Beyond the Threshold: An Adaptive Markov Decision Layer for Graded, Personalized Seizure Alerting

First Author Name [1], Second Author Name [2]*

Abstract

A persistently high false alarm rate (FAR) remains a primary barrier to the clinical adoption of EEG-based seizure prediction, even as classification accuracy has improved. Most systems convert prediction outputs into alerts using an instantaneous rule, a fixed threshold or majority vote applied at each step, ignoring signal trajectory and the asymmetric costs of missed seizures versus false alarms. This work introduces an adaptive, model-agnostic clinical decision layer formulated as a Markov Decision Process (MDP) that operates on top of any fixed, pre-trained prediction model. Instead of retraining the classifier, it replaces the instantaneous rule with a sequential, cost-aware policy that governs when and how strongly to alert, enabling graded soft/hard alerting and personalization through a single cost parameter. Evaluated on CHB-MIT under a leave-one-seizure-out protocol against threshold, k-of-n, and two-level baselines at matched sensitivity, the framework reduced FAR by 58.6%, 66.7%, and 16.8% for three of four screened patients and matched the best baseline for the fourth, without modifying the underlying model.

Keywords: Epileptic seizure prediction; Markov decision process; False alarm rate; EEG signal processing; Graded alerting; Clinical decision support.

---

[1] Undergraduate Author, Department of _____, King Fahd University of Petroleum & Minerals, Dhahran 31261, Saudi Arabia. (*Corresponding Author; Email: u20*******@kfupm.edu.sa; ORCID XXXX-XXXX-XXXX-XXXX).

[2] Research Mentor, Professor, Department of _____, King Fahd University of Petroleum & Minerals, Dhahran 31261, Saudi Arabia. (Email: author.email@kfupm.edu.sa; ORCID XXXX-XXXX-XXXX-XXXX).

Introduction

Epilepsy is a chronic neurological disorder with profound neurobiological, cognitive, and social consequences[1]. Seizure unpredictability has driven decades of research into prediction systems that warn before onset, formalized through a seizure prediction horizon (SPH) and occurrence period (SOP)[2]. EEG segments are categorized as interictal, preictal (the SOP-plus-SPH window before onset), or ictal.

Deep learning architectures have achieved strong classification accuracy on benchmarks such as CHB-MIT[3], yet the false alarm rate (FAR) remains a persistent barrier to clinical adoption that accuracy alone does not resolve[4]. Frequent false warnings erode patient trust and are repeatedly cited as a primary obstacle to clinical translation[5].

Much of this stems from how prediction outputs become alerts. Most systems apply an instantaneous rule, a fixed threshold or majority vote (k-of-n), evaluated independently at each step, ignoring signal trajectory and the differing costs of a missed seizure versus a false alarm. Such rules cannot adapt to clinician risk tolerance without manual re-tuning and offer no graded response.

This work introduces an adaptive, model-agnostic clinical decision layer, formulated as a Markov Decision Process (MDP), operating on top of any fixed prediction model to govern when and how strongly to alert. Rather than modifying the classifier, the framework replaces the instantaneous rule with a sequential, cost-aware policy supporting graded alerting and personalization. This study formalizes the decision layer, evaluates it against fixed-rule baselines under matched conditions, and characterizes when sequential decision-making offers a measurable advantage. The contributions are fourfold: a sequential, cost-aware alerting policy

replacing instantaneous rules; personalization via a single cost parameter without retraining; graded soft/hard alerting driven by the learned policy; and competitive performance atop a fixed, unmodified prediction model.

Literature Review / Related Work

2.1. EEG-Based Seizure Prediction: Classifiers and Architectures

Deep learning approaches, from CNN-LSTM and CNN-TCN hybrids[6,3] to attention-based and CBAM-based architectures[8,9,10], have improved window-level accuracy on CHB-MIT, though performance varies with preprocessing[7]. These studies report sensitivity and AUC rather than deployment-oriented criteria.

2.2. The False Alarm Rate as a Persistent Clinical Gap

FAR remains disproportionately high for real-world deployment despite improved accuracy[4], is rarely reported in standardized units[11], and is complicated by inconsistent SPH/SOP and significance-testing protocols[12,13].

2.3. Post-Processing and Decision-Layer Approaches

A smaller body of work targets FAR through the decision layer. The dominant approach remains a fixed threshold or majority-vote rule (k-of-n), originally proposed for CNN-based CHB-MIT prediction with 81.2% sensitivity and 0.16 false predictions per hour[14]. Soft-fusion post-processing under the same SPH/SOP configuration adopted here has improved on k-of-n while still reporting FAR above clinically acceptable levels[15], and artifact-suppression methods for wearable EEG complement rather than replace decision-layer refinements[16]. In all cases, the

decision layer remains a fixed rule evaluated independently at each time step, without regard to signal trajectory or the differing costs of missed seizures and false alarms.

