Active Inference as a Framework for Brain-Computer Interfaces

Syed Hussain Ather*

**Institute of Medical Science, University of Toronto, Toronto, ON M5S*

Corresponding Author:

Syed Hussain Ather,

e-mail: s.ather@mail.utoronto.ca

tel: +13177011491

Active Inference as a Framework for Brain-Computer Interfaces

Syed Hussain Ather

Institute of Medical Sciences, University of Toronto
Krembil Centre for Neuroinformatics, Centre for Addiction and Mental
Health, Toronto
`hussain.ather@camh.ca`

Abstract

As Karl Friston explained during the International Symposium on Artificial Intelligence and Brain Science 2020, active inference provides a way of using abstract rule-learning and approximate Bayesian inference to show how minimizing (expected) free energy leads to active sampling of novel contingencies. Friston elaborated how there were ways of making an optimal decision using active inference that can offer perspectives to advances in artificial intelligence. These methods of optimization within the context of active inference can also be used as a framework for improving brain-computer interfaces (BCI). This way, BCIs can give rise to artificial curiosity in the way Friston had described during his session. Using Friston's free energy principle, we can optimize the criterion a BCI uses to infer the intentions of the user from EEG observations. Under Friston's criteria for making an optimal decision, BCIs can expand their framework of optimal decision-making using active inference.

Key words: Artificial intelligence, Neuroscience, Free energy principle, Active inference, Computational Neuroscience, Mathematics

1. Introduction

1.1. Free energy principle

In the search for grand unified theories of the brain, the free energy principle states that a self-organizing principle at equilibrium with its environ-

5 ment must minimize the amount of free energy that it has. Its applica-
6 tions have spanned far and wide in explaining brain structure and function.
7 This self-organizing behavior arises from the defining characteristic of bio-
8 logical systems to resist disorder in dynamic environments[1] stemming from
9 Helmholtz’s ideas of perception[2]. A model of perceptual influence and learn-
10 ing built upon these ideas can explain how principles can resolve problems of
11 the inference of the causes underlying sensory input and learning the causal
12 structure that generates them. From there, one can study how inference and
13 learning follow. Friston has used this in the context of Empirical Bayes and
14 hierarchical models of sensory input to show how the free energy principle
15 can understand a range of cortical organization and responses.

16 The free energy principle provides a way for adaptive systems to unify
17 action, perception, and learning[3]. The theory and background of the prin-
18 ciple involves a system using a Markov blanket to minimize the difference
19 between a model of the world and its sense and the perception associated
20 with it. Through continuously correcting the world model of the system, the
21 system changes the world into an expected state while minimizing the free
22 energy of the system. Using the Bayesian idea of the brain as an “inference
23 engine,” the system can actively change the world (active inference) into the
24 expected state and minimize the free energy of the system. This holds true
25 for a wide variety of adaptive systems from animals to brains themselves
26 in understanding mental disorders and artificial intelligence, and other ap-
27 plications of the free energy principle span areas of exploration and novelty
28 seeking[4].

29 Friston outlined the motivation behind using the free energy principle as
30 a unified brain theory using the system’s tendency to resist disorder. When
31 a system resists its tendency to move towards disorder, the physiological
32 and sensory states of a system move towards configurations of low entropy.
33 Given that the number of these states is limited, the system is very likely to
34 be in these states of low entropy. By using a formulation of entropy as the
35 average amount of self-information or “surprise” (the negative log-probability
36 of a specific outcome), Friston explained how biological agents minimize the
37 long-term average of surprise (or maximize sensory evidence for an agent’s
38 existence) to keep sensory entropy low. By sampling the environment to
39 change configuration and minimize free energy this way, the system changes
40 its expectations[5]. This forms the basis of action and perception, and the
41 system’s state and structure encode an implicit and probabilistic model of
42 the environment. The nervous system in particular maintains order through

these methods, and the specific structural and functional organization is maintained by the environment’s causal structure.

