

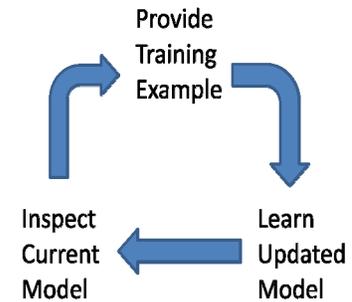
Designing for End-User Interactive Concept Learning in CueFlik

Saleema Amershi[†], James Fogarty[†], Ashish Kapoor[‡], Desney Tan[‡]

[†]Computer Science & Engineering
DUB Group, University of Washington
Seattle, WA 98195
{ samershi, jfogarty }@cs.washington.edu

[‡]Microsoft Research
One Microsoft Way
Redmond, WA 98052
{ akapoor, desney }@microsoft.com

The current information explosion fundamentally changes how people live and work with computing: vast numbers of documents and images are available on the Web; ubiquitous sensing enables near-continuous tracking and monitoring of people and objects; and inexpensive storage allows people to keep near-unlimited personal data and sensing archives. One strategy to enable effective access to and interaction with such large unstructured datasets is to support example-based iterative end-user training of machine learning systems to identify relevant concepts. These concepts can then be used specify desired manipulations. In the context of the CueFlik system for re-ranking Web image search results according to their visual characteristics, we have been examining general questions surrounding the design of end-user interactive machine learning.



By iteratively providing examples and inspecting the resulting model, end-users can train machine learning systems to recognize desired concepts in data.

One focus of our research has been *which examples* a system should present to support effective end-user inspection and interpretation of the current learned model. In our CHI 2007 paper [2], we compared (1) presenting the full ranking of a set of images according to the current model with (2) a *split* presentation of high-certainty positive and negative examples, finding that the *split* presentation led participants to train higher quality models, based on fewer examples, in less time. Hypothesizing that the split presentation was important, but that there are more effective ways to convey positive and negative regions of the space defined by a current model, we next developed example selection methods intended to provide *overviews* of the positive and negative regions. Presented at UIST 2009 [1], these overviews further improve the quality of examples provided by end-users throughout an interactive learning process. We have also found that iterative process can lead to *decay* during training, wherein an end-user continues to provide training examples that are actually hurting the performance of a learned model. In work currently submitted to CHI 2010, we help to address the problem of decay by showing that end-users naturally adopt a lightweight undo mechanism for exploring multiple potential combinations of training examples, leading to higher quality models than the traditional iterative process.

We look forward to participating in this workshop, and believe we bring a valuable perspective based in a combination of machine learning and human-computer interaction principles. Our work is based in distance metric learning and active information acquisition, and we have derived novel algorithms and techniques to build a unified framework for ranking, sampling, and automatically determining which images to display to a person as they iteratively train a model. This enables new approaches to effective end-user Web image search, and also demonstrates a powerful combination of machine learning and human-computer interaction perspectives on important problems.

References

- [1] Amershi, S., Fogarty, J., Kapoor, A. and Tan, D.S. (2009). Overview-Based Example Selection in End-User Interactive Concept Learning. *Proceedings of the ACM Symposium on User Interface Software and Technology* (UIST 2009), 247-256.
- [2] Fogarty, J., Tan, D., Kapoor, A. and Winder, S.A. (2008). CueFlik: Interactive Concept Learning in Image Search. *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI 2008), 29-38.