

Finding Place in a Design Space: Challenges for Supporting Community Design Efforts at Scale

STEPHEN MACNEIL, University of California, San Diego

ZIJIAN DING, University of California, San Diego

ASHLEY BOONE, University of Washington

ANTHONY GRUBBS, University of California, San Diego

STEVEN P. DOW, University of California, San Diego

Many organizations have adopted design processes that integrate community voices to discover the real problems that communities face. Online discussion forums offer a familiar and flexible technology that can help facilitate discussion around problems and potential solutions. However, we lack understanding about what information community members share, how that information is structured, and how social interactions affect design processes at scale. This paper presents a mixed-methods analysis of Canvas, a learning management system, which enables users to contribute to the design of the platform by sharing and deliberating on problems and solutions in a discussion forum. We collected and analyzed 1412 ideas and 18,335 associated comments shared on the Canvas discussion forum. We found that only 56.8% of posted ideas articulated a clear problem and 40.4% included a comprehensive set of information (problem, solution, and contextual details) sufficient to make design decisions. At least 11.6% of the ideas were duplicates, many of which were posted by members to garner attention for their ideas. The distributed nature of design information, the presence of duplicate ideas, and gaming behaviors made it difficult for the community to get oriented to the discussion. Finally, we reflect on how Canvas community members contribute information to a co-constructed design space and how future systems could more effectively coordinate community design efforts.

CCS Concepts: • **Human-centered computing** → **Computer supported cooperative work**.

Additional Key Words and Phrases: community-driven design, participatory design, open innovation, collaborative ideation, design spaces, problem framing, discussion forums

ACM Reference Format:

Stephen MacNeil, Zijian Ding, Ashley Boone, Anthony Grubbs, and Steven P. Dow. 2020. Finding Place in a Design Space: Challenges for Supporting Community Design Efforts at Scale. *J. ACM* 37, 4, Article 111 (August 2020), 30 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

As a discipline, design has been evolving from an exclusive field practiced by experts who tend to be disconnected from the communities they serve, to a more participatory practice. Participatory design [118] and co-design [114, 115] have set the stage for new design methods and tools that not only incorporate communities, but are entirely driven by communities themselves [62]. Such

Authors' addresses: Stephen MacNeil, University of California, San Diego, , San Diego, CA, 92122, smacneil@ucsd.edu; Zijian Ding, University of California, San Diego, , San Diego, CA, 92122, zding@eng.ucsd.edu; Ashley Boone, University of Washington, , Seattle, WA, 98195, aboone06@uw.edu; Anthony Grubbs, University of California, San Diego, , San Diego, CA, 92122, agrubbs@ucsd.edu; Steven P. Dow, University of California, San Diego, , San Diego, CA, 92122, spdown@ucsd.edu.

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 ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

0004-5411/2020/8-ART111 \$15.00

<https://doi.org/10.1145/1122445.1122456>

approaches leverage the unique assets [97, 99] and capacities [130] that community members have to offer, and have the potential to integrate a diverse collection of voices into the design pipeline. The goal is to democratize the design process, give designers access to the collective knowledge of the community, and discover the problems that matter most.

Advances in collaborative design tools have enabled these community-focused design processes to scale [22, 79, 120]. These specialized platforms provide substantial scaffolding to guide non-designers through the design process [22, 79] and tend to focus on specific aspects of design such as ideation [22, 46, 120]. Many organizations use online discussion forums to give community members a voice in ways that would be difficult to support through in-person meetings alone. Online forums have traditionally been used to support discussion and deliberation around news [11, 131], public policy [31], and education [59]. In the context of design, the familiarity of discussion forums can lower the barrier to participation, allowing communities to surface problems and ideate solutions [4, 6, 88, 108]. Discussion forums are also relatively flexible which may allow for open-ended design activity to unfold. However, without the scaffolding provided in open innovation platforms, it is unclear how community members get oriented, make contributions, and influence design decisions within an open discussion forum. While prior research sheds light on how communities discuss topics, our work explores how communities engage in design discussions.