One can evaluate free energy as a function of two things to which the agent has access: the sensory states and a recognition density encoded by its internal states (such as neuronal activity and connection strengths)[6]. In terms of the environment’s causal structure, the recognition density would be a probabilistic representation of what caused a particular sensation. The causes underlying sensory input, or the probabilistic representation of what causes sensations, are used as the recognition density. They can vary from an object in one’s field of vision or blood pressure changing the physiological state of organs.

1.2. Active Inference

A corollary of the free energy principle, active inference, arises from the way natural agents act in the context of these observations. Active inference claims natural agents act to fulfill prior beliefs about preferred observations. By adjusting sensory data (without changing recognition density) to minimize free energy, an agent chooses the sensory inputs from a sample based on prior expectations to increase the accuracy, the surprise about sensations expected under a particular recognition density, of an agent’s predictions. In the context of Bayesian inference, one may define the complexity (“Bayesian surprise”) as the difference between the prior density, which encodes beliefs about the state of the world before sensory data are assimilated, and posterior beliefs, encoded by the recognition density. In essence, the agent avoids surprising states by making active inferences.

As Friston explained during his session[7], active inference is self-evidencing in that action and perception can be cast as maximising Bayesian model evidence for the generative models of the world[8]. Using a generative model of the underlying causal structure of a system, you can explain how evidence accumulates or how a specific action was chosen. Examples of active inference for Markov decision processes include using Bayes optimal precision to predict activity in dopaminergic areas [9] and using a gradient descent on variational free energy to simulate neuronal processing [10]. This evidence can be decomposed through its accuracy and complexity.

One may even describe active inference as a method of explaining action using the idea that the brain has “stubborn predictions” that are resistant to change, such as adaptive body temperature necessary for survival, that cause the system to behave in a way to cause the predictions to come true[11].

Figuring out the etiology of stubbornness would provide insight into ways of how to change these predictions helpful for understanding the relationship between drugs and psychotherapy that have synergistic effects when used together. Other applications of active inference extend to visual foraging[12] and BCIs[13].

1.3. Active Inference Python Implementation

A good example of a Python implementation of Active Inference can be found in the “Active Inference for Markov Decision Processes” repository [14].

Using the `grid_world` environment that has 3 x 3 states, an agent can choose one of five actions for each time step. The agent can move in any of the four directions (up, down, left, right) or remain where it is. After performing inferences about its current location, the agent must figure out which location it prefers to move to and move there. Using a reward value, we can set the reward value to reward the user when choosing the direction leading to a preferred location.

```
REWARD_LOCATION = 7
env = GridWorldEnv()
env.set_reward_state(REWARD_LOCATION)
```

In it, the “Inference and action example” illustrates the process of creating a generative model consisting of a likelihood distribution A and an empirical prior distribution B for the environment.

```
likelihood_matrix = env.get_likelihood_matrix()
A = Categorical(values=likelihood_matrix)
A.remove_zeros()
plot_likelihood(A)

transition_matrix = env.get_transition_matrix()
B = Categorical(values=transition_matrix)
B.remove_zeros()
plot_empirical_prior(B)
```

With the beliefs about the current location:

```
Qs = Categorical(dims=[env.n_states])
```

113 The beliefs are then plotted from the agent's preferences:

```
114 C[REWARD_LOCATION] = 1.
115 plot_beliefs(C, title="Prior preference (C)")
```

116 By evaluating the policy, one can calculate negative expected free energy
117 for a policy:

```
118 def evaluate_policy(policy, Qs, A, B, C):
119     # store expected free energy
120     G = 0
121
122     # create copy of our state
123     Qs = Qs.copy()
124
125     # loop over policy
126     for t in range(len(policy)):
127
128         # get action
129         u = int(policy[t])
130
131         # work out expected state
132         Qs = B[u].dot(Qs)
133
134         # work out expected observations
135         Qo = A.dot(Qs)
136
137         # get entropy
138         H = A.entropy()
139
140         # get predicted divergence and uncertainty and novelty
141         divergence = F.kl_divergence(Qo, C)
142         uncertainty = H.dot(Qs)[0, 0]
143
144         G += (divergence + uncertainty)
145     return -G
```

