Minimax Theory and its Applications
Volume 11 (2026), No. 4s, 1 to 13

Minimax-Based robust optimization of cognitive engagement prediction in adaptive online learning systems

Lutfillo Abdunazarov[1], Mohammed Shamsul Hoque[2], Nirav Jadav[3], Doniyorbek Khudoyberganov[4], Nigora Kenjaeva[5], Latofat Jabbarova[6]

[1]Professor, Doctor of Philosophy (PhD), Pedagogical Sciences, Kokand University, Kokand, Uzbekistan. abdunazarovlutfullo0@gmail.comkokanduni.uz, https://orcid.org/0009-0006-0747-8260

[2]Professor, Department of English Southeast University, Dhaka, Bangladesh. hoque@seu.edu.bd, https://orcid.org/0000-0001-5695-8742

[3]Database Architect, Intuitive.AI, Palm Beach Gardens, United States. njadav@asu.edu, https://orcid.org/0009-0000-9316-5506

[4]Lecturer, Department of Psychology, Mamun University, Xorazm, Uzbekistan. xudoyberganov.doniyor@gmail.com, https://orcid.org/0009-0007-1387-7097

[5]Tashkent Pharmaceutical Institute, Tashkent, Uzbekistan. nikaken1612@gmail.com, https://orcid.org/0000-0003-4384-5137

[6]Doctor of Philosophy (PhD), Art Studies Kamoliddin Behzod National Institute of Fine Arts and Design, Tashkent, Uzbekistan. latofat.jbt.85@gmail.com, https://orcid.org/0009-0004-9608-3345

Abstract:

Adaptive online learning systems use predictions of learners' cognitive engagement to determine when to intervene, but engagement predictors are trained by average loss minimization for the majority group of learners and used on changing populations with interaction logs corrupted by device-specific noise. formalize the engagement prediction task as a hierarchical minimax problem where the learner model plays against two adversary players who alter the proportions of device-discipline groups around the observed mixture within Kullback-Leibler neighborhood and add per-learner feature-specific noise based on the test-retest variability estimate. For the case of device-aware linear scorers, both maxima in the minimax game have closed forms: the first one corresponds to a loss-weighted l1 margin penalty and the second one is the log-sum-exp of cohort losses. Thus, the saddle-point problem can be solved as a convex program that is trained as empirically risk minimization and has equilibrium by Sion's theorem. On a 16-feature learner-interaction benchmark using six cohorts, with a mobile dominant shifted distribution and attacks in the worst-case telemetry setting, and validation-based model selection for each method across ten random seeds, the proposed Hierarchical Minimax Robust predictor (HMR) increases shifted robust accuracy by 0.655 - 0.643 = 0.012 over empirical risk minimisation in all ten seeds (p=0.002) and in-distribution robust worst-group accuracy by 0.648 - 0.627 = 0.021 (p=0.004), and beats importance weighting, Group DRO and CVaR in robust accuracy by 1.05 to 1.61 points, at a cost of only 0.38 points of in-distribution accuracy and without any impact on shifted accuracy or calibration. HMR achieves performance similar to adversarial training but uses a smaller perturbation budget. Ablations demonstrate that the telemetry adversary provides the robustness benefit, the cohort adversary with regularisation maintains accuracy under distribution shift, and noise-aware bounds converge to the uniform ball with homogeneous logging noise.

Keywords: Cognitive engagement, minimax optimisation, distributionally robust optimisation, adversarial robustness, adaptive learning systems, learning analytics, distribution shift, telemetry noise.

1. Introduction

An adaptive e-learning system needs to evaluate constantly whether the learner is engaged and whether it is time to provide a hint, alter the modality of presentation, make a check-in request, or notify the teacher. All such

decisions are made using an engagement predictor that relates interaction data (session rhythm, use of video controls, assessment activity, navigation activity, and discussion activity) to the level of cognitive engagement (Baker & Hawn, 2022). Cognitive engagement is a type of mental effort and self-regulation used in the process of learning, which is a major determinant of learning outcomes and very hard to assess directly (Arunfred & Ganesan, 2025; Abidova et al., 2026).

The predictors of engagement are generally acquired by empirical risk minimisation (ERM) on aggregated data for the dominating learner population of a certain platform. There are three reasons why it might not work in the real world (Duchi & Namkoong, 2021; Henrie et al., 2015).

1. Population diversity. The same behavior may have different meanings under different circumstances. Long desktop sessions can be an indicator of studying, while long mobile sessions could mean idle browser usage. Humanities-type discussion-based classes have students showing their interest by interacting in forums, while professional-type classes have students showing interest by timely submitting assignments. Mixed models learn behavior of the prevalent population only Chi & Wylie, (2014) and Kuzilek et al., (2017).

2. Distribution shift. Devices, courses, and learners' populations may differ from one academic term to another and between institutions. The rise in mobile-first enrollment may cause distribution discrepancies between training and testing phases, thus affecting prediction accuracy (Duchi & Namkoong, 2021; Fredricks et al., 2004).

3. Metric distortion. Data is collected client-side where dropped events, background tabs, and player instrumentation can distort metrics. Even small metric distortions can flip predictions at decision borders (Henrie et al., 2015; Tuy, 2004).

In turn, there is a corresponding robust solution for each property, namely, Group DRO for population heterogeneity, f-divergence DRO for distribution shift, and adversarial training for input perturbations (Henrie et al., 2015; Jeganathan et al., 2024). In the case of engagement prediction, these properties may be interrelated, since minority groups can experience noisy telemetry. Optimization of worst-group performance without taking into account telemetry noise can lead to overfitting of noise of the minority groups, whereas optimization of worst-case perturbation without consideration of population heterogeneity can result in robustness allocation only to the majority population (Nemirovski, 2004; Piech et al., 2015).

In this paper, consider the engagement prediction problem as a nested minimax problem in which one adversary modifies the distribution of the learner group weights within a KL ball around the true group distribution, while another adversary changes telemetry of each learner within the budgets defined by estimated logging noise. In the case of device-aware linear scorers, the solutions of inner maximizations can be found analytically, and the saddle-point problem reduces to convex minimization (Nemirovski, 2004; Piech et al., 2015; Tuy, 2004; Whitehill et al., 2014).

The contributions of this paper are as follows.

1. Hierarchical minimax reformulation (HMR) that combines the robustness against group-level distribution shift with robustness against individual sample perturbations through the temperature τ, which smoothly interpolates between adversarial training ($\tau \to \infty$) and Group DRO ($\tau \to 0$).

2. An exact reduction. For the linear scorer model that operates on device-specific information, the perturbation maximization becomes an l_1 margin penalty with respect to effective weights with noise weights, while the group maximization becomes log-sum-exp over group risks. Such an objective is convex, the existence of a saddle point is guaranteed by Sion's theorem [20], and such an objective provides guarantees on loss for any perturbed mixture of groups.