Our research explores how collaborative technology—discussion forums, in particular—affects the ability of a large distributed community to participate in a design process. As an empirical setting for our research, we investigate *Canvas Studio*, a discussion forum dedicated to community discussions about the design of the Canvas learning management system (LMS). The Canvas LMS supports assignment uploads, quizzes, online discussions, and countless other features used by students, teaching assistants, and instructors. The Canvas Studio discussion forum gives these community members an opportunity to post, discuss, and vote on the ideas for how to improve Canvas. Canvas moderators facilitate these design discussions and select features to implement. Figure 1 shows how posts are presented in Canvas Studio.

To understand how design unfolds within the Canvas community, we critically examined user engagement, design contributions, and empowerment through the lens of participatory design (PD). While the Canvas forum and organizers do not explicitly talk about PD, the PD framework provides a useful analytic lens for exploring how a large community collaboratively engages in various design activities such as submitting, deliberating, and prioritizing problems and solutions. It also sheds light on issues of community empowerment, and how they relate to forum behavior.

Through a mixed-method study, we collected 1412 ideas and 18,335 associated comments shared on Canvas between Mar 18, 2015 to Aug 06, 2018 (three and a half years or 41 months). We analyzed this data to understand how community members orient themselves to different design issues, make design contributions, and influence what features get adopted. As part of our analysis, we qualitatively coded for design information contained within forum posts and found that only 56.8% of posts included a problem and 40.4% of posts included all the necessary information about the design situation (e.g.: a problem, solution, and contextual information). The fact that design information was spread across posts likely made it difficult for members to get oriented and see a full picture of the design space. An analysis of the tags showed many duplicate ideas (11.6%) that had been flagged by the community and moderators. Some members commented in their posts about deliberately re-posting their ideas to get attention and 5% of posters contributed 19.9% of ideas. Deliberately re-posting ideas, and otherwise gaming the system, may have further exacerbated the challenge of getting oriented by flooding the forum with posts from a few members while drowning out other community voices. Consequently, these behaviors create challenges for community members to feel empowered to make change and prioritize their needs.




Fig. 1. Canvas elicits design ideas using a discussion forum called Canvas Studio. The ideas contributed by individuals can be discussed and voted on by other community members to develop and triage new features.

This paper offers three contributions to the CSCW community: 1) a mixed-methods analysis of a community discussion forum dedicated to re-designing the Canvas LMS; 2) insights about hosting a large-scale design effort in an open discussion forum and the challenges members face getting oriented, making contributions to an evolving design space, and being empowered to make change; and 3) implications for how tools can better support community-driven design by collecting distributed design assets, representing the emergent design space co-created by the community, and providing training and scaffolding. As PD continues to integrate technology to support scale, these social and technological aspects need be considered. This lens highlights where the Canvas forum, and discussion forums more generally, meet and fall short of reaching the ideals of PD.

2 RELATED WORK

The discipline of design has seen a shift in power dynamics, from professionals saving the day, to supporting more community-driven [62], participatory approaches [114, 115, 118] where a range of stakeholders contribute to the process. The goal of participatory design (PD) is to involve stakeholders [114, 115, 118], often through guided activities that include creating mock-ups, prototyping, hosting workshops, and building scenarios [40]. These structured activities can help non-designers overcome a lack of design expertise. However, scaffolding enforces a specific structure on the design process that may not support the emergent and open-ended nature of design. Additionally, PD workshops are often short, facilitated, face-to-face sessions, which privileges those who have the time and means to attend [61].

Challenges with scaffolding and access have prompted researchers to question whether PD has fallen short of its ambition to democratize design [5, 18]. To broaden participation, calls have been made to adapt PD into a more open and inclusive process, often through technology-mediated online spaces [104, 112]. Open innovation platforms have emerged to support collaborative design by focusing on ideation. For example, OpenIDEO [79], Dell Idea Storm [6], and InnoCentive [80] present structured design challenges to collect ideas from online crowds. Although open innovation platforms engage contributors at scale, they often fail to adopt contributors' ideas [55, 63] and to achieve long-term involvement throughout the design process [4]. Bug reporting systems are

another example of how technology can enable large-scale participation. Unlike open innovation which often focuses solutions, bug reporting systems tend to focus on the elicitation of bugs—undesired functionality and “problems encountered while using a software” [7]. The users who report bugs, *bug reporters*, often remain anonymous and disconnected from the process of developing solutions. Bug reporters seldom engage in problem framing, an opportunity to scope and structure a problem [36, 37, 107, 117], let alone ideating on potential fixes and deciding how to prioritize what gets fixed. The resulting bug reports are often simple one-off descriptions of undesirable system behavior [7]. This limited involvement may account for the fact that bug reporters often fail to include information developers consider valuable for solving the problem [7, 23, 33, 55].

