Preferential Mixture-of-Experts: Interpretable Models that Rely on Human Expertise As Much As Possible

Melanie F. Pradier, PhD[1,2,*], Javier Zazo, PhD[1,2,*], Sonali Parbhoo, PhD[1,*],

Roy H. Perlis, MD MSc[3,4], Maurizio Zazzi, MD[5], Finale Doshi-Velez, PhD[1]

[1]School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA;
[2]Health Intelligence, Microsoft Research, Cambridge, Cambridgeshire, UK;
[3]Center for Quantitative Health, Massachusetts General Hospital, Boston, MA, USA;
[4]Harvard Medical School, Boston, MA, USA [5]University of Sienna, Italy.

arXiv:2101.05360v1 [cs.LG] 13 Jan 2021

Abstract *We propose Preferential MoE, a novel human-ML mixture-of-experts model that augments human expertise in decision making with a data-based classifier only when necessary for predictive performance. Our model exhibits an interpretable gating function that provides information on when human rules should be followed or avoided. The gating function is maximized for using human-based rules, and classification errors are minimized. We propose solving a coupled multi-objective problem with convex subproblems. We develop approximate algorithms and study their performance and convergence. Finally, we demonstrate the utility of Preferential MoE on two clinical applications for the treatment of Human Immunodeficiency Virus (HIV) and management of Major Depressive Disorder (MDD).*

1 Introduction

In the last few years, there has been a growth in the use of machine learning (ML) methods for decision-making in complex domains such as loan approvals, medical diagnosis and criminal justice. In particular, ML currently plays a key role in the healthcare sector for several tasks such as developing medical procedures [Hamid et al., 2017; Esteva et al., 2017], handling patient data and records [Pianykh et al., 2020] and treating chronic diseases [Raghu et al., 2019]. However, these algorithms typically require large amounts of data to make reasonable predictions. Additionally in the health sector, variability in practice between clinicians, patient heterogeneity, different disease prevalences, and confidentiality issues all result in final training cohorts being relatively small. Moreover, a clinician is often faced with rare events or outlier cases, where classic ML approaches suffer from insufficient training samples. In each of these scenarios, it is crucial to be able to incorporate clinical experience and domain knowledge.

Specifically, in practice, clinicians often rely on relatively simple human-based rules that reflect reasonable approaches to handle a situation. These rules can be seen as an additional source of knowledge that can be leveraged when building ML systems for clinical decision-support. For instance, clinicians treating patients with HIV tend to adhere to a list of guidelines for administering first and second-line therapies specified by several organizations [Organization, 2017; OARAC, 2017]; other well-known guidelines exist for prescribing antidepressants to address Major Depressive Disorder (MDD) [Lage et al., 2020]. Often, these rules provide benefits that are not easily formalized into a machine learning objective, for example, in terms of safety [Stone et al., 2009], or tolerability [Blumenthal et al., 2014] (e.g., not giving excitatory drugs to a patient that has insomnia). Thus, one might prefer an ML system that agrees with these human-based rules as much as possible.

Several ML methods have been proposed that combine human expertise in conjunction with training data to perform a prediction task [Mozannar and Sontag, 2020; Gennatas et al., 2020]. Some of these methods such as [Madras et al., 2018] explicitly focus on modeling the *interaction* between an automated ML model and an external decision-maker; the decision-maker determines whether to reject a particular decision made by the model based on the model's confidence and the expertise of the decision-maker. An extension to this procedure in [Mozannar and Sontag, 2020] describes when to defer decisions to a downstream decision-maker based solely on samples of the expert's decisions. In contrast to these approaches, we propose a ML system that complements human expertise only when needed, that is, it gives preference to human-based rules as much as possible, subject to explicit performance constraints in the optimization problem.

*Equal contribution.

In this work, we develop a novel mixture-of-experts (MoE) approach, called *Preferential MoE*, that explicitly incorporates human expertise in learning to provide predictions that align with human-based rules as frequently as possible without losing performance. The MoE framework allows for an intuitive way to combine ML with clinical expertise. Importantly, Preferential MoE provides a means of enforcing preference for the human decision rules, as well as an interpretable gating function that allows us to understand when data-driven or clinical expertise should be used. Specifically, we identify when a human decision rule should be followed, and when it makes more sense to provide an alternative data-driven prediction. Overall, by explicitly incorporating and optimizing for human expertise in our predictions, we obtain models that aligns better with human knowledge, making them easier to inspect, audit and trust.

2 Related Work

Human-ML decision making systems. There is a long history of approaches to incorporate human expertise in the architecture of ML systems. In particular, Towell and Shavlik [1994] and Tran and d'Avila Garcez [2018] propose methods that map rules to elements of a neural network. Wu et al. [2018] incorporates human-based knowledge gates into Recurrent Neural Networks for question answering or text matching. Closer in spirit, Wang et al. [2018] constrained a ML model to be more credible by relying as much as possible on input predictors that are intuitive for human experts. All these approaches include human expertise as input or intermediate features, whereas we assume that the expert information is available in the form of output decision rules, on which we want to rely as much as possible. Recently, Chattha et al. [2019] learns a ML system complementary to humans by modeling the residual of humans in the context of timeseries. Here we focus on classification, and additionally provide an interpretable explanation about when to rely on human-based rules. Finally, Hu et al. [2016] proposes a knowledge distillation approach, where human decisions are used as a teacher, and a student network is trained to mimic the human decisions while performing well on test data. Unlike implicitly assuming human expertise as additional ground truth labels (teacher), this work has the capacity of ignoring human rules if those are found unreliable.

Mixture of Experts. In the ML community, mixture-of-expert (MoE) models [Jacobs et al., 1991; Jordan and Jacobs, 1994] are frequently used to leverage different types of expertise in decision-making. The model works by explicitly learning a partition of the input space such that different regions of the domain may be assigned to different specialized sub-models or experts. MoEs have also been applied to several healthcare domains such as HIV [Parbhoo et al., 2017, 2018]. The proposed approach Preferential MoE is different in three regards: first, we explicitly incorporate human knowledge in the form of therapy standards and guidelines for medical decision-making. Second, our framework expresses an explicit preference for a specific expert (human-based), and trains an ML-based expert to complement the primary expert; third, we learn an interpretable gating function, which makes the model easy-to-interpret and give us information on when human-based rules should be followed.

Learning to defer approaches. [Madras et al., 2018; Mozannar and Sontag, 2020] propose MoE classification models to be used as triage tools, where only the most critical decisions are deferred to a medical expert, whilst relying on data-driven approaches the majority of the time. Specifically, these classifiers are trained based solely on the samples of an expert's decisions. Other approaches for integrating human expertise in decision-making such as [Raghu et al., 2019; Wilder et al., 2020] train a standard classifier on the data and subsequently obtain uncertainty estimates based on this classifier and the human expert. The decision is ultimately deferred to the expert with the lowest uncertainty. Unlike triage methods, we view human expertise as *complementary* to data-driven approaches and explicitly leverage these sources of knowledge to inform better predictions. That is, we optimize to rely on human expertise as much as possible, except for those regions for which human-based rules are inadequate. Our training samples consist of generic (potentially partial) rules that have been specified by humans a priori.

3 Methodology

In this section, we present Preferential MoE. The proposed approach fulfills two desiderata. First, Preferential MoE relies on the human rules as much as possible while preserving predictive performance. When the human-based rules are damaging w.r.t the prediction task, the proposed approach is able to overrule them (that is, we recover the same solution as the unconstrained standard MoE formulation). Second, the gating function is interpretable, providing information on when each human guideline is applicable.

