A Survey of Graduate Programs in Computer Science

Richard Conway

TR 78-357
November 1978

Department of Computer Science
Cornell University
Ithaca, New York 14853

This report presents the results of a survey of computer science graduate programs conducted in September and October of 1978. A questionnaire was sent to the department chairman at each of the seventy-one "Universities Offering Ph.D. programs in Computer Science", based on the 7-13-78 list from the Snowbird Conference. This list turned out to include several universities that do not, in fact, offer a Ph.D. in computer science and may have omitted some that do, but it nevertheless was the basis of the study.

Each chairman was invited to rate (not rank) each of the seventy-one programs with respect to "quality of graduate faculty", using the following scale:

5 for "distinguished"
4 for "strong"
3 for "good"
2 for "adequate"
1 for "marginal"
0 for "not sufficient"
blank for "inadequate information to rate"

They were also asked to rate each institution with respect to "effectiveness of doctoral program", using the following scale:

3 for "extremely attractive"
2 for "attractive"
1 for "acceptable"
0 for "not attractive"
blank for "inadequate information to rate"

They were also asked to estimate the first derivative of each program, using the following symbols:

+ for "gaining in quality"
0 for "not changing"
- for "declining in quality"
blank for "inadequate information to rate"

This terminology was borrowed from the Roose-Andersen "Rating of Graduate Programs" for the American Council on Education (1970), but there the similarity ends. That study was broadly based and carefully controlled--this one just went to chairmen of the programs being rated. The Roose-Andersen study can be criticized as favoring programs with many graduates; this study perhaps reduced that risk, but surely introduced many others.

In addition to ACE-like ratings, each chairman was asked to "list institutions that you would recommend as having especially strong programs in" seven different subject areas within computer science:

1. Theory of computation, complexity, algorithms
2. Programming systems, languages and methodology
3. Operating systems
4. Machine organization and hardware systems
5. Numerical analysis
6. Database systems
7. Artificial intelligence.

There are many reasons why these results ought not to be taken too seriously. To begin with, I made no attempt to achieve consistency of interpretation by trying to define "quality" or "effectiveness", or the descriptors "distinguished", "strong", "attractive", etc. I did not suggest how the ratings should be produced. It was apparent from the letters that some responses were the chairman's personal opinion, some represented a systematic departmental consensus, and in at least one case the chore of replying was delegated to an unfortunate assistant professor. I did not attempt to define what should be included in a "program" to be rated. At many institutions some relevant faculty may be in mathematics or electrical engineering departments, as well as in the nominal computer science program. Many schools also have research institutes which contribute strength to a graduate program. But I left all these uncertainties to the respondent's discretion.

The results should also be suspect just because the study was conducted by an individual at one of the institutions being rated. Such studies, if they ought to be done at all, should be carried out by some impersonal, objective organization. But such organizations do not respond quickly, and I felt that the virtue of immediate action at least partially offset the effrontery of individual effort. Cornell's rating may have suffered from the respondents' resentment of my initiative, or it may have benefited from embarassment at letting me see how little they really thought of our program.

For all these reservations, the results are undeniably interesting, and even schools that declined to participate may not be able to resist looking at a summary of their peers' opinions.

Thirty-four responses were received. Surprisingly, most responses rated most of the programs, and more than half of the responses rated all the programs.

Thirty of the responses included a rating of their own program--which provides material for an interesting substudy in itself. No school underestimated the esteem with which it is regarded by peer institutions. The extent of overestimate, in comparison to the

average of peer institution ratings, varied from 0.03 to 2.90 (on the 5-point "quality" scale). The largest difference was by the school which was ranked 70th by average quality rating, and whose self-rating would have put it in 7th place. Curiously, highly rated schools seemed almost as optimistic in their self-evaluation as lower rated schools. All self-ratings are omitted from the data given below.

Apart from self-ratings, the responses were remarkably consistent. There was rarely more than a one point range in the quality rating of a given program in different responses.

I have attempted to present the data so that it is not possible to identify which schools participated in the survey, much less determine what individual responses were. However, I should note that the responses generally came from the schools that turned out to be relatively highly-rated by the study. This suggests a certain circularity in the process, and you should be aware of this possibility as you examine the results.

