



Nice: Or What It Was Like to Be Mike's Student

Marti Hearst

There are three people who were pivotal to my success as a researcher, and Mike Stonebraker was the first of these—and also the tallest! I am very pleased to be able to share in this tribute to Mike, from my perspective as one of his former students.

Mike truly is a visionary. He not only led the way in the systems end of databases, but he was also always trying to bring other fields into DBMSs. He tried to get AI and databases to work together (OK, that wasn't the most successful effort, but it turned into the field of database triggers, which was enormously successful). He led early efforts to bring economic methods to DBMSs, and was a pioneer in the area of user interfaces and databases. I remember him lamenting back around 1993 that the database conferences would not accept papers on user interfaces when the work on Tioga [[Stonebraker et al. 1993b](#), [1993c](#)] was spurned by that community. That work eventually led to the Informix visualization interface, which again was ahead of its time. Shortly after that, Oracle had more than 100 people working in its visualization group.

Not only is Mike a visionary, but he is also inspirational to those around him. Back around 1989, when I was a graduate student who didn't find databases interesting enough for my dissertation and instead wanted to work on natural language processing (NLP), Mike said to me: "If you're going to work with text, think BIG text." I took that on as a challenge, despite the strange looks I got from my NLP colleagues, and as a result, was one of the first people to do computational linguistics work on large text corpora. That approach now dominates in the field, but was nearly unheard of at the time.

Another time, after I'd obtained significant research results with big text, Mike said something to the effect of, "What we need are keyboardless interfaces for text—why don't you solve that?" This led me to start thinking about visualization for interfaces for text search, which in turn led to several inventions for which I am now well known, and eventually to my writing the very first academic book on the topic, *Search User Interfaces* [Hearst 2009]. Again, it was Mike's vision and his special form of encouragement that led me down that path.

Mike also taught me that a world-famous professor wasn't too important to help a student with annoying logistical problems that were blocking research. In 1989, it was very difficult to get a large text collection online. I still remember Mike helping me download a few dozen *Sacramento Bee* articles from some archaic system in some musty room in the campus library, and paying the \$200 to allow this to happen.

I first met Mike when I was a UC Berkeley undergraduate who wandered into his seminar on next-generation databases. I needed a senior honors project, and even though he had just met me and I hadn't taken the databases course, Mike immediately suggested I work with him on a research project. He was the first and only CS professor I'd encountered who simply assumed that I was smart and capable. In retrospect, I think that Mike's attitude toward me is what made it possible for me to believe that I could be a CS Ph.D. student. So even though I suspect that sometimes Mike comes across as brusque or intimidating to others, toward students, he is unfailingly supportive.

As further evidence of this, in terms of number of female Ph.D. students advised and graduated from the UC Berkeley CS department, in 1995 Mike was tied for first with eight female Ph.D.s graduated (and he'd have been first if I'd stuck with databases rather than switching to NLP). I don't think this is because Mike had an explicit desire to mentor female students, but rather that he simply supported people who were interested in databases, and helped them be the very best they could be, whoever they were and whatever skills they brought with them. As a result, he helped elevate many people to a level they would never have reached without him, and I speak from direct experience.

Mike's enormous generosity is reflected in other ways. I still remember that when he converted the research project Postgres into a company (Illustra Corporation), he made a big effort to ensure that every person who had contributed code to the Postgres project received some shares in the company before it went public, even though he had no obligation whatsoever to do that. Although a few of us who contributed small amounts of code were overlooked until almost the very end, he insisted that the paperwork be modified right before the IPO so that a few more



Figure 34.1 Logo from the Postgres'95 t-shirt.

people would get shares. I find it hard to imagine anyone else who would do that, but Mike is extraordinarily fair and generous.

Mike is also very generous with assigning research credit. The aforementioned undergraduate research thesis had to do with the Postgres rules system. After I joined his team as a graduate student in 1987, Mike wrote up a couple of thought pieces on the topic [Stonebraker and Hearst 1988, Stonebraker et al. 1989] and insisted on including my name on the papers even though I don't believe I added anything substantive.

Mike's projects were very much in the tradition of the UC Berkeley Computer Science Department's research teams, consisting of many graduate students, some postdoctoral researchers, and some programming staff. Mike fostered a feeling of community, with logos and T-shirts for each project (see Figure 34.1 for one example) and an annual party at his house in the Berkeley hills at which he gave out goofy gifts. Many of his former students and staff stay in touch to various degrees, and, as is common in graduate school, many romantic relationships blossomed into eventual marriages.

So that's Mike Stonebraker in a nutshell: visionary, inspirational, egalitarian, and generous. But surely you are thinking: "Hey, that can't be the whole story! Wasn't Mike kind of scary as a research advisor?" Well, OK, the answer is yes.

I still remember the time when I was waffling around, unable to decide on a thesis topic with my new advisor, and so I made an appointment to talk with my now former advisor Mike. For some reason, we'd scheduled it on a Saturday, and

it was pouring outside. I still remember Mike appearing at the office and looking down at me from his enormous height and basically saying something like, “What’s wrong with you? Just pick a topic and do it!” From that day on, I was just fine and had no problem doing research. I have found that for most Ph.D. students, there is one point in their program where they need this “just do it” speech; I’d otherwise never have had the guts to give it to students without seeing how well this speech worked on me.

I also remember the extreme stances Mike would take about ideas—mainly that they were terrible. For instance, I was there for the ODBMS wars. I remember Mike stating with great confidence that object-oriented was just not going to cut it with databases: that you needed this hybrid object-relational thing instead. He had a quad chart to prove it. Well, he hadn’t been right about putting expert systems into databases, but he certainly ended up being right about this object-relational thing (see Chapter 6).

As with many great intellects, Mike very much wants people to push back on his ideas to help everyone arrive at the best understanding. I remember several occasions in which Mike would flatly state, “I was utterly and completely wrong about that.” This is such a great lesson for graduate students. It shows them that they have the opportunity to be the one to change the views of the important professor, even if those views are strongly held. And that of course is a metaphor for being able to change the views of the entire research community, and by extension, the world (see Chapter 3).

As I mentioned, Mike is a man of few words, at least over email. This made it easy to tell when you’d done something really, truly great. Those of you who’ve worked with him know that treasured response that only the very best ideas or events can draw out of Mike. You’d send him an email and what you’d see back would be, on that very rare occasion, the ultimate compliment:

neat.
/mike