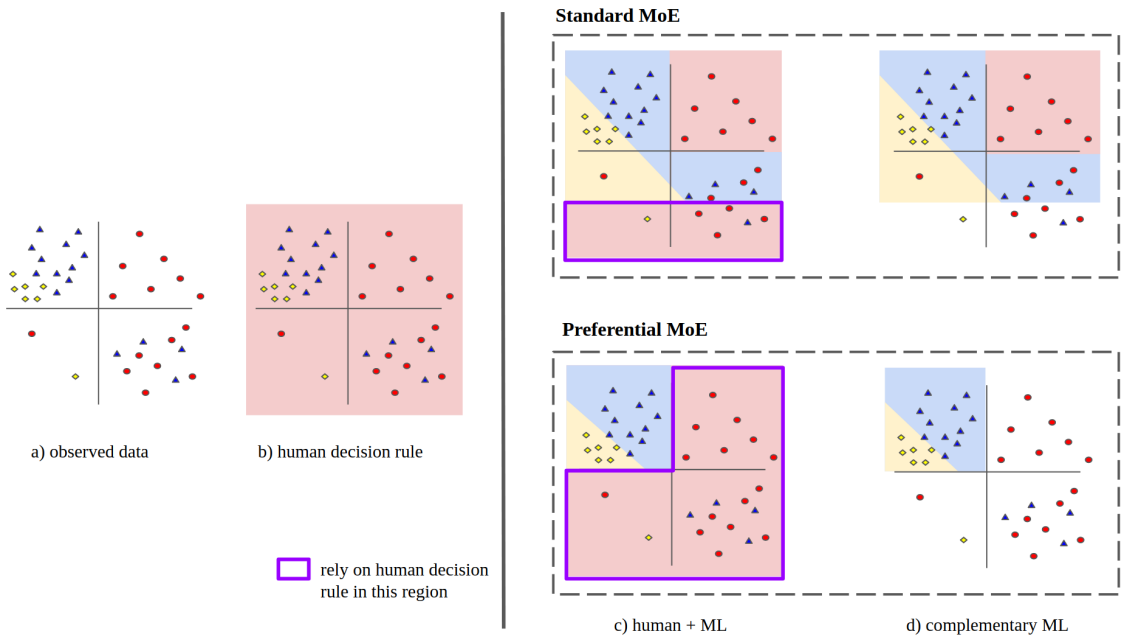

Figure 1: Preferential MoE: a mixture-of-experts (MoE) approach that relies on human decision rules subject to performance guarantees, and only disagrees with humans when a data-driven model can do better. A standard MoE exhibits similar predictive performance, but relies less on human rules. The proposed approach also provides insights on when/why human rules should be followed or not via an interpretable gating function (region highlighted in purple).

An overview of how Preferential MoE operates is illustrated in Figure 1. Colors in columns b)-d) represent predictive decision boundaries. In the proposed example, the human-based rule predicts red everywhere in the input space (sketch b). The third column (sketch c) shows the final predictions (colors) for each region of the input space. Each prediction either comes from the human decision rule, or from a data-based ML classifier. We learn a gating function (highlighted in purple) to select which classifier to rely on, as well as a complementary ML classifier to make predictions in regions outside of the purple region. In this diagram, both the standard MoE and the preferential MoE exhibit same predictive performance; however, the preferential MoE relies on humans much more often.

More formally, let $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$ be a dataset of observations where $\mathbf{x}_n \in \mathbb{R}^d$ are the covariates and $y_n \in \{1, ..., K\}$ is a categorical outcome for a specific prediction task. Let the *guideline* function $g : \mathbb{R}^d \to \{1, 0, -1\}$ be an aggregated function encoding all available human decision rules, whose input are the covariates, and the output might be any output category, or a flag (-1) indicating that the rule is not applicable (human does not know). This human guideline function g is fixed a priori by domain knowledge or well-established medical practice. In the case of not having access to an explicit human function g, but samples of past human decisions instead, we can pre-train a classifier to mimic those human decisions beforehand, and use such classifier as our human-based rules function g.

Given dataset $\mathcal{D}$, there might exist several functions that exhibit similar predictive performance, but are qualitatively different. We want to use expert knowledge (via human-based rules) to guide the optimization such that we are able to find models that have high predictive performance and agree with the human-based rules as much as possible. In order to accomplish that, we will include both objectives in the proposed optimization.

Modeling. Our goal is to make predictions that prioritize human-based rules when the data supports (or does not contradict) such knowledge, and learn to defer to another trainable ML expert when the human rules counters empirical evidence. For that, we propose a new classification model based on a mixture of experts formulation. Let $f_\theta : \mathbb{R}^d \to \Delta^K$ be a trainable ML expert parameterized by θ, where Δ^K denotes the $(K-1)$-simplex (outcome vectors of f_θ should sum to one). Our approach combines the predictions of the ML expert f_θ and the human expert g via the gating

function $\rho_w : \mathbb{R}^d \to \{0, 1\}$ parametrized by w; ρ_w is another classifier that selects which expert to rely on given the covariates $\mathbf{x}$. The prediction model of Preferential MoE is formalized as

$$\hat{y}_{\theta,w}(\mathbf{x}) = \begin{cases} (1 - \rho_w(\mathbf{x}))f_\theta(\mathbf{x}) + \rho_w(\mathbf{x})g(\mathbf{x}) & \text{if } g(\mathbf{x}) \neq -1 \\ f_\theta(\mathbf{x}) & \text{if } g(\mathbf{x}) = -1 \end{cases} \tag{1}$$

where $\hat{y}_{\theta,w}(\mathbf{x}) \in \Delta^K$, and the likelihood function is given by

$$y|\mathbf{x} \sim \text{Categorical}\left(\hat{y}_{\theta,w}(\mathbf{x})\right). \tag{2}$$

Note that the ML expert f_θ might make predictions and specialize in input regions where the human rule g is not applicable or is inaccurate. In summary, Equation (2) assumes that every data point $\mathbf{x}$ can be discriminated by f_θ or g, and ρ_w makes a deterministic decision on which expert to rely on. When learning ρ_w, we will prioritize human-based rules g during inference. Notice that (1) produces a non-convex prediction model which may be difficult to optimize. MoE classification models are notorious for converging to local optima [Jordan and Xu, 1995].

The gating function selects when a decision should rely on human-based rule or a trained expert. By making ρ_w an interpretable function, e.g., a linear classifier or a decision tree, the model learns which features are important for human-based decision making, and identifies the regions of other expert classifiers. We note that, even if the gating function ρ_w is chosen to be interpretable, our approach does not provide theoretical guarantees on identifying all the regions suitable for human decision rules. More generally, ρ_w can also be a non-interpretable function, e.g., a neural network. In such case, the gating function still identifies regions appropriate for human-based decisions, although it may miss the interpretability of the parameters w. Overall, our framework allows model constructions that balance flexibility and interpretability suitable to different applications.

Before continuing, we note that our Preferential MoE differs from other MoE approaches that learn to defer in the following ways. Madras et al. [2018] trains multiple experts and learns when to defer to a human, based on accuracy, uncertainty and fairness, via regularization of the loss function. Mozannar and Sontag [2020] proposes a novel cost sensitive function, studies its theoretical properties, and learns the regions of space where each expert is accurate. Neither of these approaches prioritize human rules, so it is possible that given an expressive expert, the mixture learns to reject human-based rules and relies solely on the trained expert. Such deferral may come at the cost of not truly identifying the input regions where human rules are adequate. Because in many applications identifiability may be critical to understand when humans' decisions are suitable, we propose a new inference procedure that defers primarily to humans over other methods while preserving pre-specified performance guarantees. Preferential MoE will maximize use of human-based rules with an explicit constraint that the performance cannot be signficantly worse than the standard MoE without preferences.

Problem formalization. Our formulation as an optimization problem needs to reflect the following criteria: (i) we want to minimize the predictive error, (ii) we want to follow human-based rules as frequently as possible without hurting performance.