2.4. Markov Decision Processes in Medical and Safety-Critical Systems

MDPs formalize sequential decision-making by minimizing expected cumulative cost over time[17] and have been applied to treatment-interval decisions[18], emergency resource allocation[19],

Gaps

Three limitations emerge from the literature above. First, existing systems convert classifier outputs into alerts via an instantaneous rule, a fixed threshold or k-of-n vote evaluated independently at each step, with no mechanism for signal trajectory or sequential decision-making[14,15]. Second, reported FAR values are inconsistent due to varying SPH/SOP conventions and evaluation protocols[4,11,12,13]. Third, no reviewed approach optimizes the trade-off between missed-seizure and false-alarm costs; alerting decisions remain binary and uniform across patients, with no mechanism for graded response or personalization.

Problem Statement

The limitations identified above converge on a single problem: no existing seizure alerting mechanism treats the alert decision as a sequential, cost-aware choice that weighs missed-seizure risk against false-alarm burden, incorporates recent signal trajectory, or personalizes to clinical preference without manual re-tuning. We address this by introducing an Adaptive Clinical Decision Layer: a Markov Decision Process operating on top of any fixed, pre-trained prediction model, governing when and how strongly to alert, with all adaptation confined to this layer while the prediction model itself remains untouched.

Methodology

5.1. System Overview

The proposed framework consists of two functionally separate layers, as shown in Figure 1. The first is a fixed, pre-trained prediction layer that outputs a continuous seizure probability at each step and is never retrained across patients or conditions. The second is an adaptive decision layer, formulated as a Markov Decision Process, that consumes the smoothed probability and its recent trajectory to issue one of three alerting actions. This separation is deliberate: all adaptation occurs exclusively in the decision layer, while the prediction model remains model-agnostic and untouched.

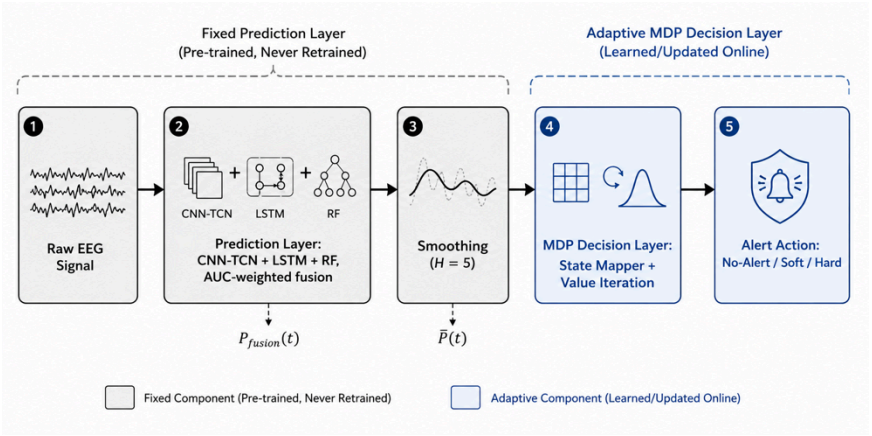

Figure 1. System architecture of the proposed two-layer framework.

5.2 Dataset and Signal Preprocessing

The proposed framework was evaluated on the CHB-MIT Scalp EEG Database, a pediatric scalp EEG collection sampled at 256 Hz[23]. Signals were band-pass filtered (0.5-40 Hz) with a 60 Hz notch filter and divided into 5-second windows, with 50% overlap during training and 0% during evaluation. Interictal segments were drawn with a minimum one-hour margin from any seizure onset, and interictal-to-preictal windows were sampled at a 3:1 ratio during training.

Evaluation followed a leave-one-seizure-out (LOSO) protocol, with a further held-out validation split within each training fold for computing fusion weights, kept separate from the test fold. The state-transition matrix T was estimated from the training portion of each fold by counting empirical state-to-state transitions and normalizing each row to sum to one, with uniform fallback for unobserved transitions.

5.3. Prediction Layer & Fusion

The prediction layer combines three independently trained classifiers, a hybrid CNN-TCN model, an LSTM network, and a Random Forest, each producing a probability estimate of the preictal state from 5-second EEG windows (256 Hz, 0.5-40 Hz bandpass). Rather than averaging these outputs uniformly or weighting them by a sensitivity-to-false-alarm ratio, which can degenerate when any single classifier performs poorly, each classifier is assigned a weight proportional to its validation AUC, computed on a held-out validation split separate from the test set used for FAR evaluation, to avoid leakage:

$$w_i = \frac{\mathrm{AUC}_i^{\mathrm{val}}}{\mathrm{AUC}_{\mathrm{cnn}}^{\mathrm{val}} + \mathrm{AUC}_{\mathrm{lstm}}^{\mathrm{val}} + \mathrm{AUC}_{\mathrm{rf}}^{\mathrm{val}}} \quad (1)$$

