Clinical Neurophysiology 158 (2024) 212–217

Contents lists available at ScienceDirect

Clinical Neurophysiology

journal homepage: www.elsevier.com/locate/clinph

Opinion Paper

# Towards real-time EEG–TMS modulation of brain state in a closed-loop approach

Dania Humaidan [a,b,1], Jiahua Xu [a,b,1], Miriam Kirchhoff [a,b], Gian Luca Romani [e], Risto J. Ilmoniemi [c,d], Ulf Ziemann [a,b,*]

[a] Department of Neurology & Stroke, University of Tübingen, Tübingen, Germany
[b] Hertie-Institute for Clinical Brain Research, Tübingen, Germany
[c] Department of Neuroscience and Biomedical Engineering, Aalto University School of Science, Espoo, Finland
[d] BioMag Laboratory, HUS Medical Imaging Center, University of Helsinki, Aalto University and Helsinki University Hospital, Helsinki, Finland
[e] Institute for Advanced Biomedical Technologies, University of Chieti-Pescara, Chieti, Italy

## 1. From neuroscience-inspired AI to AI-based personalized treatment

Artificial intelligence (AI) has been defined as the theory and algorithms that enable machines to learn, reason, and solve problems, also in the way human beings do so. Neuroscience has always had a reciprocal information flow with the field of AI. On the one hand, most AI algorithms are inspired by the human brain. The field of neuroscience has not only provided a source of inspiration for AI algorithms, but it also presented the possibility to validate certain algorithms if they were found to be used in the brain. On the other hand, AI algorithms introduced revolutionary transitions in the field of neuroscience. An important example is the highly efficient and accurate analysis of neuroimage datasets. However, a huge contribution of AI in neuroscience lies in the field of reinforcement learning (RL). This field has been inspired by animal learning, and as the name reveals, involves learning the best behavior that achieves the desired reward by reinforcing the actions that lead to higher rewards as feedback from the surrounding environment (Hassabis et al., 2017).

While RL has been extensively used for research purposes to advance our understanding in the field of neuroscience, it has many potential applications in the fields of medical neuroscience and computational neurology (Maia and Frank, 2011). After decades of traditional therapy, scientists have realized that the intra-individual differences in terms of environmental, biological, and psychosocial factors should be investigated, a concept that has been termed "precision medicine". Furthermore, these differences should be taken into account when new treatment is planned, so that it would be tailored according to the specific characteristics of the patient receiving it, a direction called "personalized medicine" (Carbonara et al., 2022). This movement has so far targeted traditional therapy, including drug and non-drug treatments. However, there is a need to extend this concept to novel therapeutic techniques such as transcranial magnetic stimulation (TMS).

TMS has emerged as a safe and painless neuromodulation technique that utilizes the electromagnetic field to induce spatiotemporal dynamics of cortical excitability in the brain (Siebner et al., 2022). Recently, TMS has been combined with electroencephalography (EEG) (Hernandez-Pavon et al. 2023), which detects the electrical neuronal population activity in the brain with electrodes on the scalp, to enable delivery of TMS pulses on targeted brain regions triggered by ongoing neuronal oscillations, called brain state-dependent stimulation, and precise modulation of brain activity (Zrenner et al., 2018).

The journey towards "personalized TMS" should begin with a clear definition of the required elements. While it is challenging to identify all aspects, one possible approach is to start from an identification of the limitations encountered with the current TMS approaches. For example, single-site TMS limiting modulation of distributed neuronal networks, and open-loop TMS uninformed of the ongoing neuronal activity in the stimulated network limiting precise temporal targeting of brain states that would be particularly receptive for causing a desired modification of this network. Hence, we envisage that the following aspects would constitute four important pillars of future TMS applications in research and therapy:

1. Technology to move beyond single-channel TMS towards multi-node brain stimulation at the network level by designing multi-locus TMS (mTMS) transducers (Nieminen et al., 2022; Sinisalo et al., 2024).
2. Algorithms for real-time analysis of EEG data to determine the desired brain state for triggering TMS or mTMS (e.g., phase and/or power of ongoing brain oscillations, brain network connectivity measures) (Marzetti et al., 2023).
3. Methods to modulate the brain state based on precise individual characteristics in an adaptive closed-loop manner.

* Corresponding author at: Department of Neurology & Stroke, and Hertie Institute for Clinical Brain Research, Eberhard Karls University Tübingen, Hoppe-Seyler-Straße 3, 72076 Tübingen, Germany.
*E-mail address:* ulf.ziemann@uni-tuebingen.de (U. Ziemann).