Quality of Graduate Faculty

The primary result of the survey concerns the quality of graduate faculty as perceived by other institutions. The average of these ratings is given in Table 1. Fifty-three schools had an average rating of 2.00 or more, corresponding to "adequate" quality. These fifty-three schools are ranked by average perceived quality, and divided into four groups with breakpoints of 3.50, 3.00, and 2.50. The detailed ranking is shown only for the first group; the other groups are listed alphabetically.

Table 1 also shows the "net derivative" for each of these schools. If the number of responses indicating a school is "gaining in quality" exceeds the number indicating "declining in quality", the net derivative is shown as a "+"; if the opposite is true the net derivative is "-". If the number is equal, a "0" is shown.

For twenty-six schools the net number of raters (difference between "gaining" and "declining") exceeded 15% of the number of responses that gave a derivative for that school. These schools are listed in four groups in Table 2. Note that these listings include several schools that do not appear in Table 1.

Recall that the questionnaire elicited the sign, but not the magnitude of the derivative. Consequently, for a school to appear in the +30% group does not necessarily imply that the raters thought it was gaining more rapidly than a school in the +15% group--only that the progress is more widely recognized. The same is true for the "declining" groups.

Table 1 Average Perceived Quality of Graduate Faculty

	Average Quality Rating	Net Derivative	Subject Area Recommendations (Table 3) Max Group	Nbr Subjects
Group Q1: 10 schools (ranked by avg)	5-3.5			
1. Stanford	4.97	0	S1	7
2. MIT	4.76	-	S1	7
3. Carnegie-Mellon	4.70	-	S2	6
4. Cornell	4.52	+	S1	5 (4)
5. Cal-Berkeley	4.47	+	S2	6
6. Illinois	4.06	-	S2	5
7. Toronto	3.93	+	S4	6 (5)
8. UCLA	3.69	+	S5	3
9. Purdue	3.65	-	S4	3
10. Waterloo	3.50	+	S4	4 (3)
Group Q2: 9 schools (alphabetical)	3.5-3.0			
Maryland		-	S5	4 (3)
Michigan		-	S5	4
NYU		-	S5	2
Princeton		-	S4	3 (2)
Texas		-	S4	6 (5)
Utah		-	S5	2
Washington		+	S5	3
Wisconsin		-	S4	4 (3)
Yale		-	S5	5 (4)
Group Q3: 16 schools (alphabetical)	3.0-2.5			
Brown		+	S6	2 (0)
Cal-Irvine		-		
Cal-Santa Cruz		0	S6	1 (0)
Cal Tech		+	S5	3 (1)
Harvard		-	S5	3 (1)
Massachusetts		-	S6	1 (0)
Minnesota		0	S6	1 (0)
North Carolina		0	S6	1 (0)
Northwestern		0	S6	1 (0)
Ohio State		-	S5	1
Penn State		-	S5	2 (1)
Pennsylvania		-	S5	2 (1)
Rochester		+	S5	1
SUNY-Buffalo		-	S6	1 (0)
SUNY-Stony Brook		+		
USC		+	S5	3 (2)
Group Q4: 18 schools (alphabetical)	2.5-2.0			
Arizona		+		
British Columbia		+	S5	1
Cal-Santa Barbara		+	S6	1 (0)
Case-Western		-		
Colorado		+	S6	1 (0)
Columbia		+		
Duke		0		
Georgia Tech		+		
Indiana		+	S6	1 (0)
Iowa		-		
Iowa State		0		
Kansas		+		
Michigan State		0		
Pittsburgh		0		
Rutgers		+	S6	1 (0)
Syracuse		-	S6	1 (0)
Virginia		+		
Wash (St. Louis)		0		

Table 2 Perceived Direction of Change

Net Percentage of Responses indicating:

"Gaining in Quality"		"Declining in Quality"	
> 30%	> 15%	> 15%	> 30%
Arizona	Brown	Harvard	Case-Western
Cal-S Barbara	Colorado	IIT	Ohio State
Cal Tech	Columbia	Massachusetts	Penn
Rochester	Delaware	Michigan	Penn State
USC	Georgia Tech	Wisconsin	
Washington	Indiana	Yale	
	Kansas		
	UBC		
	VPI		
	Waterloo		

(schools are listed alphabetically)

The remaining columns in Table 1 are derived from the subject area recommendations and are explained in the next section.