The three outputs are then combined by convex weighted fusion:

$$P_{\text{fusion}}(t) = w_{\text{cnn}} \cdot P_{\text{cnn}}(t) + w_{\text{lstm}} \cdot P_{\text{lstm}}(t) + w_{\text{rf}} \cdot P_{\text{rf}}(t) \quad (2)$$

To reduce sample-level noise before the fusion output is passed to the decision layer, a causal backward-looking moving average is applied over a window of H = 5 steps:

$$\bar{P}(t) = \frac{1}{5} \sum_{k=0}^{4} P_{\text{fusion}}(t-k) \quad (3)$$

This entire prediction layer, including all three classifiers and the fusion weights, is trained once and held fixed throughout the study; it is never retrained or recalibrated for individual patients, cost configurations, or personalization profiles. All adaptation described in subsequent sections occurs exclusively downstream, in the MDP decision layer.

5.4. MDP Decision Layer

The decision layer formulates the alerting problem as a Markov Decision Process, operating on the smoothed prediction probability P̄(t) rather than on raw signal features:

$$\mathcal{M} = \langle S, A, T, R, \gamma \rangle \quad (4)$$

Each state s ∈ S is defined by four discrete components: the quantized probability level (Level, 4 bins), the short-term trend of P̄(t) (Trend: rising, stable, or falling), the duration for which the current level has persisted (Duration: short or long), and the operational mode (Mode: Idle, Soft-Refractory, or Hard-Refractory), giving a state space of size

$$|S| = |\text{Level}| \times |\text{Trend}| \times |\text{Duration}| \times |\text{Mode}| = 4 \times 3 \times 2 \times 3 = 72 \quad (5)$$

As shown in Figure 2. At each state, the agent selects one of three actions, A = {No-Alert, Soft,

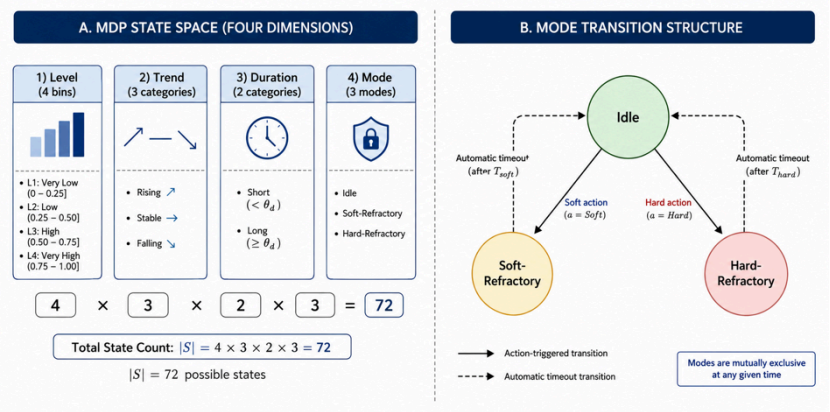

Hard}. No distinct "wait" action is retained: an earlier four-action formulation including such an action was found to select it in none of the resulting optimal policies, confirmed by an exact match between the reduced and full state-space results across all tested operating points. Mode transitions follow a deterministic refractory structure: issuing a Soft or Hard alert moves the system into the corresponding refractory mode, from which it exits automatically after a fixed clinical suppression window, while Idle mode allows any of the three actions to be selected freely at each step.

Figure 2. MDP state space and mode transition structure.

5.5 Cost Function Design

The cost function assigns a scalar penalty to each state-action pair, balancing the risk of a missed seizure against the cost of a false alarm. Let l denote the discretized probability level (0 to 3), τ = 1 minus the trend indicator (+1 for rising, 0 for stable, −1 for falling), and d the duration indicator plus one (1 for short, 2 for long). Two intermediate quantities drive the cost: a miss-risk term that grows with both level and persistence,

$$A = d \cdot l \quad (6)$$

and a false-alarm-risk term that is largest when the level is low and the duration is short,

$$B = d \cdot (3 - l) \quad (7)$$

The total cost for state s and action a combine these two terms, gated by whether an alert is currently active, together with a fixed operational cost per mode and a repeat-alert penalty applied only while the system remains in a refractory mode:

$$J(s,a) = g(m)\cdot A\cdot(\tau+2)\cdot w_{md}[a]\cdot\alpha_{md} + B\cdot(2-\tau)\cdot w_{fa}[a]\cdot\alpha_{fa} + c_{\text{mode}}(m) \quad (8)$$

where g(m) equals 1 in Idle mode and 0 during any refractory mode, so that the miss-risk term is only counted while the system is actively able to alert. When the system is in Soft-Refractory or Hard-Refractory mode and action a is Soft or Hard, an additional repeat-alert penalty is added:

$$J(s,a) \mathrel{+}= \rho \cdot \frac{w_{fa}[a]}{w_{fa}[\text{Hard}]} \quad (9)$$

which discourages issuing further alerts of either strength while a previous alert is still within its clinical suppression window. The weight vectors w_md and w_fa, the mode costs c_mode, and the repeat-penalty coefficient ρ are held fixed across all patients; only the cost multipliers α_md and α_fa are varied, providing the mechanism for clinical personalization described in Section 5.8.