[1] Contributed equally.

https://doi.org/10.1016/j.clinph.2023.12.006

4. Adaptation of points 1, 2 and 3 in clinical practice, for example to support motor recovery in stroke patients (Rösch et al., 2024).

In this paper, we discuss the third aspect towards achieving personalized TMS. We propose a method to close the loop in TMS applications, by adapting the stimulation parameters on the fly based on the neurophysiological feedback from the person receiving TMS. We then illustrate this concept using a proof-of-concept experiment to show the practical application of a novel closed-loop real-time TMS setup.

## 2. From open-loop, off-line TMS to reinforcement learning-based, closed-loop, real-time TMS

In the conventional application of TMS, a coil is placed in contact with the subject's head over the targeted region, and a predefined pulse is initiated without considering the ongoing endogenous neural oscillations. This approach, known as open-loop TMS (Bergmann, 2018), typically relies on brain activity solely for post-experiment data analysis. Meanwhile, brain state-dependent TMS, which takes into account the ongoing activity measured by real-time EEG, is still considered open-loop TMS (Fig. 1**A**). A closed-loop brain state-dependent stimulation paradigm has been proposed to overcome the limitations of ignoring the modulation outcome (Bergmann, 2018). The key differentiation between closed-loop and open-loop (EEG–TMS) systems is grounded in the direct utilization of output data (psychological or physiological) from the previous pulse, either during offline learning (Fig. 1**B**), or in conjunction with real-time monitoring of the brain state using EEG, subsequent to deliver the next pulse (Fig. 1**C**).

In a real-time machine learning-based approach, an initial set of parameters (e.g., stimulation parameter or specific brain state) define the first TMS pulses or their timing. Data collected during the ongoing stimulation give feedback to an inference model that optimizes, using machine learning algorithms, the stimulation parameters based on predefined functions or thresholds (e.g., desired brain state). Then, the output (readout of the resulting effect of the previous TMS pulse) is fed back into the stimulator to close the loop for the next round of stimulation; iterations continue until the predefined functions or thresholds are satisfied (Fig. 1**C**). This real-time inference mode-based approach has been reported by a previous study, which utilized beta and alpha power as the brain state to trigger transcranial direct current stimulation in a closed-loop manner (Leite et al., 2017). Several studies have explored the possibility of a real-time machine learning-based approach in a closed-loop TMS protocol, such as automatically optimizing the pulse amplitude and width (Alavi et al., 2023), stimulation location for the largest MEP amplitude (Tervo et al., 2020), automated hotspot hunting for the lowest motor threshold (Meincke et al., 2016) and TMS-induced current orientation for the largest TMS/EEG responses (Tervo et al., 2022). Up to now, no study has been reported to specifically target the modulation of network-based brain states in a real-time closed-loop EEG–TMS manner.

## 3. A proof-of-concept experiment on real-time closed-loop EEG–TMS

We will illustrate here an experiment that includes repetitive TMS (rTMS) of a two-node brain network and determination of the optimized value of a stimulation parameter through closed-loop EEG-informed real-time reinforcement learning.

### 3.1. Reinforcement learning (RL) algorithm

The concept of reinforcement learning involves training an algorithm (*Agent*) to make decisions (*Actions*) that will affect the surrounding medium (*Environment*). Then, based on the results of these decisions (*Observations*), a certain value is calculated and provided as feedback to the algorithm (*Reward*).

The equivalent elements in our experiment are as follows:

1) The target network to be stimulated (*Environment*):

The target network is composed of the supplementary motor area (SMA) and the primary motor cortex (M1) in a conditioning–test stimulus protocol, such that the conditioning stimulus is applied over the SMA-proper (stimulation intensity 140% of active motor threshold over M1) and the test stimulus over the hand area of M1 (stimulation intensity to elicit MEPs of on average 1 mV in peak-to-peak amplitude). In previous experiments, we and others have performed similar paired-pulse stimulation of the SMA–M1 network and demonstrated paired-pulse SMA-to-M1 facilitation as a measure of effective connectivity between these two nodes when an interstimulus interval of 6 ms was used (Arai et al., 2012; Arai et al., 2011; Neige et al., 2023). The location of M1 (hotspot) stimulation was manually searched as the point consistently resulting in largest MEP amplitudes. To determine the best location for SMA stimulation, a second coil was placed around 4 cm anterior to the vertex (Cz according to the 10–20 International EEG System) and the spot resulting in strongest SMA-to-M1 facilitation was determined. At each site of SMA stimulation, 10 paired-pulse trials with stimulation of SMA-M1 and 10 single-pulse trials with stimulation of M1 only were obtained, and SMA-to-M1 facilitation was calculated by the ratio of the mean MEP amplitude resulting from paired-pulse stimulation of SMA-M1 over the mean MEP amplitude resulting from single-pulse stimulation of M1. The location, angulation and orientation of the two coils were monitored with a neuronavigation system (Localite GmbH, Sankt Augustin, Germany) throughout the experiment. Whenever either of the two coils slipped away by >5 mm, the experiment was briefly halted and the coil was returned onto its original target position.

