Learning Representations by Humans, for Humans

Sophie Hilgard* School of Engineering and Applied Sciences Harvard University

Mahzarin R. Banaji Department of Psychology Harvard University School of Engineering and Applied Sciences Harvard University

Nir Rosenfeld*

Jack Cao Department of Psychology Harvard University David C. Parkes School of Engineering and Applied Sciences Harvard University

Abstract

We propose a new, complementary approach to interpretability, in which machines are not considered as experts whose role it is to suggest what should be done and why, but rather as advisers. The objective of these models is to communicate to a human decision-maker not *what* to decide but *how* to decide. In this way, we propose that machine learning pipelines will be more readily adopted, since they allow a decision-maker to retain agency. Specifically, we develop a framework for learning representations by humans, for humans, in which we learn representations of inputs ('advice') that are effective for human decision-making. Representationgenerating models are trained with humans-in-the-loop, implicitly incorporating the human decision-making model. We show that optimizing for human decisionmaking rather than accuracy is effective in promoting good decisions in various classification tasks while inherently maintaining a sense of interpretability.

1 Introduction

Across many important domains, machine learning algorithms have become unparalleled in their predictive capabilities. The accuracy and consistency of these algorithms has made them highly appealing as tools for supporting human decision-making [22, 46]. However, these criteria are far from comprehensive [49, 5]. Our continued reliance on humans as the final arbiters of these decisions suggests an awareness that incorporating higher-level concepts, such as risk aversion, safety, or justification, requires the exercise of human reasoning, planning, and judgment.

The field of interpretable machine learning has developed as one answer to these issues. A common view of interpretable ML is that it provides explanations [41], thereby allowing integration into the human reasoning process, and verification as to whether or not auxiliary criteria are being met. Under this framework, the algorithm is an expert whose task is to suggest *what* should be done, and, from its own perspective, *why*. The human role is reduced to that of quality control: should the algorithm's work be accepted or rejected? This role of 'computer as expert' undermines a decision-maker's sense of agency and generates information that is difficult to integrate with existing intuition. Hence, users may be reluctant to accept algorithmic suggestions or even inclined to go against them [8, 68], especially after seeing the algorithm make errors, which can lead to a degradation in performance over time [20, 15, 43, 47]. In any system in which humans make the final decisions, even highly-accurate machine outputs are only useful if and when humans make appropriate use of them; c.f. the use of risk assessment tools in the context of sentencing [60].

^{*}Equal contribution, alphabetical order.



Figure 1: Examples of visualized advice for various inputs: word highlighting for text data, customized plots for embedded data, and computerized avatars for structured data. Instead of explaining algorithmic predictions, we learn representations that directly aid in human decision-making.

Fortunately, advice that conveys *how* to decide (rather than what) can often be of great value [13]. Advice of this form can be designed to *augment* the capabilities of human decision makers, rather than replace them, which many see as a more socially-optimal role for AI [40, 21, 39, 30]. This can be achieved, for example, by highlighting certain aspects of the problem, providing additional information, presenting tradeoffs in risks and returns, or outlining possible courses of action. There is ample empirical evidence suggesting that informative advice can, by acknowledging the central role decision makers play, both enhance performance *and* retain agency [31, 34].

Motivated by the above, we advocate for a broader perspective on how machine learning can be used to support decision-making. Our work builds on a well-known observation in the social sciences, which is that the performance of humans on decision tasks depends on how problems are presented or framed [61, 12, 24, 11, 32, 9] To leverage this idea, we shift the algorithmic focus from learning to predict to learning to represent, and seek representations of inputs ('advice') that will lead to good decisions and thus good outcomes *when presented to a human decision maker*. Our framework is designed to use machine learning in a way that preserves autonomy and agency, and in this way builds trust— crucial aspects of decision-making that are easy to overlook [3, 4, 16, 43].

To successfully reframe difficult problems, we harness the main engine driving deep learning— the ability to learn useful representations. Just as deep neural networks learn representations under which classifiers predict well, we learn representations under which human decision makers perform well. Our model includes three main components: a "truncated" neural network that maps inputs into vector representations, a visualization module that maps vector representations into visual representations, and a human decision maker. Our main innovation is a human-in-the-loop training procedure that seeks to directly optimize human decision outcomes, thus promoting both accuracy *and* agency.

We demonstrate the approach on three experimental tasks, represented in Figure 1, that cover different types of decisions and different forms of computational advice, and in problems with increasing complexity. Both training and evaluation are done with the aid of real human subjects, which we argue is essential for learning credible human-supportive tools. Our results show that we can iteratively learn representations that lead to high human accuracy while not explicitly presenting a recommended action, providing users with means to reason about decisions. Together, these results demonstrate how deep learning can serve as an instrumental tool for human intelligence augmentation [40, 21, 39, 30].

1.1 Related Work

Interpretability as decision support. There are several ways in which interpretability can be used to support decision-making. In general, interpretability can help in evaluating criteria that are important for decisions but hard to quantify, fairness or safety for example, and hence hard to optimize [17]. Many methods do this by producing simplified [1, 37] or augmented [54, 59, 38] versions of the input that aids users in understanding if the data is used in ways that align with their goals or not. While some methods exist for systematically iterating over models [56, 36], these give no guarantees as to whether models actually improve with respect to user criteria. Virtually all works in interpretability focus on predictive algorithms. Our work differs in that the focus is directed at the human-decision maker, directly optimizing for better decisions by learning useful human-centric representations.

Incorporating human feedback. Our use of human-in-the-loop methods is reminiscent of work in active learning, in that humans supply labels to reduce machine uncertainty [58], and in preference-based reinforcement learning in that we implicitly encode human preferences in our evaluation [67].

However, in our work, learning a model that approximates human policy decisions is not the end goal but rather a tool to improve decisions by approximating 'decision gradients'. While this can be viewed as a form of black-box gradient estimation [27], current methods assume either inexpensive queries, noise-free gradients, or both, making them inadequate for modeling human responses.

Expertise, trust, and agency. Recent studies have shown that links between trust, accuracy, and explainability are quite nuanced [69, 52, 25]. Users fail to consistently increase trust when model accuracy is superior to human accuracy and when models are more interpretable. Expertise has been identified as a potentially confounding factor [43], when human experts wrongly believe they are better than machines, or when they cannot incorporate domain-specific knowledge within the data-driven model estimate. Agency has also been shown to affect the rate at which people accept model predictions [16], supporting the hypothesis that active participation increases satisfaction, and that users value the ability to intervene when they perceive the model as incorrect.

2 Learning Decision-Optimal Representations

2.1 Preliminaries

We consider a setting where users are given *instances* $x \in \mathcal{X}$ sampled from some *distribution* D, for which they must decide on an *action* $a \in \mathcal{A}$. For example, if x are details of a loan application, then users can choose $a \in \{\text{approve, deny}\}$. We denote by h the *human mapping* from arbitrary inputs to decisions or actions (we use these terms interchangeably). We assume that users are seeking to choose a = h(x) to minimize an incurred loss $\ell(x, a)$, and our goal is to aid them in this task. To achieve this, we can present users with machine-generated *advice* $\gamma(x)$, which we think of as a human-centric 'representation' of the input. To encourage better outcomes, we seek to learn the representation γ under which human decisions $a = h(\gamma(x))$ entail low expected loss $\mathcal{L}(\gamma) = \mathbb{E}_D[\ell(x, h(\gamma(x)))]$.