We optimize for predictive performance by minimizing the cross-entropy $L^\gamma_{\theta,w}(\mathcal{D})$ with respect to the predictions from Equation (1); this corresponds to a standard maximum log-likelihood estimator for the probabilistic model in Equation (2) with an additional regularizer. For example, if the outcome is binary we write

$$L^\gamma_{\theta,w}(\mathcal{D}) = \sum_{n=1}^{N} -\left[y_n \ln(\hat{y}_{\theta,w}(\mathbf{x}_n)) + (1 - y_n)\ln(1 - \hat{y}_{\theta,w}(\mathbf{x}_n))\right] + \gamma||w||_1, \tag{3}$$

where $\gamma \geq 0$ is a regularization weight that controls the trade-off between predictive performance and sparsity of w. A sparse w can help identify important features for the gating function $\rho_w(x)$.

We bound the cross-entropy loss $L^\gamma_{\theta,w}(\mathcal{D})$ with a prefixed optimized value for performance guarantees. Denote $L^\gamma_{\theta^*,w^*}(\mathcal{D})$ an attainable loss where θ^* and w^* are solutions of minimizing $L^\gamma_{\theta,w}(\mathcal{D})$ for the stated MoE in Equation (2). Consider a margin $\varepsilon \geq 0$ measuring an acceptable performance decrease, and consider the constraint:

$$L^\gamma_{\theta,w}(\mathcal{D}) \leq (1 + \varepsilon)L^\gamma_{\theta^*,w^*}(\mathcal{D}). \tag{4}$$

Equation (4) guarantees that the performance loss will not increase more than specified, and will maintain predictive error results. We introduce sets $\Theta \subset \mathbb{R}^q$ and $W \subset \mathbb{R}^p$ such that $\theta \in \Theta$ and $w \in W$. Variables θ and w do not need to have same dimensions and can be constructed with different model classifiers.

We present next the problem formulation for Preferential MoE:

$$\mathcal{G}: \quad \text{(player 1)} \quad \begin{aligned} \min_{w \in W} \quad & -\Sigma_{n=1}^{N} \ln(\rho_w(x)) \\ \text{s.t.} \quad & L^{\gamma}_{\theta,w}(\mathcal{D}) \leq (1+\varepsilon) L^{\gamma}_{\theta^*,w^*}(\mathcal{D}) \end{aligned} \qquad \text{(player 2)} \quad \min_{\theta \in \Theta} \quad L^{\gamma}_{\theta,w}(\mathcal{D}). \tag{5}$$

We refer to (5) as game $\mathcal{G}$. Using game theory terminology, there are 2 players and each player optimizes their own objective, variables and constraints, while taking into account the other player's decisions. Notice that $\mathcal{G}$ explicitly models our discussed goals: player 1, which is optimizing the gating function, maximizes the number of human-based decisions; player 2, which optimizes the classifier f_θ, minimizes prediction error according to the loss function (3). Note that the negative logarithm is a monotone transformation that helps obtain a convex objective for player 1. Player 1 also imposes the performance constraint and limits the classification loss.

$\mathcal{G}$ is a particular instance of a generalized Nash equilibrium problem (GNEP) [Pang and Fukushima, 2005]. Our goal is to minimize both objectives and reach an equilibrium point known as Nash equilibrium, where no player is incentivized to change its decision based on the other player's actions. In general, finding the Nash equilibrium is particularly challenging as a result of the dynamic nature of the feasibility region.

Existence of solutions is guaranteed for $\mathcal{G}$ as discussed in Proposition 1 (see Appendix B). Next, we present two algorithms for solving (5), and discuss their properties and convergence.

4 Inference Algorithms

Inference for determining θ and w from $\mathcal{G}$ proceeds in two steps:

1. **Unconstrained optimization**: we train a standard MoE model from Equation (2) by minimizing the performance loss $L^{\gamma}_{\theta,w}(\mathcal{D})$ described in Equation (3). This step yields a performance reference value of $L^{\gamma}_{\theta^*,w^*}(\mathcal{D})$ which we will aim to maintain up to a certain margin ε. We use the optimal parameters θ^* and w^* from the unconstrained problem as warm initialization for the next step.

2. **Constrained optimization**: we solve game $\mathcal{G}$ initializing from previous solution.

We discuss two algorithms for solving $\mathcal{G}$. The first proposal combines both objectives and uses a log-barrier method to approximate a solution. The second proposal takes gradient steps that minimize each objectives alternatively and projects to the feasible region. Both methods have convergence guarantees.

Log-Barrier Method. We want to approximate a solution of $\mathcal{G}$ by simplifying its formulation. We move player 1's constraint to the objective using a log-barrier penalty used in interior point methods [Boyd and Vandenberghe, 2004, Chapter 11] and combine both objectives (see Appendix A). These operations transform $\mathcal{G}$ into the following unconstrained non-convex optimization problem:

$$\min_{\theta \in \Theta, w \in W} \quad -t \sum_{n=1}^{N} \ln\left(\rho_w(x_n)\right) - \ln\left((1+\varepsilon) L^{\gamma}_{\theta^*,w^*}(\mathcal{D}) - L^{\gamma}_{\theta,w}(\mathcal{D})\right). \tag{6}$$

The first term of equation (6) corresponds to player 1's objective, and the second term to the log-barrier function $\widehat{I}(u) = -1/t \ln(-u)$ transforming its constraint, which also aligns with player 2's objective. Note that a solution for problem (6) exists, provided Assumption 1 in Appendix B holds.

The log-barrier argument is susceptible of becoming negative inside the logarithm and be a source of numerical instability, so care needs to be taken with step sizes and correct initialization (warm-start). Parameter t is a hyperparameter

that weights the satisfiability of the constraint, and the approximation improves as t grows. Note that this approximated form encourages that the difference $L^{\gamma}_{\theta^*,w^*}(\mathcal{D}) - L^{\gamma}_{\theta,w}(\mathcal{D})$ becomes large, regardless of the constraint already being satisfied. This has the desirable effect of continuously minimizing $L^{\gamma}_{\theta,w}(\mathcal{D})$. Finally, because of the non-convex nature of the problem, gradient descent methods only guarantee convergence to stationary solutions.

Projected Gradient Method. Player 2's decisions affect player 1's constraint, and player 1's affect player's 2 objective in game $\mathcal{G}$. A simple algorithm would be to alternate solving subproblems and repeat until convergence. Such schemes are only guaranteed to converge under very stringent conditions of monotonicity of the game. Monotonocity is a desirable property of multivariate mappings, informally stating that a small change in the input guarantees a bounded change in the output, therefore permitting dynamics of control towards stable solutions. We refer the reader to Scutari et al. [2012] for definitions, properties and algorithms for solving monotone games.

We present Algorithm 1 for solving $\mathcal{G}$. The algorithm makes a gradient update on each objective, and projects the result onto the feasibility region. We denote estimates on iteration k with θ^k and w^k. The feasibility region is denoted with $K_\varepsilon(\theta^{k+1})$, and is formally introduced in Appendix B. The operation $\Pi_{K_\varepsilon(\theta^{k+1})}$ denotes projection of w onto the set $K_\varepsilon(\theta^{k+1})$. The projection operation solves the following optimization problem

$$\Pi_{K_\varepsilon(\theta^{k+1})}\big(z\big) = \underset{w \in \big(W \cap K_\varepsilon(\theta^{k+1})\big)}{\arg\min} \quad \frac{1}{2}\|w - z\|^2, \tag{7}$$

whose solution can be efficiently computed via a bisection search, described in Algorithm 2. The optimization inside the while loop in Algorithm 2 can be solved via L-BFGS [Liu and Nocedal, 1989], or other efficient algorithm, since it does not have a known analytical expxression.

We state the convergence of Algorithm 1 to a Nash Equilibrium in the next theorem. Proof is provided in Appendix B.

