# 1 Project Summary

**Intellectual merit.** The dramatic processing power and large memories of modern computers, coupled with sophisticated machine learning algorithms, now makes it possible to teach computers visual concepts through example-based learning. But this process is time consuming and arduous. Often large data sets must be manually collected. Machines typically do not take advantage of previously learned knowledge when performing new tasks. And when confronted with a new situation, systems fail catastrophically. The goal of my research is to make it dramatically easier to teach vision systems new skills, and to design machines that can learn tasks faster by leveraging previously learned knowledge. In short, *I want to develop computer vision systems that are largely self-taught.* I have achieved and will continue to achieve these results by developing computer vision systems that

- learn from a small number of examples,
- use previously learned knowledge to improve performance on novel tasks,
- learn properties of one object that can be used to make inferences about other objects,
- acquire and organize information autonomously, and
- leverage interdisciplinary techniques to help relieve us from training them.

**Impact.** These capabilities are taken for granted in human beings, but represent serious shortcomings in today's computer systems. My thesis is that we cannot afford to train vision systems one problem at a time, acquiring large training sets and developing training paradigms for each task to be learned. Even if we were willing to do this, there are many scenarios in which training data are severely limited (there are limited photos of Abraham Lincoln). Furthermore, we would like computer systems to be *adaptive*, and not have to be prepared for each new task, especially when these new tasks are similar to previous ones.

This proposal discusses five examples of research in this direction:

- building a classifier of handwritten digits from a single example of each digit, using knowledge previously acquired (by the computer) about general handwriting;
- learning a notion of color constancy for most objects under most lighting conditions, with only a single object provided as training data;
- learning to recognize any particular car or face from a single example, simply by watching *other cars or faces* as they move about;
- developing software for robots to continously explore the visual world and the interactions between vision and the other senses; and
- learning to recognize typewritten text in a font never seen before, without *ANY* training examples of that font.

The common thread in these efforts is that they *relieve the burden on the teacher of the computer*. The goal of my research is ultimately to develop computers that can be taught simply and rapidly, and that explore on their own. Perhaps if we achieve this goal, then computers will be able to learn enough skills to be generally useful, and to have the long sought after adaptability we call common sense.

This approach to computer vision requires a broad interdisciplinary background. In my education plan, I address three curriculum-related issues, including undergraduate and graduate teaching as well as curriculum couseling, to make sure students have the appropriate preparation to participate and contribute to this type of research. Outreach to minority and low-income students ensure that these ideas will reach the widest possible audience.