146 And infer an action:

```

147 def infer_action(Qs, A, B, C, n_actions, policy_len):
148
149     # this function generates all possible combinations of policies
150     policies = F.generate_policies(n_actions, policy_len)
151     n_policies = len(policies)
152
153     # negative expected free energy
154     neg_G = np.zeros([n_policies, 1])
155
156     for i, policy in enumerate(policies):
157         neg_G[i] = evaluate_policy(policy, Qs, A, B, C)
158
159     # get distribution over policies
160     Q_pi = F.softmax(neg_G)
161
162     # probabilities of control states
163     Qu = Categorical(dims=n_actions)
164
165     # sum probabilities of controls
166     for i, policy in enumerate(policies):
167         # control state specified by policy
168         u = int(policy[0])
169         # add probability of policy
170         Qu[u] += Q_pi[i]
171
172     # normalize
173     Qu.normalize()
174
175     # sample control
176     u = Qu.sample()
177
178     return u

```

179 Putting everything together, we can write a loop to infer beliefs about
180 hidden states:

```

181 # number of time steps
182 T = 10

```

```

183
184 # number of actions
185 n_actions = env.n_actions
186
187 # length of policies we consider
188 policy_len = 4
189
190 # set initial state
191 env.set_initial_state(0)
192
193 # reset environment
194 o = env.reset()
195
196 # infer initial state
197 Qs = F.softmax(A[o, :].log())
198
199 # loop over time
200 for t in range(T):
201
202     # random action
203     a = infer_action(Qs, A, B, C, n_actions, policy_len)
204
205     # perform action
206     o, r = env.step(a)
207
208     # infer new hidden state
209     Qs = F.softmax(A[o, :].log() + B[a].dot(Qs).log())
210
211     # print information
212     print("Time step {}".format(t))
213     env.render()
214     plot_beliefs(Qs, "Beliefs (Qs) at time {}".format(t))

```

215 A straightforward method such as this one could be applied to BCIs to
 216 let them use active inference in optimizing their decision-making.

217 *1.4. Artificial Curiosity*

218 For artificially intelligent agents to look at their environments and learn
 219 from them in a way that a human being would, they develop a kind of artifi-

220 cial curiosity. They can examine and understand which behavior is rewarded
 221 based on the outcomes of their actions. Schmidhuber put forward a simple
 222 formal theory of fun and intrinsic motivation based on maximizing intrinsic
 223 reward for active creation or discovery of novel, surprising patterns[15].
 224 In this sense, artificially curious agents learn to become bored or tired of
 225 predictable patterns or behaviors.

226 Friston’s method of using the free energy principle and predictive coding,
 227 the method by which the brain generates and updates a mental model of the
 228 environment given sensory input[16][17], doesn’t achieve this, Schmidhuber
 229 wrote. By visiting highly predictable states, Friston argued that perception
 230 seeks to suppress prediction error by changing predictions and action changes
 231 the signals themselves. Instead of learning, Friston’s approach only teaches
 232 agents to stabilize and make things predictable. Other methods of using
 233 the free energy principle in active inference have included variational Bayes,
 234 formal account explaining the relationship between posterior expectations of
 235 hidden states, control states and precision[18]. As Friston explained during
 236 his session, artificial curiosity falls under a different form of optimality that
 237 uses Hamilton’s principle of least action.

