Selecting Influential Examples: Active Learning with Expected Model Output Changes

Code Available

since 1558

Alexander Freytag, Erik Rodner, and Joachim Denzler

Computer Vision Group, Friedrich Schiller University Jena, Germany

`http://www.inf-cv.uni-jena.de`

Friedrich Schiller University Jena

Computer Vision Group

ECCV '14

Zürich, September 6-12th, 2014

Motivation and method

Active learning: automatically select examples $\mathbf{x}' \in \mathfrak{U}$ that shall be labeled by an annotator, i.e., assigned output value $y' \in \mathcal{Y}$, and are likely to increase the accuracy of a model $f : \mathcal{X} \to \mathcal{Y}$

Main idea:

select examples that likely change future outputs of the model

Expected model output changes

$$\Delta f(\mathbf{x}') = \mathbb{E}_{y' \in \mathcal{Y}} \mathbb{E}_{\mathbf{x} \in \Omega} \left(\mathcal{L}\left(f(\mathbf{x}), f'(\mathbf{x}; \mathbf{x}', y') \right) \right)$$

active learning score for unlabeled example $\mathbf{x}'$

marginalization over yet unknown label y'

marginalization over unknown future examples

loss function measuring the differences of model outputs *e.g.*, $\mathcal{L}(z, z') = |z - z'|$

