LEARNING EVENT OWNERSHIP VIA FINITE-STATE ACTION-PATH MARGINALIZATION FOR IRREVERSIBLE COMMAND DECODING FROM TWO-CHANNEL EOG

Anonymous authors
Paper under double-blind review

ABSTRACT

Continuous facial-gesture control from few-channel electrooculography (EOG) requires distinguishing both action classes and event identities: incoming evidence may continue an existing event or begin a new one. Frame-wise classifiers with fixed post-trigger cooldowns do not explicitly model this distinction and can suppress closely spaced commands. We introduce Action-Constrained Event Decoding (ACED), which separates action classification from event ownership. A learned ownership policy selects among accumulating evidence, emitting a command, and re-arming, replacing fixed cooldown timers. For ambiguous action prefixes, supervision targets an intermediate class and defers the final label until the command grammar permits disambiguation. Training exactly marginalizes over all supervision-compatible paths satisfying an event grammar that requires one emission per supervised event and excludes premature triggers. At inference, the decoder operates causally, never revises past outputs, and guarantees at most one emission per inferred event identity. Experiments show that ACED recovers closely spaced commands suppressed by cooldown baselines and reduces false-trigger rates. The same decoding logic transfers to public datasets with different action taxonomies through task-specific retraining. We further quantify the effects of cross-wear signal drift and address it with robust training recipes.

1 INTRODUCTION

Facial gestures recorded with two periocular EOG electrodes—winks, blink patterns, jaw clenches—can drive discrete commands for assistive and wearable interfaces (Lee et al., 2017; Newn et al., 2023; Schärer et al., 2024). The signal is a continuous stream; the output is a sparse sequence of timestamped, labelled events. Between the two sits a decision that frame-wise classifiers do not make explicitly: whether the evidence arriving now belongs to an event already reported, or begins a new one. We call this decision *event ownership*. In our recordings the gap between consecutive command cues (median 2.4 s) is comparable to the gesture duration (median 1.7 s), so a wrong ownership decision either duplicates a command or silences the next one.

The common practical answer, a fixed post-trigger cooldown, suppresses all detections for an interval that cannot depend on what the classifier is seeing, so any single setting is too long for some gesture pairs and too short for others. Two-hop confirmation with hysteresis (re-arm when the score drops) is a better ownership model, but it commits to a class at confirmation—exactly the moment at which double and triple blinks are indistinguishable. Accurate decoding therefore requires resolving three separable decisions: (1) ownership—does current evidence belong to an existing event or begin a new one; (2) for blink patterns, whether causal evidence yet distinguishes double from triple; (3) when a commit becomes irrevocable. ACED addresses each explicitly: re-arming is conditioned on evidence decay rather than a fixed gap; blink resolution is deferred to a bounded causal window; and all decoders share identical per-hop logits from a frozen scorer, isolating downstream decoding from feature or training variation.

Our contributions are as follows:

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

- **A learned ownership policy (ACED-P)** fitted by exact path marginalisation over a hop-level automaton (Section 3.4): a stochastic FSM over *ready*/*owned*/*pending* states whose 18 parameters replace fixed cooldown timers with data-dependent re-arming.
- **A deferred blink resolution rule (ACED-D)** with bounded label latency (Section 3.5), aligned with a pooled-blink prefix supervision, together with a multi-axis output semantics audit covering event time, commit time, duplicates, and transfer to public datasets.
- **A controlled evaluation protocol and the ACED dataset**: all decoders scored on byte-identical per-fold logits with predictions sealed before labels are read, LOWO and LOSO folds, a pre-registered outside-emission guardrail, and full disclosure of the 17-board development history; 5 participants, 30 wears, 1442 manually verified gesture windows, released with sealed per-fold predictions.

On 30-fold LOWO evaluation, the deferred decoder raises first-correct fraction from 0.641 (hysteresis) to 0.715 on the same logits, and among the 1175 consecutive cross-class cue pairs the fraction with both commands decoded correctly from 0.30–0.32 to 0.48. The learned policy trades that recall gain for a different operating point: 649 outside emissions and 168 rest-cue false activations per hour against 752 and 296 for hysteresis, at an unchanged first-correct fraction of 0.644. No single decoder dominates; the tradeoff is shown in Section 5.4 and the full development history is disclosed in Section 5.6.

2 Related Work

Online and early detection. Online action detection scores each frame of an untrimmed video as it arrives (De Geest et al., 2016); StartNet (Gao et al., 2019) localises action *starts* from the history of per-frame scores. Early classification learns *when* to halt—EARLIEST (Hartvigsen et al., 2019) by reinforcement learning, TEASER (Schäfer & Leser, 2020) by calibrating classifiers per prefix length—and ECOTS (Achenchabe et al., 2023) extends this to open time series containing a succession of events. All of these decide when to commit; none owns events—nothing constrains the next decision to belong to a different event. Our policy is a finite-state model of that constraint.

Sequence criteria over automata. Connectionist Temporal Classification (Graves et al., 2006) marginalises over frame-level alignments via a forward–backward recursion; online variants (Hwang & Sung, 2015) and differentiable WFSTs (Hannun et al., 2020) generalise this. Our policy objective (Section 3.4) fits this family: a forward recursion over an ownership automaton whose accepted paths emit the cued class exactly once per command cue. Two differences matter: CTC is a scorer criterion, whereas we fit only the 18 policy parameters on a frozen scorer; and our automaton carries state (ownership age, pending count) that a label lattice does not.

EOG and physiological command interfaces. EOG has been used for activity recognition (Bulling et al., 2011), discreet gesture input on eyewear (Lee et al., 2017), eye-expression vocabularies (Newn et al., 2023) and on-board eye tracking (Schärer et al., 2024). These works evaluate classifiers on segmented gestures or offline streams; the decoder that turns a continuous score into a command sequence is usually a threshold with a refractory period and is rarely evaluated on its own. Object-detection formulations have been applied to long physiological recordings, e.g. a one-dimensional YOLO for ECG arrhythmia (Hwang et al., 2020), which we re-implement as a transfer baseline. We know of no prior EOG work that fits the ownership decision from data or reports decoder outputs with the timing semantics a downstream consumer experiences.

3 Method

3.1 Problem formulation and metrics

Stream and hops. A wear is a two-channel EOG recording at 250 Hz. A causal frontend produces one feature vector per *hop* of $H{=}50$ samples (0.2 s); hop t ends at sample edge e_t. A backbone maps the features of hops $\leq t$ to scores $z_t \in \mathbb{R}^5$, one per command class $\mathcal{C} = \{$`left_wink_hold`, `right_wink_hold`, `jaw_clench`, `double_blink`, `triple_blink`$\}$, and a *decoder* con-

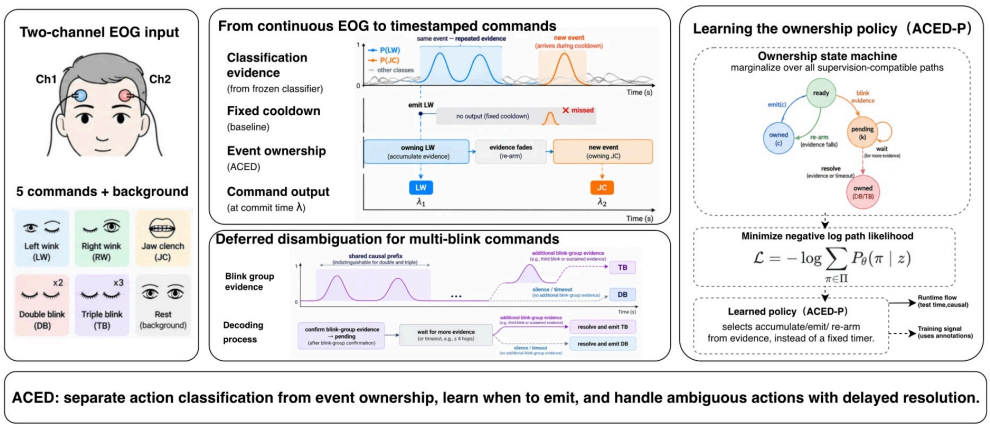


Figure 1: ACED overview. *Left:* two-channel EOG input and five command classes plus background. *Centre top:* a frozen per-hop scorer produces class probabilities; event ownership (ACED-D) re-arms on evidence decay rather than a fixed timer, enabling a new event to be detected before the previous one's evidence fully clears; commit time λ may trail event time τ by up to $0.8\,\mathrm{s}$ for blink commands. *Centre bottom:* deferred disambiguation holds a blink candidate in a *pending* state until additional causal evidence or timeout resolves the final blink class. *Right:* ACED-P fits 18 parameters by maximising the log-likelihood of all supervision-compatible state-machine paths, marginalising over unobserved timing.

sumes $z_{1:t}$ and produces *emissions*. Everything downstream of the frontend is causal and never revises an emission.

