

Learning to Trust: Understanding Editorial Authority and Trust in Recommender Systems for Education

Taha Hassan

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

Timothy Stelter

Department of Computer Science, Virginia Tech
Blacksburg, Virginia
tstelter@vt.edu

Bob Edmison

Department of Computer Science, Virginia Tech
Blacksburg, Virginia
kedmison@vt.edu

D. Scott McCrickard

Department of Computer Science, Virginia Tech
Blacksburg, Virginia
mccricks@vt.edu

ABSTRACT

Trust in a recommendation system (RS) is often algorithmically incorporated using implicit or explicit feedback of user-perceived trustworthy social neighbors, and evaluated using user-reported trustworthiness of recommended items. However, real-life recommendation settings can feature group disparities in trust, power, and prerogatives. Our study examines a complementary view of trust which relies on the editorial power relationships and attitudes of all stakeholders in the RS application domain. We devise a simple, first-principles metric of *editorial authority*, i.e., user preferences for recommendation sourcing, veto power, and incorporating user feedback, such that one RS user group confers trust upon another by ceding or assigning editorial authority. In a mixed-methods study at Virginia Tech, we surveyed faculty, teaching assistants, and students about their preferences of editorial authority, and hypothesis-tested its relationship with trust in algorithms for a hypothetical ‘Suggested Readings’ RS. We discover that higher RS editorial authority assigned to students is linked to the relative trust the course staff allocates to RS algorithm and students. We also observe that course staff favors higher control for the RS algorithm in sourcing and updating the recommendations long-term. Using content analysis, we discuss frequent staff-recommended student editorial roles and highlight their frequent rationales, such as perceived expertise, scaling the learning environment, professional curriculum needs, and learner disengagement. We argue that our analyses highlight critical user preferences to help detect editorial power asymmetry and identify RS use-cases for supporting teaching and research.

CCS CONCEPTS

• **Human-centered computing** → HCI theory, concepts and models; *User models*; *Social recommendation*; • **Information systems** → Personalization.

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.

UMAP '21, June 21–25, 2021, Utrecht, Netherlands

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8366-0/21/06...\$15.00

<https://doi.org/10.1145/3450613.3456811>

KEYWORDS

recommendation, education, context, trust, interpretation

ACM Reference Format:

Taha Hassan, Bob Edmison, Timothy Stelter, and D. Scott McCrickard. 2021. Learning to Trust: Understanding Editorial Authority and Trust in Recommender Systems for Education. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '21)*, June 21–25, 2021, Utrecht, Netherlands. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3450613.3456811>

1 INTRODUCTION

Building user trust in recommendation systems (RS) is inextricably linked to concerns of algorithmic accountability and trust in automation. User experiments with RSs typically assess the *experiential* notion of trust using the *perceived* trustworthiness of the RS output. Algorithms for recommendation also frequently incorporate explicit and implicit signals of trust, for instance, the RS user’s *interactional* awareness of their local neighborhood, or some consensus of the preferences of their self-reported, trustworthy friends in a social network at large [21][13][12]. While algorithmic awareness of a user’s neighborhood is important for producing accurate recommendations, real life recommendation tasks often involve user groups with differences in institutional, group, or task-based roles, powers and prerogatives. Building an educational RS, for example, needs us to model the interaction (or lack thereof) between teachers, teaching assistants and students in deriving value out of said RS. We argue that recognizing the editorial power relationships between stakeholders in the RS application domain is one way to begin to identify the broader context of trust in the recommendation algorithm, and to expand the interpretive power of the RS output. As a case study, we examine Virginia Tech faculty and students’ preferences of editorial authority and trust in algorithms, for a hypothetical ‘Suggested Readings’ recommender system aboard a learning management system (LMS) course site. Using a simple, first-principles metric of editorial authority (**E-Auth**), we hypothesis-test the relationship between RS editorial task distribution and stakeholder trust in algorithmic agency. We then describe the top three editorial roles (author, active viewer, viewer) allocated to students by course staff, and identify frequent themes, contexts and RS use cases using the trust relationships that define these editorial roles. Figure 1 provides an overview of our

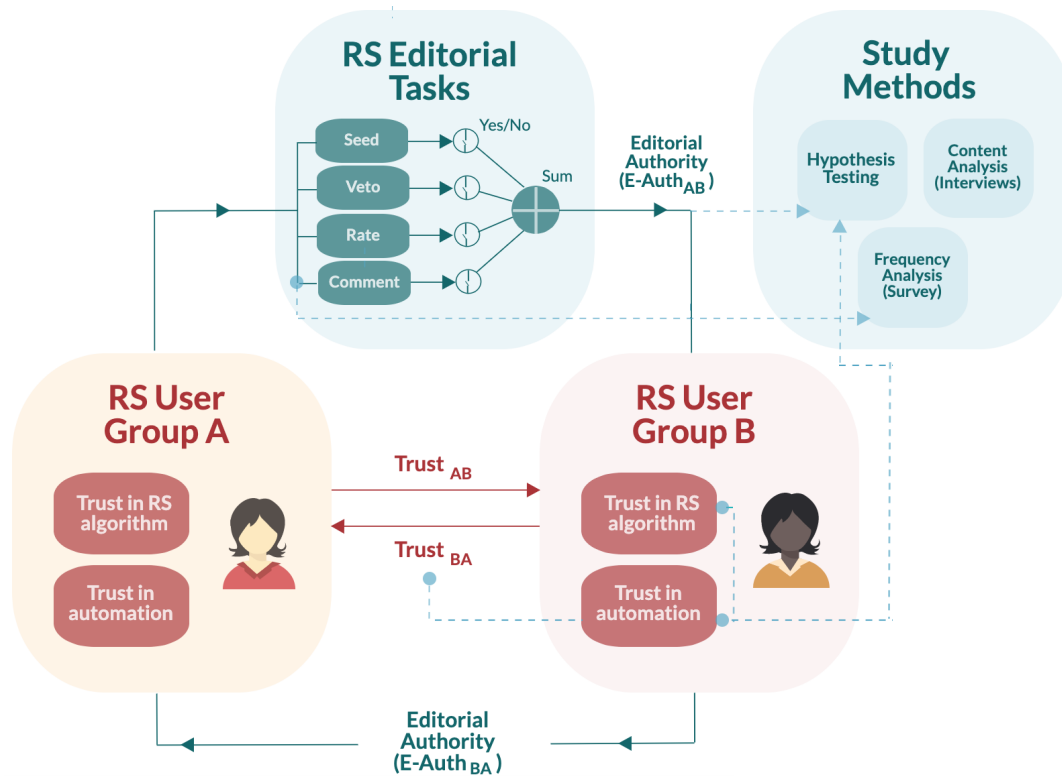


Figure 1: Study overview: recommender system (RS) user groups, trust relationships, attitudes about algorithms and automation, and mutual editorial role assignment. For our study, user groups include faculty, teaching assistants, and students.

study methodology, and the group attitudes and trust relationships we investigate.

2 RELATED WORK

There is ample prior work on both interpretive and algorithmic models of trust, in fields as diverse as recommender systems [28], adaptive web systems [3], user modeling [10][25], computer networks [15], and game theory [4]. In a cross-domain review of trust frameworks, Artz and Gil [3] describe four broad paradigms of trust analyses (policy, reputation, information web, and general). The first three deal with algorithmic, heuristic, or policy-based views of trust inference for systems with human and software agents. For instance, Massa and Avesani [20][21] develop an influential trust-aware recommendation framework which devises a *local* trust propagation strategy: predicting the trustworthiness of the RS user's neighbors based on proximity to user-labeled trusted 'peers'. A large body of work has been devoted to trust-aware collaborative filtering (CF), trust propagation, aggregation and user-similarity assessment using contextual, personal and network traits ever since [27]. Ma et al. [18], for instance, propose an ensemble fusion of the ratings and web-of-trust similarity matrices. User experiments with recommender systems [17] often identify trust as one of the personal attributes of users or user groups. For instance, Knijnenburg et al. [16] compare five interaction methods (topN, hybrid, explicit, implicit and hybrid) with a recommender system for energy saving

measures. They discover that a user's propensity to trust is linked more to their choice satisfaction and trust in the overall recommendation system, rather than their preference for additional control or customization.

The final paradigm, in comparison, consists of sociological, psychological, and game-theoretic precursors and implications of trust. Mui et al. [23] propose a computational model using the notion of 'reciprocity' (the bi-directional exchange of favors or revenge) to infer social reputation and trust. Braynov and Sandholm [4] discover that a misrepresentation of agents' trustworthiness or distrust can result in sub-optimal degrees of social welfare, profits and cooperation in a bilateral negotiation game. Note how aforementioned models often capture trust exchange between stakeholders, enhancing or undermining their *reputation*, and by implication, trustworthiness of their social neighbors in the process.

A parallel body of work on recommendation for groups [8][2] has grappled with the problem of evaluating group consensus when its various subgroups exhibit independent and overlapping preferences and roles. A subset of this research relies on social choice theory to find broad types of consensus (most popular items, least misery for group members, etc.) [5]. Seko et al. [26] combine the notions of overlapping behavioral tendencies and power-balance in the group to assign consensus. Many open questions remain, including how to evaluate power asymmetry and assess the trustworthiness of algorithmic judgments of consensus.

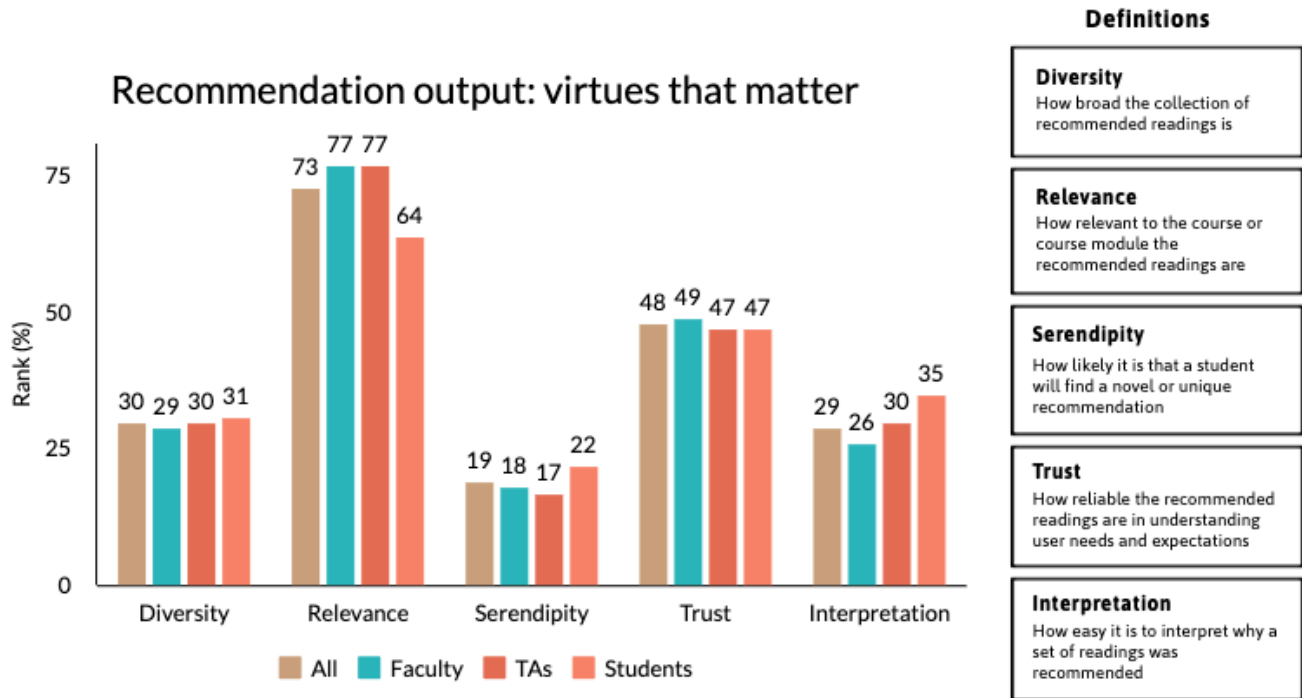


Figure 2: Virginia Tech faculty and students' perception of the relative importance of RS output virtues (with definitions). All three user roles (faculty, teaching assistants, students) rank relevance and trust consistently higher than interpretation, diversity and serendipity.

Our work seeks to advance a complementary and somewhat under-explored, actor network theory-inspired [7] approach to modeling the rich contexts in which educational RSs operate. We contend that editorial power assignment is a key early mechanism by which educational recommender systems are adopted by and appropriated for faculty, students and TAs to support teaching and research. Our approach assesses the distribution of trust preferences by RS editorial authority vested in these broad user roles. We then examine the distinctions between different stakeholders' power to source, veto, rate, and comment on recommended readings, and how their attitudes towards algorithmic agency and trust relationships with other stakeholders inform and affect these differences. We make the following contributions in this study.

- (1) We propose a first-principles view of group trust in recommendation, based in editorial task assignment,
- (2) We hypothesis-test the relationship between editorial authority (**E-Auth**) and trust in algorithmic agency,
- (3) We infer key editorial roles allocated to students for educational recommendation, as well as assess staff needs and precedents that inform the corresponding trust relationships,

3 EDITORIAL AUTHORITY, TRUST AND ALGORITHMIC AGENCY

Recommender systems for education have fulfilled a variety of tasks for the individual learner, interpretation and intervention, in-class and online [19]. However, there is a pressing need for the educational RS community to recognize platform-level changes in the domain. Recommender systems have to reckon with concerns of efficacy, trust and interpretation for *multiple stakeholders*, and *at scale* [14], especially as service-based learning management systems (LMS) become the primary infrastructure for productivity, communication and class management at institutions of higher learning [22]. Figure 2 describes five desirable aspects of RS output (relevance, trust, diversity, serendipity, and interpretation), and how faculty members and students who participated in our study felt about their relative importance. Relevance and trust were consistently ranked more important relative to the rest by all user roles.

Our study investigates whether the differences in *editorial authority* - exemplified in, say, an editor-consumer relationship between faculty and students - are related to the trust both assign to each other and the RS algorithm. It explores the most frequent RS editorial models preferred by the study participants, and their relationship with the trust exchanged between faculty, TAs, students and the RS algorithm. For instance, faculty's willingness to

Table 1: Survey questions: editorial authority and trust for a ‘Suggested Readings’ recommender system

Question Category	Question Statement
Editorial authority	Who can select seed articles for the RS? (Instructor, TA, student) Who should have the ‘veto power’? (Instructor, TA, student) Who should have the ability to rate the recommendations to influence their rank? Who should be able to submit feedback on a recommendation?
Relative trust in algorithms	In the short/long-term, who’s responsible for trustworthy recommendations? (Rank order: course staff, recommendation algorithm, students)
Automation	In the short/long-term, who’s responsible for trustworthy recommendations? (Likert: all human - all algorithmic)
Output attributes	What matters to you viz-a-viz RS output? (Rank order: diversity, relevance, serendipity, trust, interpretation)

Table 2: Key attributes of survey participants

Role	#	# Male/Female	# Departments	# STEM
Faculty	27	13/14	16	17
Teaching assistant	6	4/2	5	6
Student	9	7/2	4	7

incorporate student and TA feedback into the RS algorithm can point to a belief in editorial authority for multiple stakeholders, or regard for automation in the longer-term. Table 1 describes the tasks accounted for in editorial authority.

3.1 Definitions

We define the ‘Editorial Authority’ (**E-Auth**) allocated to a study participant as a linear aggregate of all editorial powers (seed, veto, rate, comment) they identify for their own user group relative to all the other user groups. The individual editorial powers are represented by binary (‘yes’ or ‘no’) votes. Consider equation 1 as follows.

$$\mathbf{E-Auth (Faculty \rightarrow Students)} = \mathbf{Seed Score} + \mathbf{Veto Score} + \mathbf{Rate Score} + \mathbf{Comment Score} \quad (1)$$

For example, a faculty member’s choice to let the TAs and students seed RS articles will result in a **seed score** of 2 (as in, the power to source recommended articles is shared with two user groups). If this faculty member favors exclusive veto power for faculty, but retains the rate/comment power for TAs and students, the **veto score**, **rate score** and **comment score** would be 0, 2 and 2 respectively, with a final **E-Auth** score of 6 (normalized to a percentage for our analyses). A small **E-Auth** score thus indicates that the course instructor might minimize student participation to reduce the cognitive effort of managing student feedback at scale, to prevent spam and inappropriate content, and to better align the course with learning outcomes in a given degree specialization. Section 5 uses frequent themes from semi-structured interviews to further discuss faculty rationales behind differing choices of editorial authority.

3.2 Research questions and hypotheses

In this section, we describe our study research questions which test the connection between editorial authority (**E-Auth**), trust in RS algorithm, and automation.

3.2.1 RQI: Trust in RS algorithm. How does the relative trust in a ‘Suggested Readings’ recommendation algorithm vary by job role (**RQ1a**), and how does it vary over time (**RQ1b**)?

3.2.2 RQII: Automation. How much does the relative preference for algorithmic agency and control (over recommendation sourcing, updates and removal) vary by job role (**RQ2a**), and how does it vary over time (**RQ2b**)?

3.2.3 RQIII: Editorial authority, trust and automation. What is the relationship between editorial authority assigned to students and the corresponding trust in RS algorithm (**RQ3a**) and preference for automation (**RQ3b**)?

4 EVALUATION

We conducted a small-scale, mixed-methods study to discover and interpret the Virginia Tech staff and students’ perceived efficacy and trustworthiness of a hypothetical ‘Suggested Readings’ RS. In this section, we describe the design, evaluation methods, key results, and limitations of the study.

4.1 Dataset and Methods

Our dataset consists of survey responses from 42 participants (27 faculty, 6 teaching assistants, 9 students), and transcripts and notes from follow-up semi-structured interviews with 11 (6 faculty, 3 teaching assistants, 2 students) of the aforementioned 43, all affiliated with Virginia Tech. Table 2 details their key attributes. The faculty members in our survey represent 16 departments, while the TAs and students represent 5 and 4 departments respectively. 30




Figure 3: Editorial authority: Virginia Tech faculty, TA and students' opinions on who should be allowed to seed (left) and veto (right) articles for a 'Suggested Readings' RS. 44% of surveyed faculty allocated the article sourcing task to students, but only 4% preferred students participate in article vetoing.




Figure 4: Editorial authority: Virginia Tech faculty, TA, and students' opinions on whether students should participate in rating and commenting on the output of a 'Suggested Readings' RS. For instance, 25% of surveyed faculty do not favor that students' ratings influence recommended articles, and 13% do not favor student comments.

of the survey respondents (17 faculty members, 6 TAs, 7 students) represent STEM disciplines of study. We recruited participants on a rolling basis between August and October of 2020, using convenience sampling and voluntary response sampling on departmental mailing lists and Facebook groups. For our analyses, we use a mixed-methods approach. We perform hypothesis tests regarding the relationships between our study variables (relative trust in algorithms, automation, editorial authority) using one-way ANOVA (F-test). We also perform content analysis on survey responses and interview transcripts to identify frequent themes in the study participants' commentary regarding their preferred RS editorial tasks for each user role.

4.2 Results

Figure 3 illustrates how faculty, TAs and students prefer the distribution of RS tasks by job role. Nearly all faculty members allocate sourcing and vetoing tasks to themselves, but fewer (52%) think TAs should seed RS articles and even fewer (44%) allocate this task to students. The divide is quite apparent for vetoing, as all but one (4%) faculty member suggest that students should not be able to instantly remove recommended readings. As per figure 4, about 75% of faculty (as opposed to all students) favor the idea of soft power for students: the ability to rate individual recommendations and let the algorithm update the relative location of a recommended reading using group consensus strategies. 13% of faculty members and

Table 3: Key results from hypothesis tests about the relationship between E-Auth scores (RS editorial authority assigned by faculty to students) and staff/student attitudes. ‘Course staff’ refers to faculty and teaching assistants considered together. Statistically significant effects ($F > F_{crit}, p < 0.05$) are described in bold text in the rightmost column. For instance, per H1, teachers trust the RS algorithm less than teaching assistants, but according to H2, course staff does not trust the RS algorithm any less than students do.

Question Category	Research Question	Hypothesis	Effect sizes (F, p)
Trust in RS algorithm	RQ1a	H1: Teachers trust the RS algorithm less than teaching assistants.	6.49, 0.01
	RQ1a	H2: Course staff trusts the RS algorithm less than students.	0.19, 0.66
	RQ1b	H3: All participants trust the staff more short-term compared to long-term.	10.8, <0.01
	RQ1b	H4: All participants trust the RS algorithm more long-term.	6.3, 0.01
Automation	RQ2a	H5: Course staff trusts automation less in the short-term than students.	9.7, <0.01
	RQ2b	H6: All respondents trust automation more long-term compared to short-term.	3.9, 0.06
	RQ2b	H7: Course staff trusts automation more long-term compared to short-term.	7.2, <0.01
Editorial authority	RQ3a	H8: Higher trust in students long-term is linked to higher student E-Auth .	3.7, 0.03
	RQ3a	H9: Lower average trust in RS algorithm is linked to higher student E-Auth .	6.5, 0.01
	RQ3b	H10: Favoring high RS automation is linked to higher student E-Auth .	0.52, 0.67

17% of teaching assistants suggest that they do not favor students’ ability to submit feedback about a recommended reading.

Table 3 describes the results of hypothesis tests on key claims about trust in RS algorithm, preference for automation, and aggregate editorial authority (see research questions **RQ1-3** in section 3.2). We discover that overall, course staff (faculty and teaching assistants) tends to be more risk-averse in its trust relationship with RS algorithm and automation than students, favoring substantial human intervention in the sourcing, updating and removal of recommended readings in the short-term. Here is a detailed breakdown of these results.

4.2.1 Trust in RS algorithm (RQ1, H1-H4). We discover that teachers tend to trust the RS algorithm less than teaching assistants ($F = 6.49, p = 0.01$). All user groups favor the role of course staff in ensuring trustworthy recommendations short-term more than long-term ($F = 10.8, p = 1e^{-4}$). Similarly, all user groups trust the RS algorithm more long-term compared to short-term ($F = 6.3, p = 0.01$).

4.2.2 Automation (RQ2, H5-H7). Course staff tends to favor automation in recommendations less relative to students, especially in the short-term ($F = 9.7, p = 3e^{-3}$). Similarly, course staff favors automation in the long-term more than in the short-term ($F = 7.2, p = 9e^{-3}$).

4.2.3 Editorial authority (RQ3, H8-H10). Higher trust in students is linked to preference for higher student **E-Auth** scores in the long term ($F = 3.7, p = 0.03$), hinting at a connection between perceived trust and editorial role assignment. Similarly, lower average trust in RS algorithm is linked to preference for higher student **E-Auth** scores ($F = 6.5, p = 0.01$). Preference for higher RS automation overall is however, not linked to editorial authority ($F = 0.52, p = 0.67$).

5 DISCUSSION AND LESSONS LEARNED

In this section, we describe the top three frequently recommended editorial roles (author, active viewer, viewer) for students, as allocated by our study respondents (table 4). Allocation refers to the fraction (%) of participants who chose to allocate a given editorial role (and corresponding powers and tasks) to students. For each editorial role, we discuss the frequent themes (in bold text) of the mediating trust relationships we investigate in **RQs I** through **III**. These emerge from our content analysis, and they include, and are not limited to, differences in perceived expertise, supporting learner engagement, and scaling the learning environment. These editorial roles highlight distinct - if overlapping - user preferences that can suggest RS use-cases for supporting both teaching and research.

5.1 "The author" (A): Everything except veto

According to table 4, 40% of participants overall, and 36% of course staff favored the ‘author’ (A) role for students. A-faculty members prefer sharing article sourcing, rating and feedback responsibilities with students, but retain the veto power (i.e. the ability to instantly remove a recommended reading for all RS users) for the primary course instructor, as well as for the teaching assistants in a subset of cases. Table 5 lists the frequent themes and course contexts for all RS editorial roles. In addition to citing their prerogative being in charge of facilitating student learning, A-faculty members frequently cited the **need to moderate content** and remove readings deemed irrelevant or malicious. According to one instructor: "I worry about abusive behavior, unless it was clearly tracked" and this would "ensure the integrity of the course and the suggested readings". They simultaneously acknowledged the potential utility of **student engagement** in these editorial tasks, and inquired about the ability of the RS algorithm to comprehend student needs early on (**H1, H4**), especially in graduate, research-focused courses. One faculty member said that based on her teaching experience, she perceived MS-level students to be resourceful and another commented:

Table 4: RS editorial roles allocated to students. Editorial tasks essential to each role are in bold text. Allocation refers to the number and fraction (%) of participants who chose to allocate a given editorial role to students. Average E-Auth refers to the amount of editorial authority assigned to students in each role. For instance, 39% of course staff favored students provide input as active viewers (AV), assigning them 54% of the available editorial authority.

Editorial Role	Editorial Powers/Tasks	Average E-Auth	Overall Allocation	Staff Allocation
Editor (E)	Seed, veto , rate, comment	80%	1 (2%)	1 (3%)
Author (A)	Seed , rate, comment	77%	17 (40%)	12 (36%)
Active Viewer (AV)	Rate , comment	54%	17 (40%)	13 (39%)
Viewer (V)	Comment	14%	7 (16%)	7 (21%)

".. if objectionable material was posted, I'd rather we use it as a discussion topic rather than outright reject it."

Soft power, as in the ability to rate (or 'like' and 'dislike') content as a signal to RS algorithm, was frequently favored for students as a means of further engaging them with the course content. An instructor said of soft power, "I like the way the soft power idea will engage their minds and (aid their) learning". This is consistent with hypothesis **H8** and **H9** in table 3. Teaching assistants often brought up the case of large course sections with multiple instructors, and how TAs and students having article sourcing powers could potentially help scale the teaching resources faster, and make them accessible to a larger subset of the student population. According to one teaching assistant,

".. course instructors may sometimes not know when to remove content unless they are intimately familiar with it and know it to not be useful."

Student participants in our study were, by-and-large, in agreement with course staff about exclusive veto power for faculty and teaching assistants, especially because they felt that the ability to seed, rate and comment on articles gave them ample opportunity to engage with the course materials and RS authorship policies. As per one student: "I feel anything that is against the teaching of the class should be allowed and a respectful discussion should occur; preferably in a forum setting."

5.2 "The active viewer" (AV): Rate and comment

Table 4 suggests that 40% of study participants overall, and 39% of course staff favor an 'active viewer' role for students. This allocates the RS rating and commenting tasks to students while reserving the sourcing and vetoing of suggested readings for course staff. Members of this group frequently talk about the challenges of managing and responding to feedback at **scale**, especially in undergraduate courses with multiple sections and hundreds of students. One faculty member remarked about students not being able to seed or veto suggested readings:

".. (this is) just to be able to manage with a course that has 14-15 sections of 65 students each."

This comment is not a lone anomaly. Over the last two years, the average undergraduate course at Virginia Tech has 3.7X the class size of the average graduate course (66 and 17 students on average, respectively), with many first-year, major-unrestricted courses enrolling several hundred students in any given term. It is worthwhile

noting that **AV**-faculty members do not seek a fully cooperative model of student feedback to solve this challenge in the manner of **A**-faculty. Frequent reasons posited by faculty include perceived historical patterns of **student disengagement** and exclusive attention by students to course outcomes, or the **outcome bias** [12]. This bias is also known to affect student perceptions of instructors on academic social forums like *Koofers* [9] and *RateMyProfessor* [1]. As one faculty member commented,

"My experience has been that most students do not do more than what's required of them, unfortunately. They just want to know what they can do to earn an A in class."

Another faculty member talked about the requirements of professional diplomas - dictated by the rapidly evolving demands of the job market - as one key driver of student disengagement. He noted that the "fast pace and skill-focused curriculum" of professional diplomas (as opposed to research-based degrees) made it difficult for students to spend time on optional course content. In his experience, this had led to a substantial decline in student interaction with the course LMS site, to the point where cooperative editing of recommended readings seemed ineffective. In comparison, teaching assistants and students who favored the **AV** role largely cited course staff's **need for moderating** the RS articles for malicious behavior and honor code violations. About the article rating task, one student commented that it will "allow the instructor and TAs to see what's most widely accepted". Consistent with hypothesis **H7** (table 3), all participants expected the role of automation to be significantly higher long-term in consolidating student updates to the recommendations' rank order.

5.3 "The viewer" (V): Comment only

Study participants in this group favored the least egalitarian editorial model for recommending readings to students. 16% of participants overall, and 21% of faculty and teaching assistants prefer no direct input from students in deciding the source and rank order of recommended readings. Same as **AV**, this is frequently observed for faculty members with undergraduate teaching responsibilities. This model considers student participation in article sourcing, updates and removal to be unsustainable if not counter-productive. Faculty survey respondents and interviewees frequently cited challenges of **content moderation at scale**. One instructor of an undergraduate Computer Science class described the problem of cooperative RS editing as analogous to the challenge of regulating discussion forums posts for the 200+ enrolled students in her section. Drawing

Table 5: Frequent course contexts, themes, and use-cases for RS editorial roles allocated to students. For instance, study participants favored cooperative RS editorial models (E/A) for small class sizes (typically graduate courses). Courses with a preference for lesser or no student input (AV/V) have large class sizes (often undergraduate) and time constraints defined by course content (STEM, professional diplomas).

Editorial Role	Frequent Course Context	Frequent Themes	RS Use Cases
E/A	Graduate, Non-STEM	student engagement, content moderation	teaching
AV	Undergraduate, STEM	student disengagement, scale, outcome bias, content moderation	teaching, professional development
V	Undergraduate, STEM	student disengagement, scale, content moderation	research, teaching

on instances of age-sensitive commentary by students about an automated grading software, this instructor suggested that comments on recommended readings be invisible to fellow students by default. This would allow course staff to remove malicious content and notify the students. Another course instructor in this group complained "... unfortunately, students are often testing the limits of the honor code."

An Engineering instructor cited his time in the industry as having informed his singular emphasis on problem-solving in teaching graduate courses. He mentioned he favored creating and updating course assignments, projects, and exams without serious consultation from a primary textbook. This, according to him, discouraged students' use of online solution manuals, better assessed their progress with course milestones, and ensured they learned a precise set of Engineering skills without **undue cognitive burden**. Note the parallel with AV-faculty's rationale for the limited role of a 'Suggested Readings' RS in teaching courses for professional diplomas. About restricting students to a comment-only RS editorial role, he commented:

"It's not so much that I want the control. It's more that I don't think that should be there focus from an educational standpoint. I'm trying to get them to comprehend some pretty intense Engineering and design concepts. I consider it a burden for them to go out and find other resources. I really have to get them to practice a lot of things over and over."

As per table 5, several V-faculty suggested they might use the recommender system in **supporting research**, as opposed to as-sorting readings for teaching purposes. The frequently cited use-cases for such a recommender system were discovering research articles beyond Virginia Tech library-indexed databases, and discovering topic overlap with other research fields to inform literature reviews. It is worth noting that no teaching assistants or students favored a V-role for students, hinting at a strongly asymmetric editorial relationship between faculty and students. A more rigorous evaluation of this power asymmetry is left for future work.

5.4 Editor-consumer relationships

We observe frequent power asymmetry when it comes to editorial preferences for recommendation in the education domain. 25% and 13% of faculty members surveyed did not favor students' ability to rate and comment on the recommended readings, respectively

(figure 4). As we observe earlier in this section, this frequently coincides with complaints of inevitable learner disengagement with course materials, but the course staff's comfort with automating and delegating editorial tasks increases with time. Note the connection with passive use of recommender systems observed, for instance, in music and video recommenders where the recommendation algorithm has disproportionate editorial power (both explicit and implicit) as a function of time relative to passive consumers. While outside the scope of this study, it is important to investigate the UX implications of such a disparity using trust frameworks that take into account the capacity of the recommendation system to incorporate context [11], increase meaningful use through diverse use-cases anchored in trust, provide a diverse set of interaction techniques [17], and disrupt filter bubbles [24].

6 CONCLUSIONS AND FUTURE WORK

We examine the distribution of editorial preferences and its broader roots in notions of trust, expertise and institutional and domain-specific prerogatives. We discover that power asymmetry is an important component of the user experience of educational recommender systems, and should be leveraged in algorithmic and interactional affordances for trust-aware recommender systems at scale.

The scope of our pilot study can be extended in several important ways. We review these as directions for future work as follows. The sparsity of interview evidence at this stage of the project makes it difficult to generalize our conclusions across the Virginia Tech faculty and student population at large. We reached out to course staff with a history of collaboration with departmental initiatives towards instructional design. For our future work, we would like to collect a comprehensive sample of faculty members, especially the ones who need to prioritize research over teaching, favor legacy tools, or experience a high cognitive burden-of-discovery [29] when adopting new digital tools. Given the variability of trust preferences is itself different for faculty, students and teaching assistants, we also plan to incorporate more robust multi-item trust and automation scales. We also intend to pursue a concrete definition of present bias in RS users, in order to thoroughly interpret the time variance of editorial preferences we observed in our study. Concrete experimental conditions with RS algorithm strata (learning strategies, interaction methods) and domain-specific timescales (semester, course project timeline, course add/drop/withdraw periods)

would further improve the consistency and actionability of our conclusions. Finally, we plan to measure the actual use patterns of our personalized ‘Suggested Readings’ recommendation system as a function of our inferred editorial roles. It is part of a larger effort by our team to understand the adoption of an institution-wide learning management system (LMS) and its managed apps. LMS adoption is a complex function of department-level precedents, faculty’s cognitive burden-of-discovery, task-technology fit, as well as the needs for scale and ubiquitous access [6]. In our future work, we hope to expand the interpretive power of LMS platform analytics using our multistakeholder trust models.

REFERENCES

- [1] 2019. *Rate My Professors - review teachers and professors, school reviews, college campus ratings*. Retrieved 2019 from <http://ratemyprofessors.com/>
- [2] Sihem Amer-Yahia, Senjuti Basu Roy, Ashish Chawlat, Gautam Das, and Cong Yu. 2009. Group recommendation: Semantics and efficiency. *Proceedings of the VLDB Endowment* 2, 1 (2009), 754–765.
- [3] Donovan Artz and Yolanda Gil. 2007. A survey of trust in computer science and the semantic web. *Journal of Web Semantics* 5, 2 (2007), 58–71.
- [4] Sviatoslav Braynov and Tuomas Sandholm. 2002. Contracting with uncertain level of trust. *Computational Intelligence* 18, 4 (2002), 501–514.
- [5] Iván Cantador and Pablo Castells. 2012. Group recommender systems: new perspectives in the social web. In *Recommender systems for the social web*. Springer, 139–157.
- [6] Hamish Coates, Richard James, and Gabrielle Baldwin. 2005. A critical examination of the effects of learning management systems on university teaching and learning. *Tertiary education and management* 11 (2005), 19–36.
- [7] Tara Fenwick and Richard Edwards. 2010. *Actor-network theory in education*. Routledge.
- [8] Mike Gartrell, Xinyu Xing, Qin Lv, Aaron Beach, Richard Han, Shivakant Mishra, and Karim Seada. 2010. Enhancing group recommendation by incorporating social relationship interactions. In *Proceedings of the 16th ACM international conference on Supporting group work*. 97–106.
- [9] Dan Donahoe Glynn LoPresti, Patrick Gartlan and Minhe Le. 2019. *Koofers - professor ratings, practice exams and flash cards*. Retrieved 2019 from <http://koofers.com/>
- [10] Junpeng Guo, Yanlin Zhu, Aiai Li, Qipeng Wang, and Weiguo Han. 2016. A social influence approach for group user modeling in group recommendation systems. *IEEE Intelligent systems* 31, 5 (2016), 40–48.
- [11] Taha Hasan, Naveed Arshad, Erik Dahlquist, and Scott McCrickard. 2017. Collaborative filtering for household load prediction given contextual information. In *Proceedings of the 2017 SIAM Workshop on Machine Learning for Recommender Systems (MLRec)*.
- [12] Taha Hassan. 2019. On bias in social reviews of university courses. In *Companion Publication of the 10th ACM Conference on Web Science*. 11–14.
- [13] Taha Hassan. 2019. Trust and trustworthiness in social recommender systems. In *Companion Proceedings of The 2019 World Wide Web Conference*. 529–532.
- [14] Taha Hassan, Bob Edmison, Larry Cox, Matt Louvet, Daron Williams, and D Scott McCrickard. 2020. Depth of Use: An Empirical Framework to Help Faculty Gauge the Relative Impact of Learning Management System Tools. In *Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education*. 47–53.
- [15] Sepandar D Kamvar, Mario T Schlosser, and Hector Garcia-Molina. 2003. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the 12th international conference on World Wide Web*. 640–651.
- [16] Bart P Knijnenburg, Niels JM Reijmer, and Martijn C Willemsen. 2011. Each to his own: how different users call for different interaction methods in recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*. 141–148.
- [17] Bart P Knijnenburg and Martijn C Willemsen. 2015. Evaluating recommender systems with user experiments. In *Recommender Systems Handbook*. Springer, 309–352.
- [18] Hao Ma, Irwin King, and Michael R Lyu. 2009. Learning to recommend with social trust ensemble. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*. 203–210.
- [19] Nikos Manouselis, Hendrik Drachler, Riina Vuorikari, Hans Hummel, and Rob Koper. 2011. Recommender systems in technology enhanced learning. In *Recommender systems handbook*. Springer, 387–415.
- [20] Paolo Massa and Paolo Avesani. 2004. Trust-aware collaborative filtering for recommender systems. In *OTM Confederated International Conferences“ On the Move to Meaningful Internet Systems”*. Springer, 492–508.
- [21] Paolo Massa and Paolo Avesani. 2007. Trust-aware recommender systems. In *Proceedings of the 2007 ACM conference on Recommender systems*. 17–24.
- [22] Tanya J McGill and Jane E Klobas. 2009. A task–technology fit view of learning management system impact. *Computers & Education* 52, 2 (2009), 496–508.
- [23] Lik Mui, Mojdeh Mohtashemi, and Ari Halberstadt. 2002. A computational model of trust and reputation. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*. IEEE, 2431–2439.
- [24] Sayooran Nagulendra and Julita Vassileva. 2014. Understanding and controlling the filter bubble through interactive visualization: a user study. In *Proceedings of the 25th ACM conference on Hypertext and social media*. 107–115.
- [25] Pearl Pu, Li Chen, and Rong Hu. 2012. Evaluating recommender systems from the user’s perspective: survey of the state of the art. *User Modeling and User-Adapted Interaction* 22, 4–5 (2012), 317–355.
- [26] Shunichi Seko, Takashi Yagi, Manabu Motegi, and Shinyo Muto. 2011. Group recommendation using feature space representing behavioral tendency and power balance among members. In *Proceedings of the fifth ACM conference on Recommender systems*. 101–108.
- [27] Yue Shi, Martha Larson, and Alan Hanjalic. 2014. Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges. *ACM Computing Surveys (CSUR)* 47, 1 (2014), 1–45.
- [28] Patricia Victor, Martine De Cock, and Chris Cornelis. 2011. Trust and recommendations. In *Recommender systems handbook*. Springer, 645–675.
- [29] Richard E West, Greg Waddoups, and Charles R Graham. 2007. Understanding the experiences of instructors as they adopt a course management system. *Educational Technology Research and Development* 55, 1 (2007), 1–26.