We will focus on tasks where actions are directly evaluated against some ground truth $y \in \mathcal{Y}$ associated with x and given at train time, and so the loss is of the form $\ell(y, h(\gamma(x)))$. In this way, we cover a large class of important decision problems called *prediction policy problems*, where the difficulty in decision-making is governed by a predictive component [35]. For example, the loss from making a loan depends on whether or not a person will return a given loan, and thus on being able to make this conditional prediction with good accuracy. This setting is simpler to evaluate empirically, and allows for a natural comparison to interpretable predictive approaches where $\gamma(x)$ includes a machine prediction \tilde{y} and some form of an explanation. In our experiments we have $\mathcal{Y} = \{1, \ldots, C\}$, and denote by Δ_C the C-dimensional simplex (allowing probabilistic machine prediction $\tilde{y} \in \Delta_C$).

Given a train set $S = \{(x_i, y_i)\}_{i=1}^m$, we will be interested in minimizing the *empirical loss*:

$$\min_{\gamma \in \Gamma} \sum_{i=1}^{m} \ell(y_i, a_i) + \lambda R(\gamma), \qquad a_i = h(\gamma(x_i))$$
(1)

where Γ is the advice class, R is a regularization term that can be task-specific and data-dependent, and λ is the regularization parameter. The main difficulty in solving Eq. (1) is that $\{a\}_{i=1}^{m}$ are *actual human decisions* that depend on the optimized function γ via an unknown decision mechanism h. We first describe our choice of Γ and propose an appropriate regularization R, and then present our method for solving Eq. (1).

2.2 Learning human-facing representations

Deep neural networks can be conceptualized as powerful tools for learning representations under which simple predictors (i.e., linear) perform well [6]. By analogy, we leverage neural architectures for learning representations under which humans perform well. Consider a multi-layered neural network $\mathcal{N}(x)$. Splitting the network at some layer partitions it into a parameterized representation mapping $\phi_{\theta} : \mathbb{R}^d \to \mathbb{R}^k$ and a predictor $f : \mathbb{R}^k \to \Delta_C$ such that $\mathcal{N}(x) = f(\phi_{\theta}(x))$. If we assume for simplicity that f is fixed, then learning is focused on ϕ . The challenge is that optimizing θ may improve the predictive performance of the algorithm, but may not facilitate good human decisionmaking. To support human decision makers, our key proposal is to remove f and instead plug in the human decision function h, therefore leveraging the optimization of θ to directly improve human performance. We refer to this optimization framework as "**M**o**M**", *Man Composed with Machine*, pronounced "mom" and illustrated in Fig. 2 (left).

$$x \qquad \phi \qquad z \rightarrow \overbrace{\rho}^{(A)} \bigwedge_{h}^{(A)} \rightarrow a \qquad (A) \qquad$$

Figure 2: Left: The M \circ M framework. The network learns a mapping ϕ from inputs x to representations z, such that when z is visualized through the visualization component ρ , the representation elicits good human decisions a. **Right**: The learning process. Users are queried for decisions on the current representations (A). These decisions are used to train a proxy network \hat{h} (B), that is then used to re-train representations (C). This process is repeated until convergence.

We also need to be precise about the way a human would perceive the output of ϕ . The outputs of ϕ are vectors $z = \phi_{\theta}(x) \in \mathbb{R}^k$, and not likely to be helpful as human input. To make representations accessible to human users, we add a *visualization* component $\rho : \mathbb{R}^k \to \mathcal{V}$, mapping vector representations into meaningful visual representations $v = \rho(z)$ in some class of visual objects \mathcal{V} (e.g. scatter-plots, word lists, avatars). Choosing a proper visualization is crucial to the success of our approach, and should be chosen with care to utilize human cognition (and this is in itself a research question). Combined, these mappings provide what we mean by the 'algorithmic advice':

$$\gamma(x) = \rho(\phi_{\theta}(x)) \tag{2}$$

In the remainder of the paper, we assume that the visualization component ρ is fixed, and focus on optimizing the advice by learning the mapping ϕ_{θ} . It will be convenient to fold ρ into h, using the notation $h^{(\rho)}(z) = h(\rho(z))$. Eq. (1) can now be rewritten as:

$$\min_{\theta \in \Theta} \sum_{i=1}^{m} \ell(y_i, a_i) + \lambda R(\theta), \qquad a_i = h^{(\rho)}(\phi_\theta(x_i))$$
(3)

By solving Eq. (3), we hope to learn a representation of inputs such that, when visualized, promote good decisions. In the remainder of the paper we will simply use h to mean $h^{(\rho)}(z)$.

2.3 Optimization

The difficulty in optimizing Eq. (3) is that gradients of θ must pass through h. But these are actual human decisions! To handle this, we propose to replace $h^{(\rho)}$ with a differentiable proxy $\hat{h}_n : \mathbb{R}^k \to \mathcal{Y}$ parameterized by $\eta \in H$ (we refer to this proxy as "h-hat"). A naïve approach would be to train \hat{h} to mimic how h operates on inputs z, and use it in Eq. (3). This, however, introduces two difficulties. First, it is not clear what data should be used to fit h. To guarantee good generalization, h should be trained on the distribution of z induced by the learned $\phi_{\theta}(x)$, but the final choice of θ depends on \hat{h} itself. Second, precisely modeling h can be highly unrealistic (i.e., due to human prior knowledge, external information, or unknown considerations).

То circumvent these issues, we propose a human-in-the-loop training procedure alternating between fitting \hat{h}_{η} for a fixed θ and training ϕ_{θ} for a fixed \hat{h}_{η} Fig. 2 (right) illustrates this process, and pseudocode is given in Algorithm 1. The process begins by generating representations $z_i = \phi_{\theta_0}(x_i)$ for $n \leq m$ random training inputs x_i with an initial θ_0 , and obtaining decisions a_i for each z_i generated in this way by querying human participants. Next, we take these representation-5: 6: decision pairs and create an auxiliary sample 7: set $S' = \{(z_i, a_i)\}_{i=1}^n$, which we use to fit the 8: human model \hat{h}_{η} by optimizing η . Fixing η , we then train ϕ_{θ} by optimizing θ on the empirical

	ica o	una	uuning	$\varphi \theta$	101	u	intea	1011.
Alg	gorithn	1 Al	ternating	optin	nizati	on	algorit	hm
1:	Initializ	$e \theta = \theta$	θ_0					
2:	repeat							
3:	$x_1, .$	\ldots, x_n	$\sim \mathcal{S}$	$\triangleright S$	ample	n ti	rain exar	nples
4:	$z_i \leftarrow$	$-\phi_{\theta}(x)$	$\forall i \in [n]$] ⊳0	Genera	te re	presenta	tions
5.	~ /	hlal			0	. b		

 $a_i \leftarrow h(\rho(z_i)) \ \forall i \in [n] \ \triangleright$ Query human decisions $\mathcal{S}' = \{(z_i, a_i)\}_{i=1}^n$ $\eta \leftarrow \operatorname{argmin}_{\eta} \mathbb{E}_{\mathcal{S}'}[\ell(a, \hat{h}_{\eta}(z))]$ \triangleright Train \hat{h} $\theta \leftarrow \operatorname{argmin}_{\theta'} \mathbb{E}_{\mathcal{S}}[\ell(y, \hat{h}_{\eta}(\phi_{\theta'}(x)))]$ \triangleright Train ϕ 9: until convergence



Figure 3: Visualization of 2D projection task. Points in their original 3D representation give little visual indication of class (X or O). The initial 2D projection (round 1), set to the final layer representation of a fully accurate machine-only model, is similarly unintelligible to humans. However, as training progresses, feedback from human decisions improves the learned 2D projection until the class becomes visually apparent (round 5), achieving 100% *human* accuracy.

loss of the original sample set S. We repeat this alternating process until re-training \hat{h} does not improve results. In our experiments, both ϕ and \hat{h} are implemented through neural networks. In the Appendix, we discuss practical issues regarding initialization, convergence, early stopping, and working with human inputs.

