

LIMEADE: From AI Explanations to Advice Taking

BENJAMIN CHARLES GERMAIN LEE*, University of Washington & Allen Institute for Artificial Intelligence, USA

DOUG DOWNEY, Allen Institute for Artificial Intelligence, USA

KYLE LO, Allen Institute for Artificial Intelligence, USA

DANIEL S. WELD, University of Washington & Allen Institute for Artificial Intelligence, USA

Research in human-centered AI has shown the benefits of systems that can explain their predictions. Methods that allow an AI to take advice from humans in response to explanations are similarly useful. While both capabilities are well-developed for *transparent* learning models (e.g., linear models and GA²Ms), and recent techniques (e.g., LIME and SHAP) can generate explanations for *opaque* models, little attention has been given to advice methods for opaque models. This paper introduces LIMEADE, the first general framework that translates both positive and negative advice (expressed using high-level vocabulary such as that employed by post-hoc explanations) into an update to an arbitrary, underlying opaque model. We demonstrate the generality of our approach with case studies on seventy real-world models across two broad domains: image classification and text recommendation. We show our method improves accuracy compared to a rigorous baseline on the image classification domains. For the text modality, we apply our framework to a neural recommender system for scientific papers on a public website; our user study shows that our framework leads to significantly higher perceived user control, trust, and satisfaction.

CCS Concepts: • **Information systems** → **Recommender systems**; • **Computing methodologies** → **Learning from critiques**; • **Human-centered computing** → **HCI design and evaluation methods**; **User interface management systems**.

Additional Key Words and Phrases: Explainable Recommendations, Explainable AI, Advice Taking, Interactive Machine Learning, Human-AI Interaction

1 INTRODUCTION

A long-standing vision in AI is the construction of an *advice taker*, a system whose behavior, in the words of John McCarthy, “will be improvable merely by making statements to it, telling it about its symbolic environment and what is wanted from it. To make these statements will require little if any knowledge of the program or the previous knowledge of the advice taker” [53]. Indeed, today’s guidelines for human-AI interaction dictate that ML systems should be able to explain their predictions to end-users and accept advice and corrections from them [4, 6]. Both explanation and advice-taking methods exist for transparent models, such as linear classifiers or generalized additive models (GA2Ms) [17, 86, 88], and their benefits for transparent recommenders have been demonstrated within the human-in-the-loop machine learning and human-AI interaction literature [13, 43]. These advice-taking approaches allow the human to provide high-level feedback on how specific input features should be driving the transparent model’s behavior.

*Work performed during internship at the Allen Institute for Artificial Intelligence and Ph.D. at the University of Washington.

Authors’ addresses: Benjamin Charles Germain Lee, bcgl@cs.washington.edu, University of Washington & Allen Institute for Artificial Intelligence, USA; Doug Downey, dougd@allenai.org, Allen Institute for Artificial Intelligence, USA; Kyle Lo, kylel@allenai.org, Allen Institute for Artificial Intelligence, USA; Daniel S. Weld, weld@cs.washington.edu, University of Washington & Allen Institute for Artificial Intelligence, USA.

However, opaque models, such as boosted decision forests and deep neural networks, are a different story. Because they often provide the highest performance and are widely used, numerous researchers have investigated methods for generating post-hoc explanations of opaque ML models — typically by creating a transparent approximation to the opaque model, called an explanatory model [29]. Several researchers have developed methods for translating high-level *human advice* into specific classes of differentiable, neural models [24, 47, 64, 66, 69], but to our knowledge only one paper (Schramowski *et al.* [69]) have introduced a method that works for *arbitrary* opaque models, and it is not capable of handling advice that corrects an agent’s erroneous predictions (Section 6.3).

Furthermore, even the advice-taking methods whose application is restricted to specific opaque model classes [47, 64, 66] have limited empirical evaluation, often restricted to datasets that have been artificially biased (e.g., Decoy MNIST and Iris-Cancer [66]) in a way that a simple human tip (e.g., “Ignore the artifact in the lower right corner”) can correct the problem. To demonstrate that advice-taking methods are truly useful, experiments with large, real-world domains seem essential.

Thus, two central questions for human-AI interaction remain unanswered:

- (1) Can one translate high-level human advice into a correction to an arbitrary, opaque, machine-learned model which uses a different set of features than those used to express the advice?
- (2) Do these methods allow end-users to improve the accuracy of natural, real-world models more easily than by simply annotating more examples?

This paper answers the first question affirmatively, but presents mixed results on the second. Specifically, we present LIMEADE, a general framework for updating an arbitrary, opaque machine learned model given high-level human advice, e.g. phrased in the same vocabulary used by a posthoc explanation of its behavior. As shown in Figure 1, our approach builds upon explanatory approaches such as LIME [63] and SHAP [52] that describe the local behavior of a model in the region of a given instance. Given a trained model and an instance to be classified, these post-hoc approaches output an explanation in the form of a weighted list of *interpretable* features (typically distinct from the features utilized in the opaque model) that influence the instance’s classification. With LIMEADE, a user can then provide feedback in the same high-level terms as the explanation in order to modify the original, opaque model. LIMEADE converts this user advice back into the original feature space of the opaque classifier by generating pseudo-examples representative of these features and retraining. Unlike other methods of control intended for machine learning practitioners and model developers,

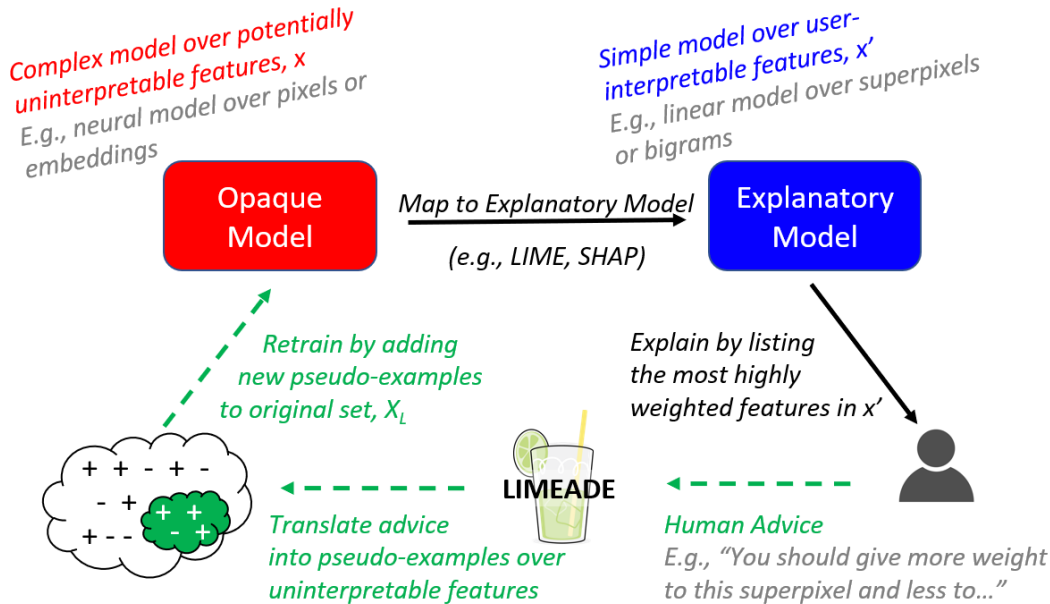


Fig. 1. LIMEADE takes a user’s advice – given in terms of features of the explanatory model – and then modifies the original, opaque model by retraining. This is challenging because the mapping from opaque to explanatory model is typically many-to-one and hence not invertible.

LIMEADE empowers end-users with little or no machine learning expertise to tune the system using high-level advice.

LIMEADE builds on the longstanding research areas of human-in-the-loop machine learning and interactive machine learning to provide a framework that is sufficiently general to address a wide range of model architectures, tasks, and modes of advice. We emphasize that LIMEADE is a general framework in three distinct senses:

- (1) LIMEADE can be utilized for a wide range of advice-taking applications, from explanatory debugging to personalized recommendation.
- (2) LIMEADE is architecture-agnostic and enables advice taking for different types of opaque machine learning models, including both classifiers and rankers.
- (3) LIMEADE accepts different types of human advice (in this paper, we focus on advice given as binary feedback in terms of high-level features).

Accordingly, we show that our framework is general by demonstrating its success on seventy real-world models across two broad domains: image classification and text ranking. For our first case study, we use LIMEADE to give advice to twenty binary image classifiers (e.g., models predicting “giraffe” or “not giraffe”) that are built on precomputed neural embeddings [32]. Our implementation of LIMEADE in the image domain translates a human’s simulated advice to the classifier in response to a LIME explanation expressed using superpixel features. Using this simulated advice, we demonstrate that this implementation significantly improves system accuracy, compared to a strong baseline, in a few shot setting. To accelerate future research, we will release our LIMEADE code and that for the image experiments.

For our second case study, we incorporated LIMEADE within Semantic Sanity, a publicly-deployed research paper recommender system with hundreds of users. While recommendations are made using an opaque neural model built on top of precomputed paper embeddings [19], LIMEADE allows humans to provide advice in terms of unigrams and bigrams (e.g., marking them as of interest or not) that are suggested by an approximate, linear explanatory model. In a simulation study based on organic user feeds in the log data, we show that explanation-based advice taking improves recommender quality, but we fail to find a significant improvement compared to adding a comparable number of labeled examples. We also perform an in-person user study showing that users feel that the ability to provide high-level feedback significantly improves their sense of trust, control and system transparency. By analyzing the logs of 300 Semantic Sanity users, our work also reveals a fundamental tension, which we call the *explanation-action trade-off*, between providing the most faithful explanations (by greedily selecting the most influential features) and providing the best affordances for advice taking (by providing a variety of possible refinements). With a faithful explanation approach that surfaces the most influential features, a user’s positive action on a feature via a LIMEADE update results in increased importance, causing the feature to dominate explanations. A more diverse explanation strategy can be accomplished by presenting more unique terms, which comes at the expense of faithfulness to the explanatory model.

Significantly, our paper leaves a number of questions surrounding the advice-taking problem unanswered. For example, we do not conclusively answer the question of whether advice-taking methods allow end-users to improve the accuracy of real-world, opaque models more easily than by annotating more examples. Moreover,

in our image domain experiment, we uncover that the effectiveness of advice-taking methods may decrease with more supervision. Lastly, it is important to further study how advice-taking fits into broader frameworks within human-in-the-loop machine learning that incorporate human interventions with design parameters, model and algorithm choice, error preferences, and beyond [4]. In many ways, we view this paper as a “Call to Action” to galvanize more researchers to study the advice-taking problem for opaque machine learners, as it is a rich area of study within human-AI interaction with many questions still to answer.

2 LIMEADE: ADVICE TAKING FOR OPAQUE MODELS

In this section, we provide a formal overview of the LIMEADE framework and detail how it can be applied to opaque machine learning models to enable advice taking. With LIMEADE, we assume that the human would like to give advice to an opaque machine learning model. By *opaque*, we mean that the model architecture may be completely unknown, or (if known), it may have too many parameters and nonlinearities for a human to understand. However, we assume that the model’s inputs and outputs are available and that the model can be retrained on new examples. We work in a semi-supervised learning setting, in which the goal is to learn a hypothesis that maps an s -dimensional real-valued input vector to a label (for classification) or a real-valued output score in $[-1, 1]$ (e.g., for recommendation). We are given a set \mathcal{X}_L of labeled training examples (x, y, w) , where $x \in \mathbb{R}^s$, y is the value to be learned, and w is the weight assigned to the example when training. Additionally, we optionally have a large, dense pool \mathcal{X}_U of unlabeled examples (x) . Our explainable machine learning problem setting closely follows that of previous work in explainable ML [52, 63]. We assume that each instance x can be represented as a binary-valued vector x' that lies in an *interpretable* space. For example, in the text domain, the dimensions of x might contain embeddings produced by a transformer, whereas the dimensions of x' would correspond to interpretable features such as term frequency-inverse document frequency (TF-IDF) values for n -grams.¹ In the image domain, the dimensions of x would be pixels, while the dimensions of x' might be superpixels [63] or fine-grained features [3, 41].

