

META-LEARNING TRANSFERABLE ACTIVE LEARNING POLICIES BY DEEP REINFORCEMENT LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Active learning (AL) aims to enable training high performance classifiers with low annotation cost by predicting which subset of unlabelled instances would be most beneficial to label. The importance of AL has motivated extensive research, proposing a wide variety of manually designed AL algorithms with diverse theoretical and intuitive motivations. In contrast to this body of research, we propose to treat active learning algorithm design as a meta-learning problem and learn the best criterion from data. We model an active learning algorithm as a deep neural network that inputs the base learner state and the unlabelled point set and predicts the best point to annotate next. Training this active query policy network with reinforcement learning, produces the best non-myopic policy for a given dataset. The key challenge in achieving a general solution to AL then becomes that of learner generalisation, particularly across heterogeneous datasets. We propose a multi-task dataset-embedding approach that allows dataset-agnostic active learners to be trained. Our evaluation shows that AL algorithms trained in this way can directly generalize across diverse problems.

1 INTRODUCTION

In many applications, supervision is costly relative to the volume of data. In these settings active query selection methods can be invaluable to predict which instances a base classifier would find it informative to label. By carefully choosing the training data, the classifier can perform well even with relatively sparse supervision. This vision has motivated a large body of work in active learning that has collectively proposed dozens of query criteria based on different theoretical or intuitive motivations, such as margin (Tong & Koller, 2002) and uncertainty-based (Kapoor et al., 2007) sampling, expected error reduction (Roy & McCallum, 2001), representative and diversity-based (Chattopadhyay et al., 2012) sampling, or combinations thereof (Hsu & Lin, 2015). It is hard to pick a clear winner all these methods, because each is based on a reasonable and appealing – but completely different – motivation; and there is no consistent winner in terms of performance across all datasets.

Rather than hand-designing a criterion and hoping that it performs well, we take a data-driven learning-based approach. We treat active learning algorithm development as a meta-learning problem and train an active learning policy represented by a neural network using deep reinforcement learning (DRL). It is natural to represent AL as a sequential decision making problem since each action (queried point) affects the context (available query points, state of the base learner) successively for the next decision. In this way the active query policy trained by RL can potentially learn a powerful and non-myopic policy. By treating the increasing accuracy of the base learner as the reward, we optimise for the actual goal: the accuracy of a classifier with a small number of labels. As the class of deep neural network (DNN) models we use includes many classic criteria as special cases, we can expect this approach should be at least as good as existing methods and likely better due to exploiting more information and non-myopic optimisation of the actual evaluation metric.

This idea of learning the best criterion within a very general function class is appealing, and other very recent research has had similar inspiration (Bachman et al., 2017). However it does not provide a general solution to AL unless the learned criterion generalizes across diverse datasets/learning problems. With DRL we can likely learn an excellent query policy

for any given dataset. But this is not necessarily useful alone: if we had the labels required to train the policy on a specific problem, we would not need to do AL on that problem in the first place. Thus the research question for AL moves from “what is a good criterion?” to “how to learn a criterion that generalizes?”. In this paper we investigate how to train AL query criteria that generalize across tasks/datasets. Our approach is to define a DNN query criterion policy that is *paramaterized by a dataset embedding*. By multi-task training of our DNN policy on a diverse batch of source tasks/datasets, the network learns how to calibrate its strategy according to the statistics of a given dataset. Specifically we are inspired by the recently proposed auxiliary network idea (Romero et al., 2016) to define a meta-network that provides paramaterized domain adaptation. The meta network generates a dataset embedding and produces the weight matrices that parameterize the main policy. Besides enabling the policy to adapt to datasets with different statistics, this also means that our policy benefits from end-to-end processing of raw features while being transferrable to datasets of any feature space dimensionality. Finally, unlike Woodward & Finn (2017); Bachman et al. (2017) our framework is agnostic to the base classifier. Treating the underlying learner as part of the environment to be optimised means our framework can be applied to improve the label efficiency of any existing learning architecture or algorithm.

