Exploratory Machine Learning with Unknown Unknowns*

Peng Zhao, Yu-Jie Zhang, Zhi-Hua Zhou
National Key Laboratory for Novel Software Technology,
Nanjing University, Nanjing 210023, China
{zhaop, zhangyj, zhouzh}@lamda.nju.edu.cn

Abstract

In conventional supervised learning, a training dataset is given with ground-truth labels from a known label set, and the learned model will classify unseen instances to known labels. In real situations, when the learned models do not work well, learners generally attribute the model failure to the inadequate selection of learning algorithms or the lack of enough labeled training samples. In this paper, we point out that there is an important category of failure, which owes to the fact that there are *unknown* classes in the training data misperceived as other labels, and thus their existence is *unknown* from the given supervision. Such problems of unknown unknown classes can hardly be addressed by common re-selection of algorithms or accumulation of training samples. For this purpose, we propose the *exploratory machine learning*, where in this paradigm once learner encounters unsatisfactory learning performance, she can examine the possibility and, if unknown unknowns really exist, deploy the optimal strategy of feature space augmentation to make unknown classes observable and be enabled for learning. Theoretical analysis and empirical study on both synthetic and real datasets validate the efficacy of our proposal.

1 Introduction

Machine learning has achieved great success in many real-world applications. The success heavily relies on the suitable learning algorithm and sufficient supervised training data. Therefore, facing model failure, the learner would always doubt the inadequate selection of algorithms and the lack of data. A common practice is to try other algorithms or accumulate more data, and such an approach could work effectively when there are no other factors leading to the failure. In this paper, however, we point out that there is an important cause of model failure always ignored before: *unknown unknowns* hidden in the training dataset.

Specifically, we attribute the unknown unknowns to the fact that some training instances of certain *unknown* classes are wrongly perceived as others, and thus appear *unknown* to the learned model with the given supervision. This is always the case when the label space is misspecified due to the insufficient feature information. Consider the task of medical diagnosis, where we need to train a model for community healthcare centers based on their patient records, to help diagnose the cause of a patient with cough and dyspnea. As shown in Figure 1, there are actually three causes: two common ones (*asthma* and *pneumonia*), as well as an unusual one (*lung cancer*) whose diagnosis crucially relies on the computerized tomography (CT) scan device, yet too expensive to purchase. Thus, the community healthcare centers are not likely to diagnose patients with dyspnea as cancer, resulting in that the class of "lung cancer" becomes invisible and hidden in the collected training dataset. As a result, the learned model will be unaware of this unobserved class.

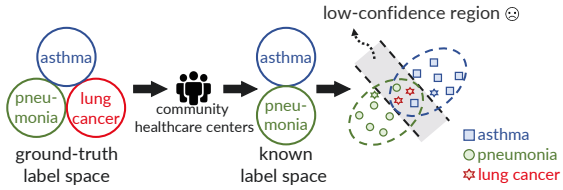

Figure 1: Unknown unknowns in medical diagnosis task

Similar phenomena occur in many other applications. For instance, the trace of a new-type aircraft was mislabeled as old-type ones until performance of aircraft detectors is found poor (i.e., capability of collected signals is inadequate), and the officer suspects that there are new-type aircrafts unknown previously. When feature information is insufficient, there is a high risk to misperceive some classes of training data as others, leading to existence of hidden classes. Especially, hidden classes are sometimes of more interest, like in above two cases. Thus, it is crucial to discover hidden unknown classes and classify known classes well simultaneously.

Conventional supervised learning (SL) cannot obtain a satisfied model when such *unknown unknowns* emerge in the training dataset, even if we could accumulate more data and re-select algorithms exhaustively. The reason lies in that the unknown factors are beyond the expressivity of training data. We thus require new ideas to tackle such unknown unknowns.

2 ExML: A New Learning Framework

The problem we are concerned with is essentially a class of *unknown unknowns*. In fact, how to deal with unknown unknowns is the fundamental question of robust artificial intelligence (Dietterich 2017), and many studies have been

*This research was supported by NSFC (61921006).

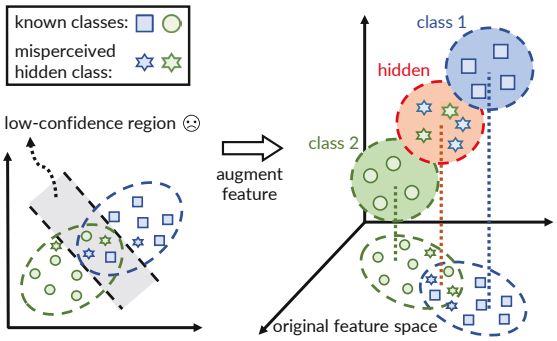

Figure 2: An example illustrates that an informative feature can substantially improve separability of low-confidence samples and make the hidden class distinguishable.

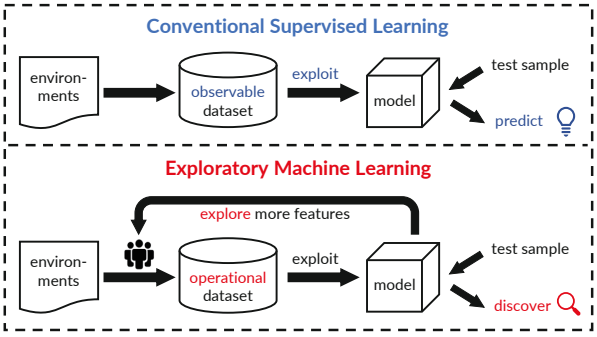

Figure 3: Comparison of two frameworks. SL exploits observable dataset for prediction. ExML explores more features based on operational dataset for discovery of hidden classes.

devoted to addressing various aspects including changing distributions (Pan and Yang 2010; Gama et al. 2014), evolvable features (Hou, Zhang, and Zhou 2017; Hou and Zhou 2018), open categories (Scheirer et al. 2013; Geng, Huang, and Chen 2018), etc. Different from them, we study a new problem setting ignored previously, that is, the training dataset is badly advised by the *incompletely perceived label space* due to the *insufficient feature information*. This problem turns out to be quite challenging, since feature space and label space are entangled and *both* of them are unreliable.