3. A telemetry-noise-aware uncertainty set where the radii of the noise balls of each feature are learned from the difference between two sessions of the same learner.

4. A controlled empirical study on a 16-feature learner-interaction benchmark with six device–discipline cohorts, a mobile-dominant shifted test population and worst-case telemetry attacks, comparing HMR with five baselines and three ablations over ten seeds with validation-based model selection for every method. HMR significantly improves robust accuracy over ERM, importance weighting, Group DRO and CVaR at no significant cost in shifted accuracy; the ablations show that the telemetry adversary carries the gain, while the group adversary and noise-aware radii add little once the scorer is device-aware.

Section 2 reviews related work. Section 3 states the problem. Section 4 presents HMR and its analysis. Sections 5 and 6 describe the experimental design and results. Section 7 discusses implications and limitations, and Section 8 concludes.

2. Related work

2.1. Measuring and predicting engagement

In most cases, engagement is considered an amalgam of behavioral, emotional, and cognitive components (Arunfred & Ganesan, 2025). According to the ICAP framework, cognitive engagement is based on observable learning behaviors that range from passive to active, constructive, and interactive levels and allow using interaction patterns as proxies for measuring engagement (Abidova et al., 2026). Studies of engagement assessment in technology-mediated learning show an increase in the use of interaction data instead of self-reported measures and difficulties related to validation of the proxies [12]. Affect-aware educational technologies use interaction and physiology to model transition among confusion, frustration, and boredom and stress the relevance of cognitive-affective states while learning (Bakaeva et al., 2026). Engagement recognition techniques based on vision categorize learner engagement using facial expressions (Bertsimas & Sim, 2004). Gamified collaborative learning frameworks consider student engagement and interaction among peers within online educational systems (Chi & Wylie, 2014). The Open University Learning Analytics Dataset offers logs of virtual learning environments for big data analysis in education (D'Mello & Graesser, 2012). Prediction techniques in education include optimized regression models and deep knowledge tracing which make predictions about academic proficiency or knowledge acquisition but not cognitive engagement based on learner information (Fredricks et al., 2004; Rockafellar & Uryasev, 2000).

2.2. Fairness and subgroup validity in educational prediction

Predictive models for students' behavior can show different performance levels when applied to different demographics and contexts (Duchi & Namkoong, 2021). Evaluation of subgroups analyzes predictive results for each subgroup of learners separately, and thus allows one to see some differences that might not be revealed otherwise due to the aggregate level of performance assessment. Adaptive learning algorithms allow for differentiated education of learners with diverse needs.

2.3. Robust and distributionally robust learning

The approach of robust optimization makes decisions robust to realizations from an uncertainty set, balancing robustness and optimization (Tuy, 2004). In classifiers, robustness to bounded perturbation in inputs can be achieved using regularization techniques, whereas adversarial training can deal with worst-case perturbations using the concept of minimax optimization. Distributionally robust optimization deals with uncertainties in data distributions and tries to find models that maintain consistent prediction performances across different populations (Henrie et al., 2015). Optimization based on conditional value at risk (CVaR) focuses on losses, giving a mathematical foundation to worst-case prediction risks (Jeganathan et al., 2024). Group distributionally robust optimization (Group DRO) focuses on worst group risk, while distributionally robust learning considers uncertainties in population and subgroup performances (Henrie et al., 2015). Minimax theory is the mathematical foundation for analyzing problems with opposing goals and uncertainties (Nemirovski, 2004). Convex-concave saddle point optimization helps with equilibrium analyses and first order optimization methods with proven convergence properties (Piech et al., 2015). These concepts drive robust engagement prediction under population changes and telemetry perturbations.

2.4. Gap addressed

As far as know, no previously published engagement predictor is optimized for group shift and telemetry perturbation simultaneously, nor has the budget on the latter been correlated with the logging noise. HMR closes this gap by providing an exactly solvable formulation for the scorer type used by adaptive systems.

3. Problem formulation

3.1. Data and groups

Each learner-session record is a triple (x, y, g). The vector $x \in \mathbb{R}^d$ holds $d = 16$ standardised telemetry features (Table 1). The label $y \in \{-1, +1\}$ indicates low or high cognitive engagement. The group $g \in \{1, \ldots, G\}$ is the device–discipline cohort; here $G = 6$, formed by two device classes (desktop and mobile) and three course families (STEM, humanities and professional). Group membership is known at prediction time, because the platform knows the client device and the course.

Table 1: Telemetry features used by the engagement predictor.

Category	Features
Session rhythm	sessions per week; log mean session duration; idle-time ratio; coefficient of variation of inter-click intervals
Video interaction	video completion ratio; pause rate; rewind rate
Assessment	quiz attempts; first-attempt quiz accuracy; self-assessment use
Discussion	forum posts; mean reply depth
Self-regulation	help requests; note-taking events; assignment submission latency; navigation entropy

3.2 Scorer class

Adaptive systems need predictors that are fast, auditable and stable under retraining. using a device-aware linear scorer

$$f_\theta(x, g) = \big(w + m_g\, v\big)^\top x + b_g, \qquad \theta = (w, v, b_1, \dots, b_G) \qquad (1)$$

Where $m_g = 1$ for mobile cohorts and $m_g = 0$ otherwise. The effective weight vector of group g is $u_g = w + m_g v$. The interaction vector v allows the model to reverse the sign of a feature on mobile, for example session duration, while sharing statistical strength across cohorts, and the group intercepts b_g absorb differences in engagement base rates.

3.3 Two sources of uncertainty

Group-mix uncertainty. Let $p = (p_1, \dots, p_G)$ be the training group proportions and $R_g(\theta)$ the expected loss on group g. The deployment mix q is unknown but plausibly near p. model it as any q in the probability simplex Δ_G, penalised by its divergence $\mathrm{KL}(q \parallel p)$ from p.

Telemetry uncertainty. Each observed feature vector may differ from the true behaviour by a perturbation δ, bounded feature-wise:

$$\mathcal{U}_\varepsilon(s) = \big\{\delta \in \mathbb{R}^d :\ |\delta_j| \le \varepsilon\, s_j,\ j = 1, \dots, d\big\} \qquad (2)$$

Where ε is a global budget and s_j is the relative logging noise of feature j. Platforms observe several sessions per learner, so logging noise can be measured directly as test–retest variation. With two sessions $x_i^{(1)}$ and $x_i^{(2)}$ of learner i, whose underlying engagement is the same,

$$\hat{\sigma}_j = \sum_{g=1}^{G} p_g\, \frac{\mathrm{sd}_{i \in g}\big(x_{ij}^{(1)} - x_{ij}^{(2)}\big)}{\sqrt{2}}, \qquad s_j = \frac{\hat{\sigma}_j}{\frac{1}{d}\sum_k \hat{\sigma}_k} \qquad (2a)$$

on standardised features, so that the radii average to one. Features logged noisily relative to their spread receive larger radii.