2 PRELIMINARIES

Reinforcement Learning (RL) In a general model-free reinforcement learning setting, an agent interacts with an environment \mathcal{E} over a number of discrete time step t . At each time step, the agent receives the state $s_t \in \mathcal{S}$ from environment and select an action $a_t \in \mathcal{A}$ based on its policy $\pi(a_t|s_t)$ which is a mapping from state to action. Then, agent will receive a new states s_{t+1} and immediate reward r_t from \mathcal{E} . The aim of RL is to maximise the return $R = \sum_{t=1}^{\infty} \gamma^{t-1} r_t$ where the return is the accumulated immediate reward from time step t with discount factor $\gamma \in (0, 1]$. There are multiple approaches to learning the policy π (Kober & Peters, 2009; Mnih et al., 2015). We use direct policy search based RL, which learns π by gradient ascent on the objective function $J_{\pi}(\theta) = \sum_{s \in \mathcal{S}} d(s) \sum_{a \in \mathcal{A}} \pi_{\theta}(a|s) R$, where $d(s)$ is stationary distribution of Markov chain for π_{θ} .

Active Learning (AL) A dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ contains N instances $\mathbf{x}_i \in \mathbb{R}^D$ and labels $y_i \in \{1, 2\}$, most or all of which are unknown in advance. In active learning, at any moment the data is split between a labelled set \mathcal{L} and unlabelled set $\mathcal{U} = \mathcal{D} \setminus \mathcal{L}$ where $|\mathcal{L}| \ll |\mathcal{U}|$ and a classifier f has been trained on \mathcal{L} so far. In each iteration, a pool-based active learner τ selects an instance from unlabelled pool \mathcal{U} to query its label $\tau : \{(\mathcal{L}, \mathcal{U}, f) \rightarrow i\}$, where $i \in \{1, \dots, |\mathcal{U}|\}$. Then the selected instance i is removed from the unlabeled set \mathcal{U} and added to the labeled set \mathcal{L} along with its label, and the classifier f is retrained based on the updated \mathcal{L} .

Connection between RL and AL In order to go beyond the many existing heuristic criteria, we propose to model an active learning algorithm as a neural network, and formalize discovery of the ideal criterion as a deep reinforcement learning problem. Let the state of the world s_t consist of a featurisation of the dataset and the state of the base classifier $s_t = \{\mathcal{L}_t, \mathcal{U}_t, f\}$. Let an active learning criterion be a policy $\pi(a_i|s_t)$ where index of actions $i \in \{1, \dots, |\mathcal{U}|\}$ select a point in unlabelled set to query. Upon querying a point the world state is updated to s_{t+1} as that point is moved from \mathcal{U} to \mathcal{L} and f is updated as the base classifier is retrained. Assume the policy is a neural network paramaterized by weights θ , that selects actions as $\pi(a_i|s_t) \propto \exp^{\Phi_{\theta}(a_i, s_t)}$, where $i \in \{1, \dots, |\mathcal{U}|\}$ is the index of the unlabelled instances. Finally, we define the reward of an episode to be the quantity we wish to maximize. E.g., If the budget is N queries and we only care about the accuracy after the N th query, then we let $R = Acc_N$ where Acc_N is the accuracy after the N th query. Alternatively, if we care about the performance during all the N queries, we can use $R = \sum_{t=1}^N \gamma^{t-1} Acc_t$. (This illustrates an important advantage of the learning active learning approach: we can tune the learned criterion to suit the requirements of the AL application.) By training θ to maximize the objective $J(\theta)$ we obtain the optimal active learning policy. In interpreting AL criterion learning as a DRL problem, there is the special consideration that unlike general RL problems, each action can only be chosen once in an

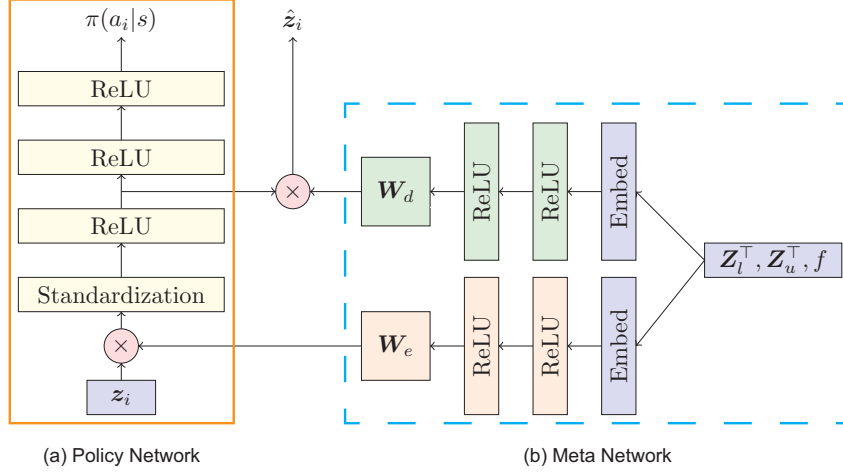


Figure 1: Policy and Meta Network architecture for deep reinforcement learning of a task-agnostic active query policy. (a) Policy network; (b) Meta network. Blue rectangles are inputs and red circles indicates the product operations.

episode. We will achieve this by defining a fully convolutional policy network architecture where the dimensionality of the output softmax $\pi(a_i|s_t)$ can vary with t .