Theorem 1 (Convergence)**.** *Assume Assumptions 1 to 3 are satisfied (see Appendix B), and step-size α is small enough. Then, $\mathcal{G}$ is monotone and Algorithm 1 converges to a Nash Equilibrium of game $\mathcal{G}$.*

Algorithm 1: Projected Gradient Descent

Input: $\mathcal{D}$, ε, $L^{\gamma}_{\theta^*,w^*}$, $\{\,\alpha^k\,\}$
Output: θ and w.
Initialization: $\theta^0 \leftarrow \theta^*$, $w^0 \leftarrow w^*$, $k \leftarrow 0$;
while *stopping criteria not satisfied* **do**
  $\theta^{k+1} \leftarrow \theta^k - \alpha^k \frac{\partial}{\partial\theta} L^{\gamma}_{\theta^k,w^k}$;
  $w^{k+1} \leftarrow \Pi_{K_\varepsilon(\theta^{k+1})}\big(w^k - \alpha^k \frac{\partial}{\partial\theta} L^{\gamma}_{\theta^{k+1},w^k}\big)$;
  $k \leftarrow k+1$
end

Algorithm 2: $\Pi_{K_\varepsilon(\theta^{k+1})}$ (bisection search)

Input: θ^{k+1}, ε , $z \in \mathbb{R}^p$, $\overline{\lambda}$; Initialization: $\underline{\lambda} \leftarrow 0$
Output: w^{k+1}.
while $(\overline{\lambda} - \underline{\lambda}) \geq$ *tolerance* **do**
  $\lambda \leftarrow (\overline{\lambda} + \underline{\lambda})/2$;
  $w \leftarrow \arg\min_{w \in W} \quad \frac{1}{2}\|w - z\|^2 + \lambda L^{\gamma}_{\theta,w}(\mathcal{D})$;
  if $L^{\gamma}_{\theta,w}(\mathcal{D}) - (1+\varepsilon) L^{\gamma}_{\theta^*,w^*}(\mathcal{D}) > 0$ **then** $\underline{\lambda} \leftarrow \lambda$;
  else $\overline{\lambda} \leftarrow \lambda$;
end

5 Results

We compare the performance of Preferential MoE against several baselines for two medical tasks for the treatment of Human Immunodeficiency Virus (HIV), or pharmacological management of Major Depressive Disorder (MDD). Our baselines include using predictions a) based on a human expert alone; b) a logistic regression ML expert alone; c) a standard mixture-of-experts model (standard MoE); d) the learn-to-defer model in [Madras et al., 2018]; and e) a learn-to-defer model from [Mozannar and Sontag, 2020]. For the standard MoE and Preferential MoE, we train models either assuming discrete $\rho(x)$ values to begin with, or assuming continuous $\rho(x)$ values and the discretizing at the end, exploring all operating points for the threshold of the gating function. Here we report the latter, which seems to work better in practice.

Hyperparameter selection. For both prediction tasks, we explore different learning rates for both, the unconstrained and constrained optimization steps, in the range of $\{10^{-4}, 10^{-3}, 0.01, 0.1\}$. We also explore a range of regularization parameters $\gamma \in \{0.0, 0.001, 0.01, 0.05, 0.1, 1.0\}$ for the gating function, and select those that maximize predictive performance in a validation set. For the psychiatry dataset, we additionally regularize the ML classifier with an L1 penalty to avoid overfitting due to the high-dimensionality of the input space. We fix the margin $\varepsilon = 0.1$, and the trade-off parameter $t = 5.0$ for the log-barrier penalty in Equation (6). Our results were stable to perturbations of these parameters. Intuitively, t can be matched to existing interior-point algorithms and is quite robust with appropriate gradient step sizes. The margin ϵ affects model´s accuracy, but even if there is no direct mapping from its value to a desired performance level, its impact was similar in the range $\epsilon \in [1e^{-2}, 2e^{-1}]$. Setting ϵ too small can make the model not move from the initialization point, and its solution stay similar to the standard MoE's.

Evaluation metrics. To evaluate Preferential MoE and other baselines, we measure performance as Area-Under-the-operating-ROC-Curve (AUC), as well as predictive accuracy (percentage of correct predictions) for a fine-grid of threshold values, both for the gating function and final predictions. Note that all thresholds are chosen by cross-validation, we thus guarantee that the right thresholds (w.r.t the most adequate metric for each downstream task) are selected, in a data-driven manner. We report *coverage* as a measure of how frequently (in percentage) each model relies on the human-based guideline function g. More specifically, we define *soft-coverage* and *hard-coverage*(t) for a given gating function threshold t as follows:

$$\text{soft-coverage} = 100.0 \times \mathbb{E}[\rho(x)] \qquad \text{hard-coverage}(t) = 100.0 \times \mathbb{E}\big[\mathbb{1}[\rho(x) \geq t]\big]. \tag{8}$$

5.1 Human Immunodeficiency Virus (HIV) Therapy Outcome Prediction

HIV currently affects more than 36 million people worldwide. The life-long use of combinations of antiretrovirals has largely helped combat the virus in most parts of the world and has transformed the virus from a life-threatening condition to a chronic illness. However, administering therapies is tricky as patients frequently suffer from drug resistance, viral relapses or spikes, as well as adherence issues and several other side-effects from use of antiretrovirals.

We identified individuals between 18-72 years of age from the EuResist database comprising of genotype, phenotype and clinical information of over 65 000 individuals in response to antiretroviral therapy administered between the years 1983 and 2018. We focus on a subset of 36 780 of these patients who received at least 3 prior treatments and base our predictions on the genotype, phenotype, clinical and demographic information of these individuals. The curated dataset contains a total of 384 such features. Our goal is to predict short-term therapy success where viral suppression is maintained for at least 40 days after a therapy is administered.

Table 1 shows predictive performance and coverage results for the proposed approach and competing baselines. Compared to other approaches, Preferential MoE exhibits highest soft coverage while either retaining or improving predictive performance. Figure 3 compares the accuracy relative to hard thresholding of the coverage for each of the MoE models. In the HIV setting, both variants of the Preferential MoE outperform the standard MoE approach at various coverage values. At 60% coverage, the methods all seem to perform relatively similarly in terms of accuracy.

	AUC		soft coverage (%)	
Baselines	mean	CI	mean	CI
ML only	0.64	[0.63-0.65]	0.00	[0.00-0.00]
Learn-to-defer[Madras et al., 2018]	0.71	[0.68-0.72]	54.07	[48.18 - 55.63]
Consistent Learn-to-defer [Mozannar and Sontag, 2020]	0.66	[0.62-0.69]	56.81	[50.02 - 57.62]
Standard MoE (unconstrained)	0.69	[0.69-0.70]	52.87	[51.19-54.55]
Preferential MoE (log barrier)	**0.74**	[0.72-0.76]	62.06	[60.8-63.32]
Preferential MoE (projected gradient)	**0.74**	[0.73-0.75]	**63.18**	[61.7-64.66]

Table 1: Performance vs Coverage (HIV): Preferential MoE relies much more often on human expertise while preserving predictive performance. Predictive performance measured by Area-Under-the-operating-ROC-Curve (AUC); Reliance on human decision rules based on soft coverage, as defined in Equation (8).