5.6. Optimal Policy via Value Iteration

The optimal policy is obtained by solving the Bellman optimality equation for cost minimization. Starting from V(s) = 0 for all states, the value function is updated iteratively:

$$V_{k+1}(s) = \min_{a \in A} \left[J(s,a) + \gamma \sum_{s' \in S} T(s' \mid s,a)\, V_k(s') \right] \quad (10)$$

where γ = 0.95 and T(s'|s,a) is estimated from training data, combined deterministically with the mode-transition structure in Section 5.4. Iteration continues until convergence:

$$\max_{s \in S} |V_{k+1}(s) - V_k(s)| < \varepsilon \quad (11)$$

with ε = 10^{-6}. Once V has converged, the optimal policy π* is recovered by selecting, at each state, the action that minimizes the one-step cost plus discounted future cost:

$$\pi^*(s) = \arg\min_{a \in A} \left[J(s,a) + \gamma \sum_{s' \in S} T(s' \mid s,a)\, V^*(s') \right] \quad (12)$$

A diagnostic check confirmed π* differs from a naive greedy policy (ignoring the future term) for a non-trivial subset of states, verifying the sequential formulation is active rather than collapsing to an instantaneous rule.

5.7 Patient Screening Protocol

Each patient's fixed prediction model was screened using a pre-registered criterion: only patients with validation AUC ≥ 0.70 were included, fixed in advance to prevent tuning the screening toward the proposed method. One patient, chb18, fell marginally below this threshold (AUC = 0.699) and was excluded; a secondary check confirmed a simple fixed threshold already outperformed any decision-layer configuration for this patient, independently supporting the exclusion. All patients meeting the criterion were retained regardless of subsequent decision-layer performance, including one for which the framework only matched the best fixed-rule baseline.

5.8. Uniform Parameter Refinement

Two cost parameters were refined through a systematic sweep applied identically across all patients: w_fa[Hard] and the level_bins upper edge, both originally conservative (5.0 and 0.85). A pre-specified rule adopted a candidate only if it improved at least one patient without harming any other. Six parameters were tested (Table 1); only w_fa[Hard] (1.5-8.0) and level_bins (0.75-0.85) satisfied the criterion, yielding 2.0 and 0.75 respectively, verified via clean training-stream re-runs across all four patients. All other parameters (w_md[Hard], w_fa[Soft], mode_cost[Hard], repeat_penalty) were tested individually and left unchanged.

The full adopted configuration: level_bins = [0.30, 0.60, 0.75], w_md = [4.0, 2.0, 1.0], w_fa = [0.0, 2.0, 2.0], mode_cost = [0.0, 0.3, 0.6], ρ = 2.0, α_md = 1.0, γ = 0.95, ε = 10^{-6}.

Parameter Tested	Candidate Value(s) Tested	Outcome	Justification
$w_{fa}[Hard]$	1.5, 2.0, 3.0, 5.0, 8.0	✓ Adopted (5.0 → 2.0)	Improved chb10 and chb20 sensitivity/FAR without harming chb01 or chb03
$level_{bins}$ upper edge	0.75, 0.78, 0.80, 0.82, 0.85	✓ Adopted (0.85 → 0.75)	Improved chb10 and chb20 without harming chb01 or chb03; verified via clean training-stream re-run
$w_{md}[Hard]$	0.5, 1.0, 1.5, 2.0, 3.0	⊗ Rejected — no improvement	No consistent gains across patients
$w_{fa}[Soft]$	1.0, 1.5, 2.0, 2.5, 3.0	⊗ Rejected — no improvement	No consistent gains across patients
$mode_{cost}[Hard]$	0.2, 0.4, 0.6, 0.9, 1.2	⊗ Rejected — no improvement	No consistent gains across patients
$repeat_{penalty}$ (ρ)	0.5, 1.0, 2.0, 3.0, 4.0	⊗ Rejected — no improvement	No consistent gains across patients

Table 1: Uniform parameter refinement tested configurations and acceptance outcomes

5.9. Evaluation Protocol

Performance was evaluated using a seizure prediction horizon (SPH) of 5 minutes and a seizure occurrence period (SOP) of 30 minutes. Three metrics were used: the false alarm rate,

$$\pi^*(s) = \arg\min_{a \in A} \left[J(s,a) + \gamma \sum_{s' \in S} T(s' \mid s,a)\, V^*(s') \right] \qquad (12)$$

sensitivity at the seizure level,

$$\text{Sensitivity} = \frac{N_{\text{TP}}}{N_{\text{TP}} + N_{\text{MS}}} \quad (14)$$

and the mean warning time,

$$\overline{WT} = \frac{1}{N_{\text{TP}}} \sum_{i=1}^{N_{\text{TP}}} (t_{\text{onset},i} - t_{\text{alert},i}) \quad (15)$$

An alert was counted as a true positive if it occurred within the SOP-plus-SPH window preceding a seizure onset, and as a false alarm otherwise; a missed seizure was recorded when no alert was .issued within this window prior to onset

Statistical significance was assessed against a chance predictor generating alerts at the same :mean FAR as the evaluated method, following the analytical framework of Winterhalder et al.[24]

The framework was compared against three fixed-rule baselines on the identical fusion output: a fixed threshold swept across six values (0.30-0.85), a k-of-n majority-vote rule following the CNN-based CHB-MIT configuration[14], and a two-level threshold rule capable of soft and hard .alerts, the only baseline with graded alerting comparable to the proposed framework

Results

6.1 Patient Screening Outcomes

Five patients were considered for evaluation. Applying the pre-registered AUC ≥ 0.70 screening criterion (Section 5.7) retained four patients (chb01, chb03, chb10, chb20) and excluded chb18, whose validation AUC of 0.699 fell marginally below threshold, as shown in (Table 2).

Patient ID	Validation AUC	Screening Outcome
chb01	0.954	Included
chb03	0.748	Included
chb10	0.793	Included
chb18	0.699	**Excluded**
chb20	0.976	Included

6.2 MDP vs. Fixed-Rule Baselines

Across the four screened patients, the proposed framework was compared against the best-performing fixed-rule baseline for each patient at matched sensitivity, selected from the full sweep of threshold, k-of-n, and two-level rules (Section 5.9), as shown in (Table 3). The framework reduced FAR relative to the best fixed baseline in three of four patients while matching sensitivity, and matched the best baseline in the fourth. For chb01, FAR fell from 1.223 (best baseline, two-level) to 0.506 at 100% sensitivity, a 58.6% reduction; for chb10, from 0.696 (threshold 0.70) to 0.232 at 86% sensitivity, a 66.7% reduction; and for chb20, from 0.250 (threshold 0.85) to 0.208 at 100% sensitivity, a 16.8% reduction. For chb03, both the framework and the best baseline reached FAR = 0.000 at 60% sensitivity, an exact tie. No patient showed degraded performance relative to the best available fixed rule.

Patient	Best Fixed Baseline	Baseline FAR/h	Baseline Sensitivity	MDP FAR/h	MDP Sensitivity	Outcome

chb01	2-level (0.60/0.85)	1.223	100%	**0.506**	**100%**	Win (−58.6% FAR)
chb03	Threshold 0.70	0.000	60%	**0.000**	**60%**	Tie
chb10	Threshold 0.70	0.696	86%	**0.232**	**86%**	Win (−66.7% FAR)
chb20	Threshold 0.85	0.250	100%	**0.208**	**100%**	Win (−16.8% FAR)

6.3 Operating-Point Spectrum

Because a single cost parameter, α_fa, governs the sensitivity-FAR trade-off without retraining or re-estimating the transition matrix, each patient admits a continuous spectrum of operating points rather than one fixed configuration, as shown in (Fig.3). At a conservative setting (α_fa = 16), chb10 reaches 100% sensitivity, capturing a seizure missed by fixed thresholds, at the cost of a substantially higher FAR (3.676/h); at the adopted operating point (α_fa = 96), the same patient achieves 86% sensitivity at FAR = 0.232/h, as reported in Section 6.2. This illustrates that the adopted configurations represent a deliberate choice along a continuum, not the only achievable trade-off, and that personalization occurs entirely within the decision layer without touching the prediction model.

