

On Bias in Social Reviews of University Courses

Taha Hassan

Computer Science Department, Virginia Tech
Blacksburg, VA
taha@vt.edu

ABSTRACT

University course ranking forums are a popular means of disseminating information about satisfaction with the quality of course content and instruction, especially with undergraduate students. A variety of policy decisions by university administrators, instructional designers and teaching staff affect how students perceive the efficacy of pedagogies employed in a given course, in class and online. While there is a large body of research on qualitative driving factors behind the use of academic rating sites, there is little investigation of the (potential) implicit student bias on said forums towards desirable course outcomes at the institution level. To that end, we examine the connection between course outcomes (student-reported GPA) and the overall ranking of the primary course instructor, as well as rating disparity by nature of course outcomes, for several hundred courses taught at Virginia Tech based on data collected from a popular academic rating forum. We also replicate our analysis for several public universities across the US. Our experiments indicate that there is a discernible albeit complex bias towards course outcomes in the professor ratings registered by students.

KEYWORDS

bias; academic forums; social ranking; student satisfaction; course outcomes

ACM Reference Format:

Taha Hassan. 2019. On Bias in Social Reviews of University Courses. In *11th ACM Conference on Web Science Companion (WebSci '19 Companion)*, June 30–July 3, 2019, Boston, MA, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3328413.3328416>

1 INTRODUCTION

Online forums for rating university course instructors, like *Rate-MyProfessors* [1] and *Koofers* [7] have inspired considerable research attention over the years [4] [2] [11] [10]. An aggregate of content sharing, networking and recommendation services, these forums cater to a number of their contributors' needs, including but not limited to information seeking, gratification, and convenience [10]. Prior studies have identified several broad themes in the student

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebSci '19 Companion, June 30–July 3, 2019, Boston, MA, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-6174-3/19/06...\$15.00

<https://doi.org/10.1145/3328413.3328416>




Figure 1: Correlation between the average student-reported GPA for a course and overall instructor rating, as a function of the minimum number of ratings considered

feedback sampled from these forums, including teacher personality, aptitude and preparation, ease of access to help and feedback from course staff, and perceived practicality of the course rubric [8]. There is however, lesser attention devoted in this literature, to institution-level correlates of student perception. Empirical investigations of the potential sources of student bias in said perception are often divided in their conclusions, because of limitations of sample sizes or meta-variable space [12] [3] [6]. Assessing its reliability at scale can lend insights to instructional designers, department administrators and instructors alike on the limitations of existing pedagogies. It can also potentially extend the utility of university-managed end-of-semester course evaluations, and help improve the usability, relevance, accessibility and trustworthiness [11][9] of its host forums.

2 APPROACH

We pursue preliminary evidence of what appears to be a modest to strong relative connection between aggregate course outcomes and student perception of the course instructor. Figure 1 visualizes the correlation between these two for 478 courses taught at Virginia Tech. As we increase the minimum number of ratings per course considered towards the correlation, this correlation increases and achieves a steady average value of about 0.47 beyond 20 ratings per course. We explore this further by examining the disparity of instructor ranking between high, medium and low GPA student groups, as well as regressing student approval against course outcomes and perceived difficulty of various course instruments.

Table 1: Key counts for courses in the historical ratings dataset

Institution	Courses	Departments	Ratings
MSU	888	110	67086
VT	597	80	42503
CAL-POLY	578	67	36900
CSU	430	74	20506
UGA	367	99	22992
UNC-CH	84	34	2180
GMU	75	37	1866
JMU	72	33	2158
UKY	69	30	1537
NCSU	61	29	1923

VCU	39	20	939
UVA	24	14	581
ASU	13	10	236

3 EVALUATION

3.1 Datasets

We scraped course metadata from Koofers [7], a popular forum for sharing course content and instructor reviews (see table 1). For a university, we consider all courses with a minimum of 10 ratings. A course can have a multitude of instructors and offerings. Instructor ratings on Koofers are on a 0 to 5 scale, GPA reports are on a 4.0 scale. We define the minimum acceptable use of the forum (at the institution level) as an excess of 1000 total ratings, with the course count at least twice that of the department count.

3.2 Methods

We significance-test the disparity in professor ratings by GPA groups using one-way ANOVA (F-test, table 2). We also use the Python package *statsmodels.OLS* towards regression analysis of instructor ratings as a function of student GPA, as well as their self-reported ease of course instruments (exams, quizzes, projects and homeworks).

3.3 Results

Table 2 lists the correlations between average professor rating and student GPA for all institutions considered, as well as average ratings for each GPA group. The top four institutions with the largest set of ratings (MSU, VT, CAL-POLY, and CSU) register a modest correlation (between 0.2 and 0.35) and large group disparity by GPA. MSU, for instance, reports the highest disparity ($F(2, 885) = 14.7, p = 5e - 7$) against $F_{crit} = 3$. The next four institutions (UGA, UNC, GMU and JMU) report stronger correlations (between 0.4 and 0.5), with weaker but significant group disparities. It is also instructive to consider the differences between the two groups of courses. The group with minimum acceptable use of the forum almost consistently registers a correlation between outcomes and student perception of the primary instructor, as well as a disparity between ratings of the GPA groups. The group with less than minimum use of the forum generally does not. Overall and

department-level sample sizes decline fairly rapidly, especially for the final five institutions in table 2, which we believe, contributes to the larger uncertainty in the corresponding measurements, empty GPA groups and smaller effects sizes if any.