Given an instance x to explain, our approach uses an *explanatory model* g in the interpretable space that locally approximates the opaque classifier f , i.e., $g(h'(z)) \approx f(z)$ for z' nearby x' . The model g can be any interpretable model, such as a decision tree or linear model, produced using LIME or a comparable method. We refer to the method that produces g as $\text{EXPLAIN}(f, x, h')$.

Algorithm 1 details LIMEADE’s approach to enabling a model to take advice, and Figure 2 illustrates a concrete example of applying LIMEADE on the paper recommendation domain. Given an instance of interest, x , we obtain an explanation $g(x')$ of the model’s output $f(x)$ using $\text{EXPLAIN}(f, x, h')$. The human can then provide a label on a feature of x' . Informally, a positive label on feature j of x' represents the human’s assessment that examples z' near

x' should tend to be positive when $z'[j] = 1$. For example, a user of our paper recommendation system might give a positive label to the term “BERT” in a natural language processing paper to indicate interest in papers about the technique.

Algorithm 1 Enabling an opaque model to take advice using LIMEADE. Given a set of required inputs, LIMEADE solicits human advice in response to an explanation of a classified instance and retrain the opaque model accordingly. EXPLAIN is a function that generates an explanation for a given model and instance.

Inputs:

$\mathcal{X}_L, \mathcal{X}_U$	// sets of labeled and unlabeled instances
$f_t : \mathbb{R}^s \rightarrow [-1, 1]$	// opaque classifier, version at time t
$x \in \mathbb{R}^s, x' \in \{0, 1\}^{s'}$	// instance & instance in interpretable space
$h' : \mathbb{R}^s \rightarrow \{0, 1\}^{s'}$	// mapping s.t. $x' = h'(x)$
$\pi_{x'} : \{0, 1\}^s \rightarrow \mathbb{R}_+$	// weighting based on distance
$k \in \mathbb{N}$	// number of pseudo-examples

- 1: $g_t = \text{EXPLAIN}(f_t, x, h')$ // obtain explanatory model
- 2: $\text{DISPLAY}(g_t, x')$ // display key features of $g_t(x')$ to end-user,
- 3: // who then selects one feature (indexed j) as + or – indicator of instance
receive $a \in \{-1, 1\}$ **and** $j \in \{1, \dots, s'\}$
- 4: // select k instances, label them using action a , and weight according to distance from x'
 $\mathcal{N}_x \leftarrow \{\}$
- 5: **for** $1, \dots, k$ **do**
- 6: $\tilde{x} = \text{GETINSTANCE}(x, x', \mathcal{X}_U)$
- 7: **if** $h'(\tilde{x})[j] = 1$ **then**
- 8: $\mathcal{N}_x \leftarrow \mathcal{N}_x \cup \{(\tilde{x}, a, \pi_{x'}(h'(\tilde{x})))\}$
- 9: **end if**
- 10: **end for**
- 11: $\mathcal{X}_L \leftarrow \mathcal{X}_L \cup \mathcal{N}_x$
- 12: $f_{t+1} \leftarrow \text{RETRAIN}(\mathcal{X}_L, f_t)$
- 13: **return** f_{t+1}

LIMEADE uses the human’s action to improve the opaque model f by creating a set of k training pseudo-examples with repeated calls to $\text{GETINSTANCE}(x, x', \mathcal{X}_U)$. We experiment with two implementations of GETINSTANCE : sampling and generative. Sampling from the unlabeled pool is effective when the unlabeled pool is relatively dense, meaning one can acquire many examples with interpretable features similar to those of x' . Generative approaches can be helpful when data are less dense. For example, with images, LIMEADE can create synthetic pseudo-examples by greying out random subsets of the superpixels in the input image, essentially reversing LIME’s process for generating the explanatory model, g . The generative approach also works in the textual domain, e.g., by creating a synthetic document with nothing but the tokens selected by the user.

LIMEADE only retains the pseudo-examples that contain the acted-upon feature j , i.e. those \tilde{x} for which $h'(\tilde{x})[j] = 1$. LIMEADE then

¹Term frequency-inverse document frequency, or TF-IDF, is a method of text featurization. Each document is featurized according to a fixed set of n -grams, where the feature value for the n -gram is given by the n -gram’s frequency in the specific document in question, times a weight that puts more emphasis on terms that are less common across the corpus [74].




Fig. 2. LIMEADE updates an arbitrary opaque ML model by creating pseudo-examples. Here, we consider a recommender system for papers. Small black o’s and +’s show the original training set (here, a user’s ratings of papers), and shaded regions denote the complex boundary of the opaque classifier f_t . In order to explain a prediction, $h(x')$, the system generates a locally faithful explanatory model using LIME or an alternative method. This is g_t , shown as a purple dotted line. In practice, the explanatory model likely has many more than the two dimensions shown above, but suppose ‘Turing’ and ‘agents’ are highly weighted terms, hence used in the explanation. When the human specifies feature-level advice, e.g., ‘I want more papers about “agents”’, it could be used to directly alter a linear explanatory model (creating the new purple dotted line g_{t+1}); however, no simple update exists for an arbitrary, opaque classifier, which may be nonlinear and use completely different features, such as word embeddings. Instead, LIMEADE generates positive pseudo-examples (shown as blue +’s) that have the acted-upon feature and are similar to the predicted example. The pseudo-examples are weighted (shown by relative size) by their distance to the predicted example x' that was used to elicit feedback. By retraining on this augmented dataset, LIMEADE produces an opaque classifier that has taken the advice, shown as a changed nonlinear decision boundary f_{t+1} .

assigns a value to each pseudo-example according to the user action: +1 if the user assigned a positive feature label, and −1 otherwise.

LIMEADE assigns each pseudo-example a weight based on its proximity to x' , with examples more similar to x' given higher weight.² The reasons to weight local examples more highly are twofold: the explanatory method may only be locally correct [63], and the human actions may only be locally applicable. For example, the positive label on “BERT” discussed earlier is helpful within the local scope of natural language processing papers, but could become misleading if applied globally—in biology papers for example, the term “BERT” often refers to a different meaning (the “BERT gene”). After selecting and weighting the pseudo-examples, LIMEADE can optionally condense the selections (e.g., collapsing the examples into a single centroid). Finally, LIMEADE adds the resulting pseudo-examples to the labeled training set \mathcal{X}_L and calls RETRAIN to train the classifier f on the new data set.

While Algorithm 1 is written in terms of binary classification, our approach generalizes naturally to the multiclass setting. This

²We measure proximity in the interpretable space, but it is equally possible to measure in the original space instead.

would entail that step 3 in Algorithm 1 solicit not only which feature was a positive or negative indicator, but also for which class—pseudo-examples would then be labeled in step 8 with respect to the chosen class. In the case of negative indicators in the multiclass setting, the pseudo-examples could be assigned random classes other than the chosen class.

We reiterate that LIMEADE is general in many senses: our framework is model-agnostic, applies to a diverse range of advice-taking applications, and enables different forms of advice taking. In the next sections, we present two case studies that highlight the general applicability of LIMEADE.

3 CASE STUDY 1: LIMEADE FOR IMAGE CLASSIFICATION

We now present our evaluation of LIMEADE in the image domain in order to study whether LIMEADE allows humans to update real-world models more effectively than simply labeling more examples. In particular, we use LIMEADE to enable updates based on simulated end-user advice for twenty deep neural image classifiers, e.g., a skateboard detector or fire hydrant detector. In Figure 3, we illustrate an example of how LIMEADE is used to process high-level

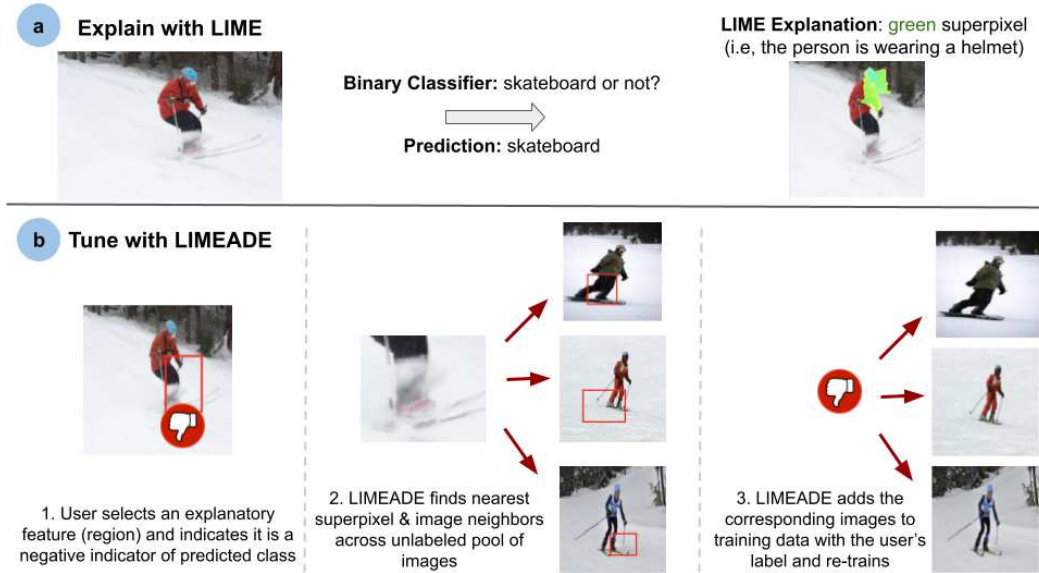


Fig. 3. **a)** Suppose an opaque classifier incorrectly classifies an image of a skier as a positive example of skateboarding. Suppose further that LIME returns an explanation showing a superpixel containing the skier’s helmet as a positive indicator of the skateboarding class. Having seen this explanation, the user realizes that the classifier is predicting “skateboard” based on a spurious confound and should be looking elsewhere (we note that the end-user, such as a crowd worker, must understand the classification task but needs neither domain-specific knowledge nor an understanding of machine learning). **b)** While a helmet is an appropriate positive indicator for the skateboarding class, the user gives the advice that another superpixel, containing skis and ski poles, is a negative indicator. LIMEADE translates this advice by updating the opaque model and retrieving unlabeled images containing superpixels most similar to this ski superpixel. These images are then added to the training data — with negative labels — and the model is retrained, completing the LIMEADE update. In general, a false positive classification will lead to negative feedback, and a false negative classification will lead to positive feedback (see Figure 2).

advice in this context. With this simulated image domain experiment, we wanted to study the following research questions:

- (1) Does advice taking with LIMEADE further improve classifier performance as compared to a rigorous baseline of adding more labeled examples?
- (2) How do LIMEADE-powered improvements change as a function of supervision?