The initial training of \hat{h} makes it match h as best as possible on the distribution of z induced by θ_0 . In the next step, however, optimizing θ causes the distribution of z to drift. As a result, forward passes push out-of-distribution samples into \hat{h} , and \hat{h} may no longer be representative of h (and with no indication of failure). Fortunately, this discrepancy is corrected at the next iteration, when \hat{h} is re-trained on fresh human-annotated samples drawn from the distribution induced by the new parameters θ . In this sense, our training procedure literally includes humans-in-the-loop.

In order for performance to improve, it suffices that \hat{h} induces gradients of the loss that approximate those of h. This is a weaker condition than requiring \hat{h} to match h exactly. In the Appendix we show how even simple \hat{h} models that do not fit h well are still effective in the overall training process.

3 Experiments

We conduct a series of experiments on data-based decision-making tasks of increasing complexity. Each task uses the general algorithmic framework presented with a different, task-appropriate class of advice representations. Each experiment is also successively more sophisticated in the extent of human experimentation that is entailed. The appendix includes further details on each experiment.

3.1 Decision-compatible 2D projections

High-dimensional data is notoriously difficult for humans to handle. One way to make it accessible is to project points down to a low dimension where they can be visualized (e.g., with plots). But neither standard dimensionality reduction methods nor the representation layer of neural networks are designed to produce visualizations that support human decision-making. PCA, for example, optimizes a statistical criterion that is agnostic to how humans visually interpret its output.

Our M \circ M framework suggests to learn an embedding that directly supports good decisions. We demonstrate this in a simple setting where the goal of users is to classify *d*-dimensional point clouds, where d > 2. Let V be a linear 2D subspace of \mathbb{R}^d . Each point cloud is constructed such that, when orthogonally projected onto V, it forms one of two visual shapes— an 'X' or an 'O' —that determine its label. All other orthogonal directions contain similarly scaled random noise. We use M \circ M to train an orthogonal 2D projection (ϕ) that produces visual scatter-plots (ρ). Here, ϕ is a 3x3 linear model augmented with an orthogonality penalty $\phi^T \phi - \mathbb{I}$, and \hat{h} is a small single-layer 3x3 convolutional network that takes as inputs a soft (differentiable) 6x6 histogram over the 2D projections.

In each task instance, users are presented with a 2D visualization of a point cloud and must determine its shape (i.e., label). Our goal is to learn a projection under which point clouds can be classified by humans accurately, immediately, and effortlessly. Initially, this is difficult, but as training progresses, user performance feedback gradually "rotates" the projection, revealing class shapes (see Fig. 3). Importantly, *users are never given machine-generated predictions*. Rather, progress is driven solely

Great one.... Is it for the atmosphere of the movie? If yes, it is for sure the utmost human one... I can imagine most of the people who hate to watch slow movies suffering during the movie.... Great one.... Is it for the atmosphere of the movie? If yes, it is for sure the utmost human one... I can imagine most of the people who hate to watch <u>slow</u> movies suffering during the movie... Believable sets and costumes, good sword fighting and a beautiful female warrior, Mila make this an enjoyable watch.

M∘M

LIME

Figure 4: Examples of word sets selected by MoM and by LIME. Color indicates machine-perceived sentiment (green for positive, red for negative). The explanation generated by LIME includes many words with no intuitive sentiment (e.g., 'movie'). While LIME can be useful for identifying words that may not be desirable as predictive features (e.g., 'female'), MoM works in a different way, by directly adjusting itself to how humans make decisions.

by the performance of users on algorithmically "reframed" problem instances (i.e., projections), achieving 100% *human* accuracy in only 5 training rounds with at most 20 queries each.

3.2 Decision-compatible feature selection

In some applications, inputs are composed of many discrete elements, such as words or sentences in a document, or objects in an image. A useful form of advice in this setting is to 'summarize' inputs by highlighting a small subset of important elements or features. Consider, for example, a task of determining text sentiment, where the summary would be relevant words. The MoM framework suggests that models should be trained to choose summaries (representations) that are effective in helping humans make good decisions.

In this section, we consider the task of determining text sentiment using the IMDB Movie Review Dataset [44]. We compare M \circ M with the LIME [54] method, which learns a *post hoc* summarization to best explain the predictions of black-box predictive models. LIME chooses a subset of words for an input x by training a simpler model to match the black-box prediction in the neighborhood of x. The summarization selected by LIME may therefore give insight to the model's internal workings, but seems only likely to build trust to the extent that the "explanation" matches human intuition. And when it does not, the advice offered by LIME is unlikely to help users to form their own opinion.

In our experiment, we implement a subset-selection mechanism in ϕ as a *Pointer Network* [65], a neural architecture that is useful in learning mappings from sets to subsets. In particular, we model ϕ as a pair of "dueling" Pointer Network advisers, one for 'positive sentiment' and one for 'negative sentiment'. The learning objective is designed to encourage each adviser to give useful advice by competing for the user's attention, with the idea of giving the user a balanced list of "good reasons" for choosing the each of the possible alternatives (see Appendix for details). The visualizer ρ simply presents the chosen words to the user, and the goal of users is to determine the sentiment of the original text from its summary. In this experiment we trained using simulated human responses via queries to a word sentiment lexicon, which proved to be cost effective, but as in all other experiments, evaluation was done with real humans. For LIME we use a random forest black-box predictor and a linear 'explainable' model, as in the original LIME paper.

Results. The black-box random forest classifier is fairly accurate, achieving 78% accuracy on the test set when trained and evaluated on full text reviews. However, when LIME summaries composed of the top and bottom three words with highest coefficients were given as input to humans, their performance was only 65%. Meanwhile, when given summaries generated by $M \circ M$, human performance reached 76%, which almost matches machine performance *but using summaries alone*. Examples of summaries generated by $M \circ M$ and LIME are given in Figure 4.

MoM creates summaries that are more diverse and nuanced; LIME uses half the number of overall unique words, five of which account for 20% of all word appearances. Words chosen by LIME do not necessarily convey any sentiment— for instance, the word 'movie' is LIME's most frequent indication of negative sentiment (7.4%), and the word 'female' is chosen to convey negative sentiment. This artifact may be helpful in revealing spurious correlations used by the black-box algorithm to achieve high accuracy, but is uninformative as input as input to a human decision maker.



