What Makes a Good Focal Problem?

Paul R. Cohen

April 23, 2008

Computing today is more integrative and problem-driven than the "textbook publisher's model" of twenty years ago, in which individual researchers would specialize in operating systems, databases, graphics, robotics, etc. At some universities, computing is being reorganized around research areas that integrate many of the old divisions of computer science with other areas of science, engineering and the arts. For instance, bioinformatics integrates molecular biology, genomics and proteomics with machine learning and data mining, computational geometry, databases, and so on.

I think we can do a bit better, in terms of focused, coherent research and the department's reputation, by identifying not only areas but *focal problems* within areas. Examples of focal problems are cars that drive themselves, providing high-quality K12 education over the web in many languages, tracking habitat loss and change, robotic companions for the elderly, increasing the carrying capacity of roads, fixing the healthcare record keeping system, marshaling the efforts of millions of recreational computer users to build a common-sense knowledge base for AI programs, killing spam, designing robust carbon trading schemes, water management for Arizona, and so on.

So, to me, good focal problems for 21st Century Computer Science have these attributes:

- They afford the opportunity for basic research in computer science.
- They are important to society. They are not proxies for important problems (the way Robot Soccer is a proxy for more important things to do with teams of mobile robots) but are *themselves* important problems.
- They can be defined with a few, nontechnical words.
- They have criteria and metrics for success, and, ideally, they have intermediate forms or require well-defined steps to solve them, so progress can be charted continuously.
- Substantial funding will remain available for several years to solve them.
- They establish mutual dependencies between different kinds of computer scientists and between computer science and other disciplines.

Three other attributes are pragmatic: A department must have researchers to do the work, the cost of entry should be low, and the problem affords leadership opportunities. Cars that drive themselves would probably fail on these criteria at Arizona: The department doesn't have roboticists, the cost of participation is very high (although much of the equipment is donated), and CMU and Stanford (and their industrial partners) are very far ahead. In other words, it would probably be difficult for the University of Arizona to excel at building cars that drive themselves.

My favorite focal problem is fixing K12 education, of which the International Internet Classroom is an element. It meets all the criteria: It affords opportunities for basic research in data mining, student modeling, curriculum planning, temporal databases, very large scale, robust systems, peer-to-peer content delivery, user interfaces, and so on. It is important to society. Several K12 problems can be concisely described in few words: Integrating English Language Learners, access to high quality K12 education, raising test scores in Arizona, and so on. The criteria for success are also easy to state in nontechnical terms: More kids go to college, fewer flunk high school, Arizona's math scores improve, ten million users, worldwide, and so on. Improving K12 education isn't a single problem, like beating Kasparov at chess, but has many facets and intermediate steps. For example, getting more kids through Algebra 1 is a really important step because in California, at least, kids can't graduate without this course, and flunking it is a strong predictor of dropping out. As to funding, it is pretty good, the Bush Administration has not been disastrous, and, perhaps more importantly, it diversifies a portfolio that, in Computer Science at Arizona, arguably depends too much on NSF. Much education funding is through the Institute of Education Sciences. For Computer Science at Arizona to play in this area, it will have to establish mutual dependencies with other faculty on campus in psychology, education and cognitive science, as well as with domain experts in the subjects we're trying to teach, such as Bill McCallum's organization. Pragmatically, the cost of entry is relatively low, in part because of Beal and Cohen's previous work and current funding, there are plenty of researchers in Computer Science who can contribute to the International Internet Classroom, and, because these ideas seem to be novel, the University of Arizona could establish leadership.

Whichever focal problems we choose, we must look outward to the university and beyond, and deliberately build networks of mutually-dependent researchers. These networks are both necessary to make progress on socially-important problems and help to focus attention — of sponsors, donors, other researchers, and the press — on the University of Arizona as the hub of important activities. Our faculty and students will benefit from the increased visibility and contacts afforded by these networks.