3 METHODS

Recall that our aim is to obtain the parameters θ of an effective dataset-agnostic active query policy $\pi_\theta(a|s)$. The two key challenges are how to learn such a policy given that: (i) the testing dataset statistics may be different from training dataset statistics, and moreover (ii) different datasets have different feature dimensionality d . This challenge is addressed by defining the overall policy $\pi_\theta(a|s)$ in terms of two sub-networks – a policy network and meta network – described as follows.

Policy Network Overall the policy network π inputs all N unlabelled instance $\mathbf{Z}_u \in \mathbb{R}^{N \times d}$ and its output is an N -way softmax distribution for selecting choice of instance to query. We assume the policy models actions via the softmax $\pi(a_i|s) \propto \exp^{\Phi_{\theta_p}(\mathbf{W}_e^T \mathbf{z}_i)}$, where $\mathbf{z}_i \in \mathbb{R}^d$ is the i th unlabelled instance in \mathbf{Z}_u and $\mathbf{W}_e \in \mathbb{R}^{d \times k}$ encodes the pool of instances. Although dimensionality d varies by dataset, the encoding $\mathbf{u}_i = \mathbf{W}_e^T \mathbf{z}_i \in \mathbb{R}^k$ does not, so the rest of the policy network $\pi(a_i|s) \propto \exp^{\Phi_{\theta_p}(\mathbf{u}_i)}$ is independent of dataset dimension. The key is then how to obtain encoder \mathbf{W}_e which will be provided by the meta network. Following previous work (Bachman et al., 2017; Konyushkova et al., 2017) we also allow the instances to be augmented by instance-level expert features so $\mathbf{Z} = [\mathbf{X}, \boldsymbol{\xi}(\mathbf{X})]$ where \mathbf{X} are the raw instances and $\boldsymbol{\xi}(\mathbf{X})$ are the expert features of each raw instance.

Meta Network The encoding parameters \mathbf{W}_e of the policy network is obtained from the meta network: $\Psi_{\theta_m^e} : \{(\mathcal{L}, \mathcal{U}, f) \rightarrow \mathbf{W}_e; \theta_m^e\}$. The meta network inputs a featurisation of \mathcal{L}, \mathcal{U} and f and produces $\mathbf{W}_e \in \mathbb{R}^{d \times k}$ to allow the policy network to process d -dimensional inputs into a fixed k -dimensional hidden representation. Following Romero et al. (2016) we also synthesise and use the $\mathbf{W}_d \in \mathbb{R}^{k \times d}$ dimensional decoder $\Psi_{\theta_m^d} : \{(\mathcal{L}, \mathcal{U}, f) \rightarrow \mathbf{W}_d; \theta_m^d\}$ to regularise this process by reconstructing the input features. The meta network synthesises these weight matrices based on dataset-embeddings of \mathbf{Z}^T described in the following section.

3.1 ACHIEVING CROSS DATASET GENERALIZATION

The idea of auxiliary networks to predict weights for a target network was recently used in Romero et al. (2016). There the auxiliary network inputs an embedding of \mathbf{X}^T and predicts the weights for a main network that inputs \mathbf{X} , with the purpose of reducing the

total number of parameters if \mathbf{X} is high dimensional. In Romero et al. (2016) all the training and testing is performed on the same dataset. Here we are inspired by this idea in proposing a meta-network strategy for achieving end-to-end learning of multiple-domains. By multi-task training on multiple datasets, the meta-network learns to generate dataset-specific weights for the policy network such that it performs effectively on all training problems and generalises well to new testing problems based on their embedding.