3.4. Objective

With the logistic loss $\ell(z) = \ln(1 + e^{-z})$ on the margin z, the robust risk of group g is

$$R_g^\varepsilon(\theta) = \mathbb{E}_{(x,y)\sim P_g}\Big[\max_{\delta \in \mathcal{U}_\varepsilon(s)} \ell\big(y\, f_\theta(x + \delta, g)\big)\Big] \qquad (3)$$

and the learning problem is the hierarchical minimax

$$\min_\theta \max_{q \in \Delta_G} \sum_{g=1}^{G} q_g\ R_g^\varepsilon(\theta) - \tau\, \mathrm{KL}(q \parallel p) + \frac{\lambda}{2}\, \Omega(\theta) \qquad (4)$$

The outer player is the learner. The first adversary, nature over groups, chooses the deployment mix; the second, nature over telemetry, chooses per-record perturbations. The temperature $\tau > 0$ prices departures from the observed mix, and $\Omega(\theta) = \| w \|_2^2 + 3 \| v \|_2^2$ is a ridge penalty that shrinks the mobile interaction more strongly because fewer mobile records are available.

4. Proposed Method: Hierarchical Minimax Robust Engagement Prediction (HMR)

4.1. Closing the inner maximisations

Proposition 1 (telemetry adversary). For the scorer (1) and any non-increasing loss ℓ,

$$\max_{\delta \in \mathcal{U}_\varepsilon(s)} \ell\left(y\left(u_g^\top(x+\delta)+b_g\right)\right) = \ell\left(y\left(u_g^\top x + b_g\right) - \varepsilon \sum_{j=1}^{d} s_j \left|u_{g,j}\right|\right) \qquad (5)$$

Attained at $\delta_j^* = -\varepsilon\, s_j\, y\, \mathrm{sign}\left(u_{g,j}\right)$.

Proof. The margin term $y\, u_g^\top \delta$ is linear in δ and is minimised over the box by setting each coordinate at its bound with sign opposite to $y\, u_{g,j}$, which gives $-\varepsilon \sum_j s_j \left|u_{g,j}\right|$. Because ℓ is non-increasing, minimising the margin maximises the loss.

The telemetry adversary therefore acts as a margin requirement scaled by a noise-weighted ℓ_1 norm of the group's effective weights, extending the robustness–regularization. to group-specific weights and heterogeneous radii. A feature that is logged unreliably must earn its weight through a larger margin.

Proposition 2 (group adversary). For fixed θ with robust group risks $R_g = R_g^\varepsilon(\theta)$,

$$\max_{q \in \Delta_G} \sum_g q_g\, R_g - \tau\, \mathrm{KL}(q \parallel p) = \tau \ln \sum_g p_g\, e^{R_g/\tau}, \qquad q_g^* = \frac{p_g\, e^{R_g/\tau}}{\sum_h p_h\, e^{R_h/\tau}} \qquad (6)$$

Proof. This is the Gibbs variational principle (Donsker–Varadhan duality). The objective is strictly concave in q, and the stationarity condition of its Lagrangian over the simplex yields the stated softmax form.

The temperature controls the adversary. As $\tau \to \infty$ the objective tends to the population-weighted robust risk $\sum_g p_g\, R_g$, which is adversarial training. As $\tau \to 0$ it tends to $\max_g R_g$, which is Group DRO on robust losses. An intermediate τ keeps the adversary's weight on small groups bounded by their likelihood under p, preventing the collapse onto a single noisy.

4.2 The reduced problem and its properties

Substituting both propositions gives the HMR objective

$$\Phi_\tau(\theta) = \tau \ln \sum_{g=1}^{G} p_g \exp\left(\frac{\hat{R}_g^\varepsilon(\theta)}{\tau}\right) + \frac{\lambda}{2}\, \Omega(\theta), \qquad \hat{R}_g^\varepsilon(\theta) = \frac{1}{n_g} \sum_{i:\, g_i = g} \ell\left(y_i f_\theta(x_i, g) - \varepsilon\, s^\top \left|u_g\right|\right) \qquad (7)$$

Where n_g is the number of training records in group g and $\left|u_g\right|$ is taken element-wise.

Proposition 3 (convexity and equilibrium). Φ_τ is convex in θ. Problem (4) has a saddle point (θ^*, q^*), its min–max and max–min values coincide, and θ^* minimises Φ_τ.

Proof. The robust margin is the sum of a function linear in θ and $-\varepsilon s^\top \left|u_g\right|$, which is concave in θ, so the margin is concave. Because ℓ is convex and non-increasing, ℓ composed with a concave function is convex, and each $\hat{R}_g^\varepsilon$ is convex. The log-sum-exp function is convex and increasing in each argument, so the composition is convex, and adding the ridge term preserves convexity. The game payoff $\sum_g q_g\, \hat{R}_g^\varepsilon - \tau\, \mathrm{KL}(q \parallel p)$ is convex in θ and concave in q over the compact simplex, so Sion's minimax theorem gives equality of values and the existence of an equilibrium.

Proposition 4 (shift certificate). For any deployment mix $q \in \Delta_G$ and any telemetry perturbation in $\mathcal{U}_\varepsilon(s)$,

$$\sum_g q_g\, \mathbb{E}_{P_g}\left[\ell\left(y\, f_\theta(x+\delta, g)\right)\right] \;\le\; \Phi_\tau(\theta) - \frac{\lambda}{2}\Omega(\theta) + \tau\, \mathrm{KL}(q \parallel p) \qquad (8)$$

Proof. By Proposition 1 the left side is at most $\sum_g q_g\, R_g^\varepsilon(\theta)$, and the Donsker–Varadhan inequality gives $\sum_g q_g\, R_g \le \tau \ln \sum_g p_g\, e^{R_g/\tau} + \tau\, \mathrm{KL}(q \parallel p)$.

The certificate makes τ interpretable for practitioners. For the mobile-dominant test mix used in Section 5, $\mathrm{KL}(q \parallel p) = 0.74$ nats; with $\tau = 0.02$ the certified increase in loss above the training objective is at most 0.015.

4.3. Algorithm

Φ_τ is minimised by full-batch subgradient descent with Adam step sizes. Each iteration performs four steps.

1. Compute the robust margins and per-record losses using the closed form (5).
2. Average the losses within groups to obtain $\hat{R}_g^{\varepsilon}$.
3. Compute the adversary weights q_g^* from (6).
4. Form the subgradient $\sum_g q_g^* \nabla \hat{R}_g^{\varepsilon} + \frac{\lambda}{2}\nabla\Omega$, using $\mathrm{sign}(u_{g,j})$ for the ℓ_1 term, and take an Adam step.

Each iteration costs $O(nd)$ operations, the same order as one epoch of ERM, and no inner attack loop is needed, unlike projected-gradient adversarial training. We use $\lambda = 10^{-3}$, learning rate 0.05 and $1{,}000$ iterations; $\varepsilon \in \{0.05, 0.1, 0.2\}$ and $\tau \in \{0.01, 0.02, 0.05, 0.1\}$ are selected on a validation split (Section 5.2). After training, a single temperature T is fitted on a held-out validation split by minimising negative log-likelihood to calibrate predicted probabilities; T rescales scores and leaves every classification and ranking decision unchanged.