Two times per emission. An emission is a triple (τ, λ, c): the *event time* τ (a hop edge) is the timestamp the decoder attaches to the event, fixed once written; the *commit time* $\lambda \geq \tau$ is the hop edge at which the class c is final. For most decoders $\lambda = \tau$; for the deferred decoder $\lambda - \tau \leq 0.8\,\mathrm{s}$ by construction (Section 3.5), and we score it with both timestamps (Section 5.3). Scoring uses the annotated gesture window $[\mathrm{on}_n, \mathrm{off}_n]$: an emission is *inside* window n if $\mathrm{on}_n \leq \tau \leq \mathrm{off}_n$; rest cues and inter-cue gaps are background, and emissions there are *outside* emissions.

Grammar and metrics. We use *grammar* for structural constraints the decoders enforce: (G1) an emission at hop t depends only on samples up to e_t; (G2) emissions are irrevocable; (G3) for threshold decoders, an emission is preceded by a confirmation run of m consecutive hops above threshold (m=2 in 29 of 30 folds); (G4) at most one emission per *inferred* identity. (G4) concerns the decoder's state, not annotated gestures, so we measure duplicates rather than assume their absence. The primary metric is *first-correct fraction*: the fraction of the 1442 windows whose first emission with τ inside the window has the correct label. *Outside emissions* are those with τ in no window; a pre-registered guardrail declares any decoder with more than 900 ineligible. We additionally report outside emissions per background hour, *rest-cue false activations* per hour, and *duplicate windows* (containing two or more emissions).

3.2 Overview

ACED has two levels. The lower level is a causal per-hop scorer (frontend and backbone), shared by every decoder in this paper. The upper level is the ownership decoder: either the learned policy (ACED-P) or the deferred-resolution rule (ACED-D); Figure 1 summarises both.

3.3 Causal scorer

Frontend and backbone. Each channel is filtered with two causal third-order Butterworth bandpasses, 0.3–12 Hz (eye and eyelid movement) and 20–80 Hz (masseter EMG during jaw clenches). Per hop and channel, six compressed primitives (median, change, maximum, minimum of the low band; peak-to-peak amplitude and high-band power relative to trailing lower-quartile references)

give 12 values; the current and 11 preceding hops are concatenated (144 values) and 12 blink-chain descriptors from a trailing 2 s buffer complete the 156-dimensional input; nothing looks ahead (Appendix A). The scorer is a per-hop MLP, $156 \to 128 \to 64 \to 5$ with LayerNorm, GELU and dropout 0.2, on features standardised with source-wear statistics. Its outputs are *independent* logits $z_t \in \mathbb{R}^5$, read as probabilities $p_t = \sigma(z_t)$ by the threshold decoders (no softmax, no explicit background class: background is the state in which all five probabilities are low). A per-fold bias $b \in \{-0.25, 0, 0.25\}^5$ selected on inner-validation wears is added to z_t before decoding.

Training: an adjacent-pair bag loss. Annotation gives, for each command cue, a window and a class, not hop labels. We treat the hops whose edges fall inside the window as a *bag* with class y_n (about 1390 command bags per fold—cues without a usable window are dropped—and 290 rest bags). The loss asks the backbone to place two *consecutive* high logits of the right class somewhere in the bag:

$$\ell_n \;=\; \mathrm{softplus}\Big(-\max_{t:\, t,t+1\in T_n} \min\big(z_t^{y_n}, z_{t+1}^{y_n}\big)\Big) \;+\; 0.6\cdot\frac{1}{4|T_n|}\sum_{c\neq y_n}\sum_{t\in T_n}\mathrm{softplus}\big(z_t^c\big), \qquad (1)$$

where T_n is the set of valid hops of the bag; a rest bag contributes the mean of $\mathrm{softplus}(z_t^c)$ over all five classes. The loss is a multiple-instance surrogate aligned with the two-hop confirmation the threshold decoders use, but it does not guarantee a valid event-output path and says nothing about ownership; that is the decoder's job. Training uses AdamW (learning rate and weight decay 10^{-3}), batches of 64 bags, 80 epochs on the inner-fit wears with the epoch chosen by validation bag loss, then a refit on all 29 source wears for that many epochs (3–33 across folds). All decoders consume the same refit logits.

Blink-prefix supervision. Double and triple blink share an identical causal prefix up to second-blink confirmation; the discriminating evidence—whether a third peak arrives—has not yet appeared. For `double_blink` and `triple_blink` bags we therefore split supervision at a *split point* (first valid hop edge plus two hops, i.e. 0.4 s from onset): hops *before* the split use a pooled-blink loss targeting $\max(z_t^{\texttt{double_blink}}, z_t^{\texttt{triple_blink}})$, hops *at and after* the split use the standard leaf-class loss of Eq. 1. The 0.4 s split is a conservative estimate, not a causal guarantee that every bag's suffix contains discriminating evidence; during the shared prefix the scorer need only signal "blink group". For all other classes the loss is unchanged.

3.4 Learned ownership policy (ACED-P)

States and transitions. The policy is a stochastic automaton over 34 states: *ready*; *owned*(c, a) for $c \in \mathcal{C}$ and age $a \in \{0, \ldots, 4\}$; and *pending*(k) for $k \in \{0, \ldots, 7\}$ (at most 1.6 s of deferral). Let z_t be the biased logits at hop t (clipped to $[-8, 8]$), $m_{t,c} = z_t^c - \max_{i\neq c} z_t^i$ the margin, and $r_{t,c} = z_t^c - z_{t-1}^c$ the rise. From *ready* the policy chooses among "none" and "emit c" with a softmax over scores

$$s_{t,c} = \theta_{\text{emit}} + \theta_c + \theta_{\text{logit}}\, z_t^c + \theta_{\text{margin}}\, m_{t,c} + \theta_{\text{rise}}\, r_{t,c},$$
$$P(\text{emit } c) = \frac{e^{s_{t,c}}}{1+\sum_{c'} e^{s_{t,c'}}}, \qquad P(\text{none}) = \frac{1}{1+\sum_{c'} e^{s_{t,c'}}}. \qquad (2)$$

Emitting a non-blink class moves to *owned*$(c, 0)$ with an emission at e_t. The two blink classes are *not* emitted from *ready*: their combined probability $P(\text{emit } \texttt{double_blink}) + P(\text{emit } \texttt{triple_blink})$ is the probability of entering *pending*(0) without an emission—this is the intermediate "blink-pair" class of the abstract. From *owned*(c, a) the policy re-arms with probability

$$g_{t,c,a} = \sigma\big(\theta_{\text{rearm}} - \theta_{\text{own}}\, z_t^c - \theta_{\text{oth}}\, m_{t,c} + \theta_{\text{age}}\, a\big), \qquad (3)$$

and otherwise moves to *owned*$(c, \min(a{+}1, 4))$; while owned it cannot emit. From *pending*(k), with $d_t = z_t^{\texttt{triple_blink}} - z_t^{\texttt{double_blink}}$, a three-way softmax over "wait" (score 0, disallowed at $k = 7$), "triple" ($\theta_{\text{tri}} + \theta_d\, d_t + \theta_{\text{tri},k}\, k$) and "double" ($\theta_{\text{dbl}} - \theta_d\, d_t + \theta_{\text{dbl},k}\, k$) either waits or resolves; resolving *is* the emission (timestamp e_t, so $\lambda = \tau$ for this policy) and moves to *owned*$(c, 0)$. In total $\theta \in \mathbb{R}^{18}$ (4 emission, 4 re-arm, 5 class-bias, 5 resolution weights). A 13-parameter variant without the pending states (blink classes are emitted directly from *ready*) is the "ownership-only" policy ACED-P13.