2) The stimulation parameter to be optimized (*Action*):

We used rTMS, which is capable of inducing aftereffects in the human brain that outlast the period of stimulation. These aftereffects are thought to reflect plasticity processes in the human brain (Cooke and Bliss, 2006; Hoogendam et al., 2010; Ziemann et al., 2008). Different phases of human EEG oscillations represent different excitability states of the brain (Buzsaki and Draguhn, 2004). Triggering rTMS at specific EEG-informed excitability states leads to differential plasticity-like effects (Baur et al., 2020; Gordon et al., 2022; Zrenner et al., 2018), with high potential to enhance the effects of rTMS-induced plasticity. The phase of the endogenous sensorimotor μ-rhythm is the stimulation parameter that we tried to adaptively optimize. We defined eight discrete phase bins that were sampled from $[-\pi, \pi]$ as follows: $[-\frac{7}{8}\pi, -\frac{5}{8}\pi, -\frac{3}{8}\pi, -\frac{1}{8}\pi, \frac{1}{8}\pi, \frac{3}{8}\pi, \frac{5}{8}\pi, \frac{7}{8}\pi]$.

3) The primary output (*Observation*):

We aimed at maximizing SMA-to-M1 facilitation, as measured by the ratio of the MEP amplitude resulting from paired-pulse SMA–M1 stimulation over the MEP amplitude resulting from single-pulse stimulation of M1.

