

MarkovType: A Markov Decision Process Strategy for Non-Invasive Brain-Computer Interfaces Typing Systems

Elifnur Sunger¹, Yunus Bicer¹, Deniz Erdogmus¹, Tales Imbiriba²

¹Northeastern University, Boston, MA, USA

²University of Massachusetts Boston, Boston, MA, USA

{sunger, bicer, erdogmus}@ece.neu.edu, tales.imbiriba@umb.edu

Abstract

Brain-Computer Interfaces (BCIs) help people with severe speech and motor disabilities communicate and interact with their environment using neural activity. This work focuses on the Rapid Serial Visual Presentation (RSVP) paradigm of BCIs using noninvasive electroencephalography (EEG). The RSVP typing task is a recursive task with multiple sequences, where users see only a subset of symbols in each sequence. Extensive research has been conducted to improve classification in the RSVP typing task, achieving fast classification. However, these methods struggle to achieve high accuracy and do not consider the typing mechanism in the learning procedure. They apply binary target and non-target classification without including recursive training. To improve performance in the classification of symbols while controlling the classification speed, we incorporate the typing setup into training by proposing a Partially Observable Markov Decision Process (POMDP) approach. To the best of our knowledge, this is the first work to formulate the RSVP typing task as a POMDP for recursive classification. Experiments show that the proposed approach, MarkovType, results in a more accurate typing system compared to competitors. Additionally, our experiments demonstrate that while there is a trade-off between accuracy and speed, MarkovType achieves the optimal balance between these factors compared to other methods.

Code — <https://github.com/neu-spiral/MarkovType>

1 Introduction

Brain-Computer Interfaces (BCIs) allow people with severe speech and motor disabilities to communicate and interact with their environments (Orhan et al. 2012). Recent BCI research spans applications like typing tasks (Bigdely-Shamlo et al. 2008; Lawhern et al. 2018; Smedemark-Margulies et al. 2023) and prosthetic arm control (Bright et al. 2016). Common BCI paradigms include steady-state visual evoked potentials (SSVEP) (Norcia et al. 2015), imagined motor movements (MI) (Wierzgała et al. 2018), rapid serial visual presentation (RSVP) (Acqualagna et al. 2010), and matrix layout-based P300-speller (Krusienski et al. 2008), which uses P300 event-related potentials observed in EEG to enable individuals to communicate without relying on voluntary muscle activity (Farwell and Donchin 1988; Mak et al.

2011). In certain cases, one paradigm may be more suitable than others. For example, individuals with severe speech and motor impairments, such as those with locked-in syndrome, ALS, or cerebral palsy, might not have the precise gaze control required for eye-tracking systems and matrix layout-based P300 spellers. However, precise gaze control is not required in RSVP to distinguish between symbols (Acqualagna et al. 2010; Orhan et al. 2012).

This paper focuses on the RSVP paradigm, where users type a target symbol from a predefined alphabet without physical movement, while the system rapidly presents candidate symbols (a query) and collects continuous EEG data. Users perform multiple sequences, and the model updates symbol probabilities based on their responses to each query. Typing stops when the probability of any symbol exceeds a confidence threshold or when the maximum number of sequences is reached. In RSVP setups, users do not observe the entire set of symbols at once. Instead, they observe a subset of symbols at each sequence. This makes RSVP suitable for formulation as a partially observable Markov decision process (POMDP). In POMDPs, the true state of the environment is unknown, and the agent or neural network receives only partial observations at each step (Kaelbling, Littman, and Cassandra 1998).

We introduce “MarkovType”, which formulates the RSVP typing task as a POMDP and enhances the recursive typing performance. We state our contributions as follows:

- We propose a Markov Decision Process for non-invasive BCI typing systems (MarkovType) and formulate the BCI typing procedure as a Partially Observable Markov Decision Process (POMDP), incorporating the typing mechanism into the learning procedure. To the best of our knowledge, this is the first time that the BCI typing procedure has been formalized as a POMDP and used for the typing task classification.
- We evaluate MarkovType in a simulated typing task and show it outperforms previous approaches using Recursive Bayesian Estimation in classification accuracy and information transfer rate.
- We also show that there is a trade-off between accuracy and the number of sequences in the typing task. This opens the door for future improvements and explorations.

2 Related Work

Extensive research exists on RSVP classification using EEG data. Prior work on classification has focused on using logistic regression or linear discriminant analysis (Meng et al. 2012; Bigdely-Shamlo et al. 2008), and applying the Recursive Bayesian rule for classification (Higger et al. 2013). In addition to classification, Recursive Bayesian methods are used for query optimization (Koçanaoğulları, Erdoğan, and Akçakaya 2018; Koçanaoğulları et al. 2019; Marghi et al. 2022), and designing stopping criteria for typing tasks (Koçanaoğulları, Akçakaya, and Erdoğan 2021).

There are also prior works on using convolutional neural networks (CNNs) for single-trial EEG RSVP classification (Shamwell et al. 2016; Zang et al. 2021; Lawhern et al. 2018). However, these methods cannot perform recursive classification. Smedemark-Margulies et al. (2023) focuses on a recursive Bayesian classification based on a discriminative probability for each query, which they approximate using neural network classifiers such as EEGNet (Lawhern et al. 2018) and hand-crafted CNNs. They show that discriminative models using hand-crafted CNNs outperform generative models such as linear discriminant analysis and logistic regression. They do not include the typing procedure in training and assume responses within a query are conditionally independent. Consequently, they perform binary classification, predicting each response in a single query as either a “target” or “non-target”. Although the proposed method performs well in binary classification, symbol classification can be improved by including the typing task in training and reconsidering the independence assumption to allow classification over the alphabet. We propose to do this by formulating the RSVP typing task as a partially observable Markov decision process (POMDP).

POMDP is used in many applications, such as visual attention networks (Mnih et al. 2014), human-robot interaction (Chen et al. 2018), autonomous driving (Bai et al. 2015), and medical diagnosis (Ayer, Alagoz, and Stout 2012; Hauskrecht and Fraser 2000). In the BCI domain, numerous works use POMDP-based models for symbol classification. Park and Kim (2012); Park, Kim, and Jo (2010) employed a POMDP-based model for classification using a matrix speller, while Tresols, Chanel, and Dehais (2023) proposed a POMDP-based model for high-level decision making across three BCI modalities: SSVEP (Norcia et al. 2015), MI (Wierzgała et al. 2018), and code-modulated visual evoked potentials (CVEP) (Sutter 1992). However, to the best of our knowledge, no prior work formulate the RSVP typing task as a POMDP.