Exact path likelihood. For a source wear with cue boundaries and window classes, define the set Π of hop-level trajectories of the automaton that (i) emit the cued class exactly once inside each command cue—inside the annotated window when one exists, otherwise inside the App cue—and emit nothing else there, (ii) emit nothing inside rest cues, and (iii) are unconstrained elsewhere. The objective is the exact negative log-probability of Π under the policy, averaged over cues, with a quadratic prior on θ:

$$\mathcal{L}(\theta) = -\frac{1}{N_{\text{cue}}} \sum_{w} \log \sum_{\pi \in \Pi_w} P_\theta(\pi \mid z^{(w)}) \; + \; 0.05 \, \|\theta - \theta_0\|_2^2. \tag{4}$$

The inner sum is computed by an exact forward recursion (no sampling, no beam) over the 34 states augmented with a one-bit "already emitted" flag: transitions that would emit a wrong class, emit twice, or emit in a rest cue are dropped, and at the end of a command cue only the flagged mass survives. $\mathcal{L}$ is minimised with L-BFGS-B (at most 60 iterations) on the logits of the 25 inner-fit wears; the scorer is not trained through Eq. 4.

Decoding. At inference the policy is run deterministically: in *ready* it emits (or enters pending) when the largest of $\{P(\text{emit } c)\}$ exceeds a threshold, in *owned* it re-arms when $g \geq$ a threshold after a minimum age, in *pending* it resolves when the larger of the two resolution probabilities exceeds a threshold or at $k = 7$. The five thresholds are selected per fold from a grid of 1680 settings on the four inner-validation wears by the utility first-correct $- \, 0.25 \cdot$ outside (Appendix B).

3.5 Deferred blink resolution (ACED-D)

ACED-D is a hand-designed decoder over the same state families with hard rules; it has no trained parameters beyond the backbone bias. Let $m_t = \max(p_t^{\texttt{double_blink}}, p_t^{\texttt{triple_blink}})$ be the blink-group peak.

Confirmation and ownership. Non-blink classes use two-hop hysteresis: the top class must exceed p_{on} on m consecutive hops (selected p_{on}=0.9, m=2 in 29 of 30 folds); the emission is stamped at the second hop, and the decoder is then *owned*: it can emit again only after one hop with the top probability below p_{off}=0.45, after a refractory gap of 3 or 7 hops, or by switching to a class that beats the owner by a margin. The blink group has its own run gate: a hop extends a blink run if a blink class is on top and $m_t \geq 0.75$, or $0.62 \leq m_t < 0.75$ with a margin of 0.05 over the runner-up; two run hops confirm and emit at the second, and the decoder re-arms immediately, so a following gesture is not blocked by the pending window.

Pending window. A blink emission is written with its timestamp τ but an open label, fixed by a second causal pass over hops $\tau{+}1, \ldots, \tau{+}4$. Let n count hops with $m \geq 0.75$, starting at n=1 if $m_\tau \geq 0.75$. A hop k with $m_k \geq 0.75$ increments n only if it passes a *no-decay test*, $m_k \geq m_{k-1} - 0.10$ and $m_k \geq \max_{\tau \leq x < k} m_x - 0.12$; a hop that fails it is neither evidence nor silence. When n reaches 4 the label is `triple_blink`. A hop with $m_k < 0.75$ is silence; three silent hops resolve to `double_blink`, or, if the top class at that hop is a non-blink class with $p \geq p_{\text{on}}$, transfer the emission to that class (relabelled, not duplicated). If the window closes undecided the label is `double_blink`. Hence $\lambda - \tau \in \{0.6, 0.8\}$ s. The no-decay test stops the decaying tail of a second blink from being counted as a third, but it bounds per-hop decrements rather than requiring a resumption: a span such as $0.95, 0.90, 0.86, 0.83$ passes every hop and resolves to a triple without any rise; 6 of the 331 LOWO triples are of this kind (Section 5.3).

3.6 Reference decoders

Fixed cooldown: top class above p_{on} for m hops, then a timer of Δ seconds during which confirmed candidates are dropped, plus one quiet hop to re-arm; $(p_{\text{on}}, \Delta, m)$ selected per fold from $5 \times 8 \times 2$ settings including Δ=0 and 6.4 s. *Hysteresis*: the non-blink rule of ACED-D applied to all five classes. *Blink-run gate*: ACED-D without the pending window (run gate at 0.75, class fixed at confirmation). Each reference decoder receives its own nested selection of thresholds and bias on the same validation wears.

Table 1: 30-fold LOWO on identical per-fold logits. FC = first-correct fraction (windows of 1442); Out = outside-window emissions (guardrail ≤ 900) and per background hour; Rest = false activations per rest-cue hour; Dup = windows with ≥ 2 emissions; per-class first-correct fractions on the right. Sealed board = the board fixed after 17 development iterations, with per-fold mode selection (29 folds ACED-D, one fold blink-peak aggregation).

Decoder	FC	Out	Out/h	Rest/h	Dup	LW	RW	JC	DB	TB
Fixed cooldown	0.619 (893)	904	493	432	33	.843	.541	.815	.653	.254
Hysteresis	0.641 (924)	752	410	296	66	.867	.704	.811	.439	.385
Blink-run gate	0.651 (939)	763	416	216	58	.860	.731	.811	.172	.677
ACED-D (deferred)	**0.715** (1031)	899	490	288	66	.850	.820	.808	.407	.687
sealed board	0.718 (1035)	887	484	288	66	.850	.820	.808	.435	.674
ACED-P13 (no pending)	0.644 (928)	**649**	**354**	**168**	**25**	.864	.813	.762	.056	.715
ACED-P18 (pending)	0.621 (895)	822	448	232	31	.853	.769	.759	.154	.564

4 THE ACED DATASET

Participants and hardware. Five participants (P001, P003, P004, P006, P007) each recorded six wears, re-attaching the electrodes between wears; 30 wears of 4.9–5.6 min, 2.54 h of valid signal, two periocular channels at 250 Hz. P001 and P003 used wet gel electrodes, P004, P006 and P007 dry conductive-foam pads, so electrode type is *confounded with participant* and the wet/dry breakdown in Appendix E is descriptive only.

Cued protocol. Every wear follows the same App-driven script of 60 trials: 50 command cues (10 per class) and 10 rest cues. A cue is displayed for 1.0–3.0 s (median 1.5 s) and cues are separated by 1.6–2.8 s (median 2.4 s). The protocol is *cued*, not self-paced, but the cue schedule is known only to the script: all decoders run without cue information.

Manual windows. For each command cue an annotator marked the onset and offset of the gesture, giving 1442 windows (of 1500 cues; 58 cues have no usable gesture). Windows last $0.82/1.68/2.89$ s at the 10th/50th/90th percentile. Annotation was *candidate-seeded*: an automatic detector proposed candidate onsets; the annotator confirmed, shifted, rejected or added each. Seeding can bias window placement towards what a detector finds easy; we flag this as a limitation shared by every decoder compared.