Figure 5: **Right**: different learned avatars conveying algorithmic advice through facial expressions (see Appendix for more examples). **Left**: Human accuracy in the algorithmic advice condition ('avatar advice') consistently increases over rounds. Performance quickly surpasses the 'data only' condition, and steadily approaches performance of users observing algorithmic predictions ('predictive advice'), which in itself is lower than machine-only performance. When faces are shuffled within predicted labels of \hat{h} , accuracy falls, suggesting that faces convey important multi-variate information.

3.3 Decision-compatible algorithmic avatars

Our main experiment focuses on the problem of approving loans using the *Lending Club* dataset.² Given details of a loan application, the task of a decision maker is to decide whether to approve the loan or not. This can be done by first predicting the conditional outcome of giving a loan, and then determining an appropriate course of action. Predicting accurately is important but not sufficient, as in reality, decision makers must also justify their decisions. Our goal in this task is twofold: aid decision makers in making good decisions, and provide them with means to reason about their choices.

The standard algorithmic approach to assisting users would be to give them predictions or risk scores, perhaps along with an 'explanation'. This, however, reduces the rich data about an application to a single number. Instead, we propose to give a decision maker 'just right' high-dimensional advice—compressed enough to be managable, yet rich enough to preserve multi-variate aspects of the input —crucial for retaining users' ability to reason about their decisions [51].

For this task, we augment inputs with algorithmic advice in the form of an 'avatar' framed as conveying through its facial expression information that is relevant to the conditional outcome of giving a loan. Facial expressions have been used successfully to represent and augment multivariate data [57, 63, 10], but by manually mapping features to facial components (whereas we learn this mapping). We use realistic-looking faces, with the goal of harnessing innate human cognitive capabilities— immediate, effortless, and fairly consistent processing of facial signals [26, 33, 62, 23] —to successfully convey complex high-dimensional information (see Fig. 5 and Appendix for details).

Setup. We split the data 80:20 into a train set and a held-out test set, which is only used for the final evaluation. To properly assess human decisions we include only loans for which we know the resolution in the data (either repay in full or default), and accordingly set $\ell(y, a) = \mathbb{1}_{\{y=a\}}$ where $y \in \{0, 1\}$ indicates the ground truth (1 = repay, 0 = default), and $a \in \{0, 1\}$ indicates the decision (1 = approve, 0 = deny). Following MoM we use the train set to optimize the representation ϕ , and at each round, use the outputs of ϕ (parametrizations of faces) to fit \hat{h} using real human decisions (i.e., approve or deny) gathered from mTurk.³ We set ϕ and \hat{h} to be small fully connected networks with 1 25-hidden unit layer and 2 20-hidden unit layers, respectively. The visualizing unit ρ turns the vectorized outputs of ϕ into avatars by morphing seven 'facial dimensions' from various sources [18, 62] using the *Webmorph software* [14]. To prevent mode collapse, wherein faces "binarize" to two prototypical exemplars, we add a reconstruction regularization term $R(x) = ||x - \psi(\phi(x))||_2^2$ to the objective, where ψ is a decoder implemented by an additional neural network. In the Appendix we give a detailed description of the learning setup, training procedure, mTurk experimental environment, and the unique challenges encountered when training with turkers in the loop.

Evaluation. We are interested in evaluating both predictive performance and the capacity of users for downstream reasoning. We compare between the following conditions: (1) no advice, (2) predictive advice: $\gamma(x) = \tilde{y} \in [0, 1]$ is a predictive probability by a pre-trained predictive model $\mathcal{N}(x)$, (3)

²https://www.kaggle.com/wendykan/lending-club-loan-data

³All experiments were approved by the Harvard University IRB.

representational advice: $\gamma(x) = v$, where $v = \rho(\phi(x))$ is an avatar, and (4) a 'shuffled' condition which we will soon describe. In all conditions, this advice is given to users in addition to the five most informative features of each example (given by the regularization path of a LASSO model). Since users in the experiment are non-experts, and because there is no clear incentive for them *not* to follow predictive advice, we expect the predictive advice condition to give an upper bound on human performance in the experiment; this artifact of the experimental environment should not necessarily hold in reality. We benchmark results with the accuracy of \mathcal{N} (having architecture equal to $\hat{h} \circ \phi$).

Results. Fig. 5 shows the training process and resulting test accuracies⁴ (the data is fairly balanced so chance ≈ 0.5). Initially, the learned representation ϕ produces arbitrary avatars, and performance in the avatar condition is lower than in the no advice condition. This indicates that users take into account the (initially uninformative) algorithmic advice. As learning progresses, user feedback accumulates, and accuracy steadily increases. After six training rounds, accuracy in the avatar condition reaches 94% of the accuracy in the predictive advice condition. Interestingly, performance in the predictive advice condition does not reach the machine accuracy benchmark, showing that even experimental subjects do not always follow predictive advice. This resonates well with our arguments from Sec. 1.

In addition to accuracy, our goal is to allow users to reason about their decisions. This is made possible by the added reconstruction penalty R, designed to facilitate arguments based on *analogical reasoning*: "x will likely be repaid because x is similar to x', and x' was repaid" [42, 29]. Reconstruction serves two purposes. First, it ensures that reasoning in 'avatar-space' is anchored to the similarity structure in input space, therefore encouraging sound inference, as well as promoting fairness through similar treatment of similar people [70]. Second, reconstruction ensures the high dimensionality of the avatar advice representation, conveying rich information. To demonstrate the importance of using high-dimensional advice, we add a condition where avatars are "shuffled" within predicted classes according to \hat{h} (i.e., examples with $\hat{y} = 0$ and with $\hat{y} = 1$ are shuffled separately). Results show a drop in accuracy, confirming that avatars support decision-making by conveying more than unidimensional predictive information. Clearly, this cannot be said of scalar predictive advice, and in the Appendix we show how in this condition reasoning becomes impractical.

In regard to the gap between the avatar and predictive advice conditions, note that (1) R is a penalty term, and introduces a tradeoff between accuracy and reasoning capacity, and (2) users on mTurk have nothing at stake and are more likely to follow predictive advice where professionals would not.

4 Discussion

Our paper presents a novel learning framework for supporting human decision-making. Rather than viewing algorithms as omniscient experts asked to explain their conclusions, we position algorithms as advisors whose goal is to help humans make better decisions while retaining agency. Our framework leverages the power of representation learning to find ways to provide advice promoting good decisions. By tapping into innate cognitive human strengths, learned representations can aid decision-making by prioritizing information, highlighting alternatives, and correcting biases.

The broader MoM framework is motivated by the many professional settings, such as health, education, justice, and business, in which people make data-dependent decisions. We also believe it applies to everyday decisions of a personal, social, or financial nature. Without access to professional decision makers, a challenge we have faced is that we've needed to limit our experimental focus to decision tasks that are governed by a prediction problem. But the framework itself is not limited to these tasks, and we hope to stimulate further discussion and motivate future research initiatives.