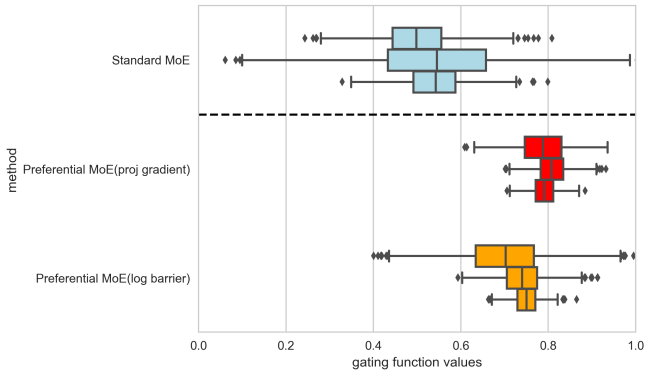

Figure 2: $\rho(x)$ values in the test set for HIV. Preferential MoE pushes up the values for the gating functions, favoring human decision rules more frequently in the input space. Each box plot corresponds to a different random seed (we report 3 different initializations per method).

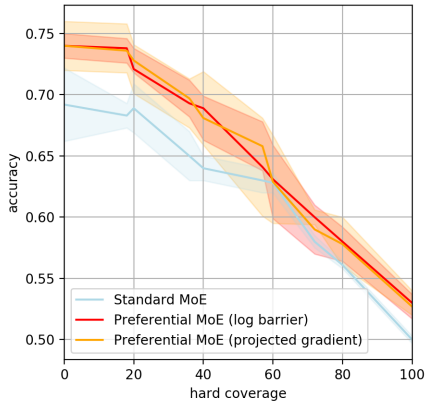

Figure 3: Accuracy-coverage trade-off. Preferential MoE (trained by the log barrier or projected gradient method) for HIV either relies more on human rules for the same predictive accuracy, or gets higher accuracy for the same coverage with human rules.

Importantly, Preferential MoE allows us to incorporate human expertise into the prediction task and provides us with insights of when it makes sense to follow the rules based on the gating function. Table 2 provides a sparse list of predictors and corresponding weights averaged over 10 random seeds for the gating function. These predictors are associated with regions where it makes sense to follow human intuition. While Standard MoE identifies blood count data, certain mutations and a patient's risk group as meaningful factors, Preferential MoE identifies a significantly different set of predictors. Notably, many of the predictors identified in the latter correspond to cases where patients have additional conditions such as lipodystrophy or side effects to medication where it is preferable to rely on human judgement to determine how to treat these individuals. Figure 2 compares the gating function values $\rho(x)$ in the test set for HIV. Unsurprisingly, Preferential MoE shows a higher preference for relying on human rules.

Weight w	Covariate Description
+0.1612 $\pm$ 0.014	CD8+ cell count (cells/ml)
-0.1161 $\pm$ 0.002	Reverse Transcriptase Mutation 67N
-0.0310 $\pm$ 0.025	Protease Mutation 20M
0.0280 $\pm$ 0.001	Blood count; complete (CBC)
-0.0195 $\pm$ 0.005	Co-infection of Hepatitis C
-0.0156 $\pm$ 0.001	Stavudine
-0.0124 $\pm$ 0.011	Reverse Transcriptase Mutation 215YF
+0.0121 $\pm$ 0.020	Nevirapine
-0.0068 $\pm$ 0.031	Risk group MSM
-0.0055 $\pm$ 0.005	Age

Weight w	Covariate Description
+0.0359 $\pm$ 0.022	CD4 + cell count (cells/ml)
-0.0236 $\pm$ 0.027	Baseline Viral Load
+0.0151 $\pm$ 0.030	High Adherence
+0.0150 $\pm$ 0.001	Number of Prior Treatment Lines
+0.076 $\pm$ 0.007	Pregnancy
-0.0055 $\pm$0.016	Reverse Transcriptase Mutation 184V
-0.0035 $\pm$0.002	Race black
-0.0026 $\pm$0.001	Lamivudine
+0.0025 $\pm$ 0.003	Anaemia
+0.0012 $\pm$ 0.007	Lipodystrophy

Table 2: Interpretation of gating function (HIV). Sparse list of predictors describing the regions where human decision rules are followed. We report weight parameters averaged across 10 different random seeds, and for regularization γ=0.1). (left) Standard MoE (predictors after step 1 in training); (right) preferential MoE (predictors after step 2 in training). Highlighted in red/green are those predictors that disappear/pop-up after step 2 in training.

5.2 Prediction of Antipsychotic for Major Depressive Disorder (MDD)

Antidepressant prescription for MDD often involves trial and error. Roughly 2/3 of individuals diagnosed with MDD do not yield remission with their initial treatment, and 1/4 of patients is expected to dropout against clinical advice

before finishing their treatment [Hughes et al., 2020; Pradier et al., 2020]. The list of potential side-effects translates in tolerability and safety concerns that need to be taken into account while prescribing antidepressants. Here we focus on predicting prescription of antipsychotics, which is a class of medication primarily used to manage psychosis, but often used as an adjunctive treatment in the pharmacological management of MDD. The guideline function g for this prediction task is as follows: if the patient has anxiety or insomnia, promote antipsychotic (predict positive label), if the patient has overweight, avoid antipsychotic (predict negative label).

We identified individuals age 18-80 years drawn from the outpatient clinical networks of two academic medical centers in New England, Massachusetts General Hospital and Brigham and Women's Hospital. These patients had received at least one electronically-prescribed antidepressant between March 2008 and December 2017 with a diagnosis of MDD or depressive disorder at the nearest visit to that prescription. The goal is to predict prescription of antipsychotic based on demographic information (gender, race) as well as diagnostic and procedure codes. Race and gender were self-identified features and were included as a proxy for socio-economic variables. The curated dataset consists of 3,865 individuals and 1,680 features.

Table 3 shows predictive performance and soft coverage results for the proposed approach and competing baselines, averaged across 5 random initializations. We encountered issues training the Learn-to-defer approaches to this data (probably due to its high-dimensionality), so we only include the other baselines. Preferential MoE exhibits highest soft coverage (reliance on human rules) while maintaining (or even slightly improving) predictive performance.

	AUC		soft coverage (%)	
Baselines	mean	CI	mean	CI
ML only	0.70	[0.69-0.71]	0.00	[0.00-0.00]
Standard MoE (unconstrained)	0.71	[0.70-0.71]	31.41	[28.48-34.56]
Preferential MoE (log barrier)	**0.72**	[0.71-0.73]	**48.34**	[46.24-51.74]
Preferential MoE (projected gradient)	0.72	[0.71-0.72]	45.06	[42.85-46.70]

Table 3: Performance vs Coverage (Psychiatry): Preferential MoE relies much more often on human expertise while preserving predictive performance. Predictive performance measured by Area-Under-the-operating-ROC-Curve (AUC); Reliance on human decision rules based on soft coverage, as defined in Equation (8).

Preferential MoE gives us additional information on when to follow such human rules by inspecting the gating function. Table 4 presents the sparse list of predictors for the gating function, associated to regions where human decision rules are followed. By regularizing the gating function classifier with an L1-penalty, we get concise list of predictors to describe those regions. The list on the left correspond to Standard MoE (unconstrained optimization), and the list on the right correspond to Preferential MoE (constrained optimization maximizing reliance on humans). In both lists, most predictors corresponding to general patient care (examination, hospital care, etc) are negatively-correlated: this can be interpreted as higher reliance on humans in the absence of patient care related codes. In the case of Preferential MoE, additional covariates coding for cardiovascular risk factors (highlighted in green) are positively-correlated with reliance on human rules. Such information can be used to explore refinements of the human-based rules.

Figure 4 compares the histogram of the gating function values $\rho(x)$ in the test set. As expected, Preferential MoE pushes those values up, reflecting a preference for relying on human rules when possible. Although these values are continuous, we can discretize them using a specific threshold v calibrated on the validation set. Each threshold v yields a different trade-off between accuracy and coverage. Figure 5 shows the trade-off between accuracy and hard coverage reachable by these models. As a reference point, the human decision rules have an accuracy of 49.87% for this prediction task. The curves are averaged over 10 different random seeds, each curve is obtained by changing the thresholds for the gating function and final decision. Overall, Preferential MoE is able to reach better trade-offs, either better accuracy for a given fixed hard coverage, or more hard coverage for a given accuracy level.