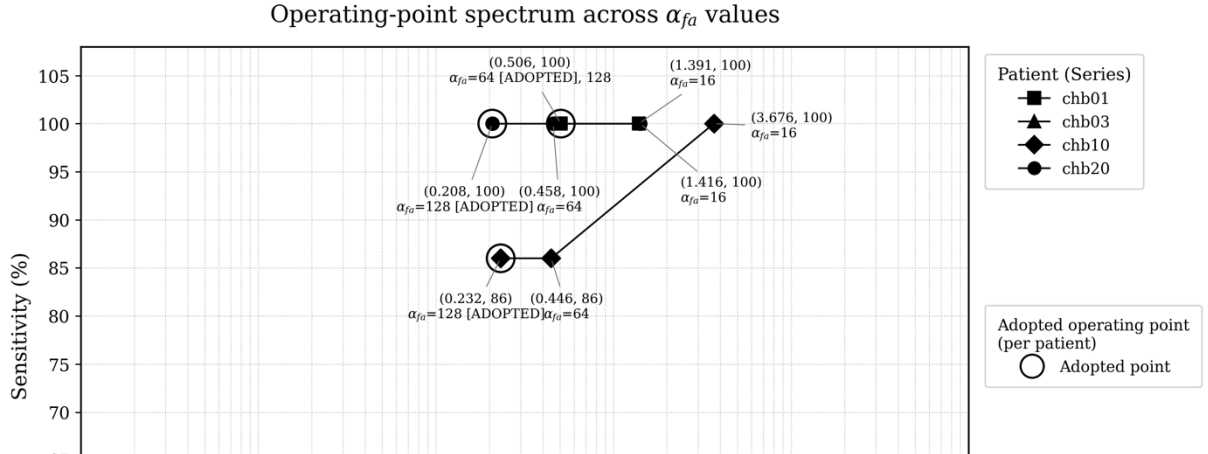

Figure 3. Operating-point spectrum across alpha_fa values

6.4 Illustrative Clinical Profiles

Three configurations illustrate this spectrum: High Safety (α_fa = 16), which maximizes sensitivity at the cost of tolerating more alarms; Balanced (α_fa = 64), a representative mid-range setting; and Low Alarm (α_fa = 128), which minimizes alarm burden, applied uniformly across all four patients as shown in (Table 4). It is important to emphasize that these three profiles are explanatory examples only, chosen to demonstrate the range of achievable behavior; they are not a fixed menu of options. Any value of α_fa yields a complete, immediately usable policy, so the operating point can be tuned continuously and specifically to an individual patient's seizure severity, alerting history, and the treating clinician's risk preference, without any change to the underlying prediction model.

Patient	High Safety (α = 16)*Sensitivity / FAR (h^{-1})*	Balanced (α = 64)*Sensitivity / FAR (h^{-1})*	Low Alarm (α = 128)*Sensitivity / FAR (h^{-1})*
chb01	100% / 1.391	100% / 0.506	100% / 0.506
chb03	60% / 0.125	60% / 0.000	60% / 0.000
chb10	100% / 3.676	86% / 0.446	86% / 0.232
chb20	100% / 1.416	100% / 0.458	100% / 0.208

Table 4: Illustrative clinical profiles across all patients

Note: The three columns are a uniform illustrative sweep (α_fa = 16, 64, and 128) applied identically to all four patients to display the achievable operating range; they are not the

adopted per-patient operating points. Each patient's adopted operating point and its corresponding false alarm rate are those reported in Section 6.2 (Table 3).

6.5 Graded Alerting Behavior

Unlike the fixed baselines, which issue a single alert type and therefore always report zero Hard-FAR, the proposed policy allocates a distinct fraction of states to the Hard action across all four patients, as shown in (Table 5). For chb01, the optimal policy assigns Hard alerts to 8% of states and Soft alerts to 1%, with the remainder in No-Alert; the other three patients converge on a similar structural split (92% No-Alert, 1% Soft, 7% Hard). This convergence reflects the uniform cost function shared across patients: most states are resolved by the immediate cost term alone, while the patient-specific transition matrix affects only a minority of states, so the abstract policy structure is similar even though the empirical FAR differs substantially by patient. The states escalated to Hard are consistently those combining high probability level, active mode, and a persistent or rising trend, indicating that escalation is conditioned on the joint state rather than probability alone, a distinction unavailable to any single-threshold rule.

Patient	αfa	No-Alert (%)	Soft (%)	Hard (%)
chb01	64	90	1	8
chb03	64	92	1	7
chb10	96	92	1	7
chb20	96 *	92	1	7

Table 5: Graded alerting behavior, policy action distribution per patient

6.6 Where Sequential Decision-Making Helps and Where It Does Not

The magnitude of the framework's advantage varies systematically with classifier separability. For chb01 and chb10, where the classifier's interictal probability distribution overlaps substantially with the alerting region (a wide "gray zone" above the lower decision threshold), the sequential policy yields large FAR reductions (58.6% and 66.7%). For chb20, where classifier separation is comparatively sharper, the advantage is smaller but still positive (16.8%). For chb03, where separation is nearly complete, the framework exactly matches rather than exceeds the best fixed baseline, both reaching FAR = 0.000 at 60% sensitivity. This pattern is consistent with a ceiling effect: sequential decision-making adds value primarily where instantaneous rules cannot resolve genuine ambiguity in the classifier output, and that value diminishes, without becoming negative, as classifier separability approaches its limit.