Classes. `left_wink_hold`, `right_wink_hold` (held winks), `jaw_clench` (masseter activity visible in the 20–80 Hz band), `double_blink` and `triple_blink` (two or three rapid voluntary blinks). Windows per class: 286, 294, 286, 285, 291. The two blink classes share their first two blinks and are the reason for the pending state.

5 EXPERIMENTS

5.1 PROTOCOL

Protocol. The main evaluation is 30-fold leave-one-wear-out (LOWO): one wear is the target, the other 29 the source; four inner-validation wears are drawn from the source by a salted hash of the wear names, the remaining 25 are the inner-fit wears. LOWO measures generalisation to a *new wear of a known participant*; leave-one-subject-out (LOSO, Section 5.4) holds out all six wears of a participant with the same inner rule. Per fold, the backbone epoch and every decoder's thresholds, temporal mode and per-class bias are chosen on the inner wears only (Appendix B); target predictions are written and hashed before any target window is read. All decoders in Table 1 consume *byte-identical* refit checkpoints in every fold, so rows differ only in the decoder and its own nested selection.

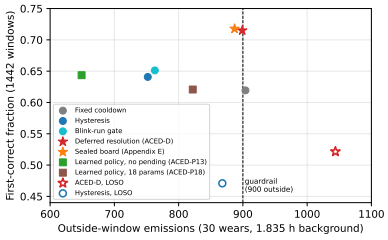


Figure 2: The first-correct/outside frontier. Filled markers: LOWO on identical logits (Table 1); hollow markers: LOSO (Table 4). No decoder is best on both axes.

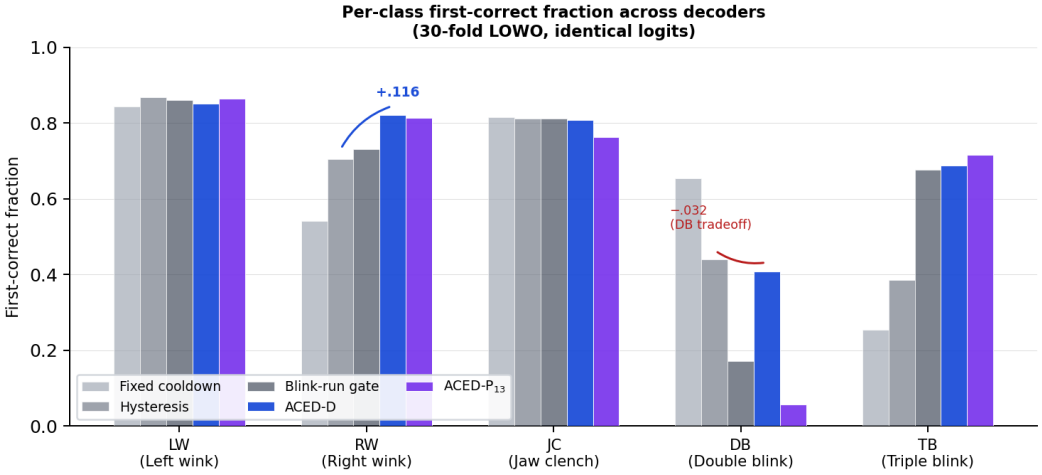


Figure 3: Per-class first-correct fraction for all decoders on identical logits (Table 1). ACED-D gains most on right wink (+.116 vs. hysteresis) and triple blink (+.433 vs. fixed cooldown) but trades off double blink (−.246 vs. fixed cooldown). ACED-P13 suppresses double blink entirely (.056) while improving triple blink.

5.2 Main Results

Closely spaced commands are recovered by deferred resolution. On the same logits, ACED-D decodes 1031 windows first-correct against 924 for hysteresis (+107 windows), 939 for the blink-run gate and 893 for fixed cooldown (Wilcoxon vs. hysteresis $p = 1.3 \times 10^{-5}$). The gain is where ownership is hardest: `right_wink_hold` (.704 → .820) and `triple_blink` (.385 → .687) (Figure 3); `double_blink` is *not* improved over hysteresis (.439 → .407) and falls below fixed cooldown (.653), whose short timers simply re-fire on the second blink; ACED-D trades doubles for triples. The pair analysis makes the claim concrete: among the 1175 consecutive cross-class cue pairs, both commands are first-correct for 0.478 under ACED-D against 0.303 (hysteresis) and 0.323 (fixed cooldown) (Figure 8, Appendix F); the ablation (Figure 4) shows the largest single step (+.064) comes from adding the pending window. Fixed cooldown (selected timers 0.5–1.5 s) blocked 223 confirmed candidates inside its timer.

False triggers are reduced by the learned policy, not by deferred resolution. ACED-D emits *more* outside events than hysteresis (899 vs. 752) and as many as fixed cooldown (904). The learned policy without a pending state (ACED-P13) is the only decoder that lowers every false-trigger measure at once—649 outside emissions, 168 rest-cue false activations per hour, 25 duplicate windows—at the first-correct fraction of hysteresis (0.644). Utility-matched hysteresis reaches 632 outside emissions, so much of the reduction arises from calibration; the learned policy retains advantages in rest-cue activations. It re-arms on evidence rather than a timer: the fitted re-arm probability falls with the owner's current logit (median weight 0.45) while the age weight stays at 0

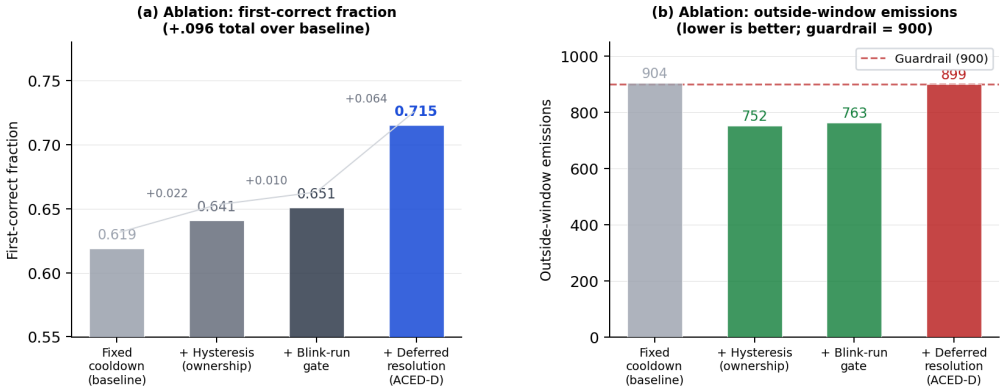


Figure 4: Ablation: each decoder adds one mechanism over the previous. (a) First-correct fraction rises $+.096$ from fixed cooldown to ACED-D; the largest single step ($+.064$) comes from adding the pending window. (b) Outside emissions: adding hysteresis already reduces outside events (-152), but the pending window trades that back (+147 vs. hysteresis) to recover closely-spaced commands; ACED-P13 (not shown here) achieves 649 outside with nearly the same first-correct as hysteresis.

(Appendix C). The price is `double_blink` (.056): of the 285 double-blink windows, 172 receive no emission inside the window (107 committed just after the offset, mostly as triples) and 63 are missed outright. Adding the pending node (ACED-P18) balances the two blink labels (44 doubles, 44 triples, 191 empty) but loses on every other axis (0.621, 822 outside): its learned resolution is weaker than the hand-set four-hop rule of ACED-D. No single decoder dominates; Figure 2 shows the frontier, and the choice between ACED-D and ACED-P13 is an application decision.

5.3 Output semantics audit

Event time versus commit time. Of the 2209 ACED-D emissions, 1195 are blink emissions whose label is open at τ and resolved 0.6 s later in 634 cases and 0.8 s later in 561 (331 triples, 360 doubles by silence, 362 by window cap, 142 transfers); Figure 5(a) shows the emission-time ECDF and (b) the fitted re-arm probability. For the 320 first-correct blink emissions (124 double, 196 triple; onset–offset window) the label became available a median 1.34 s after onset and a median 0.18 s after offset (66% after the offset). Re-scoring with λ gives first-correct 0.542 with 1236 outside emissions (`double_blink` .018, `triple_blink` .247); non-blink classes change negligibly. We report τ as primary because it is the decoder's event time and the latency is bounded and disclosed; a deployment that scores on λ should expect the second set.