The idea of seeking to optimize for human decisions should not be considered lightly. In our work, the learning objective was designed to align with and support the goals of users. Ideally, by including humans directly in the optimization pipeline, we can augment human intelligence as well as facilitate autonomy, agency, and trust. It is our belief that a responsible and transparent deployment of models with "h-hat-like" components should encourage environments in which humans are aware of what information they provide about their thought processes. Unfortunately, this may not always be the case, and ethical, legal, and societal aspects of systems that are optimized to promote particular kinds of human decisions must be subject to scrutiny by both researchers and practitioners. Decision support methods can also be applied in a biased way to induce persuasion [28], and strategies for

⁴Results are statistically significant under a one-way ANOVA test, F(3, 97) = 9.8, p < 1e - 5.

effecting influence that are learned in one realm may be transferable to others [19]. Of course, these issues of algorithmic influence are not specific to our framework, consider news ranking, social content promotion, product recommendation, and targeted advertising, for example.

Looking forward, we think there is good reason to be optimistic about the future of algorithmic decision support. Systems designed specifically to provide users with the information and framing they need to make good decisions can seek to harness the strengths of both computer pattern recognition and human judgment and information synthesis. Through this, we can hope that the combination of man and machine can do better than either one by themselves. The ideas presented in this paper serve as a step toward this goal.

References

- [1] Elaine Angelino, Nicholas Larus-Stone, Daniel Alabi, Margo Seltzer, and Cynthia Rudin. Learning certifiably optimal rule lists. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 35–44. ACM, 2017.
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [3] Albert Bandura. Human agency in social cognitive theory. *American psychologist*, 44(9):1175, 1989.
- [4] Albert Bandura. Self-efficacy. The Corsini encyclopedia of psychology, pages 1–3, 2010.
- [5] Chelsea Barabas, Karthik Dinakar, Joichi Ito, Madars Virza, and Jonathan Zittrain. Interventions over predictions: Reframing the ethical debate for actuarial risk assessment. *arXiv preprint arXiv:1712.08238*, 2017.
- [6] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798– 1828, 2013.
- [7] Andreea Bobu, Eric Tzeng, Judy Hoffman, and Trevor Darrell. Adapting to continuously shifting domains. 2018.
- [8] Jack W Brehm. A theory of psychological reactance. 1966.
- [9] Jeffrey R Brown, Jeffrey R Kling, Sendhil Mullainathan, and Marian V Wrobel. Framing lifetime income. Technical report, National Bureau of Economic Research, 2013.
- [10] Lawrence A Bruckner. On chernoff faces. In *Graphical representation of multivariate data*, pages 93–121. Elsevier, 1978.
- [11] Jack Cao, Max Kleiman-Weiner, and Mahzarin R Banaji. Statistically inaccurate and morally unfair judgements via base rate intrusion. *Nature Human Behaviour*, 1(10):738, 2017.
- [12] Leda Cosmides and John Tooby. Cognitive adaptations for social exchange. *The adapted mind: Evolutionary psychology and the generation of culture*, 163:163–228, 1992.
- [13] Reeshad S Dalal and Silvia Bonaccio. What types of advice do decision-makers prefer? *Organizational Behavior and Human Decision Processes*, 112(1):11–23, 2010.
- [14] LM DeBruine and BP Tiddeman. Webmorph, 2016.
- [15] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1):114, 2015.
- [16] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, 64(3):1155–1170, 2016.
- [17] Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608, 2017.
- [18] Shichuan Du, Yong Tao, and Aleix M Martinez. Compound facial expressions of emotion. *Proceedings of the National Academy of Sciences*, 111(15):E1454–E1462, 2014.

- [19] Dean Eckles, Doug Wightman, Claire Carlson, Attapol Thamrongrattanarit, Marcello Bastea-Forte, and BJ Fogg. Social responses in mobile messaging: influence strategies, self-disclosure, and source orientation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1651–1654. ACM, 2009.
- [20] Avshalom Elmalech, David Sarne, Avi Rosenfeld, and Eden Shalom Erez. When suboptimal rules. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [21] Douglas C Engelbart. Augmenting human intellect: A conceptual framework. *Menlo Park, CA*, 1962.
- [22] 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, 2017.
- [23] Jonathan B Freeman and Kerri L Johnson. More than meets the eye: Split-second social perception. *Trends in cognitive sciences*, 20(5):362–374, 2016.
- [24] Gerd Gigerenzer and Ulrich Hoffrage. How to improve bayesian reasoning without instruction: frequency formats. *Psychological review*, 102(4):684, 1995.
- [25] Ben Green and Yiling Chen. Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 90–99. ACM, 2019.
- [26] Carroll E Izard. Innate and universal facial expressions: evidence from developmental and cross-cultural research. 1994.
- [27] Alon Jacovi, Guy Hadash, Einat Kermany, Boaz Carmeli, Ofer Lavi, George Kour, and Jonathan Berant. Neural network gradient-based learning of black-box function interfaces. arXiv preprint arXiv:1901.03995, 2019.
- [28] Anthony Jameson, Bettina Berendt, Silvia Gabrielli, Federica Cena, Cristina Gena, Fabiana Vernero, Katharina Reinecke, et al. Choice architecture for human-computer interaction. *Foundations and Trends* (R) in Human–Computer Interaction, 7(1–2):1–235, 2014.
- [29] Phillip N Johnson-Laird and Bruno G Bara. Syllogistic inference. Cognition, 16(1):1–61, 1984.
- [30] Michael Jordan. Artificial intelligence the revolution hasn't happened yet. Medium, Apr 2018.
- [31] Daniel Kahneman, Andrew M Rosenfield, Linnea Gandhi, and Tom Blaser. Noise: How to overcome the high, hidden cost of inconsistent decision making. *Harvard business review*, 94(10):38–46, 2016.
- [32] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*, pages 99–127. World Scientific, 2013.
- [33] Nancy Kanwisher, Josh McDermott, and Marvin M Chun. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of neuroscience*, 17(11):4302– 4311, 1997.
- [34] Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. Human decisions and machine predictions. *The quarterly journal of economics*, 133(1):237–293, 2017.
- [35] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. Prediction policy problems. *American Economic Review*, 105(5):491–95, 2015.
- [36] Isaac Lage, Andrew Ross, Samuel J Gershman, Been Kim, and Finale Doshi-Velez. Human-inthe-loop interpretability prior. In Advances in Neural Information Processing Systems, pages 10159–10168, 2018.
- [37] Himabindu Lakkaraju, Stephen H Bach, and Jure Leskovec. Interpretable decision sets: A joint framework for description and prediction. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1675–1684. ACM, 2016.
- [38] Tao Lei, Regina Barzilay, and Tommi Jaakkola. Rationalizing neural predictions. *arXiv preprint arXiv:1606.04155*, 2016.
- [39] Fei-Fei Li. How to make a.i. that's good for people. The New York Times, Mar 2018.