3 Methods

3.1 Problem Statement

This work focuses on an RSVP typing task with N independent sequences, where the user tries to type the symbol t from an alphabet of A letters. At each sequence, the user sees a query of symbols indexed by $\mathbf{q}_n \in \mathbb{Z}^K$ with K different symbols, where $K \leq A$ and the index of target letter t may or may not be in \mathbf{q}_n . We observe the EEG responses to the queried symbols, $\mathbf{E}_n \in \mathbb{R}^{K \times c \times f}$, where c is the number

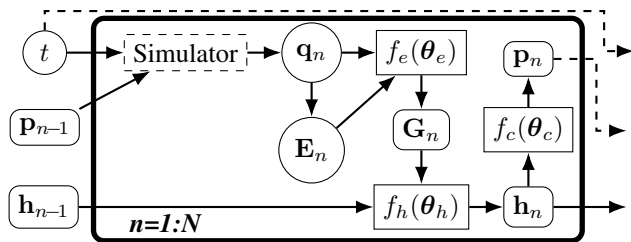


Figure 1: MarkovType Model Architecture. It consists of a simulator, a feature extractor f_e , a core network f_h and a classification network f_c . A single typing trial includes N sequences. Given the target symbol t and the prior over the alphabet $\mathbf{p}_{n-1} \in \mathbb{R}^A$, the simulator extracts a query with K unique symbols $\mathbf{q}_n \in \mathbb{Z}^K$ and EEG responses $\mathbf{E}_n \in \mathbb{R}^{K \times c \times f}$, where A is the length of the alphabet, c is the number of EEG channels and f is the feature length. f_e extracts features from \mathbf{E}_n and maps them over the alphabet using \mathbf{q}_n as in Eq. (1) to obtain $\mathbf{G}_n \in \mathbb{R}^{A \times L}$, where L is the feature length. f_h forms the hidden states of the current sequence \mathbf{h}_n using the alphabet features \mathbf{G}_n and the hidden states of the previous sequence (\mathbf{h}_{n-1}), which provides knowledge of the responses to previous queries.

of EEG channels and f is the feature-length. Note that the user never sees all the symbols in the alphabet simultaneously; they only see a subset of symbols appearing sequentially on the screen within a sequence.

To perform classification on this typing task, previous works assume that the EEG responses \mathbf{E}_n to the K symbols in the query q_n are independent (Smedemark-Margulies et al. 2023; Koçanaoğulları, Erdoğan, and Akçakaya 2018). As a result, they classify each symbol in the query as either “target” or “non-target”, without incorporating the recursive typing task into training.

Here, we include the RSVP typing task in the training procedure, where the model has partial observations (never receiving responses for all symbols in the alphabet simultaneously) and jointly learns EEG feature extraction and a fusion process for recursive classification over multiple sequences. To achieve recursive classification from partial observations and to incorporate the typing process in both training and testing, we formulate the typing process as a POMDP.

We propose a recursive classification model for the RSVP typing task, inspired by the recursive classification network of the Recurrent Attention Model (RAM) proposed by Mnih et al. (2014). In our setup, as illustrated in Figure 1, the agent (which is the model itself) learns to extract features from a query and incorporates hidden states recursively to perform classification over the alphabet A , while receiving a scalar reward that it aims to maximize. During this process, the agent chooses a query \mathbf{q}_n from the prior over the alphabet \mathbf{p}_{n-1} , which is updated with each sequence.

At the first sequence, we assume a uniform prior over the alphabet $\mathbf{p}_0 \in \mathbb{R}^A$, which is then updated after each sequence. Note that non-uniform priors can be incorporated, often employing language models (Orhan et al. 2013; Oken et al. 2014). However, we opted to analyze the performance

of MarkovType without the influence of language models. Details on the prior update are given in Section 3.2.

3.2 MarkovType Model Architecture

The overall model architecture, as shown in Figure 1, includes a simulator, a feature extractor f_e , followed by a core network f_h and a classification network f_c . In this setup, a single typing trial consists of N sequences. Given the target symbol t and the prior over the alphabet \mathbf{p}_{n-1} , the simulator extracts the query \mathbf{q}_n and EEG responses \mathbf{E}_n . f_e extracts features $\mathbf{G}_n \in \mathbb{R}^{A \times L}$ from \mathbf{E}_n and maps it over the alphabet using \mathbf{q}_n as in Eq. (1), where L is the feature length.

$$\mathbf{G}_{n,i} = \begin{cases} f_e(\mathbf{E}_{n,i}) & \text{if } i \in \mathbf{q}_n, \\ \mathbf{0}_L & \text{otherwise.} \end{cases} \quad (1)$$

$\mathbf{0}_L$ is a zero array with length L , $\mathbf{E}_{n,i}$ is the EEG response, and $\mathbf{G}_{n,i}$ are the features for the symbol $i \in A$ at sequence n . The hidden states of the current sequence $\mathbf{h}_n \in \mathbb{R}^v$ (where v is the length of the hidden feature) are formed by f_h using the alphabet features \mathbf{G}_n and the hidden states of the previous sequence (\mathbf{h}_{n-1}), which provide information on previous responses. Please see Section 4.1 for model details.

Actions At each sequence, the agent (neural network) performs an environment action $\mathbf{q}_n \in \mathbb{Z}^K$, where each element of \mathbf{q}_n is sampled from $\mathbf{p}_n = f_c(\mathbf{h}_n; \theta_c)$. This action affects the next state of the environment (\mathbf{E}_{n+1}).

Reward The setup explained in Sections 3.1 and 3.2 is a particular case of the POMDP. In the POMDP setup for this problem, the user’s EEG responses to all symbols in the alphabet are unknown and the agent (in our case, the neural network) receives partially observed sets of responses \mathbf{E}_n to the query \mathbf{q}_n at each sequence.

