Technical Opinion | Nigel Ward

The (Un)Predictability of Computer Science Graduate School Admissions

Providing a general calculator for students to compare where they are likely to get accepted for graduate school.

Today, most applicants to graduate programs cannot accurately predict where they will be accepted. To assist them in this process we have created a tool for students, the Acceptance Estimator, which is available at www.cs.utep.edu/admissions/.

THREE APPLICANT STORIES
KN had always planned to go on to graduate school, but only after working for a few years first. When a reassignment took him away from his chosen career path, he began to apply to graduate schools. As he was working 65 hour weeks, he didn't make a fuss; rather he just applied to four well-known departments where he met all the published criteria. His plan was to stay on at work until his project was completed and then deftly transition to graduate school with no down time. Unfortunately, all four schools rejected him.

Later, after he left his job, he had time to come see me to discuss what had gone wrong. Obviously, it was the 390 Verbal score that effectively disqualified him at any top-50 school. But this was only obvious

to me; he had had no way to know that. (As a result he had to sit out until the next application cycle, spending the time doing part-time jobs and studying for the GRE.)

The problem for KN, as for thousands of other applicants, is that many departments provide little or no useful information on their admissions criteria. Even those departments that do give quantitative information regarding GPAs and GREs use such a variety of reporting methods—including averages, minimums, cut-offs, soft minimums, nominal scores, desired scores, median scores, average percentiles, minimum sums of scores, and so on—that it is the rare applicant who can figure out what these imply for the only real question: who will be accepted.

Why do departments fail to publish useful information? Is it beneath their dignity? Or do they strategically encourage the wrong people to apply, in order to be able to boast of high rejection rates? Maybe sometimes. However, departments also face an honest dilemma.

MM was an undergraduate

whose work habits, motivation, and accomplishments indicated to her advisor that she would do well in graduate school, even though her GRE scores were far below the norms published on the department Web site. Her advisor strongly encouraged her to apply to the graduate program. After some discussion among the faculty about the relative importance of GREs and other factors, MM was admitted. (She did very well.) Learning from the experience, the department decided it was too risky to rely only on faculty champions to pull in such students, and instead got the word out to the undergraduate population that no one should be discouraged from going to graduate school because of grades and GREs alone.

About a year later, FW applied. He didn't know the particulars, and being somewhat shy he hadn't actually checked with the graduate advisor, but he had heard that people with low numbers were getting accepted, and he knew that one of his professors liked him so he was confident about getting a nice letter of recommendation. When he was rejected he felt that he had been misled. (In the end he found a good

job, albeit six months later than he could have.)

The dilemma is this: on the one hand departments don't want to discourage atypical but desirable applicants by appearing to stress the qualitative factors, but on the other hand they don't want to be so vague about criteria as to mislead students who are not cut out for graduate school.

TOWARD A METRIC OF APPLICANT STRENGTH

Clearly what is needed is a metric of applicant strength that includes not only the GREs but also subjective factors. With such a metric departments could avoid the dilemma and provide specific and accurate information about their admissions policies.

Is such a metric possible? At the University of Texas at El Paso we decided to attempt to develop one. Doing so involved a lot of grunt work—doing data mining and knowledge elicitation to select values for a few dozen detailed parameters encoding answers to such questions as: Is the GRE Q more important than the GRE AW? How many GRE points are worth one GPA point? How informative are GPAs from Indian schools? What GPA adjustment is reasonable for non-CS majors?— but also generated some interesting problems.

First was the question of how to combine the qualitative factors. One option is to use a sum of GRE scores, based on the idea that a strength on one dimension can compensate for a weakness on another. The other is to just measure "height" above the minimum requirement on each dimension, based on the idea that the weakest skill will be the limiting factor. There are good arguments for using each approach, but fortunately we found a middle road, an "order-weighted average" operator [1] between the average and the minimum.

Second was the question of how to quantify the impact of letters of recommendation. While this is arguably impossible in principle, admissions committees do somehow manage to weigh GREs versus letters. A model that does not do the same will fail to give any guidance on whether an applicant's letters will overcome a weak GPA and GREs, or conversely. We decided that three factors were needed: the warmth, the believability of the recommender, and the recommender's basis for judgment, with the latter two multiplied to give the weight of the letter compared to the other factors. While applicants may still find it difficult to estimate the warmth of the letter, the other two factors are relatively objective, which allows the model to at least give a useful upper bound on how much the letters are likely to count.

Did the resulting model work? On a test set of 55 applicant packets, it correctly predicted 50 accept/reject decisions. One of the incorrect predictions was a borderline case that could have gone either way. Another was due to special circumstances not handled by the model. In the remaining three cases, all due to the same parameter, the model predicted rejection but the committee had accepted. However, follow-up showed that those three students had later all dropped out, so in a sense the model was an improvement on the collective seasoned wisdom of the admissions committee.

While the model turned out too complex for casual users to work through, fortunately it lent itself to implementation in the form of a calculator on the Web, available at the URL listed at the beginning of this column. Potential applicants are encouraged to enter their GRE scores and GPA and use the pull-down menus to estimate the impact of the letters of recommendation, and other factors. Those receiving scores of -25 or higher should apply, since they will almost certainly be accepted.

TOWARD A SOLUTION

Of course, some students are interested in graduate study at schools other than the University of Texas at El Paso. Having the model, we decided to attempt predictions of acceptance decisions at other departments.

Using information on the Web, we estimated the threshold score for each of the 73 other departments that publish useful quantitative information [2]. This was incorporated into the Estimator, enabling a potential applicant to quickly get a list of departments where acceptance is likely, without needing to know how to interpret oblique statements referring to average percentile, soft minimum, average sum, and other obscure statistics.

Of course, this is a stopgap; there is only so much that can be done without inside information on each department's admissions decision making. However, anecdotal reports suggest the Estimator is used and useful.

My hope is that more departments will find ways to better inform potential applicants of their chances; perhaps by referring to this Web site, or by adapting the metric and providing their own calculators, or even just by publishing better information. Doing so will enable more students to make smoother transitions to graduate school. **c**

REFERENCES
1. Carlsson, C., Fuller, R. and Fuller, S. OWA operators for doctoral student selection problem. In R.R. Yager and J. Kacprzyk, Eds., *The Ordered Weighting Averaging Operators: Theory and Applications*. Kluwer 1997, 167–177.
2. Ward, N. Towards a model of computer science graduate admissions decisions. *Journal of Advanced Computational Intelligence and Intelligent Informatics 10* (2006), 372–383.

NIGEL WARD (nigel@utep.edu) is an associate professor of Computer Science at the University of Texas at El Paso.