Duplicates and transfers. The (G4) guarantee holds per inferred identity, not per gesture: 66 of the 1442 windows (4.6%) contain two or more ACED-D emissions (hysteresis 66, fixed cooldown 33, ACED-P13 25). The pending window does not block later events (26 non-blink emissions were committed inside one), and 36 of the 142 ownership transfers relabelled an emission to a class that was also emitted separately. Six of the 331 triples were resolved from spans whose blink peak never rose after τ, the failure case of the no-decay test (Section 3.5). Causality check: re-running every decoder on prefixes of each wear cut at 30%, 60% and 85% of its length reproduced every emission whose label-available hop lay inside the prefix (0 failures over 90 cuts).

5.4 Generalisation across wears and participants

ACED-D achieves the highest LOWO first-correct fraction for four of five participants (Appendix E); under LOSO both decoders lose ≈ 0.19 absolute and ACED-D's 1044 outside emissions exceeds the 900 guardrail. Replacing the trailing-quartile reference with a fixed floor reduces first-correct from 0.715 to 0.651 and increases outside emissions from 899 to 1034: wear-relative normalisation is load-bearing.

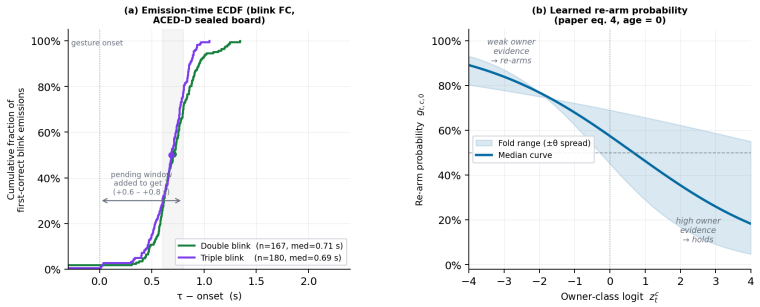


Figure 5: **(a)** ECDF of emission time $\tau - \text{onset}$ for first-correct double-blink (DB, n=167, median 0.71 s) and triple-blink (TB, n=180, median 0.69 s) emissions on the sealed board (cue-window scoring; the onset–offset window gives n=124 DB and n=196 TB, reported in Section 5.3); the shaded band marks the pending window, and adding 0.6–0.8 s gives the commit time λ. **(b)** Analytic re-arm probability $g_{t,c,0} = \sigma(\theta_{\text{rearm}} - \theta_{\text{own}}\, z_t^c)$ (eq. 3, age=0) as a function of the owner-class logit, for the median fold parameters (solid curve) and the fold-parameter range (ribbon).

5.5 Transfer to a public dataset

We apply the unchanged ownership decoder to two public datasets. On PhysioMotion (Yang et al., 2026) (blink vs. chewing, subjects 7–30 as test, Appendix J), ACED trades recall for precision: under a 10/h budget it covers 33.9% of blocks at 0.6 false activations/h vs. 65.9% at 3.5/h for the CNN; under a 60/h budget ACED covers 83.8% at 9.4/h vs. 65.9% (CNN) and 54.1% at 25/h (YOLO-1D (Hwang et al., 2020)). On OPPORTUNITY (Chavarriaga et al., 2013) (4 subjects, 17 gesture classes, 30 Hz), a task-adapted causal MLP drives ACED-simple to 45.2% event recall (1135/2511) vs. hysteresis 33.7% (847/2511) at 105 vs. 260–270 background firings per fold, confirming the ownership mechanism transfers through scorer retraining alone.

5.6 Development history and what it implies

The 17-board outer development loop makes ACED-D's LOWO number optimistic by an unknown amount; the concern is bounded by three facts: the reference decoders received the same nested selection on the same logits; the decoder ordering is reproduced by the single-shot LOSO run; and the semantics audit of Section 5.3 ran once, after the last board.

6 Limitations and Conclusion

Limitations. Five participants is a small population; electrode type is confounded with participant; the cued protocol makes gesture timing more regular than free use. Annotation is candidate-seeded and single-annotator; the deferred decoder is the product of 17 LOWO boards, so its LOWO number is optimistic by an unknown amount and its LOSO number fails our own guardrail. Transfer is shown on two public datasets, neither involving free-paced use, against threshold decoders on the same scorer rather than learned or end-to-end models.

Conclusion. Treating event ownership as an explicit decision changes what a continuous EOG command decoder can do: on identical scores, deferring the double/triple decision recovers 107 of 1442 commands and raises consecutive-cue accuracy from 0.30 to 0.48; the learned policy reaches 14% fewer false triggers than hysteresis, though utility-matched hysteresis shows much of that reduction is attributable to calibration. We release the data and sealed predictions so that the frontier in Figure 2 can be moved by others.

Ethics statement

Local recordings were collected with consent from the authors and released for research; public datasets are used under their published licences.

Reproducibility statement

Every number in Tables 1 and 4 is computed from sealed per-fold prediction files (SHA-256 hashes recorded before target windows were read). Scorer checkpoints, selection logs, policy parameters, the LOSO run, and the development board list will be released with the data.

AI use statement

Large language models assisted with drafting, revision, literature search, and experimental-design feedback; a generative image tool assisted with the schematic in Figure 1. The authors are responsible for all methodology, results, claims, and the final manuscript.

References

Y. Achenchabe, A. Bondu, A. Cornuéjols, and V. Lemaire. When to classify events in open times series? In *Proceedings of the Asian Conference on Machine Learning*, volume 189 of *Proceedings of Machine Learning Research*, 2023. URL https://proceedings.mlr.press/v189/achenchabe23a.html.

Andreas Bulling, Jamie A. Ward, Hans Gellersen, and Gerhard Tröster. Eye movement analysis for activity recognition using electrooculography. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(4):741–753, 2011. doi: 10.1109/TPAMI.2010.86.

Ricardo Chavarriaga, Hesam Sagha, Alberto Calatroni, Sundara Tejaswi Digumarti, Gerhard Tröster, José del R. Millán, and Daniel Roggen. The opportunity challenge: A benchmark database for on-body sensor-based activity recognition. *Pattern Recognition Letters*, 34(15): 2033–2042, 2013.

R. De Geest, E. Gavves, A. Ghodrati, Z. Li, C. Snoek, and T. Tuytelaars. Online action detection. In *European Conference on Computer Vision*, 2016. URL https://arxiv.org/abs/1604.06506.

M. Gao, M. Xu, L. S. Davis, R. Socher, and C. Xiong. StartNet: Online detection of action start in untrimmed videos. In *IEEE/CVF International Conference on Computer Vision*, 2019. URL https://openaccess.thecvf.com/content_ICCV_2019/papers/Gao_StartNet_Online_Detection_of_Action_Start_in_Untrimmed_Videos_ICCV_2019_paper.pdf.

Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd International Conference on Machine Learning*, pp. 369–376. ACM, 2006.

Awni Hannun, Vineel Pratap, Jacob Kahn, and Wei-Ning Hsu. Differentiable weighted finite-state transducers. arXiv preprint arXiv:2010.01003, 2020.

T. Hartvigsen, C. Sen, X. Kong, and E. Rundensteiner. Adaptive-halting policy network for early classification. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019. URL https://web.cs.wpi.edu/~xkong/publications/papers/kdd19.pdf.

Kyuyeon Hwang and Wonyong Sung. Online sequence training of recurrent neural networks with connectionist temporal classification. arXiv preprint arXiv:1511.06841, 2015.