238 **2. Body**

239 *2.1. Brain-Computer Interfaces*

240 The BCI is a system that lets a human brain interact directly with an
 241 external device for restoring movement [19] or communication[20], for as-
 242 sisting and optimizing tasks[21], and for rehabilitating by enabling the self-
 243 regulation of brain activity for therapeutic purposes[22]. The BCI works by
 244 collecting information through electroencephalography (EEG) to update an
 245 internal model of the user, optimizing interactions via inference about the
 246 user, and optimizing interactions by acting upon the user. These behaviors
 247 parallel that of an Active Inference system. With Active Inference, the BCI
 248 can choose desired outcomes based on the noisy, variable EEG input signal
 249 through optimization. By minimizing entropy to anticipate future outcomes,
 250 one can use an Active Inference framework to model sequential learning and
 251 decision-making in the brain guided by an objective function to minimize
 252 free energy when adjusting to signal variability during the signal processing
 253 pipeline and modeling and acting in accordance with the causal structure the
 254 model determines. The model was shown in the context of a P300-speller,
 255 a BCI for communication that uses event-related potentials (ERPs) such as

the P300 ERP, an EEG positive deflection around 300 ms after a rare and related event.

By inferring the user's intentions of which letter to spell (the observation) given a 6 x 6 grid of characters (the action), the machine studies how a user focuses their visual attention and how their brain reacts with a P300. This lets the machine detect a particular ERP in response. The machine flashes the items in repetition until it becomes confident about figuring out which stimulus corresponds to a targeted item[23].

The P300-based BCI has shown improvement with statistical methods have included Bayesian sensor fusions for event detection[24], a classifier with both Fisher's and Bayesian linear discriminant analysis[25], and a support vector machine classifier[26].

2.2. Brain-Computer Interfaces

Figure 1: Active Inference for the P300 speller. With free energy as a function of hidden states, the states can be either long term or short term reactions to stimuli. The machine's actions are based on the generative model of the user. The free energy is, then, a function of hidden states and actions $F_m(u_t, s_t)$.

With this method of Active Inference, one can use a generative Bayesian model of priors that encode predictions over future outcomes with a model m , a joint probability distribution over hidden states S , control states (actions) U , observations (sensations) O and model parameters S , U , and O . The

hidden states are internal representations of the BCI about the etiology of the observations such as the letter a user thinks of that causes the P300 ERP. The Markov model uses the precision parameter γ with preferences over future outcomes C .

The model works by figuring out which hidden states are most likely to occur by optimizing the expectations with respect to the free energy of the observations. This way, the fully adaptive P300 speller can learn and act optimally in real time by automatically and optimally updating an internal model of the environment. Then, given the speller's optimized flashing and the corresponding error potentials, the BCI can choose to spell the next probable letter or continue flashing to increase evidence for the target letter.

The finite set of hidden states would be: $S = s^{(1)}, s^{(2)}, \dots, s^{(n)}$ with $|S| = n$; in which s maps each trial t onto one element from finite set S such that $s(t) = s_t \in S, \forall t = 1, \dots, T$.

If n represents the number of possible states (cardinality of S at every trial t), T is the final trial, and t is the current one. Only one state out of n possible ones can take place at a time or trial t .

By sampling actions from control beliefs and inferring them from observations as well as making the assumption that agents know the realized actions, one can define U , a finite set of control states or actions such as which letter to flash next as $U = u^{(1)}, u^{(2)}, \dots, u^{(r)}$, with $|U| = r$ such that u maps t onto an element from the finite set U . The control states (actions) are then given by $u(t) = u_t \in U$ and $\forall t = 1, \dots, T$ in which r represents the number of possible states (cardinality of U at trial t). This lets the agent choose an action out of r possible ones.

Similarly, the observations are, then $O = o^{(1)}, o^{(2)}, \dots, o^{(z)}$, with $|O| = z$ such that o maps the trial t onto an element from O with $o(t) = o_t \in O, \forall t = 1, \dots, T$ for z , the number of possible observations (cardinality of O at trial t).

