the proposed approach was demonstrated in online continuous and discrete feedback tasks.

Introduction

Noninvasive brain–computer interfaces (BCI) based on electroencephalography (EEG) have proven efficient in applications such as neurorehabilitation (1, 2), robotics (3, 4), communication (5, 6), or virtual reality (7, 8). Motor imagery (MI)—mental rehearsal of a limb movement without execution—is a common EEG–BCI modality. MI elicits distinct sensorimotor rhythms (SMR) for different movements (9, 10); however, online decoding suffers from the nonstationary nature of EEG. Although complex machine learning (ML) models can alleviate this problem, a significant portion of subjects often exhibits a classification performance close to

chance level (11), thus, subject's learning of the BCI skill—generation of distinctive SMR—also seems crucial to operate brain-controlled devices (12–15). Consequently, mutual learning—building ML models that promote subject's acquisition of BCI skills—has gained increasing attention (13, 15–21) and remains an open problem in BCI.

Training a BCI subject customarily starts with an offline calibration session to collect data to build an individual decoder. Apart from being time-consuming, this initial decoder might be inefficient as subjects do not receive feedback that helps them to elicit proper SMR during calibration. A solution is to leverage

Competing Interest: The authors declare no competing interest.

Received: May 5, 2023. **Accepted:** February 5, 2024

© The Author(s) 2024. Published by Oxford University Press on behalf of National Academy of Sciences. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial reproduction and distribution of the work, in any medium, provided the original work is not altered or transformed in any way, and that the work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

pre-recorded data to build subject-independent MI–BCI decoders. In the BCI field, this phenomenon is commonly referred to as inter-subject transfer learning (22), in analogy with the more common understanding of transfer learning in the ML field literature as transferring the knowledge of a model trained on one domain to another domain (23). Inter-subject transfer learning strategies, however, need to cope with across-subject SMR variability and require an expensive data collection effort, which may include subjects with poor BCI performance due to nondiscriminable SMR. Here we propose a Riemannian incremental domain adaptation framework, which performs statistical matching between SMR distributions across subjects (22, 24), for inter-subject transfer learning based on pre-recorded data from a single expert BCI subject (Fig. 1). We hypothesized that our framework supports longitudinal MI–BCI training and promotes learning of naïve subjects.

Our framework assumes that different subjects have their SMR covariance features shifted on the Riemannian manifold (22) (Fig. 1a). Our approach matches the data distributions of the expert and the naïve subjects in real time using first-order statistics (Fig. 1b). As a result, the BCI provides contingent robust feedback, thus enabling participants to operate the brain-controlled device immediately and acquire their BCI skills via longitudinal training. A key element is that the matching, or recentering, process is unsupervised and takes place continuously for the naïve subject, thus coping also with the intra-subject SMR variability over time—within and across BCI sessions. Nevertheless, it could be the case that recentering might not suffice to handle a strong inter- and intra-subject SMR variability. Therefore, we tested two variants of our inter-subject transfer learning framework. The first one, Generic Recentering (GR), keeps the decision boundary of the expert decoder fixed throughout the entire experiment (Fig. 1c). Our second framework, Personally Adjusted Recentering (PAR), extends GR by tuning the parameters of the expert decoder with a small chunk of the incoming naïve subject data (Fig. 1d).

To test our hypothesis, we recruited 18 BCI-naïve, healthy volunteers ($N = 9$ for each GR and PAR group) to participate in a 5-day training program (Fig. 1e). We assessed the efficacy of GR and PAR in two different settings (Fig. 1f): a standard bar task and the Cybathlon car racing game (25), a more realistic application. In the first task, controlled synchronously, BCI feedback was continuous, while the game provided discrete feedback when subjects asynchronously delivered a turning command for the car. In each training session, participants completed multiple runs with the bar task followed by the racing game. Of note, the expert data was acquired in the bar task and the expert never played the car racing game. We show that subjects could learn to operate their MI–BCI in both GR and PAR frameworks following the longitudinal training for both tasks. Moreover, we provide evidence that improvement in BCI control was partly due to subjects acquiring the skill of producing increasingly discriminant neurophysiological features, which not necessarily matched those of the original expert subject. Finally, contrary to the popular belief that personalization or supervised tuning of decoders would lead to better BCI performance (24), subjects in the GR and PAR groups reached statistically similar performance in both tasks.

Results

Acquiring BCI control: the bar task

In a trial subjects performed left- or right-hand MI. A trial ended when the cumulative probability for a certain MI class exceeded

a predefined threshold for command delivery or when a timeout period elapsed. During the task execution period of each trial, subjects received continuous visual feedback reflecting the accumulated evidence by the decoder. We used Cohen's kappa (26) ($\kappa \in [-1, 1]$, chance level = 0) to characterize command delivery performance, which was then adjusted for the number of timeout trials [normalized kappa value (NKV)]. Since the first run of the PAR framework uses the ground truth labels of the trials to adjust the decoder parameters, that run was excluded from analysis. We also measured the command latency (CL) to quantify how long it took to subjects to deliver a correct command.

Participants achieved significant NKV improvement over training in both GR (start: 0.2636 ± 0.1351 , end: 0.4694 ± 0.2293 , $N = 9$, $P = 0.02$) and PAR frameworks (start: 0.4045 ± 0.2857 , end: 0.6802 ± 0.2526 , $N = 9$, $P = 0.001$). Moreover, NKV improvement exhibited a statistically significant increasing trend for GR and PAR (Fig. 2a and b). Although subjects in the PAR group showed generally higher NKV scores on average over the training sessions, the difference between the two groups never reached statistical significance (Supplementary material, Complementary statistical analysis). Similar to NKV, subjects improved their CL over sessions and exhibited a significantly decreasing trend in both frameworks (Fig. 2c and d). CL between GR and PAR were statistically similar across sessions (Supplementary material, Complementary statistical analysis).

BCI control in a realistic scenario: car racing

While the bar task is controlled synchronously and provides continuous feedback, it does not reflect real-life settings that may involve discrete feedback only upon the delivery of asynchronous BCI commands. Therefore, the performance of the participants was further assessed in the Cybathlon car racing game, used in other studies (27–29). The racing track was modified to consist only of right and left turn patches. Since a player with higher efficacy (i.e. better accuracy and shorter latency) in BCI command delivery finishes the race faster, we used the race competition timing (RCT) as the primary metric. We also report the NKV scores and CL for the sake of completeness.

Participants finished the races significantly faster after longitudinal training for both GR (start: 203.63 ± 18.07 s, end: 162.61 ± 22.15 s, $N = 9$, $P = 0.00012$) and PAR (start: 192.23 ± 23.22 s, end: 164.92 ± 35.86 s, $N = 9$, $P = 0.018$). RCT values showed a significantly negative trend across sessions for the two frameworks (Fig. 3a and b). RCTs of all subjects were significantly lower when compared to the upper bound of the metric (378.88 s); i.e. when no or only incorrect commands are delivered during race completion.

Similar to RCT, participants also exhibited a significantly increasing trend of their NKV scores over the sessions for both GR and PAR (Fig. 4a and b). Moreover, subjects improved their CL over sessions, showing a statistically significant decreasing trend in both the frameworks (Fig. 4c and d).

RCT, NKV, and CL scores indicate that the longitudinal training with both the GR and PAR frameworks promotes acquisition of MI–BCI control in a realistic scenario with high cognitive demand (i.e. playing racing games). While all performance metrics (RCT, NKV, and CL) showed insignificantly better scores in the PAR framework at the beginning of the training, the cross-group difference in those scores was considerably minimized at the end of the training, and eventually reverted slightly for RCT (Supplementary material, Complementary statistical analysis).




Fig. 1. Longitudinal frameworks for inter-subject BCI transfer learning. a) Differences in individual EEG results in covariate shifts of MI sample distributions across subjects. A decoder built on expert data leads to biased classification on naive subject data. b) Domain adaptation to match the sample distributions of the expert and naive subjects improves classification. c1.1) Generic Recentering (GR) framework: the samples of the naive subject are matched with those of the expert in real time during an online session. The transformed naive subject samples are classified using a fixed decoder that is trained on the transformed expert samples. c1.2) Incremental domain adaptation gets better at matching the sample distributions with more incoming data. c2) The decoder decision boundary is kept fixed throughout a session. d1.1) Personally Adjusted Recentering framework: after GR, the incoming labeled data of the naive subject is used to adjust the parameters of the expert decoder. d1.2) The parameters of the decoder are updated in real-time using labeled data from the first run of a session exclusively. d2) After the first run in each session, the decoder is fixed for all remaining runs, and only incremental recentering still continues to remove the nonstationarities due to covariate shift. e) Experimental protocols for the longitudinal training frameworks (GR and PAR). Each online session consists of multiple runs in which subjects operate a BCI that provides feedback on their MI skill level. f) The two BCI tasks: a bar with continuous feedback (left) and car racing with discrete feedback (right).