Won Hee Hwang, Chan Hee Jeong, Dong Hyun Hwang, and Young Chang Jo. Automatic detection of arrhythmias using a YOLO-based network with long-duration ECG signals. In *Proceedings of the 7th International Electronic Conference on Sensors and Applications (ECSA-7), Engineering Proceedings*, volume 2, pp. 84, 2020. doi: 10.3390/ecsa-7-08229.

Juyoung Lee, Hui-Shyong Yeo, Murtaza Dhuliawala, Jedidiah Akano, Junichi Shimizu, Thad Starner, Aaron Quigley, Woontack Woo, and Kai Kunze. Itchy nose: discreet gesture interaction using EOG sensors in smart eyewear. In *Proceedings of the 2017 ACM International Symposium on Wearable Computers (ISWC)*, pp. 94–97, 2017.

Joshua Newn, Sophia Quesada, Baosheng James Hou, Anam Ahmad Khan, Florian Weidner, and Hans Gellersen. Exploring eye expressions for enhancing EOG-based interaction. In *Proceedings of the 19th IFIP TC13 International Conference on Human-Computer Interaction (INTERACT 2023), Part IV*, pp. 68–79, 2023.

P. Schäfer and U. Leser. TEASER: Early and accurate time series classification. *Data Mining and Knowledge Discovery*, 2020. doi: 10.1007/s10618-020-00690-z.

Nicolas Schärer, Federico Villani, Aishwarya Melatur, Steven Peter, Tommaso Polonelli, and Michele Magno. ElectraSight: smart glasses with fully onboard non-invasive eye tracking using hybrid contact and contactless EOG. arXiv preprint arXiv:2412.14848, 2024.

Chunfeng Yang, Jiangwei Yu, Aonan He, Wentao Xiang, Xi Wang, Guangquan Zhou, Yudong Zhang, Miao Cao, Yang Chen, and Juan M. Gorriz. PhysioMotion artifact: A task-driven EEG dataset with point-wise motion artifact annotations. *Scientific Data*, 13:807, 2026. doi: 10.1038/s41597-026-07120-7.

A FRONTEND PRIMITIVES

Both channels are filtered causally (third-order Butterworth, second-order sections) into a low band (0.3–12 Hz) and a high band (20–80 Hz). Per hop of 50 samples and per channel, six primitives are computed from the hop's samples and trailing references (Table 2); the 12 values of the current hop and of the 11 preceding hops form the 144-dimensional history block. The blink-chain block uses a 0.5–10 Hz causal filter on the trailing 500 samples: peaks are detected with a minimum distance of 45 samples and a prominence of at least $\max(40,\ 1.5 \times \text{quiet level})$, the longest chain of peaks with inter-peak gaps in $[0.18, 0.55]$ s is found, and for each channel and polarity three values are stored (chain length divided by 3, largest prominence relative to threshold divided by 5, recency of the last peak). Compression functions ($\operatorname{arcsinh}$, $\log(1+x)$) keep the dynamic range bounded across electrodes.

Table 2: Per-hop, per-channel primitives (12 per hop for two channels).

Primitive	Definition
low median	$\operatorname{arcsinh}(\operatorname{median}(\text{low})/400)$
low change	$\operatorname{arcsinh}(\Delta\,\text{low}/200)$ (last minus first sample of the hop)
low max, low min	$\operatorname{arcsinh}(\max(\text{low})/400)$, $\operatorname{arcsinh}(\min(\text{low})/400)$
relative peak-to-peak	$\log(1 + \text{p2p}(\text{low})/\max(25, q_{25}))$, q_{25} = lower quartile of p2p over the preceding 20 hops
relative high power	$\log(1 + \text{rms}(\text{high})/\max(3, q'_{25}))$, q'_{25} = lower quartile over the preceding 20 hops

B SELECTION GRIDS AND UTILITIES

Threshold decoders. The hysteresis configuration grid has 120 settings: $p_{\text{on}} \in \{.35, .5, .65, .8, .9\}$, $p_{\text{off}} = p_{\text{on}} \times \{.5, .8\}$, refractory gap $\in \{0, 3, 7\}$ hops, confirmation $m \in \{1, 2\}$, switch margin $\in \{0, .15\}$. The configuration is selected on the inner-validation wears by the rank (rest-cue false activations $= 0$, hits, $-$rest false activations, $-$wrong-first, $-$overfire); the temporal mode is then selected by hits $+\,0.5\times$blink hits; the per-class bias is selected by two rounds of coordinate ascent over $\{-0.25, 0, 0.25\}^5$ with the same rank and a tie-break towards smaller $\|b\|_1$. Selected in 29 of 30 folds: p_{on}=0.9, p_{off}=0.45, m=2; gap 3 or 7 hops. The fixed-cooldown grid has 80 settings ($p_{\text{on}} \in \{.35, .5, .65, .8, .9\}$, $\Delta \in \{0, .2, .5, 1, 1.5, 2.5, 4, 6.4\}$ s, $m \in \{1, 2\}$; $p_{\text{off}} = \min(0.8p_{\text{on}}, p_{\text{on}} - 0.05)$), selected by the same rank.

Learned policy. Decoding thresholds are selected from $\{.08, .12, .18, .25, .35, .5, .65\}$ (emit) $\times\{.3, .5, .7, .85\}$ (re-arm) $\times\{0, 1, 2, 4\}$ (minimum owned hops) $\times\{.5, .6, .75, .9, 1.01\}$ (resolve) $\times\{1, 1.4, 1.8\}$ (blink-entry scale) by the utility first-correct $- 0.25\times$outside on the inner-validation wears. For ACED-P13 the grid instead includes an inference-time deferral of a blink emission by at most $h \in \{1, 2, 3, 4, 6\}$ hops until $z^{\texttt{triple_blink}} - z^{\texttt{double_blink}} \geq t \in \{0, .5, 1, 2\}$, or no deferral; no deferral was selected in 14 folds, h=1 in 13 and h=2 in 3. Both policies add the per-fold bias b that the nested threshold-decoder selection of the first development board chose for that fold. ACED-P13 was fitted with App-cue constraints (the cued class must be emitted exactly once inside the cue), ACED-P18 with manual-window constraints for command cues; on App-cue constraints ACED-P18 scores 0.606 (874) with 801 outside emissions, and with the bias of the iteration-15 board instead it scores 0.627 (904) with 820.

C FITTED POLICY PARAMETERS

Table 3 lists the median, minimum and maximum of each fitted parameter across the 30 folds. L-BFGS-B converged in a median of 15–17 iterations, reducing the per-cue objective from 2.03 to 0.93 (ACED-P13) and from 1.72 to 0.79 (ACED-P18). Margin and age terms of the emission and re-arm scores stayed at or near their lower bound of zero in ACED-P13: the fitted policy is essentially "emit when the logit is high, re-arm when the owner's logit has fallen".

Table 3: Fitted parameters (median [min, max] over folds). Bounds in parentheses.

Parameter	ACED-P13	ACED-P18
θ_{emit} $(-8, 2)$	-2.58 $[-2.94, -2.40]$	-2.81 $[-3.17, -2.62]$
θ_{logit} $(0, 4)$	0.68 [0.57, 0.93]	0.80 [0.73, 0.89]
θ_{margin} $(0, 3)$	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]
θ_{rise} $(0, 2)$	0.03 [0.00, 0.14]	0.02 [0.00, 0.31]
θ_{rearm} $(-8, 4)$	-1.40 $[-1.54, -1.30]$	-1.20 $[-1.31, -1.12]$
θ_{own} $(0, 3)$	0.45 [0.26, 0.69]	0.50 [0.42, 0.65]
θ_{oth} $(0, 3)$	0.00 [0.00, 0.10]	0.17 [0.08, 0.23]
θ_{age} $(0, 1.5)$	0.00 [0.00, 0.01]	0.06 [0.00, 0.19]
θ_c (LW, RW, JC, DB, TB) $(-3, 3)$	$-.11$, $.10$, $-.09$, $-.41$, $-.03$	$-.17$, $-.11$, $-.27$, $-.15$, $-.10$
$\theta_{\text{tri}}, \theta_d, \theta_{\text{tri},k}$	—	-1.01, 1.15, 0.31
$\theta_{\text{dbl}}, \theta_{\text{dbl},k}$	—	-0.80, 0.30

D CLASS CONFUSION MATRIX

Figure 6 shows the row-normalised confusion matrix of ACED-D on the sealed board. Left-wink, right-wink, and jaw-clench each achieve recall above 0.80; double-blink is the main source of confusion, with 61 of 285 windows mislabelled as triple-blink inside the cue.