Many organizations leverage discussion forums—a familiar technology for online communication—to lower barriers to participation for those with limited digital literacy skills [95, 128]. Discussion forums have the potential to scale to support discussion by a much larger community than a typical in-person workshop. While discussion forums are popular, scalable, and flexible, more research is needed to understand whether they support an equitable and effective design process.

To ground our analysis of the Canvas discussion forum, we present related work about how discussants make sense of information in discussion forums, how their contributions to a discussion forum constitute legitimate design behaviors, and the challenges individuals can face when participating in community discussions about design and civic issues.

2.1 Getting oriented to ongoing online discussions

Online discussion forums have the potential to connect people across the world. They have been widely used to support online communities [48], to coordinate work [29], to share information [94], to deliberate policy decisions [27], to facilitate help seeking [25], and to facilitate learning online [1, 103]. Though discussion forums have the potential to connect people at scale, discussants often struggle with information overload as they try to orient themselves around the various conversations unfolding in parallel throughout a discussion forum [119]. To address information overload in forums, research has focused on developing technologies and affordances for forums to help people get oriented, including collaborative tagging [51], thread recommendation [82], recursive summarization [129, 133], and moderation carried out by humans [14, 42, 56] or machines [69]. These approaches help discussants by presenting only the most important information and filtering out irrelevant information. However, these orientation approaches also face challenges. Collaborative tagging systems often fail due to inconsistent tagging [25]. Summarization techniques typically only handle a single discussion thread [129, 133].

Anti-social behaviors—such as super-posting [66, 123], and signal boosting [20, 68]—further complicate the discussion landscape. Signal boosting (e.g.: super-posting and self-promotion) can flood the forum with content that only represents a minority of perspectives, effectively drowning out other community members’ voices.

In Canvas, community members not only need to navigate the social landscape and the substantial amount of information that has been shared, but must also understand how this information contributes to an evolving design process. Though previous research has explored challenges related to information overload in large discussions and anti-social behaviors, less is known about how design conversations unfold in large open community discussions. It is unclear which of the known challenges persist, which are exacerbated, and whether new challenges emerge when large groups of people engage in a shared design process. This prompts us to ask the research question:

RQ1 What challenges do community members face when orienting themselves to ongoing design conversations in a discussion forum?

2.2 Considering posts and comments as expressions of designerly behaviors

The act of designing can be thought of as the exploration, definition, and evaluation of co-evolving problem and solution spaces [38, 91]. Dorst and Cross describe design as a problem-solving activity where designers iteratively analyze, synthesize, and evaluate the corresponding problem and solution space [38]. Meanwhile, Bijker et al. describe how designers employ strategies to define aspects and options that constrain the emerging design space [10]. Finally, Schön describes how designers explore a design space through *design moves* that consist of trying, testing, and reflecting on a series of design experiments [117]. Through each of these different conceptualizations, design is about how actions serve to explore, define, and evaluate the space of possible designs.

Design processes structure and sequence design activities around an evolving design space [38, 50, 91]. For example, the Stanford d.school's Design Thinking Model¹ describes how designers move through phases of *Empathize*, *Define (the problem)*, *Ideate*, *Prototype*, and *Test*. Researchers have developed models of the design process by observing designers' behaviors [3, 30, 35, 52, 73, 100, 108, 125]. Atman et al. developed a formal coding scheme for observing design behaviors: *identifying needs, defining problems, gathering information, generating ideas, modeling, analyzing feasibility, evaluating, decision making, and communicating* [3]. These designerly behaviors are echoed in other coding schemes developed for observing problem framing in education [125], comparing novice and expert design processes [30], and analyzing online design discussions [73, 108].