4.4. Ablated variants

There are three versions where each individual component is isolated. For HMR-noAdvFeat, we use ε = 0 and use only the group adversary. For HMR-noKL, we replace the KL-regularized adversary with the non-regularized exponentiated gradient adversary from Group DRO (step size of 0.05) and keep the telemetry adversary. Finally, for HMR-uniform, we replace the noise aware radius by s_j=1 for all j, i.e., the standard l_∞ ball. In both HMR-noAdvFeat and HMR-uniform, we use the same hyperparameters as HMR.

5. Experimental design

5.1. Learner-interaction benchmark

The experiments involve an engagement model where the interaction between the learner and the test item is modeled using a structural generative process where the underlying cognitive engagement state generates observables via cohort-based measurement models. These experiments capture three characteristics that have been found to be true in engagement literature, i.e., behavior metrics load on engagement in a cohort-based manner, telemetry reliability varies depending on client device, and minority cohorts are few in the training data. Due to the generative process, we can control the group shift and telemetry corruption accurately.

For learner i in cohort g, cognitive engagement is $c_i \sim \mathcal{N}(0,1)$ and the binary label is

$$y_i = \begin{cases} +1, & c_i + 0.35\,\xi_i > t_g \\ -1, & \text{otherwise} \end{cases}, \qquad \xi_i \sim \mathcal{N}(0,1) \qquad (9)$$

with cohort thresholds $t = (0,\ 0.15,\ -0.10,\ 0.25,\ 0.35,\ 0.20)$ for the six cohorts in the order of Table 2, and 5% of labels flipped uniformly at random to represent annotation noise. Telemetry is generated as

$$x_i = c_i\,\boldsymbol{\lambda}_g + \mathbf{o}_g + \boldsymbol{\sigma}_g \odot \boldsymbol{\zeta}_i, \qquad \boldsymbol{\zeta}_i \sim \mathcal{N}(\mathbf{0}, I_d) \qquad (10)$$

where $\boldsymbol{\lambda}_g$ is the cohort loading vector, $\mathbf{o}_g$ a cohort offset and $\boldsymbol{\sigma}_g$ the cohort noise scale. Base loadings are positive for productive behaviours (video completion, quiz accuracy, reply depth, note-taking) and negative for submission latency, navigation entropy, idle ratio and inter-click variability. Cohort structure enters as follows.

- Cohorts on mobile: the loading of log session duration changes sign (factor -0.4), accounting for idle tabs; pause, rewind, navigation-entropy, idle-ratio, and inter-click loadings are reduced by factors ranging between 0.40 to 0.60; noise in telemetry data is 1.1 to 2.0 times larger than on the desktop and largest for idle ratio, inter-click variability and session duration; and mean session duration is lowered.
- Cohorts in STEM and humanities disciplines: STEM disciplines increase quiz and help-seeking loadings (factor 1.4), while humanities disciplines increase discussion and note-taking loadings (1.6) and reduce quiz loadings (0.4); professional courses increase latency and self-assessment loadings (1.7) and reduce discussion loadings (0.3).

- Common across all cohorts: o_g offsets are generated once per cohort from N (0, 〖0.35〗^2) for each feature and kept constant across seeds.

- Cohort with mobile-professional: audio-first study decreases all video loadings (factor 0.2) and increases submission latency loadings (1.5).

Table 2: Cohort composition of the training and shifted test populations.

Cohort	Training share p_g	Shifted test share q_g
Desktop–STEM	0.36	0.10
Desktop–Humanities	0.24	0.08
Desktop–Professional	0.20	0.07
Mobile–STEM	0.10	0.30
Mobile–Humanities	0.07	0.25
Mobile–Professional	0.03	0.20

In every run, there are 20,000 training samples where every learner has two sessions (with the first session being used for training and both sessions for test-retest radii (2a)), 5,000 validation samples and 20,000 in-distribution test samples from the training split. In the shifted test dataset of 20,000 samples, the mobile-dominated split from Table 2 is used along with two deployment effects: a cohort-level covariate shift, where every mobile cohort gets an offset vector that has standard deviation of 0.6 per feature and a 30% increase in telemetry noise in mobile clients. Standardization is done using training statistics. Each experiment is run for ten different random seeds (from 0 to 9).

5.2. Methods compared

They all share the same device-aware linear scorer (1), the same ridge penalty, the same optimiser and the same number of iterations; hence the only thing that will make their performance differ is the training objective. All those methods that have hyperparameters are selected based on the validation split from Table 3 grid using the same criterion, which is the average of validation worst-group accuracy and validation robust worst-group accuracy at ε_att=0.2. The test sets are not used for selection.

Table 3: Methods compared.

Method	Training objective
ERM	Mean logistic loss over all records
IW	Group-balanced importance weighting, each cohort weighted equally
AT	Adversarial training with uniform radii, inner maximum solved exactly by (5) with $s_j = 1$; $\varepsilon \in \{0.05,0.1,0.2\}$
CVaR	Mean loss over the worst fraction α of records [16, 19]; $\alpha \in \{0.1,0.2,0.5\}$
Group DRO	Worst-cohort loss with exponentiated-gradient adversary [15]; step $\eta \in \{0.01,0.05,0.2\}$
HMR	Objective (7) with test–retest radii s; $\varepsilon \in \{0.05,0.1,0.2\}$, $\tau \in \{0.01,0.02,0.05,0.1\}$
HMR-noAdvFeat	HMR with $\varepsilon = 0$, τ as selected for HMR
HMR-noKL	HMR with the unregularised Group DRO adversary; $\varepsilon \in \{0.05,0.1,0.2\}$
HMR-uniform	HMR with $s_j = 1$, (ε, τ) as selected for HMR

5.3 Evaluation metrics

- Accuracy and WGA on the in-distribution and shifted datasets; WGA is the minimum of accuracies over all six demographic groups.

- AUC, the area under the ROC curve.

- Robust accuracy and robust WGA under the worst-case telemetry perturbation in U_ε_att (s) with ε_att=0.2 except when noted otherwise, determined analytically via Proposition 1. This threat model uses the test-retest radii that were actually observed in the data for every model, including models trained using uniform radii.

- Expected Calibration Error (ECE) with ten equally sized confidence bins, and post-validation-set temperature scaling for every model.

- Intervention Recall Rate. In our adaptive algorithm, we flag the top 30% of learners with the lowest engagement scores using an alert budget. We give the percentage of disengaged learners that were flagged in total, and in the worst-served demographic group.

The results are expressed as mean ± standard deviation for ten seeds. The difference between HMR and each competitor is evaluated using a two-tailed Wilcoxon signed-rank test for paired data over seeds. With ten paired samples, the smallest possible p-value is 0.002.