Dimension Embedding Strategy The auxiliary meta-network requires a feature embedding that produces a fixed size description of each dimension across all datasets. The meta network takes $(\mathcal{L}, \mathcal{U}, f)$ as input, treating each feature as an example. It extracts an embedding from each input (feature) and then predicts the policy network’s weights for the corresponding feature. All together, the auxiliary network predicts the weight matrix $\mathbf{W}_e \in \mathbb{R}^{d \times k}$, which the policy network can use to map each feature dimension to a k dimensional embedding, as

$$(\mathbf{W}_e)_j = \Psi\left([e_j^1(\mathbf{Z}_u^T), e_j^1(\mathbf{Z}_l^T), e_j^2([\mathbf{Z}_u^T, \mathbf{Z}_l^T], f)]\right). \quad (1)$$

Here e is a non-linear feature embedding, j indexes features, selecting the j th embedded feature and the j th row of \mathbf{W}_e , and Ψ is the non-linear mapping of the meta-network, which outputs a vector of dimension k . Similarly, the meta-network also predicts the weight matrix \mathbf{W}_d used for auto-encoding reconstruction (Fig 1). Although d is dataset dependent, the meta network generates weights for a policy network of appropriate dimensionality ($d \times k$) to the target problem. The specific embeddings used are explained next.

Choice of Embeddings We use two ‘representative’ and ‘discriminative’ histogram style embeddings. The dimension-level embedding is to embed each feature dimension into a h histogram. **Representative** ($e_j^1(\mathbf{Z}_u^T)$ and $e_j^1(\mathbf{Z}_l^T)$) For the representative embedding, we encode each feature dimension as a histogram over the instances in that dimension. For example, we may rescale the i th dimension features into $[0, 1]$ and divide the dimension into 10 bins. Then we count the proportion of labelled and unlabelled data for each bin. This would give a 1×20 embedding for each dimension. **Discriminative** ($e_j^2([\mathbf{Z}_u^T, \mathbf{Z}_l^T], f)$) In this case we create a 2-D histogram of 10 bins per dimension. In this histogram we count the frequency of instances with feature values within each bin as per the previous embedding jointly with the frequency of instances with posterior values within each bin. Finally we get counts in a 10×10 grid, which we vectorize to 1×100 . Concatenating these two kind of embeddings we have we have that $[e_j^1(\mathbf{Z}_u^T), e_j^1(\mathbf{Z}_l^T), e_j^2([\mathbf{Z}_u^T, \mathbf{Z}_l^T], f)]$ provides a $E = 120$ dimensional representation of each feature dimension for processing by the meta network.

Training for Cross Dataset Generalization We train policy networks and meta networks using the policy gradient method REINFORCE (Williams, 1992) to ensure that the generated policies maximise the return (active learning accuracy) with the desired reward discounting. To ensure that our pair of networks achieve the desired dataset (active learning problem) invariance, we perform multi-task training on multiple source datasets: (i) In every mini batch we sample a random subset of source datasets, and set the return to the average return over all the sampled datasets. Thus achieving a high return means the meta network has learned to synthesise suitable per-dataset weights for the policy network based on the dataset embedding, and that together they generalize across multiple tasks/datasets. (ii) To further promote cross-dataset generalization, we apply the baseline method to standardize the return from each episode which compensates the variant return scale in different datasets. This relative return alleviates the risk of domination by a single dataset with large return due to differing scale of accuracy increments among datasets of varying difficulty. The overall training algorithm is summarised in Alg. 1.

3.2 REINFORCEMENT LEARNING TRAINING AND OBJECTIVE FUNCTIONS

The ideal active learner should query the instance that maximally improves the base learner performance. The reward that reflects the quantity we care about is therefore the increase of test split accuracy $r_t = \text{Acc}_t - \text{Acc}_{t-1}$. To optimise this quantity non-myopically, we define the return of an active learning session as $J(\theta) = \mathbb{E}[\sum_{t=1}^{\infty} \gamma^{t-1} r_t(s, \pi_\theta(\cdot, s))]$. We then use policy gradient to train the policy and meta-networks to optimise the objective $J(\theta)$.