Subject Area Recommendations

The recommendations of "especially strong programs" in seven subject areas of computer science are summarized in Table 3. Schools are grouped according to the percentage of responses in which they were recommended. All self-recommendations were excluded from the data.

Groups S1 to S5 are in equal increments of 20%. Group S6 was added to segregate schools that were mentioned on only one response. Recall that the questionnaire did not suggest how many schools might be recommended in a given area, and there was significant variation between responses in this regard. There was also variation between different subject areas. Note that the responses were not asked to rank the schools in a given subject--only to recommend or not. Consequently, the percentages in Table 3 indicate only how widely a school is well-regarded, and not how highly it is regarded in comparison to other schools.

The results of Table 3 are summarized in Table 1 in two different ways. The "Max Group" column indicates the highest percentage group for any subject area. This might suggest how widely recognized the school's best subject is at other schools. The "Nbr Subjects" column indicates in how many of the seven subject areas the school was recommended. (A second parenthesized figure is given for the number of subjects if Group S6 is excluded.) This provides some indication of the perceived breadth of quality.

The subject area recommendations provide an interesting check on the rationality of the general quality ratings. Both the maximum group and number of subjects rather strongly support the ranking by average quality rating.

Table 3 is also instructive with regard to the state of the various subject areas. For example, the "systems" side of computer science does not enjoy nearly as clear a consensus as to which are the strongest programs as does the "theory" side. (It was also remarkable how many schools recommended themselves in systems but did not have this recommendation seconded by any other institution.) The scattered recommendations in database systems are not surprising, considering how recently this subject has claimed the attention of computer science departments, but the relatively scattered recommendations in numerical analysis were unexpected.

Effectiveness of Doctoral Program

Following the pattern of the ACE study, the questionnaire asked for ratings of the "effectiveness of the doctoral program" on a three-point scale. Forty-eight schools had an average rating of 1.00 or greater, corresponding to "acceptable" on the questionnaire. These schools were ranked by their average effectiveness rating, and divided into four groups, as shown in Table 4.

In general, the effectiveness ranking roughly coincides with the quality ranking. However, five schools--Colorado, Minnesota, North Carolina, Syracuse and Wisconsin--ranked significantly higher in perceived effectiveness than in perceived quality of faculty.

Responses from Q1, Q2 and Q3 Schools

Many people suggested that the schools strongly committed to research might have a different perception of quality than the community in general. To test this supposition, I partitioned the schools based on the average quality ranking given in Table 1. Groups Q1, Q2 and Q3 totaled thirty-five schools--just half the total list. These schools provided twenty-one of the thirty-four reponses.

Table 5 gives the results with respect to average quality, if only the responses from the Q1, Q2 and Q3 schools (Table 1) are considered. The "net derivative" in Table 5 has the same interpretation as in Table 1. The rightmost column indicates the school's ranking in Table 5 relative to Table 1.

The actual values of the average quality rankings in Table 5 are interesting. Overall, the ratings are slightly harsher--only forty-four schools have an average rating of 2.00 or higher, as compared to fifty-three in Table 1. Although there were interesting

Table 3 Recommendations of Strong Programs by Subject

Schools grouped by percentage of responses in which they were recommended
(listed alphabetically within groups)