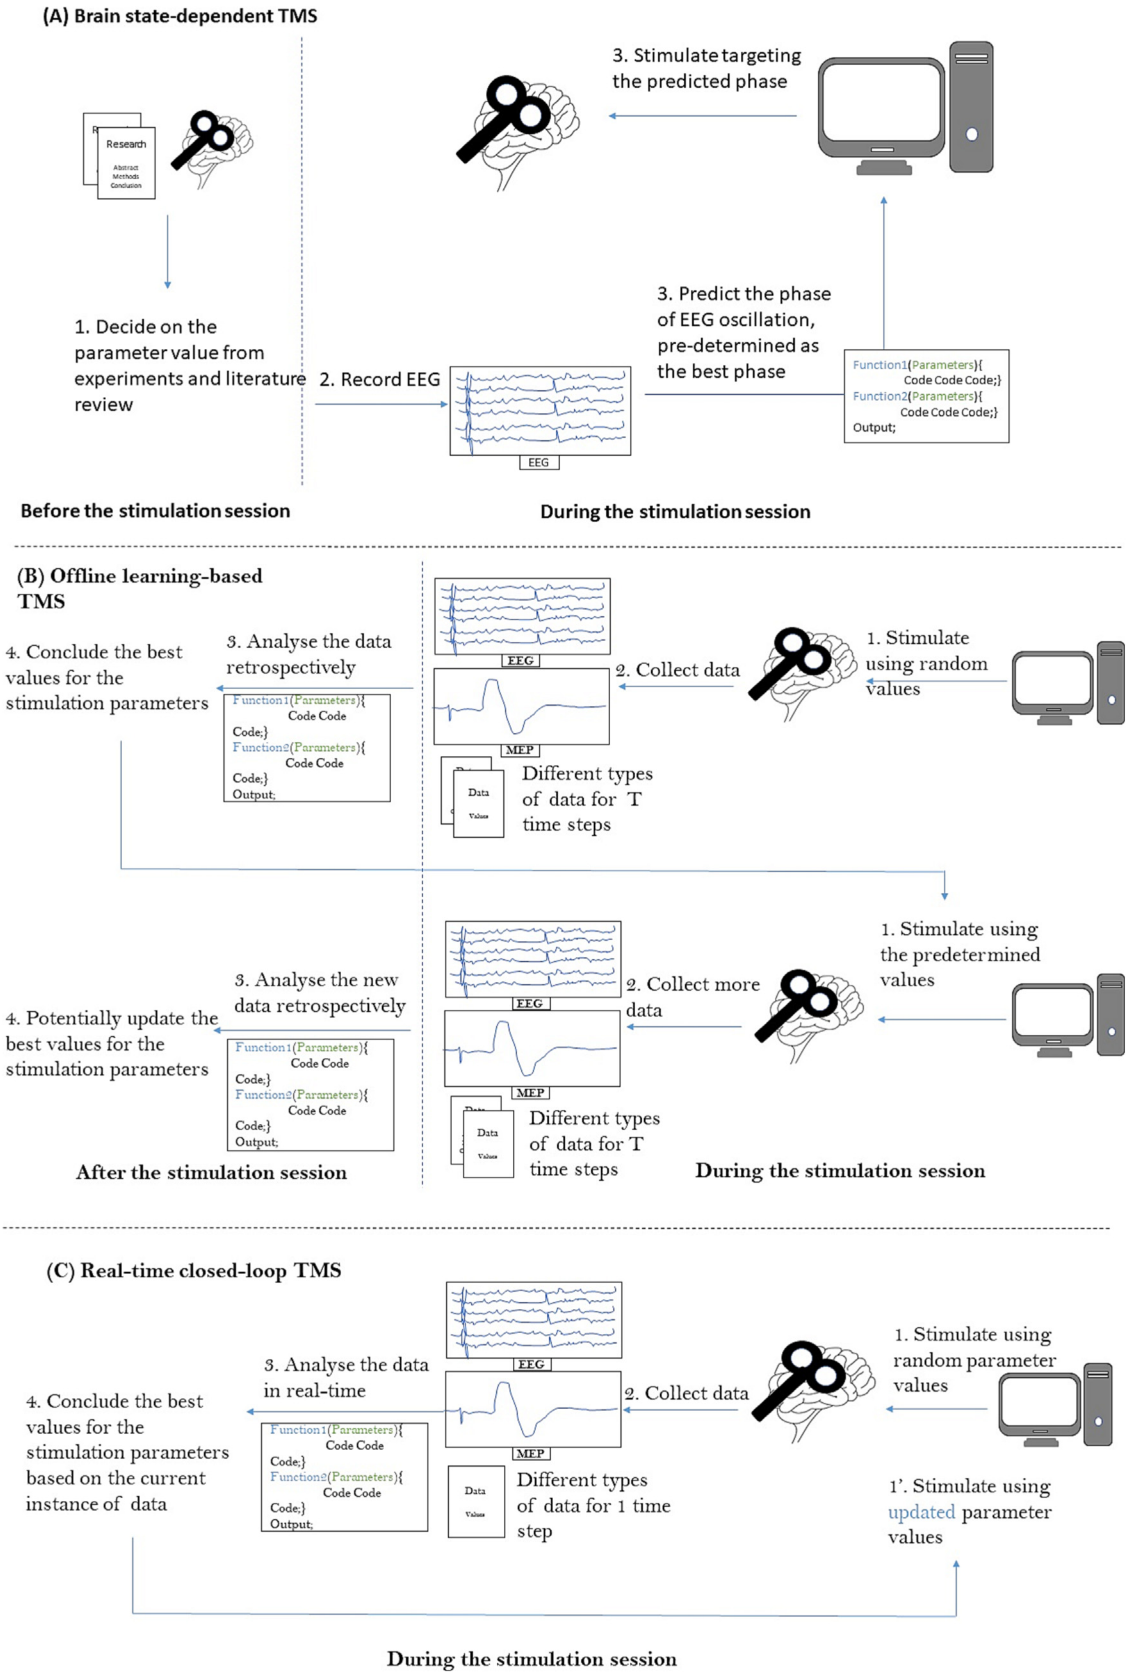


**Fig. 1.** A schematic illustration of open-loop brain state-dependent repetitive transcranial magnetic stimulation (rTMS) (**A**), offline learning-based brain state-dependent rTMS (**B**) and closed-loop real-time brain state-dependent rTMS (**C**). TMS: transcranial magnetic stimulation, EEG: electroencephalography, MEP: motor evoked potential.

4) The optimization algorithm (*Algorithm*):

We employed an RL algorithm according to the method of Deep-Q Learning (DQN) (Watkins and Dayan, 1992). This algorithm implements a deep neural network to approximate the query function under the framework of reinforcement learning. The agent in DQN learns by providing rewards/penalties for the actions taken via interacting with the surrounding environment. The agent operates by selecting decisions for the next actions to achieve a higher cumulative reward.