Algorithm 1 Reinforcement Learning of a Transferrable Query Policy

```

Input:
1: for < each iteration > do                                     ▷ 1 ... 20,000
2:   for < each episode > do                                       ▷ Collect batch
3:     Pick source dataset randomly
4:     Initialise label and unlabelled pool
5:     for < each time step to time T > do
6:       Sample action  $\pi(a_i|s) \propto \exp^{\Phi_{\theta_p}(\mathbf{W}_e^T \mathbf{z}_i)}$ 
7:       Update the  $\mathbf{Z}_u, \mathbf{Z}_l$  and base learner  $f$ 
8:       Record the triplet  $\langle \mathbf{Z}_u, a, r \rangle$                                ▷ state, action, reward
9:     end for
10:    Standardize episode-collected return
11:  end for
12:  Update Policy with standardized return
13: end for
14: return Trained Active Query Policy

```

Auxiliary Regularization Losses Besides optimising the obtained reward, we also optimise for two auxiliary regularisation losses. **Reconstruction:** The policy network should reconstruct the unlabeled input data using \mathbf{W}_d predicted by the meta-network (Romero et al., 2016). **Entropy Regularization:** Following Mnih et al. (2016), we also prefer a policy that maintains a high-entropy posterior over actions so as to continue to explore and avoid pre-emptive convergence to an over-confident solution.

Integrating the main RL and two auxiliary supervised tasks together, we train both networks end-to-end. We maximize the whole objective function F by reversing the sign of reconstruction loss:

$$F = J_{\theta}(\Phi) - A_{\theta_m^d}(\mathbf{Z}_u) + \lambda \mathcal{H}(\pi_{\theta}(\mathbf{a}|\mathbf{Z}_u)) \quad (2)$$

where $\theta = \{\theta_p, \theta_m^e\}$. The network (Fig. 1) trained by Eq. 2 using Alg. 1 learns to synthesise policies that are effective active query criteria (high return J) on any domain/dataset (synthesising domain specific network parameters via auxiliary network), adapting to the statistics of the dataset and independent of the dimensionality of the dataset.

4 EXPERIMENTS

4.1 DATASETS AND SETTINGS

Datasets Overall we experiment with a diverse set of 14 datasets from UCI machine learning repository. These include *austra*, *heart*, *german*, *ILPD*, *ionospheres*, *pima*, *wdbc*, *breast*, *diabetes*, *fertility*, *fourclass*, *habermann*, *livers*, *planning*. We split these into 7 for multi-task training, and 7 for testing.

Architecture The auxiliary network for encoder has fully connected layers with of size 120, 100, 100 ($E = 120, k = 100$) and decoder auxiliary network has analogous structure. The policy network has layers of size $N \times d$ ($N \times d$ input matrix \mathbf{Z}_u), $N \times 100$, $N \times 50$, $N \times 10$, $N \times 1$ (N -way output). All penultimate layers use ReLU activation. Transition of the input to first hidden layer of policy network is provided by the auxiliary network. Thereafter for efficient implementation with few parameters and to deal with the variable sized input and output, it is implemented convolutionally. We convolve a $h_1 \times h_2$ sized matrix across the N dimension of each $N \times h_1$ matrix shaped layer to obtain the next $N \times h_2$ layer.

Experiment Settings We train using Adam optimiser with initial learning rate 0.001 and hyperparameters set to $\alpha = \lambda = 0.005$ and discount factor $\gamma = 0.99$. During RL training, we use two tricks to stabilize the policy gradient. 1) We use a relatively large batch size of 32 episodes. 2) We smooth the gradient by accumulated time-step $G_t = (1 - \alpha)G_{t-1} + \alpha g_t$ where g_t is the gradient of the a_t in time step t and the G_t is the accumulated gradient. Intuitively, the accumulated gradient G_t puts more emphasis on early time step actions. We

train for a total of 20,000 iterations and use perform active learning over a time horizon (budget) of 20. All results shown are averages over 100 trials of training and testing datasets. **Expert Features:** To enhance the low-level feature of each instance in \mathbf{X} we define expert features $\xi(\mathbf{X})$ to include distance furthest first and uncertainty as the augmented feature. In most experiments we use linear SVM as the base learner.

Alternatives We compare our learning approach to AL with two classic approaches uncertainty/margin-based sampling (US) (Tong & Koller, 2002; Kapoor et al., 2007) and furthest-first-based sampling (DFF) (Baram et al., 2004), as well as to random sampling (RAND) as a lower bound. Uncertainty sampling is a simple deterministic approach that queries the instance with minimum certainty (maximum entropy). While simple, and not the most state of the art criteria, it is consistently very competitive with more sophisticated criteria and more robust in the sense of hardly ever being a very poor criteria. As a representative more sophisticated approach, we compare with QUIRE (Huang et al., 2010). We denote our method meta-learned policy for general active learning (MLP-GAL). As a related alternative we propose SingleRL. This is our RL approach, but without the meta-network, so a single model is learned over all datasets. Without the meta-network it can only use expert features $\xi(\mathbf{X})$ so that dimensionality is fixed over datasets. To give SingleRL an advantage we concatenate some extra global features to the input space¹. This method can also be seen as a version of one of the few state of the art learning-based alternatives (Konyushkova et al., 2017). But upgraded in that we learn it with reinforcement learning instead of the more myopic supervised learning used in (Konyushkova et al., 2017).