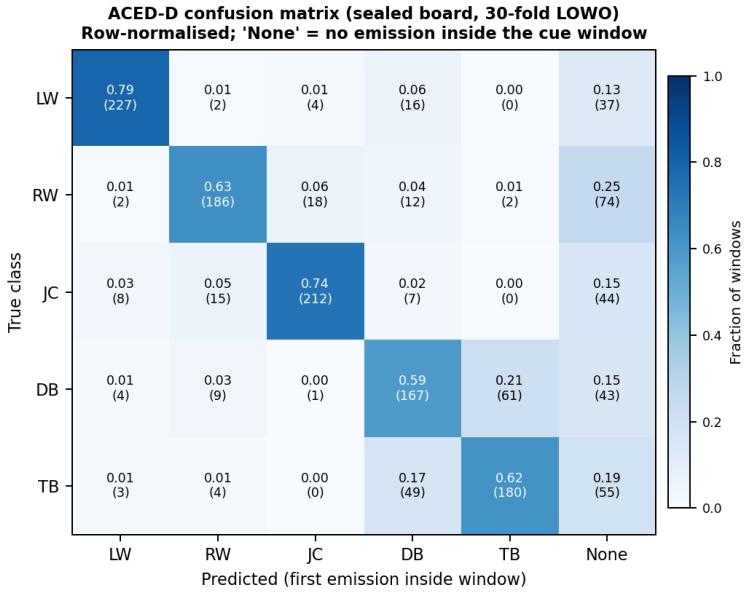


Figure 6: Row-normalised confusion matrix of ACED-D on the sealed board (30-fold LOWO). Each cell shows the fraction and raw count; "None" = no emission inside the cue window.

E PER-PARTICIPANT RESULTS

Under LOSO, the per-class first-correct fractions of ACED-D are `left_wink_hold` .710, `right_wink_hold` .514, `jaw_clench` .640, `double_blink` .277 and `triple_blink` .467.

F CONSECUTIVE-PAIR SUCCESS RATE VS INTER-CUE GAP

Figure 8 shows the consecutive-pair success rate (both commands first-correct) among the 1175 cross-class pairs with inter-cue gaps of 1.6–2.8 s: ACED-D reaches 0.478 against 0.303 for hysteresis; the dashed reference shows the single-event first-correct (0.674), and the remaining gap reflects emission-time and ownership effects that apply even to well-separated commands.

Table 4: Per-participant first-correct fraction. LOWO: identical logits as Table 1. LOSO: the full recipe retrained with one participant held out (single run, no iteration). W = wet gel, D = dry foam; electrode type is confounded with participant.

			LOWO			LOSO	
Participant	Electr.	Fixed cd.	Hysteresis	ACED-P13	ACED-D	Hysteresis	ACED-D
P001	W	.645	.652	.692	.677	.468	.525
P003	W	.557	.582	.604	.704	.357	.404
P004	D	.617	.621	.614	.655	.441	.497
P006	D	.653	.656	.687	.780	.577	.601
P007	D	.622	.689	.622	.756	.505	.575
All (FC)		.619	.641	.644	.715	.471	.522
Outside		904	752	649	899	868	1044

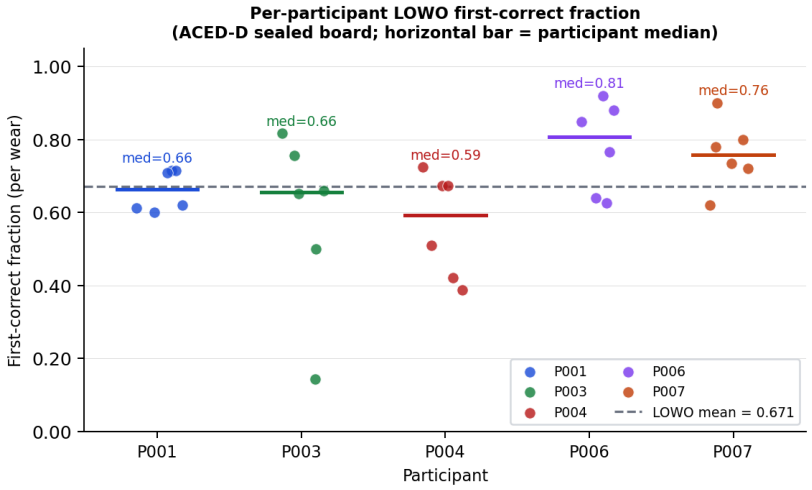


Figure 7: Per-wear first-correct fraction for each participant under 30-fold LOWO (ACED-D, sealed board). Each dot is one held-out wear; horizontal bars mark the participant median. The figure shows wear-to-wear variability within each participant; cross-decoder comparisons are in Table 4.

Figure 9 stratifies consecutive-pair success by the inter-cue gap (time from previous cue end to next cue start). Pairs with gaps below 1.5 s are rare (n=0 in the sealed board). For the most common gaps of 1.5–3.5 s, ACED-D recovers roughly 45–51% of pairs. Notably, the rate does not continue rising for gaps above 3.5 s (roughly 45%; the >5 s bin has n=2 and should not be over-interpreted), suggesting that gap alone does not determine pair success and that per-command emission-time effects set a ceiling independent of ownership state.

G Blink window outcome breakdown

Figure 10 shows how the 285 double-blink and 291 triple-blink windows end: the dominant failure mode for double blinks is mislabelling inside the window (61 windows, 21%), while for triple blinks the dominant failure is a late commit — the decoder emits outside the annotated window (36 windows, 12%), which reduces first-correct even though the event was detected. Table 1 reports per-class first-correct fractions; this figure gives the full outcome decomposition.

H Class-pair consecutive success

Figure 11 stratifies the consecutive-pair success rate by the classes of the two commands. Among cells with at least five consecutive pairs, LW→JC achieves the highest rate (0.63, n=68) and RW→DB the lowest (0.36, n=67); any-class pairs ending in DB are consistently hard. The diagonal

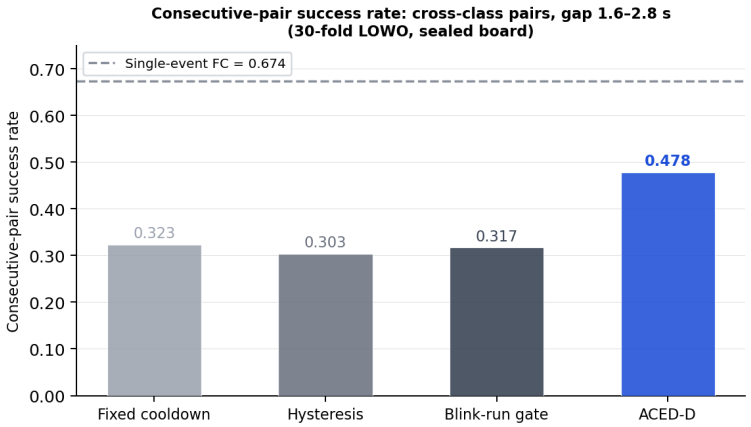


Figure 8: Consecutive-pair success rate (both commands first-correct) among the 1175 cross-class pairs with inter-cue gaps of 1.6–2.8 s. ACED-D reaches 0.478, a gain of +.175 over hysteresis.