Subject	Group S1 > 80%	Group S2 ≥ 60%	Group S3 > 40%	Group S4 > 20%	Group S5 ≤ 20%	Group S6 single mention
Theory, complexity, algorithms	Cornell Stanford	Berkeley	MIT	Carnegie Princeton Toronto	Illinois Michigan Penn State Washington Waterloo Yale	Harvard Minnesota Syracuse USC Wisconsin
Programming, languages, methodology			Carnegie MIT	Cornell Stanford Waterloo	Berkeley Harvard Illinois Michigan NYU Purdue Texas Toronto UCLA USC Yale	Brown Indiana Maryland Penn Princeton Santa Cruz Syracuse
Operating systems			Carnegie MIT	Purdue	Berkeley Michigan Stanford Texas Toronto UCLA USC Washington Wisconsin	Cal Tech Cornell UNC Penn State S Barbara Waterloo Yale
Machine org., hardware		Illinois	Carnegie	MIT	Berkeley Cal Tech Stanford UCLA Utah	Mass NWestern Texas Toronto Washington
Numerical analysis		Stanford		Illinois Texas Wisconsin	Berkeley Carnegie Cornell Maryland MIT NYU Purdue Toronto Waterloo Yale	Brown SUNY-Buf
Database systems				Berkeley Toronto	Cornell Maryland Michigan MIT Ohio State Penn Princeton Stanford Texas Utah	Colorado Harvard Syracuse
Artificial intelligence	MIT Stanford	Carnegie			Illinois Maryland Rochester Texas UBC Utah Wisconsin Yale	Cal Tech Penn Rutgers Santa Cruz

Table 4 Average Perceived Effectiveness of Doctoral Program

(Schools marked * are ranked significantly higher than for Quality of Graduate Faculty.)

Group E1: 9 schools 3-2.0 (ranked by avg)

1.	Stanford	2.76
2.	MIT	2.72
3.	Carnegie-Mellon	2.69
4.	Cornell	2.53
5.	Cal-Berkeley	2.45
6.	Toronto	2.31
7.	Illinois	2.22
8.	Purdue	2.07
9.	*Wisconsin	2.00

Group E2: 9 schools 2.0-1.5 (alphabetical)

- Maryland
- Michigan
- *Minnesota
- Princeton
- Texas
- UCLA
- Washington
- Waterloo
- Yale

Group E3: 12 schools 1.5-1.3 (alphabetical)

- Cal Tech
- *Colorado
- Harvard
- Massachusetts
- *North Carolina
- NYU
- Ohio State
- Rochester
- SUNY-Stony Brook
- *Syracuse
- USC
- Utah

Group E4: 18 schools 1.3-1.0 (alphabetical)

- Arizona
- British Columbia
- Brown
- Cal-Irvine
- Cal-Santa Barbara
- Cal-Santa Cruz
- Duke
- Georgia Tech
- Indiana
- Iowa
- Iowa State
- Michigan State
- Northwestern
- Penn State
- Pennsylvania
- Rutgers
- SUNY-Buffalo
- Washington (St. Louis)

Table 5 Average Perceived Quality of Graduate Faculty

(based on Table 1 Q1, Q2 and Q3 reponses)

	Average Quality Rating	Net Derivative	Change relative to Table 1
Group R1: 9 schools (ranked by avg)	5.0-3.5		
1. Stanford	5.00	0	same
2. MIT	4.86	-	same
3. Carnegie-Mellon	4.85	-	same
4. Cornell	4.60	+	same
5. Cal-Berkeley	4.50	+	same
6. Toronto	4.00	+	higher
7. Illinois	3.85	-	lower
8. Yale	3.55	-	higher
8. UCLA	3.55	+	same
Group R2: 8 schools (alphabetical)	3.5-3.0		
Maryland		-	higher
Michigan		-	lower
Princeton		-	lower
Purdue		-	lower
Texas		-	lower
Washington		+	higher
Waterloo		+	same
Wisconsin		-	higher
Group R3: 11 schools (alphabetical)	3.0-2.5		
Brown		+	higher
Harvard		-	lower
Massachusetts		-	lower
Minnesota		0	higher
North Carolina		0	higher
NYU		-	lower
Pennsylvania		-	lower
Rochester		+	higher
SUNY-Stony Brook		+	same
USC		+	higher
Utah		-	lower
Group R4: 16 schools (alphabetical)	2.5-2.0		
Arizona		+	lower
British Columbia		+	higher
Cal-Irvine		-	higher
Cal-Santa Cruz		-	lower
Cal Tech		+	higher
Colorado		+	lower
Duke		+	higher
Iowa State		-	lower
Northwestern		-	lower
Ohio State		-	lower
Penn State		-	lower
Pittsburgh		0	higher
Rutgers		+	higher
SUNY-Buffalo		-	higher
Syracuse		0	lower
Washington (St. Louis)		0	higher

local rearrangements in the list, the general pattern in Table 5 is not much different from Table 1. However, eight schools--Alberta, British Columbia, Delaware, Duke, Hawaii, North Carolina, Washington (St. Louis) and Yale are significantly higher-ranked in the Table 5 data.