4.2 MULTI-TASK TRAINING EVALUATION

In the first experiment we verify that it is indeed possible to learn a single policy that generalizes across multiple training datasets. Here we train on a random split of all seven of the training datasets, and test on a random split of the same seven datasets. The results in Table 1 shows that our approach learns an effective criterion that outperforms the given baselines. Although there are different random splits there is possibly some overfitting here as the policy has seen all these datasets at some during training (datasets randomly selected within minibatches). However it is reassuring that it works because it means that it is possible to learn a single query policy that performs well on a diverse set of datasets.

4.3 CROSS-TASK GENERALIZATION

In the next experiment we apply our multi-task trained method to the held out datasets. According to the results in Table 2, we see that our MLP-GAL approach is comparable or better than alternatives, even when applied on a diverse suite of datasets different to those it was trained on. Fig. 2 shows examples of the resulting learning curves we can see that our MLP-GAL can indeed learn query policies that improve the base learner visibly more rapidly than alternatives, even on novel datasets. SingleRL is generally quite effective compared to prior methods, showing the efficacy of training a policy with RL. However it does not benefit from a meta network, so is not as effective as our MLP-GAL. From the table it is also interesting to see that while sophisticated methods such as QUIRE sometimes perform very well, they also sometimes perform very badly scoring even worse than random (red). Meanwhile although the simple and classic uncertainty-sampling is never best, it is also never worst: performing more consistently than more elaborate QUIRE. This illustrates that the challenge in building sophisticated AL algorithms that generalize to datasets that they were not engineered on. In contrast, although our approach has not seen these datasets it performs consistently well due to adapting to each dataset via the meta-network.

An advantage of our approach compared to related methods such as Bachman et al. (2017); Woodward & Finn (2017) is that it treats the base learner as part of the environment to be optimised against rather than tying the user to a particular method. To demonstrate this, we repeat the previous experiment using RBF-SVM as the base learner rather than linear SVM. After multi-task training on source datasets, the results for held out testing datasets

¹Variance of classifier weight, proportion of labelled pos/neg instances, proportion of predicted unlabelled pos/neg instances², proportion of budget used (Konyushkova et al., 2017)

	Australian	heart	germannumber	ILPD	ionospheres	pima	wdbc	Mean
MLP-GAL	0.8086	0.7544	0.6834	0.6817	0.7454	0.6752	0.9014	0.75
DRLexp	0.7727	0.7072	0.653	0.6519	0.7149	0.6375	0.8951	0.7189
US	0.7819	0.7332	0.6473	0.6354	0.6889	0.6452	0.9014	0.7191
DFF	0.7496	0.734	0.6284	0.6659	0.5736	0.5515	0.8309	0.6763
RAND	0.7515	0.7134	0.6368	0.6551	0.6906	0.6347	0.8667	0.707
QUIRE	0.7273	0.7091	0.651	0.6596	0.574	0.59	0.8606	0.6816

Table 1: AUC comparison of active learning algorithms on seen datasets. Linear SVM. Red indicates worse than random policy.

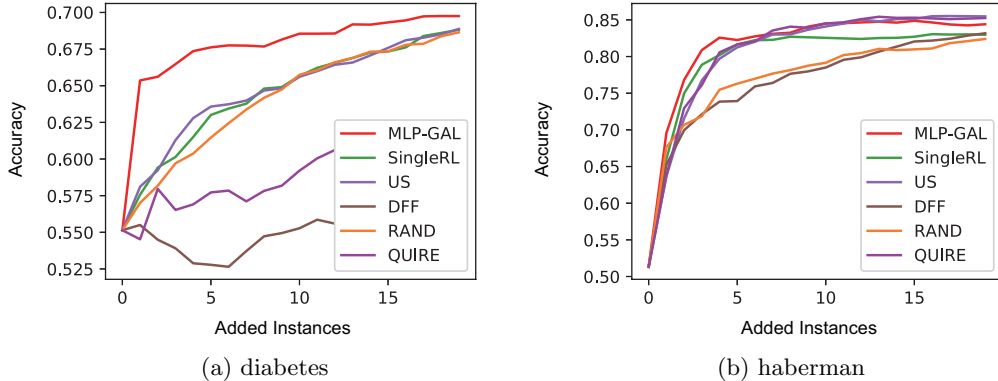


Figure 2: Illustrative active learning curves from evaluating our learned policy on held out UCI datasets (Linear SVM).

are shown in Table 3. We can see that the results of hard-crafted policies are similar to linear SVM (expected given the difficulty of learning a non-linear model in a budget of 20 points). However our learning-based approach is again consistently high performing and effective overall – it is able to learn a policy customised for this new type of base learner.