The first challenge is that when the learning performance is undesired we do not know whether the issue is caused by the hidden unknown classes or not. To tackle that, we may accumulate more training data and re-select the learning algorithms. If the model failure persists, we would suspect the existence of unknowns. The second challenge is how to recognize the hidden unknown classes. Notably, it is infeasible to merely pick out instances with low predictive confidence as hidden classes, since we can hardly distinguish (i) instances from hidden classes that suffer from low-confidence predictions owing to the incomplete label space; (ii) instances from known classes that suffer from low-confidence predictions because of insufficient feature information. This characteristic reflects intrinsic hardness of learning with unknown unknowns due to feature deficiency, and thus it is necessary to ask for external feature information.

2.1 Exploratory Machine Learning

To handle unknown unknowns caused by feature deficiency, we resort to the human in the learning loop to interact with environments for enhancing the data collection, more specifically, actively augmenting the feature space. The idea is that when a learned model remains performing poorly even fed with much more data, learner will suspect existence of hidden classes and subsequently seek several candidate features to augment. Figure 2 shows a straightforward example that learner receives a dataset and observes that there are two classes with poor separability, resulting in a noticeable low-confidence region. After a proper feature augmentation, learner will then realize that there exists an additional class hidden in training data previously due to feature deficiency.

Enlightened by the above example, we introduce a new learning framework called *exploratory machine learning* (ExML), which explores more feature information to deal with unknown unknowns caused by feature deficiency. The terminology of exploratory learning is originally raised in the area of education, defined as an approach to teaching and learning that encourages learners to examine and investigate new material with the purpose of discovering relationships between existing background knowledge and unfamiliar content and concepts (Njoo and De Jong 1993; Spector et al. 2014). In the context of machine learning, our proposed framework encourages learners to *examine and investigate the training dataset via exploring new feature information, with the purpose of classifying known classes and discovering potentially hidden classes*. Figure 3 compares the proposed ExML to conventional supervised learning (SL). Conventional SL views the training dataset as an observable representation of environments and exploits it to train a model to predict the label. By contrast, ExML considers the training dataset is *operational*, where learners can examine and investigate the dataset by *exploring* more feature information, and thereby *discover* unknown unknowns due to feature deficiency.

We develop an approach to implement ExML, consisting of three ingredients: rejection model, feature exploration, and model cascade. The rejection model identifies suspicious instances that potentially belong to hidden classes. Feature exploration guides which feature should be explored, and then retrains the model on the augmented feature space. Model cascade allows a layer-by-layer processing to refine the selection of suspicious instances. Theoretical analysis is provided to justify the superiority of the proposed framework. We present empirical evaluations on synthetic data to illustrate the idea and further validate the effectiveness on real datasets.

2.2 Problem Formulation

Training Dataset. The learner receives a training dataset $\widehat{D}_{tr} = \{(\widehat{\mathbf{x}}_i, \widehat{\mathbf{y}}_i)\}_{i=1}^{m}$, where $\widehat{\mathbf{x}}_i \in \widehat{\mathcal{X}} \subseteq \mathbb{R}^d$ is from the *observed* feature space, and $\widehat{\mathbf{y}}_i \in \widehat{\mathcal{Y}}$ is from the *incomplete* label space with N known classes. We consider the binary case for simplicity. Note that there exist training samples that are actually from hidden classes yet wrongly labeled as

known classes due to feature deficiency.

Candidate Features and Cost Budget. Besides the training dataset, the learner can access a set of candidate features $\mathcal{C} = \{c_1, \ldots, c_K\}$, whose values are unknown before acquisition. For the example of medical diagnosis (Figure 1), a feature refers to signals returned from CT scan devices, only available after patients have taken the examination. Moreover, a certain cost will be incurred to acquire any candidate feature for any sample. The learner aims to identify top k informative features from the pool under a given budget B. For convenience, the cost of each acquisition is set as 1 uniformly and the learner desires to find the best feature, i.e., $k = 1$.

Testing Stage. Suppose the learner identifies the best feature as c_i, he/she will then augment the testing sample with this feature, leading to the augmented feature space $\mathcal{X}_i = (\widehat{\mathcal{X}} \cup \mathcal{X}^i) \subseteq \mathbb{R}^{d+1}$ where $\mathcal{X}^i$ is the feature space of c_i. The learned model requires to predict the label of the augmented testing sample, either classified to one of known classes or discovered as the hidden classes (abbrev. hc).

We finally note that several assumptions are introduced for simplicity, with the aim of avoiding distractions of an over-complicated setting and better understanding the essence of the new problem. Actually, our proposal still works when relaxing these assumptions by borrowing well-known techniques such as multi-class rejection (Zhang, Wang, and Qiao 2018), learning with non-uniform cost (Seldin et al. 2014). We emphasize that above aspects are not the current focus. These extensions will be considered as future works.

3 A Practical Approach

Due to the feature deficiency, the learner might be even unaware of the existence of hidden classes based on the observed training data. It is thus necessary to introduce the assumption that *instances with high predictive confidence are safe, i.e., they will be correctly predicted as one of known classes*. Learner will suspect the existence of hidden classes when the learned model performs badly.

We justify the necessity of above assumption. Actually, there are some previous works studying the problem of high-confidence false predictions without considering the issue of feature deficiency (Attenberg, Ipeirotis, and Provost 2015; Lakkaraju et al. 2017), in which there exist some instances wrongly predicted with high confidence. Since the model's performance is highly unreliable, to rectify that, they assume the existence of an oracle providing ground-truth labels for the given query. However, in present of feature deficiency as in our scenario, the problem would not be tractable unless there is an oracle able to provide ground-truth labels based on the insufficient feature information, which turns out to be an even stronger assumption that does not hold in reality generally. We leave high-confidence unknown unknowns due to the insufficient feature as future work to explore.