6 Limitations

First, human guidelines may not be available for specific applications directly. In such cases, an auxiliary classifier can be trained to predict human labels, as done in Madras et al. [2018]; Mozannar and Sontag [2020]. Such a classifier

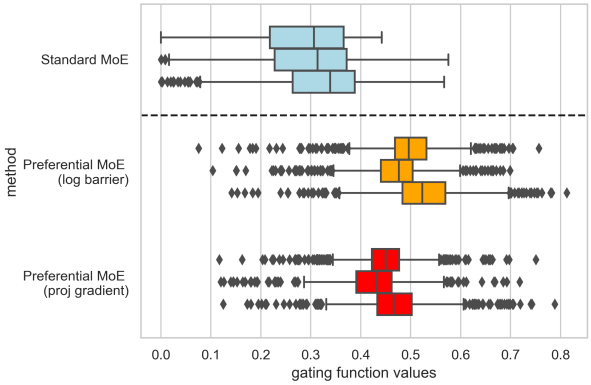

Figure 4: Histograms for $\rho(x)$ in the test set. Preferential MoE pushes up the values for the gating function, favoring relying on human decision rules more frequently. Each box plot corresponds to a different random seed (3 per method).

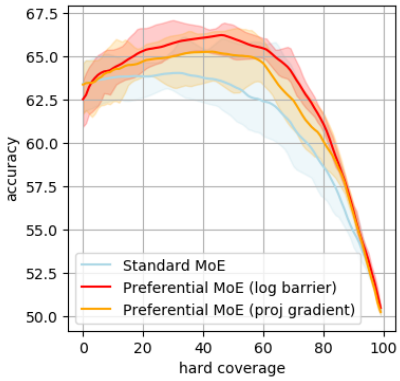

Figure 5: Accuracy-coverage trade-off. Preferential MoE either relies more on human rules for the same predictive accuracy, or gets higher accuracy for the same coverage with human rules.

Weight w	Covariate Description
-0.0303 $\pm$ 0.0143	Subsequent hospital care
-0.0242 $\pm$ 0.0152	MDD, recurrent episode
-0.0235 $\pm$ 0.0126	Psychiatric examination
-0.0208 $\pm$ 0.0071	Depressive disorder
-0.0153 $\pm$ 0.0106	Anxiety state
-0.0117 $\pm$ 0.0089	Office or outpatient visit
-0.0083 $\pm$ 0.0033	Radiologic examination
-0.0073 $\pm$ 0.0062	Trazodone
-0.0068 $\pm$ 0.0049	Emergency department visit
-0.0031 $\pm$ 0.0062	race white

Weight w	Covariate Description
-0.0202 $\pm$ 0.0043	Subsequent hospital care
-0.0115 $\pm$ 0.0075	MDD, recurrent episode
-0.0110 $\pm$ 0.0107	Psychiatric examination
0.0103 $\pm$ 0.0041	Office or outpatient visit
-0.0070 $\pm$ 0.0112	Depressive disorder
0.0068 $\pm$ 0.0022	General medical examination
0.0061 $\pm$ 0.0022	Type II diabetes
0.0037 $\pm$ 0.0016	Hypertension
-0.0035 $\pm$ 0.0046	Anxiety state
-0.0034 $\pm$ 0.0083	Trazodone

Table 4: Interpretation of gating function. Sparse list of predictors describing the regions where human decision rules are followed. We report weights averaged across 10 different random seeds, and for a regularization parameter γ=0.1). (left) Standard MoE (predictors after unconstrained step 1 in training); (right) Preferential MoE (predictors after step 2 in training). Highlighted in red/green are those predictors that disappear/pop-up after step 2 in training.

can then be used as a proxy for a human expert in the MoE. Second, MoEs suffer from local optima, and our methods may converge to different sets of experts and gating functions, depending on initialization. We advice running multiple initializations and average results (as reported in Section 5) or select the best runs based on performance. Third, our method relies on data that might reflect societal biases, and it may thus suffer from these undesired effects. This is not a limitation specific to our approach; further analysis and exploration is actively pursuit by the community. Finally, we empirically showed that human rules were most useful in regions where data was less prevalent and predictions were less accurate, but our approach does not provide theoretical guarantees in that regard; we let this as future work.

7 Conclusion

We presented *Preferential MoE*, a mixture of experts that learns and combines a ML general classifier with a human expert, prioritizing the human-based rules. We presented a game formulation of two objectives, which we solve by a log-barrier method or alternating projected gradient descent. We evaluate both approaches in the prediction of HIV therapy success, and prescription of antipsychotic for MDD. Both algorithms preserve performance and maximize coverage of human-based decisions compared to other baselines, assuming soft and hard decision assignments of the gating function. Future work will further explore other MoE formulations balancing performance and global optimality of the MoE formulation.

References

1. Rina Foygel Barber and Wooseok Ha. Gradient descent with non-convex constraints: local concavity determines convergence. *Information and Inference: A Journal of the IMA*, 7(4):755–806, 2018.

2. Dimitri P Bertsekas. *Convex optimization theory*. Athena Scientific Belmont, 2009.

3. Sarah R. Blumenthal, Victor M. Castro, Caitlin C. Clements, Hannah R. Rosenfield, Shawn N. Murphy, Maurizio Fava, Jeffrey B. Weilburg, Jane L. Erb, Susanne E. Churchill, Isaac S. Kohane, Jordan W. Smoller, and Roy H. Perlis. An electronic health records study of long-term weight gain following antidepressant use. *JAMA psychiatry*, 71(8):889–896, August 2014. ISSN 2168-6238. doi: 10.1001/jamapsychiatry.2014.414.

4. Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, March 2004.

5. Muhammad Ali Chattha, Shoaib Ahmed Siddiqui, Muhammad Imran Malik, Ludger van Elst, Andreas Dengel, and Sheraz Ahmed. KINN: Incorporating Expert Knowledge in Neural Networks. *arXiv:1902.05653*, February 2019. URL `http://arxiv.org/abs/1902.05653`. arXiv: 1902.05653.

6. Andre Esteva, Brett Kuprel, Roberto A Novoa, Justin Ko, Susan M Swetter, Helen M Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115–118, 2017.

7. Francisco Facchinei and Jong-Shi Pang. *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media, 2007.

8. Stefan Feltenmark and Krzysztof C Kiwiel. Dual applications of proximal bundle methods, including lagrangian relaxation of nonconvex problems. *SIAM Journal on Optimization*, 10(3):697–721, 2000.

9. Anders Forsgren and Philip E Gill. Primal-dual interior methods for nonconvex nonlinear programming. *SIAM Journal on Optimization*, 8(4):1132–1152, 1998.

10. Efstathios D. Gennatas, Jerome H. Friedman, Lyle H. Ungar, Romain Pirracchio, Eric Eaton, Lara G. Reichmann, Yannet Interian, José Marcio Luna, Charles B. Simone, Andrew Auerbach, Elier Delgado, Mark J. van der Laan, Timothy D. Solberg, and Gilmer Valdes. Expert-augmented machine learning. *Proceedings of the National Academy of Sciences*, 117(9):4571–4577, 2020. ISSN 0027-8424. doi: 10.1073/pnas.1906831117. URL `https://www.pnas.org/content/117/9/4571`.

11. Kanza Hamid, Amina Asif, Wajid Abbasi, Durre Sabih, et al. Machine learning with abstention for automated liver disease diagnosis. In *2017 International Conference on Frontiers of Information Technology (FIT)*, pages 356–361. IEEE, 2017.

12. Zhiting Hu, Xuezhe Ma, Zhengzhong Liu, Eduard Hovy, and Eric Xing. Harnessing deep neural networks with logic rules. *arXiv:1603.06318*, 2016.