The total reward function according to POMDP is given in Eq. (2), where r_n is the reward assigned at sequence n .

$$R = \sum_{n=1}^N r_n d(n) \quad (2)$$

r_n is equal to 1 if the predicted label $\hat{t} = \arg \max_{i \in A} (\mathbf{p}_{n,i})$ matches the correct target label t at sequence n , and 0 otherwise, where $\mathbf{p}_{n,i}$ is the probability of symbol i . $d(n)$ is the discount factor. It creates a trade-off between finding the actions with the largest immediate reward and those with the largest sum of rewards. See Section 4.2 for details on $d(n)$.

3.3 Learning Rule

The neural network needs to learn a stochastic policy $\pi(\mathbf{q}_n | \mathbf{s}_{1:n}; \theta)$ which maps the history of past interactions with the environment to a distribution over actions for the current sequence. The history of past interactions $\mathbf{s}_{1:n} = \mathbf{E}_1, \mathbf{q}_1, \dots, \mathbf{E}_{n-1}, \mathbf{q}_{n-1}, \mathbf{q}_n, \mathbf{E}_n$ where \mathbf{E}_n is the observation produced by the environment, and \mathbf{q}_n are the actions. In this paper’s setup, the policy π is defined by the model, and the history \mathbf{s}_n is summarized in the hidden state \mathbf{h}_n .

Model parameters are learned by maximizing the total reward the model expects when interacting with the environment, using the REINFORCE rule (Wierstra et al. 2007).

POMDP - REINFORCE RULE The REINFORCE rule (Wierstra et al. 2007) follows the policy gradient framework, where model weights θ are updated directly by estimating the gradient in the direction of higher reward. $\mathbf{s}_{1:n}$ is the observed history, but we also need to define the complete history \mathbf{H} , which includes the unobserved (future) states. It is expressed by $\mathbf{H}_N = \langle \mathbf{s}_N, \mathbf{G}_{1:N} \rangle$. The probability of a history for given policy-defining weights θ is $p(\mathbf{H} | \theta)$, and $R(\mathbf{H})$ is a measure of the total reward achieved during history \mathbf{H} , which we aim to maximize in Eq. (3).

$$J(\theta) = \int_{\mathbf{H}} p(\mathbf{H} | \theta) R(\mathbf{H}) d\mathbf{H}. \quad (3)$$

We rewrite Eq. (3) using gradient ascent to update θ :

$$\nabla_{\theta} J = \int \nabla_{\theta} p(\mathbf{H} | \theta) R(\mathbf{H}) d\mathbf{H}. \quad (4)$$

Viewing the problem as a POMDP setup, we can first apply the “likelihood-ratio trick” (also considering the fact that $\nabla_{\theta} R = 0$ for a single, fixed \mathbf{H}), as follows:

$$\begin{aligned} \nabla_{\theta} J &= \int \frac{p(\mathbf{H} | \theta)}{p(\mathbf{H} | \theta)} \nabla_{\theta} (p(\mathbf{H} | \theta)) R(\mathbf{H}) d\mathbf{H}, \\ &= \int p(\mathbf{H} | \theta) \nabla_{\theta} \log(p(\mathbf{H} | \theta)) R(\mathbf{H}) d\mathbf{H}. \end{aligned} \quad (5)$$

We take the sample average with the Monte Carlo approximation (MC) in Eq (6), where M is the number of episodes.

$$\begin{aligned} \nabla_{\theta} J &= \mathbb{E}_{\mathbf{H}} [\nabla_{\theta} \log(p(\mathbf{H} | \theta)) R(\mathbf{H})] \\ &\approx \frac{1}{M} \sum_{m=1}^M \nabla_{\theta} \log(p(\mathbf{H}^m | \theta)) R(\mathbf{H}^m). \end{aligned} \quad (6)$$

To calculate $\log p(\mathbf{H} | \theta)$, we state that the probability of a particular history is the product of all actions and observations given the sub-histories as follows:

$$\begin{aligned} p(\mathbf{H}_N | \theta) &= p(\mathbf{E}_1, \mathbf{G}_1) \prod_{t=2}^N [p(\mathbf{E}_t, \mathbf{G}_t | \mathbf{s}_{1:t-1}, \mathbf{q}_{t-1}, \mathbf{G}_{1:t-1}) \\ &\quad \times \pi(\mathbf{q}_{t-1} | \mathbf{s}_{1:t-1}; \theta)], \end{aligned} \quad (7)$$

Then, taking the log-derivate of Eq. (7) results in:

$$\nabla_{\theta} \log p(\mathbf{H}_N | \theta) = \sum_{n=1}^N \nabla_{\theta} \log \pi(\mathbf{q}_n | \mathbf{s}_{1:n}; \theta). \quad (8)$$

Therefore, the gradient approximation can be written as:

$$\nabla_{\theta} J \approx \frac{1}{M} \sum_{m=1}^M \sum_{n=1}^N \nabla_{\theta} \log \pi(\mathbf{q}_n^m | \mathbf{s}_{1:n}^m; \theta) R_n^m, \quad (9)$$

where R_n^m is the sum of rewards at sequence n and episode m . Note that $\nabla_{\theta} \log \pi(\mathbf{q}_n^m | \mathbf{s}_{1:n}^m; \theta)$ is the gradient of the model and can be computed by standard backpropagation. Also, to overcome the high variance problem of the gradient due to MC, a baseline b is added (that may depend on $\mathbf{s}_{1:n}^m$ but

not on the actions \mathbf{q}_n^m as explained by Mnih et al. (2014). Therefore, we rewrote Eq. (9) as:

$$\nabla_{\theta} J \approx \frac{1}{M} \sum_{m=1}^M \sum_{n=1}^N \nabla_{\theta} \log \pi(\mathbf{q}_n^m | \mathbf{s}_{1:n}^m; \theta) (R_n^m - b_n). \quad (10)$$

Hybrid Supervised Loss We follow the supervised loss described by Mnih et al. (2014), as defined in Eq. (11).

$$\mathcal{L}_{\text{action}} = -\log(\mathbf{p}_{N,t}), \quad (11a)$$

$$\mathcal{L}_{\text{baseline}} = \frac{1}{N} \sum_{n=1}^N (R_n - b_n)^2, \quad (11b)$$

$$\mathcal{L}_{\text{reinforce}} = -\frac{1}{M} \sum_{m=1}^M \sum_{n=1}^N \log \pi(\mathbf{q}_n^m | \mathbf{s}_{1:n}^m; \theta) (R_n^m - b_n), \quad (11c)$$

$$\mathcal{L} = \mathcal{L}_{\text{action}} + \lambda (\mathcal{L}_{\text{baseline}} + \mathcal{L}_{\text{reinforce}}). \quad (11d)$$

$\mathcal{L}_{\text{action}}$ is the negative log-likelihood loss, where $\mathbf{p}_{N,t}$ is the predicted probability of the target t at the last sequence, and $\lambda \geq 0$ is a regularization parameter. $\mathcal{L}_{\text{baseline}}$ and $\mathcal{L}_{\text{reinforce}}$ increase the reward and reduce bias in gradient estimation.