5 RELATED WORK

Active Learning by Learning A few papers have very recently appeared that also approach finding an AL criterion as a learning problem. Konyushkova et al. (2017) proposes to learn a criterion based on a vector of expert features (e.g., classifier confidence, label imbalance). However by using expert features, this misses the chance to learn the representation from raw features as in our approach; and by using supervised rather than reinforcement learning to train the policy, it is not optimally non-myopic. Bachman et al. (2017) and Woodward & Finn (2017) use RL to train a single model that provides both the base classifier and the active learner. This tight integration has the drawback that the frameworks are constrained to a specific base learner, so cannot be used to improve the training of an arbitrary base learner as per our framework. More importantly, while these methods learn effective non-myopic policies, they are trained and tested on different classes within the same dataset, so the generalization challenge and evaluation is minimal. There is no mechanism to ensure effective transfer across datasets of different statics or to allow any transfer at all across datasets of different dimensionality.

Active Learning Ensembles Different AL algorithms perform well on different datasets, or at different learning stages. For this reason studies have proposed heuristics to switch criteria from early to late stage learning (Donmez et al., 2007; Baram et al., 2004), or use multi-armed bandit (MAB) approaches to estimate the best criterion for a given dataset within an ensemble (Hsu & Lin, 2015). But aside from being myopic, MAB learners do not learn transferrable knowledge: They perform all their learning within a single rollout, and their need to explore/learn online is fundamentally at odds with active learning. Chu

	breast	diabetes	fertility	fourclass	haberman	livers	planning	Avg.
MLP-GAL	0.9591	0.6755	0.8111	0.7173	0.6766	0.552	0.6201	0.7159
SingleRL	0.9574	0.6435	0.7951	0.6973	0.6875	0.5473	0.6298	0.7083
US	0.9556	0.6452	0.8033	0.7191	0.6748	0.5325	0.6154	0.7066
DFF	0.9383	0.5515	0.7635	0.7313	0.6557	0.5534	0.6312	0.6893
RAND	0.9455	0.6396	0.7675	0.6955	0.6712	0.5319	0.6027	0.6934
QUIRE	0.9603	0.59	0.8055	0.7372	0.6648	0.5498	0.6718	0.7113

Table 2: Comparison of active learning algorithms on unseen datasets. AUC averages over 100 trials. (Linear SVM). Red indicates worse than random policy.

	breast	diabetes	fertility	fourclass	haberman	liver	planning	Avg.
MLP-GAL	0.8994	0.6339	0.8515	0.6841	0.6791	0.5509	0.6792	0.7112
SingleRL	0.9137	0.6043	0.8525	0.6866	0.6757	0.5285	0.6710	0.7046
US	0.8556	0.5827	0.8483	0.6710	0.6693	0.5242	0.6659	0.6882
DFF	0.4849	0.5804	0.8470	0.5493	0.6641	0.5196	0.6679	0.6162
RAND	0.8308	0.6128	0.8451	0.6310	0.6725	0.5208	0.6634	0.6823
QUIRE	0.8747	0.5501	0.8475	0.6623	0.6672	0.5344	0.6893	0.6893

Table 3: Comparison of active learning algorithms on unseen datasets. AUC averages over 100 trials. (RBF SVM). Red indicates worse than random policy.

& Lin (2016) ameliorate this somewhat with regularization, but still need dataset-specific learning. Our approach can address these issues: Besides non-myopic policy learning with RL, a DNN has capacity to encode multiple criteria and apply different ones at different stages of learning. By learning a meta-policy that parameterizes a dataset-specific policy, it customizes the overall active learning strategy to the target dataset; thus transferring knowledge for immediate efficacy on a new dataset without dataset specific learning.