3.1 Experimental Setup

In order to determine whether LIMEADE can support advice taking in the image domain, we evaluated on binary image classifiers, each comprising a logistic regression model trained on pre-computed image embeddings. As a base image dataset, we utilized 20,000 images from the COCO dataset [45]. In order to create superpixel features for LIMEADE feedback, we leveraged the same segmentation algorithm [57] used by LIME to process all 20,000 images. To generate embeddings for all images and superpixels, we retrieved their representations from the penultimate layer of a ResNet-50 backbone pre-trained on ImageNet [23, 32].

In order to ensure that our embeddings had not already been trained on the target classes in our experiment, we tested binary classifiers only on all 20 classes that are in COCO but not in ImageNet-1000.³ We wanted to measure the performance of a LIMEADE update relative to a baseline update, so we completed 100 randomized

³The 20 classes are: baseball glove, snowboard, giraffe, carrot, surfboard, fork, sink, cow, donut, toothbrush, knife, bed, horse, cake, motorcycle, frisbee, skateboard, fire hydrant, scissors, and suitcase.

initial configurations for each class. Moreover, for each configuration, we randomly constructed an initial training set of one positive and one negative instance (experiments in the 10-shot setting were less-effective, as described in Section 5.1).

We evaluated the two-shot accuracy of a logistic regression model on a held-out validation set and then performed one of the following two updates with both a randomly-drawn positive instance and a randomly-drawn negative instance simultaneously to preserve class balance:

- (1) **Baseline:** We update the model by adding the positive and negative instances to the training data and retraining
- (2) **LIMEADE:** First, we generate LIME explanations of the opaque classifier for both the positive and negative instance. In the positive case, we simulate a human’s advice in response to the explanation by utilizing the COCO segmentation masks to automatically give the superpixel(s) indicative of the class a positive label (i.e., in the case of “giraffe,” we select all superpixels containing giraffes using the COCO segmentation masks in the image labeled as “giraffe”). In the negative case, we give the superpixel most influencing the LIME explanation a negative label. We then generate embeddings of these labeled regions and retrieve the nearest superpixels and full images across the unused pool of 19,996 images. We append the corresponding embeddings to the training data along with + and – labels, respectively, and retrain. This simulated approach to human advice enabled us to study the effectiveness

Class	2-Shot Accuracy	Δ Baseline	Δ LIMEADE	p-value	Winner
Baseball Glove	76.73%	10.64% \pm 1.25%	10.75% \pm 1.52%	0.91	LIMEADE
Snowboard	75.97%	10.67% \pm 1.08%	10.74% \pm 1.42%	0.93	LIMEADE
Giraffe	74.54%	12.30% \pm 1.28%	16.42% \pm 1.51%	9.3×10^{-8}	LIMEADE*
Carrot	72.82%	7.37% \pm 1.15%	9.76% \pm 1.39%	4.9×10^{-4}	LIMEADE*
Surfboard	72.59%	8.23% \pm 1.19%	8.96% \pm 1.36%	0.34	LIMEADE
Fork	72.14%	7.70% \pm 1.24%	6.96% \pm 1.64%	0.53	Baseline
Sink	71.55%	10.38% \pm 1.12%	10.44% \pm 1.49%	0.95	LIMEADE
Cow	69.90%	8.66% \pm 1.00%	11.53% \pm 1.07%	5.5×10^{-6}	LIMEADE*
Donut	67.51%	8.23% \pm 0.96%	9.65% \pm 1.11%	0.093	LIMEADE
Toothbrush	65.85%	5.26% \pm 0.82%	4.86% \pm 1.02%	0.63	Baseline
Knife	65.47%	7.31% \pm 1.10%	7.43% \pm 1.43%	0.86	LIMEADE
Bed	65.16%	9.50% \pm 0.93%	11.73% \pm 1.16%	5.7×10^{-3}	LIMEADE*
Horse	63.66%	8.49% \pm 0.90%	10.20% \pm 1.31%	0.050	LIMEADE*
Cake	63.48%	8.53% \pm 1.02%	9.07% \pm 1.33%	0.49	LIMEADE
Motorcycle	63.16%	9.37% \pm 0.98%	15.97% \pm 1.08%	3.7×10^{-11}	LIMEADE*
Frisbee	62.67%	7.80% \pm 0.85%	7.03% \pm 1.18%	0.32	Baseline
Skateboard	61.09%	6.93% \pm 0.82%	7.52% \pm 1.03%	0.48	LIMEADE
Fire Hydrant	59.21%	6.31% \pm 0.75%	8.38% \pm 0.92%	8.7×10^{-3}	LIMEADE*
Scissors	57.38%	6.51% \pm 0.79%	4.92% \pm 1.05%	0.037	Baseline*
Suitcase	55.65%	4.04% \pm 0.63%	4.23% \pm 0.67%	0.71	LIMEADE
Total	66.83%	8.21%	9.33%	2.3×10^{-9}	LIMEADE*

Table 1. Updates using LIMEADE boost the accuracy of an opaque image classifier more than the baseline. Results are shown for 20 classes averaged over 100 randomly-initialized runs each, and the accuracy boosts are reported relative to an average initial, 2-shot accuracy on a test set. For the updates, standard errors are reported, and a * indicates p -value ≤ 0.05 . LIMEADE outperforms the baseline on 16 of 20 classes and provides an overall boost of 9.33%, as opposed to the baseline’s overall boost of 8.21%.

of LIMEADE updates by testing many initial configurations across a range of image classes.

We wanted to evaluate LIMEADE across different hyperparameter settings, so we varied the number of nearest neighbors included in the update ($n_{neighbors} = \{1, 5, 10, 25, 50, 100\}$), as well as the relative sample weight of the update ($w = \{0.25, 0.5, 1, 2, 4\}$), and performed a grid search. We evaluated performance on a balanced, held-out validation set of 400 positive examples and 400 negative examples for each class and selected the hyperparameters with the highest validation accuracy. This process yielded a relative sample weight of 0.25, as well as 50 nearest neighbors included in the update. With these hyperparameters selected, we then evaluated final performance on a separate, held-out test set of 400 positive examples and 400 negative examples for each class.

3.2 LIMEADE Feedback is More Effective than the Baseline

In Table 1, we report the net changes in classifier accuracy when making updates with LIMEADE and the baseline across all 20 classes and 100 runs per class, as evaluated on the test set. We find that LIMEADE updates with simulated advice outperform the baseline for 16 of 20 classes, giving an average boost of 9.33% compared to the baseline’s average boost of 8.21%. These results are statistically significant: a paired t-test of LIMEADE against the baseline yields a p -value of 2.3×10^{-9} across all 2,000 runs.

3.3 Diminishing Returns as Supervision Increases

While conducting our case study with the image domain, our empirical results indicated that the LIMEADE-powered improvements decrease as a function of more supervision. For example, repeating the experiments in the 10-shot setting, we find that LIMEADE gives an overall boost of 0.63%, whereas the baseline gives an overall boost of 0.88%. It is important to note, however, that there is a fundamental entanglement between training data and supervision with respect to LIMEADE. LIMEADE is most valuable when the original supervisory data is subject to spurious correlations (i.e., when teaching “cat,” if all cats seen in the training data happen to be black, a LIMEADE update communicating that color does not matter has high utility). If the training data is representative (because the training data contains more examples or because the examples themselves are better-selected), we expect a LIMEADE update to provide less utility, as there are fewer potential spurious correlations for a human to correct via a LIMEADE update. Indeed, as the quality of the originally-trained classifier approaches perfection, the value of LIMEADE goes to 0, much as the value of more training data also decreases. Our empirical evidence thus agrees with our intuition that a LIMEADE update is most valuable in the low supervision setting.

4 CASE STUDY 2: LIMEADE FOR PAPER RECOMMENDATION

For our second domain, we selected text ranking both for variety and importance. The overwhelming influx of new scientific publications poses a daily challenge for researchers [12, 26, 33, 37, 72]. However, based on Beel *et al.* [10]’s survey of 185 publications on academic paper recommendation, only a few systems explain why papers have been recommended or respond to user feedback other than liking/disliking specific papers, and all such systems rely on interpretable recommenders [8, 15, 38, 56, 82]. The ability to explain and take advice for higher-performance paper recommenders, therefore, fills an important void.

Furthermore, a complete evaluation of a human-AI interaction approach requires testing it with real users in the loop [6]. For LIMEADE, we wanted human users who were authentically motivated to understand and improve an ML classifier. In this regard, we built Semantic Sanity, a computer science (CS) research-paper recommender system based on Andrej Karpathy’s arXiv Sanity Pre-server [39]. Deployed as a publicly-available platform, Semantic Sanity enables users to curate feeds from over 150,000 CS papers recently published on arXiv.org. With this testbed, users are implicitly incentivized to understand and improve the recommender system powering their feed in order to receive more interesting papers. Note further that each user is a task expert, since the users determine their own preferences.

4.1 Neural Recommender

To generate individual recommendations, we utilize a neural model consisting of a linear SVM on top of neural paper embeddings pre-trained on a similar papers task [18]. Each paper is represented by the first vector (i.e., the [CLS] token typically chosen for text classification) after encoding the paper title and abstract using SciBERT [11]. The neural embedding model is finetuned on a triplet loss $\mathcal{L} = \max(0, v_+^T v_+ - v_i^T v_- + m)$ where m is a margin hyperparameter and v_i , v_+ and v_- are the vectors representing a query paper, a similar paper to the query paper, and a dissimilar paper to the query paper, respectively. The similar paper triples are heuristically defined using citations from the SEMANTIC SCHOLAR corpus [7], treating cited papers as more similar than un-cited papers. Recommendations are generated by training the model on a user’s annotation history, with additional negative examples randomly drawn from the full corpus of unannotated papers.

A user begins the process of curating their feed by either selecting a specific arXiv CS category or issuing a keyword search and then rating a handful of the resulting papers. A feed consists of a list of recommended papers sorted by predicted recommendation score (see Figure 4). Each paper can be rated using traditional “More like this” or “Less like this” buttons underneath each paper description.

4.2 Implementation of Explanations and Feedback

The UI for Semantic Sanity (Figure 4) displays a list of recommended papers and adorns each with an explanation comprising four terms; to the left of each term are thumbs-up and thumbs-down buttons, enabling the user to not only rate the papers themselves but also *give advice* in response to the explanation and indicate if they would

like to see more or fewer papers related to that term. The explanatory terms are generated using a simple, explanatory model (LIMEADE’s EXPLAIN function), which we implement as a linear model over uni- and bigram features. In particular, our linear model is defined as $g(x') = w_0 + \sum_i w_i x'_i$, and the explanation surfaced for $g(x')$ consists of high-impact terms in the model, i.e., those with high values for the product $w_i x'_i$. Specifically, we select the 20,000 features with the highest term frequency across our corpus. Our approach of using a post-hoc explanatory model is similar to that used by LIME, except to enable real-time performance our explanatory model is trained as a global, rather than local, approximation of the neural model [63]. This global approximation was chosen because testing on an early prototype revealed that generating explanations for a feed using LIME was too computationally expensive, since LIME requires sampling nearby instances and training a model for *each* recommendation on the page; this latency negatively impacted the recommendation experience.⁴

Given the explanatory model, LIMEADE’s DISPLAY function identifies explanations to display by computing each term’s contribution to the output of the linear model for the given paper, which is equal to the product of the term’s TF-IDF value for the paper with the term’s feature weight in the linear model. We note here that even though the explanatory model is a global approximation, the explanations are local ones, as this product encodes instance-specific information on why the paper has been recommended.⁵ Next to each explanatory term are thumbs up and down buttons (see Figure 4). When the user provides feedback with these buttons, LIMEADE generates pseudo-examples and retrains the neural recommender. We use a generative approach within GETINSTANCE that leverages the unlabeled pool of papers. We select the top 100 papers from the full corpus with the highest TF-IDF value for the feedback term and generate a single synthetic pseudoexample (i.e., we use $k = 1$) equal to the centroid of these papers’ embeddings with a weight of 1. The example is appended to the user’s history and labeled with the user’s annotation of the term (+/-).