Canvas Studio community members may exhibit similar designerly contributions in discussion posts by describing problems they face, gather information by asking questions and sharing their own experiences, or generate ideas for new features. Canvas provides guidelines for submitting ideas in the documentation for Canvas Studio. However, these guidelines (described in detail in Section 3.1) are not integrated into the workflow for submitting an idea and it is unclear whether every community member reviewed them before posting. This lack of knowledge about what design information is communicated and how it is structured raises the research question:

RQ2 What design information do community members contribute to a large and evolving design discussion?

2.3 Democratizing participation and decision-making in design

Participatory design (PD) has the potential to shift power from a small group of expert designers to an entire community of people who can draw from lived experiences to surface important problems. However, PD does not reach its full potential unless everyone can participate, be heard, and influence decision-making. Prior research has offered technology approaches that support this more democratic approach to design, including distributed decision making [43, 77], consensus building [85], and decentralized leadership [70, 86, 87]. For example, ConsensusUs helps groups reflect on each others' perspectives and to weigh multiple criteria simultaneously [85]. For instance, ConsiderIt visualizes community-sourced design trade-offs in an argumentation structure to support reflection [77]. Recent work on digital juries shown that people prefer decisions based on deliberation compared to voting alone [43]. These interventions suggest that voting is seldom enough and that deep reflection on the criteria and positions of others involved is necessary to support effective decision making.

Previous work has also explored different leadership structures for supporting design in online communities [70, 86, 87]. Leadership is not only about who makes the decisions but also about how different members assume specific roles or complete specific tasks. Luther et al. explored how leadership can be distributed in a creative community, where different members take the lead at different points in a creative process [87]. More recently, Mechanical Novel explored how

¹<https://dschool.stanford.edu/resources/get-started-with-design>

people can cycle between setting goals and doing lower level tasks to accomplish those goals, in the context of co-writing short fiction stories [70]. These leadership models show how community members can move in and out of positions of leadership over time. Consequently, there are times when individuals have more or less power to influence the design process.

Despite these attempts to support consensus building and decisions making, challenges remain when supporting equitable participation in PD initiatives [5, 18, 61]. There are challenges associated with who gets to attend, who gets to make decisions, and who gets to lead the design process [18, 61]. When introducing technology, such as discussion forums, additional challenges can emerge related to anti-social behaviors. Super-posting (i.e.: super-participants [54] or super-contributors [123]), where some community members contribute disproportionately more than others [66, 123], is common in large online discussions [66, 88] and open collaboration systems [45, 105, 110]. By dominating the conversation, super-posters can intimidate newcomers and have more opportunities for their voice to be heard and for their contributions to be recognized. More broadly, signal boosting is the process of amplifying one's voice through self-tagging and self-promotion [20]. Members can boost their contributions by posting frequently and by abusing features for social tagging [76, 101] and voting [20, 68]. These negative behaviors, which are problematic for most reputation systems [64], can limit opportunities for equitable participation by community members.

For each idea posted in Canvas, there is a discussion section where that idea can be deliberated and voted on by the community. Through discussion and voting, the community has an opportunity to reframe and prioritize problems and associated solutions. However, Canvas reserves the right to make final decisions about which ideas are ultimately implemented. This unique aspect of community voting combined with Canvas leadership provides an interesting case study to understand how discussion forums support or hinder PD practices which inspires the research question:

RQ3 What barriers do community members face when trying to influence prioritization and decisions making?

3 BACKGROUND ON CANVAS AND CANVAS STUDIO

Canvas is a learning management system that supports student assignment uploads, attendance taking, quizzes, online discussions, and countless other features used by students, teaching assistants, and instructors. The platform currently has over 30 million members at over 4000 academic institutions around the world [67]. To improve their platform and respond to community members' various needs across the globe, Canvas hosts *Canvas Studio*², which allows its end-users to "participate in Canvas product development priorities" by contributing to a design-oriented discussion forum (shown in Figure 1). This end-user community, referred to in this paper as *community members*, primarily consisted of instructors, but also includes students, teaching assistants, administrators, and technologists. Participation also spans multiple schools and grade-levels (k-12 and higher education). In Canvas Studio, members can post, follow, discuss, or vote on community ideas. Canvas uses these ideas during the discovery phase of the development cycle to understand its users' goals and needs. Community members can discuss ideas throughout phases of development, testing, and implementation. Canvas also employs a 'Community Team' responsible for moderating posts, responding to requests, and advocating for popular or impactful ideas from the community.