output of the model when $(\mathbf{x}', y')$ is added as a training example

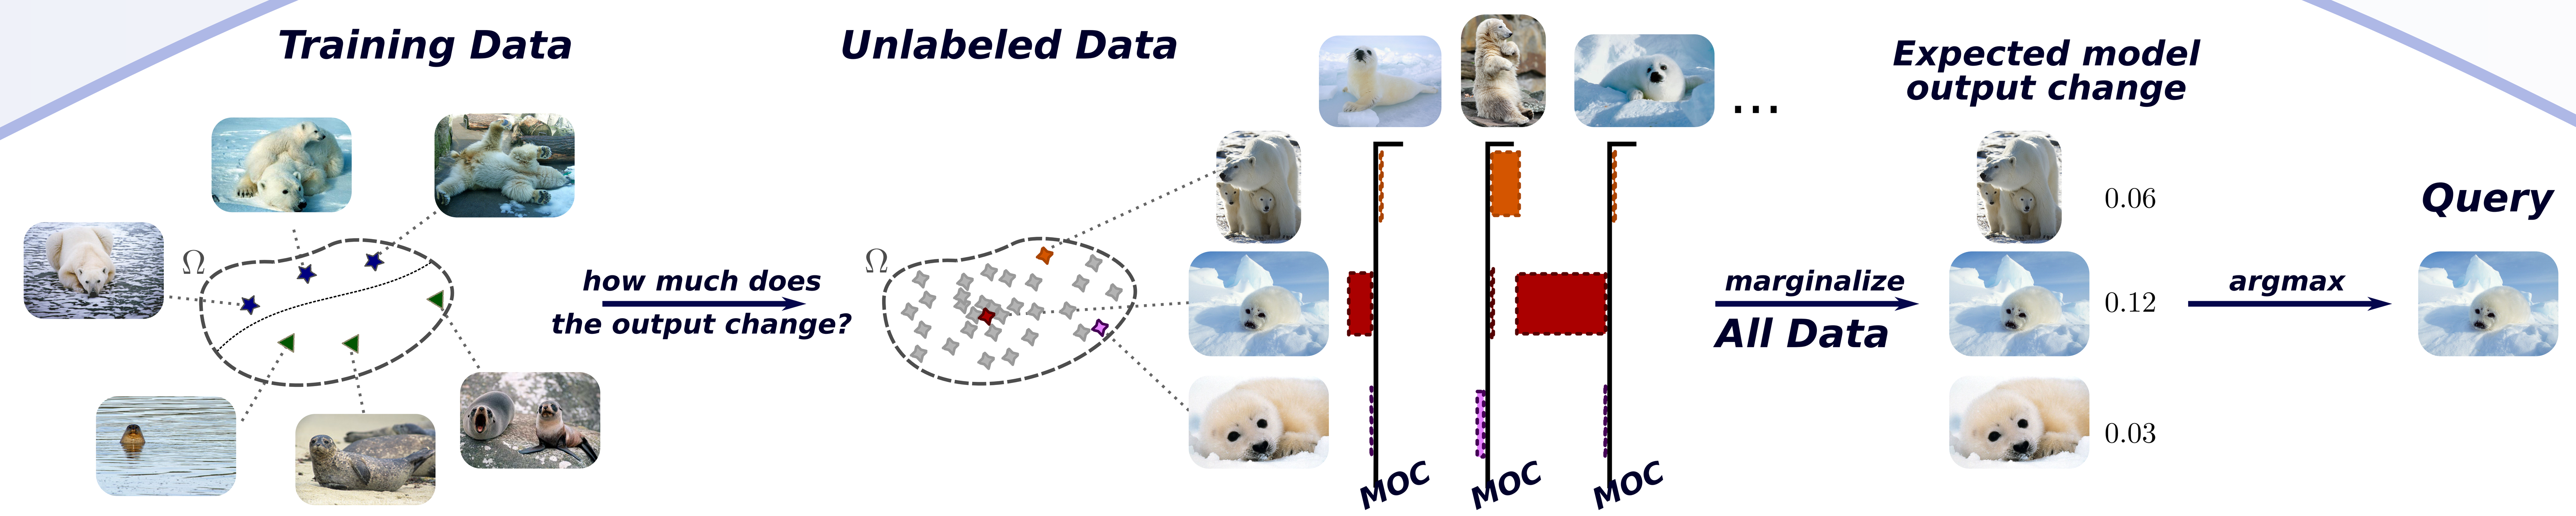

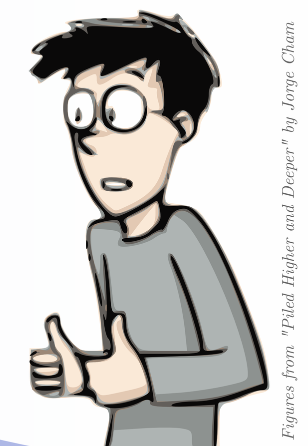

What kind of examples would you like to label, examples that change future classification decisions or examples with no impact at all?

Properties of our method

- strong ties to other active learning strategies (often a generalization)
- upper bound for expected loss reduction (proof in the paper)
- efficient calculation in the case of Gaussian process regression
- further approximation leads to fast evaluation times

Qualitative results

Changes in classification decisions

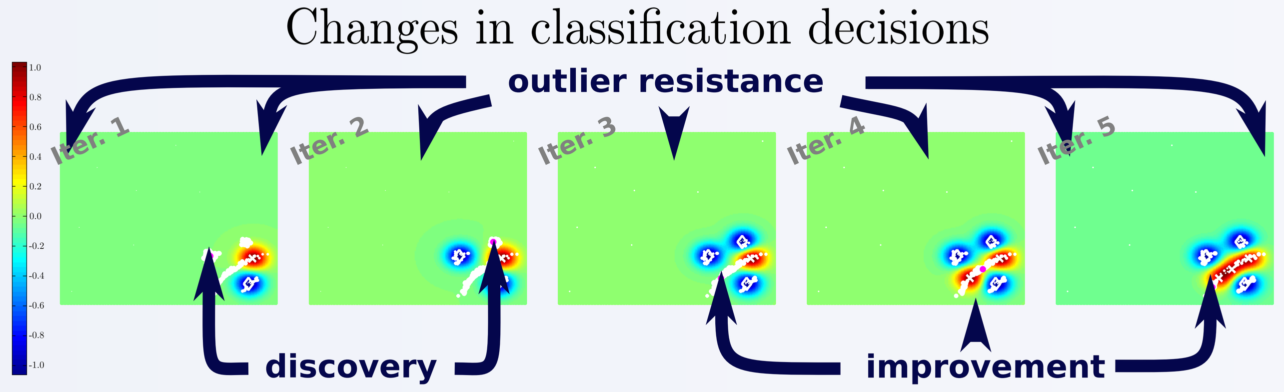

Active learning of an airplane detector

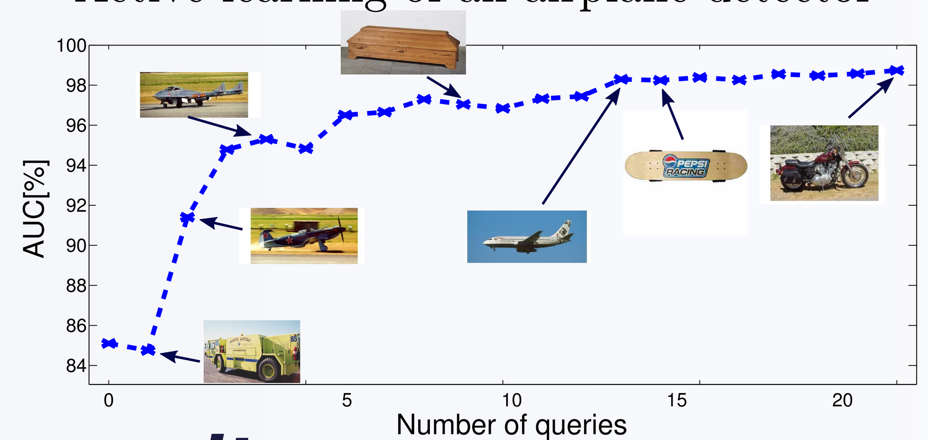

Quantitative results

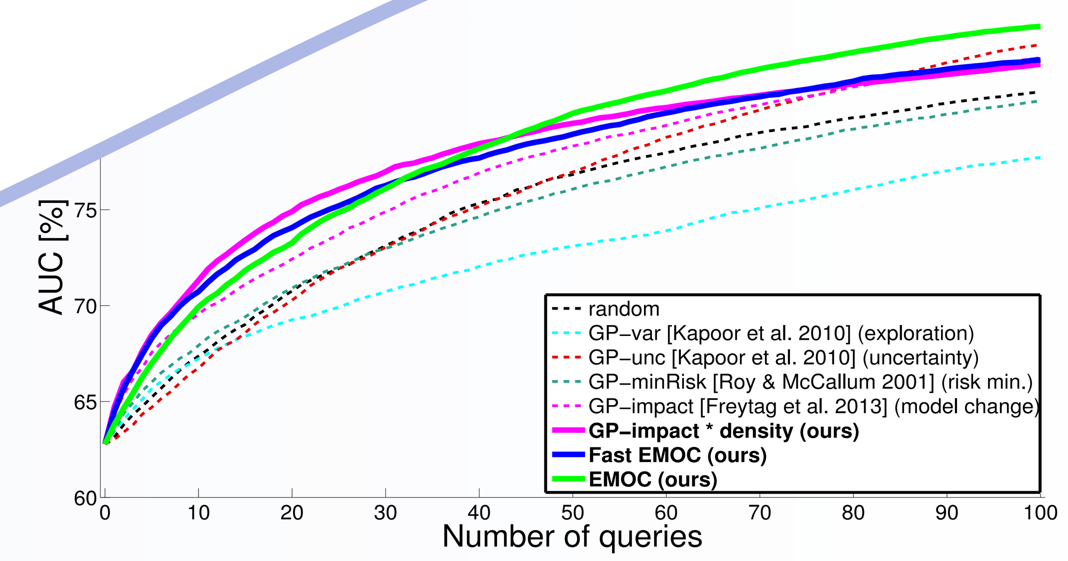

Strategy	ImageNet	Caltech 256
Random	76.83	81.70
GP-var (Kapoor et al., IJCV'10)	73.11	77.06
GP-unc (Kapoor et al., IJCV'10)	76.97	84.31
GP-minRisk (Roy and McCallum, ICML'01)	76.08	84.59
GP-impact (Freytag et al., GCPR'13)	78.32	84.56
EMOC strategy (Ours)	**80.03**	**85.88**

Kapoor et al., *"Gaussian processes for object categorization,"* (IJCV 2010)

Roy and McCallum, *"Toward optimal active learning through sampling estimation of error reduction,"* (ICML 2001)

Freytag et al., *"Labeling examples that matter: Relevance-based active learning with gaussian processes,"* (GCPR 2013)