4.3 Online Traffic

In the next section, we describe a controlled user study comparing Semantic Sanity with and without LIMEADE. However, since its public launch, Semantic Sanity has also attracted considerable organic traffic: users with accounts have constructed 2,478 feeds and have logged 21,713 paper annotations and 1,320 topic annotations (we note that annotating topics was only possible after the LIMEADE-based implementation was introduced on November 11, 2019, five months after the initial launch of Semantic Sanity). The target user base was computer science researchers, and the platform was advertised through social media and email lists. In Sections 4.7 and 5.2, we analyze a subset of the organic user logs as a complementary part of our evaluation.

⁴Note that in a LIME-style approach to instance-level explanation, one needs to either sample neighbors around the instance (which may be quite distant) or generate synthetic examples, which in our case requires generating paper embeddings on the fly. Both approaches introduced unacceptable latency.

⁵In a later implementation of Semantic Sanity, the platform also included a global explanation surfaced at the top of each feed, which is the subject of another study [61].

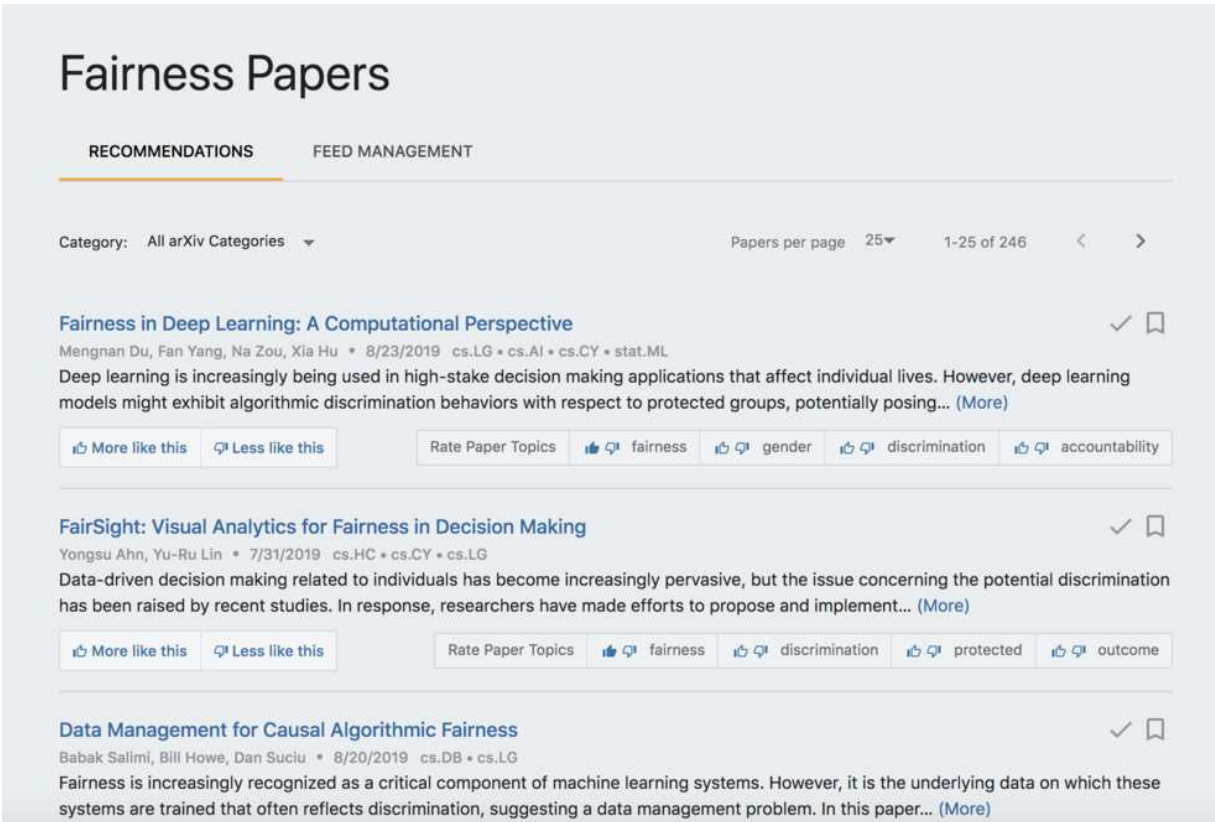


Fig. 4. The UI for a feed in Semantic Sanity. Users can rate the papers themselves with the “More like this” and “Less like this” buttons, a standard feed affordance. Under each paper, the system also presents four terms to explain why it was recommended and solicits feedback with “Rate Paper Topics” — by clicking thumbs up or down, the user can give advice by requesting that the feed include more or less of the specified topic.

4.4 User Study: Experimental Setup

In order to evaluate the effectiveness of our LIMEADE-based system for recommending papers with real users, we performed an in-person user study. With this user study, we wanted to address the following research questions:

- (1) Do participants prefer LIMEADE over a baseline of just explanations according to self-reported ratings of trust, control, transparency, intuition, paper coverage, and the overall system?
- (2) Does LIMEADE increase participants’ feed quality, evaluated quantitatively with blind ratings of recommended papers?
- (3) How do participants utilize the topic-rating affordances powered by LIMEADE?
- (4) What constructive feedback do participants have surrounding our particular instantiation of LIMEADE with Semantic Sanity?

We recruited 21 participants through a public university’s computer science email lists. All participants were adults who reported experience with reading computer science research papers in our screening questionnaire. Each session lasted one hour, and each participant was compensated with a \$25 Amazon gift card.

Participants were asked to curate feeds of computer science papers pertaining to a topic of their choice using two different recommendation user interfaces (UIs), one that used LIMEADE to provide advice-taking explanations, and one that did not present explanations, instead only allowing users to rate the papers themselves (the baseline); other than this difference, the UIs were the same. The participants were asked to choose a topic that they were interested in following over time as new papers are added to the arXiv, but not so general that it is already covered by an existing arXiv CS category (e.g., artificial intelligence). Once a topic was selected, each participant was asked to name the desired feed, which served as the goal for curation using both UIs.

Each participant began curation by selecting exactly three seed papers that were then used to initialize the feeds in both UIs. Both systems surfaced the same initial recommendations in response to the participant’s three seed papers and thus had identical initial states. Each participant was then presented with one of the two UIs and given instructions on how to use it. 11 participants received the baseline system first, and 10 received the LIMEADE system first. They were then presented with the second UI. For both UIs, the participants were told to use as many or as few annotations as desired until their feed was curated to their liking, or a maximum of

10 minutes was reached. We recorded the participants’ annotations for both feeds. After using each UI, the participants were asked to complete a short survey. They were then asked to rate a blind list of combined recommendations from the two feeds that they had curated, according to whether they would like to see each paper in their desired feed. These recommendations were generated on a held-out paper corpus, disjoint from the papers available within the feed UIs.

Data were successfully collected for all 21 participants. The participants’ chosen topics varied greatly, including “Spiking Neural Networks,” “Moderation of Online Communities,” and “Dialogue System Evaluation.”

4.5 User Study: Quantitative Results

4.5.1 User Experience: Participants Prefer LIMEADE. In the surveys administered after using each UI, we asked each participant to provide overall ratings for each system and to state which system they preferred along dimensions such as trust and intuitiveness. The results are summarized in Tables 2 and 3.⁶

Overall, participants rated our approach significantly higher than the baseline. They also rated it significantly higher on trust, control, and transparency, and on confidence that their recommendations were not missing relevant papers. Understandably, our LIMEADE system appeared less intuitive to participants than the baseline due to the increased complexity of the UI, though this result was not statistically significant. Finally, while not statistically significant due to small sample size, participants indicated more likelihood to use our system again over the baseline. In aggregate, these results indicate a higher-quality user experience with the LIMEADE system than with the baseline system.

4.5.2 Mixed Results for Feed Curation Time. In analyzing the times required by each participant to complete feed curation using the two systems, we observe that eight participants finished feed curation with the baseline system first, seven finished with the LIMEADE system first, and the remaining six utilized all ten minutes for the curation of both systems.

4.5.3 Most Participants used Both Paper and Topic Ratings. To explore the breakdown of participants’ rating habits with the baseline system and the LIMEADE system, we present Figure 5. In the left plot in Figure 5, we observe that participants displayed a high degree of variance in the number of ratings applied during feed curation, ranging from 7 annotations to 61 annotations with the baseline system. Comparing the total number of annotations made using the system with LIMEADE vs. the number of annotations made with the baseline, we find a best-fit slope of 0.913. This suggests that the participants made approximately the same number of annotations across both systems.

In the right plot, we observe that there is significant diversity in how participants applied topic annotations, ranging from 2 annotations to 27. However, most participants utilized a combination of paper and topic ratings, with more paper ratings than topic ratings on average. Interestingly, five out of the twenty-one participants

⁶For all statistical significance tests, we report adjusted p -values using the Holm-Bonferroni procedure for multiple comparisons [35] in R’s P.ADJUST library [60].

Which system...	Baseline	LIMEADE	p -value
...trust more?	4	17*	0.043
...more control?	0	21*	≈ 0
...more transparent?	3	18*	0.012
...more intuitive?	12	9	0.664
...not missing relevant papers?	3	18*	0.012

Table 2. Among 21 participants, most prefer our system over the baseline when prompted with these questions. (*) indicates a statistically significant result under a two-sided binomial test against a null hypothesis of no preference between the systems.

Likert scale rating	Baseline	LIMEADE	p -value
Overall system	3.38 ± 0.59	$3.85 \pm 0.57^*$	0.043
Would use again?	3.38 ± 1.16	3.90 ± 0.94	0.257

Table 3. Mean \pm Standard Deviation of 21 participant ratings of each system. Ratings were on a scale from 1 (worst/no) to 5 (best/yes). (*) indicates a statistically significant result under a two-sided paired t -test against a null hypothesis of zero mean difference between the systems.

provided more negative paper ratings than positive ones in the baseline; when presented with the LIMEADE affordances, no participants provided more negative paper ratings than positive ones, but four participants applied more negative topic ratings than positive ones.

4.5.4 Blind Ratings of Recommendations: No Significant Difference in Feed Quality. We also investigated whether the topic-level feedback provided by LIMEADE measurably increased the quality of participants’ feed. We showed participants the top 20 recommendations generated by both systems on the held-out corpus of papers and measured their ratings. Specifically, we computed the discounted cumulative gain (DCG)⁷ and average precision (AP), common metrics for assessing recommendation feed quality. For DCG, we observe a mean difference of 0.259 in favor of the baseline system recommendations; however, the corrected p -value for the two-sided, paired t -test for mean differences is 0.218, indicating no significant difference in feed quality between the two systems under DCG. For AP, we observe a mean difference of 0.0412 in favor of the baseline system, with a corrected p -value of 0.257, also indicating lack of significant difference in quality under AP. Based on the constructive feedback that we received, we speculate that this result could be improved by making implementation-specific adjustments to Semantic Sanity.