Table 3 and table 4 report the coefficients, errors and significance tests for regression analysis with average instructor rating as a function of average student-reported GPA **X1**, average perceived difficulty of exams, quizzes, projects and homeworks (**X2** through **X5**). Course outcomes outweigh the perceived ease of course evaluations in their aggregate effect on instructor ratings ($t = 6.0, t = -2.3$ for **X1** and **X3**, respectively). Difficulty of quizzes appears modestly relevant in ascertaining the overall student satisfaction with the course, with higher difficulty linked to lower student approval (t-statistic is negative). An in-depth analysis of frequent course rubrics and their student approval is left for future work.

4 DISCUSSION

Several studies over the last few decades have reported a modest correlation between student evaluations and their grades, albeit for individual students or a course [15] [3] [5] [6]. A review by Feldman [6] reports this correlation to be somewhere between +0.1 and +0.3. Our preliminary inquiry at the institution level demonstrates that this correlation matches, and often exceeds said figures. The aforementioned work notes that given the learning acquired by students during the course of a given class or academic term, all of this observed correlation can not necessarily be a result of implicit, time-invariant student bias towards course outcomes. The multi-faceted nature of student perception can affect this connection. This complexity is echoed in a comparative study of in-class assessments, and pre- and post-assessment ratings on RateMyProfessors [11]. The study reported how pre-assessment course ratings on instructor clarity were significantly lower than both in-class and post-assessment reviews. However, instructor easiness was reviewed lower in-class relative to online. Our study attempts to initiate a line of large-scale contextual inquiry of these ratings across institutions that can potentially help consolidate these differing interpretations. To that end, we discuss some limitations of our study design, and plans for future research as follows.

4.1 Limitations

We intend to expand our analysis by considering the time order of the instructor ratings in our dataset. While the magnitude of the observed correlation between course outcomes and student rankings is nearly consistent across institutions (with minimum aggregate forum use) we tested for, it is harder to argue about the directionality of this correlation without said data. Another critical deficiency in our approach is the assumption of linearly independent course characteristics. Fixes include dimensionality reduction [14] and modeling the joint reliability of course features using simple Bayesian networks [13] and are left for future work.

4.2 Future Work

Beyond algorithmic improvements, we are working on expanding our dataset to include a larger number of academic institutions, and

Table 2: Hypothesis-testing the relationship b/w course outcomes and instructor rankings

Institution	Corr, p	$F(df_1, df_2), p$	F_{crit}	μ_{ov}, N_{ov}	μ_{hi}, N_h	μ_{med}, N_m	μ_{low}, N_l
MSU	0.24, 3.1e-13*	14.7 (2, 885), 5e-7 [†]	3.0	3.82, 888	3.88, 64	3.85, 741	3.46, 83
VT	0.33, 4e-17*	15.3 (2, 594), 3.1e-7 [†]	3.01	3.83, 597	4.15, 30	3.86, 492	3.52, 75
CAL-POLY	0.23, 2.8e-8*	16.7 (2, 575), 8.8e-8 [†]	3.01	3.89, 578	4.39, 17	3.92, 397	3.78, 164
CSU	0.27, 6.8e-9*	15.6 (2, 427), 2.7e-7 [†]	3.01	3.72, 430	4.14, 16	3.79, 291	3.5, 123
UGA	0.41, 3e-16*	19.4 (2, 364), 9e-9 [†]	3.02	3.84, 367	4.32, 43	3.87, 302	3.59, 22
UNC-CH	0.48, 4.8e-6*	6.6 (2, 81), 2e-3 [†]	3.1	3.69, 84	4.65, 2	3.73, 68	3.35, 14
GMU	0.5, 4.5e-6*	8 (1, 73), 6e-3 [†]	3.12	3.78, 75	-	3.95, 42	3.57, 33
JMU	0.46, 4.2e-5*	3.8 (2, 69), 2e-2 [†]	3.13	3.74, 72	4.05, 1	3.82, 45	3.58, 26
UKY	-0.09, 0.49	0.13 (2, 66), 0.71	3.14	3.54, 69	3.17, 1	3.52, 41	3.6, 27
NCSU	0.39, 2e-3*	4.3 (2, 58), 1.7e-2 [†]	3.15	3.72, 61	4.43, 3	3.75, 39	3.53, 19

VCU	0.1, 0.53	3.5 (1, 37), 0.06	3.26	3.82, 39	-	4.04, 9	3.75, 30
UVA	0.14, 0.52	0.62 (1, 22), 0.43	3.44	3.73, 24	4.62, 1	3.7, 23	-
ASU	0.74, 3e-3*	0.07 (1, 11), 0.78	3.98	3.91, 13	-	4.14, 7	3.64, 6