4 Experiments

4.1 Dataset and Model Details

We perform all experiments on a large, publicly available RSVP benchmark dataset containing over 1 million stimulus presentations (Zhang et al. 2020) that is suitable for a simulated typing task. In this study, 64 healthy subjects performed a target image detection task involving street-view images, either with humans (“target”) or without (“non-target”), during 64-channel electroencephalogram (EEG) data recording. We follow the preprocessing steps outlined by Smedemark-Margulies et al. (2023). We perform all experiments using data pooled from all subjects, repeating each experiment across 5 randomized train/test splits (with seeds from 0 to 4) and using 10% training set for validation, as described by Smedemark-Margulies et al. (2023). Details of the simulator, feature extractor f_e , the core network f_h , baseline b_n and classification network f_c are provided below.

Simulator One can use any type of typing simulator or language model. We use the simulator from Smedemark-Margulies et al. (2023). This simulator randomly selects K unique symbols from the alphabet based on the prior over the alphabet \mathbf{p}_{n-1} and samples the query of symbols indexed \mathbf{q}_n . This affects the state of the environment at sequence n , i.e., the EEG responses \mathbf{E}_n to the queried symbols.

Feature Extractor For feature extraction, we use a hand-designed one-dimensional CNN with 5 convolutional layers, applying 1D convolutions only across the time axis.

Core Network The core network f_h is used to update $\mathbf{h}_n = \text{LayerNorm}(\text{Rect}(\text{Linear}(\mathbf{h}_{n-1}) + \text{Linear}(\mathbf{G}_n)))$, where *LayerNorm* is a normalization layer, *Rect* is a rectified linear unit layer, and *Linear* is a linear layer.

Baseline We follow the baseline calculation from the image classification example given by Mnih et al. (2014) and calculate the baseline at n as $b_n = \text{Linear}(\mathbf{h}_n)$.

Classification Network It is formulated using a softmax output: $f_c(\mathbf{h}_n) = \text{Softmax}(\text{Linear}(\mathbf{h}_n))$. \mathbf{p}_0 is uniformly initialized, and \mathbf{h}_0 is initialized as an all-zero feature vector.

4.2 Implementation Details

We ran experiments on 2.4 GHz Intel E5-2680 v4 CPUs and 2.1 GHz Intel Xeon Platinum 8176 CPUs, using PyTorch (Paszke et al. 2019) 2.3.0. We train the models using the Adam optimizer (Kingma and Ba 2014) with a learning rate of 0.001 for 200 epochs, decaying the learning rate by a factor of 0.97 after each epoch, and use a batch size of 28. We set the number of symbols shown in a query K to 10, the number of sequences N to 10, the number of episodes M to 1 and the alphabet length A to 28. We perform hyper-parameter tuning on a validation set consisting of a held-out 10% of the training set, checking λ values of $\{0.01, 0.02, 0.03, \dots, 0.1\}$.

Our analyses include different discount factors of: $d(n) = \{\frac{2N-n-1}{N}, \frac{1}{n}, \frac{1}{n^2}, \frac{1}{n^3}\}$. These vary in how much importance they give to early correct classifications. Specifically, the importance increases in the given order, from $\frac{2N-n-1}{N}$ to $\frac{1}{n^3}$. The final λ values used for models with these discount factors are $\{0.02, 0.02, 0.01, 0.1\}$, respectively.

4.3 Competing Methods

We compare our model against the two best-performing discriminative methods described by Smedemark-Margulies et al. (2023), which are 1D and 2D CNNs. Similarly, these methods are trained with the Adam optimizer (Kingma and Ba 2014) using a learning rate of 0.001 for 25 epochs, decaying the learning rate by a factor of 0.97 after each epoch, as described by Smedemark-Margulies et al. (2023). We refer to these methods as “RB - 1D CNN” and “RB - 2D CNN” throughout this paper. Note that during classification, RB methods do not classify \mathbf{E}_n directly over the alphabet (number of A classes). Instead, they classify symbols as “target” or “non-target”. During testing, they receive queries of length K , perform binary classification using the Recursive Bayesian rule and map the results to the alphabet, explained by Smedemark-Margulies et al. (2023).

4.4 Experimental Setup

Algorithm 1 outlines the testing procedure for the simulated RSVP task using MarkovType¹. The threshold τ is used for early stopping. If any posterior probability exceeds this threshold or if the maximum number of sequences (set to 10) is reached, the classification sequence for target t is terminated. Since we are simulating the typing process, we generate \mathbf{E}_n by sampling from the “target” and “non-target” examples given the query \mathbf{q}_n and the symbol t . For RB methods, a recurrent Bayesian update occurs instead of line 8 of Algorithm 1, which does not use \mathbf{h}_{n-1} as inputs. See Section 4.5 for Information Transfer Rate (ITR) calculations.

¹<https://github.com/neu-spiral/MarkovType>

Algorithm 1: Test procedure

Input: Trained model $f(\cdot)$, Pos. and Neg. Test Data \mathcal{E}^+ , \mathcal{E}^- , Symbols per query K , Sequences per symbol N , Alphabet size A , Decision threshold τ .