5) The feedback to the algorithm (*Reward*):

At each time step, the agent receives an observation, which is the SMA-to-M1 facilitation (see above, for definition). This value is then used to calculate the reward of choosing this action. We arbitrarily set a target facilitation effect of 1.5 times the baseline SMA-to-M1 facilitation (i.e., prior to learning), and defined the reward function as follows:

$$Reward = current\,SMA - to - M1\,facilitation \; - 1.5 * baseline\,SMA - to - M1\,facilitation$$

Based on this reward value, the agent decides on the next action, i.e., picking a μ-rhythm phase bin out of the eight options to trigger in the next pulse to increase the reward.

### *3.2. Running the experiment*

The experiment was approved by the local ethics committee of the medical faculty of the University of Tübingen (Project Number 525/2021BO2); subjects including those of pilot sessions were tested only after written informed consent was provided. Scalp EEG was recorded from a TMS compatible 64-channel Ag/AgCl sintered ring electrode cap (EasyCap GmbH, Germany), and surface EMG was recorded through bipolar electromyography (EMG) adhesive hydrogel electrodes (Kendall, Covidien) over the first dorsal interosseus muscle of the right hand. TMS of the left SMA proper was delivered using a Cool-B35 HO coil connected to a MagVenture MagPro X100 stimulator including MagOption with a monophasic current waveform, while the test pulse over the left M1 was delivered through a Cool-B35 HO coil connected to a MagVenture MagPro R30 stimulator with a biphasic current waveform. The site of M1 stimulation was determined according to standard hot spot search methods, with the coil oriented 45° away from the midline to induce a posterior-to-anterior current in M1 for the second phase of the biphasic stimulus (Groppa et al., 2012). The site of SMA stimulation was determined with the coil oriented so that the monophasic stimulus was oriented right-to-left towards the targeted left-hemispheric SMA (Arai et al., 2012) and by mapping several sites for the optimal SMA-M1 MEP facilitation. The experiment started by initiating the training of the RL agent (Reinforcement Learning Toolbox, MATLAB R2021a), which initially picked a random phase of the sensorimotor μ-rhythm to trigger the next pulse. Phase targeting was enabled by EEG real-time analysis and forward prediction using the Brain Oscillation State Sensor (BOSS) device (sync2brain, Germany). At each learning step of the RL agent, 2 paired pulses to SMA-M1 were triggered, and every 4 steps, 2 single pulses to M1 were triggered. The inter-trial interval (ITI) was set to 2–3 s. Then, the SMA-to-M1 facilitation was obtained for this step and used to calculate the reward value. This was then provided to the RL agent. As iterations continued, improved actions that induced a higher SMA-to-M1 facilitation resulted. At the end of the learning, the agent kept targeting the phase yielding the highest SMA-to-M1 facilitation. The agent was trained over a period of approximately 1 hour for ten episodes in total, with each episode containing 40 steps (i.e., 10 episodes × 40 steps × 2 trials = 800 trials for training). The BOSS device monitored the ongoing EEG activity and enable triggering of the magnetic stimulators only, if a predetermined power threshold of the sensorimotor μ-rhythm was exceeded (Zrenner et al., 2018).

We tested the pipeline on simulated data, where targeting a specific phase would yield the highest SMA-to-M1 facilitation. We found that the algorithm is able to find the optimal phase. On the simulated data, we compared two RL algorithms, a simple Q-learning table and a deep Q-learning agent. We found that the simple Q-learning table agent increased the reward faster, by more quickly identifying the phase associated with highest SMA-to-M1 facilitation, while the deep Q-learning agent needed a longer learning time while still managing to identify the best state (Humaidan et al., 2021). The results are shown in Fig. 2**A.**

Then, we ran the DQN algorithm on real data in an experiment on one subject, as detailed above. Based on our experience with simulated data, we updated the RL parameters when applying the algorithm on real data. The used values are listed in Table 1.

We demonstrated that, similar to the simulation, the algorithm also learned on real data within 5 episodes (i.e., 400 trials) and then maintained the performance over the remainder of the learning period (Fig. 2**B-C**).