This paper explores how to engage large online communities in a long-term, sustained design process. Canvas is interesting to study because it features a large end-user community that contributes to the design process with minimal scaffolding. Community members need to make sense of the ideas that have been posted to orient themselves and contribute effectively. In addition, community members are not constrained to contribute to a single phase of the design process,

²<https://community.canvaslms.com/community/ideas/content>

as is often the case in open innovation. Instead, the community participates in a wide range of designerly activities that includes sharing, discussing, and prioritizing problems and solutions. By focusing on the use of a discussion forum—a ubiquitous and familiar technology—to structure and share design activity, we hope the resulting design insights may inform participatory design practitioners who want to equitably and authentically engage communities at scale.

3.1 Getting oriented, making contributions, and voting in Canvas Studio

In Canvas, technology guides the process by providing forms to submit ideas and threaded discussions to guide the associated conversation. Participation is open to any community member who has an account with Canvas which includes students, teaching assistants, and instructors. Community members can submit an idea by filling out two submission fields: “title your idea” and “add some persuasive details.” Members can tag other members, add keywords, or attach files to their posts. Canvas provides guidelines for submitting ideas in the documentation for Canvas Studio. The documentation asks community members to: 1) “Have one idea per submission for clarity in voting and discussion.”, 2) “Choose a title that describes the feature idea.”, 3) “Be specific about your suggestion.”, and 4) “Express your ideas clearly and concisely.” The guidelines also suggest including screenshots, use cases, desired outcomes, and goals. These guidelines appear to be helpful; however, these guidelines are not integrated into the workflow for submitting an idea and it is unclear whether community members review them before posting.

Canvas has a process for accepting and implementing suggested features. Our research team created the diagram presented in Figure 2 to summarize the process described by Canvas³. Once an idea is submitted, it is reviewed by Canvas community team. If the idea is incomplete or requires further clarification, it will be categorized as ‘moderating.’ Otherwise the idea will become ‘open-for-voting.’ At the end of the six-month voting period, ideas in the top 10% will remain open-for-voting, while the bottom 90% of ideas are archived. Being in the top 10% ensures that the idea will not be archived, but does not guarantee that the idea will be implemented. Ideas chosen for development are marked as such and enter the development cycle. After development and beta testing, the feature is integrated into Canvas and marked as ‘complete.’ During this development period, communities have little ability to observe or influence the process. Between Mar 18, 2015 to Aug 06, 2018, 474 ideas from the Canvas community were been accepted and implemented.

³<https://community.canvaslms.com/docs/DOC-14903-how-does-the-feature-idea-process-work-in-the-canvas-community>




Fig. 2. A process diagram that was created by our research team to describe what happens when an idea is posted on Canvas Studio. Throughout this process, community members use the discussion forum to elaborate on the proposed idea.

3.2 Researcher Positionality Statement

To acknowledge the potential for biases and subjectivity that is introduced by researchers, many have called for reflexivity in research [15, 28, 41, 111]. In keeping with these important considerations, we were careful with how we obtained and analyzed the data. To obtain the data, the lead author developed a relationship with Canvas through a series of emails and one Skype meeting. After the Skype meeting, the Canvas community team shared their data (which is also publicly accessible). To deliberately reduce bias when handling the data, the lead author made no promises about the use of the data or the desired outcomes from its analysis. For this same reason, we did not consult with Canvas about the results found in this analysis. Furthermore, the researchers received no funding or incentives from Canvas to carry out this research. As opposed to ethnography [12, 16] and participatory action research [98] in which researchers are situated at the site or active participants, and collectively shape the content and methods of the work, we focus entirely on an content analysis with no interactions between researchers and participants in the discussion forums. We did not want to impose our research agenda on Canvas community members who use the site to address problems, not to participate in research studies. Similarly, we also did not conduct interviews with community members because Canvas has a ‘no solicitation’ policy.