Output: ITR (bits per symbol), ITR_{n_τ} (bits per sequence)

```
1:  $C \leftarrow 0, N_\tau \leftarrow 0$  // *Correct symbol count, Total number of sequences taken
2: for  $t \leftarrow 1 : T$  // *Target symbols do
3:    $\mathbf{h}_0 \leftarrow (0, \dots, 0)$  // *Hidden feature initialization
4:    $\mathbf{p}_0 \leftarrow (\frac{1}{A}, \dots, \frac{1}{A})$  // *Uniform symbol prior
5:   for  $n \leftarrow 1 : N$  // *Number of sequences do
6:      $\mathbf{q}_n \sim \mathbf{p}_{n-1}$  // *Sample query
7:      $\{\mathbf{E}_n[k] \sim \mathcal{E}^+ \text{ if } \mathbf{q}_n[k] = t \text{ else } \mathbf{E}_n[k] \sim \mathcal{E}^-\}_{k=1}^K$ 
8:      $\mathbf{p}_n, \mathbf{h}_n \leftarrow f(\mathbf{E}_n, \mathbf{q}_n, \mathbf{h}_{n-1})$  // *Model outputs
9:      $\text{ind}, \text{val} \leftarrow \arg \max(\mathbf{p}_n), \max(\mathbf{p}_n)$ 
10:    if  $\text{val} \geq \tau$  or  $n = N$  then
11:      if  $\text{ind} = t$  then
12:         $C \leftarrow C + 1$ 
13:      end if
14:       $N_\tau \leftarrow N_\tau + n$ 
15:      break
16:    end if
17:  end for
18: end for
19:  $n_\tau \leftarrow N_\tau / T$  // *Mean number of sequences
20: return  $\text{ITR}(A, C/T), \text{ITR}(A, C/T)_{n_\tau}$ 
```

Our analyses of test time include classifications both with and without a threshold $\tau = 0.8$, using 1000 target symbols from an alphabet of length 28.

4.5 Evaluation Metrics

We compare the accuracy and information transfer rate (ITR) (Shannon 1948) of different models. ITR per selection (Smedemark-Margulies et al. 2023) is calculated as:

$$\text{ITR}(A, P) := \log_2(A) + P \log_2(P) + (1 - P) \log_2\left(\frac{1 - P}{A - 1}\right), \quad (12)$$

where P is the accuracy and A is the alphabet length. However, this formulation does not account for the time to type a symbol. We use ITR per sequence given in Eq. (13) to include the time spent typing a symbol, where n_τ is the mean number of sequences taken for making a decision.

$$\text{ITR}(A, P)_{n_\tau} = \frac{1}{n_\tau} \text{ITR}(A, P). \quad (13)$$

5 Results

We evaluate the performance of MarkovType and RB methods both with and without using the threshold τ . When the threshold τ is applied, the classification sequence for a symbol terminates if any posterior probability exceeds τ or the maximum number of sequences is reached. This allows us to assess the time spent making a decision (i.e., the number of sequences taken), as well as accuracy and ITR. Without the threshold, early stopping does not occur, and we evaluate

performance as a function of the number of sequences, with classification occurring at every sequence. This also demonstrates the upper limit for classification performance.

5.1 With Threshold τ

Here, we report the model performances when the testing process for the simulated typing uses the threshold τ for early stopping. Table 1 includes a comparison of MarkovType with RB methods using the following metrics: *number of parameters*, *ITR per selection*, *accuracy*, *mean number of sequences per selection* (n_τ), and *ITR per sequence*. We observe that all MarkovType models perform better than RB models in terms of ITR per selection and accuracy. However, RB methods have a lower n_τ , indicating they make decisions faster than MarkovType methods, but with lower performance in terms of ITR per selection and accuracy. This highlights a trade-off between speed and accuracy.

To analyze ITR with the time spent on decisions, we incorporate the number of sequences into the ITR calculation (ITR per sequence). We observe that despite requiring a higher n_τ , all MarkovType models achieve better overall performance. When comparing the effect of discount factors, we observe that going from $\frac{2N-n-1}{N}$ to $\frac{1}{n^2}$ results in a decrease in n_τ and an increase in ITR per sequence. However, this trend does not continue when moving to $\frac{1}{n^3}$. As explained in Section 4.1, the discount factors vary in the importance they assign to early correct classifications; i.e., the importance increases in the given order from $\frac{2N-n-1}{N}$ to $\frac{1}{n^3}$.

Figure 2 shows the distributions of correct and incorrect decisions made at each sequence. Similar to the observations in Table 1, RB methods make most of their decisions at earlier sequences and classify all 1000 symbols before sequence 8. MarkovType methods make decisions at later sequences, with some cases classified after sequence 7, but they achieve a higher accuracy rate at each sequence. We observe that going from $\frac{2N-n-1}{N}$ to $\frac{1}{n^2}$ results in more decisions being made in the first 4 sequences. However, this trend does not occur when moving from $\frac{1}{n^2}$ to $\frac{1}{n^3}$.

To analyze the accuracy of decisions in more detail statistically, Figure 3 examines the accuracy of decisions made at each sequence, showing performance over 5 data splits. The lines represent the mean accuracy, with confidence bars indicating the standard deviation across the 5 data splits. Figure 3 clearly demonstrates the accuracy difference between MarkovType and RB methods, with MarkovType consistently achieving much higher accuracy than RB methods. In most sequences, the accuracy difference is about 0.4.

5.2 Without Threshold τ

In this section, we evaluate accuracy performance as a function of the number of sequences, ranging from 1 to 10. We present model performances without using the threshold τ , with all 1000 symbols in the test set classified at each sequence and no early stopping. Without a threshold, the model can correct incorrect predictions in subsequent sequences, resulting in differences in accuracy between cases with and without a threshold. Figure 4 shows the accuracy of all symbols across different sequences, with the lines representing mean accuracy and confidence bars indicating

Model	$d(n)$	Num. Par	ITR w/ τ per Selection (\uparrow)	Accuracy w/ τ (\uparrow)	Num. Sequences per Selection (\downarrow)	ITR w/ τ per Sequence (\uparrow)
MarkovType	$\frac{2N-n-1}{N}$	479293	3.635\pm0.031	0.870\pm0.004	4.956 \pm 0.183	0.734 \pm 0.023
MarkovType	$\frac{1}{N}$	479293	3.550 \pm 0.110	0.859 \pm 0.015	4.832 \pm 0.213	0.735 \pm 0.026
MarkovType	$\frac{1}{n}$	479293	3.450 \pm 0.143	0.845 \pm 0.020	4.524 \pm 0.243	0.763\pm0.012
MarkovType	$\frac{1}{n^2}$	479293	3.502 \pm 0.074	0.852 \pm 0.010	4.638 \pm 0.142	0.756 \pm 0.031
RB - 1D CNN	-	542210	1.230 \pm 0.061	0.457 \pm 0.013	2.321\pm0.064	0.531 \pm 0.037
RB - 2D CNN	-	1603042	1.377 \pm 0.058	0.489 \pm 0.012	2.564 \pm 0.040	0.537 \pm 0.020