4.6 User Study: Qualitative Feedback

We analyze participants’ text responses and provide a sample of quotes that complement the quantitative results. After using each system, participants were asked to provide free text responses to the question, “Would you like to share anything else about using the system?” At the end of the study, participants were also asked, “Do you have any last thoughts that you would like to share regarding actionable explanations?” Overall, participants found the

⁷We did not use NDCG because the participants liked different numbers of papers between the two systems.

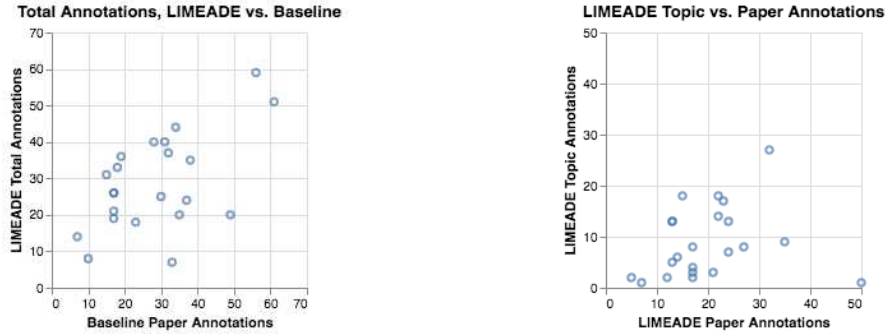


Fig. 5. Scatter plots showing (left) the total number of annotations used to curate a feed with LIMEADE (paper annotations + topic annotations) vs. the number of baseline paper annotations per user, and (right) the number of topic annotations vs. the number of paper annotations in the LIMEADE system. Most participants used LIMEADE-powered topic-level feedback as well as paper-level feedback.

advice-taking affordance granted by our system helpful: “*The explanations here were especially useful in their capacity as decisions rather than just explanations. I would have found them really really annoying if they were presented only as an explanation of why you thought I would like a paper, rather than an attribute I could ask for more or less of.*” In particular, participants stressed the importance of the LIMEADE affordance as a filtering mechanism: “*The topics feature was excellent, because there are many papers which cover *some* topics I like but also some that I don’t, and this let me pick that out.*”

The constructive feedback received in the qualitative responses illustrate a number of implementation-specific improvements that could be made to Semantic Sanity. The most common category of constructive feedback concerned the quality of terms in the explanations, mentioned by 10 participants. Though we utilized stemming to eliminate these redundancies in each paper explanation, we did not eliminate synonyms from the list of terms. For example, three of the ten participants specifically requested that abbreviations in explanations be removed or linked to full terms. These issues reflect the negative consequences of utilizing 20,000 TF-IDF terms for our explanatory model featurization. In addition, five of these users also stated that the terms were too general. We speculate that the term quality in the explanations negatively impacted the users’ ability to give advice to the model via the LIMEADE affordance.

Similarly, three participants directly addressed what we term the explanation-action tradeoff in the next section, noting that the lack of diversity of terms in the explanations was limiting. One participant commented: “*After a few minutes, almost all the same terms that I had liked were coming up, so there were few new terms for me to thumbs up or down. I think if the system could focus on bringing up relevant papers that have a new term or two to which I can react, that would make the curation even better.*” This suggests tuning the system to favor more explanation diversity even more than we did in our initial implementation.

Interestingly, two users believed that the set of topics surfaced was too restrictive, one thought that the terms were too diverse, and one thought the diversity was a good feature. This provides

some evidence that different users have different preferences for explanation diversity, suggesting that it should perhaps be tuned in a user-specific manner. Additionally, four participants commented on topic annotation strength, all of whom indicated that it was too potent, revealing the importance of empirically evaluating the optimal strength of an update. Based on this feedback, we reduced the annotation strength in our application following the evaluation.

4.7 Feed Quality Revisited Using Log Data: LIMEADE Improves Performance

We also investigated the effect of high-level advice on a different set of users — those who used Semantic Sanity in the wild, rather than as part of laboratory a user study — using the log data of the on-line deployment. Specifically, we compiled a data set of 1,636 rated papers across 30 feeds, where each feed had at least one annotated explanation (the average number of annotations for these feeds was 4.4 terms). We evaluate two recommenders: a baseline ranker that uses only the rated papers, and a LIMEADE ranker that uses both the rated papers and the annotated terms processed by LIMEADE. We evaluate at three different training sizes (2, 5, and 10 labeled papers), and for each feed and size compute the average normalized discounted cumulative gain (NDCG) ranking performance for up to ten sampled training sets. The average of these statistics across feeds is our final evaluation measure.

Table 4 shows that LIMEADE does improve performance over the baseline, but the benefits of the annotated explanations diminish as the number of rated papers increases. The individual differences shown in the table are not statistically significant, but the aggregate performance over all three sizes shows LIMEADE performing significantly better than the baseline (p-value 0.017, two-tailed paired t-test, after Holm-Bonferoni correction). Notably, LIMEADE with 2 and 5 annotated papers performs comparably to the baseline with 5 and 10 annotated papers, respectively, suggesting that during the early phases of giving advice to a recommender, term annotations can serve as a low-cost substitute for assessing additional papers. Experiments with more users and feeds are necessary to further validate these claims.

Number of labeled papers	Baseline	LIMEADE
2	0.884	0.899
5	0.901	0.910
10	0.908	0.913

Table 4. Simulated evaluation of ranking performance (NDCG) based on log data from actual usage in case study 2. LIMEADE improves performance over a baseline that does not use the annotated explanations.

5 DISCUSSION

Evaluating on real-world domains with real human interactions is crucial in order to make progress in human-centered AI broadly, and for advice taking in particular. This section considers broader questions of the effectiveness of human feedback, as well as interactions between the fidelity of explanations and the affordances provided for action.

5.1 When Does High-level Advice Improve Learner Accuracy?

When tested on numerous domains, we obtained positive to indeterminate results about the effectiveness of LIMEADE processing high-level human advice. Does this reflect a weakness in the LIMEADE approach or limitations of our LIME explanations? Or is it intrinsic — perhaps human-interpretable vocabulary is simply too dissimilar to the features learned by modern neural methods for *any* human advice to be useful. Maybe getting more data is the only or the most effective way for humans to help out?

One thing seems clear — in order to answer this question, the research community must conduct more experiments on real world domains, rather than toy domains with artificial confounds, such as Decoy MNIST.

While they only simulate interactions, our image domain experiments (Section 3) reflect actual human judgements about which regions contain the object in question. LIMEADE-processed advice about which regions contained an object significantly improved classifier accuracy in the two shot case. However, when we conducted similar experiments after training the twenty classifiers with ten examples, we found no significant improvement. Perhaps this is because the model had already learned where the objects were located. More likely, it had found the context imparted from background information to be useful in the classification decision. It also may stem from the segmentation algorithm that induced the ‘advice vocabulary’ or perhaps the LIMEADE method weighted examples incorrectly.

While users clearly liked the ability to provide high-level advice and felt it increased their sense of trust and control, we found mixed results with respect to improving ranking accuracy as measured with DCG. Our controlled study over 21 users (Section 4.5.4) found no significant difference between feed accuracy incorporating LIMEADE advice vs. feeds created with simple labeled examples. In contrast, we did find significant improvements stemming from LIMEADE advice in our simulated log study on 30 different users (Section 4.7). The differences could stem from our LIMEADE mechanism, the bi-gram vocabulary chosen as features in the explanatory model, the size of our study, or some other reason.

We strongly believe that much more research should be devoted to this important question. LIMEADE is an important first step, but our paper should be considered a “Call to Action” for more investigation. To this end, we will release the code for LIMEADE and our image experiments, including our modified COCO dataset with the precomputed superpixel vocabulary and corresponding embeddings.

It is also important to contextualize our findings within prior work on advice taking. While studies such as [42] demonstrated a clear improvement in classifier accuracy in the setting of explanatory debugging, other studies have found the opposite. Of particular interest are the results from [1] and [84], which demonstrated that tunability for search and recommendation tasks can negatively impact feed quality when it takes the form of adding or removing terms from the featurization. Likewise, [21] shows that advice taking with interpretable models can lower accuracy. Lastly, [88] is another datapoint that indicates that letting people into the interactive machine learning loop can be problematic. These concerns are especially problematic, given users’ clear expectations that feedback will lead to ML improvement [73]. For this reason, we reiterate that our paper is a “Call to Action” for more research surrounding high-level advice and learner accuracy.

5.2 Exposing the Explanation-Action Tradeoff

Semantic Sanity chooses explanations to display by computing each term’s contribution to the output of the linear model for the given paper, which is equal to the product of the term’s TF-IDF value for the paper with the term’s feature weight in the linear model. The canonical explanation choice is to surface the terms with the highest-magnitude contributions in the linear model [63]; we call this a *greedy* approach. However, comments from early users of our paper recommender indicated that there is a tradeoff between using the greedy approach and providing affordances for feedback, which we call the *explanation-action tradeoff*. In particular, user action on a feature will lead the model to place increasing importance on it and correlated features. With the greedy approach, these terms will begin to dominate the explanations, limiting the number of unique explanation terms and thus opportunities for feedback. For example, ‘thumbs-up’ing the term “fairness” causes papers about fairness to rise in the feed; under the greedy approach, these papers will contain the term “fairness” in their explanations, thereby crowding out new terms for the user to act on. Conversely, a *diversity-biased* strategy would present more unique terms to provide more opportunities for feedback, but would diverge from using the terms with the highest contributions to the explanatory model. When creating an advice-taking system, the explanation-action tradeoff is an important consideration from a user perspective.

Based on the feedback we received, our final implementation of DISPLAY in Semantic Sanity uses a diversity-biased approach that samples explanatory features controlled by a parameter γ . We sample terms proportionally to the magnitude of term contribution, raised to the γ power (higher values of γ result in a more greedy approach; lower values increase diversity). We selected $\gamma = 4$ for our




Fig. 6. A scatter plot showing the number of unique explanation terms in the first page of the feed vs. the number of actions taken by the user in order to give advice to their their feed. Orange dots correspond to diversity-biased explanations currently used in the system. Blue dots correspond to greedy explanations, where the most important terms are surfaced without stochasticity. The size of each dot corresponds to the number of feeds in that bin. Note that greedy explanations (blue) display a stronger negative correlation between unique terms and term annotations than diversity-biased explanations (orange). Thus, the greedy approach limits opportunities for advice taking with topics as the feed curation process evolves, while the diversity-biased approach continues to facilitate advice taking with topics.

implementation. To further reduce term redundancy in each recommended paper’s explanation, we used the Python NLTK PorterStemmer [49] to deduplicate terms with the same stems (e.g., “fair” and “fairness”) from each explanation.

To illustrate the impact of the explanation-action tradeoff and the distinction between our diversity-biased approach and the canonical greedy method, we perform an analysis on the logs of 300 users’ feeds from Semantic Sanity’s online deployment. For each user, we compute (i) the total number of actions the user has taken on displayed explanatory terms, and (ii) the number of unique explanation terms among the latest top eight recommended papers under our diversity-biased DISPLAY implementation. We then repeat (ii) but with DISPLAY with $\gamma = \infty$ to simulate what explanatory terms the users would see today under a greedy approach.

In accordance with the explanation-action tradeoff, we observe in Figure 6 that the number of unique explanation terms (i.e. advice-taking affordances) tends to be lower under a greedy approach. Furthermore, this effect grows stronger as users give advice to their feeds to be increasingly specific to a particular topic.⁸ In contrast, the number of affordances remains relatively constant under our diversity-biased approach. Though some explanation terms with

⁸Figure 6 likely understates the impact of the tradeoff, as users had been exposed to explanations under the diversity-biased approach prior to this analysis. Had they been exposed to explanations under the greedy approach for their entire sessions, we likely would observe an even stronger crowding-out effect.

lower contribution weight are included within the explanatory model, our diversity-biased approach thus successfully mitigates the crowding effect observed with the greedy approach.