4 DATA COLLECTION AND ANALYSES

To evaluate our research questions, we conducted a mixed-methods analysis of data collected from Canvas Studio, a discussion forum where end-users contribute to an ongoing design process through posts and comments. To analyze the posts, we developed a coding scheme that we applied to sentences and phrases within each post. This helped us to understand what information community members contributed to the design process. Additionally, we analyzed the metadata of the posts to understand what effect affordances like tags, comments, votes, and views had on the process. Finally, we thematically analyzed posts and comments to identify challenges that community members faced, and used a concordance analysis to understand the prevalence of these themes.

Our data set includes posts and comments submitted to Canvas between Mar 18, 2015 to Aug 06, 2018 (three and a half years, or 41 months). This time frame represents all the data available on the platform at the time when the data was collected. This data was obtained directly from the Canvas Community team and included the metadata for every idea on their platform (10,854 posts total). This metadata includes ID, idea title, author, date created, view counts, up/down vote counts, URL, implementation stage, social tags, and the author’s account age. We augmented this data with the post content for 1418 design ideas and the associated comments for those posts. After removing 6 outliers that consisted predominantly of source code, we were left with 1412 design ideas. 171 posts did not have any comments and the remaining 1241 had 18,335 comments (14.8 comments per post). The comment data included the author, comment content, and date. Our analyses focus on these posts and the related comments as well as the metadata.

4.1 RQ1: Analysis of the Tags and the Canvas Tagging System

With the 10,854 ideas on the platform (including archived posts), members have to sift through many posts and discussion threads to get an idea of how to contribute. Most posts had been tagged with keywords to help organize them. This additional organization theoretically makes it easy to search and explore the large number of ideas. Trending tags can also bring attention to new and growing topics within the posts. Members can navigate by browsing through a threaded discussion forum or by clicking on a tag cloud visualization that leverages the tags applied by community members and moderators. These tags also contain information about the number of duplicate ideas on the platform. In this section, we analyze the tags applied to posts.

4.1.1 Using tags to identify duplicate posts. Based on the documentation, Canvas community members and moderators used tags to mark duplicate posts for removal. Duplicate posts were tagged with both *poezd* and *duplicate*. We used both tags because each tag on its own can have multiple meanings. According to Canvas documentation, the *poezd* tag indicates that a post is either a “duplicate idea”, “off-topic”, an idea with “no response”, or an idea that is “no longer relevant due to changes in Canvas.” Similarly, *duplicate* could be used to label duplicate ideas or features and needs in Canvas, such as ‘duplicating questions in a quiz.’ Consequently, only posts that contained both tags were considered by our research team as being actual duplicates. We analyze these tagged duplicates through descriptive statistics and a thematic analysis of posts.

4.1.2 Analyzing tags to see what is being organized. To understand the potential that social tagging has for helping community members get oriented to the design discussion, we analyzed the tags contributed by community members (*community tags*) and moderators (*administrative tags*). We observed that many of the tags were semantically similar, such as *grade*, *grader*, *grades*, and *grading*. We also observed that moderators were tagging posts for administrative purposes. To make sense of the 2701 distinct tags on the platform, we aggregated the tags semantically using the Dedupe python library⁴, a fuzzy matching and entity-resolution library based on semi-supervised clustering to find and aggregate similar tags [8]. By aggregating similar tags, we reduced the set of distinct tags to 703. Next, we separated the administrative tags by extracting tags that represented either a date or tags that had 0% de-duplicated (every tag in that aggregated cluster was the same, e.g.: *radar_stage*). These administrative tags, applied by moderators, related to the implementation process rather than the semantic content of the idea. This process allowed us to differentiate and compare tags.

4.2 RQ2: Analysis of Community Design Contributions in Discussion Posts

To understand how designers contributed to the discussion forum and what information they included in their contributions, we developed and applied a qualitative coding scheme. The coding scheme was then applied at the phrase and sentence level.

4.2.1 Developing a qualitative coding scheme for designerly contributions. Several existing coding schemes describe design behaviors. In some cases, these have been used to compare the design behaviors of novice and expert designers [3, 30, 125]. To enable comparisons, these coding schemes generally focus on the individual. Coding schemes can also focus on collaborative design. Kou and Gray’s coding scheme examines design in online creative communities, but focuses on critique [73]. Paulini et al. [108] present a coding scheme with a similar context to ours (Quirky’s open innovation platform), introducing the behaviors *ideation*, *social communication*, *evaluation*, *referencing*, and *qualifications*. However, these codes focus on ideation. They also include codes about how people refer to their domain expertise, which was uncommon in our context because most community members were not trained in design. Given the limitations of existing coding schemes in relation to our observations, we developed a coding scheme to identify the units of information contained within each post.

We developed our coding scheme inductively by reading through a set of 100 posts to identify themes. We used both open coding and in vivo coding [49, 113]. The themes were discussed within the research team and negotiated to improve agreement. Following a series of training phases, the codes were adjusted to better cover the data and improve inter-rater reliability (IRR). After this process we finalized the code book with the following codes: Solution (So), Workaround (W), Problem, (P), Situation (Si), and Acknowledgement (A). These codes, summarized in Table 1, also relate to codes identified in other coding schemes. The code *Acknowledgement* relates to Jennifer

⁴<https://dedupe.io/>

Code	Description
So	Solution description, design details, or outcomes
W	Makeshift workarounds that were not intended by Canvas
P	Problem description or an outcome of a problem
Si	Situational context related to a class or to the Canvas system
A	Acknowledge others: reference posts or people, request votes, social talk [124]

Table 1. An overview of the codes applied to the segmented posts.

(P11697) "Hello, [A] We would like to request a feature that would allow a teacher to mute an assignment in Canvas without muting the feedback that comes from student peer reviews. [So] How things currently work: If you mute an assignment that has peer review enabled, the feedback from peer review will also be muted. [Si] This means that muting hides both grades from the teacher and feedback from other students. [Si] How things should work: It should be possible to, when muting an assignment, allow for feedback from peer reviews to remain visible while still hiding from students grades that the teacher may be working on. [So] Our reasoning: These are two separate concepts and should not be affected by the same feature. [P] It is very reasonable for a teacher to want to work on grades and keep this aspect muted at the same time that students provide and view feedback between each other. [Si] This feedback does not influence their grades. [Si] and therefore it should be possible to mute grades without muting feedback. [So] Thank you! [A]"	Acknowledgement Situation Problem Solution Workaround
---	---

Fig. 3. A coded post shows how the coding scheme was applied and how the data was unitized.

Stromer-Galley's conceptualization of *Social Talk* [124] and to Paulini et al.'s *Social Communication* code [108]. *Problem* is similar to the concept of *problem talk* [124] and *problem definition* [2].

4.2.2 Segmenting the posts and applying the qualitative coding scheme. To apply our coding scheme, we segmented each post into single units of analysis. Segmenting (i.e.: unitization) is the process of breaking data into smaller units to be analyzed. In our case, we applied a single code to each segmented unit. An example of how we segmented and coded a single post is shown in Figure 3. Segmenting and obtaining inter-rater reliability (IRR) for continuous data is challenging because coders need to agree on the bounds of the segment along a continuum. For these reasons, some researchers have advocated for segmenting on natural sentences where the punctuation indicates the beginning and end of a segment [32]. Daubler et al. have shown that segmenting on natural sentences improves IRR without reducing the value of the resulting codes [32]. However, this simplified approach does not work if multiple codes appear within a single sentence. More rigorous methods require coders to agree on both the segmentation and also on the coding [96]. Mathet et al. propose the Gamma measure [96] which computes agreement along a continuum to ensure that codes agree on the start and stop characters of a segment. Similarly, Krippendorff et al. extended their Alpha measure to account for start and stop characters, introducing Krippendorff's α_u [78].