Table 1: Number of parameters, ITR per selection, accuracy, mean number of sequences per selection (n_τ), and ITR per sequence are evaluated with the threshold for MarkovType using different discount functions $d(\cdot)$, and RB - 1D and 2D. Note that classification sequences terminate when any posterior probability exceeds the threshold or when the maximum number of sequences (in this case, 10) is reached. Therefore, lower sequences per symbol with high accuracy are desired for good typing performance. We observe a trade-off between accuracy and the number of sequences per symbol. Even though RB methods perform quicker classification, MarkovType methods still achieve better ITR per sequence, where the mean number of sequences per symbol is included in the ITR calculation.

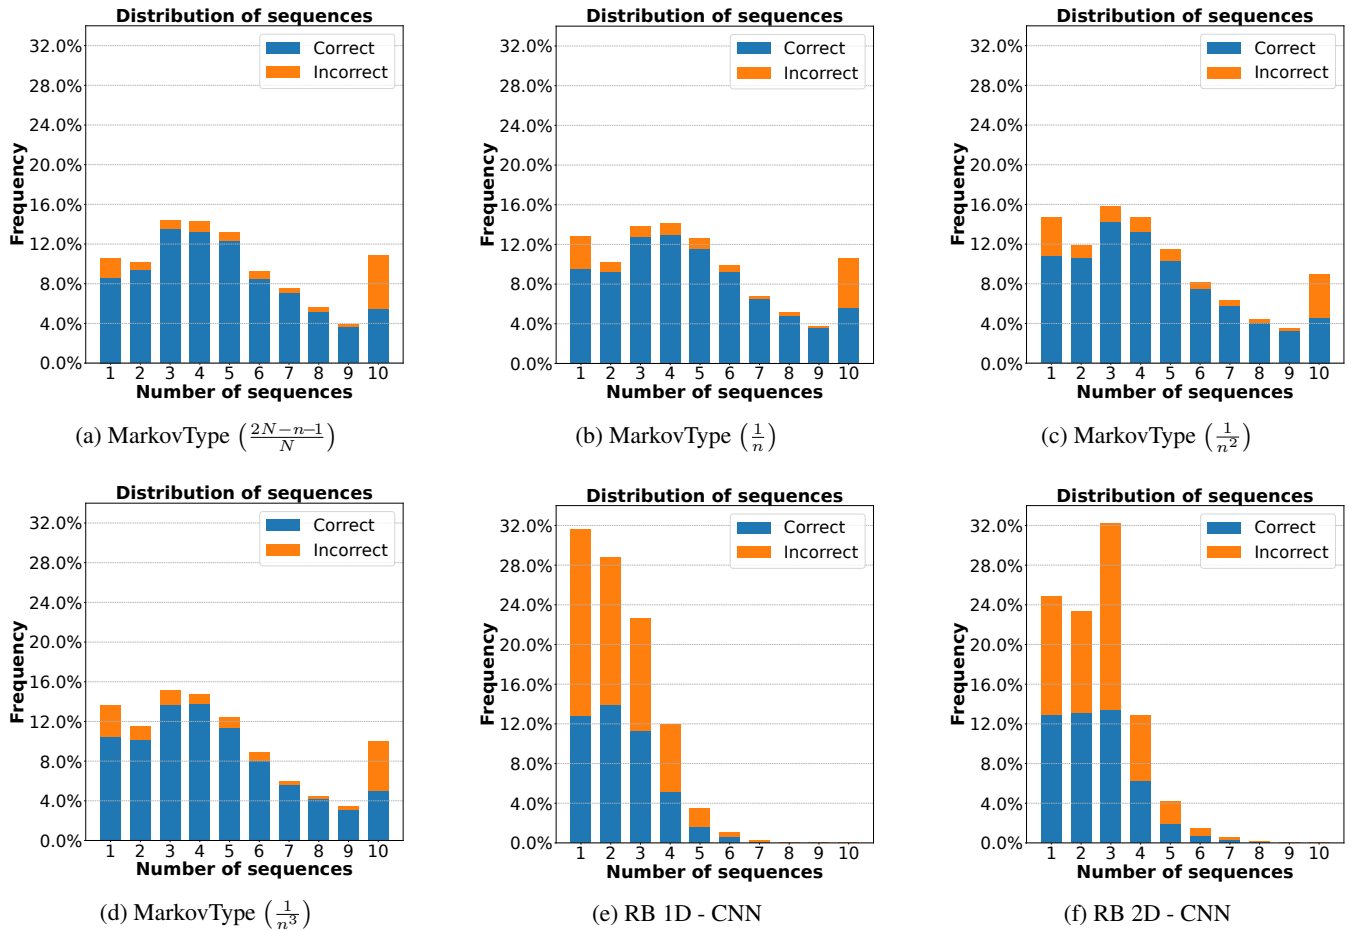


Figure 2: Distribution of correctly and incorrectly classified symbols over 5 data splits across sequences (1-10). Note that classification sequences terminate when any posterior probability exceeds the threshold or when the maximum number of sequences (in this case, 10) is reached. At each sequence, the blue bar represents correct decisions, and the orange bar represents incorrect decisions. Together, they show the total number of decisions made at each sequence. We observe that although RB methods (e) and (f) make most of their decisions early on, these decisions are mostly incorrect. In contrast, MarkovType methods (a)-(d) tend to make decisions later than those using recursive Bayesian methods but achieve a higher ratio of correct decisions.