## 4. Conclusion, discussion, and perspective

From efficient analysis of neuroimaging datasets to accurate analysis of EEG signals to predict epileptic seizures, artificial intelligence has revolutionized theoretical and applied neuroscience. This conceptual paper aims to introduce a vision of artificial intelligence-based personalized TMS. Using a real-time closed-loop EEG–TMS setup, TMS can be most efficient by adaptively assigning the stimulation parameters to achieve the desired maximum effect for the specific treated individual at the current conditions. We illustrated this vision through a practical application of a truly adaptive closed-loop EEG–TMS setup to enhance effective connectivity between two brain areas in healthy individuals. This was achieved by utilizing reinforcement learning algorithms to train an agent that could automatically select the optimal phase of an ongoing local brain oscillation to induce a higher facilitation effect between the two nodes in the targeted brain network. This closed-loop approach allows tailoring the stimulation parameters based on the specific feedback received from each subject towards achieving the most optimal effect as defined by the research question. This will address inter-individual and even intra-individual variability as major limitations of current open-loop TMS applications (Hamada et al., 2013). Also, a general model for a given individual can be saved and used for later TMS sessions, but may need to be fine-tuned in subsequent sessions of the same individual, as there might be significant inter-session differences due to modifiable factors such as time of day, prior activity, or alertness at the time of the TMS experiment (Ridding and Ziemann, 2010).

The real-time approach allows us to overcome the need for large amounts of pre-training data and reduces the effect of priors from the professionals. However, an extension of the proposed proof-of-concept closed-loop TMS experiment necessitates addressing further limitations. On the one hand, optimizing the stimulation parameters on the fly within the limited training time means that the inference model will be using only limited data, which can result in instability and difficulties in learning with complicated algorithms. In addition, in a truly closed-loop system, data are processed and managed in real-time scenarios, which poses significant challenges given current software and hardware limitations. For instance, high-dimensional EEG data must be ana-