- [40] Joseph Carl Robnett Licklider. Man-computer symbiosis. *IRE transactions on human factors in electronics*, (1):4–11, 1960.
- [41] Zachary C Lipton. The mythos of model interpretability. *arXiv preprint arXiv:1606.03490*, 2016.
- [42] Geoffrey Ernest Richard Lloyd and Geoffrey Ernest Richard Lloyd. *Polarity and analogy: two types of argumentation in early Greek thought.* Hackett Publishing, 1992.
- [43] Jennifer Marie Logg. Theory of machine: When do people rely on algorithms? 2017.
- [44] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting* of the Association for Computational Linguistics: Human Language Technologies, pages 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics.
- [45] A. J. Moss and L. Litman. How do most mturk workers work?, Mar 2019.
- [46] David W Nickerson and Todd Rogers. Political campaigns and big data. *Journal of Economic Perspectives*, 28(2):51–74, 2014.
- [47] Gali Noti, Noam Nisan, and Ilan Yaniv. An experimental evaluation of bidders' behavior in ad auctions. In *Proceedings of the 23rd international conference on World wide web*, pages 619–630. ACM, 2014.
- [48] Nikolaas N Oosterhof and Alexander Todorov. The functional basis of face evaluation. Proceedings of the National Academy of Sciences, 105(32):11087–11092, 2008.
- [49] Ravi B Parikh, Ziad Obermeyer, and Amol S Navathe. Regulation of predictive analytics in medicine. *Science*, 363(6429):810–812, 2019.
- [50] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pages 1532–1543, 2014.
- [51] Richard E Petty and John T Cacioppo. The elaboration likelihood model of persuasion. In Communication and persuasion, pages 1–24. Springer, 1986.
- [52] Forough Poursabzi-Sangdeh, Daniel G Goldstein, Jake M Hofman, Jennifer Wortman Vaughan, and Hanna Wallach. Manipulating and measuring model interpretability. *arXiv preprint arXiv:1802.07810*, 2018.
- [53] 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.
- [54] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144. ACM, 2016.
- [55] Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. *arXiv preprint arXiv:1802.09127*, 2018.
- [56] Andrew Slavin Ross, Michael C Hughes, and Finale Doshi-Velez. Right for the right reasons: Training differentiable models by constraining their explanations. *arXiv preprint arXiv:1703.03717*, 2017.
- [57] P Wesley Schultz, Jessica M Nolan, Robert B Cialdini, Noah J Goldstein, and Vladas Griskevicius. The constructive, destructive, and reconstructive power of social norms. *Psychological science*, 18(5):429–434, 2007.
- [58] Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- [59] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. arXiv preprint arXiv:1706.03825, 2017.
- [60] Megan Stevenson and Jennifer Doleac. Algorithmic risk assessment tools in the hands of humans. 2018.
- [61] Peter Thompson. Margaret thatcher: a new illusion. Perception, 1980.

- [62] Alexander Todorov, Chris P Said, Andrew D Engell, and Nikolaas N Oosterhof. Understanding evaluation of faces on social dimensions. *Trends in cognitive sciences*, 12(12):455–460, 2008.
- [63] Yehonatan Turner and Irith Hadas-Halpern. The effects of including a patient's photograph to the radiographic examination. In *Radiological Society of North America scientific assembly and annual meeting. Oak Brook, Ill: Radiological Society of North America*, volume 576, 2008.
- [64] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *science*, 185(4157):1124–1131, 1974.
- [65] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In Advances in Neural Information Processing Systems, pages 2692–2700, 2015.
- [66] Whittaker M. West, S.M. and K. Crawford. Discriminating systems: Gender, race and power in ai., 2019.
- [67] Christian Wirth, Riad Akrour, Gerhard Neumann, and Johannes Fürnkranz. A survey of preference-based reinforcement learning methods. *The Journal of Machine Learning Research*, 18(1):4945–4990, 2017.
- [68] Michael Yeomans, Anuj Shah, Sendhil Mullainathan, and Jon Kleinberg. Making sense of recommendations. *Journal of Behavioral Decision Making*, 2017.
- [69] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. Understanding the effect of accuracy on trust in machine learning models. 2019.
- [70] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *International Conference on Machine Learning*, pages 325–333, 2013.

Appendix A General Optimization Issues

Initialization. Because acquiring human labels is expensive, it is important to initialize ϕ to map to a region of the representation space in which there is variation and consistency in human reports, such that gradients lead to progress in subsequent rounds. In some representation spaces, such as our 2D projections of noisy 3D rotated images, this is likely to be the case (almost any 3D slice will retain some signal from the original 2D image). However, in 4+ dimensions, as well as with the subset selection and avatar tasks, there are no such guarantees. To minimize non-informative queries, we adopt two initialization strategies:

1. Initialization with a computer-only model: In scenarios in which the representation space is a (possibly discrete) subset of input space, such as in subset selection, the initialization problem is to isolate the region of the input space that is important for decision-making. In this situation, it can be useful to initialize with a computer-only classifier. This classifier should share a representation-learning architecture with ϕ but can have any other classifying architecture appended (although simpler is likely better for this purpose). This should result in some ϕ which at least focuses on the features relevant for classification, if not necessarily in a human-interpretable format. For an example, see Table 1, which shows how a machine-initialized Pointer Network selects relevant words but uses an idea of '1'(top set) and '0'(bottom set) which is not discernible to human users.

Table 1: Examples of words chosen by the Ptr-Net when initialized with machine

init								
pos	pos	neg						
deplorable	admirable	abomination						
marcel	clever	sullen						
gifted	nighttime	occasion						
aroused and	academia by	abomination sullen						
rumors	Gabon	said						

2. Initialization to a desired distribution with a WGAN: In scenarios in which the initialization problem is to isolate a region of representation space into which to map all inputs, as in the avatar example, in which we wish to test a variety of expressions without creating expression combinations which will appear overly strange to participants, it can be useful to hand-design a starting distribution over representation space and initialize ϕ with a Wasserstein GAN [2]. In this case, we use a Generator Network with the same architecture as ϕ but allow the Discriminator Network to be of any effective architecture. As with the previous example, this results in an ϕ in which the desired distribution is presented to users, but not necessarily in a way that reflects any human intuitive concept.

Convergence. As is true in general of gradient descent algorithms, our framework is not guaranteed to find a global optimum but rather is likely to end up at a local optimum dependent on both the initialization of ϕ and \hat{h} . In our case, however, the path of gradient descent is also dependent on the inherently stochastic selection and behavior of human users. If users are inconsistent or user groups at different iterations are not drawn from the same behavior distribution, it is possible that learning at one step of the algorithm could result in convergence to a suboptimal distribution for future users. It remains future work to test how robust machine learning methods might be adapted to this situation to mitigate this issue.

Regularization/Early Stopping As mentioned in Section 2, training ϕ will in general shift the distribution of the representation space away from the region on which we have collected labels for \hat{h} in the previous iterations, resulting in increasing uncertainty in the predicted outcomes. We test a variety of methods to account for this, but developing a consistent scheme for choosing how best to maximize the information in human labels remains future work.

• **Regularization of** \hat{h} : We test regularization of \hat{h} both with Dropout and L2 regularization, both of which help in preventing overfitting, especially in early stages of training, when the representation distribution is not yet refined. As training progresses and the distribution

 $\phi_{\theta}(x)$ becomes more tightly defined, decreasing these regularization parameters increases performance.