13. Michael C. Hughes, Melanie F. Pradier, Andrew Slavin Ross, Thomas H. McCoy, Roy H. Perlis, and Finale Doshi-Velez. Assessment of a Prediction Model for Antidepressant Treatment Stability Using Supervised Topic Models. *JAMA Network Open*, 3(5), May 2020. ISSN 2574-3805. doi: 10.1001/jamanetworkopen.2020.5308. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7240354/`.

14. Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.

15. Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994.

16. Michael I Jordan and Lei Xu. Convergence results for the em approach to mixtures of experts architectures. *Neural networks*, 8(9):1409–1431, 1995.

17. Isaac Lage, Melanie F. Pradier, Thomas McCoy, Roy Perlis, and Finale Doshi-Velez. Do clinicians follow heuristics in prescribing antidepressants? *submitted*, 2020.

18. Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.

19. David Madras, Toni Pitassi, and Richard Zemel. Predict responsibly: improving fairness and accuracy by learning to defer. In *Advances in Neural Information Processing Systems*, pages 6147–6157, 2018.

20. Hussein Mozannar and David Sontag. Consistent estimators for learning to defer to an expert. *arXiv:2006.01862*, 2020.

21. OARAC. Guidelines for the use of antiretroviral agents in adults and adolescents with hiv. *Panel on Antiretroviral Guidelines for Adults and Adolescents*, 2017. URL `https://aidsinfo.nih.gov/contentfiles/lvguidelines/adultandadolescentgl.pdf`.

22. World Health Organization. Tackling hiv drug resistance: trends, guidelines and global action. Technical report, World Health Organization, 2017. URL `https://apps.who.int/iris/bitstream/handle/10665/255881/WHO-HIV-2017.21-eng.pdf?sequence=1`.

23. Jong-Shi Pang and Masao Fukushima. Quasi-variational inequalities, generalized nash equilibria, and multi-leader-follower games. *Computational Management Science*, 2(1):21–56, 2005.

24. Sonali Parbhoo, Jasmina Bogojeska, Maurizio Zazzi, Volker Roth, and Finale Doshi-Velez. Combining kernel and model based learning for hiv therapy selection. *AMIA Summits on Translational Science Proceedings*, 2017:239, 2017.

25. Sonali Parbhoo, Omer Gottesman, Andrew Slavin Ross, Matthieu Komorowski, Aldo Faisal, Isabella Bon, Volker Roth, and Finale Doshi-Velez. Improving counterfactual reasoning with kernelised dynamic mixing models. *PloS one*, 13(11):e0205839, 2018.

26. Oleg S Pianykh, Steven Guitron, Darren Parke, Chengzhao Zhang, Pari Pandharipande, James Brink, and Daniel Rosenthal. Improving healthcare operations management with machine learning. *Nature Machine Intelligence*, 2 (5):266–273, 2020.

27. Melanie F. Pradier, Thomas H. McCoy Jr, Michael Hughes, Roy H. Perlis, and Finale Doshi-Velez. Predicting treatment dropout after antidepressant initiation. *Translational Psychiatry*, 10(1):1–8, February 2020. ISSN 2158-3188. doi: 10.1038/s41398-020-0716-y. URL `https://www.nature.com/articles/s41398-020-0716-y`. Number: 1 Publisher: Nature Publishing Group.

28. Maithra Raghu, Katy Blumer, Greg Corrado, Jon Kleinberg, Ziad Obermeyer, and Sendhil Mullainathan. The algorithmic automation problem: Prediction, triage, and human effort. *arXiv preprint arXiv:1903.12220*, 2019.

29. Gesualdo Scutari, Daniel P Palomar, Francisco Facchinei, and Jong-Shi Pang. Monotone games for cognitive radio systems. In *Distributed decision making and control*, pages 83–112. Springer, 2012.

30. Gesualdo Scutari, Francisco Facchinei, and Lorenzo Lampariello. Parallel and distributed methods for constrained nonconvex optimization—part i: Theory. *IEEE Transactions on Signal Processing*, 65(8):1929–1944, 2016.

31. Marc Stone, Thomas Laughren, M. Lisa Jones, Mark Levenson, P. Chris Holland, Alice Hughes, Tarek A. Hammad, Robert Temple, and George Rochester. Risk of suicidality in clinical trials of antidepressants in adults: analysis of proprietary data submitted to US Food and Drug Administration. *BMJ (Clinical research ed.)*, 339:b2880, August 2009. ISSN 1756-1833. doi: 10.1136/bmj.b2880.

32. Geoffrey G. Towell and Jude W. Shavlik. Knowledge-based artificial neural networks. *Artificial Intelligence*, 70(1-2): 119–165, October 1994. ISSN 00043702. doi: 10.1016/0004-3702(94)90105-8. URL `http://linkinghub.elsevier.com/retrieve/pii/0004370294901058`.

33. Son N. Tran and Artur S. d'Avila Garcez. Deep Logic Networks: Inserting and Extracting Knowledge From Deep Belief Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 29(2):246–258, February 2018. ISSN 2162-237X, 2162-2388. doi: 10.1109/TNNLS.2016.2603784. URL http://ieeexplore.ieee.org/document/7738566/.

34. Jiaxuan Wang, Jeeheh Oh, Haozhu Wang, and Jenna Wiens. Learning Credible Models. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2417–2426, July 2018. doi: 10.1145/3219819.3220070. URL http://arxiv.org/abs/1711.03190. arXiv: 1711.03190.

35. Bryan Wilder, Eric Horvitz, and Ece Kamar. Learning to complement humans. *arXiv:2005.00582*, 2020.

36. Yu Wu, Wei Wu, Can Xu, and Zhoujun Li. Knowledge enhanced hybrid neural network for text matching. In *AAAI Conference on Artificial Intelligence*, 2018.

A Derivation of Log-Barrier Method:

We derive the log-barrier algorithm presented in Section 4 first by approximating player 1's subproblem into an unconstrained problem using interior point methods [Boyd and Vandenberghe, 2004, Chapter 11].

$$\text{(player 1):} \quad \min_{w \in W} \quad -t \sum_{n=1}^{N} \rho_w(x_n) - \ln\Big((1+\varepsilon) L^{\gamma}_{\theta^*, w^*}(\mathcal{D}) - L^{\gamma}_{\theta, w}(\mathcal{D})\Big). \tag{9}$$

We transform player 2's objective by taking negative logarithm:

$$\text{(player 2)} \quad \min_{\theta \in \Theta} \quad -\ln\Big((1+\varepsilon) L^{\gamma}_{\theta^*, w^*}(\mathcal{D}) - L^{\gamma}_{\theta, w}(\mathcal{D})\Big) \tag{10}$$

Both player 2's objective in (5) and (10) are equivalent because the operation performed is monotone, and (10) remains convex in θ [Boyd and Vandenberghe, 2004, Equation 3.10]. Even though we enforced that $L^{\gamma}_{\theta, w}(\mathcal{D}) \leq (1+\varepsilon) L^{\gamma}_{\theta^*, w^*}(\mathcal{D})$ in the objective, we know that such solution is non-empty (see Assumption 2).

After transforming both player's objective from (5), we can optimize a Pareto solution that combines both objectives and minimizes all variables. Since both subproblems have common terms, they simplify and yield a similar formulation as the one presented in Equation (6). Therefore, the log-barrier method searches for a Pareto solution of game $\mathcal{G}$ approximately, using interior-point methods to enforce player 1's constraint.

B Assumptions and Proofs

Assumptions and existence of solutions. Game $\mathcal{G}$ models complex relations between variables in both objectives and constraints. We make some simplifying assumptions and provide sufficient guarantees to establish existence of solutions. These assumptions will also permit to develop algorithms that attain such solutions.