Domain Generalisation and Adaptation Our task-agnostic AL goal is related to Domain Generalisation (DG) (Muandet et al., 2013) and Domain Adaptation (DA) (Ganin & Lempitsky, 2015) in supervised learning in that we would like to train on one dataset and perform well when testing on another dataset. Our framework has aspects of DG (multi-task training to increase generality) and DA (adapting to target data, via dataset embedding meta network) methods. But we are not aware of any dataset embedding approaches to achieving DA within supervised learning.

Related Methods Models that predict the parameters of other models are increasingly widely used (Ha et al., 2017). In a reinforcement learning context, this is known as contextual (Kupcsik et al., 2013) or parameterised (Stulp et al., 2013) policies. Romero et al. (2016) used auxiliary networks for parameter reduction when training and testing within one dataset.

6 DISCUSSION

We have proposed a learning-based perspective on the problem of active query criteria design. Such learning-based algorithm design elegantly obtains AL models by optimising the ultimate goal of classification performance with few labels. However aside from the widely-shared questions of good network architecture and RL training algorithms, the key challenge is then whether general enough policies can be learned to be widely useful in different applications, rather than requiring dataset-specific training which contradicts the motivation of AL. Our key contribution is to provide the first solution to this issues through multi-task training of a meta-network that synthesises dataset-specific active query policies.

Our study thus far has the main limitation that we have only evaluated our method on a binary base classifier (binary assumption shared by Konyushkova et al. (2017)). In future work we would like to evaluate our method on deep multi-class classifiers by designing embeddings which can represent the state of such learners, as well as explore application to the stream-based AL setting.

REFERENCES

- Philip Bachman, Alessandro Sordoni, and Adam Trischler. Learning algorithms for active learning. *ICML*, 2017.
- Yoram Baram, Ran El-Yaniv, and Kobi Luz. Online choice of active learning algorithms. *J. Mach. Learn. Res.*, 5, December 2004.
- Rita Chattopadhyay, Zheng Wang, Wei Fan, Ian Davidson, Sethuraman Panchanathan, and Jieping Ye. Batch mode active sampling based on marginal probability distribution matching. *KDD*. ACM, 2012.
- H. M. Chu and H. T. Lin. Can active learning experience be transferred? In *ICDM*, 2016.
- Pinar Donmez, Jaime G. Carbonell, and Paul N. Bennett. Dual strategy active learning. In *ECML*, 2007.
- Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015.
- David Ha, Andrew M. Dai, and Quoc V. Le. Hypernetworks. In *ICLR*, 2017.
- Wei-Ning Hsu and Hsuan-Tien Lin. Active learning by learning. In *AAAI*, 2015.
- Sheng-jun Huang, Rong Jin, and Zhi-hua Zhou. Active learning by querying informative and representative examples. In *NIPS*. 2010.
- Ashish Kapoor, Kristen Grauman, Raquel Urtasun, and Trevor Darrell. Active learning with gaussian processes for object categorization. In *ICCV*, 2007.
- Jens Kober and Jan R. Peters. Policy search for motor primitives in robotics. In *NIPS*. 2009.
- Ksenia Konyushkova, Raphael Sznitman, and Pascal Fua. Learning active learning from real and synthetic data. *CoRR*, abs/1703.03365, 2017.
- Andras Gabor Kupcsik, Marc Peter Deisenroth, Jan Peters, and Gerhard Neumann. Data-efficient generalization of robot skills with contextual policy search. In *AAAI*, 2013.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmarajan Kumar, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 2015.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *ICML*, 2016.
- Krikamol Muandet, David Balduzzi, and Bernhard Scholkopf. Domain generalization via invariant feature representation. In *ICML*, 2013.
- Adriana Romero, Pierre Luc Carrier, Akram Erraqabi, Tristan Sylvain, Alex Auvolet, Etienne Dejoie, Marc-André Legault, Marie-Pierre Dube, Julie G. Hussin, and Yoshua Bengio. Diet networks: Thin parameters for fat genomics. *ICLR*, 2016.
- Nicholas Roy and Andrew McCallum. Toward optimal active learning through sampling estimation of error reduction. In *ICML*, 2001.
- F. Stulp, G. Raiola, A. Hoarau, S. Ivaldi, and O. Sigaud. Learning compact parameterized skills with a single regression. In *Humanoids*, 2013.
- Simon Tong and Daphne Koller. Support vector machine active learning with applications to text classification. *J. Mach. Learn. Res.*, 2, March 2002.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 1992.
- Mark Woodward and Chelsea Finn. Active one-shot learning. *CoRR*, abs/1702.06559, 2017.