- Training \hat{h} with samples from previous iterations: We also found it helpful in early training iterations to reuse samples from the previous human labeling round in training \hat{h} , as inspired by [7]. We weight these samples equally and use only the previous round, but it may be reasonable in other applications to alter the weighting scheme and number of rounds used.
- Early stopping based on Bayesian Linear Regression: In an attempt to quantify how the prediction uncertainly changes as θ changes, we also implement Bayesian Linear Regression, found in [55] to be a simple but effective measure of uncertainty, over the last layer of $\hat{h}(\phi_{\theta})$ as we vary θ through training. We find that in early iterations of training, this can be an effective stopping criterion for training of ϕ . Again, as training progresses, we find that this mostly indicates only small changes in model uncertainty.

Human Input. Testing on mTurk presents additional challenges for our application:

- In some applications, such as loan approval, Mturk users are not experts. It is therefore difficult to convince them that anything is at stake (we found that bonuses did not meaningfully affect performance), It is also difficult to directly measure effort, agency, trust, or autonomy, all of which result in higher variance in responses.
- In many other applications, the ground truth is generated by humans to begin with (for example, sentiment analysis). Since we require ground truth for training, in these task it cannot be expected of humans to outperform machines.
- As the researchers found in [36], there can be large variance in the time users take to complete a given task. Researchers have found that around 25% of mTurk users complete several tasks at once or take breaks during HITs [45], making it difficult to determine how closely Turkers are paying attention to a given task. We use requirements of HIT approval rate greater than 98%, US only, and at least 5,000 HITs approved, as well as a simple comprehension check.
- Turker populations can vary over time and within time periods, again leading to highly variate responses, which can considerably effect the performance of learning.
- Recently, there have been concerns regarding the usage of automated bots within the mTurk communiy. Towards this end, we incorporated in the experimental survey a required reading comprehension task and a captcha task, and filtered users that did not succeed in these.

Appendix B Experimental Details

B.1 Decision-compatible 2D projections

In the experiment, we generate 1000 examples of these point clouds in 3D. The class of ϕ is a 3x3 linear layer with no bias, where we add a penalization term on $\phi^T \phi - \mathbb{I}$ during training to constrain the matrix to be orthogonal. Humans are shown the result of passing the points through this layer and projecting onto the first two dimensions. The class of \hat{h} is a small network with 1 3x3 convolutional layer creating 3 channels, 2x2 max pooling, and a sigmoid over a final linear layer. The input to this network is a soft (differentiable) 6x6 histogram over the 2D projection shown to the human user.

In an interactive command line query and response game we tested ourselves, ϕ was consistently able to find a representation that allowed for 100% accuracy. Many times this was the projection that appeared to be an 'x' and 'o' shown in Figure 6, but occasionally it was user-specific. For example, a user who associates straight lines with the 'x' may train the network to learn any projection for 'x' that includes many points along a straight line.

The architecture of ϕ and \hat{h} are described in Section 3. For training, we use a fixed number of epochs (500 for \hat{h} and 300 for ϕ) with base learning rates of .07 and .03, respectively, that increase with lower accuracy scores and decrease with each iteration. We have found these parameters to work well in practice, but observed that results were not sensitive to their selection. The interface allows the number of rounds and examples to be determined by the user, but generally 100% accuracy can be achieved after about 5 rounds of 10 examples each.



Figure 6: Images of x-o interface

B.2 Decision-compatible feature selection

The input to our pointer network is a sequence of 100-dimensional GloVe[50] embeddings of words. The outputs to \hat{h} are one-hot vectors of the selected words' GloVe embeddings multiplied by the softmax probabilities output by the attention mechanism. This allows for differentiable subset selection. The outputs to the human user are subsets of words.

Here \hat{h} attempts to replicate the evaluation of the human user on each individual word selected by mapping the GloVe embedding for the word to the human value for that word. With real humans in the loops, we would allow this to be -1 (negative sentiment), 0 (neutral sentiment), or 1 (positive sentiment).

In this experiment, to isolate the performance of the Pointer Network with feedback from \hat{h} and because hand-labeling examples without access to a crowd is time-consuming, these evaluations were made by a simple simulation of how a human might make decisions. In our reference task of sentiment classification, the simulation assigns positive and negative weights to all words, with explicitly positive words receiving a weight $w_p \sim U[.5, 1]$, explicitly negative words receiving a weight $w_n \sim U[-1, -.5]$, and all other neutral words receiving a weight $w_i \sim U[-.1, .1]$. The positive and negative weights are fixed for any given word throughout a dataset, so for example "good" has the same value every time it appears in an example. While this represents a very rough approximation of human text evaluation, it has the clear benefit of being able to be queried many times , which allows us to test whether or not the Pointer Network can succeed in combination with \hat{h} before proceeding to tests with real human users. Note that while training was performed with a simulator, evaluation on the test set was done using real human queries and therefore represent human performance.

Datasets are generated by taking the first 40 alphanumeric non-stop words from the IMDB review dataset [44] for examples with at least 40 such words.

We additionally use LIME to explain a Random Forest Classifier with 500 estimators and max depth 75 on the bag of words transformation of the dataset.

B.3 Decision-compatible algorithmic avatars

B.3.1 Data Preprocessing

We use the *Lending Club* dataset, which we filter to include only loans for which we know the resolution (either default or paid in full, not loans currently in progress) and to remove all features that would not have been available at funding time. We additionally drop loans that were paid off in a single lump sum payment of at least 5 times the normal installment. This results in a dataset that is 49% defaulted and 51% repaid loans. Categorical features are transformed to one-hot dummy variables. There are roughly 95,000 examples remaining in this dataset, of which we split 20% into the test set.

B.3.2 Learning architecture and pipeline

The network ϕ takes as input the standardized loan data. Although the number of output dimension are \mathbb{R}^9 , ϕ outputs vectors in \mathbb{R}^{11} . This is because the some facial expressions do not naturally coexist as compound emotions, i.e., happiness and sadness [18]. Hence, we must add some additional constraints to the output space, encoded in the extra dimensions. For example, happiness and sadness are split into two separate parameters (rather than using one dimension with positive for happiness and negative for sadness). The same is true of "happy surprise", which is only allowed to coincide



Figure 7: Visualization of reconstruction component

with happiness, as opposed to "sad surprise". For parameters which have positive and negative versions, we use a tanh function as the final nonlinearity, and for parameters which are positive only, we use a sigmoid function as the final nonlinearity.

These parameters are programmatically mapped to a series of Webmorph [14] transformation text files, which are manually loaded into the batch transform/batch edit functions of Webmorph. We use base emotion images from the CFEE database [18] and trait identities from [48]. This forms ρ for this experiment.

The network ϕ is initialized with a WGAN to match a distribution of parameters chosen to output a fairly uniform distribution of feasible faces. To achieve this, each parameter was chosen to be distributed according to one of the following: a clipped $\mathcal{N}(0,4)$, $\mathcal{U}[0,1]$, or Beta(1,2). The choice



Figure 8: \hat{h} does not necessarily have to match h well to lead to an increase in accuracy

of distribution was based on inspection as to what would give reasonable coverage over the set of emotional representations we were interested in testing. In this initial version of ϕ , x values end up mapped randomly to representations, as the WGAN has no objective other than distribution matching.

In the first experiment, we collect approximately 5 labels each (with minor variation due to a few mTurk users dropping out mid-experiment) for the LASSO feature subset of 400 training set x points and their ϕ_0 mappings (see Figure 10). a is taken to be the percentage of users responding "approve" for each point.