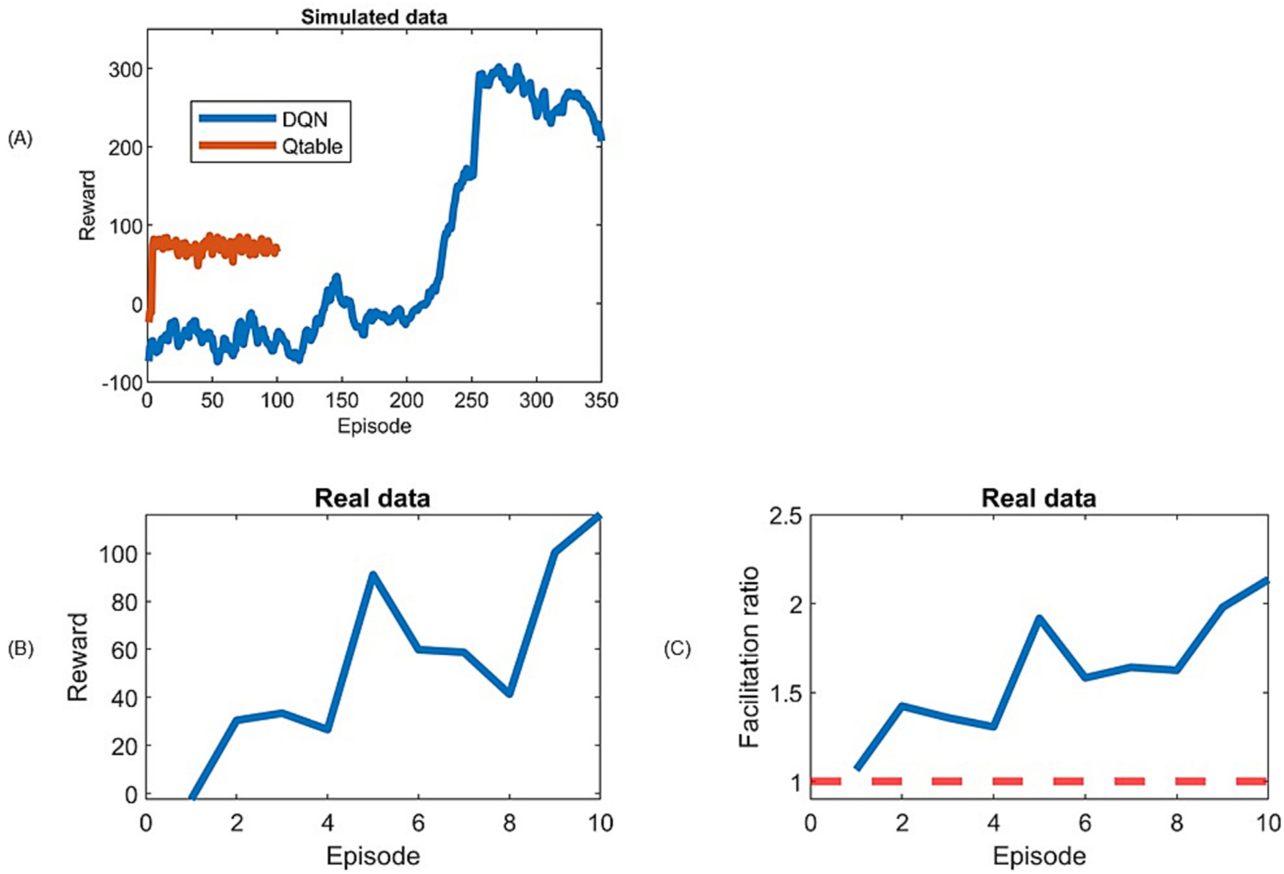


**Fig. 2.** (**A**) The learning curves of DQN (blue) and simple Q-table (red) algorithms on simulated data. Please note that 350 episodes were used for learning. (**B**) The learning curve of the agent on real data from one subject. (**C**) The change in the SMA-to-M1 facilitation in the course of the learning. Note that the targeted increase of a 1.5-fold increase in SMA-to-M1 facilitation compared to baseline was reached and then maintained already after 5 episodes. DQN: Deep-Q Network, SMA: supplementary motor area, M1: primary motor cortex.

**Table 1**
The parameter values used in the reinforcement learning (RL) algorithm.

| Parameter | Value |
|---|---|
| Deep-Q Network (DQN) | Three fully connected layers (neurons: 12, 8, 8) and one Long Short-term Memory (LSTM) unit (10 neurons) (Hochreiter and Schmidhuber, 1997) |
| Activation function | Rectified linear unit function (ReLU) |
| Learning rate | 0.001 |
| MiniBatchSize | 32 |
| ExperienceBufferLength | 50 |
| EpsilonDecay | 0.005 |

lyzed in real time to extract the desired brain state parameters at the network level. Moreover, the high-dimensional EEG data demand complex tools for denoising, and some methodologies, such as source reconstruction and functional connectivity, require high computational resources that can only be provided by hardware support such as AI chips and GPUs (Graphic Processing Units). On the other hand, the tens of milliseconds duration of brain state status changes leave only a narrow time window for preparing all the necessary parameters for the next round of stimulation. In clinical practice, clinicians may find retrospective data analysis for optimization of stimulation parameters more intuitive and explainable rather than black-box style updates within the adaptive closed-loop system. Lastly, an easy-to-use graphical user interface and an efficient API (Action Programming Interface) are indispensable in clinical settings. Despite various issues in this application, the implementation of this paradigm has gained extreme attention and continues to evolve towards providing precise, personalized, and intelligent noninvasive brain stimulation therapies in the treatment of neurological disorders.

## Disclosures

RI has patents and patent applications on TMS technology. All other authors have no conflicts of interest.

## Acknowledgements

This study is part of the ConnectToBrain project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 810377). This concept paper should be considered in conjunction with the other three concept papers illustrating the ongoing activities that are carried out and that, altogether, constitute the overall investigation of the ConnectToBrain project.

## References

Alavi SMM, Vila-Rodriguez F, Mahdi A, Goetz SM. Closed-loop optimal and automatic tuning of pulse amplitude and width in EMG-guided controllable transcranial magnetic stimulation. Biomed Eng Lett 2023;13:119–27.
Arai N, Lu MK, Ugawa Y, Ziemann U. Effective connectivity between human supplementary motor area and primary motor cortex: a paired-coil TMS study. Exp Brain Res 2012;220:79–87.

Arai N, Müller-Dahlhaus F, Murakami T, Bliem B, Lu MK, Ugawa Y, et al. State-Dependent and timing-dependent bidirectional associative plasticity in the human SMA-M1 network. J Neurosci 2011;31:15376–83.

Baur D, Galevska D, Hussain S, Cohen LG, Ziemann U, Zrenner C. Induction of LTD-like corticospinal plasticity by low-frequency rTMS depends on pre-stimulus phase of sensorimotor μ-rhythm. Brain Stimul 2020;13:1580–7.

Bergmann TO. Brain state-dependent brain stimulation. Front Psychol 2018;9:2108.

Buzsaki G, Draguhn A. Neuronal oscillations in cortical networks. Science 2004;304:1926–9.

Carbonara K, MacNeil AJ, O'Leary DD, Coorssen JR. Profit versus Quality: the enigma of scientific wellness. J Pers Med 2022:12.

Cooke SF, Bliss TV. Plasticity in the human central nervous system. Brain 2006;129:1659–73.

Gordon PC, Belardinelli P, Stenroos M, Ziemann U, Zrenner C. Prefrontal theta phase-dependent rTMS-induced plasticity of cortical and behavioral responses in human cortex. Brain Stimul 2022;15:391–402.