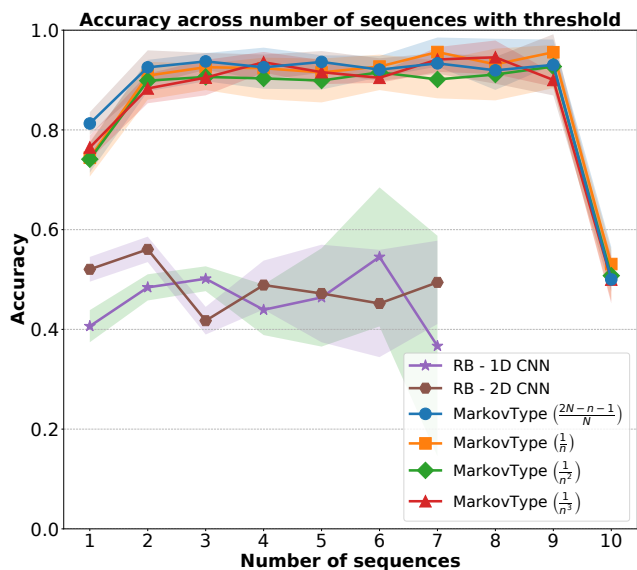


Figure 3: Accuracy across sequences with threshold τ . Lines represent the mean accuracy and confidence bars indicate the standard deviation across the 5 data splits. This plot shows the accuracy of decisions made across sequences. Classification sequences terminate when any posterior probability exceeds the threshold or when the maximum number of sequences is reached. With RB methods, no symbols remain to be classified after the 7th sequence. MarkovType consistently achieves higher accuracy than RB methods.

the standard deviation across the 5 data splits. MarkovType methods have higher accuracy across all sequences, except the first, where all methods exhibit similar accuracy. By the 10th sequence, which demonstrates the limits of all methods, MarkovType achieves higher accuracy, with the method using the $\frac{1}{n^2}$ discount factor performing the best.

6 Conclusion

This work proposes a Markov Decision Process for non-invasive BCI typing systems (MarkovType) that formulates the RSVP typing task as a Partially Observed Markov Decision Process (POMDP) for recursive classification. To the best of our knowledge, this is the first work to formalize the RSVP typing task as a POMDP. Formulating the RSVP typing task as a POMDP is appropriate because the task involves multiple sequences where partial subsets of symbols from an alphabet are presented. We perform our analyses on a simulated typing task and compare them with Recursive Bayesian-based methods. Unlike MarkovType, they assume that responses to each query are independent and do not include the typing procedure during training. Our experiments show that MarkovType significantly outperforms Recursive Bayesian methods in terms of accuracy and information transfer rate per selection and symbol, with the symbol-based calculation accounting for the time spent on selection. However, we observe a trade-off between accuracy and speed, specifically the number of sequences required for a

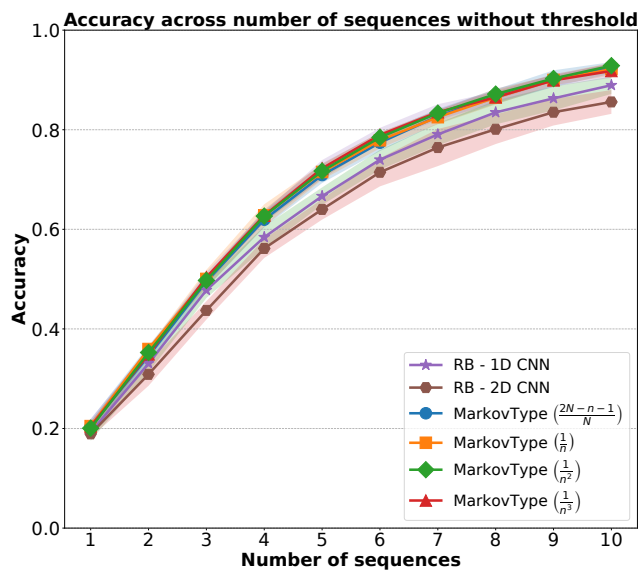


Figure 4: Accuracy across sequences without threshold τ , with lines representing the mean accuracy and confidence bars indicating the standard deviation across the 5 data splits. This plot shows the accuracy of symbols with varying numbers of sequences from 1 to 10. Note that when the algorithm includes a threshold, early stopping occurs. Without early stopping, we can observe the upper limit of accuracy with different numbers of sequences. For instance, with a threshold, the model might make an incorrect prediction without an opportunity to correct it in subsequent sequences, resulting in differences in accuracy between cases with and without a threshold. At each sequence, the classification of 1000 symbols in the test set is performed. We see that MarkovType methods perform better at each sequence compared to RB methods, with the MarkovType method using the $\frac{1}{n^2}$ discount factor achieving the highest accuracy.

decision when a confidence threshold is used to terminate sequences. This opens the door for future improvements in achieving both fast and accurate classification. Additionally, all experiments are conducted on a large RSVP benchmark.

Future work might focus on the real-world adaptation of MarkovType with limited data and user-specific training. In this work, we aim to keep MarkovType simple yet effective to facilitate its potential real-world use, with fewer parameters than the baselines. If the model proves too complex for real-world applications with limited data, simplifying the architecture could be considered. Our method is also trained globally in this work. For user-dependent cases, transfer learning approaches should be considered to adapt the model to different users in real-time.

Acknowledgments

This work was supported by NIH grant R01DC009834 as part of CAMBI. We thank all the members of CAMBI for their insightful discussions and valuable contributions.