6. Results and analysis

6.1. Clean accuracy under distribution shift

Table 4 shows good results. In distribution, all approaches achieve 81.4–82.0% accuracy, whereas in the mobile-dominated shifted population, the accuracy drops to 76.3–77.1% (worst group accuracy 72.0–72.9%). The mobile-professional sub-population that accounts for 3% of the training set and 20% of the shifted population, represents the worst served sub-population for most approaches; under ERM, the accuracy in the shifted population is 73.1% versus 82.5% in the desktop-STEM population. The HMR approach incurs a cost of 0.38 percentage point in terms of in-distribution accuracy compared to ERM (0.816 against 0.820; lower in all ten seeds, p = 0.002), while it does not differ significantly in terms of shifted accuracy and worst group accuracy in the shifted population compared to ERM, importance weighting, CVaR and Group DRO (Table 6). Post-hoc temperature scaling of all approaches yields well-calibrated models, with shifted ECE between 0.035 and 0.044.

Table 4: Clean performance on the in-distribution (ID) and shifted test sets (mean ± SD over ten seeds). WGA: worst-group accuracy; ECE: expected calibration error (lower is better).

Method	ID accuracy	ID WGA	Shifted accuracy	Shifted WGA	Shifted AUC	Shifted ECE
ERM	0.820 ± 0.003	0.779 ± 0.017	0.770 ± 0.015	0.723 ± 0.042	0.849 ± 0.012	0.035 ± 0.015
IW	0.820 ± 0.003	0.781 ± 0.015	0.771 ± 0.014	0.728 ± 0.038	0.849 ± 0.012	0.035 ± 0.014
AT	0.814 ± 0.005	0.775 ± 0.017	0.765 ± 0.015	0.720 ± 0.030	0.843 ± 0.012	0.039 ± 0.017
CVaR	0.817 ± 0.011	0.775 ± 0.019	0.768 ± 0.013	0.723 ± 0.036	0.844 ± 0.015	0.044 ± 0.034
Group DRO	0.818 ± 0.005	0.780 ± 0.015	0.769 ± 0.012	0.725 ± 0.035	0.846 ± 0.011	0.041 ± 0.016
HMR	0.816 ± 0.003	0.781 ± 0.013	0.768 ± 0.016	0.726 ± 0.029	0.846 ± 0.012	0.038 ± 0.017
HMR-noAdvFeat	0.819 ± 0.003	0.781 ± 0.017	0.771 ± 0.014	0.725 ± 0.039	0.849 ± 0.011	0.036 ± 0.013
HMR-noKL	0.814 ± 0.003	0.781 ± 0.009	0.763 ± 0.018	0.725 ± 0.028	0.842 ± 0.013	0.041 ± 0.017
HMR-uniform	0.817 ± 0.003	0.782 ± 0.014	0.769 ± 0.015	0.729 ± 0.030	0.847 ± 0.011	0.038 ± 0.015

6.2. Robustness to telemetry perturbation

Table 5 and Table 6 show the main effect of the minimax formulation. Under worst-case telemetry perturbation at $\varepsilon_{att} = 0.2$, HMR raises shifted robust accuracy from 0.643 under ERM to 0.655, a gain of 1.16 points observed in all ten seeds ($p = 0.002$), and raises shifted robust worst-group accuracy by 1.23 points ($p = 0.010$). In distribution, the gain in robust worst-group accuracy is 2.07 points (0.648 against 0.627; $p = 0.004$). HMR also improves shifted robust accuracy over importance weighting (+1.05), Group DRO (+1.16) and CVaR (+1.61), each in ten of ten seeds. Methods that re-weight data without a telemetry adversary therefore do not protect against corrupted logging.

Table 5: Robust performance under the worst-case telemetry perturbation with $\varepsilon_{att} = 0.2$, and recall of disengaged learners under a 30% alert budget (mean ± SD over ten seeds).

Method	ID robust accuracy	ID robust WGA	Shifted robust accuracy	Shifted robust WGA	Shifted recall	Shifted worst-cohort recall
ERM	0.671 ± 0.003	0.627 ± 0.015	0.643 ± 0.015	0.600 ± 0.034	0.477 ± 0.005	0.294 ± 0.095
IW	0.670 ± 0.004	0.636 ± 0.012	0.644 ± 0.014	0.608 ± 0.032	0.477 ± 0.005	0.299 ± 0.090
AT	0.691 ± 0.005	0.640 ± 0.016	0.657 ± 0.013	0.612 ± 0.030	0.474 ± 0.006	0.284 ± 0.100
CVaR	0.665 ± 0.021	0.624 ± 0.022	0.639 ± 0.019	0.598 ± 0.032	0.475 ± 0.009	0.259 ± 0.118
Group DRO	0.668 ± 0.005	0.638 ± 0.014	0.643 ± 0.014	0.606 ± 0.031	0.476 ± 0.006	0.290 ± 0.081
HMR	0.687 ± 0.006	0.648 ± 0.017	0.655 ± 0.015	0.612 ± 0.032	0.475 ± 0.006	0.298 ± 0.094
HMR-noAdvFeat	0.669 ± 0.004	0.636 ± 0.014	0.643 ± 0.015	0.604 ± 0.034	0.477 ± 0.005	0.295 ± 0.094
HMR-noKL	0.685 ± 0.007	0.649 ± 0.016	0.652 ± 0.015	0.616 ± 0.027	0.473 ± 0.007	0.305 ± 0.089
HMR-uniform	0.682 ± 0.006	0.649 ± 0.017	0.653 ± 0.014	0.616 ± 0.029	0.476 ± 0.006	0.298 ± 0.088

Table 6: Paired differences HMR minus competitor over ten seeds (percentage points), with, in brackets, the number of seeds in which HMR is higher and the two-sided Wilcoxon signed-rank p-value.