Groppa S, Oliviero A, Eisen A, Quartarone A, Cohen LG, Mall V, et al. A practical guide to diagnostic transcranial magnetic stimulation: Report of an IFCN committee. Clin Neurophysiol 2012;123:858–82.

Hamada M, Murase N, Hasan A, Balaratnam M, Rothwell JC. The role of interneuron networks in driving human motor cortical plasticity. Cereb Cortex 2013;23:1593–605.

Hassabis D, Kumaran D, Summerfield C, Botvinick M. Neuroscience-inspired artificial intelligence. Neuron 2017;95:245–58.

Hernandez-Pavon JC, Veniero D, Bergmann TO, Belardinelli P, Bortoletto M, Casarotto S, et al. TMS combined with EEG: recommendations and open issues for data collection and analysis. Brain Stimul 2023;16:567–93.

Hochreiter S, Schmidhuber J. Long short-term memory. Neural Comput 1997;9:1735–80.

Hoogendam JM, Ramakers GM, Di Lazzaro V. Physiology of repetitive transcranial magnetic stimulation of the human brain. Brain Stimul 2010;3:95–118.

Humaidan D, Vetter DE, Metsomaa J, Ermolova M, Ziemann U. Reinforcement machine learning for closed-loop rTMS stimulation of brain networks. Brain Stimul 2021;16:1696.

Leite J, Morales-Quezada L, Carvalho S, Thibaut A, Doruk D, Chen CF, et al. Surface EEG-transcranial direct current stimulation (tDCS) closed-loop system. Int J Neural Syst 2017;1750026.

Maia TV, Frank MJ. From reinforcement learning models to psychiatric and neurological disorders. Nat Neurosci 2011;14:154–62.

Marzetti L, Makkinayeri S, Pieramico G, Guidotti R, D'Andrea A, Roine T, et al. Towards real-time identification of large-scale brain states for improved brain state-dependent stimulation. Clin Neurophysiol. 2024;158:196–203.

Meincke J, Hewitt M, Batsikadze G, Liebetanz D. Automated TMS hotspot-hunting using a closed loop threshold-based algorithm. NeuroImage 2016;124:509–17.

Neige C, Vassiliadis P, Ali Zazou A, Dricot L, Lebon F, Brees T, et al. Connecting the dots: harnessing dual-site transcranial magnetic stimulation to quantify the causal influence of medial frontal areas on the motor cortex. Cereb Cortex 2023. https://doi.org/10.1093/cercor/bhad370.

Nieminen JO, Sinisalo H, Souza VH, Malmi M, Yuryev M, Tervo AE, et al. Multi-locus transcranial magnetic stimulation system for electronically targeted brain stimulation. Brain Stimul 2022;15:116–24.

Ridding MC, Ziemann U. Determinants of the induction of cortical plasticity by non-invasive brain stimulation in healthy subjects. J Physiol 2010;588:2291–304.

Rösch J, Vetter DE, Baldassarre A, Souza VH, Lioumis P, Roine T, et al. Individualized treatment of motor stroke: a perspective on open-loop, closed-loop and adaptive closed-loop brain state-dependent TMS. Clin Neurophysiol. 2024;158:204–11.

Siebner HR, Funke K, Aberra AS, Antal A, Bestmann S, Chen R, et al. Transcranial magnetic stimulation of the brain: What is stimulated? - A consensus and critical position paper. Clin Neurophysiol 2022;140:59–97.

Sinisalo H, Rissanen I, Kahilakoski O-P, Souza VH, Tommila T, Laine M, et al. Modulating brain networks in space and time: multi-locus transcranial magnetic stimulation. Clin Neurophysiol 2024;158:218–224.

Tervo AE, Metsomaa J, Nieminen JO, Sarvas J, Ilmoniemi RJ. Automated search of stimulation targets with closed-loop transcranial magnetic stimulation. NeuroImage 2020;220 117082.

Tervo AE, Nieminen JO, Lioumis P, Metsomaa J, Souza VH, Sinisalo H, et al. Closed-loop optimization of transcranial magnetic stimulation with electroencephalography feedback. Brain Stimul 2022;15:523–31.

Watkins CJCH, Dayan P. Q-learning. Mach Learn 1992;8:279–92.

Ziemann U, Paulus W, Nitsche MA, Pascual-Leone A, Byblow WD, Berardelli A, et al. Consensus: motor cortex plasticity protocols. Brain Stimul 2008;1:164–82.

Zrenner C, Desideri D, Belardinelli P, Ziemann U. Real-time EEG-defined excitability states determine efficacy of TMS-induced plasticity in human motor cortex. Brain Stimul 2018;11:374–89.