The explanation-action tradeoff is related to, but distinct from, the classical *explore-exploit tradeoff* faced by recommender systems and other machine learners [75]. The explore-exploit tradeoff entails deliberately passing up a known reward in the hopes of learning more about the reward structure in order to have better long-term gains. Thus, the explore-exploit tradeoff encourages taking a chance in executing an action in the hopes that it will provide a big reward, leading to frequent execution of the action in the future. The explanation-action tradeoff is similar to the explore-exploit tradeoff, in the sense that it entails deliberately declining to provide the most accurate explanation in the hope that providing an affordance for the user to execute a feedback action will lead to better long term recommendations. However, with the explanation-action tradeoff, even if the system is fortunate when taking a chance by providing a less faithful explanation that successfully solicits user feedback, the system will *never* want to repeat the specific explanation-action in the future. We therefore highlight the explanation-action tradeoff as an important consideration when implementing an advice-taking system.




Fig. 7. Plots showing participants’ Likert scale evaluations of our overall LIMEADE system (left) and the likelihood that they would use our system in the future (right) as functions of the number of topic annotations made when using our LIMEADE system. The red triangles show the median number of annotations for each rating level.

5.3 Decoupling the Effect of Explanations & Advice Taking

Previous studies have shown that users prefer recommendations with explanations over recommendations alone [77, 90]. In our user study, we did not include an “explanations only” baseline, which would have helped to isolate the contribution of explanations in the preference for our LIMEADE system among participants. However, we did analyze the user study results post-hoc to investigate this question. In particular, we studied the results in Tables 2 and 3 in order to assess whether participants’ self-reported preferences for our LIMEADE system over the baseline system correlated with utilization of the LIMEADE affordance for rating topics. The participants who voted LIMEADE higher on trust, transparency, intuitiveness, and confidence in not missing papers performed 5.4, 4.6, -0.5 , and 3.8 more topic annotations, on average, than those who voted the baseline higher, respectively. This suggests that the positive outcomes for those metrics were *not* a result of the explanations alone, but were influenced by the advice-taking affordance of LIMEADE.

In Figure 7, we investigate how the number of topic ratings used by each participant varies as a function of their Likert scale ratings in Table 3. We find that a higher overall rating of our LIMEADE system and a higher self-reported likelihood of using our LIMEADE system in the future are correlated with using more topic annotations (i.e., giving more advice). This indicates that more usage of the LIMEADE affordance correlates with a more positive perception of the LIMEADE system.

6 RELATED WORK

Space precludes a discussion of work on explanation generation; we focus our description of prior work on approaches for incorporating human advice in machine learning models and on approaches for creating pseudo-examples by labeling features. Some work transcends these distinctions, however; Smith-Renner *et al.* [73] show that many users expect that an ML model will improve over time, and that users are frustrated with imperfect AI systems that provide explanations without supporting the ability to receive corresponding feedback.

6.1 Enabling Machine Learners to Take Human Advice: Interpretable Models

Research from interactive machine learning and human-AI interaction has shown the benefits of enabling learning models, including recommender systems and beyond, to take advice from humans [5, 71]. For example, Lou *et al.* [50], Lou *et al.* [51], Caruana *et al.* [17], and Wang *et al.* [86] have demonstrated the value of GAMs and GA^2 Ms, which can be directly modified by humans via the alteration of shape functions. Likewise, Kulesza *et al.* [42] have shown the power of explanatory debugging of models. However, this research has focused on transparent, interpretable models, where the models can be adjusted directly [87]. LIMEADE extends the paradigm of interactive machine learning and advice taking to opaque models. Moreover, these evaluations often focus on user ratings rather than quantitative demonstrations of a model’s improvement via advice taking. As argued by [76], benchmarking and evaluation remain open problems in leveraging explanations in interactive machine learning, one form of advice taking. In our second case study, we directly quantitatively evaluate LIMEADE improvements relative to a baseline.

Recommender systems are a common domain for studying explainability and advice taking due to the feedback loop and interactivity essential to the task of recommendation [2, 16, 31, 48, 59, 78–80, 84, 90]. Some recommender systems take a human’s advice via affordances other than rating content [31]. The majority of these systems enable advice taking in response to a global explanation of the system’s behavior [8, 13–15, 28, 36, 38, 40, 55, 56, 65, 67, 81]. Others enable advice taking in response to instance-level explanations or no explanations at all [1, 30, 43, 83]. The combined affordances of advice taking and explainability can lead to a higher degree of user satisfaction [34]; more trust in and perceived control of the system [20, 34, 58, 82]; and better mental models, without significantly increasing the cognitive load [42, 44, 65]. In contrast to LIMEADE, however, all of this work either relies on interpretable models or implements advice taking in an algorithmic-specific fashion that is not extensible to an arbitrary opaque machine-learned model.

6.2 Enabling Machine Learners to Take Human Advice: Architecture-Specific Models

Other work has explored the extension of advice taking to specific classes of opaque models, such as neural architectures. Like LIMEADE, the methods proposed by both Rieger *et al.* [64] and Ross *et al.* [66] accept human input in response to advice given in terms of an explanatory vocabulary, but their methods are restricted to differentiable models whose gradients can be accessed. Rieger *et al.* modify the loss function in order to incorporate a “contextual decomposition explanation penalization” that encodes a human’s domain knowledge in response to an explanation; and Ross *et al.* modify the loss function through input gradient penalization as a form of regularization. However, both methods are largely evaluated with simulated experiments on small, artificial datasets, where the confounds are often synthetically generated. With DECOY-MNIST, for example, the training data is artificially colored systematically, leading a learner to recognize color rather than shapes. Some methods can effectively adjust the loss function to reflect advice like “ignore color” yielding more robust behavior, but in the real world confounds are much more complex, and it is not clear that these methods generalize well, even for their specific architecture classes.

Liu and Avci [47] present an NLP-specific method that allows a developer to introduce a term into the loss function that can counteract biases exposed by explanations. Specifically, the method can be used to guide a hate-speech detection model away from overly relying on tokens (such as ‘gay’) associated with protected groups. This is different from the feedback accepted by LIMEADE, since it says “Ignore this feature,” rather than “Consider this feature to be positive/negative,” but it is an important type of high-level advice. Liu and Avci tested their approach on both a synthetic and real-world domain, showing modest improvement on the latter. Unlike LIMEADE, however, their approach works only for neural models and has only been tested on an NLP toxicity detection task.

In computer vision models, researchers have created methods for analyzing the behavior of specific neurons, *e.g.* discovering one that produces foliage in a generative model; follow-on research has developed methods for similarly editing these models by rewriting the behavior of those neurons [9, 54]. While impressive, these models are both domain and architecture specific, and require great expertise on the part of the user — far from McCarthy’s dream.

6.3 Enabling Machine Learners to Take Human Advice: Arbitrary Opaque Models

Dasgupta *et al.* [22] consider the problem of teaching an opaque learner whose representation and hypothesis class are unknown. The authors show that by interacting with the black-box learner, a teacher can efficiently find a good set of teaching examples. However, Dasgupta *et al.*’s approach is highly theoretical and assumes a noiseless version-space formulation of learning, where the concept is perfectly learnable. Most importantly, in contrast to LIMEADE, their method doesn’t enable the teacher to provide advice in a high-level language.

Broadly speaking, the advice-taking interaction in LIMEADE is similar to classical human-in-the-loop active learning (AL) [70], which

includes techniques that are applicable to opaque models. However, LIMEADE is distinct from typical AL in that the user is not limited to labeling examples, but can give advice on how the interpretable features should be driving model behavior (which are converted into pseudo-labeled examples using our approach). Further, AL work focuses on algorithms to select informative examples for labeling, whereas LIMEADE creates affordances for feedback on top of explanations that *the user* may choose to act upon.

Closest to our work, Schramowski *et al.* [69] present a method for adding a user into the ML training loop in order to see the AI’s explanations and provide feedback to improve decision making. Like LIMEADE, their method works with an arbitrary opaque classifier, requiring only the ability to add new instances to the training set. Furthermore, they also interpret human feedback in the vocabulary used in an arbitrary, explanatory model, such as that produced by LIME [63]. However, unlike our work, Schramowski *et al.* do not provide a way for the human to explain to the AI why it made a mistake. Instead, they focus on corrections for when the model is “right for the wrong reason.” Like LIMEADE, their method generates pseudo-examples, called “counter-examples,” that are created by altering the selected feature of the explained example in order to reduce confounds (including through randomization, a change to an alternative value, or a substitution with the value for that component appearing in other training examples of the same class). Furthermore, Schramowski *et al.* include only a single experiment to demonstrate their model-agnostic method: on a version of the toy MNIST dataset that was artificially biased to include decoy pixels (Table 1a [69]); their other experiments used a version of Ross *et al.*’s neural-specific loss [66].

6.4 Labeling Features and Creating Pseudo-examples

While canonical methods of feedback involve providing additional labeled examples [85], one approach to semi-supervised learning involves training a machine learner on labeled examples as well as labeled *features* [25, 27, 46, 62, 68, 89]. In the text classification setting, this often takes the form of labeling *n*-gram features. These features are then used to construct pseudo-examples (*e.g.*, documents containing just the labeled *n*-gram itself, labeled according to the feature’s assigned label) or to power methods such as the generalized expectation criteria [89]. LIMEADE extends this semi-supervised approach by translating feature labels in an explanatory model into pseudo-examples for retraining a much more complex opaque model, which is represented using different features.

7 CONCLUSION & A CALL TO ACTION

To be effective partners in a human-AI team, an AI system must be able to not only explain its decisions but also take advice given by humans in terms of that explanation. While interpretable classifiers such as GAMs support explanation-based advice taking, and post-hoc methods such as LIME provide *explanations* for opaque ML models, we present the first method for *updating* an arbitrary opaque model using positive and negative advice given in terms of a high-level vocabulary (such as the featurization of an explanatory model). Furthermore, we are the first to evaluate such a method

on a large number (70) of real-world domains and with user studies. In our first case study, we used LIMEADE to implement advice taking on twenty image classification domains. We showed significant improvement over a strong baseline in the two-shot case. In our second case study, we incorporated LIMEADE into Semantic Sanity, a publicly-available computer science research paper recommender. Our user study over 21 participants demonstrated that users strongly prefer our advice-taking system, lauding perceived control and their sense of trust. While we failed to show improved accuracy of the resulting recommender for these users, as measured with DCG, a study of the long-term logs of 30 different, organic users did show significantly improved NDCG. Furthermore, another log study uncovered a fundamental tension between canonical explanation approaches that greedily select the most influential features and those that provide the best affordances for advice taking.

Much work remains to be done. We hope to develop improved methods for interpreting human advice and better understand when such advice is useful. Experiments using different explanatory vocabularies would also be useful. Additional questions, such as simultaneous advice taking from multiple people in the non-personalized setting, are worth pursuing. Furthermore, developing other forms of advice taking remains a fruitful area for exploration. For example, enabling humans to give advice by adding features, or communicating through natural language or other forms of communication, are understudied challenges. Moreover, understanding various “hyperparameters” surrounding advice taking, such as the proper strength of an update, remains an important question both empirically and theoretically. For example, in the case of recommenders, should the strength of an advice-taking update be personalized? Should it change during different stages of updating a model? While we did not evaluate LIMEADE according to improvements in fairness, robustness, or model compliance, advice taking could be used for these purposes, and another compelling direction of research concerns refining and evaluating advice-taking frameworks in this context. Lastly, it is important to further investigate the entanglement between training data and supervision with respect to advice taking, as described in Section 3.3.

