

Learning a Policy for Opportunistic Active Learning

Aishwarya Padmakumar, Peter Stone and Raymond J. Mooney

Department of Computer Science

University of Texas at Austin

{aish, pstone, mooney}@cs.utexas.edu

A Supplemental Material

A.1 Sampling predicates for label queries

A large number of perceptual predicates can be used to describe objects. When choosing a predicate whose classifier is to be improved by active learning, the simplest way to choose between predicates is to favor those for which the agent currently has poor classifiers. However, if the number of possible predicates is much larger than the number of queries available for learning, it becomes necessary to focus on a small number of predicates, possibly stopping the improvement on a predicate once the classifier for it has been sufficiently improved.

We use the following distribution to obtain probability weights for predicates as a function of estimated classifier F1. Let $w(p_i)$ be the weight for predicate p_i with estimated F1 $C(p_i)$. Weights start at w_{min} for $C = 0.0$ and increase linearly to w_{max} at some $C_{max} \in (0, 1)$. For $C > C_{max}$, weights again linearly decrease to w_{min} for $C = 1.0$. That is, for $C \leq C_{max}$,

$$w(p_i) = \frac{C(p_i)}{C_{max}}(w_{max} - w_{min})$$

For $C > C_{max}$,

$$w(p_i) = \frac{1.0 - C(p_i)}{1.0 - C_{max}}(w_{max} - w_{min})$$

The weights are then normalized to obtain a probability distribution. A beam of label queries can then be sampled from this.

A.2 Results of Individual Feature Ablation

Table 1 contains the results of ablation of individual features. We use the notation $P_A = \{p_1, p_2, \dots, p_k\}$ for the predicates extracted from the current description. For each predicate $p \in P_A$, we have the estimated F1 of the classifier

$C(p)$, and for each object o in the active test set, we have a decision $d(p, o) \in \{-1, 1\}$ from the classifier. We refer to $s(p, o) = d(p, o) * C(p)$ as the score of the classifier of p for object o . The best classifier for the current interaction is the one with maximum estimated F1, that is, the classifier for $p_{best} = \operatorname{argmax}_{p \in P_A} C(p)$. The second best classifier is $p_{sec} = \operatorname{argmax}_{p \in P_A - p_{best}} C(p)$

Feature Ablated	Success rate	Average Dialog Length
None	0.44	12.95
Number of system turns used - normalized	0.41	3.8
Density of object in label query	0.4	12.89
Fraction of previous dialogs using predicate in query that have succeeded	0.39	5.46
Score (normalized) of top region	0.39	6.3
Fraction of k nearest neighbors of the object in label query, which are unlabeled	0.39	10.41
Indicator for guess action	0.38	7.21
Minimum value of $C(p)$ for $p \in P_A$	0.37	6.37
Decision of p_{sec} for object with highest score	0.37	11.21
Difference between decision of p_{best} for object with highest score, and the average of its decisions for objects in the active test set	0.36	2.78

Feature Ablated	Success rate	Average Dialog Length
Indicator of whether the question is on-topic	0.36	5.25
Is decision of p_{best} same for objects with top two scores	0.36	5.32
Indicate of whether the predicate in the query has a classifier	0.36	13.97
Frequency of use of the predicate in query - normalized	0.35	4.48
Indicator for the action of asking a positive example	0.35	5.36
Second highest value of $C(p)$ for $p \in P_A$	0.35	5.9
Current estimated F1 for classifier of the predicate in query	0.35	6.53
Decision of p_{best} for object with highest score	0.34	3.85
Indicator for the action of asking a label	0.34	7.05
Average value of $C(p)$ for $p \in P_A$	0.34	7.63
Margin of object in label query	0.34	8.08
Maximum value of $C(p)$ for $p \in P_A$	0.33	6.31
Difference between decision of p_{sec} for object with highest score, and the average of its decisions for objects in the active test set	0.33	8.84
Difference between top two scores in the active test set	0.32	8.05
Is decision of p_{best} same for objects with top two scores	0.32	10.01
Difference between top score and average score in the active test set	0.31	7.18
Baseline	0.29	16

Table 1: Results of individual feature ablation. Bold-face indicates that the difference in that metric with respect to *Static* is statistically significant according to an unpaired Welch t-test with $p < 0.05$.