Table 6 summarizes the perceived changes in quality in a manner analogous to Table 2. It is interesting that Table 6 shows more highly-ranked schools perceived as "declining in quality" than does Table 2.

Comparison to Previous Bitner Survey

A year ago James Bitner of the University of Texas asked the chairmen of twenty-one computer science departments to rank those twenty-one departments. He distributed the responses in August, 1978. His ranking is shown in Table 7, along with the corresponding rankings from Table 1 and Table 5.

For the particular schools Bitner selected, there is general agreement even though the rankings were produced by quite different processes. An exception is the University of Washington, which fares much better in current study.

The origin of Bitner's list of schools is not obvious. In comparison with Table 1, he omitted Toronto and Waterloo from Q1, NYU, Princeton and Utah from Q2, Brown, Cal-Irvine, Cal-Santa Cruz, Cal Tech, Massachusetts, Minnesota, North Carolina, Rochester,

Table 6 Perceived Direction of Change

(based on Q1, Q2 and Q3 responses)

Net Percentage of Responses indicating:

"Gaining in Quality"		"Declining in Quality"	
> 30%	> 15%	> 15%	> 30%
Arizona	Alberta	Cal-Irvine	Case-Western
Brown	Cal-Berkeley	Carnegie	Ohio State
Cal-S Barbara	Columbia	Harvard	Penn
Cal Tech	Delaware	Illinois	Penn State
Colorado	Georgia Tech	IIT	Texas
Rochester	UBC	Michigan	Yale
USC	UCLA	MIT	
Washington	Waterloo	Purdue	
		Utah	
		Wisconsin	

(schools are listed alphabetically)

SUNY-Buffalo and SUNY-Stony Brook from Q3, and similarly, many schools from Q4. In fact, more than twenty-one additional schools must be interleaved into Bitner's list to match Table 1, so he somehow selected fewer than half of the stronger schools for his ranking.

Table 7 "Bitner Ranking" of Selected Schools

		Ranking from: Table 1		Table 5	
1.	Stanford	Q1	1	R1	1
2.	Carnegie	Q1	3	R1	3
3.	MIT	Q1	2	R1	2
4.	Cornell	Q1	4	R1	4
5.	Cal-Berkeley	Q1	5	R1	5
6.	Illinois	Q1	6	R1	7
7.	Texas	Q2		R2	
8.	Purdue	Q1	9	R2	
9.	UCLA	Q1	8	R1	8
10.	Yale	Q2		R1	8
11.	Michigan	Q2		R2	
12.	Wisconsin	Q2		R2	
13.	Maryland	Q2		R2	
14.	Ohio State	Q3		R4	
15.	Harvard	Q3		R3	
16.	Penn State	Q3		R4	
17.	USC	Q3		R3	
18.	Washington	Q2		R2	
19.	Pennsylvania	Q3		R3	
20.	Northwestern	Q3		R4	
21.	Michigan State	Q4			

Conclusions

Some people apparently feel that this type of information is not useful, but at least it provides prospective graduate students and prospective faculty members with an interesting check on the prejudices of their local advisors. Others may feel that however useful, such a study should be conducted only under the aegis of some objective professional society. I agree, and hereby retire from the survey field having, I hope, achieved the objective of replacing the Bitner ranking with a broader and fairer comparison. For the benefit of my successors, if any, I suspect it is possible to invest a great deal more effort in such a study with only marginal improvement in the utility or credibility of the results.

Let me underscore once more that such a study addresses only perceived quality. But the perceptions of other institutions are not unimportant. I suspect many of us will discover from these results that we have not done an adequate job of informing other schools of our true distinction. In fact, several schools have already embarked on an educational campaign. But note that this seeks to change the perception rather than the quality.

There is always the possibility that in a few isolated cases, the perception might be closer to the truth than self-evaluation. If just a few schools recognize that possibility and use this information to justify strengthening their program, then this whole exercise will have been worthwhile.