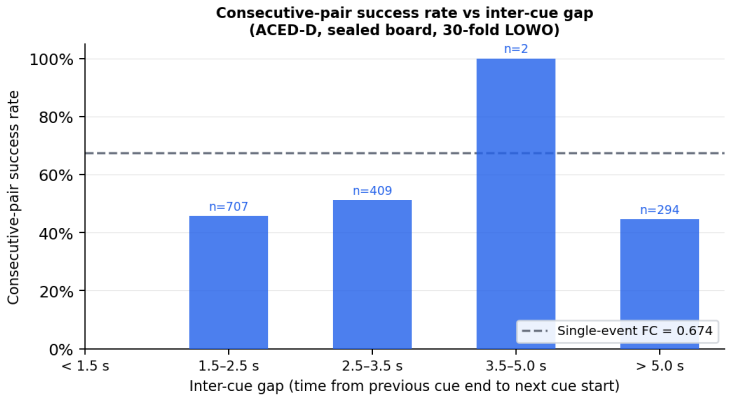


Figure 9: Consecutive-pair success rate (both events correct) as a function of inter-cue gap, ACED-D sealed board (all consecutive command pairs). The dashed line marks the single-event first-correct fraction for reference; sample counts appear above each bar. Success rate does not increase with gap beyond 1.5 s (the >5 s bin has n=2), indicating that emission-time effects rather than ownership state are the binding constraint for well-separated commands.

(same-class repeats) is sparse because the protocol discourages consecutive identical gestures; cells with $n \leq 10$ are shown but their rates should be read with caution.

I Raw EOG waveforms

Figure 12 shows one representative action-phase segment per gesture class from participant P001, wear 1. The low-band (0.3–12 Hz) component captures the slow corneoretinal potential shifts that drive the scorer; the high-band RMS envelope (shown for jaw clench) captures the EMG burst that distinguishes jaw clench from blink events. The shaded region marks the annotated action phase; dashed vertical lines show its boundaries.

J Transfer experiment

PhysioMotion (Yang et al., 2026) contains 30 subjects with six recordings each and point-wise annotations of motion artifacts, among them blink blocks and chewing blocks. We treat block detection as a two-class command task: the ownership decoder is unchanged, the scorer is retrained per ar-

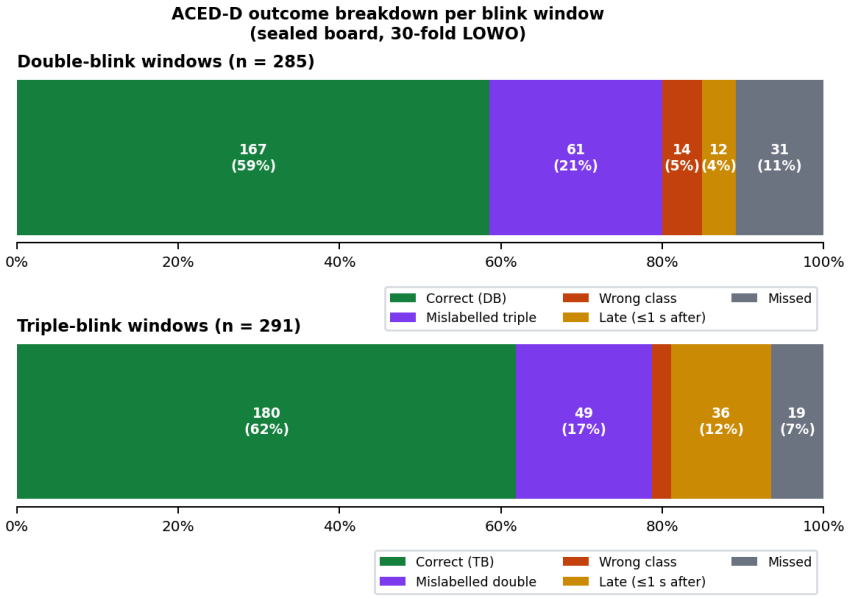


Figure 10: Outcome breakdown for double-blink (top, n=285) and triple-blink (bottom, n=291) windows, ACED-D sealed board. "Mislabelled" means the other blink class was emitted inside the window; "Late" means a blink emission arrived within 1 s after the window closed; "Wrong class" means a non-blink class was emitted inside the window.

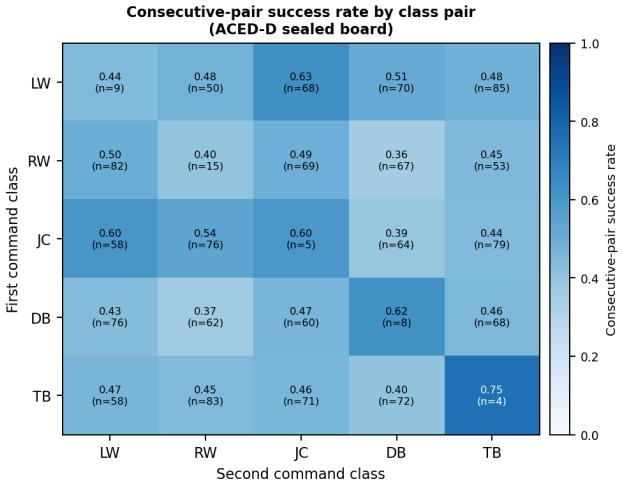


Figure 11: Consecutive-pair success rate by class pair (row = first command, column = second command), ACED-D sealed board (all consecutive command pairs, no gap restriction). Each cell shows the rate and sample count; white cells have no consecutive pairs of that type in the data. Among cells with $n \geq 5$, LW→JC achieves the highest rate (0.63, n=68) and RW→DB the lowest (0.36, n=67); pairs whose second command is DB are consistently harder than other combinations.

chitecture on subjects 1–6, and thresholds are selected on those subjects under a rest false-activation budget of 10/h (primary) or 60/h (secondary). Subjects 7–30 (144 files) form the fixed test set. Coverage is macro block coverage (at least one correct-class emission inside a block); wrong-first is the fraction of blocks whose first emission has the wrong class.

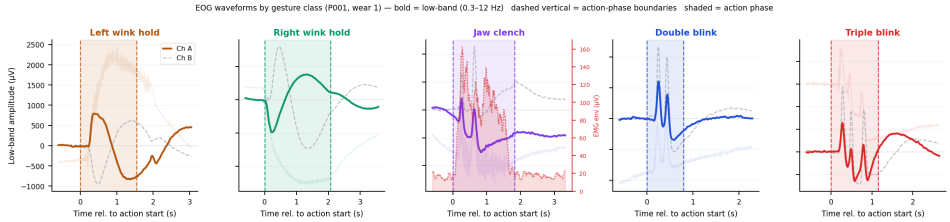


Figure 12: Representative EOG waveforms for each gesture class (P001, wear 1). Bold traces: low-band filtered (0.3–12 Hz, causal Butterworth); faint traces: raw signal; shaded region: annotated action phase; dashed vertical lines: action-phase boundaries. For jaw clench, the right axis shows the high-band (20–80 Hz) RMS envelope, which distinguishes it from blink classes. Time is relative to the action-phase start as marked in `markers.csv`.

Table 5: PhysioMotion block detection (subjects 7–30, 144 files).

Scorer	Budget	Coverage	Wrong-first	Blink	Chewing	Rest FA/h
ACED	10/h	0.339	0.000	184/720	610/1445	0.6
ACED	60/h	**0.838**	0.003	595/720	1227/1445	9.4
CNN1D	10/h	**0.659**	0.000	462/720	978/1445	3.5
CNN1D	60/h	0.659	0.000	462/720	978/1445	3.5
YOLO-1D	10/h	0.034	0.000	1/720	96/1445	0.8
YOLO-1D	60/h	0.541	0.005	474/720	613/1445	25.0