Assumption 1. *Sets Θ and W are compact, convex and non-empty.*

Assumption 1 is mild and practical, constraining that θ and w will take finite values, and that the constraint sets of the optimization procedure will be convex. Compactness is also a necessary requirement. Consider for example a binary classification problem where all points are from a single class type, and all points are located in a bounded region of space. A logistic regression classifier would allocate all points on one side of the hyperplane, and then maximize its distance from the points without bound and goes to infinity and a solution would not exist. Compactness prevents this undesirable behavior.

Define $K_\varepsilon(\theta) = \{\, w \in W \mid L^{\gamma}_{\theta, w}(\mathcal{D}) \leq (1+\varepsilon) L^{\gamma}_{\theta^*, w^*}(\mathcal{D}) \,\}$. $K_\varepsilon(\theta)$ is a set over w given some $\theta \in \Theta$.

Assumption 2. *$K_\varepsilon(\theta)$ is non-empty for every $\theta \in \Theta$.*

Assumption 2 is technical, whose purpose is to guarantee that a solution exists (see Proposition 1 below). This assumption can always be fulfilled, since $\Theta = \{\, \theta^* \,\}$ and $W = \{\, w^* \,\}$ satisfy it by construction, although more interesting problems arise when solving over larger sets. In practice, the iterates obtained from an alternating optimization process over the player's subproblems in $\mathcal{G}$ will satisfy Assumption 2 at every iteration, and no subproblem will be empty.

Assumption 3. *Functions $f_\theta(x)$ and $\rho_w(x)$ are log-concave.*

Assumption 3 guarantees that both subproblems of $\mathcal{G}$ are convex and can be solved optimally on each variable. Logistic regression classifiers satisfy the assumption, i.e., $f_\theta(x) = \sigma\big(\theta^T x\big)$ and $\rho_w(x) = \sigma\big(w^T x\big)$, where $\sigma(z) = 1/(1 + \exp(-z))$ is the sigmoid function. We used these functions to illustrate our results in Section 5.

Proposition 1. *Given Assumptions 1 to 3, a Nash Equilibrium of game $\mathcal{G}$ exists.*

Proposition 1 guarantees that $\mathcal{G}$ is well posed and have a solution. The proof is given below.

Proof of Proposition 1. Existence of solution for game $\mathcal{G}$ can be established by [Pang and Fukushima, 2005, Theorem 2]. Convexity of the set $K_\varepsilon(\theta)$ is guaranteed by construction (the cross-entropy loss is a convex function of w for fixed θ). By Assumption 2 the set is feasible and non-empty. By Assumption 1 on sets W and Θ the sets are compact and non-empty. Assumption 3 guarantees each subproblem is convex. Finally, the constraint satisfies the necessary constraint qualifications for dual variables to exist, and a Nash Equilibrium solution must exist.

We state the following result regarding the monotonicity of game $\mathcal{G}$:

Proposition 2 (Monotonicity)**.** *Given Assumptions 1 to 3. Assume that $f_\theta(x)$ and $\rho_w(x)$ are strongly log-concave functions. Then, $\mathcal{G}$ is monotone.*

Proposition 2 is an important intermediate result that allows us to develop Algorithm 1 and prove its convergence. The strong log-concavity requirement may be relaxed, see the proof below.

Proof of Proposition 2 The monotonicty of the game can studied via an equivalent variational inequality (VI) derived from the game subproblems [Scutari et al., 2012, Proposition 4.1]. First, notice that game (5) can be transformed via Lagrangian relaxation of the first player's subproblem into the following three player game:

$$\begin{array}{lll} (player1) & \min\limits_{w\in W} & -\Sigma_{n=1}^{N} -c\ln(\rho_w(x)) + \lambda L^{\gamma}_{\theta,w}(\mathcal{D}) \\ (player2) & \min\limits_{\theta\in\Theta} & (1+\lambda)L^{\gamma}_{\theta,w}(\mathcal{D}) \\ (player3) & \min\limits_{\lambda\geq 0} & -\lambda\big(L^{\gamma}_{\theta,w}(\mathcal{D}) - (1+\varepsilon)L^{\gamma}_{\theta^*,w^*}(\mathcal{D})\big). \end{array} \tag{11}$$

The third suproblem of (11) is derived from the complementary slackness condition of player 1, which has added the constraint multiplied by the dual variable λ. The second player multiplied its objective by positive constant $(1+\lambda)$, where $\lambda \geq 0$, that does not modify the problem. Neither does c. Clearly, game (11) is equivalent to (5).

Game (11) has three unconstrained subproblems, so we can now construct an equivalent VI. Following [Scutari et al., 2012, Proposition 4.1], we obtain VI($\mathbf{F}, \mathcal{Q}$) with $\mathbf{F} : \mathbb{R}^{p+q+1} \to \mathbb{R}^{p+q+1}$, $\mathbf{F} = ((\nabla_w f_1)^T, (\nabla_\theta f_2)^T, (\nabla_\lambda f_3)^T)^T$, f_i correspond to the objective function of players $i \in \{1,2,3\}$ in Equation (11); and $\mathcal{Q} = W \times \Theta \times \mathbb{R}_+$.

Monotonicity can be ascertained by establishing that the symmetric part of the Jacobian of $\mathbf{F}$ is positive semidefinite [Scutari et al., 2012, Section 4.2.3]. Indeed,

$$J\mathbf{F} = \begin{pmatrix} cr_{ww} & \lambda\frac{\partial^2}{\partial w\partial\theta}L^{\gamma}_{\theta,w}(\mathcal{D}) & \frac{\partial}{\partial w}L^{\gamma}_{\theta,w}(\mathcal{D}) \\ (1+\lambda)\frac{\partial^2}{\partial\theta\partial w}L^{\gamma}_{\theta,w}(\mathcal{D}) & (1+\lambda)r_{\theta\theta} & \frac{\partial}{\partial\theta}L^{\gamma}_{\theta,w}(\mathcal{D}) \\ -\frac{\partial}{\partial w}L^{\gamma}_{\theta,w}(\mathcal{D}) & -\frac{\partial}{\partial\theta}L^{\gamma}_{\theta,w}(\mathcal{D}) & 0 \end{pmatrix}, \tag{12}$$

where $r_{ww} \succ 0$ and $r_{ww} \succ 0$ because of the strong convexity assumption on the objectives. Notice that the third row and column are antisymmetric, and those terms cancel when studying positive semidefiniteness of $J\mathbf{F}$. Therefore, we only need to show that the upper left corner is positive semi-definite.

We compute the Schur complement of (symmetric) $J\mathbf{F}$ [Boyd and Vandenberghe, 2004, Section A.5.5], i.e.,

$$S = r_{\theta\theta} - \frac{1+2\lambda}{2c}\Big(\frac{\partial^2}{\partial w\partial\theta}L^{\gamma}_{\theta,w}(\mathcal{D})\Big)^T r_{ww}^{-1}\Big(\frac{\partial^2}{\partial w\partial\theta}L^{\gamma}_{\theta,w}(\mathcal{D})\Big) \tag{13}$$

$J\mathbf{F}$ is positive semi-definite if S is positive definite. The maximum value of λ is bounded, because W and Θ are compact. Therefore, there exists $c > 0$ that necessarily makes $S \succeq 0$, proving that $J\mathbf{F}$ is positive semi-definite and that game (11) is monotone. Because (11) and (5) are equivalent, both games are monotone.

The strong convexity assumption may be relaxed by showing $S \succeq 0$ with singular Hessians r_{ww} or $r_{\theta\theta}$, through a more convoluted but relaxed requirements [Boyd and Vandenberghe, 2004, Section A.5.5].

Proof of Theorem 1 By Proposition 2 game (5) is monotone. Using [Facchinei and Pang, 2007, Theorem 12.1.2] the projected gradient descent algorithm converges.