Competitor	ID acc.	Shifted acc.	Shifted WGA	ID robust WGA	Shifted robust acc.	Shifted robust WGA
ERM	-0.38 [0/10; 0.002]	-0.28 [2/10; 0.154]	+0.27 [3/10; 0.625]	+2.07 [9/10; 0.004]	+1.16 [10/10; 0.002]	+1.23 [8/10; 0.010]
IW	-0.36 [0/10; 0.002]	-0.33 [2/10; 0.078]	-0.13 [2/10; 0.322]	+1.16 [8/10; 0.027]	+1.05 [10/10; 0.002]	+0.46 [8/10; 0.322]
AT	+0.18 [7/10; 0.193]	+0.30 [7/10; 0.432]	+0.57 [5/10; 0.695]	+0.82 [7/10; 0.193]	-0.20 [4/10; 0.375]	-0.02 [5/10; 0.922]
CVaR	-0.10 [1/10; 0.084]	-0.07 [3/10; 0.432]	+0.27 [4/10; 1.000]	+2.33 [10/10; 0.002]	+1.61 [10/10; 0.002]	+1.43 [9/10; 0.006]
Group DRO	-0.17 [2/10; 0.105]	-0.15 [3/10; 0.336]	+0.17 [5/10; 0.770]	+1.02 [8/10; 0.027]	+1.16 [10/10; 0.002]	+0.65 [8/10; 0.322]
HMR-noAdvFeat	-0.33 [0/10; 0.002]	-0.32 [3/10; 0.168]	+0.11 [2/10; 0.432]	+1.17 [9/10; 0.004]	+1.13 [10/10; 0.002]	+0.82 [8/10; 0.105]
HMR-noKL	+0.18 [6/10; 0.160]	+0.41 [9/10; 0.020]	+0.09 [7/10; 0.557]	-0.16 [4/10; 0.492]	+0.24 [6/10; 0.375]	-0.40 [4/10; 0.375]
HMR-uniform	-0.09 [0/10; 0.002]	-0.12 [4/10; 0.232]	-0.27 [3/10; 0.322]	-0.13 [3/10; 0.461]	+0.17 [8/10; 0.193]	-0.38 [3/10; 0.275]

HMR and adversarial training are statistically indistinguishable on every metric in Table 6. The two methods reach the same robustness by different routes. Validation selected $\varepsilon = 0.05$ for HMR in five seeds and $\varepsilon = 0.1$ in the other five, while it selected $\varepsilon = 0.1$ or 0.2 for AT in nine of ten seeds. The group adversary concentrates the robustness budget on the mobile cohorts, so HMR needs a smaller global budget to reach the same robust accuracy, and its in-distribution accuracy is 0.18 points higher on average, a difference that is not significant ($p = 0.19$).

Figure 1 traces robust accuracy as the attack budget grows. The gap between the robust methods and the non-robust ones widens with the budget: at $\varepsilon_{att} = 0.5$, shifted robust accuracy is 0.464 for HMR and 0.473 for AT, against 0.428 for ERM and 0.429 for Group DRO, and shifted robust worst-group accuracy is 0.414 for HMR and 0.410 for AT, against 0.376 for ERM.

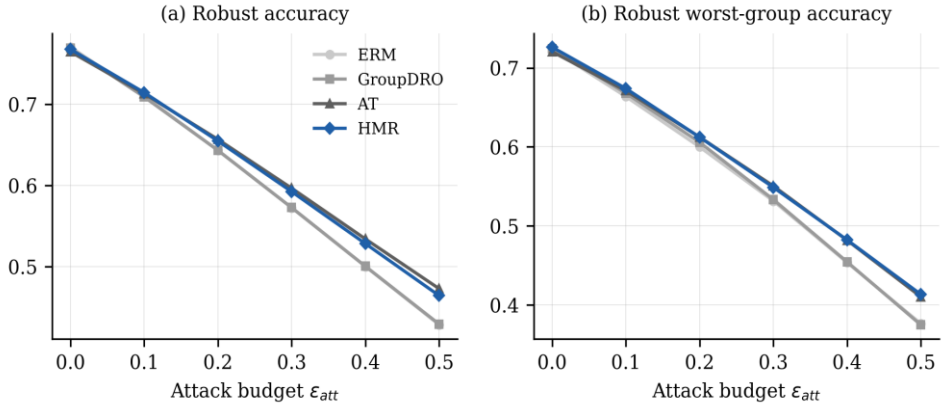


Figure 1: Robust accuracy (a) and robust worst-group accuracy (b) on the shifted test set as a function of the attack budget, averaged over ten seeds. ERM and Group DRO overlap.

Figure 2 illustrates the breakdown of shifted robust accuracy per cohort. The most significant gain for HMR relative to ERM can be seen on the mobile-professional cohort, which is the least populated cohort, where the robust accuracy goes up from 0.604 to 0.624; AT achieves 0.620 in that cohort and Group DRO 0.620. In the desktop cohorts, AT slightly outperforms HMR (for example, 0.707 versus 0.701 on desktop-STEM).

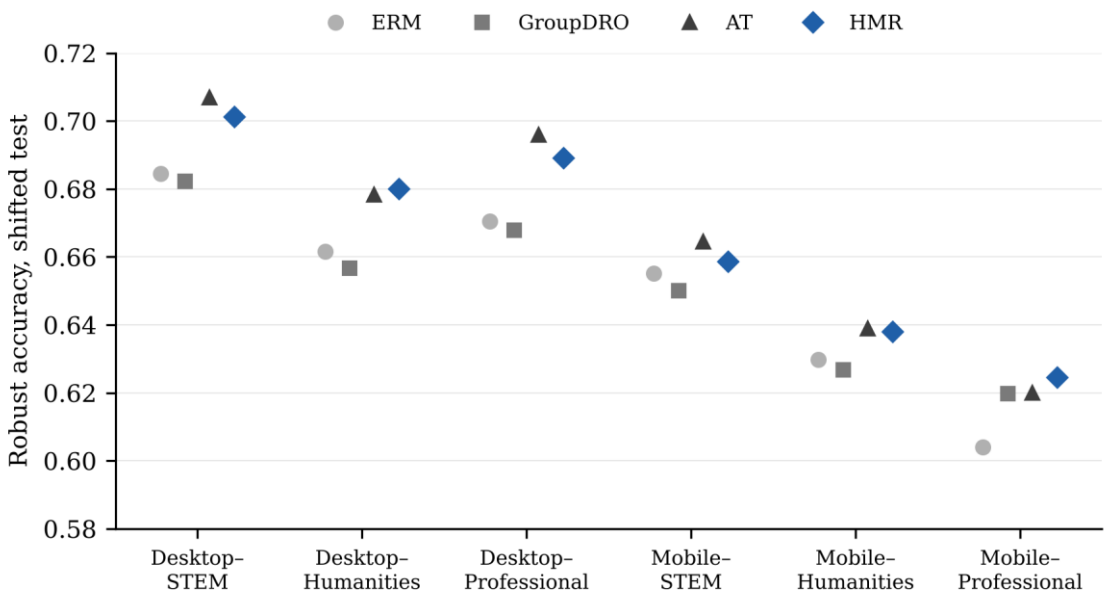


Figure 2: Shifted robust accuracy by cohort at $\varepsilon_{att} = 0.2$, averaged over ten seeds.

6.3. Downstream intervention

With the alert budget identifying the 30% of students with the lowest performance, every strategy identifies around 47.5% of disengaged learners for the shifted distribution (see Table 5). As seen in the cohort with the lowest recall, the variance of recall across seeds is large (0.08 to 0.12) and all methods perform no better than HMR. The increased robustness from Section 6.2 thus guards against errors near the decision boundary while not affecting who the policy targets in clean data.

6.4. Ablation: which component carries the gain

The ablations in Table 6 isolate the three components.