On the other hand, we emphasize that the introduced assumption does not trivialize the problem because low-predictive instances are not necessarily from hidden classes (as explained at the beginning of Section 2), which necessities more efforts. Following the methodology of ExML (examining the training dataset via exploring new feature information), we design a novel approach, which consists of three components: rejection model, feature exploration, and model cascade. Figure 4 illustrates main procedures, and we will describe details of each component subsequently.

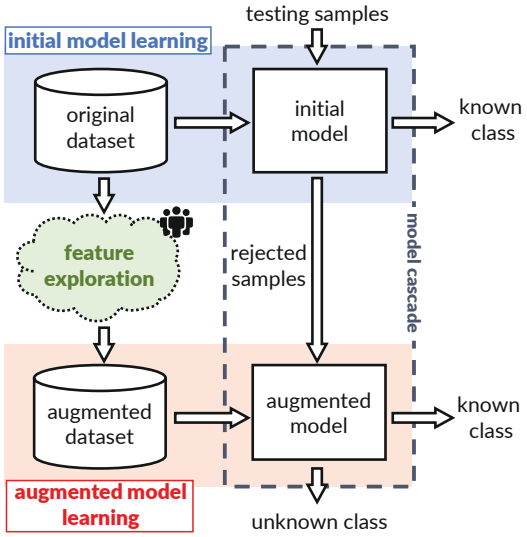

Figure 4: Overall procedure of ExML. Our approach begins with an initial model (blue part), followed by exploring the best candidate feature (green part). Afterwards, a model is retrained based on the augmented dataset, and finally cascaded with the initial model to discover the hidden class (red part).

3.1 Rejection Model

As shown in Figure 4, the learner starts from training an initial model on the original dataset, with capability of identifying low-confidence instances. As emphasized previously (cf. the beginning of Section 2), these low-confidence instances could come from either known or hidden classes, so they are only detected as suspicious and will be further refined.

We realize this goal by the learning with rejection technique (Cortes, DeSalvo, and Mohri 2016), where the learned model will abstain from predicting instances whose maximum conditional probability lower than a given value $1 - \theta$. More precisely, we learn a function pair $f = (h, g)$, where $h : \widehat{\mathcal{X}} \mapsto \mathbb{R}$ is the *predictive* function for known classes and $g : \widehat{\mathcal{X}} \mapsto \mathbb{R}$ is the gate function to *reject* the hidden class. The sample $\widehat{\mathbf{x}}$ is identified to the hidden class if $g(\widehat{\mathbf{x}}) < 0$, and otherwise to the class of $\mathtt{sign}(h(\widehat{\mathbf{x}}))$. Such rejection models can be trained via optimizing the following objective:

$$\min_f \ \mathbb{E}_{(\widehat{\mathbf{x}},\widehat{y})\sim\widehat{\mathcal{D}}}[\ell_{0/1}(f, \widehat{\mathbf{x}}, \widehat{\mathbf{y}}; \theta)], \tag{1}$$

where $\ell_{0/1}(f, \widehat{\mathbf{x}}, \widehat{\mathbf{y}}; \theta) = \mathbb{1}_{\widehat{y}\cdot h(\widehat{\mathbf{x}})<0} \cdot \mathbb{1}_{g(\widehat{\mathbf{x}})>0} + \theta \cdot \mathbb{1}_{g(\widehat{\mathbf{x}})\leq 0}$ is the 0-1 loss of the rejection model f parameterized by the threshold $\theta \in (0, 0.5)$ and $\widehat{\mathcal{D}}$ is the data distribution over $\widehat{\mathcal{X}} \times \widehat{\mathcal{Y}}$. A smaller θ will lead to more rejections but a higher predictive accuracy on known classes. To tackle the difficulty of non-convex optimization arising from the indicator function, Cortes, DeSalvo, and Mohri (2016) introduce the following calibrated surrogate loss function

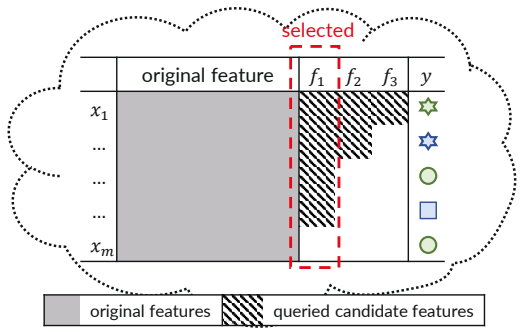

Figure 5: Feature Exploration

$\ell_{surr} := \ell_{surr}(f, \widehat{\mathbf{x}}, \widehat{y}; \theta)$ defined as

$$\ell_{surr} = \max\Big\{1 + \frac{1}{2}\big(g(\widehat{\mathbf{x}}) - \widehat{y} \cdot h(\widehat{\mathbf{x}})\big), \theta\big(1 - \frac{g(\widehat{\mathbf{x}})}{1-2\theta}\big), 0\Big\}$$

to approximate the original $\ell_{0/1}$ loss. Since the distribution is unknown we cannot directly measure the risk, we choose the model that minimizes the empirical risk:

$$\min_{f \in \mathbb{H} \times \mathbb{H}} \frac{1}{m} \sum_{i=1}^{m} \ell_{surr}(f, \widehat{\mathbf{x}}_i, \widehat{y}_i; \theta) + C_h \|h\|_{\mathbb{H}}^2 + C_g \|g\|_{\mathbb{H}}^2, \quad (2)$$

where C_h and C_g are regularization parameters, and $\mathbb{H}$ is the RKHS induced by kernel $K : \widehat{\mathcal{X}} \times \widehat{\mathcal{X}} \mapsto \mathbb{R}$. By the representer theorem (Schölkopf and Smola 2002), the optimizer of (2) is in the form of $h(\widehat{\mathbf{x}}) = \sum_{i=1}^{m} u_i K(\widehat{\mathbf{x}}, \widehat{\mathbf{x}}_i)$ and $g(\widehat{\mathbf{x}}) = \sum_{i=1}^{m} w_i K(\widehat{\mathbf{x}}, \widehat{\mathbf{x}}_i)$, where u_i and w_i are coefficients to learn. So (2) can be reformulated as quadratic programming and solved efficiently. For more details we refer the reader to the seminal work of Cortes, DeSalvo, and Mohri (2016).