A secondary analysis examined why the frequency of Hard escalation differs across patients despite the shared cost structure (Section 6.5). State purity, the fraction of upper-level states visited exclusively by preictal or exclusively by false-alarm windows, did not explain the pattern: chb01, the patient with the most frequent Hard usage, had the lowest purity rate among the four patients. Instead, the determining factor was the absolute volume of preictal training data at the highest probability level: chb01 contributed 405 preictal windows at Level 3, compared to 20 for chb10, giving the transition matrix T substantially more evidence from which to estimate confident escalation behavior at that level. This indicates that the policy's willingness to commit to a Hard alert is governed by the statistical reliability of the patient-specific transition estimate, not by the separability of the underlying signal at that state.

Two structural conditions were verified independently for chb03: the tie was confirmed across five separate cost-parameter sweeps (Section 5.8), all converging on the same FAR floor, ruling out a tuning artifact; and it holds specifically at the classifier's near-complete separation for this patient, not from a defect in the decision layer itself. A related limitation applies when the refractory suppression window is long relative to the inter-seizure interval: under this condition, the decision effectively reduces to a single threshold crossing, the same choice a fixed rule would make, so the two approaches converge.

7. Conclusions

This work introduced an Adaptive Clinical Decision Layer, a Markov Decision Process operating on top of a fixed, unmodified seizure prediction model, as a sequential and cost-aware alternative to instantaneous alerting rules. Evaluated against threshold, k-of-n, and two-level baselines under matched sensitivity on four screened CHB-MIT patients, the proposed framework reduced FAR by 58.6%, 66.7%, and 16.8% for three patients and matched the best baseline exactly for the fourth, without retraining the underlying classifier. A single cost parameter enabled continuous, patient-specific operating-point selection, and the policy allocated Hard alerts selectively to states combining high probability, active mode, and persistent or rising trend, a distinction unavailable to single-threshold rules. The magnitude of improvement was found to depend on classifier separability, consistent with a ceiling effect in which sequential decision-making adds the most value where the underlying signal remains genuinely ambiguous.

Several limitations bound these conclusions. The evaluation is restricted to four patients from a single public dataset, limiting statistical power and generalizability to broader clinical populations. The near-complete separation observed for one patient (chb03) illustrates that the framework's advantage is not universal and diminishes as classifier quality improves. When the refractory suppression window is long relative to the inter-seizure interval, the sequential

formulation converges toward a single threshold decision, reducing its practical distinction from fixed rules. The transition matrix governing state dynamics is estimated once from training data and held fixed thereafter; it does not adapt as new observations accumulate for a given patient. Finally, statistical significance testing was constrained by the limited number of seizures per patient available in the dataset.

Future work should pursue three directions arising directly from these limitations. First, the transition matrix could be updated prospectively as new patient observations accumulate, allowing the decision layer to track evolving seizure dynamics without recomputing the policy from scratch. Second, transition matrices estimated from data-rich patients could serve as informative priors for patients with limited seizure counts, extending the framework's applicability to the low-data regime that currently constrains screening. Third, pairing the decision layer with a higher-accuracy prediction model would provide a direct empirical test of the ceiling-effect hypothesis identified in Section 6.6, clarifying whether the framework's advantage narrows as predicted when the underlying signal becomes less ambiguous. Prospective, closed-loop clinical evaluation would further establish whether the offline gains reported here translate into real-world alerting benefit.

Acknowledgment

The authors gratefully acknowledge the Department of Artificial Intelligence and Robotics Research and the Joint SDAIA Research Center for Artificial Intelligence at the university for their support of this work. The authors also thank Dr. Ali Nasser, whose supervision and guidance were invaluable throughout this research. Finally, the authors thank the providers of the CHB-MIT Scalp EEG Database, made publicly available through PhysioNet, whose recordings

made this study possible. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

[1] R. S. Fisher et al., "ILAE official report: A practical clinical definition of epilepsy," Epilepsia, vol. 55, no. 4, pp. 475–482, 2014. https://pubmed.ncbi.nlm.nih.gov/24730690

[2] S. Ramgopal et al., "Seizure Detection, Seizure Prediction, and Closed-Loop Warning Systems in Epilepsy," Epilepsy & Behavior, 2014. https://www.sciencedirect.com/science/article/pii/S1525505014002334

[3] Y. Alotaibi et al., "Hybrid CNN-TCN Modeling of EEG Temporal Signals for Epileptic Seizure Prediction," IEEE AMLDS, 2025. https://doi.org/10.1109/AMLDS63918.2025.11159432

[4] Y. Wu et al., "A Review of Machine Learning and Deep Learning Trends in EEG-Based Epileptic Seizure Prediction," IEEE Access, 2025. https://e-space.mmu.ac.uk/642057/1/A_Review_of_Machine_Learning_and_Deep_Learning_Trends_in_EEG-Based_Epileptic_Seizure_Prediction.pdf

[5] C. Moutonnet et al., "Clinical Translation of Machine Learning Algorithms for Seizure Detection in Scalp EEG: A Systematic Review," arXiv:2404.15332, 2024. https://arxiv.org/pdf/2404.15332