- Telemetry adversary. Eliminating it (HMR-noAdvFeat) results in a decrease in shifted robust accuracy of 1.13 points in all ten seeds (p=0.002) and in-distribution robust worst-group accuracy of 1.17 points (p=0.004), whereas in-distribution accuracy increases by 0.33 points. This feature generates robustness.

- Regularised Group DRO adversary. Substituting the regularised Group DRO adversary by the unregularised variant (HMR-noKL) keeps robustness measures stable, but decreases shifted accuracy by 0.41 points (p=0.020, HMR better in nine out of ten seeds). Restricting the adversary's divergence from the observed distribution helps preserve in-distribution accuracy in the shift setting, as predicted by Proposition 2.

- Noise-aware radius definition. HMR-uniform is statistically identical to HMR in terms of robustness metrics, and has 0.09 points higher in-distribution accuracy (p=0.002). Indeed, according to the observed test–retest feature-wise radii, standardizing the features results in relative noise values of all but one feature between 0.83 and 1.14, with only video completion ratio being clearly reliable at 0.72. As soon as the logging noise distribution is almost homogeneous relative to feature dispersion, the noise-aware ball is close to uniform.

Telemetry adversary. Omitting it (HMR-noAdvFeat) lowers shifted robust accuracy by 1.13 (all ten seeds, p=0.002) and in-distribution robust worst-group accuracy by 1.17 (p=0.004), but raises in-distribution accuracy by 0.33. This is the element generating the robustness boost.

Kullback-Leibler regularization of the group adversary. Omitting the KL-regularized adversary and using just the unregularized Group DRO adversary (HMR-noKL) does not affect robust metrics but decreases shifted accuracy by 0.41 (p=0.020; HMR is better in nine of ten seeds). Constraining the deviation of the adversary from the observed distribution mix preserves clean accuracy against shifts, as expected in Proposition 2.

Noise-aware radii. HMR-uniform is not statistically different from HMR in robust metrics, and HMR-uniform's in-distribution accuracy is 0.09 higher (p=0.002). This follows from the measured test–retest radii: normalized,

all features have relative noise between 0.83 and 1.14 except for video completion ratio (0.72), which is the most reliable feature. With nearly equal relative logging noise compared to the feature range, the noise-aware set is close to the uniform ball.

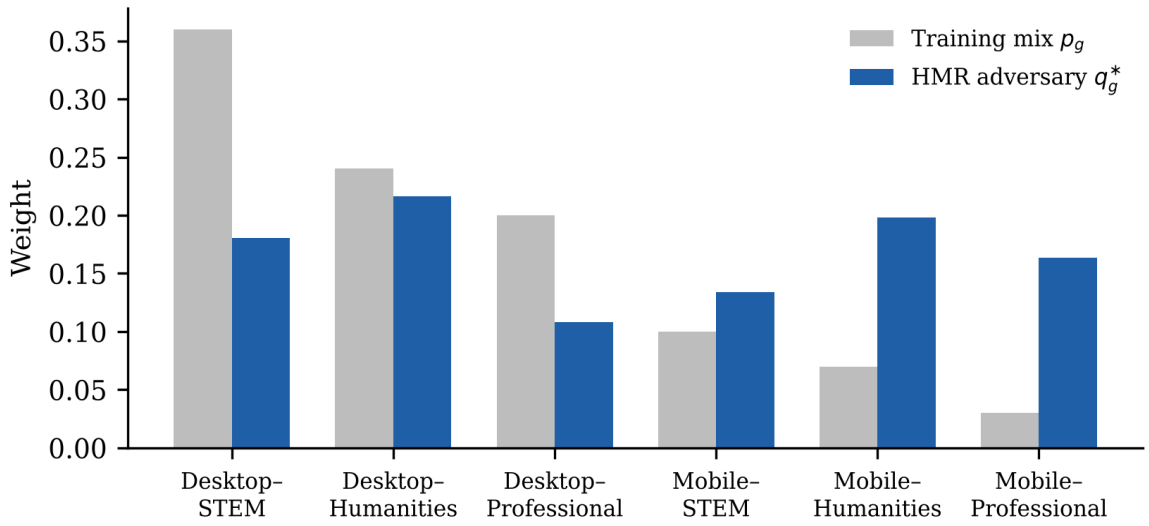


Figure 3: Training group mix and the HMR adversary's equilibrium weights, averaged over ten seeds.

6.5. Sensitivity to the budget and the temperature

The fourth figure visualizes the metrics across the hyper-parameter space. The budget is responsible for controlling the accuracy-robustness tradeoff since an increase in ε from 0.05 to 0.2 causes the in-distribution accuracy to drop from around 0.818 to 0.802 and increased robust worst-group accuracy to rise from 0.613 to 0.621 at τ=0.02. The temperature shows a relatively small but steady effect since robust worst-group accuracy is maximized at $\tau \in \{0.01, 0.02\}$ at all budgets and drops as τ increases toward 0.1, making the objective close to adversarial training on the training mix. The validation metric is maximized at $\varepsilon \in \{0.05, 0.1\}$ and $\tau \leq 0.05$, the zone selected by the procedure.

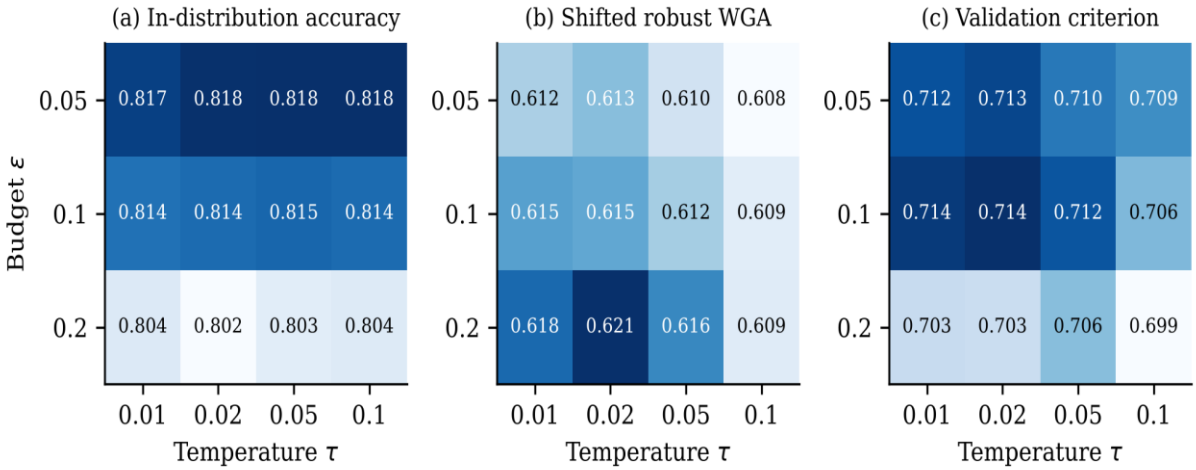


Figure 4: HMR test metrics over the hyperparameter grid, averaged over ten seeds: (a) in-distribution accuracy, (b) shifted robust worst-group accuracy and (c) the validation selection criterion.