We consider our paper a “Call to Action” for researchers in human-AI interaction to study the advice-taking problem for opaque machine learners. From search & recommendation to image recognition to medical diagnosis, opaque machine learners are ubiquitous. End-users deserve new methods for adjusting these machine learning systems by giving advice in terms of a high-level vocabulary. To aid future research, we will release the code written for our image domain experiments, including our our modified COCO dataset with the precomputed superpixel vocabulary and corresponding embeddings, as well as our functioning implementation of LIMEADE. We hope that this work will contribute to opening a new direction of research in human-AI interaction devoted to this challenging and pressing problem.

ACKNOWLEDGMENTS

The authors would like to thank Sam Skjonsberg and Daniel King for their help in setting up the initial Semantic Sanity prototype;

Matt Latzke for his interface design work; Chelsea Haupt and Sebastian Kohlmeier for their management of platform development; Alex Schokking for his work on scaling Semantic Sanity for public launch; Arman Cohan and Sergey Feldman for providing the neural paper embeddings; Yogi Chandrasekhar and Chris Wilhelm for their guidance with implementing LIMEADE for Semantic Sanity; Chris Wilhelm for his help in implementing functionality necessary for the user study; Ani Kembhavi for his advice regarding LIMEADE in the image domain; and Iz Beltagy, Jonathan Bragg, Krzysztof Gajos, Rao Kambhampati, and Marco Tulio Ribeiro for their helpful feedback. We would also like to thank Andrej Karpathy for his arXiv Sanity Preserver platform, which served as the basis for Semantic Sanity. Lastly, we thank the many users of Semantic Sanity for their feedback. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant DGE-1762114, by ONR grant N00014-18-1-2193, NSF RAPID grant 2040196, the WRF/Cable Professorship, and the Allen Institute for Artificial Intelligence (AI2).

REFERENCES

- [1] Jae-wook Ahn, Peter Brusilovsky, Jonathan Grady, Daqing He, and Sue Yeon Syn. 2007. Open User Profiles for Adaptive News Systems: Help or Harm?. In *WWW '07* (Banff, Alberta, Canada). ACM, 11–20. <https://doi.org/10.1145/1242572.1242575>
- [2] Jae-wook Ahn, Peter Brusilovsky, and Shuguang Han. 2015. Personalized Search: Reconsidering the Value of Open User Models. In *IUI '15* (Atlanta, USA). ACM, 202–212. <https://doi.org/10.1145/2678025.2701410>
- [3] Zeynep Akata, Scott E. Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. 2015. Evaluation of output embeddings for fine-grained image classification. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)*, 2927–2936.
- [4] Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. 2014. Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine* 35, 4 (2014), 105–120. <https://doi.org/10.1609/aimag.v35i4.2513>
- [5] Saleema Amershi, James Fogarty, and Daniel Weld. 2012. Regroup: Interactive Machine Learning for On-demand Group Creation in Social Networks. In *CHI '12* (Austin, USA). ACM, 21–30. <https://doi.org/10.1145/2207676.2207680>
- [6] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fournery, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul Bennett, Kori Inkpen, and et al. 2019. Guidelines for Human-AI Interaction. In *CHI '19* (Glasgow, Scotland). ACM, Article 3. <https://doi.org/10.1145/3290605.3300233>
- [7] Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, and et al. 2018. Construction of the Literature Graph in Semantic Scholar. In *NAACL-HLT '18* (New Orleans, USA). ACL, 84–91. <https://doi.org/10.18653/v1/N18-3011>
- [8] Fedor Bakalov, Marie-Jean Meurs, Birgitta König-Ries, Bahar Sateli, René Witte, Greg Butler, and Adrian Tsang. 2013. An Approach to Controlling User Models and Personalization Effects in Recommender Systems. In *IUI '13* (Santa Monica, CA). ACM, 49–56. <https://doi.org/10.1145/2449396.2449405>
- [9] David Bau, Steven Liu, Tongzhou Wang, Jun-Yan Zhu, and Antonio Torralba. 2020. Rewriting a Deep Generative Model. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [10] Joeran Beel, Bela Gipp, Stefan Langer, and Corinna Breiteringer. 2016. Research-paper recommender systems: a literature survey. *International Journal on Digital Libraries* 17, 4 (2016), 305–338. <https://doi.org/10.1007/s00799-015-0156-0>
- [11] Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A Pretrained Language Model for Scientific Text. In *EMNLP '19*.
- [12] Chandra Bhagavatula, Sergey Feldman, Russell Power, and Waleed Ammar. 2018. Content-Based Citation Recommendation. In *ACL '18*. ACL, New Orleans, LA, 238–251. <https://doi.org/10.18653/v1/N18-1022>
- [13] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. TasteWeights: A Visual Interactive Hybrid Recommender System. In *RecSys '12* (Dublin, Ireland). ACM, 35–42. <https://doi.org/10.1145/2365952.2365964>
- [14] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2013. LinkedVis: exploring social and semantic career recommendations. In *IUI '13* (Santa Monica, USA). ACM, 107. <https://doi.org/10.1145/2449396.2449412>
- [15] Simon Bruns, André Calero Valdez, Christoph Greven, Martina Ziefle, and Ulrik Schroeder. 2015. What Should I Read Next? A Personalized Visual Publication Recommender System. In *Human Interface and the Management of Information. Information and Knowledge in Context*. Springer, 89–100.

- [16] Peter Brusilovsky, Marco de Gemmis, Alexander Felfernig, Pasquale Lops, John O'Donovan, Giovanni Semeraro, and Martijn C. Willemsen. 2020. Interfaces and Human Decision Making for Recommender Systems. In *RecSys '20* (Virtual Event, Brazil). 613–618. <https://doi.org/10.1145/3383313.3411539>
- [17] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-Day Readmission. In *KDD '15* (Sydney, NSW, Australia). ACM, 1721–1730. <https://doi.org/10.1145/2783258.2788613>
- [18] Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. SPECTER: Document-level Representation Learning using Citation-informed Transformers. *arXiv:2004.07180* [cs.CL]
- [19] Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel S. Weld. 2020. SPECTER: Document-level Representation Learning using Citation-informed Transformers. In *ACL*.
- [20] Henriette Cramer, Vanessa Evers, Satyan Ramal, Maarten van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. 2008. The effects of transparency on trust in and acceptance of a content-based art recommender. *UMUAI* 18, 5 (2008), 455. <https://doi.org/10.1007/s11257-008-9051-3>
- [21] Shubhomoy Das, Travis Moore, Weng-Keen Wong, Simone Stumpf, Ian Oberst, Kevin Mcintosh, and Margaret Burnett. 2013. End-User Feature Labeling: Supervised and Semi-Supervised Approaches Based on Locally-Weighted Logistic Regression. *Artif. Intell.* 204 (nov 2013), 56–74. <https://doi.org/10.1016/j.artint.2013.08.003>
- [22] Sanjoy Dasgupta, Daniel Hsu, Stefanos Poulis, and Xiaojin Zhu. 2019. Teaching a black-box learner. In *ICML '19*. PMLR, Long Beach, USA, 1547–1555.
- [23] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- [24] Finale Doshi-Velez and Been Kim. 2017. Towards A Rigorous Science of Interpretable Machine Learning. <http://arxiv.org/abs/1702.08608>
- [25] Gregory Druck, Gideon Mann, and Andrew McCallum. 2008. Learning from Labeled Features Using Generalized Expectation Criteria. In *SIGIR '08* (Singapore, Singapore). 595–602. <https://dl.acm.org/doi/10.1145/1390334.1390436>
- [26] Michael Ekstrand, Praveen Kannan, James Stemper, John Butler, Joseph Konstan, and John Riedl. 2010. Automatically Building Research Reading Lists. In *RecSys '10* (Barcelona, Spain). ACM, 159–166. <https://doi.org/10.1145/1864708.1864740>
- [27] Shantanu Godbole, Abhay Harpale, Sunita Sarawagi, and Soumen Chakrabarti. 2004. Document Classification Through Interactive Supervision of Document and Term Labels. In *PKDD '04*, Jean-François Boulicaut, Floriana Esposito, Fosca Giannotti, and Dino Pedreschi (Eds.). 185–196.
- [28] Brynjar Gretarsson, John O'Donovan, Svetlin Bostandjiev, Christopher Hall, and Tobias Höllerer. 2010. SmallWorlds: Visualizing Social Recommendations. *Computer Graphics Forum* 29, 3 (2010), 833–842. <https://doi.org/10.1111/j.1467-8659.2009.01679.x>
- [29] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. 2018. A Survey of Methods for Explaining Black Box Models. *ACM Comput. Surv.* 51, 5, Article 93 (2018). <https://doi.org/10.1145/3236009>
- [30] F. Maxwell Harper, Funing Xu, Harmanpreet Kaur, Kyle Condiff, Shuo Chang, and Loren Terveen. 2015. Putting Users in Control of their Recommendations. In *RecSys '15* (Vienna, Austria). ACM, 3–10. <https://doi.org/10.1145/2792838.2800179>
- [31] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications* 56 (2016), 9–27. <https://doi.org/10.1016/j.eswa.2016.02.013>
- [32] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep Residual Learning for Image Recognition. *arXiv preprint arXiv:1512.03385* (2015).
- [33] Qi He, Jian Pei, Daniel Kifer, Prasenjit Mitra, and Lee Giles. 2010. Context-aware Citation Recommendation. In *WWW '10* (Raleigh, USA). ACM, 421–430. <https://doi.org/10.1145/1772690.1772734>
- [34] Jonathan Herlocker, Joseph Konstan, and John Riedl. 2000. Explaining collaborative filtering recommendations. In *CSCW '00* (Philadelphia, USA). ACM, 241–250. <https://doi.org/10.1145/358916.358995>
- [35] Sture Holm. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70.
- [36] Yucheng Jin, Nava Tintarev, and Katrien Verbert. 2018. Effects of personal characteristics on music recommender systems with different levels of controllability. In *RecSys '18* (Vancouver, Canada). ACM, 13–21. <https://doi.org/10.1145/3240323.3240358>
- [37] Anshul Kanakia, Zhihong Shen, Darrin Eide, and Kuansan Wang. 2019. A Scalable Hybrid Research Paper Recommender System for Microsoft Academic. In *WWW '19* (San Francisco, USA). ACM, 2893–2899. <https://doi.org/10.1145/3308558.3313700>
- [38] Antti Kangasrääsiö, Dorota Glowacka, and Samuel Kaski. 2015. Improving Controllability and Predictability of Interactive Recommendation Interfaces for Exploratory Search. In *IUI '15* (Atlanta, USA). ACM, 247–251. <https://doi.org/10.1145/2678025.2701371>
- [39] Andrej Karpathy. 2015. Arxiv Sanity Preserver. <http://www.arxiv-sanity.com/>
- [40] Bart Knijnenburg, Niels Reijmer, and Martijn Willemsen. 2011. Each to His Own: How Different Users Call for Different Interaction Methods in Recommender Systems. In *RecSys '11* (Chicago, USA). ACM, 141–148. <https://doi.org/10.1145/2043932.2043960>
- [41] Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. 2020. Concept Bottleneck Models. *ArXiv abs/2007.04612* (2020).
- [42] Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In *IUI '15* (Atlanta, USA). ACM, 126–137. <https://doi.org/10.1145/2678025.2701399>
- [43] Todd Kulesza, Simone Stumpf, Margaret Burnett, and Irwin Kwan. 2012. Tell me more?: the effects of mental model soundness on personalizing an intelligent agent. In *CHI '12* (Austin, USA). ACM, 1. <https://doi.org/10.1145/2207676.2207678>
- [44] Todd Kulesza, Simone Stumpf, Margaret Burnett, Sherry Yang, Irwin Kwan, and Weng-Keen Wong. 2013. Too much, too little, or just right? Ways explanations impact end users' mental models. In *VL/HCC '13*. 3–10. <https://doi.org/10.1109/VLHCC.2013.6645235>
- [45] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. 2014. Microsoft COCO: Common Objects in Context. <http://arxiv.org/abs/1405.0312>
- [46] Bing Liu, Xiaoli Li, Wee Sun Lee, and Philip S. Yu. 2004. Text Classification by Labeling Words. In *AAAI '04* (San Jose, California). 425–430.
- [47] Frederick Liu and Besim Avci. 2019. Incorporating Priors with Feature Attribution on Text Classification. In *ACL '19* (Florence, Italy). ACL, 6274–6283.
- [48] Benedikt Loepp, Catalin-Mihai Barbu, and Jürgen Ziegler. 2016. Interactive Recommending: Framework, State of Research and Future Challenges. In *Proceedings of the 1st Workshop on Engineering Computer-Human Interaction in Recommender Systems*. 3–13. <http://ceur-ws.org/Vol-1705/02-paper.pdf>
- [49] Edward Loper and Steven Bird. 2002. NLTK: The Natural Language Toolkit. In *Proceedings of the ACL Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics*. ACL '02 (Philadelphia, USA).
- [50] Yin Lou, Rich Caruana, and Johannes Gehrke. 2012. Intelligible Models for Classification and Regression. In *KDD '12* (Beijing, China). ACM, 150–158. <https://doi.org/10.1145/2339530.2339556>
- [51] Yin Lou, Rich Caruana, Johannes Gehrke, and Giles Hooker. 2013. Accurate Intelligible Models with Pairwise Interactions. In *KDD '13* (Chicago, USA). ACM, 623–631. <https://doi.org/10.1145/2487575.2487579>
- [52] Scott Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *NIPS '17*. Curran Associates, Inc., 4765–4774. <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>
- [53] John McCarthy. 1968. *Programs with common sense*. 403–418.
- [54] Jesse Mu and Jacob Andreas. 2020. Compositional Explanations of Neurons. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 17153–17163. <https://proceedings.neurips.cc/paper/2020/file/c74956fb38ba48ed6ce977af6727275-Paper.pdf>
- [55] John O'Donovan, Barry Smyth, Brynjar Gretarsson, Svetlin Bostandjiev, and Tobias Höllerer. 2008. PeerChooser: Visual Interactive Recommendation. In *CHI '08* (Florence, Italy). ACM, 1085–1088. <https://doi.org/10.1145/1357054.1357222>
- [56] Denis Parra and Peter Brusilovsky. 2015. User-controllable personalization: A case study with SetFusion. *International Journal of Human-Computer Studies* 78 (2015), 43–67. <https://doi.org/10.1016/j.ijhcs.2015.01.007>
- [57] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [58] Pearl Pu and Li Chen. 2006. Trust Building with Explanation Interfaces. In *IUI '06* (Sydney, AU). ACM, 93–100. <https://doi.org/10.1145/1111449.1111475>
- [59] Pearl Pu, Li Chen, and Rong Hu. 2011. A User-centric Evaluation Framework for Recommender Systems. In *RecSys '11* (Chicago, USA). ACM, 157–164. <https://doi.org/10.1145/2043932.2043962>
- [60] R Core Team. 2018. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. <https://www.R-project.org/>
- [61] Marissa Radensky, Doug Downey, Kyle Lo, Zoran Popovic, and Daniel S Weld. 2022. Exploring the Role of Local and Global Explanations in Recommender Systems. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI EA '22*). Association for Computing Machinery, New York, NY, USA, Article 290, 7 pages. <https://doi.org/10.1145/3491101.3519795>

- [62] Hema Raghavan and James Allan. 2007. An Interactive Algorithm for Asking and Incorporating Feature Feedback into Support Vector Machines. In *SIGIR '07* (Amsterdam, The Netherlands). 79–86. <https://dl.acm.org/doi/10.1145/1277741.1277758>
- [63] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *KDD '16* (San Francisco, USA). ACM, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- [64] Laura Rieger, Chandan Singh, William Murdoch, and Bin Yu. 2020. Interpretations are Useful: Penalizing Explanations to Align Neural Networks with Prior Knowledge. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 8116–8126. <http://proceedings.mlr.press/v119/rieger20a.html>
- [65] Stephanie Rosenthal and Anind Dey. 2010. Towards Maximizing the Accuracy of Human-Labeled Sensor Data. In *IUI '10*. 259–268.
- [66] Andrew Slavin Ross, Michael C. Hughes, and Finale Doshi-Velez. 2017. Right for the Right Reasons: Training Differentiable Models by Constraining their Explanations. In *IJCAI '17*. 2662–2670. <https://doi.org/10.24963/ijcai.2017/371>
- [67] James Schaffer, Tobias Höllerer, and John O'Donovan. 2015. Hypothetical Recommendation: A Study of Interactive Profile Manipulation Behavior for Recommender Systems. <https://www.aaai.org/ocs/index.php/FLAIRS/FLAIRS15/paper/view/10444>
- [68] Robert E. Schapire, Marie Rochery, Mazin G. Rahim, and Narendra Kumar Gupta. 2005. Boosting with prior knowledge for call classification. *IEEE Transactions on Speech and Audio Processing* 13 (2005), 174–181.
- [69] Patrick Schramowski, Wolfgang Stammer, Stefano Teso, Anna Brugger, Xiaoting Shao, Hans-Georg Luigs, Anne-Katrin Mahlein, and Kristian Kersting. 2020. Making deep neural networks right for the right scientific reasons by interacting with their explanations. arXiv:2001.05371 [cs.LG]
- [70] Burr Settles. 2009. Active learning literature survey. (2009).
- [71] Patrice Simard, Saleema Amershi, David Chickering, Alicia Pelton, Soroush Ghosh, Christopher Meek, Gonzalo Ramos, Jina Suh, Johan Verwey, Mo Wang, and John Wernsing. 2017. Machine Teaching: A New Paradigm for Building Machine Learning Systems. <http://arxiv.org/abs/1707.06742>
- [72] Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-June (Paul) Hsu, and Kuansan Wang. 2015. An Overview of Microsoft Academic Service (MAS) and Applications. In *WWW '15* (Florence, Italy). ACM, 243–246. <https://doi.org/10.1145/2740908.2742839>
- [73] Alison Smith-Renner, Ron Fan, M. Keith Birchfield, Tongshuang Sherry Wu, Jordan L. Boyd-Graber, Daniel S. Weld, and Leah Findlater. 2020. No Explainability without Accountability: An Empirical Study of Explanations and Feedback in Interactive ML. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020).
- [74] Karen Sparck Jones. 1988. *A Statistical Interpretation of Term Specificity and Its Application in Retrieval*. Taylor Graham Publishing, GBR, 132–142.
- [75] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement learning: an introduction* (second edition ed.). The MIT Press, Cambridge, Massachusetts.
- [76] Stefano Teso, Öznur Alkan, Wolfgang Stammer, and Elizabeth Daly. 2022. Leveraging Explanations in Interactive Machine Learning: An Overview. <https://doi.org/10.48550/ARXIV.2207.14526>
- [77] Nava Tintarev and Judith Masthoff. 2007. A Survey of Explanations in Recommender Systems. In *2007 IEEE 23rd International Conference on Data Engineering Workshop*. 801–810.
- [78] Nava Tintarev and Judith Masthoff. 2011. Designing and Evaluating Explanations for Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul Kantor (Eds.). Springer, Boston, MA, 479–510. https://doi.org/10.1007/978-0-387-85820-3_15
- [79] Chun-Hua Tsai and Peter Brusilovsky. 2018. Beyond the Ranked List: User-Driven Exploration and Diversification of Social Recommendation. In *IUI '18* (Tokyo, Japan). ACM, 239–250. <https://doi.org/10.1145/3172944.3172959>
- [80] Chun-Hua Tsai and Peter Brusilovsky. 2019. Explaining Recommendations in an Interactive Hybrid Social Recommender. In *IUI '19* (Marina del Ray, USA). ACM, 391–396. <https://doi.org/10.1145/3301275.3302318>
- [81] Chun-Hua Tsai and Peter Brusilovsky. 2020. The effects of controllability and explainability in a social recommender system. *User Modeling and User-adapted Interaction* (2020), 1–37.
- [82] Katrien Verbert, Denis Parra, Peter Brusilovsky, and Erik Duval. 2013. Visualizing Recommendations to Support Exploration, Transparency and Controllability. In *IUI '13* (Santa Monica, USA). ACM, 351–362. <https://doi.org/10.1145/2449396.2449442>
- [83] Jesse Vig, Shilad Sen, and John Riedl. 2012. The Tag Genome: Encoding Community Knowledge to Support Novel Interaction. *ACM TIS* 2, 3 (2012), 1–44. <https://doi.org/10.1145/2362394.2362395>
- [84] Annika Wærn. 2004. User Involvement in Automatic Filtering: An Experimental Study. *UMUAI* 14, 2 (2004), 201–237. <https://doi.org/10.1023/B:USER.0000028984.13876.9b>
- [85] Eric Wallace, Pedro Rodriguez, Shi Feng, Ikuya Yamada, and Jordan Boyd-Graber. 2019. Trick Me If You Can: Human-in-the-Loop Generation of Adversarial Examples for Question Answering. *Transactions of the Association for Computational Linguistics* 7 (2019), 387–401. https://doi.org/10.1162/tacl_a_00279
- [86] Zijie Jay Wang, Alex Kale, Harsha Nori, Peter Stella, Mark E. Nunnally, Duen Horng Chau, Mihaela Vorvoreanu, Jennifer Wortman Vaughan, and Rich Caruana. 2021. GAM Changer: Editing Generalized Additive Models with Interactive Visualization. *ArXiv abs/2112.03245* (2021).
- [87] Daniel Weld and Gagan Bansal. 2019. The Challenge of Crafting Intelligible Intelligence. *CACM* 62, 6 (2019), 70–79. <https://doi.org/10.1145/3282486>
- [88] Tongshuang Wu, Daniel S. Weld, and Jeffrey Heer. 2019. Local Decision Pitfalls in Interactive Machine Learning: An Investigation into Feature Selection in Sentiment Analysis. *ACM Trans. Comput.-Hum. Interact.* 26, 4, Article 24 (jun 2019), 27 pages. <https://doi.org/10.1145/3319616>
- [89] Xiaoyun Wu and Rohini Srihari. 2004. Incorporating Prior Knowledge with Weighted Margin Support Vector Machines. In *KDD '04* (Seattle, USA). 326–333. <https://doi.org/10.1145/1014052.1014089>
- [90] Yongfeng Zhang and Xu Chen. 2018. Explainable Recommendation: A Survey and New Perspectives. *Arxiv* (2018). <http://arxiv.org/abs/1804.11192>