[6] S. Bhattacharya, "Automatic Seizure Prediction Using CNN and LSTM," arXiv:2211.02679, 2022. https://arxiv.org/pdf/2211.02679

[7] A. Esmaeilpour et al., "Deep learning-based seizure prediction using EEG signals: A comparative analysis on CHB-MIT," Engineering Reports, 2024. DOI: 10.1002/eng2.12918 https://onlinelibrary.wiley.com/doi/full/10.1002/eng2.12918

[8] "EEG-Based Two-Stage Channel-Aware Set Transformer Network for Seizure Prediction," arXiv:2507.15364, 2025. https://arxiv.org/html/2507.15364v1

[9] "Spatio-Temporal Attention Network for Epileptic Seizure Prediction," arXiv:2511.02846, 2025. https://arxiv.org/pdf/2511.02846

[10] X. Lu, A. Wen, L. Sun, H. Wang, Y. Guo, Y. Ren, "An Epileptic Seizure Prediction Method Based on CBAM-3D CNN-LSTM Model," IEEE Journal of Translational Engineering in Health and Medicine, vol. 11, pp. 417–423, 2023. DOI: 10.1109/JTEHM.2023.3290036 https://doi.org/10.1109/JTEHM.2023.3290036

[11] "Machine and Deep Learning-Based Seizure Prediction: Scoping Review on Temporal and Spectral Features," Applied Sciences, 2025. https://www.mdpi.com/2076-3417/15/11/6279

[12] "Real-Time EEG-Based Epileptic Seizure Prediction Using Artificial Intelligence: A Systematic Review," medRxiv, 2025. https://www.medrxiv.org/content/10.1101/2025.10.09.25337692v1.full

[13] Y. Bai et al., "Epileptic Seizure Detection Using Machine Learning: A Systematic Review and Meta-Analysis," Brain Sciences, 2025. https://pmc.ncbi.nlm.nih.gov/articles/PMC12190198

[14] N. D. Truong et al., "Convolutional neural networks for seizure prediction using intracranial and scalp electroencephalogram," Neural Networks, vol. 105, pp. 104–111, 2018. DOI: 10.1016/j.neunet.2018.04.018 https://doi.org/10.1016/j.neunet.2018.04.018

[15] A. Balamas et al., "Epileptic Seizure Prediction on CHB-MIT EEG Using Soft Fusion Post-Processing," Advances in AI and Machine Learning, vol. 5, no. 4, 2025. https://www.oajaiml.com/uploads/archivepdf/711854254.pdf

[16] L. Warmerdam et al., "Minimizing artifact-induced false-alarms for seizure detection in wearable EEG devices with gradient-boosted tree classifiers," Scientific Reports, vol. 14, 2024. https://www.nature.com/articles/s41598-024-52551-0

[17] M. L. Puterman, Markov Decision Processes: Discrete Stochastic Dynamic Programming, Wiley, 2014. https://www.wiley.com/en-us/Markov+Decision+Processes%3A+Discrete+Stochastic+Dynamic+Programming-p-9780471727828

[18] M. DeRoos & M. Lavieri, "Finding the Maximum Safe Treatment Interval Using MDP," Production and Operations Management, 2025. https://journals.sagepub.com/doi/10.1177/10591478241281077

[19] "MDP Model for Medical Service Resources in Emergency Departments," Scientific Reports, 2025. https://pmc.ncbi.nlm.nih.gov/articles/PMC11758071

[20] A. Nasir et al., "Emotional Interaction Between Human and Robot: MDP for Socially Assistive Robots," IEEE SSD, 2025. DOI: 10.1109/SSD64182.2025.10989870 https://doi.org/10.1109/SSD64182.2025.10989870

[21] M. M. Quamar, A. Nasir & S. Elferik, "A Novel MDP Decomposition Framework for Scalable UAV Mission Planning," Vehicular Communications, 2026. DOI: 10.1016/j.vehcom.2026.101036 https://doi.org/10.1016/j.vehcom.2026.101036

[22] A. Nasir et al., "Optimized Task Assignment and Predictive Maintenance for Industrial Machines using MDP," Operational Research, 2025. DOI: 10.1007/s12351-025-00976-4 https://doi.org/10.1007/s12351-025-00976-4

[23] J. Guttag, "CHB-MIT Scalp EEG Database (version 1.0.0)," PhysioNet, 2010. https://physionet.org/content/chbmit/1.0.0

[24] M. Winterhalder et al., "The seizure prediction characteristic: a general framework to assess and compare seizure prediction methods," Epilepsy & Behavior, vol. 4, pp. 318–325, 2003. DOI: 10.1016/S1525-5050(03)00105-7 https://doi.org/10.1016/S1525-5050(03)00105-7