The Bayesian model is then:

$$P(\tilde{o}, \tilde{u}, \tilde{s}, \gamma | m) = P(\tilde{o} | \tilde{s}, m) P(\tilde{s}, \tilde{u} | \gamma, m) P(\gamma | m) \quad (1)$$

for $\tilde{o} = o_1, \dots, o_T \in O, \tilde{s} = s_1, \dots, s_T \in S, \tilde{u} = u_1, \dots, u_T \in U$ broken down into the likelihood matrix $P(\tilde{o} | \tilde{s}, m)$, probabilistic transition matrix $P(\tilde{s}, \tilde{u} | \gamma, m)$, and precision parameter $P(\gamma | m)$.

The likelihood matrix:

$$P(\tilde{o}, \tilde{u}, \tilde{s}, \gamma | m) = \prod_{i=1}^T P(\tilde{o}_i | \tilde{s}_i, \gamma, m) \quad (2)$$

307 with

$$P(o_i = k | s_i = h) = A_k, h \quad (3)$$

308 where $A \in \mathbb{R}^{z \times n}$ so that, for each $h = 1, \dots, n$ states with a probability to
 309 get the $k = 1, \dots, z$ observation. With the likelihood, the Bayesian model can
 310 predict the probabilities of new observations based on prior experience. The
 311 probabilistic transition matrix:

$$P(\tilde{s}, \tilde{u} | \gamma, m) = P(u_t | \gamma, m) \prod_{i=1}^t P(\tilde{s}_i + 1 | s_i, u_i, m) \quad (4)$$

312 with

$$P(s_{t+1} = w | s_t = q, u_t) = \mathbf{B}(u_t)_{w,q} \quad (5)$$

313 for $w, q = 1, \dots, n$ and $\mathbf{B}(u_t) \in \mathbb{R}^{n \times n}$ for number of hidden states n . The
 314 putative action, a specific control state that minimizes expected free energy,
 315 u_t under policy $\pi \in 1, \dots, K$. A policy indexes a specific sequence of control
 316 states $(\tilde{u} | \pi) = (u_t, \dots, u_\tau | \pi)$:

$$\ln(P(\tilde{u} | \gamma, m)) = \gamma \cdot \mathbf{Q}(\pi) = \gamma \cdot (\mathbf{Q}(u_{t+1} | \pi) + \dots + \mathbf{Q}(u_\tau | \pi)) \quad (6)$$

$$\mathbf{Q}(u_\tau | \pi) = E_{Q(o_\tau | \pi)}[\ln(P(o_\tau | m))] + E_{Q(o_\tau | \pi)}[D_{KL}[Q(s_\tau | o_\tau, \pi) | Q(s_\tau | \pi)]] \quad (7)$$

317 with extrinsic value $E_{Q(o_\tau | \pi)}[\ln(P(o_\tau | m))]$ and epistemic value $E_{Q(o_\tau | \pi)}[D_{KL}$
 318 $[Q(s_\tau | o_\tau, \pi) | Q(s_\tau | \pi)]]$, probability distribution Q . The precision parameter
 319 γ is used to minimize expected free energy, D_{KL} as the Kullback-Leibler
 320 (KL) divergence (relative entropy). $E_{Q(o_\tau | \pi)}$ is the expectation of a future
 321 outcome, o_{tau} , given policy π . The extrinsic value is the maximized preferred
 322 final outcome, and the epistemic value is the information maximized achieved
 323 by minimizing relative entropy. This way, we maximize information to get
 324 to the preferred future outcome.

325 Using C_τ , the prior distribution over the outcomes $P(o_\tau | m)$, we can de-
 326 termine the preference for future outcomes.

