

MOTIVATION

• Existing image datasets contain much more information than what is conveyed by their labels.

• Need to search for new manually annotated images when learning the visual model of a new class.

• Avoid discarding useful information and make the learning process as cheap as possible.



transfer learning

- new task: small amount of annotated data
- discover relation between old and new tasks

zero-shot learning

- new task: no annotated data
- known relation between old and new tasks

active learning

• new task: no annotated data

ACTIVE LEARNING

• Unlabeled samples $\{m{x}_i\}_{i=0}^N \in \mathbb{R}^d$ belonging to C classes. Binary setting $y_i \in \{-1, +1\}$.

• Model vector: $\boldsymbol{w} \in \mathbb{R}^d$. Prediction score: $f(\boldsymbol{x}) = \boldsymbol{w} \cdot \boldsymbol{x}$. Final annotation: y = sign(f(x)).

• $\gamma_i^t \in \{0,1\}$ indicates whether at time step t the label y_i has been queried. *B*: annotation budget.

$$\max_{\alpha^{t},\gamma^{t}} \sum_{i} \gamma_{i}^{t} \lambda_{i}^{t} \alpha_{i}^{t} - \frac{1}{2} \sum_{i,j} \alpha_{i}^{t} \alpha_{j}^{t} \gamma_{i}^{t} \gamma_{j}^{t} y_{i} y_{j} \boldsymbol{x}_{i} \cdot \boldsymbol{x}_{j}$$
s.t.
$$\sum_{i} \gamma_{i}^{t} \alpha_{i}^{t} y_{i} = 0$$

$$0 \leq \alpha_{i}^{t} \leq C, \ \forall i ,$$

$$\gamma_{i}^{t} \geq \gamma_{i}^{t-1}, \ \forall i ,$$

$$\sum_{i} \gamma_{i}^{t} = \sum_{i} \gamma_{i}^{t-1} + B .$$

ACTIVE TRANSFER LEARNING WITH ZERO-SHOT PRIORS: REUSING PAST DATASETS FOR FUTURE TASKS

E. GAVVES¹, T. MENSINK², T. TOMMASI³, C.G.M. SNOEK¹, T. TUYTELAARS⁴ ¹QUVA LAB, UNIVERSITY OF AMSTERDAM, ²UNIVERSITY OF AMSTERDAM, ³UNC CHAPEL HILL, ⁴KU LEUVEN

CONDITIONS & PRIORS

Maximum Conflict

Query the sample i^* such that (*a*) its label y_{i^*} has an opposite sign from its classification score at (t-1), while (b) the classifier score is as high as possible.

Label Equality

The number of positive and negative examples in the training set should be balanced, *i.e.* $\sum_{i} \gamma_{i}^{t} [y_{i} = 1] = \sum_{i} \gamma_{i}^{t} [y_{i} = -1].$

Zero-Shot Priors

Given a set of known concepts \mathcal{K} , the zero-shot prediction for a new class can be written as

$$f^{zs}(\boldsymbol{x}) = \sum_{k \in \mathcal{K}} \beta_{ck} \, \boldsymbol{w}_k \cdot \boldsymbol{x}_i \; ,$$

We modify the prediction score of active learning

$$f^t(\boldsymbol{x}) = \eta^t f^{zs}(\boldsymbol{x}) + \boldsymbol{w}^t \cdot \boldsymbol{x}.$$

QUERY SAMPLING



• At each iteration t compute the likelihood of sampling positive and negative labels from \mathcal{F}_+ and \mathcal{F}_0 :

$$p(l|\mathcal{F}_{+}, t) \sim p(l|\mathcal{F}_{+}, t-1) + \frac{\sum_{r=1}^{t-1} [l_r = 1] \cdot [\mathbf{x}^r \in \mathcal{F}_{+}^r]}{t-1}$$

• Normalize the probabilities to sum up to one, and measure the label equality

$$\rho^{t-1} = \frac{\sum_{r=1}^{t-1} [l_r = 1]}{\sum_{r=1}^{t-1} [l_r = 1] + [l_r = -1]} \cdot \frac{\left[l_r = 1\right]}{\left[l_r = 1\right] + \left[l_r = -1\right]}$$

• Sample such that









RESULTS

Datasets

HSUN: multi-class, multi-label dataset, covering object and scene categories. 107 classes, 8634 images, split into training/test set of 4367/4317 images. MS-COCO: multi-class, multi-label dataset containing 80 object categories. 123,287 images in total. We used the full dataset and a subset with 4000 images in training and test.



	HSUN (All samples: 0.383 mAP)						
lo. of queries	0	50	100	150	200	250	300
ICLE (Dataset prior) ICLE (External prior)	0.255 0.270	0.315 0.289	0.337 0.327	0.348 0.341	0.346 0.336	0.355 0.348	0.361 0.358
BAL [1] liearchical Sampling [3] P Mean [4] P Variance [4] P Impact Bayes [2] P EMOC Bayes [2]	0.158 0.089 0.154 0.154 0.154 0.154	0.241 0.156 0.282 0.201 0.251 0.277	0.276 0.199 0.319 0.206 0.286 0.310	0.309 0.221 0.340 0.216 0.305 0.320	0.322 0.234 0.350 0.226 0.316 0.328	0.328 0.230 0.361 0.240 0.327 0.332	0.325 0.246 0.365 0.244 0.345 0.337

Comparison with state-of-the-art active learning methods. For HSUN the external prior comes from COSTA learned on MSCOCO, whereas for MSCOCO the external COSTA prior is learned on HSUN. For Small MSCOCO, MCLE reaches the full mAP within 300 samples, even with external priors.



REFERENCES

[1] S. Vijayanarasimhan, et al. Far-sighted active learning [3] S. Dasgupta and D. Hsu. Hierarchical sampling on a budget for image and video recognition. CVPR, 2010. for active learning. ICML, 2008. [2] A. Freytag, et al. Selecting influential examples: Active [4] A. Kapoor et al. Gaussian processes for object catlearning with expected model output changes. ECCV, 2014. egorization. IJCV, 2010.







KU LEUVEN

Choosing the Zero-Shot Prior

COSTA: object co-occurrence statistics; Attribute-like: binarized class-to-class relations from class attributes;

Image search: model trained with the first 12 images returned by Google image search as positive examples.

Maximum Conflict - Label Equality constant prior: $\eta^t = 1, \forall t$ vanilla prior: $\eta^0 = 1$, $\eta^{t \neq 0} = 0$ Our method automatically adapts to the different conditions providing always higher or equal results than selecting only from F_+ or F_0 .

Small COCO (All samples: 0.460 mAP) 300 250 0.4480.383 0.460).381 0.197 0.391 0.427 0.442 0.457 0.293 0.436 0.395 0.408 0.168 0.250 0.365 0.076 0.182 0.287 0.331 0.309 0.438 0.431 0.309 0.326 0.186 0.291 0.284 0.263 0.430 0.393 0.436 0.346 0.417 0.186 0.399 0.405 0.375 0.388 0.397 0.336

> Active learning queried samples on MSCOCO with HSUN zero-shot priors. Positive/Negative samples are in green/red. Observe the balanced label distribution, while the visual similarity reveals the maximum conflict selection of the next sample.