* stat. significant, $\alpha = 0.05$
[†] stat. significant, $F > F_{crit}$, $\alpha = 0.05$

Table 3: Regression analysis: overall professor rating, function of student GPA (X1), and perceived ease of course instruments (X2 - X5)

	coef	std. error	t	p
intercept	2.6	0.318	8.3	0.00*
X1: GPA	0.5	0.085	6.0	0.00*
X2: exams	-0.03	0.041	-0.8	0.398
X3: quizzes	-0.08	0.037	-2.3	0.02*
X4: projects	-0.04	0.027	-1.7	0.08
X5: homework	-6e-3	0.032	-0.19	0.84

Table 4: Regression analysis (cont.)

stat.	val.
R-squared	0.193
Adj R-squared	0.182
F-statistic	17.8
Prob (F-statistic)	8e-16*
Log-likelihood	-244.4
AIC	500.8
BIC	524.4

mine review data to determine whether the following factors affect aggregate student perception of course instructors.

- university mandate (teaching vs. research, public vs. private),
- course modality (STEM/non-STEM, undergraduate/graduate, in-class/online)
- technology use (LMS and third-party apps for course management, testing and assessment)

- logistics (instructional design training, number of TAs, etc.)
- forum features - content creation (editorial control, authentication using university credentials)
- forum features - interaction (accessibility, review search, social tagging)
- forum features - content management and bias-correction (spam filtering, detection of cyberbullying and defamation)

An important step in realizing this contextual inquiry is designing metrics that summarize the observed disparity between effects of course rubric, content quality, interaction fidelity of host forums as well as course outcomes on the overall instructor ranking.

5 CONCLUSION

We present a preliminary quantitative analysis of sources of potential bias on academic forums. We find that for frequent users of one such academic forum, aggregate student ratings of course instructors gravitate, in many cases disproportionately, towards course outcomes and away from student perception of the relative ease of course materials, content and evaluations. We intend to generalize this analysis into a robust approach of isolating and correcting for bias on academic forums.

REFERENCES

- [1] 2019. Rate My Professors - review teachers and professors, school reviews, college campus ratings. (2019). Retrieved 2019 from <http://ratemyprofessors.com/>
- [2] Michael J Brown, Michelle Baillie, and Shawndel Fraser. 2009. Rating RateMyProfessors.com: A comparison of online and official student evaluations of teaching. *College Teaching* 57, 2 (2009), 89–92.
- [3] John A Centra. 2003. Will teachers receive higher student evaluations by giving higher grades and less course work? *Research in Higher Education* 44, 5 (2003), 495–518.
- [4] YunJeong Chang and Seung Won Park. 2014. Exploring Students' Perspectives of College STEM: An Analysis of Course Rating Websites. *International Journal of Teaching and Learning in Higher Education* 26, 1 (2014), 90–101.
- [5] Kenneth A Feldman. 1989. The association between student ratings of specific instructional dimensions and student achievement: Refining and extending the

- synthesis of data from multisection validity studies. *Research in Higher Education* 30, 6 (1989), 583–645.
- [6] Kenneth A Feldman. 2007. Identifying exemplary teachers and teaching: Evidence from student ratings. In *The scholarship of teaching and learning in higher education: An evidence-based perspective*. Springer, 93–143.
- [7] Dan Donahoe Glynn LoPresti, Patrick Gartlan and Minhe Le. 2019. Koofers - professor ratings, practice exams and flash cards. (2019). Retrieved 2019 from <http://koofers.com/>
- [8] Katherine B Hartman and James B Hunt. 2013. What RateMyProfessors.com reveals about how and why students evaluate their professors: A glimpse into the student mind-set. *Marketing Education Review* 23, 2 (2013), 151–162.
- [9] Taha Hassan and D Scott McCrickard. 2019. Trust and Trustworthiness in Social Recommender Systems. *Companion Proceedings of the 2019 World Wide Web Conference (WWW '19 Companion)* (May 2019).
- [10] Jeannette Kindred and Shaheed N Mohammed. 2005. "He will crush you like an academic ninja!": Exploring teacher ratings on ratemyprofessors.com. *Journal of Computer-Mediated Communication* 10, 3 (2005), JCMC10314.
- [11] Angela M Legg and Janie H Wilson. 2012. RateMyProfessors.com offers biased evaluations. *Assessment & Evaluation in Higher Education* 37, 1 (2012), 89–97.
- [12] Herbert W Marsh. 1984. Students' evaluations of university teaching: Dimensionality, reliability, validity, potential biases, and utility. *Journal of educational psychology* 76, 5 (1984), 707.
- [13] Marina Meila. 1999. An accelerated Chow and Liu algorithm: fitting tree distributions to high dimensional sparse data. (1999).
- [14] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. 1997. Kernel principal component analysis. In *International conference on artificial neural networks*. Springer, 583–588.
- [15] Stephen A Stumpf and Richard D Freedman. 1979. Expected grade covariation with student ratings of instruction: Individual versus class effects. *Journal of Educational Psychology* 71, 3 (1979), 293.