327 The extrinsic value contains $P(o_\tau|m)$ which is the prior distribution over
 328 future outcomes, referred to as C_τ . So, let C_τ be the preference of future
 329 outcomes $o_\tau \in O$. As part of extrinsic value, it influences the choice of action
 330 to reach such desired outcomes. If we consider all available observations from
 331 set O as future outcomes then $o_\tau = o^{(z)}$:

$$C_\tau = \sigma([C(o^{(1)}), C(o^{(2)}), \dots, C(o^{(z)})])^T \quad (8)$$

332 For σ softmax (normalized exponential function) of final outcomes that
 333 can transform observations to prior probabilities of future outcomes.

334 The precision parameter is, then:

$$P(\gamma|m) = \Gamma(\alpha, \beta) \quad (9)$$

335 where $\Gamma(\alpha)$ is a gamma distribution of scale parameter a and rate param-
 336 eter b . If a random variable X follows a gamma distribution then:

$$f(x; \alpha, \beta) = \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)} \quad (10)$$

337 for $x > 0$ and $\alpha, \beta > 0$.

338 where $\Gamma(\alpha)$ is the gamma function, as defined:

$$\Gamma(\alpha) = (\alpha - 1)! \quad (11)$$

339 The active inference framework can then be implemented using optimal
 340 stopping and flashing while classifying error potential for future outcomes.
 341 This way, the machine can act the way the user wants to.

342 3. Conclusion

343 Active inference provides a way for BCIs to improve their performance
 344 in predicting user actions based on how the user interacts with a genera-
 345 tive model of the user and their interactions. The general nature of the
 346 free energy principle provides a grand, unifying way of understanding how
 347 self-organization emerges. The corollary, active inference, then can be used
 348 for calculating the epistemic reward, extrinsic value, and other parameters
 349 in minimizing free energy and choosing the optimal decision for how the
 350 user wants to act. One can describe evidence of specific behavior for active

inference in terms of simplicity and accuracy. Paraphrasing Einstein, everything should be made as simple as possible, but not simpler. Minimizing free energy provides a way of achieving this simplicity.

References

- [1] K. Friston, Life as we know it, *Journal of the Royal Society Interface* (2013).
- [2] K. Friston, J. Kilner, L. Harrison, A free energy principle for the brain, *Journal of Physiology-Paris* (2006).
- [3] K. Friston, The free-energy principle: a unified brain theory?, *Nature Reviews Neuroscience* 10 (2010) 127–138.
- [4] P. Schwartenbeck, T. FitzGerald, R. J. Dolan, K. Friston, Exploration, novelty, surprise, and free energy minimization, *Frontiers in Psychology* 4 (2013).
- [5] K. Friston, J. Kilner, L. Harrison, Free-energy and the brain, *Synthese* (2007).
- [6] K. Friston, Embodied inference: or “i think therefore i am, if i am what i think” (2010).
- [7] F. K., M. Lin, C. Frith, G. Pezzulo, A. Hobson, J., S. Ondobaka, Active inference and artificial curiosity, *Neural Computation* 29 (2017) 3434–3445.
- [8] T. Parr, K. Friston, Generalised free energy and active inference, *Biological Cybernetics* 113 (2019) 495–513.
- [9] P. Schwartenbeck, T. H. B. FitzGerald, C. Mathys, R. Dolan, K. Friston, The dopaminergic midbrain encodes the expected certainty about desired outcomes, *Cerebral Cortex* 10 (2015) 3434–3445.
- [10] K. Friston, M. Lin, C. Frith, G. Pezzulo, A. Hobson, J., S. Ondobaka, Active inference, curiosity and insight, *Neural Computation* 10 (2017) 2633–2683.
- [11] D. Yon, F. P. de Lange, C. Press, The predictive brain as a stubborn scientist, *Trends in Cognitive Sciences* 23 (2019) 6–8.

