

Multi-Resolution Learning for Knowledge Transfer

Eric Eaton

University of Maryland Baltimore County
Department of Computer Science and Electrical Engineering
1000 Hilltop Circle, Baltimore, MD 21250
EricEaton@umbc.edu

Abstract

Related objects may look similar at low-resolutions; differences begin to emerge naturally as the resolution is increased. By learning across multiple resolutions of input, knowledge can be transferred between related objects. My dissertation develops this idea and applies it to the problem of multitask transfer learning.

Thesis Overview

Consider a child learning about farm animals. A common early mistake for children is confusing cows with horses, dogs with cats, and so forth. There are many similarities between these animals from a general perspective, enough so that a lot of knowledge learned about one animal can transfer to other similar animals (cows and horses both have four legs, are large, eat grass). At more detailed perspectives, we begin to notice significant differences between these animals that requires learning about each animal individually (horses have manes, cows have udders). The relationship between objects emerges naturally from viewing those objects at varying resolutions of detail.

Inspired by this idea, my dissertation research will focus on *learning across multiple resolutions* specifically for the purpose of *knowledge transfer*. I will apply this technique to the problem of knowledge transfer in *multitask learning*.

Multitask learning seeks to facilitate the learning of some task t_n by using information gained by first learning *related* tasks $t_1 \dots t_{n-1}$. It has a primary goal of learning task t_n either more rapidly or to a higher competence after learning the related tasks, as compared to learning task t_n in isolation. Multitask learning also has a secondary goal of creating *chunks* of general knowledge that apply to related tasks.

I claim that by exploiting the similarities between objects inherent at lower levels of detail, multi-resolution learning will facilitate transfer in multitask scenarios. The primary contributions of my dissertation will be: a general framework and various techniques for multi-resolution learning—specifically, a tree-based multi-resolution classifier and a multi-resolution extension to support vector machines; and the application of this framework to multitask transfer learning problems.

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

Proposed Approach and Justification

Learning at multiple resolutions has already been shown to significantly improve generalization, but only in one limited case by Liang and Page (1997) for predicting time series with neural networks trained on only three different resolution inputs. The framework of multi-resolution learning has not been applied to other machine learning techniques, which will be a major contribution of my dissertation.

Research has also already shown the success of examining objects at multiple resolutions to detect the similarities and differences between them (e.g. (Lowe 2004)). Low resolution representations obscure the details of objects, and the similarities in the high-level structure emerge naturally.

Because low-resolution representations are simple, they are easier to learn, although the *value* of what can be learned from low-resolution representations is limited. Higher resolutions add complexity and the value of what can be learned is higher, but learning is much more difficult.

Learning at the lower levels of resolution may provide only a limited amount of knowledge, but the knowledge will often transfer to all similar objects on a multitask problem. This low-resolution knowledge will facilitate learning at higher levels of resolution for related objects, since it provides a foundation for future learning. It also can be considered a chunk of general knowledge applicable to a class of objects, thereby serving both goals of multitask learning.

One proposed approach to multi-resolution learning will use an ensemble of classifiers, where ensemble members would focus on different resolutions. Such an ensemble classifier would be embedded in a decision tree, which will use increasingly higher resolutions of an object to determine the branches of the tree to follow. Learning with a tree-based multi-resolution classifier would use ideas from boosting, in that it would refine higher resolution classifiers based on the mistakes made by the lower levels. This would focus the higher resolution classifiers on the *differences* between objects related at lower levels, refining the chunks of knowledge learned at the lower levels based on the new information available at higher resolutions. Classifying a single object would involve looking at multiple resolutions of the object and combining the predictions from the various ensemble members at those levels.

Another approach will integrate multi-resolution learning into support vector machines (SVMs) with the development

of a multiple-resolution kernel. This kernel would represent an object at multiple resolutions and allow learning at each level simultaneously, using a hierarchical SVM. Like the ensemble method described above, this method would focus on the differences between the resolutions.

My dissertation advocates multi-resolution learning as a paradigm, and will provide multiple methods for implementing it, instead of focusing on one type of classifier. This framework of learning general-to-specific representations is related to existing case-based learning literature (Leake 1996), and also has foundations in the psychology literature (e.g. Piaget’s theory of assimilation/accomodation).