3.2 Feature Exploration

If the initial model is unqualified (for instance, it rejects too many samples for achieving desired accuracy), the learner will suspect the existence of hidden classes and explore new features to augment. In our setting, the learner requires to select the best feature from K candidates and retrain a model based on the augmented data, as shown in Figure 5.

We emphasize that the conventional feature selection is not feasible here, because it requires to know values of candidate features, while these values are unknown before acquisitions. To address the challenge, we propose a novel procedure—*feature exploration*—to adaptively identify the most informative feature under the cost budget, *without* requiring feature values in advance. To address the issue, there are two fundamental questions to answer:

(1) how to measure the quality of candidate features?

(2) how to allocate the budget to identify the best feature?

We answer the above two questions in the following.

Feature quality measure. Denote by $\mathcal{D}_i$ the data distribution over $\mathcal{X}_i \times \widehat{\mathcal{Y}}$, where $\mathcal{X}_i$ is the augmented feature space of the i-th candidate feature. We use the *Bayes risk* on $\mathcal{D}_i$ as feature quality measure:

$$R_i^* = R_i(f_i^*) = \min_f \mathbb{E}_{(\mathbf{x}, \widehat{\mathbf{y}}) \sim \mathcal{D}_i} \big[\ell_{0/1}(f, \mathbf{x}, \widehat{\mathbf{y}}; \theta)\big], \quad (3)$$

where $R_i(f)$ is the expected $0/1$ risk of function f over $\mathcal{D}_i$, and f_i^* minimizes $R_i(f)$ over all measurable functions. The Bayes risk essentially reflects the minimal error that any rejection model can attain on the augmented data distribution, whose value will be smaller when the selected augmented feature improves the separability more significantly (and is believed more informative).

Due to the inaccessibility of the underlying distribution $\mathcal{D}_i$, we approximate the Bayes risk by its empirical version over the augmented data $D_i = \{(\mathbf{x}_j, \widehat{\mathbf{y}}_j)\}_{j=1}^{n_i} \sim \mathcal{D}_i$,

$$\widehat{R}_{D_i} = \widehat{R}_i(\widehat{f}_i) = \sum\nolimits_{j=1}^{n_i} \ell_{0/1}(\widehat{f}_i, \mathbf{x}_j, \widehat{y}_j; \theta), \quad (4)$$

where $\mathbf{x}_j \in \mathcal{X}_i, \widehat{\mathbf{y}}_j \in \widehat{\mathcal{Y}}$, and $\widehat{f}_i$ is the rejection model learned by empirical risk minimization over surrogate loss (2) on augmented dataset D_i. Based on the feature quality measure (3) and its empirical version (4), we now introduce the budget allocation strategy to identify the best candidate feature.

Budget allocation strategy. Without loss of generality, suppose features are sorted according to their quality, i.e., $R_1^* \leq \cdots \leq R_K^*$. Our goal is to identify the best feature within the limited budget, and meanwhile the model retrained on augmented data should have good generalization ability.

We first propose the uniform allocation strategy as follows, under the guidance of criterion (3).

Uniform Allocation For each candidate feature c_i, $i \in [K]$, learner allocates $\lfloor B/K \rfloor$ budget and obtains an augmented dataset D_i. So we can compute the empirical feature measure by (4), and select the feature with the smallest risk. The above strategy is simple yet effective, which can provably identify the best feature with high probability (Theorem 1).

Median Elimination We further propose another variant inspired by the bandit theory to improve the budget allocation efficiency. Specifically, we adopt the technique of *median elimination* (ME) (Even-Dar, Mannor, and Mansour 2006), which removes one half of poor candidate features after every iteration and only the best one remains in the end. As a result, the algorithm can avoid allocating too many budgets on poor features. More specifically, the elimination proceeds in $T = \lceil \log_2 K \rceil$ episodes, in each episode, $\lfloor B/T \rfloor$ budget is allocated uniformly to all remaining candidate features, and the learner could query their values for updating the corresponding augmented datasets D_i. Then, the score $\widehat{R}_{D_i}$ is calculated on the current augmented datasets D_i and the half features with high $\widehat{R}_{D_i}$ are eliminated. In the last, only one candidate feature i_s will be left and its augmented dataset D_{i_s} contains around $\lfloor B/\log K \rfloor$ samples, which is the largest one among all the candidate features. Algorithm details are presented in the full version (Zhao, Zhang, and Zhou 2021).

As shown in Figure 5, poor features are eliminated earlier, budget left for the selected feature is thus improved from $\lfloor B/K \rfloor$ to $\lfloor B/\log K \rfloor$ by ME, which ensures better generalization ability of the learned model. Meanwhile, median elimination can explore the best candidate feature more efficiently than uniform allocation, as shown in the bandit theory (Even-Dar, Mannor, and Mansour 2006). We finally remark that our paper currently focuses on the best feature, and the framework is ready for identifying the top k features ($k > 1$) by

introducing more sophisticated techniques (Kalyanakrishnan et al. 2012; Chen, Li, and Qiao 2017).

3.3 Model Cascade

After feature exploration, learner will retrain a model on augmented data. Considering that the augmented model might not always be better than the initial model, particularly when the budget is not enough or candidate features are not quite informative, we propose the *model cascade* mechanism to cascade the augmented model with the initial one. Concretely, high-confidence predictions are accepted in the initial model, the rest suspicious are passed to the next layer for feature exploration, those augmented samples with high confidence will be accepted by the augmented model, and the remaining suspicious continue to the next layer for further refinements.

Essentially, our approach can be regarded as a *layer-by-layer processing for identifying instances of hidden classes*, and the procedures can be stopped until human discovers remaining suspicious are indeed with certain hidden structures. For simplicity, we only implement a two-layer architecture.

4 Theoretical Analysis

This section presents theoretical results. We first investigate the attainable excess risk of supervised learning, supposing that the best feature were *known* in advance. Then, we provide the result of ExML to demonstrate the effectiveness of our proposed criterion and budget allocation strategies.

For each candidate feature c_i, we denote the corresponding hypothesis space as $\mathcal{H}_i, \mathcal{G}_i = \{\mathbf{x} \mapsto \langle \mathbf{w}, \Phi_i(\mathbf{x}) \rangle \mid \|\mathbf{w}\|_{\mathbb{H}_i} \leq \Lambda_i\}$, where Φ_i and $\mathbb{H}_i$ are induced feature mapping and RKHS of kernel K_i in the augmented feature space.

Supervised learning with known best feature. Suppose the best feature were known in advance, we could obtain B samples augmented with this particular feature. Let f_{SL} be the model learned by supervised learning via minimizing (2). From learning theory literatures (Bousquet, Boucheron, and Lugosi 2003; Cortes, DeSalvo, and Mohri 2016), for any $\delta > 0$, with probability at least $1 - \delta$, we have

$$R_1(f_{\text{SL}}) - R_1^* \leq \mathcal{O}\left(\sqrt{\frac{(\kappa_1 \Lambda_1)^2}{B}} + \sqrt{\frac{\log(1/\delta)}{2B}}\right) + R_{ap}, \tag{5}$$

where $R_{ap} = C_\theta \left(\inf_{f \in \mathcal{H}_1 \times \mathcal{G}_1} R_1^{surr}(f) - \inf_f R_1^{surr}(f)\right)$ is the approximation error measuring how well hypothesis spaces $\mathcal{H}_1, \mathcal{G}_1$ approach the target, in terms of the expected surrogate risk $R_1^{surr}(f) = \mathbb{E}_{(\mathbf{x}, \widehat{\mathbf{y}}) \sim \mathcal{D}_1}[\ell_{surr}(f, \mathbf{x}, \widehat{\mathbf{y}}; \theta)]$. The constant factor is $C_\theta = 1/((1-\theta) \cdot (1-2\theta))$.

The above result theoretically reveals that if the best feature were *known* in advance, the excess risk of supervised learning would converge to the inevitable approximate error in the rate of $\mathcal{O}(1/\sqrt{B})$, with a given feature budget B.

Exploratory learning with unknown best feature. In our setting, the best feature is unfortunately unknown ahead of time. More importantly, since values of K candidate features are unavailable, it is *infeasible* to perform the feature selection. We show that by means of ExML (feature exploration), the excess risk also converges, in the rate of $\mathcal{O}(\sqrt{K/B})$, yet *without* requiring to know the best feature.

Theorem 1. *Let c_{i_s} be the identified feature and $\widehat{f}_{i_s}$ be the augmented model returned by ExML with uniform allocation. Then, with probability at least $1 - \delta$, we have*

$$R_{i_s}(\widehat{f}_{i_s}) - R_1^* \leq \mathcal{O}\left(\sqrt{\frac{(\kappa\Lambda)^2}{\lfloor B/K \rfloor}} + \sqrt{\frac{\log(3/\delta)}{2\lfloor B/K \rfloor}}\right) + R_{ap}, \tag{6}$$

where $\Lambda = \max_{i \in [K]} \Lambda_i$, $\kappa = \max_{i \in [K]} \sup_{\mathbf{x} \in \mathcal{X}_i} K_i(\mathbf{x}, \mathbf{x})$.

Remark. Comparing the excess risk bounds of (5) and (6), we can observe that ExML exhibits a similar convergence tendency to SL with *known* best feature, yet *without* requiring to know the best feature. An extra $\sqrt{K}$ times factor is paid for exploration of the best feature. We note that under certain mild technical assumptions, the dependence can be further reduced to $\sqrt{\log K}$ by median elimination (Even-Dar, Mannor, and Mansour 2006), as poor candidate features have been removed in the earlier episodes.

5 Experiments

In this section, we conduct experiments to examine empirical performance of the proposed exploratory machine learning (ExML). Specifically, we provide evaluations on synthetic data for visualizing the superiority of ExML to conventional supervised learning in handling unknown unknowns. Then, we report results on real-world datasets to demonstrate the effectiveness of the overall method, as well as the usefulness of feature exploration and model cascade modules. In all experiments, we denote by $B = b \cdot mk$ the feature exploration budget, where m is number of training samples, K is number of candidate features, $b \in [0, 1]$ is the budget ratio.

5.1 Synthetic Data for Illustration

We first illustrate the advantage of exploratory machine learning over conventional supervised learning in discovering the hidden classes on the synthetic data.

Setting. Following the illustrative example in Figure 1, we generate a 3-dim dataset containing 3 classes, whose ground-truth distribution is shown as Figure 6(a). However, as shown in Figures 6(b), only the first two dimensions are observable in the training stage, resulting in a hidden class (hc) located in the intersection area of known classes (kc1 and kc2).

Specifically, we generate instances of each class from Gaussian distributions. Means and variances are $[-a, 0, -z]$ and $\sigma \cdot \mathbf{I}_{3\times 3}$ for the first known class, $[a, 0, z]$ and $\sigma \cdot \mathbf{I}_{3\times 3}$ for the second known class as well as $[0, 0, 0]$ and $\sigma/2 \cdot \mathbf{I}_{3\times 3}$ for the hidden class, where $\mathbf{I}_{3\times 3}$ is a 3×3 identity matrix. We set $\sigma = 3a$ and $z = 5a$. In the training stage, the third-dim is unobservable and the hidden class (hc) is randomly labeled as another two. Each class contains 100 instances in the training data. Besides, we generate 9 candidate features in various qualities, whose angle to the horizon varies from 10° to 90°, the larger the better. Figure 6(c) plots the augmented feature space via t-SNE. The budget ratio is $b = 20\%$. In the testing stage, the learner requires to predict on the 3-dim data, where the third dimension is the selected candidate features.

Contenders. We compare ExML to SL (with rejection model). For all rejection models, we employ the Gaussian

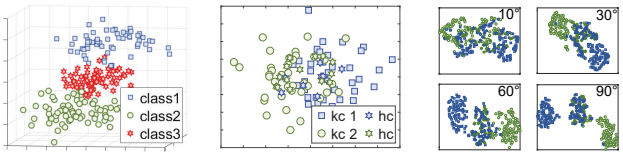

(a) ground-truth (b) training data (c) candidates

Figure 6: Visualization of synthetic data: (a) ground-truth distribution; (b) training data (only first two dimensions are observable); (c) t-SNE of candidate features with various qualities (a larger angle implies a better feature quality).

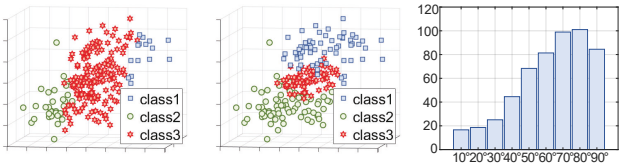

(a) SL (b) ExML (c) allocation

Figure 7: Visualization of results: (a) SL; b) ExML; (c) budget allocation of ExML with median elimination.

kernel $K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2/\gamma)$ with bandwidth $\gamma = \text{median}_{\mathbf{x}_i,\mathbf{x}_j \in D}(\|\mathbf{x}_i - \mathbf{x}_j\|_2^2)$ and set C_h, C_g to 1.

- **SL** trains the rejection model (Cortes, DeSalvo, and Mohri 2016) with the given dataset on the original feature space, following the paradigm of conventional supervised learning. The threshold θ is choose as one achieving best accuracy on the testing data from the pool $[0.1, 0.2, 0.3, 0.4]$.
- **ExML** is our proposal with cascade models and using median elimination for feature exploration. The threshold for the initial rejection model is selected by cross validation to ensure 95% accuracy on high-confidence samples. The threshold θ for the augmented rejection model is choose as one achieving best accuracy on the testing data from the pool $[0.1, 0.2, 0.3, 0.4]$. The budget ratio is 20%.

Results. We first conduct SL to train a rejection model based on the 2-dim training data, and then perform ExML to actively augment the feature within the budget to discover unknown unknowns. Figures 7(a) and 7(b) plot the results, demonstrating a substantial advantage of ExML over SL in discovering the hidden class and predicting known classes. Furthermore, Figure 7(c) reports budget allocation of each candidate feature over 50 times repetition. We can see that the allocation clearly concentrates to more informative features (with larger angles), which validates the effectiveness of median elimination for the best feature exploration.

5.2 Benchmark Data for Evaluation

Dataset and Setting. We further evaluate on a UCI benchmark dataset *Mfeat* (van Breukelen et al. 1998), which is a multi-view dataset[1] containing 2000 samples and 6 views of features extracted by various methods. Their semantic information and statistics are:

- Fac: profile correlations, 216-dim;
- Pix: pixel averages in 2×3 windows, 240-dim;
- Kar: Karhunen-Love coefficients, 64-dim;
- Zer: Zernike moments, 47-dim;
- Fou: Fourier coefficients of the character shapes, 76-dim;
- Zer: morphological features, 6-dim.

The domain knowledge sorts the features by their quality as: Fac $>$ Pix $>$ Kar $>$ Zer $>$ Fou $>$ Mor, in a descending order.

In the training stage, we randomly sample 600 instances to form the labeled training data. This procedure repeats 10 times to generate different configurations. Since Mfeat is a multi-class dataset, we randomly sample 5 configurations to convert it into the binary classification task, where each known class and hidden class contain three original classes, and the instances from the hidden class are randomly mislabeled as one of known classes. There are in total 50 random configurations for training. As for the candidate features, each one of six views (features) is taken as original feature and the rest are prepared in the candidate set. Before training, we normalize all the features to the range $[0, 1]$. We evaluate all contenders on the testing data containing 1400 instances.

Contenders. Apart from SL, we include two ExML variants: $\text{ExML}_{\text{csd}}^{\text{UA}}$ and $\text{ExML}_{\text{aug}}^{\text{ME}}$ for ablation studies. Here aug/csd denotes the final model is only the augmented or cascaded with the initial model; UA/ME refers to feature exploration by uniform allocation or median elimination.

- $\textbf{ExML}_{\textbf{csd}}^{\textbf{UA}}$ is our proposal with cascade model and using *uniform allocation* for feature exploration.
- $\textbf{ExML}_{\textbf{aug}}^{\textbf{ME}}$ is our proposal *without* cascade model and using median elimination for feature exploration.

All ExML-type methods use the same parameters. SL and ExML are configured by the same setting as those in synthetic experiments. The budget ratio b varies from 10% to 30%.

Measure. We measure the performance of all the methods by the classification. Additionally, we introduce the *recall* to measure the effectiveness of feature exploration, defined as the ratio of the number of cases when identified feature is one of its top 2 features to the total number.

- **Accuracy**: the mean and standard deviation of the predictive accuracy on testing dataset over 50 configurations, where the true label of hidden classes are observable.
- **Recall**: the ratio of the number of cases when identified feature is one of top 2 features to the total number, where features quality is measured by the accuracy of augmented model trained on whole data with this particular feature.

Results. Table 1 reports mean and std of the predictive accuracy, and all features are sorted in descending order by their quality. We first compare SL to (variants of) ExML. When the original features are in high quality (Kar, Pix, Fac), SL could achieve favorable performance and there is no need to explore new features. However, in the case where uninformative original features are provided, which is of more interest for ExML, SL degenerates severely and $\text{ExML}_{\text{aug}}^{\text{ME}}$ (the single ExML model without model cascade) achieves better performance even with the limited budget. Besides, from the last column, we can see that informative candidates (top 2) are selected to strengthen the poor original features, which validates the effectiveness of the proposed budget allocation

[1] http://archive.ics.uci.edu/ml/datasets/Multiple+Features

Fea.&Budget		SL	$\text{ExML}^{\text{ME}}_{\text{aug}}$	$\text{ExML}^{\text{UA}}_{\text{csd}}$	ExML	Recall
Fac	10%	**93.4±1.7**	71.8±9.6	92.4±2.8	92.4±2.8	48%
	20%	**93.4±1.7**	82.3±7.5	92.0±3.3	92.0±3.3	46%
	30%	**93.4±1.7**	89.3±4.7	92.2±3.3	92.5±2.9	44%
Pix	10%	**92.2±2.5**	70.5±8.3	90.5±6.3	90.6±6.3	58%
	20%	**92.2±2.5**	81.7±7.2	90.8±6.2	90.9±6.1	54%
	30%	**92.2±2.5**	88.7±4.1	90.5±5.7	**91.8±4.3**	68%
Kar	10%	**86.9±3.4**	70.3±10	85.6±4.9	85.9±4.9	56%
	20%	**86.9±3.4**	81.5±6.9	85.2±5.5	**86.5±4.8**	54%
	30%	86.9±3.4	86.0±5.4	86.5±4.7	**88.2±3.6**	56%
Zer	10%	73.8±8.8	69.6±11	73.0±10	**76.2±8.5**	82%
	20%	73.8±8.8	**80.9±8.0**	77.3±7.9	**81.7±7.3**	82%
	30%	73.8±8.8	**86.1±5.5**	81.1±6.8	**86.3±5.0**	86%
Fou	10%	68.7±9.1	69.4±9.7	68.9±12	**75.9±8.8**	82%
	20%	68.7±9.1	82.1±6.5	77.9±8.3	**85.0±4.4**	88%
	30%	68.7±9.1	**89.9±3.7**	82.5±5.2	**89.4±3.9**	92%
Mor	10%	57.5±15	69.1±11	66.6±13	**71.1±11**	80%
	20%	57.5±15	**79.6±10**	73.6±8.9	**79.8±9.9**	84%
	30%	57.5±15	**87.4±7.3**	78.3±9.0	**87.0±7.1**	90%

Table 1: Evaluation on Mfeat dataset. Features are sorted by descending qualities. Bold font indicates algorithms outperform others (paired t-test at 5% significance level).

strategy (namely, the median elimination mechanism).

Since the $\text{ExML}^{\text{ME}}_{\text{aug}}$ is not guaranteed to outperform SL, particularly with the limited budget on poor candidate features, we propose the cascade structure. Actually, ExML approach (aka, $\text{ExML}^{\text{ME}}_{\text{csd}}$) achieves roughly *best-of-two-worlds* performance, in the sense that it is basically no worse or even better than the best of SL and $\text{ExML}^{\text{ME}}_{\text{aug}}$. It turns out that even $\text{ExML}^{\text{UA}}_{\text{csd}}$ could behave better than $\text{ExML}^{\text{ME}}_{\text{aug}}$. These results validate the effectiveness of the model cascade component.

5.3 Real Data of Activities Recognition

We additionally examine the effectiveness on a real-world dataset called *RealDisp*[2], which is an activities recognition task (Baños et al. 2012). There are 9 on-body sensors used to capture various actions of participants. Each sensor is placed on different parts of the body and provides 13-dimensional features including 3-dim from acceleration, 3-dim from gyro, 3-dim from magnetic field orientation and another 4-dim from quaternions. Hence, we have 117 features in total.

Dataset. Three types of actions (*walking*, *running*, and *jogging*) are included to form the dataset containing 2000 instances, where 30% of them are used for training and the remaining 70% for testing. In the training data, one sensor is deployed and the class of jogging is randomly misperceived as walking or running. The learner would explore the rest eight candidate features to discover the unknown unknowns. Thus, there are 9 partitions, and each is repeated for 10 times by sampling the training instances randomly.

Results. Figure 8 shows the mean and std of accuracy, our approach ExML (aka, $\text{ExML}^{\text{ME}}_{\text{csd}}$) outperforms others, validating

[2]http://archive.ics.uci.edu/ml/datasets/REALDISP+Activity+Recognition+Dataset

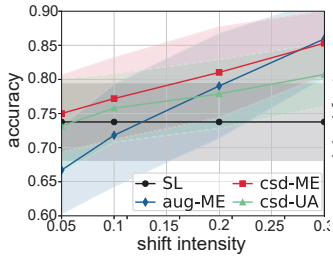

Figure 8: Performance comparisons of all contenders.

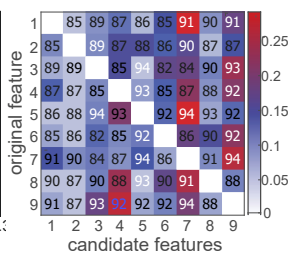

Figure 9: Budget allocation (median elimination).

the efficacy of our proposal. In addition, Figure 9 illustrates the budget allocation when the budget ratio $b = 30\%$. The i-th row denotes the scenario when the i-th sensor is the original feature, and patches with colors indicate the fraction of budget allocated to each candidate feature. The number above a patch means the attainable accuracy of the model trained on the whole training dataset with the particular feature. We highlight the top two candidate features of each row in white, and use blue color to indicate selected feature is not in top two. The results show that ExML with median elimination can select the top two informative features to augment for all the original sensors. The only exception is the 9-th sensor, but quality of the selected feature (91.8%) does not deviate too much from the best one (93.6%). These results reflect the effectiveness of our feature exploration strategy.

6 Conclusion

In this paper, we identify that aside from the inadequate selection of learning algorithms or the lack of enough labeled training samples, unknown unknowns could also lead to the model failure. In particular, we are concerned with the scenario where some instances in the training dataset belong to an unknown hidden class but are wrongly perceived as known classes, due to the insufficient feature information. To address this issue, we propose the *exploratory machine learning* (ExML) to encourage the learner to examine and investigate the training dataset by exploring more features to discover potentially hidden classes. Following this idea, we design an approach consisting of three procedures: rejection model, feature exploration, and model cascade. By leveraging techniques from bandit theory, we prove the rationale and efficacy of the feature exploration procedure. Experiments validate the effectiveness of our approach.

There remain many directions for future investigations. For instance, as mentioned in Section 2.2, we can borrow more advanced techniques to further relax some model assumptions introduced in the current work (such as binary known classes, uniform cost, best feature exploration, etc). In particular, it is interesting to consider a personalized cost for each candidate feature, since we usually need to pay a higher price to obtain more informative features in real-world applications. Moreover, in addition to the feature exploration proposed in this paper, we argue that there are many other possibilities for ExML to deal with unknown unknowns, by means of adaptive interactions with environments.

References

Attenberg, J.; Ipeirotis, P.; and Provost, F. 2015. Beat the Machine: Challenging Humans to Find a Predictive Model's Unknown Unknowns. *ACM Journal of Data and Information Quality* 1–17.

Baños, O.; Damas, M.; Pomares, H.; Rojas, I.; Tóth, M. A.; and Amft, O. 2012. A benchmark dataset to evaluate sensor displacement in activity recognition. In *Proceedings of 12th ACM Conference on Ubiquitous Computing*, 1026–1035.

Bousquet, O.; Boucheron, S.; and Lugosi, G. 2003. Introduction to Statistical Learning Theory. In *Advanced Lectures on Machine Learning (Machine Learning Summer Schools 2003)*, 169–207.

Chen, L.; Li, J.; and Qiao, M. 2017. Nearly Instance Optimal Sample Complexity Bounds for Top-k Arm Selection. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 101–110.

Cortes, C.; DeSalvo, G.; and Mohri, M. 2016. Learning with Rejection. In *Proceedings of International Conference on Algorithmic Learning Theory*, 67–82.

Dietterich, T. G. 2017. Steps Toward Robust Artificial Intelligence. *AI Magazine* 3–24.

Even-Dar, E.; Mannor, S.; and Mansour, Y. 2006. Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems. *Journal of Machine Learning Research* 7: 1079–1105.

Gama, J.; Zliobaite, I.; Bifet, A.; Pechenizkiy, M.; and Bouchachia, A. 2014. A survey on concept drift adaptation. *ACM Computing Surveys* 46(4): 44:1–44:37.

Geng, C.; Huang, S.-J.; and Chen, S. 2018. Recent Advances in Open Set Recognition: A Survey. *ArXiv preprint* arXiv:1811.08581.

Hou, B.-J.; Zhang, L.; and Zhou, Z.-H. 2017. Learning with Feature Evolvable Streams. In *Advances in Neural Information Processing Systems 30*, 1417–1427.

Hou, C.; and Zhou, Z.-H. 2018. One-Pass Learning with Incremental and Decremental Features. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40(11): 2776–2792.

Kalyanakrishnan, S.; Tewari, A.; Auer, P.; and Stone, P. 2012. PAC Subset Selection in Stochastic Multi-armed Bandits. In *Proceedings of the 29th International Conference on Machine Learning*, 227–234.

Lakkaraju, H.; Kamar, E.; Caruana, R.; and Horvitz, E. 2017. Identifying Unknown Unknowns in the Open World: Representations and Policies for Guided Exploration. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2124–2132.

Njoo, M.; and De Jong, T. 1993. Exploratory learning with a computer simulation for control theory: Learning processes and instructional support. *Journal of research in science teaching* 821–844.

Pan, S. J.; and Yang, Q. 2010. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering* 1345–1359.

Scheirer, W. J.; de Rezende Rocha, A.; Sapkota, A.; and Boult, T. E. 2013. Toward Open Set Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1757–1772.

Schölkopf, B.; and Smola, A. J. 2002. *Learning with Kernels: support vector machines, regularization, optimization, and beyond*. The MIT Press.

Seldin, Y.; Bartlett, P. L.; Crammer, K.; and Abbasi-Yadkori, Y. 2014. Prediction with Limited Advice and Multiarmed Bandits with Paid Observations. In *Proceedings of the 31th International Conference on Machine Learning*, 280–287.

Spector, J. M.; Merrill, M. D.; Elen, J.; and Bishop, M. 2014. *Handbook of Research on Educational Communications and Technology*. Springer, fourth edition.

van Breukelen, M.; Duin, R. P. W.; Tax, D. M. J.; and den Hartog, J. E. 1998. Handwritten digit recognition by combined classifiers. *Kybernetika* 381–386.

Zhang, C.; Wang, W.; and Qiao, X. 2018. On reject and refine options in multicategory classification. *Journal of the American Statistical Association* 730–745.

Zhao, P.; Zhang, Y.-J.; and Zhou, Z.-H. 2021. Exploratory Machine Learning with Unknown Unknowns. *ArXiv preprint* arXiv:2002.01605.