- 381 [12] M. B. Mirza, R. A. Adams, C. Mathys, K. Friston, Scene construc-
382 tion, visual foraging, and active inference, *Frontiers in Computational*
383 *Neuroscience* 10 (2016) 1–16.
- 384 [13] J. Mladenovic, J. Frey, M. Joffily, E. Maby, F. Lotte, J. Mattout, Active
385 inference as a unifying, generic and adaptive framework for a p300-based
386 bci, *Journal of Neural Engineering* 17 (2020).
- 387 [14] A. Tschantz, C. Heins, B. Klein, Active inference for markov decision
388 processes, <https://github.com/jkbren/infer-actively>,
389 2019.
- 390 [15] J. Schmidhuber, Formal theory of creativity, fun, and intrinsic motiva-
391 tion (1990-2010), *IEEE Transactions on Autonomous Mental Develop-*
392 *ment* 2 (2010).
- 393 [16] K. Friston, J. Daunizeau, J. Kilner, S. J. Kiebel, Action and behavior:
394 a free-energy formulation, *Biological Cybernetics* 102 (2010) 227–260.
- 395 [17] R. P. N. Rao, D. H. Ballard, Predictive coding in the visual cortex:
396 a functional interpretation of some extra-classical receptive-field effects,
397 *Nature Neuroscience* 2 (1999) 79–87.
- 398 [18] K. Friston, P. Schwartenbeck, T. FitzGerald, M. Moutoussis,
399 T. Behrens, R. J. Dolan, The anatomy of choice: dopamine and decision-
400 making, *Philosophical Transactions of the Royal Society B: Biological*
401 *Sciences* 369 (2014) 24–29.
- 402 [19] J. del R. Milan, J. M. Carmena, Invasive or noninvasive: Under-
403 standing brain-machine interface technology [conversations in bme],
404 *IEEE Engineering in Medicine and Biology Magazine* 29 (2010) 16–22.
405 doi:10.1109/MEMB.2009.935475.
- 406 [20] L. A. Farwell, E. Donchin, Talking off the top of your head: toward
407 a mental prosthesis utilizing event-related brain potentials, *Electroen-*
408 *cephalography and Clinical Neurophysiology* (1988).
- 409 [21] T. O. Zander, C. Kothe, Towards passive brain–computer interfaces: ap-
410 plying brain–computer interface technology to human–machine systems
411 in general, *Journal of Neural Engineering* (2011).

- 412 [22] J.-M. Batail, S. Bioulac, F. Cabestaing, C. Daudet, D. Drapier,
413 M. Fouillen, T. Fovet, A. Hakoun, R. Jardri, C. Jeunet, F. Lotte,
414 E. Many, J. Mattout, T. Medani, J.-A. Micoulaud-Franchi, J. Mladen-
415 novic, L. Perronet, L. Pillette, T. Ros., F. Vialatte, N. group, Eeg
416 neurofeedback research: A fertile ground for psychiatry?, *L'Encéphale*
417 45 (2019) 245–255.
- 418 [23] E. W. Sellers, D. J. Krusienski, D. J. McFarland, V. T. M., J. R. Wol-
419 paw, A p300 event-related potential brain-computer interface (bci): the
420 effects of matrix size and inter stimulus interval on performance, *Bio-*
421 *logical Psychology* 73 (2006) 242–252.
- 422 [24] G. Pires, M. Castelo-Branco, U. Nunes, Visual p300-based bci to steer
423 a wheelchair: A bayesian approach, in: 2008 30th Annual International
424 Conference of the IEEE Engineering in Medicine and Biology Society,
425 2008, pp. 658–661. doi:10.1109/IEMBS.2008.4649238.
- 426 [25] R. C. Panicker, S. Puthusserypady, Y. Sun, Adaptation in p300
427 brain-computer interfaces: A two-classifier cotraining approach, *IEEE*
428 *Transactions on Biomedical Engineering* 57 (2010) 2927–2935.
- 429 [26] M. Thulasidas, C. Guan, J. Wu, Robust classification of eeg signal for
430 brain-computer interface, *IEEE Transactions on Neural Systems and*
431 *Rehabilitation Engineering* 14 (2006) 24–29.

Declaration of interests

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

☐ The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: