

Dear editors & reviewers,

**RE: Decision on your manuscript VISI-D-12-02746, "Active Rare Class Discovery and Classification using Dirichlet Processes" by Tom S. F. Haines and Tao Xiang.**

We would like to thank the reviewers for their additional constructive feedback and suggestions. The manuscript has been revised accordingly. Details of our changes and our responses (in italics) are elaborated below. We hope that our revised manuscript is acceptable for publication.

**Editors:**

1. Please address the reviewers' recommendations for specific revisions. Please also stress the computer vision contribution in your revision in light of the comments made by Reviewer 1.

*We have addressed both reviewers' recommendations thoroughly. The paper has been revised to emphasis its computer vision contribution.*

**Reviewer 1:**

1. (1a) The authors clearly state in their response that "we are the first and only authors to use [Dirichlet Processes] for the purpose of active learning." This is not true, see [1] below (I also found several papers very close to the subject, but not quite close enough to warrant inclusion here).

[1] Vlachos, Andreas, Zoubin Ghahramani, and Ted Briscoe. "Active learning for constrained Dirichlet process mixture models." Proceedings of the 2010 Workshop on GEometrical Models of Natural Language Semantics. Association for Computational Linguistics, 2010.

*Thanks for pointing out this reference. We have read this paper carefully and reached the conclusion that despite its title, it differs significantly from our paper and does not undermine the novelty of our work. In the revised manuscript we have include this paper in the references and detailed its differences to our work.*

*Specifically, there are two key differences: (1) the work of Andreas et al. uses Dirichlet processes for the purpose of clustering, and feeds the sampled probabilities into an active learning method based on the standard entropy approach - no new active learning method is proposed. This is not the same thing as using a Dirichlet process for active learning directly, by developing a novel active learning criteria as we have done. (2) Their paper concerns the active classification problem, whilst we are concerned with a different problem - active discovery and classification with rare classes. As explained in our paper, the latter is a very different problem with many applications, particularly in computer vision.*

2. (1b) Dirichlet Processes are commonly used to estimate (eg) the number of components in a mixture model. The active learning element in the paper amounts to computing the probability that an instance is wrongly classified. Since this is given by the DP under two simplifying assumptions made by the authors, they are correct to write "It is then a simple matter to calculate the probability of incorrectly classifying an instance, ". Even if we disregard the

existence of [1], the degree of novelty is too small, in my view, for IJCV. Selecting the worst-case instance as the next candidate does make sense; in fact it is a heuristic I have used myself. However, I regard that heuristic as being too obvious for a strong claim of novelty to be made.

*We disagree that a simple solution equates to a small degree of novelty. It is true that the actual implementation of the method is simple, but simple methods that work are usually considered to be advantageous. We consider the complexity of an approach to be irrelevant when judging its novelty.*

*We also disagree that it is a heuristic – the method has been developed in a principled manner with a sound theoretical foundation.*

3. Previously I was concerned that the paper does not match the journal. The authors responded by writing "[the] ability to solve more general (machine learning) problems should preclude [our method] from continuing to be considered as a computer vision paper." Of course, I agree with them - but they miss my points.

My points are:

- (a) that ML papers that claim to be general are best reviewed by the general ML community,
- (b) ML papers in CV should be applied to CV problems.

I made one suggestion before regarding segmentation, but active visual-object class discovery is another. As it is, the authors continue to regard the data as a passive entity in the sense that it is already given by some external process. I am not suggesting the authors invent new CV, only that they apply their active learning method to a CV problem in a more intimate way.

*Regarding point (a): We understand your concern that the paper should be more CV-focused for publication on IJCV. In the revised version we have done our best to focus on the computer vision aspects of the paper. This includes revising the introduction to highlight the relevance of the work to CV and its contribution to CV, and updating the experiments section to give more insight on the application of the method to computer vision problems.*

*We would like to take this opportunity to highlight that (1) The proposed method has been formulated for and tested on six different CV problems (segment, pageblocks, letters, gait, digits and faces), which clearly shows that it can be applied to a variety of computer vision tasks. Whilst it is equally valuable to many other fields, we think this demonstrates its relevance to the computer vision community, and thus its suitability for publication in IJCV. (2) This paper was invited for submission by the editor, and both of the other reviewers are satisfied that the paper is of sufficient quality for publication in IJCV. This re-enforces our view that this paper matches the scope of IJCV.*

*Regarding your point (b): we agree, and as mentioned above, the proposed method has been applied to six different CV problems. Indeed, there are many other CV problems where the proposed method can be also applied. However, we do not think it will add much value if we were to exhaustively test on more CV problems. In the revised manuscript we have now discussed the segmentation problem and the issue of more intimate integration with human annotation.*

### Reviewer 3:

1. - P3, L12-18: Tong and Koller (2000) did not introduce the concept of "version space" (it was introduced earlier by T. Mitchell, Artificial Intelligence, 1982). I suggest rewriting that sentence to something like, "Tong and Koller (2000) introduce a margin-based active learning method for SVMs, which attempts to minimize the version space of the SVM at each iteration..."

*Done, with the correct citation used for version spaces.*

2. - P7, L31-32: While this may be true, there is no empirical evidence in the paper to support this vague claim. One suggestion to revise would be to point out that typically an estimate of confidence can still be obtained with non-Bayesian classifiers.

*Done.*

3. - P7, L51-55: The section on "Cost-benefit analysis" is unclear. (Should "maximized" be "minimized"?) Please elaborate.

*Have revised this section, and corrected it to minimised. As it is an application of cost benefit analysis from economics a suitable reference has been included and an example scenario presented.*

4. - P11, footnote 10: It would be better to omit this; it does not add anything to the paper.

*Agree - have removed.*

5. Minor comments:

- P2, L15: its  $\rightarrow$  it is
- P6, L19: italicized  $P(\text{data})$  (other places it's not italicized)
- P8, L23: 2(a)  $\rightarrow$  Figure 2(a)

*All of the above changes have been made.*