References

- Acqualagna, L.; Treder, M. S.; Schreuder, M.; and Blankertz, B. 2010. A novel brain-computer interface based on the rapid serial visual presentation paradigm. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, 2686–2689. IEEE.
- Ayer, T.; Alagoz, O.; and Stout, N. K. 2012. OR Forum—A POMDP approach to personalize mammography screening decisions. *Operations Research*, 60(5): 1019–1034.
- Bai, H.; Cai, S.; Ye, N.; Hsu, D.; and Lee, W. S. 2015. Intention-aware online POMDP planning for autonomous driving in a crowd. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 454–460. IEEE.
- Bigdely-Shamlo, N.; Vankov, A.; Ramirez, R. R.; and Makeig, S. 2008. Brain activity-based image classification from rapid serial visual presentation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 16(5): 432–441.
- Bright, D.; Nair, A.; Salvekar, D.; and Bhisikar, S. 2016. EEG-based brain controlled prosthetic arm. In *2016 Conference on Advances in Signal Processing (CASP)*, 479–483. IEEE.
- Chen, M.; Nikolaidis, S.; Soh, H.; Hsu, D.; and Srinivasa, S. 2018. Planning with trust for human-robot collaboration. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 307–315.
- Farwell, L. A.; and Donchin, E. 1988. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and clinical Neurophysiology*, 70(6): 510–523.
- Hauskrecht, M.; and Fraser, H. 2000. Planning treatment of ischemic heart disease with partially observable Markov decision processes. *Artificial intelligence in medicine*, 18(3): 221–244.
- Higger, M.; Akcakaya, M.; Orhan, U.; and Erdogmus, D. 2013. Robust classification in RSVP keyboard. In *Foundations of Augmented Cognition: 7th International Conference, AC 2013, Held as Part of HCI International 2013, Las Vegas, NV, USA, July 21-26, 2013. Proceedings 7*, 443–449. Springer.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2): 99–134.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Koçanaoğulları, A.; Akcakaya, M.; and Erdoğan, D. 2021. Stopping criterion design for recursive Bayesian classification: analysis and decision geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9): 5590–5601.
- Koçanaoğulları, A.; Erdoğan, D.; and Akçakaya, M. 2018. On analysis of active querying for recursive state estimation. *IEEE signal processing letters*, 25(6): 743–747.
- Koçanaoğulları, A.; Marghi, Y.; Akçakaya, M.; and Erdoğan, D. 2019. An active recursive state estimation framework for brain-interfaced typing systems. *Brain-Computer Interfaces*, 6(4): 149–161.
- Krusienski, D. J.; Sellers, E. W.; McFarland, D. J.; Vaughan, T. M.; and Wolpaw, J. R. 2008. Toward enhanced P300 speller performance. *Journal of neuroscience methods*, 167(1): 15–21.
- Lawhern, V. J.; Solon, A. J.; Waytowich, N. R.; Gordon, S. M.; Hung, C. P.; and Lance, B. J. 2018. EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces. *Journal of neural engineering*, 15(5): 056013.
- Mak, J.; Arbel, Y.; Minett, J. W.; McCane, L. M.; Yuksel, B.; Ryan, D.; Thompson, D.; Bianchi, L.; and Erdogmus, D. 2011. Optimizing the P300-based brain-computer interface: current status, limitations and future directions. *Journal of neural engineering*, 8(2): 025003.
- Marghi, Y. M.; Koçanaoğulları, A.; Akçakaya, M.; and Erdoğan, D. 2022. Active recursive Bayesian inference using Rényi information measures. *Pattern Recognition Letters*, 154: 90–98.
- Meng, J.; Meriño, L. M.; Shamlo, N. B.; Makeig, S.; Robbins, K.; and Huang, Y. 2012. Characterization and Robust Classification of EEG Signal from Image RSVP Events with Independent Time-Frequency Features. *PLOS ONE*, 7(9): 1–13.
- Mnih, V.; Heess, N.; Graves, A.; et al. 2014. Recurrent models of visual attention. *Advances in neural information processing systems*, 27.
- Norcia, A. M.; Appelbaum, L. G.; Ales, J. M.; Cottareau, B. R.; and Rossion, B. 2015. The steady-state visual evoked potential in vision research: A review. *Journal of vision*, 15(6): 4–4.
- Oken, B.; Orhan, U.; Roark, B.; Erdogmus, D.; Fowler, A.; Mooney, A.; Peters, B.; Miller, M.; and Fried-Oken, M. 2014. Brain-computer interface with language model-electroencephalography fusion for locked-in syndrome. *Neurorehabilitation and Neural Repair*, 28(4): 387–394.
- Orhan, U.; Erdogmus, D.; Roark, B.; Oken, B.; and Fried-Oken, M. 2013. Offline analysis of context contribution to ERP-based typing BCI performance. *Journal of Neural Engineering*, 10(6): 066003.
- Orhan, U.; Hild, K. E.; Erdogmus, D.; Roark, B.; Oken, B.; and Fried-Oken, M. 2012. RSVP keyboard: An EEG based typing interface. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 645–648. IEEE.
- Park, J.; and Kim, K.-E. 2012. A POMDP approach to optimizing P300 speller BCI paradigm. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(4): 584–594.
- Park, J.; Kim, K.-E.; and Jo, S. 2010. A POMDP approach to P300-based brain-computer interfaces. In *Proceedings of the 15th international conference on Intelligent user interfaces*, 1–10.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.

- Shamwell, J.; Lee, H.; Kwon, H.; Marathe, A. R.; Lawhern, V.; and Nothwang, W. 2016. Single-trial EEG RSVP classification using convolutional neural networks. In *Micro- and nanotechnology sensors, systems, and applications VIII*, volume 9836, 373–382. SPIE.
- Shannon, C. E. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3): 379–423.
- Smedemark-Margulies, N.; Celik, B.; Imbiriba, T.; Kocanaogullari, A.; and Erdoğan, D. 2023. Recursive estimation of user intent from noninvasive electroencephalography using discriminative models. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Sutter, E. E. 1992. The brain response interface: communication through visually-induced electrical brain responses. *Journal of Microcomputer Applications*, 15(1): 31–45.
- Tresols, J. J. T.; Chanel, C. P.; and Dehais, F. 2023. POMDP-BCI: A benchmark of (re) active BCI using POMDP to issue commands. *IEEE Transactions on Biomedical Engineering*.
- Wierstra, D.; Foerster, A.; Peters, J.; and Schmidhuber, J. 2007. Solving deep memory POMDPs with recurrent policy gradients. In *International conference on artificial neural networks*, 697–706. Springer.
- Wierzgała, P.; Zapała, D.; Wojcik, G. M.; and Masiak, J. 2018. Most popular signal processing methods in motor-imagery BCI: a review and meta-analysis. *Frontiers in neuroinformatics*, 12: 78.
- Zang, B.; Lin, Y.; Liu, Z.; and Gao, X. 2021. A deep learning method for single-trial EEG classification in RSVP task based on spatiotemporal features of ERPs. *Journal of Neural Engineering*, 18(4): 0460c8.
- Zhang, S.; Wang, Y.; Zhang, L.; and Gao, X. 2020. A Benchmark Dataset for RSVP-Based Brain–Computer Interfaces. *Frontiers in Neuroscience*, 14.