Fig. 2. BCI performance over bar training sessions. a, b) Subject-wise NKV of command delivery in GR and PAR framework, respectively. Each point corresponds to the average NKV across runs of a session for a given subject (GR: $N = 4$ and $N = 3$ in online session 1 and online sessions 2–5, respectively; PAR: $N = 3$ and $N = 2$ in online session 1 and online sessions 2–5, respectively). The colored line represents the grand average across all subjects over the online sessions. Correlation analysis using linear mixed effect modeling of GR NKV ($r^2 = 0.6795$, $P = 0.002$, $N = 45$), PAR NKV ($r^2 = 0.7853$, $P = 0.00002$, $N = 45$). c) GR command latency (CL) (start: 2.3540 ± 0.3801 s, end: 2.0833 ± 0.5151 s, $r^2 = 0.38938$, $P = 0.0180$, $N = 45$), and d) PAR CL (start: 2.5994 ± 0.7290 s, end: 2.2327 ± 0.68085 s, $r^2 = 0.60879$, $P = 0.0249$, $N = 45$) across online sessions. Each marker point in the corresponding plot denotes the NKV or CL for a specific subject in an online session. The colored line represents the grand average across all subjects over the online sessions. The statistical analysis results (linear mixed effect modeling) are reported in the same color.




Fig. 3. Brain-controlled car races. a, b) Subject-wise race completion timing (RCT) in the GR and PAR frameworks, respectively. Each point corresponds to the average RCT across races of a session for a given subject. The colored line represents the grand average across all subjects over online sessions. Correlation analysis using linear mixed effect (LME) modeling of GR RCT ($r^2 = 0.7620$, $P < 1 \times 10^{-5}$, $N = 45$). Pearson Correlation between online sessions and PAR RCT ($r^2 = 0.1007$, $P < 0.05$, $N = 45$). † denotes that correlation modeling didn't undertake into account the random effects of subjects as normality criteria was not satisfied on the LME.

Improved discriminatory of neurophysiological features

Although Riemannian geometry classifiers have been used for MI-BCIs (22, 24, 30–32), no previous work provided a neurophysiological interpretation of their learned covariance matrices. For this purpose, we define the electrode discriminatory score (EDS) as a novel measure that characterizes the contribution of each cortical location for effective discrimination of the two MI classes.

Figure 5 presents the topoplots visualizations of z-transformed EDS values for the expert subject data (Fig. 5a), as well as for the GR (Fig. 5b) and PAR (Fig. 5c) groups in the bar task and car racing game. For the expert case, electrodes with a statistically significant EDS were C4 ($z = 2.17$) and FC6 ($z = 2.55$); contralateral channels C3 and FC5 showed higher than average, yet insignificant, EDS values. For the GR group, electrodes with statistically significant EDS were C4 (bar: $z = 3.46$, car racing: $z = 1.83$) and C3 (bar:




Fig. 4. Closed-loop BCI performance throughout races. Grand average visualization of the group-level distribution of a) GR NKV (start: 0.1354 ± 0.1288 , end: 0.4128 ± 0.2328 , $r^2 = 0.7807$, $P < 2 \times 10^{-5}$, $N = 45$), b) PAR NKV (start: 0.2875 ± 0.1319 , end: 0.4535 ± 0.3355 , $r^2 = 0.804$, $P = 0.0013$, $N = 45$), c) GR command latency (CL) (start: 7.6361 ± 1.4602 s, end: 4.7100 ± 1.8369 s, $r^2 = 0.66656$, $P = 2.8737 \times 10^{-6}$, $N = 45$), d) PAR CL (start: 6.3719 ± 2.7751 s, end: 4.4406 ± 2.2214 s, $r^2 = 0.75553$, $P = 2.5537 \times 10^{-4}$, $N = 45$) during the races of the online sessions. Each marker denotes the NKV (or CL) averaged across races for a given subject in an online session. The colored line represents the grand average across all subjects over online sessions for corresponding metric.

$z = 2.40$, car racing: $z = 3.34$). For the PAR framework, average significant EDS were at C4 ($z = 3.48$) and FC5 ($z = 2.01$) in the bar task and at C4 ($z = 1.90$), and F8 ($z = 2.80$) for the car racing. Importantly, all the electrodes with significant EDS correspond to neurophysiologically relevant locations recruited during MI of left- or right-hand movements (9–11, 33). Remarkably, EDS patterns for naïve subjects in both groups differed from the original EDS pattern for the expert subject (see also Figs. S9–S12).

With evidence supporting the neurophysiological relevance of the covariance features, we assess the evolution of their discriminability [feature distinctiveness (FD)] to substantiate subject learning (13, 15, 21, 34). Subjects in both groups exhibited a significantly increasing trend in FD for both tasks (Fig. 6). There was no statistical difference in FD between the GR and PAR groups for any of the tasks in any session (except for first online session in races, see also Fig. S13).

Pseudo-online inter-subject transfer learning

We compared our proposed inter-subject transfer learning approaches, termed *Expert-Subject* below, against two other subject-specific BCI decoders in a post hoc pseudo-online setting:

- **Subject-Specific-Fix:** We built a Riemannian MDM decoder using the calibration data of the corresponding subject, which then remained fixed during the online sessions.
- **Subject-Specific-Adapt:** We also built this Riemannian MDM decoder from the calibration session of each of the subjects, which was then adapted during the online sessions according to the corresponding adaptation framework of the subject group; i.e. GR or PAR.

Comparison against *Subject-Specific-Fix* decoders allows to assess the importance of domain adaptation through the

proposed incremental recentering approach. Comparison against *Subject-Specific-Adapt* decoders allows to quantify the benefits of a subject-specific BCI with domain adaptation. We compared the three decoders on the bar task at the sample level. The reason to evaluate performance at the sample level instead of the trial level (delivery of a BCI command) is to avoid tuning the evidence accumulation hyperparameters that, if not done properly for each algorithm, could bias results. We limit the analysis to the bar task because in the car racing task it is not known when subjects start MI to deliver a BCI command.

To test the statistical significance, we performed paired t-tests (across all subjects and all online runs) between the three types of decoders. For the GR framework (Fig. 7a), a paired t-test or Wilcoxon signrank test (where normality was not satisfied) revealed a statistically significant difference between *Expert-Subject* and *Subject-Specific-Fix* ($P = 6.18 \times 10^{-8}$, 0.3777 ± 0.2908 vs. 0.2345 ± 0.2884 , $N = 144$). Furthermore it shows a strong trend towards statistical significance between *Expert-Subject* and *Subject-Specific-Adapt* ($P = 0.051$, 0.3777 ± 0.2908 vs. 0.3283 ± 0.3262 , $N = 144$). Finally, we observe a statistically significant difference between *Subject-Specific-Adapt* and *Subject-Specific-Fix* decoder ($P = 1.17 \times 10^{-8}$, 0.3283 ± 0.3262 vs. 0.2345 ± 0.2884 , $N = 144$). For the PAR framework (Fig. 7b), we observe a similar trend where both *Expert-Subject* (0.6249 ± 0.3312 , $N = 99$) and *Subject-Specific-Adapt* (0.6149 ± 0.3432 , $N = 99$) were significantly better than *Subject-Specific-Fix* (0.2929 ± 0.3289 , $N = 99$) ($P = 6.86 \times 10^{-15}$ and $P = 8.34 \times 10^{-15}$, paired t-test and Wilcoxon signrank test, respectively). However, we do not find any statistical difference between *Expert-Subject* and *Subject-Specific-Adapt*.

Importantly, performance of the *Expert-Subject* decoder was competitive since the very first session (Fig. 7), thus delivering




Fig. 5. Neurophysiological interpretation of covariance features. Topoplots show z-score values of electrode discriminancy scores (EDS) corresponding to each electrode. a) z-statistics of the EDS for the expert subject. b, c) z-statistics of the EDS for the GR and PAR groups ($N = 9$) in each BCI task (bar and race), respectively. EDS were averaged across the last two online sessions of each subject, followed by averaging across the subjects in the corresponding framework (GR or PAR) in each BCI task (bar and race).

consistent feedback that enabled subjects to acquire BCI control. For the GR framework (Fig. 7c), a paired t -test revealed a statistically significant difference between Subject-Specific-Adapt and Subject-Specific-Fix decoder ($P = 0.0012$, 0.2855 ± 0.2557 vs. 0.1869 ± 0.2320 , $N = 36$), while Expert-Subject (0.2576 ± 0.2498 , $N = 36$) performed similarly to Subject-Specific-Adapt and to Subject-Specific-Fix. For the PAR framework (Fig. 7d), both Expert-Subject (0.4574 ± 0.3803 , $N = 27$) and Subject-Specific-Adapt (0.4970 ± 0.3747 , $N = 27$) were significantly better than Subject-Specific-Fix (0.2468 ± 0.2864 , $N = 27$) ($P = 0.0012$ and $P = 0.0005$, paired t -test and Wilcoxon signrank test, respectively). However, we do not find any statistical difference between Expert-Subject and Subject-Specific-Adapt.

GR transfer learning with common spatial patterns

Having demonstrated transfer learning with Riemannian MDM decoders, we also investigated the adaptability of the GR framework to common spatial pattern-based (CSP) classifiers, which is widely used in MI-BCIs (35, 36).

We name CSP-Expert-Subject the decoder trained on the expert subject data using the proposed GR-CSP framework. Figure 7e and f presents the sample-wise kappa value comparison between the Riemannian MDM-Expert-Subject (originally referred as Expert-Subject) and the formulated CSP-Expert-Subject the first online session as well as over all the online sessions. We observe that there is no statistical difference in online session 1 (MDM-Expert-Subject: 0.2576 ± 0.2498 vs. CSP-Expert-

Subject: 0.2919 ± 0.2189 , $N = 36$). We did not observe either any statistically significant difference over all the online sessions (MDM-Expert-Subject: 0.3777 ± 0.2908 vs. CSP-Expert-Subject 0.3727 ± 0.2924 , $N = 144$).

Finally to highlight the importance of adaptation and compatibility of the proposed incremental GR-CSP framework, we mimicked the scenario of subject-specific adaptation and a CSP decoder without any adaptation. For that purpose, we use the following two decoders:

- CSP-Subject-Specific-Adapt: We built this GR-CSP decoder from the calibration session of each of the subjects in the GR framework, which was then adapted during the online sessions.
- CSP-Subject-Specific-Fix: We built a CSP decoder using the calibration data of the corresponding subject, which then remained fixed during the online sessions.

In Fig. 7g, we observe that in first online session CSP-Expert-Subject is statistically significantly better than CSP-Subject-Specific-Fix ($P = 0.0013$, 0.2919 ± 0.2189 vs. 0.1521 ± 0.1889 , $N = 36$). However, there was no significant difference between CSP-Expert-Subject and CSP-Subject-Specific-Adapt (0.2072 ± 0.2229 , $N = 36$). Figure 7h shows the comparison of three decoders over all the online sessions. Interestingly, we observe that CSP-Expert-Subject (0.3727 ± 0.2924 , $N = 144$) significantly outperforms both the CSP-Subject-Specific-Fix (0.3727 ± 0.2924 , $N = 144$, 1.05×10^{-10})




Fig. 6. Multisession longitudinal closed-loop BCI training enables user learning. Correlation analysis of feature distinctiveness (FD) in, a) GR bar task ($r^2 = 0.7001$, $P = 0.0001$, $N = 45$), b) PAR bar task ($r^2 = 0.7145$, $P = 0.0019$, $N = 45$), c) GR car racing ($r^2 = 0.8055$, $P = 0.00082$, $N = 45$), and d) PAR car racing ($r^2 = 0.7109$, $P = 0.0081$, $N = 45$). Each data point is the FD between the two MI classes (left-hand MI and right-hand MI) in a session in the corresponding framework and task condition. The colored line represents the grand average across all subjects over online sessions for corresponding metric. ζ denotes that Linear mixed effect modeling was performed on the log transformed feature distinctiveness as normality of residuals is satisfied on log transformed response variable.

and CSP-Subject-Specific-Adapt (0.3727 ± 0.2924 , $N = 144$, $P = 7.22 \times 10^{-5}$) decoders. Finally, we also observe that CSP-Subject-Specific-Adapt is significantly better than CSP-Subject-Specific-Fix ($P = 0.005$).

Our pseudo-online results with the GR-CSP framework emphasize the significance of using the GR framework for incremental unsupervised adaptation in longitudinal BCIs. Both of the adaptive frameworks, namely CSP-Expert-Subject and CSP-Subject-Specific-Adapt outperformed the nonadaptive decoder CSP-Subject-Specific-Fix (Fig. 7b). More importantly, the adaptive decoders demonstrate superior performance compared to fixed CSP decoders right from the first online session (Fig. 7e). Furthermore, there is no observable difference in performance between the MDM-Expert-Subject and CSP-Expert-Subject methods across both the initial online session and all subsequent online sessions (Fig. 7). We postulate that our integration of the GR framework with the CSP approach results in a stable estimation of spatial filters in a recentered space and the online incremental adaptation, following the alignment of signal covariances, aids in mitigating nonstationarities, ultimately improving classification.

Discussion

We proposed two inter-subject transfer learning frameworks exploiting domain adaptation for longitudinal BCI training, and we demonstrated how they promoted acquisition of individual BCI skills. Indeed, participants exhibited the two pivotal elements of skill learning, namely increased accuracy and faster execution of BCI commands over the sessions (Figs. 2–4). Contrarily to existing transfer learning approaches, ours only require data from one

single expert subject to build the initial decoder. Both frameworks proved to be highly efficient, as participants were able to immediately operate brain-controlled devices BCI-naïve subjects not only enhanced their control significantly in a laboratory setting but also in a more complex scenario (car racing game). This further highlights that the decoding model was transferable across tasks as the expert subject never played the car racing game. Furthermore, we introduced a novel method for untangling the neurophysiological relevance of the features modulated by participants to control the BCI. This provided evidence of subject learning via the consistently improving separability/discriminability of features (Fig. 6), which differentiated from those of the initial expert subject (Fig. 5) and became subject specific (Figs. S9–S12).

Is inter-subject transfer learning superior to a subject-specific BCI?

Although our results strongly support that the typical subject-specific calibration session can be omitted, a fundamental question arises: how does the proposed inter-subject transfer learning framework perform with respect to a subject-specific BCI? To answer this question, we compared our framework against two other BCI decoders built with the data of the calibration session of each subject in a post hoc pseudo-online setting of the bar task. The first decoder remained fixed during the sessions, while the second one was adapted during the sessions according to the corresponding adaptation framework of the subject group (GR or PAR).

Although comparing decoders in a pseudo-online setting might bias the results towards the decoder used online—as subjects attempt to generate patterns aligned with the feedback they receive—the expert decoders in GR and PAR performed statistically




Fig. 7. Post hoc pseudo-online analysis. Box plot distribution of the sample-wise classification kappa value of different decoders. Comparison of Expert-Subject, Subject-Specific-Fix, and Subject-Specific-Adapt for the two proposed frameworks, a) GR All online sessions, b) PAR All online sessions, c) GR first online session, and d) PAR first online session. Decoder comparison of CSP-Expert-Subject and MDM-Expert-Subject over, e) the first online session, and f) all the online sessions. Sample wise performance of Expert-Subject, Subject-Specific-Fix and Subject-Specific-Adapt with GR-CSP, g) the first online session and h) all the online sessions. Each data point corresponds to the kappa value of an online run for a subject. Compared to GR (All sessions: $N = 144$, First Session: $N = 36$), in the PAR framework (All sessions: $N = 99$, First Session: $N = 27$), the first run in each session is removed from the analysis as it was used for the supervised update. The box edges are 75th and 25th percentiles, and the horizontal line corresponds to the median of the distribution. Note that in the PAR case, the first run in each session is removed from the analysis as it was used for the supervised update. Numbers above the bars correspond to the P-value between the distributions of two groups (A superscript indicates scenarios when normality of data was violated and thus Wilcoxon signrank test was used, while paired t-tests were used otherwise).

better than the subject-specific BCI whose decoder remained fixed after calibration but had similar performance to the adaptive, subject-specific decoder (Fig. 7a and b). Furthermore, subjects achieved statistically similar or better performances with their corresponding transfer learning framework than with the subject-specific decoders already in the first session (Fig. 7c and d). Immediate competitive performance of the transfer learning approaches is critical to enable subjects to acquire BCI control. Moreover, our proposed integration of GR framework with CSP reveals that transferring CSP decoders between novice and expert subjects using GR framework results in a statistically similar performance to subject-specific CSP decoders even from first online session. However, despite CSP and MDM-based classification approaches yielding similar classification performance, we argue that MDM-based approaches are better suited for online BCIs due to their ability to perform real-time decoder parameter updates (i.e. the PAR framework). In contrast, updating spatial filters in CSP and subsequently incrementally updating the linear discriminant classifier would be a challenging task.

These results confirm that our inter-subject transfer learning approaches can eliminate the need for subject-specific calibration sessions.

Mutual learning

Training users to acquire MI skills in a longitudinal, closed-loop BCI setting seems critical for the successful operation of a brain-controlled device (4, 12–16, 18, 21, 27, 28, 37). We

demonstrated here that our longitudinal inter-subject transfer learning frameworks also promoted subject learning. Furthermore, there is compelling evidence that acquisition of BCI control is grounded on cortical reorganization (12, 13, 15, 16, 18, 37). Since model parameters were kept fixed throughout the experiment in the GR framework and only feature spaces were matched using incremental recentering, the subjects had to modulate their brain patterns distinctively for both classes to produce patterns distinguishable by the fixed decoder—thus receiving concurrent positive feedback. We refer to neuroplasticity when subjects improved their distinctive modulation of brain patterns during the training course. Our findings reveal that subjects using the GR framework exhibited a significant and increasing trend in discriminability within the feature spaces (Fig. 6a), indicating that the generated patterns became more distinguishable over time. As another indication of lasting changes in neural activity, we analyzed the strengthening in event-related desynchronization (ERD) in the fifth online session compared to the first online session for subjects in the GR group in the bar task (Fig. S2). Even though the extent of the difference did not reach the level of statistical significance after multiple comparisons adjustment, it is observed that participants were able to evoke a stronger ERD at the end compared to the beginning of their training, further illustrating persisting changes in neural activity. In the PAR group of the bar task, even though the decoding model parameters were adapted using a small amount of data (first run of bar task) at the beginning of each session, they remained fixed throughout the rest of the online session each day. Our

incremental recentering framework also helped by aligning the feature spaces to ameliorate nonstationarities. However, subjects still needed to produce discriminative brain patterns on the correct side of the hyperplane. Similar to the GR framework, in PAR we again observed a significant and consistent improvement in feature space discriminability (Fig. 6b). Furthermore, Fig. S3 also indicates enhanced ERD over the contralateral side during the fifth online session compared to the first online session.

Our results also illustrate a neuroplasticity process whereby initial features of an expert subject evolved and became subject specific. This latter result—i.e. our inter-subject transfer learning frameworks promote acquisition of individual BCI skills—, although hypothesized by longitudinal mutual learning, is nevertheless surprising. Since BCIs are based on neurofeedback and operant conditioning (5, 38), then naïve subjects would have been expected to modulate their SMR to match the target expert's brain (EDS) patterns upon which feedback is based. However, compared to usual neurofeedback approaches and early BCI systems that use univariate analysis, our BCI works with a highly dimensional covariance space that offers subjects the possibility to modulate their individual SMR differently than the original subject. In other words, the embedded EDS subspaces of the expert and naïve subjects do not need to be similar as far as they lie on the same side of the classification hyperplane after adaptive matching. To the best of our knowledge, this is the first demonstration of this theoretical result in a BCI study.

Our results illustrate an even more radical manifestation of this property of our transfer learning approaches to evolve from an initial expert discriminant EEG pattern to individual ones. Indeed, subjects acquired BCI skills grounded on distinct EDS patterns for the bar task and car racing (Figs. S9 vs. S11 for the GR group, and Figs. S10 vs. S12 for the PAR group) despite both tasks being trained at every online session. Notwithstanding that it is customary assumed that BCI skills are transferred across tasks or enhanced by additional control dimensions, especially from basic training to control of sophisticated devices (13–15), our results illustrate that humans can learn to modulate SMR of the same MI movements distinctively for each task and do so in parallel as all the subjects practiced both tasks in every online session. It is worth noting the emergence of a more frontal EDS pattern for the car racing task, especially in the PAR framework. Frontal areas under electrodes F7 and F8 are recruited not only when performing MI of left- and right-hand movements (see, in particular, (33)) but also in working memory tasks (39, 40). Since feedback during car racing is discrete, subjects must hold relevant information about the MI process and task in working memory for longer. Emergence of a more prominent frontal EDS pattern for car racing in the PAR framework with respect to GR is probably due to the additional update of the decoder parameters to better fit the new samples at the beginning of each online session. Another plausible explanation might be that frontal activity can indicate EOG contamination of neural signals; therefore, we performed additional analyses to confirm or confute this possibility. Results indicate that BCI control was driven primarily by the activity of motor channels, and that EOG activity had a low correlation with frontal electrodes (Figs. S1–S4).

Some experiments (28, 37, 41–43) have demonstrated how online adaptation of decoder parameters increase classification performance. Yet, they only partially showed that online decoder adaptation supports BCI skill acquisition as just selected subjects were discussed in terms of their final EEG patterns. Nevertheless, other studies found that online decoder adaptation did not support subject learning even if BCI performance improved (44, 45).

Intermittent decoder recalibration using data from previous sessions also promoted subject learning (13, 15, 29), although they required longer number of training sessions than the approaches reported in this work. Animal and human studies involving implanted BCIs have also illustrated how decoder adaptation (either continuously or intermittently) leads to subject learning via neuroplasticity (16, 18, 46, 47). Like our PAR framework, these implanted BCI studies modified the decoder infrequently and required subjects to practice with a fixed decoder during long periods. Interestingly, our GR framework also yields these desirable properties despite not applying supervised decoder adaptation.

Benaroch et al. (27) also employed a Riemannian approach with adaptive matching of the feature distributions across sessions to train a tetraplegic subject to control the same BCI devices used in our study, namely, bar task and car racing game. While both studies found a significant improvement in BCI control for the bar task, Benaroch et al. did not observe improvements in the car racing game for their subject. The success of our transfer learning methods can be attributed to two distinguishing elements. First, whereas we used data from an expert, Benaroch et al. relied on the conventional acquisition of calibration data of the tetraplegic subject to build the initial decoder. Second, our transfer learning approaches enabled subjects to evolve the initial EDS pattern of the expert to individual EDS patterns that differed for the bar and car racing tasks. This was facilitated because our incremental adaptation transform was estimated independently for each task. However, Benaroch et al. attempted to carry over the discriminant patterns from the bar task to car racing using adaptive matching, which proved not to be successful. In this respect, another BCI study that also trained a different tetraplegic subject to operate the bar task and car racing game found a large shift of the EEG feature distributions associated to the two tasks for their subject (28). Another element that might have contributed to the training differences of BCI subjects in the car racing game for the two studies is that in (27) the subject had to operate a 4-class BCI, while in our case subjects trained with a 2-class BCI. It remains to be proven whether our approach scales to a multiclass BCI. Nevertheless, it is worth noting that the winners of the three Cybathlon BCI competitions in 2016 (13), 2019, and 2020 (29) used a 2-class BCI to control games that required multiple commands—in particular, the car racing competition in which Benaroch's (27) and Hehenberger's subject (28) also participated. In summary, for a practical BCI it seems more efficient to detect a limited number of classes with high accuracy than a larger number of classes with lower accuracy.

Inter-subject transfer learning

Despite being a popular topic in the BCI field, inter-subject transfer learning has mainly been explored in offline studies. Actual MI–BCI experiments—i.e. involving closed-loop feedback—are scarce (43, 44, 48). All three previous studies relied on a subject-independent binary classifier built with data from a pool of expert subjects. The parameters of the initial classifier were either adapted online as EEG samples of the individual participants were processed (43, 44) or kept fixed for the whole duration of the experiment (48). Vidaurre et al. (43) performed a single-session study, where online adaptation was supervised at the beginning and then unsupervised. Most participants, 9 out of 14, achieved a good BCI control (>70% accuracy) of a bar task at the end of the session. Perdakis et al. (44) proposed an unsupervised, context-aware adaptive method that performed similarly to a fully supervised method. Nine participants run a single session

where they directly operated a MI-BCI speller, successfully writing several words. Nevertheless, as discussed before, adaptation was detrimental for participants to learn to modulate their sensorimotor rhythms. Ray et al. (48) trained five participants over 3 days to modulate EEG rhythms associated to MI vs. happy emotional imagery based on feedback from a fixed subject-independent BCI classifier. Participants exhibited an increasing trend in BCI accuracy in a bar task, reaching better-than-random performance.

While these previous works demonstrate the feasibility of inter-subject transfer learning for BCI based on the modulation of EEG rhythmic activity, MI in particular, collecting data from multiple expert subjects is cumbersome and time-consuming. In our study, we used only one expert subject and showed how this simple subject-independent approach fosters acquisition of individual BCI skills over a longitudinal mutual learning setup, enabling participants to operate devices in synchronous as well as asynchronous realistic scenarios such as car racing games.

GR or PAR?

Supervised approaches that use labeled samples when recalibrating the BCI decoder are popular (16, 18, 41, 43, 45–47), as ML theory predicts that they would lead to better model fitting and performance—a result verified in a cross-over BCI study comparing supervised vs. fully unsupervised approaches for online decoder adaptation (44).

Our bar task results support this view (Fig. 2a–d), where subjects in the PAR group showed higher NKV scores over the training sessions. Although the difference between the two groups never reached statistical significance (Fig. 7a), the PAR framework seems to facilitate faster and better acquisition of the BCI skill. Indeed, subjects in the PAR group exhibited higher FD scores and stronger trend over their training (Figs. 6 and S13). Nevertheless, the advantage of PAR over GR in the bar task vanishes when comparing another critical element of skill acquisition, namely CL, as GR subjects delivered faster correct commands (Fig. 2e and f).

PAR, as compared to GR, also appears to be advantageous at the beginning of subjects' training in the car racing game for every performance metric—i.e. RCT, NKV, and CL—, which were insignificantly better (Figs. 3 and 4). Nevertheless, the between-group differences were minimized at the end of the training (Fig. S7b) and eventually reverted slightly for RCT. Importantly, all subjects of the GR group exhibited a decreasing trend of RCT, while this was only the case for seven subjects (out of nine) in the PAR group.

In summary, our results demonstrate that the unsupervised GR framework can yield performances comparable to those of the supervised PAR framework (see also [Supplementary material, Complementary statistical analysis](#)). Thus, we favor the GR framework over PAR during use of brain-controlled devices because the former is completely unsupervised. In fact, supervised approaches are unfeasible during activities of daily living, where ground truth labels are unavailable. Moreover, since PAR framework involves the updation of decoding parameters, its integration with other decoding methods (4, 35, 49, 50) would be computationally expensive rendering it challenging for online closed-loop BCIs.

Potential improvements and prospects

While our results demonstrate the significance of longitudinal training for improved BCI control in a complex car racing game because of subjects' enhanced sensorimotor modulations, practice also helps subjects to anticipate and timely deliver BCI commands

for optimal gameplay which is crucial for subjects to drive better (51). To disentangle these two components of practice (BCI proficiency and anticipation), it will be necessary that subjects train over longer periods of time than in our current experiment in order to observe when sensorimotor rhythms become stable, and yet race time continues to improve.

Our inter-subject transfer learning approaches could be extended along several directions. Firstly, it would be desirable to enlarge the degrees of freedom and control principles of the BCI. In particular, in our current experimental setup a major limitation is that subjects performed a binary task relying on conventional left vs. right hand MI. For robust control of more complex brain-controlled devices (4), the expert subject should be able to control a multiclass BCI with high accuracy, which has been demonstrated to be feasible (41), also in a Riemannian geometry framework (24), and to incorporate continuous control of the device (52). It remains to be demonstrated that our framework succeeds in these settings.

A second direction is to replace the supervised version of our PAR approach, which requires labeled samples, with an unsupervised version that predicts sample labels and avoids the initial recalibration run of every session. Two relevant methods are pseudo-labeling (53, 54) and context-aware learning (44, 55). Both methods have been demonstrated to perform competitively to a supervised strategy.

Although we have demonstrated the feasibility of integrating the GR framework with CSP-based decoders, it would be interesting to extend the GR framework with source-space-based approaches as we could get better information on the precise anatomical localization of regions relevant for MI decoding (4). Additionally, we plan to explore more data-driven transfer learning approaches, leveraging a database of expert subjects using deep learning techniques like EEGNet (49) and TSMNet (50).

Finally, we speculate that our inter-subject transfer learning approaches could be beneficial for BCI applications targeting people with disabilities who might be unable to modulate properly their SMR and provide good calibration data to build their initial decoder. Our proposed frameworks might pave the way to exploit pre-recorded data of healthy subjects to train patients to operate assistive and neurorehabilitation BCI-controlled devices.

Methods

Participants: Eighteen able-bodied healthy participants with normal or corrected-to-normal vision participated in the study (23.22 ± 3.59 years old, seven females). The experimental protocol was approved by the local ethics commission (2020-03-0073, The University of Texas at Austin, TX, USA). All participants provided written informed consent before conducting the experiment, and were compensated for their participation to the study.

EEG and EOG recordings: Physiological activity was recorded from 22 EEG and 3 electrooculogram (EOG) channels at 512 Hz (ANT Neuro eego mylab with bipolar EOG box, waveguard EEG cap). EEG electrodes were positioned at F7, F3, Fz, F4, F8, FC5, FC1, FC2, FC6, C3, Cz, C4, CP5, CP1, CP2, CP6, P7, P3, Pz, P4, P8, and POz locations according to the international 10–10 montage, with the ground and reference electrodes placed at AFz and CPz, respectively. The three EOG electrodes were placed on the eye canthi and the forehead. Participants were asked to minimize eye and facial movements during the BCI trials. However, whenever an artifact was detected (see “Artifact detection” below), the output of BCI was blocked. EEG signals were processed as described in “Feature extraction” below.

BCI training: Subjects were instructed to mentally rehearse the kinesthetics, not the visualization, of a movement without overtly causing any contraction of their muscles. Training was based on a binary-class BCI that differentiated MI of left-hand and right-hand movements. Subjects started with an offline calibration session in which they were pseudo-randomly cued to perform MI of either of the two movement classes. The data collected from this session was later used for benchmarking our proposed inter-subject transfer learning frameworks against the standard calibration-based subject-specific approach. In the following consecutive five online sessions, subjects received feedback on their BCI performance. Each online session comprised several runs consisting of multiple MI trials in which subjects practiced two tasks, namely bar task followed by car racing. Subjects performed four runs of the bar task and played at least four car races, each race corresponding to a run, in the first online session. Afterwards, in the second to fifth online sessions, subjects practiced three runs of the bar task and at least five car races (except for one subject due to technical problems with the EEG amplifier) to emphasize and assess whether they could operate their BCI in a more realistic scenario.

Bar task: This paradigm is a standard feedback method to train users to operate MI BCIs (10, 11). A run consists of 10 trials of each MI class (left and right hand) in a pseudo-random order. In our study, each trial started with a fixation cross which lasted for 1 s. Next, the cue corresponding to a MI class was presented for 1.5 s. Following the cue, online visual feedback using a 1D bar was presented for a maximum of 7 s during the online runs and for 5 s during the offline MI runs. Each trial ended with the presentation of the result of the trial for 2 s, followed by an inter-trial rest of 1.5 s. During the offline runs, the bar always moved towards the correct side. In contrast, during the online runs, the bar moved both ways proportionally to the accumulated evidence of the BCI for each MI class (see “BCI controller” below).

Car racing: The car racing game was adopted from the International Cybathlon competition (25). The game was adjusted for a two-class BCI such that the race track was divided into pseudo-random patches of right or left wide turns. The car would accelerate through the patches only if the MI of the corresponding movement (right or left) was successfully performed—i.e. correct command delivery. In the event of an incorrect command delivery, the car hit the curb and slowed down. The state of the car (high or low speed if a correct or incorrect command was delivered, respectively) was reset after each patch, and subjects received discrete feedback whenever a correct command was delivered through the right and left headlights of the car.

Feature extraction: We use covariance matrices of EEG signals to characterize the spatio-spectral components in bandpass filtered EEG. Covariance matrices are commonly used as features or to estimate spatial filters for building machine learning models that discriminate MI classes (35). In our processing pipeline, we bandpass filter the EEG signals using a second-order Butterworth filter in [8, 30] Hz. We then estimate features on 1-s sliding windows (referred to as samples) with 1/16 s of step size. Let $X \in \mathbb{R}^{N_c \times N_t}$ denote an EEG sample where N_c is the number of channels and N_t is the number of time points. Its sample wise normalized covariance C , which is used as a feature in the Riemannian geometry framework, is expressed as:

$$C_i = \frac{XX^T}{\text{trace}(XX^T)}. \quad (1)$$

We used a shrinkage-based covariance estimator for computing the covariance features C_i (Ledoit and Wolf method (56)) in order to avoid any numerical problems.

Riemannian geometry framework: Within the Riemannian geometry framework, a spatial covariance matrix (1) of EEG signals is a feature on a Riemannian manifold. The latter is a smooth, differentiable manifold with each point in its domain having a tangent space that is a finite dimensional Euclidean space. Being symmetric positive definite (SPD), the covariance matrices lie on a Riemannian manifold (30). This allows building a Riemannian geometry classifier as a simple, distance-based classifier that uses the notion of distances and means over the manifold for classification.

The parametric equation of shortest path (called *geodesic* for curved spaces) between two matrices C_1 and C_2 on the Riemannian manifold can be written as (30):

$$\gamma(C_1, C_2, t) = C_1^{\frac{1}{2}}(C_1^{-\frac{1}{2}}C_2C_1^{-\frac{1}{2}})^t C_1^{\frac{1}{2}} \quad t \in [0, 1]. \quad (2)$$

Using the geodesic definition in (2), the distance between two matrices on the manifold can be estimated as:

$$\delta_r(C_1, C_2) = \|\log(C_1^{-\frac{1}{2}}C_2C_1^{-\frac{1}{2}})\|_{\mathcal{F}}. \quad (3)$$

Interestingly, the above-mentioned distance metric is invariant to affine transformations, also commonly known as Affine Invariant Riemannian Metric (31). Let W be an affine invariant transformation, then the Riemannian distance between affine transformed covariances and their original raw covariances is preserved:

$$\delta_r(C_1, C_2) = \delta_r(W^T C_1 W, W^T C_2 W). \quad (4)$$

Riemannian mean: It is simply the notion of center of mass of covariances constrained on a Riemannian manifold. Similar to the estimation of the Euclidean mean, where the sum of squared Euclidean distances from the Euclidean mean estimator is minimized, the center of mass/karcher mean in a Riemannian framework is estimated as:

$$\bar{C} = \arg \min_{C \in \mathcal{P}_n} \sum_{i=1}^N \delta_r^2(C_i, C), \quad (5)$$

where \mathcal{P}_n is the space of n -dimensional SPD matrices. To solve the optimization problem in (5), we relied on the approach discussed in Ref. (30).

MDM classifier: The BCI decoder corresponds to a minimum distance to mean (MDM) Riemannian geometry classifier, which is similar to a distance based nearest neighbor approach. The following steps are used for training and testing:

- *Training:* The covariance class prototypes (\bar{C}_k , $k \in (1, 2)$) of MI classes are estimated using the Riemannian mean of the expert subject’s covariance samples corresponding to each class from Eq. 5 (see “Domain adaptation” below).
- *Prediction:* For an incoming EEG sample, its spatial covariance (C_i) is estimated using Eq. 1 and classified according to:

$$k_i = \arg \min_{k \in (1, 2)} \delta_r(C_i, \bar{C}_k). \quad (6)$$

- *Prediction probability:* For an incoming EEG sample (C_i), we predict the class probability for both classes according to:

$$p_k = \frac{e^{-(\delta_r(C_i, \bar{C}_k))^2}}{e^{-(\delta_r(C_i, \bar{C}_1))^2} + e^{-(\delta_r(C_i, \bar{C}_2))^2}} \quad k \in [1, 2]. \quad (7)$$

Domain adaptation: From a data-driven machine learning perspective, the proposed inter-subject transfer learning framework is a classic scenario of domain adaptation. Simply put, the data

of expert and naïve subjects are two different domains. To build a robust classifier, we attempt to match the distributions of the naïve subject to that of the expert subject. Using the affine invariant property of the Riemannian distance in Eq. 9, we transform the covariances of the expert subject and the naïve subject to a common space. Since this transform is an affine transformation, it preserves the distances between all the covariances (22).

For the expert subject E , the affine invariant recentering transform T^{train} is estimated using (5) on all her/his available covariance samples irrespective of the class. Then, we train an MDM decoder on the transformed covariance samples.

For naïve subjects, we use online incremental domain adaptation in the spirit of the offline method proposed in (31) to estimate the recentering transform. For online classification of an incoming sample C_i^s for naïve subject s , first we estimate the incremental recentering affine transformation T_i from (8), which builds on (2):

$$T_i^{\text{test}} = \begin{cases} \gamma(T_{i-1}^{\text{test}}, C_i^s, \frac{1}{i+1}) & i > 1 \\ C_1^s & i = 1 \end{cases}. \quad (8)$$

Next, we use the transformation matrix T_i^{test} to convert the sample C_i^s to \tilde{C}_i^s such that transformed covariance matches the space of expert subject using (9):

$$\tilde{C}_i^s = T_i^{\text{test}^{-1}} C_i^s T_i^{\text{test}}. \quad (9)$$

Finally, the affine transformed covariance sample is then fed to the MDM decoder built with data from the expert subject. The affine transformation (9) has been shown previously to be similar to a first-order mean matching (31) and hence reduces the drift in the data between the training and test distributions (24). Figure S14 depicts a pictorial overview of the affine transformation procedure.

Generic Recentering (GR) framework: In the GR framework, we match the distributions of the expert and naïve subjects using the domain adaptation technique. For naïve subjects, we set their sample counter $i = 1$ in the first training run on each day for each task (bar or car racing) separately. The counter i is incremented by 1 as new samples are recorded. The reason why we have separated counters for each task is that we hypothesize that subjects will modulate sensorimotor rhythms differently for each task because of their dissimilar feedback (richer and discrete in the case of the car racing). In the case of the bar task, the sample counter is carried over to all the runs of an online session. However, for the car racing task, the counter is reset to 1 at the beginning of each race so that race times can be compared across an online session. Indeed, if the sample counter i were not reset, subjects could finish later races faster than earlier races simply because the transformation matrix T_i^{test} is better estimated as using more samples.

Personally Adjusted Recentering (PAR) framework: In contrast to the GR framework where we keep the decoding model fixed in the recentered space and only update the recentering transform incrementally, in the PAR framework, after recentering and predicting an incoming sample, we also update the parameters of the decoding model—i.e. class prototypes—, according to the ground truth label of the incoming sample.

The adaptation scheme to update the class prototype is governed by the following geodesic interpolation:

$$\tilde{C}_k = \gamma(\tilde{C}_k, \tilde{C}_i^s, \eta) \quad i_{\text{ground}} = k \quad (10)$$

We kept the adaptation parameter $\eta = 0.001$ to be constant. The choice for this value was based on a pseudo-online simulation of the PAR framework using existing in-house BCI datasets. We

simulated the classification performances for different values of η on a logarithmic scale on datasets of two experiments where four subjects, different from those who participated in this study, operated the bar task using MI of left and right hand movements (Fig. S15). We observe that for both the datasets, $\eta = 0.001$ yields better performance than other values of the learning rate. This observation is in line with literature suggesting mid-range learning rates to be optimal in adaptation frameworks (57).

In each session, we do supervised updates of class prototypes on the first run only. The class prototypes are then kept fixed throughout the remaining runs of the online session. Note that the supervised adaptation of the class prototypes is applied in the recentered space—i.e. after the recentering step of the GR framework, where the sample counter i is reset to 1 at the beginning of each session for the bar task and at the start of each car race.

BCI controller: The BCI decoder estimates the probability distribution of an EEG sample to belong to each MI class. The BCI processes a 1-s sliding window of EEG signals every 62.5 ms. To deliver a command (Right, Left), the BCI accumulates the estimated probabilities over time for each class until the accumulated evidence for one of the classes reaches a configurable threshold. If none of the thresholds is crossed within a certain timeout period, the trial is considered rejected. We use an exponential smoothing approach to accumulate evidence for all classes (11):

$$\text{Prob}_i = 0.95 \cdot \text{Prob}_{i-1} + 0.05 \cdot p_i, \quad (11)$$

where Prob_i is the accumulated probability for a given class and P_i is the posterior probability of that class for the i th sample.

The feedback provided by the bar is synchronous in the sense that the accumulated evidence for the two MI classes are reset to uniform probability distribution at the beginning of each trial. The thresholds for the bar task were set after observing the bar dynamics in an online mock run before the first online run on each session. The thresholds for consecutive bar runs were adjusted so that subjects operating the BCI could maximize the number of correct command deliveries and minimize incorrect commands or timeouts.

For car races, the probabilities were integrated asynchronously until the accumulated evidence of one class crosses its corresponding threshold within the timeout period. If one of the thresholds is reached, a UDP message is sent to the game interface to execute the corresponding command and the accumulated evidence for the two MI classes are reset to uniform distribution for subsequent command delivery. In the case of timeout, the accumulated evidences are also reset to uniform distribution, but no command is sent to the game. Also, we set a refractory period of 1 s for evidence accumulation between consecutive commands or timeouts. Notably, evidence accumulation is not reset after the car exits a patch and enters the next one. The thresholds for consecutive car races were adjusted in a similar fashion to the bar runs.

For both the bar and car racing runs, the BCI controller blocks the evidence accumulation process in two scenarios:

- Presence of an artifact: The EEG sample window contains an EOG or facial artifact.
- Uncertainty of prediction: The highest prediction probability of either of the classes does not exceed a preset minimum (0.55).

Artifact detection: Our BCI uses a simple artifact detection criterion for ocular and facial movements based on the analysis of three EOG electrodes placed on the eye canthi (EOG₁, EOG₂) and the forehead (EOG₃) while being referenced by an electrode on the mastoid. The EOG signals were synchronously recorded with

EEG using a bipolar box through the eego amplifier and filtered to [1, 10] Hz using a second-order Butterworth filter. The presence of an artifact is detected whenever the values resulting in EOG channels due to horizontal or vertical eye movements, exceed a pre-configured threshold. Also, the average of all three EOG channels is monitored to check for facial contractions and eye movements (13). In the event that an artifact is detected over a 1-s EEG sample, the output of the classifier is disregarded and the evidence accumulation is paused for that sample.

Adaptive recentering in common spatial pattern: Building upon previous research (58, 59), we have integrated the GR framework into the classical CSP method and evaluated it in a pseudo-online setting. Details of this GR-CSP framework are as follows:

1. Let $X \in \mathbb{R}^{N_c \times N_t}$ denote an EEG sample where N_c is the number of channels and N_t is the number of time points. Its sample-wise normalized covariance C_i , which is used as a feature in the Riemannian geometry framework, is estimated from Eq. 1. Note that to mimic the actual experiment we used a shrinkage-based covariance estimator for computing the covariance features C_i .
2. To alleviate the problem of nonstationarities (across sessions and subjects), we transform all the offline covariances using the Riemannian mean of offline covariances through an affine transformation (Eq. 9).
3. In the transformed space, we estimate the mean covariance of both MI classes by computing the Euclidean mean of the transformed covariances corresponding to each MI class, denoted as Σ_1 and Σ_2 for Class 1 and Class 2, respectively.
4. The spatial filter matrix, denoted as W having dimensions $N_c \times N_c$, is then estimated using the eigenvalue decomposition of $\Sigma_1^{-1}\Sigma_2$, where N_c represents the number of channels. We selected a total of $K=6$ spatial filters (three from each side of the eigenspectrum) to create the spatial filter matrix $\bar{W} = [\bar{w}_1, \bar{w}_2, \dots, \bar{w}_6] \in \mathbb{R}^{6 \times N_c}$.
5. Since the GR framework operates directly on covariance matrices, we perform feature extraction in the covariance space. For each trial i , we use its transformed covariance C_i^{trans} to estimate the variance of the spatially filtered signal k as $\text{var}(k) = \bar{w}_i C_i^{\text{trans}} \bar{w}_i^T$.
6. Following the approach of Zhang et al. (60), for trial i , we create a log-normalized feature vector of spatially filtered transformed covariances, denoted as f_i . This feature vector is estimated as:

$$f_i = \left[\frac{\text{var}_1}{\sum_{k=1}^6 \text{var}_k}, \frac{\text{var}_2}{\sum_{k=1}^6 \text{var}_k}, \dots, \frac{\text{var}_6}{\sum_{k=1}^6 \text{var}_k} \right] \quad (12)$$

This feature vector is then used to build a linear discriminant classifier.

7. During the inference stage in pseudo-online simulations, for an incoming online sample i^{on} , we first estimate its trace-normalized covariance using Eq. 1. Then, the trace-normalized covariance is transformed using the incremental affine transform (Eq. 9). Subsequently, we estimate a feature vector, denoted as f_i^{on} , which is fed into the linear discriminant classifier for prediction.

Performance metrics: We have computed a number of metrics to exhaustively characterize the performance of the proposed inter-subject transfer learning frameworks.

- *Sample level classification:* We predict class label on each of the 1-s covariance samples and compute the confusion matrix using the ground truth labels and predicted labels. Note that artifact contaminated/rejected samples (see “BCI controller”) are not included in the computation of the confusion matrix. From the confusion matrix, we estimate the kappa value as a metric of classification performance (26). A kappa value of 1 represents a perfect classification and a kappa value of 0 corresponds to chance level performance.
- *Command level classification:* This metric represents the accuracy of delivered BCI commands (after evidence accumulation) that subjects experience as a feedback.

We create a confusion matrix using all trials in a session that finished with a correct or incorrect command delivery—i.e. timeout trials are not included in the computation of the confusion matrix—and estimate the Cohen’s kappa. Then, we normalize the kappa value (NKV) to incorporate timeout trials as a penalty: $\text{Normalized-Kappa} = \text{kappa} * (1 - \frac{n_{\text{timeouts}}}{\text{total-trials}})$.
- *Race completion time (RCT):* Time taken by the subject to finish one race run.
- *Command latency (CL):* We use this metric to quantify the latency of correct command delivery. The time taken to deliver a correct command in a trial in the bar task is estimated as the time difference between the start of visual feedback and correct command delivery. In the car racing game, for each patch we take the time difference between the start of that patch and the moment the subject delivered the correct command for that patch. Then, CL is the average of the correct command delivery times across all the trials (bar task) or patches (car racing) in an online sessions.
- *Electrode discriminancy score (EDS):* We use a backward elimination approach to estimate the contribution of each of the electrodes to the separability of the two classes in the Riemannian manifold. Assume \bar{C}_1^{22} , \bar{C}_2^{22} are the prototypes of the two classes in the 22D feature space. We estimate the prediction probability P_{22} of the class prototypes using (7). We then remove the i th row and column of the two prototype covariance matrices corresponding to channel i and compute the prediction probability in this 21D space as P_{21} . Finally we estimate the discriminant score e_i corresponding to electrode i as $e_i = P_{22} - P_{21}$. We interpret this discriminant score as the increase in posterior prediction probability after the inclusion of channel i when moving from the 21D space to the 22D space. This backward elimination is repeated independently for each channel to estimate their EDS. Of note, EDS values are computed on the unmatched feature distributions for subjects—including the expert.
- *Feature distinctiveness (FD):* In the spirit of (34, 61), we used the distinctiveness metric between the unmatched feature distribution of the two MI classes to quantify the feature space discriminancy (or separability) between the two MI classes. FD utilizes the Riemannian distance between the class prototypes, normalized to the variance of the feature distributions of both classes.

Statistical analyses: To compare the statistical difference at the beginning and at the end of the longitudinal training for each of the two inter-subject transfer learning frameworks, we used a paired t-test paired across subjects. For comparison across the different frameworks we utilized an unpaired t-test on each of the online sessions. To further scrutinize the trend within each of

the frameworks, we performed correlation analysis through linear mixed effect modeling (LME). We formulated the model as $\text{metric} - 1 + (1 | \text{subjects}) + \text{group}$. Here metric is the response variable, group is fixed effect and subjects account for random effects. The used metric , group pairs for tracking longitudinal learning within each of the proposed frameworks are {Bar NKV, Race NKV, RCT, FD} and {online session 1, 2, 3, 4, 5}, respectively. Normality in each of the statistical analyses was assessed with a Lilliefors test. When the normality of the residuals was violated for LME modeling, we first applied a logarithmic transformation to the response variable ({Bar NKV, Race NKV, RCT, FD}) and then used LME again. In cases where, after the log transformation, normality of the residuals still remained violated, we instead estimated the correlation coefficient without fitting the LME model—Pearson (parametric), or Spearman (nonparametric) depending on whether the normality of the response variable was satisfied or violated, respectively.

Supplementary Material

Supplementary material is available at PNAS Nexus online.

Funding

This work was partially supported by the Coleman Fung Foundation and the Sinclair Foundation.

Author Contributions

S.K., H.A., and J.d.R.M. conceived and designed the experimental protocols. S.K., H.A., and R.F. implemented the experimental protocols. S.K. was responsible for data acquisition. S.K. and J.d.R.M. performed the analysis. S.K., H.A., F.S.R., and J.d.R.M. interpreted results and prepared the draft manuscript. All authors reviewed the results and approved the final version of the manuscript.

Data Availability

There was no public dataset used in this study. The acquired EEG data for different subjects would be anonymized and released upon publication (<https://zenodo.org/records/10694880>). Codes used for analysis will be provided in a public GitHub repository upon publication (<https://github.com/neurosatya/TransLearning/>).

References

- Ramos-Murguialday A, et al. 2013. Brain-machine interface in chronic stroke rehabilitation: a controlled study. *Ann Neurol*. 74(1):100–108.
- Biasiucci A, et al. 2018. Brain-actuated functional electrical stimulation elicits lasting arm motor recovery after stroke. *Nat Commun*. 9(1):1–13.
- Leeb R, et al. 2015. Towards independence: a BCI telepresence robot for people with severe motor disabilities. *Proc IEEE*. 103(6): 969–982.
- Edelman BJ, et al. 2019. Noninvasive neuroimaging enhances continuous neural tracking for robotic device control. *Sci Robot*. 4(31):eaaw6844.
- Birbaumer N, et al. 1999. A spelling device for the paralysed. *Nature*. 398(6725):297–298.
- Perdikis S, et al. 2014. Clinical evaluation of BrainTree, a motor imagery hybrid BCI speller. *J Neural Eng*. 11(3):036003.
- Leeb R, et al. 2007. Brain-computer communication: motivation, aim, and impact of exploring a virtual apartment. *IEEE Trans Neural Syst Rehabil Eng*. 15(4):473–482.
- Velasco-Álvarez F, Ron-Angevin R, Blanca-Mena MJ. 2010. Free virtual navigation using motor imagery through an asynchronous brain-computer interface. *Presence*. 19(1):71–81.
- Pfurtscheller G, Lopes da Silva FH. 1999. Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clin Neurophysiol*. 110(11):1842–1857.
- Pfurtscheller G, Neuper C. 2001. Motor imagery and direct brain-computer communication. *Proc IEEE*. 89(7):1123–1134.
- Leeb R, et al. 2013. Transferring brain-computer interfaces beyond the laboratory: successful application control for motor-disabled users. *Artif Intell Med*. 59(2):121–132.
- Ganguly K, Carmena JM. 2009. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol*. 7(7):e1000153.
- Perdikis S, Tonin L, Saeedi S, Schneider C, Millán JdR. 2018. The Cybathlon BCI race: successful longitudinal mutual learning with two tetraplegic users. *PLoS Biol*. 16(5):e2003787.
- Silversmith DB, et al. 2021. Plug-and-play control of a brain-computer interface through neural map stabilization. *Nat Biotechnol*. 39(3):326–335.
- Tonin L, et al. 2022. Learning to control a BMI-driven wheelchair for people with severe tetraplegia. *iScience*. 25(12):105418.
- Collinger JL, et al. 2013. High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet*. 381(9866):557–564.
- Lotte F, Larrue F, Mühl C. 2013. Flaws in current human training protocols for spontaneous brain-computer interfaces: lessons learned from instructional design. *Front Hum Neurosci*. 7:568.
- Orsborn AL, et al. 2014. Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron*. 82(6):1380–1393.
- Millán JdR. 2015. Brain-machine interfaces: the perception-action closed loop: a two-learner system. *IEEE Syst Man Cybern Mag*. 1(1):6–8.
- Chavarriaga R, Fried-Oken M, Kleih S, Lotte F, Scherer R. 2017. Heading for new shores! Overcoming pitfalls in BCI design. *Brain-Comput Interfaces*. 4(1-2):60–73.
- Perdikis S, Millán JdR. 2020. Brain-machine interfaces: a tale of two learners. *IEEE Syst Man Cybern Mag*. 6(3):12–19.
- Zanini P, Congedo M, Jutten C, Said S, Berthoumieu Y. 2017. Transfer learning: a Riemannian geometry framework with applications to brain-computer interfaces. *IEEE Trans Biomed Eng*. 65(5):1107–1116.
- Neyshabur B, Sedghi H, Zhang C. 2020. What is being transferred in transfer learning? *Adv Neural Inf Process Syst*. 33:512–523.
- Rodrigues PLC, Jutten C, Congedo M. 2018. Riemannian procrustes analysis: transfer learning for brain-computer interfaces. *IEEE Trans Biomed Eng*. 66(8):2390–2401.
- ETH Zurich and Inert Coin. 2020. BrainDriver, BCI Race of Cybathlon 2020. Zurich, Switzerland.
- Cohen J. 1960. A coefficient of agreement for nominal scales. *Educ Psychol Meas*. 20(1):37–46.
- Benaroch C, et al. 2021. Long-term BCI training of a tetraplegic user: adaptive Riemannian classifiers and user training. *Front Hum Neurosci*. 15:118.
- Hehenberger L, et al. 2021. Long-term mutual training for the Cybathlon BCI race with a tetraplegic pilot: a case study on inter-session transfer and intra-session adaptation. *Front Hum Neurosci*. 15:70.
- Tortora S, et al. 2022. Neural correlates of user learning during long-term BCI training for the Cybathlon competition. *J Neuroeng Rehabil*. 19:69.

- 30 Barachant A, Bonnet S, Congedo M, Jutten C. 2010. Riemannian geometry applied to BCI classification. In: International conference on Latent Variable Analysis and Signal Separation; St. Malo, France. Springer. p. 629–636.
- 31 Kumar S, Yger F, Lotte F. 2019. Towards adaptive classification using Riemannian geometry approaches in brain-computer interfaces. In: 7th International winter conference on Brain-Computer Interface; Gangwon-do, South Korea. IEEE. p. 1–6.
- 32 Bougrain L, Rimbart S, Rodrigues PLC, Canron G, Lotte F. 2021. Guidelines to use transfer learning for motor imagery detection: an experimental study. In: 10th International IEEE/EMBS conference on Neural Engineering; Virtual. IEEE. p. 5–8.
- 33 Pfurtscheller G, Pregenzer M, Neuper C. 1994. Visualization of sensorimotor areas involved in preparation for hand movement based on classification of μ and central β rhythms in single EEG trials in man. *Neurosci Lett*. 181(1-2):43–46.
- 34 Lotte F, Jeunet C. 2018. Defining and quantifying users' mental imagery-based BCI skills: a first step. *J Neural Eng*. 15(4):046030.
- 35 Blankertz B, Tomioka R, Lemm S, Kawanabe M, Müller K-R. 2007. Optimizing spatial filters for robust EEG single-trial analysis. *IEEE Signal Process Mag*. 25(1):41–56.
- 36 Lotte F, Congedo M, Lécuyer A, Lamarche F, Arnaldi B. 2007. A review of classification algorithms for eeg-based brain-computer interfaces. *J Neural Eng*. 4(2):R1.
- 37 Wolpaw JR, McFarland DJ. 2004. Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proc Natl Acad Sci USA*. 101(51):17849–17854.
- 38 Sitaram R, et al. 2017. Closed-loop brain training: the science of neurofeedback. *Nat Rev Neurosci*. 18(2):86–100.
- 39 Sauseng P, et al. 2004. Theta coupling in the human electroencephalogram during a working memory task. *Neurosci Lett*. 354(2):123–126.
- 40 Roberts BM, Hsieh L-T, Ranganath C. 2013. Oscillatory activity during maintenance of spatial and temporal information in working memory. *Neuropsychologia*. 51(2):349–357.
- 41 McFarland DJ, Sarnacki WA, Wolpaw JR. 2010. Electroencephalographic (EEG) control of three-dimensional movement. *J Neural Eng*. 7(3):036007.
- 42 Vidaurre C, Kawanabe M, von Büna P, Blankertz B, Müller K-R. 2010. Toward unsupervised adaptation of LDA for brain-computer interfaces. *IEEE Trans Biomed Eng*. 58(3):587–597.
- 43 Vidaurre C, Sannelli C, Müller K-R, Blankertz B. 2011. Co-adaptive calibration to improve BCI efficiency. *J Neural Eng*. 8(2):025009.
- 44 Perdikis S, Leeb R, Millán JdR. 2016. Context-aware adaptive spelling in motor imagery BCI. *J Neural Eng*. 13(3):036018.
- 45 Cunha JD, Perdikis S, Halder S, Scherer R. 2021. Post-adaptation effects in a motor imagery brain-computer interface online co-adaptive paradigm. *IEEE Access*. 9:41688–41703.
- 46 Taylor DM, Tillery SIH, Schwartz AB. 2002. Direct cortical control of 3D neuroprosthetic devices. *Science*. 296(5574):1829–1832.
- 47 Gilja V, et al. 2012. A high-performance neural prosthesis enabled by control algorithm design. *Nat Neurosci*. 15(12):1752–1757.
- 48 Ray AM, et al. 2015. A subject-independent pattern-based brain-computer interface. *Front Behav Neurosci*. 9:269.
- 49 Lawhern VJ, et al. 2018. EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces. *J Neural Eng*. 15(5):056013.
- 50 Kobler R, Hirayama J-i, Zhao Q, Kawanabe M. 2022. SPD domain-specific batch normalization to crack interpretable unsupervised domain adaptation in EEG. *Adv Neural Inf Process Syst*. 35:6219–6235.
- 51 Galán F, et al. 2008. A brain-actuated wheelchair: asynchronous and non-invasive brain-computer interfaces for continuous control of robots. *Clin Neurophysiol*. 119(9):2159–2169.
- 52 Tonin L, Bauer FC, Millán JdR. 2020. The role of the control framework for continuous teleoperation of a brain-machine interface-driven mobile robot. *IEEE Trans Robot*. 36(1):78–91.
- 53 Lee D-H, et al. 2013. Pseudo-label: the simple and efficient semi-supervised learning method for deep neural networks. In: ICML workshop on Challenges in Representation Learning; Atlanta, GA. p. 1–6.
- 54 Arazo E, Ortego D, Albert P, O'Connor NE, McGuinness K. 2020. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In: 2020 International joint conference on Neural Networks; Glasgow, UK. IEEE. p. 1–8.
- 55 Perdikis S, Leeb R, Chavarriaga R, Millán JdR. 2020. Context-aware learning for generative models. *IEEE Trans Neural Netw Learn Syst*. 32(8):3471–3483.
- 56 Ledoit O, Wolf M. 2004. A well-conditioned estimator for large-dimensional covariance matrices. *J Multivar Anal*. 88(2):365–411.
- 57 Müller JS, et al. 2017. A mathematical model for the two-learners problem. *J Neural Eng*. 14(3):036005.
- 58 He H, Wu D. 2019. Transfer learning for brain-computer interfaces: a Euclidean space data alignment approach. *IEEE Trans Biomed Eng*. 67(2):399–410.
- 59 Barachant A, Bonnet S, Congedo M, Jutten C. 2010. Common spatial pattern revisited by Riemannian geometry. In: 2010 IEEE international workshop on Multimedia Signal Processing; St. Malo, France. IEEE. p. 472–476.
- 60 Zhang Y, Zhou G, Jin J, Wang X, Cichocki A. 2015. Optimizing spatial patterns with sparse filter bands for motor-imagery based brain-computer interface. *J Neurosci Methods*. 255:85–91.
- 61 Mouriño J, Chiappa S, Jané R, Millán JdR. 2002. Evolution of the mental states operating a brain-computer interface. In: 2nd European Medical and Biological Engineering Conference; Vienna, Austria. Technischen Universität Graz. Vol. 3. p. 600–601.