Note that the “multi-resolution learning” described here is very different from the “multiresolution instance-based learning” described by Deng and Moore (1995), which uses a multi-resolution data structure to group fixed-resolution input for efficient lookup.

Plan of Research

The first major step is to identify scalable representations and methods for scaling standard fixed-scale representations. I have identified several candidates for scalable representations: standard wavelet decomposition, Mallat’s (1989) tree-based wavelet representation, and a difference-of-Gaussians representation (Lowe 2004). There are several promising methods for reducing the resolution of a fixed-scale representation, including multiple applications of a low-pass filter and standard feature selection methods.

I am currently beginning the second phase of this plan, which is to integrate the multi-resolution approach with standard machine learning classifiers. My focus is on developing online versions of the multi-resolution classifiers to allow for continuous learning, since the multitask problems will be online and continuous.

The final phase will be to use the developed multi-resolution learning methods in multitask learning scenarios. I plan to extend established multitask learning frameworks, such as those developed by Thrun (1995) and Maurer (2006). This phase will address several open questions—most importantly, whether we need to explicitly order the learning tasks to achieve knowledge transfer. Using a specific task ordering to “teach” is a common approach to multitask learning, building new tasks on top of previously learned tasks. However, multi-resolution learning is designed to automatically abstract away detail from each training example and relate the example to others at lower resolutions, so the relation to previously learned knowledge will be automatic, if the technique works as anticipated.

Evaluation Methodology

I will evaluate the multi-resolution classifiers against state-of-the-art classifiers on a variety of image recognition tasks, including single-object recognition on a plain background at a variety of orientations, and object recognition from satellite imagery. The satellite imagery will present a real-world environment where the images are naturally at varying resolutions. I will use discrete and continuous versions of these tasks, where the continuous version will use streams of contiguous images for training. I will also evaluate the multi-

resolution classifiers on standard high-dimensional data sets, such as protein data sets.

The evaluation on multitask environments will compare the multi-resolution learning against learning t_n in isolation on the object recognition tasks described above. The other objective of multitask learning is to learn *general knowledge* applicable to a class of objects. I will evaluate this general knowledge base by evaluating the multitask learning on abstracted images of objects, such as cartoon depictions or children’s illustrations. Multi-resolution learning will also be compared against existing methods for multitask learning, including those by Thrun (1995) and Caruana (1997).

Other Implications

High-resolution data is expensive to obtain, compared to low-resolution data (e.g. satellite photos). Through knowledge transfer between related tasks, multi-resolution classifiers can reduce the overall amount of high-resolution training data needed to obtain competency. Active learning extensions of this work could focus on obtaining costly high-resolution data only for tasks that cannot be generalized from existing training data.

By examining the similarities revealed by learning across multiple resolutions, a natural hierarchy of related objects may emerge from applying this technique to a multitask problem. This relationship hierarchy could be used to provide explanations for *why* an object has a particular set of properties through analogy with related objects or an abstract representation of the object’s class (such as an illustration of the object).

Acknowledgements

This work was supported by NSF ITR #0325329 and a fellowship from the Goddard Earth Sciences and Technology Center at UMBC. Thanks to Marie desJardins and Tim Oates for helpful discussions, and the DC panel for their feedback.

References

- Caruana, R. 1997. *Multitask Learning*. Ph.D. Dissertation, School of Computer Science, Carnegie Mellon University.
- Deng, K., and Moore, A. 1995. Multiresolution instance-based learning. In *IJCAI*, 1233–1239. San Francisco: Morgan Kaufmann.
- Leake, D. B., ed. 1996. *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. AAAI Press.
- Liang, Y., and Page, E. W. 1997. Multiresolution learning paradigm and signal prediction. *IEEE Trans. Signal Process.* 45(11):2858–2864.
- Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. *IJCV* 60(2):91–110.
- Mallat, S. G. 1989. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.* 11(7):674–693.
- Maurer, A. 2006. Bounds for linear multi-task learning. *JMLR* 7:117–139.
- Thrun, S. 1995. *Explanation-Based Neural Network Learning: A Lifelong Learning Approach*. Ph.D. Dissertation, University of Bonn, 53117 Bonn, Germany.