7. Discussion

7.1. Implications for adaptive learning systems

Robustness to telemetry corruption is cheap when solved exactly. For the device-aware linear scorer, the telemetry adversary reduces to a noise-weighted ℓ_1 margin term, so HMR trains at the cost of ERM and buys 1.2 points of robust accuracy for 0.4 points of clean in-distribution accuracy. Because intervention decisions are taken near the decision boundary, this is the regime in which corrupted logging changes what an adaptive system does.

Device-aware modelling already absorbs most cohort heterogeneity. The group DRO, importance weights and the group adversary of HMR affected the shifted clean accuracy of less than half a point. The interaction term v enables the scorer to invert the effect of the session duration for the mobile clients, thus nullifying the major error cause of the cohort without re-weighting. In practice, one ought to provide the contextual features for the behaviour predictor and then train on the rest.

Regularise the group adversary. In absence of KL penalty, the group adversary was penalized 0.4 point on shifted accuracy with no added robustness advantage. In addition, the temperature τ offers a deployment-related certificate (Proposition 4), where, for the mobile-predominant society, the upper bound loss increase guarantee is 0.015 for $\tau = 0.02$.

Measure logging noise before shaping the uncertainty set. Justification for noise-aware radii arises from differences in reliability between the features. In our experiment, the test–retest radii for 15 out of 16 features were near one, and thus the noise-aware radii had similar behavior as the uniform ball. If different platforms have a larger difference in terms of instrumentation between the features, such that certain events are instrumented server-side.

7.2. Limitations

The analysis relies on a structural generative benchmark instead of a log from an operational system. This enables the control and precise quantification of group shift and telemetry corruption; however, the effect sizes are loadings-, noise level- and shift-dependent and it is necessary to verify it against institutional log data with reliable engagement labels, e.g. self-reporting or observation protocols. The scoring classifier is linear; for nonlinear scorers, the inner optimization problem does not have a closed form solution anymore and needs to be approximated, e.g. using projected gradient steps or convex relaxations; thus, weakening the strength of Propositions 1 and 3. Labels are binary, but the notion of engagement is multidimensional. Finally, the threat model constrains the perturbations in a pointwise manner; therefore, any correlations among features resulting from the same event would lead to a non-rectangular uncertainty set.

8. Conclusion

In this work, cognitive engagement prediction in adaptive online learning systems is posed as a hierarchical minimax problem where there is a cohort adversary which re-weighs the cohorts of learners in a KL neighborhood of the observed mix and there is also a telemetry adversary that perturbs interaction features in a measured noise neighborhood. For device-aware linear scorers, both sub-problems have closed forms, the resulting objective function is convex, and a saddle point exists.

In a comparison across a benchmark of six cohorts with a mobile-dominant shifted population, HMR yielded superior robustness to worst-case telemetry perturbations than ERM, importance weighting, Group DRO and CVaR in all ten runs, by 1.05 to 1.61, and in-distribution robust worst-group accuracy by 2.1 than ERM, but at the expense of 0.4 of in-distribution accuracy and without affecting the shifted accuracy. The technique achieved the same level as adversarial training but with a smaller perturbation budget. The ablation study revealed that the telemetry adversary was responsible for the performance increase, that using KL divergence regularization of the group adversary maintains accuracy under the shift, and that noisy radii are important only if measured logging noise is heterogeneous. Minimax training is thus a low overhead guarantee for the engagement prediction models upon which adaptive intervention relies.

References

1. Arunfred, N., & Ganesan, A. (2025). Talent development in global organizations through structural equation modeling of leadership and employee engagement. International Academic Journal of Organizational Behavior and Human Resource Management, 12(4), 1–8.

2. Abidova, N., Khurramov, R., Djuraev, D., Otamurodova, S. Q., Kalandarova, G., Karabayev, X., & Rahmatova, F. (2026). A multi-agent reinforcement learning and diffusion model hybrid for real-time personalized curriculum optimization in intelligent education platforms. International Academic Journal of Science and Engineering, 13(1), 470–480. https://doi.org/10.71086/IAJSE/V13I1/IAJSE1344

3. Bakaeva, F., Kasimov, S., Abdullajanova, D., Abbasova, D., Sabirova, N., Djurayeva, D., & Omonova, M. (2026). Designing an advanced gamified collaborative learning model to foster student engagement and peer interaction in online education. Indian Journal of Information Sources and Services, 16(2), 704–712. https://doi.org/10.51983/ijiss-2026.16.2.70

4. Baker, R. S., & Hawn, A. (2022). Algorithmic bias in education. International Journal of Artificial Intelligence in Education, 32(4), 1052–1092.

5. Bertsimas, D., & Sim, M. (2004). The price of robustness. Operations Research, 52(1), 35–53.

6. Chi, M. T., & Wylie, R. (2014). The ICAP framework: Linking cognitive engagement to active learning outcomes. Educational Psychologist, 49(4), 219–243.

7. D'Mello, S., & Graesser, A. (2012). Dynamics of affective states during complex learning. Learning and Instruction, 22(2), 145–157.

8. Duchi, J. C., & Namkoong, H. (2021). Learning models with uniform performance via distributionally robust optimization. The Annals of Statistics, 49(3), 1378–1406.

9. Fredricks, J. A., Blumenfeld, P. C., & Paris, A. H. (2004). School engagement: Potential of the concept, state of the evidence. Review of Educational Research, 74(1), 59–109.

10. Henrie, C. R., Halverson, L. R., & Graham, C. R. (2015). Measuring student engagement in technology-mediated learning: A review. Computers & Education, 90, 36–53.

11. Jeganathan, S., Lakshminarayanan, A. R., Parthasarathy, S., Khan, A. A. A., & Sathick, K. J. (2024). OptCatB: Optuna hyperparameter optimization model to forecast the educational proficiency of immigrant students based on CatBoost regression. Journal of Internet Services and Information Security, 14(2), 111–132. https://doi.org/10.58346/JISIS.2024.I2.008

12. Kuzilek, J., Hlosta, M., & Zdrahal, Z. (2017). Open university learning analytics dataset. Scientific Data, 4(1), Article 170171.

13. Nemirovski, A. (2004). Prox-method with rate of convergence O(1/t) for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. SIAM Journal on Optimization, 15(1), 229–251.

14. Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L. J., & Sohl-Dickstein, J. (2015). Deep knowledge tracing. In Advances in Neural Information Processing Systems (Vol. 28).

15. Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. Journal of Risk, 2, 21–42.

16. Tuy, H. (2004). Minimax theorems revisited. Acta Mathematica Vietnamica, 29(3), 217–229.

17. Whitehill, J., Serpell, Z., Lin, Y. C., Foster, A., & Movellan, J. R. (2014). The faces of engagement: Automatic recognition of student engagement from facial expressions. IEEE Transactions on Affective Computing, 5(1), 86–98.