To train \hat{h} , we generate 15 different training-test splits of the collected $\{z, a\}$ pairs and compare the performance of variations of \hat{h} in which it is either initialized randomly or with the \hat{h} from the previous iteration, trained with or without adding the samples from the previous iteration, and ranging over different regularization parameters. We choose the training parameters and number of training epochs which result in the lowest average error across the 15 random splits. In the case of random initialization, we choose the best out of 30 random seeds over the 15 splits.

To train ϕ , we fix \hat{h} and use batches of 30,000 samples per epoch from the training set, which has 75,933 examples in total. In addition to the reconstruction regularization term $||x - \psi(\phi(x))||_2^2$ (see Figure 7) and the binary cross entropy accuracy loss, ϕ here also features a constraint penalty that prevents co-occurrence of incompatible emotions.

We train ϕ for 2,000 epochs with the Adam optimizer for a variety of values of α , where we use α to balance reconstruction and accuracy loss in the form $\mathcal{L}_{total} = \alpha \mathcal{L}_{acc} + (1 - \alpha) \mathcal{L}_{rec}$. We choose the value of α per round that optimally retains x information while promoting accuracy by inspecting the accuracy vs. reconstruction MSE curve. We then perform Bayesian Linear Regression over the final

layer of the current \hat{h} for every 50th epoch of ϕ training and select the number of epochs to use by the minimum of either 2,000 epochs or the epoch at which accuracy uncertainty has doubled. In all but the first step, this resulted in using 2,000 epochs.

At each of the 2-5th epochs, we choose only 200 training points to query. In the 6th epoch we use 200 points from the test set.

B.4 Results by user type

In the end of the survey, we ask users to report their decision method from among the following choices:

- I primarily relied on the data available
- I used the available data unless I had a strong feeling about the advice of the computer system
- I used both the available data and the advice of the computer system equally
- I used the advice of the computer system unless I had a strong feeling about the available data
- I primarily relied on the advice of the computer system
- Other

The percentage of users in each of these groups varied widely from round to round. We consider the first two conditions to be the 'Data' group, the third to be the 'Equal' group, and the next two to be the 'Computer Advice' group. While the groups are too small to draw many conclusions from this data, we find that users who report only or primarily using the data increase in mean accuracy from .51 in round 1 to .65 in round 6 (p < .001).

This implies at least one of the following: users misreport their decision method; users believe they are not influenced by the advice but in fact are; as the algorithmic evidence becomes apparently better, only the population of users who are comparatively skilled at using the data continue to do so.



Figure 8: Richness of avatar representation. A visualization of 200 avatars randomly sampled from the held-out test set, grouped by either human (top) or machine (bottom) predictive probability (0.2 in blue, 0.8 in orange, with a tolerance of 0.05). Avatars are positioned based on a 1D PCA dimensionality reduction of their corresponding feature vectors *z*, along which a 'gradient' of facial changes can be observed. **Top**: Here avatars are grouped by human predictive probability. The figure shows how for the same human decisions, learning results in avatars of varied and complex facial expressions, conveying rich high-dimensional information. Interestingly, avatars corresponding to loan denial exhibit more variance, suggesting that there may be more 'reasons' for denying a loan than for approving one. **Bottom**: Here avatars are grouped by machine predictive probability. Since all examples in each group have the same predictive probability, they are equally similar, which does not facilitate a clear notion for reasoning. In contrast, avatars maintain richness in variation, and can be efficiently used for reasoning (e.g., via similarity arguments) and other downstream tasks.



Figure 9: Results by Reported User Type

B.5 Diversity in avatar representation

We believe the additional dimensionality of the avatar representation relative to a numerical or binary prediction of default is useful for two reasons. Most importantly, high dimensionality allows users to retain an ability to reason about their decisions. In particular, avatars are useful because people likely have at least two inherent mental reference points for what they believe to be 'good' and 'bad' faces. Moreover, users who have a more sophisticated mental reference space than this either inherently or because they have undergone training with the algorithm may be able to teach the advising system to match specific reasoning patterns to specific characteristics over time. Additionally, when the advising system does not have a strong conviction about a prediction, presenting neutral advice should encourage the user to revisit the data, whereas percentages above or below the either base rate of default or 50% may suffer from the anchoring effect [64, 53].

Appendix C Notes on Facial Avatars

We are aware of the many concerning ways in which faces can have been used in AI systems in discriminatory ways [66]. Ours is *not* a paper about bias, and we have aimed to minimize these concerns to the extent possible, e.g., by restricting to variations on the image of a single person. Given current generative flow model technology, it is feasible that a similar experiment could be conducted using other abstract out-of-domain representation, such as landscapes, scenes, or even abstract color splashes generated according to latent parameters. Among these, we chose faces primarily for the following reasons:

• Humans have some pre-existing, shared representations in facial emotion space. This holds to a larger extent when for populations of higher homogeneity (i.e., our testing group, which included only Americans). This is convenient, as with the other representations we would have had to have workers undergo a training round so that they would have some shared conception of the representation space.

1 o	f 40				
Applicant ID: 761163 Interest Rate: 10.59% Number of Payments: 36 months Debt to Income Ratio: 14.57 Derogatory Public Records: 0 Annual Income: \$55000	Computer-Generated Advice:	Applicant ID: 759933 Interest Rate: 10.99% Number of Payments: 36 months Debt to Income Ratio: 12.79 Derogatory Public Records: 0 Annual Income: \$85000	Computer-Generated Prediction of Default Probability: 20%		
Based on the available inform I believe the b	ation and the system's advice, est option is to:	Based on the available information and the system's advice, I believe the best option is to:			
○ Approve the loan		O Approve the loan			
O Deny the loan		O Deny the loan			
(:	a)	(b)			

Figure 10: Images from mTurk questionnaire

- Humans are capable of perceiving, processing, and inferring faces at almost effortlessly and with remarkable speed. Inferences are consistent and, to some extent, universal. This is made possible due to innate and dedicated neural circuitry for face perception found in human brains, playing the role of 'brain GPU' in our learning framework.
- There are many pre-existing tools for facial morphing and face recognition, which can be useful as reliable components in the training pipeline.

We emphasize that this is merely a convenient example of a broader space of potential representations and not an important component of our framework.

Moreover, the expressions of the facial avatar developed here are only intended to be used in the context of the present system, to provide a suitable representation of the data that is relevant to a given individual and helps with decision making. The facial avatar is not intended to be used to drive decision making in other contexts, and indeed, its very generation requires access to a particular set of covariates for an individual.

Appendix D Select Turker quotes

- "I wasn't always looking at just happiness or sadness. Sometimes the expressions seemed disingenuously happy, and that also threw me off. I don't know if that was intentional but it definitely effected my gut feeling and how I chose."
- "In my opinion, the level of happiness or sadness, the degree of a smile or a frown, was used to represent applications who were likely to be payed back. The more happy one looks, the better the chances of the client paying the loan off (or at least what the survey information lead me to believe)."
- "I was more comfortable with facial expressions than numbers. I felt like a computer and I didn't feel human anymore. Didn't like it at all."