The data was coded by three independent coders. Each coder segmented the data and applied their codes based on the coding scheme. Coders tried to segment the data at natural sentences when possible (relying on the earlier insight of Daubler et al. [32]). For training, we applied our coding scheme to ten posts at a time and then reviewed the codes and segmentation. After mediating disagreements and creating heuristics to improve agreement, the coders moved on to another set of sentences. When we reached agreement ($\alpha_u=0.75$) on both the codes and segments (as measured by Krippendorff's α_u [78]), we moved on to a testing phase. In our testing phase we obtained inter-rater reliability of $\alpha_u=0.70$. This indicates moderate to substantial agreement between the three coders [83]. To reiterate, Krippendorff's α_u was used to account for multiple raters, how they each segment the data, and the agreement of their codes. It is extremely rigorous because raters have to agree on both segmentation and coding. By both segmenting and coding the data, we are

able to compute how much focus is placed on each code (number of words and number of times the code is applied in a single post).

The research team independently applied this coding scheme on a subset of the posts (354 randomly selected posts, 25% of the total data) and used memos to track observed trends [9]. We calculated the sample size for our population of 1412 posts using sample size analysis (i.e. population proportion [93]) with a 95% confidence interval, 5% error tolerance, and a conservative 50% sample proportion, resulting in 303 posts. We added an additional 51 randomly selected posts to reach the round number of 25% of the data (354 posts in total).

4.2.3 Thematic Analysis of Segments Coded as Problems or Solutions. Problems and solutions encode a lot of information about the design space. They introduce constraints, articulate goals, and describe expected outcomes. After applying our coding scheme, we further explored how problems and solutions were structured using collaborative affinity diagramming [60]. This helped us to understand not just what was shared but also how it was structured. We randomly extracted a subset of 120 segments that were coded as either problems ($n=60$) or solutions ($n=60$) randomly from the coded posts. We conducted a deeper thematic analysis of each of these aspects because we observed that people presented solutions and problems in a variety of ways during our coding and memoing process. Including these components helped us understand what people focus on in their posts, and digging deeper helped us understand how they present and structure these two aspects of their posts. Finally, problems and solutions are core aspects of the design process and we want to better understand the designerly behaviors of community members.

4.3 RQ3: Analysis of Idea Adoption and Social Behaviors

To understand whether and how community members were able to influence decision making, we analyzed the effects of community engagement (e.g.: views, votes, etc) on an idea's status, the effects of anti-social behaviors such as superposting, and the reactions of community members when ideas are not adopted.

4.3.1 Analyzing correlations between community influence and idea status. To understand the potential that discussion affordances have to drive idea adoption, we analyzed correlations between social behaviors like commenting, viewing, and voting with the status of an idea. Discussion posts often have zero associated comments, views, and votes [81]. We used Kendall's Tau-b to compute correlations because our dependent variables (views, votes, and comments) are continuous and zero-inflated (i.e.: contain many zero-valued data points) [34, 109]. Zero-inflated data are seldom normally distributed and can not be log-transformed, due to the zero values. Kendall's Tau-b, like Spearman's Rho, does not assume normality and treats the data as rank-ordered. However, unlike Spearman's Rho, it is robust to ties, which occur frequently at zero for zero-inflated data.

4.3.2 Analyzing superposting behavior. As discussed in the related work section, super-posting is common in online discussions [66, 88]. Superposting is one way that discussion participants get more attention and have more opportunities to influence the discussion. Based on the definition by Huang et al., we defined a superposter as someone in the top 5% of the posters [66]. This 5% accounted for 51 posters of the total 1015 posters in our dataset and 19.9% of the total posts. We analyzed the amount of attention received by these superposters compared with other posters on the platform using an accumulation plot for views, votes, and comments.

4.3.3 Thematic Analysis of Discussion Comments. To better understand the role of the community in the Canvas design process, we also analyzed the comments associated with the posted ideas. These comments were contributed to a threaded discussion underneath each submitted idea. To analyze these discussions, we targeted ideas that had been open to contributions for an extended

period of time by filtering for posts tagged with *dobrij*, a label used by the Canvas community team to indicate that an idea had been in the *open for voting* stage for at least 6 months. The ideas were sorted by number of votes and discussions associated with 25% of these ideas were included in our analysis. Posts were sorted by votes and uniformly sampled to get a representative sample of ideas that receive many votes and few votes. This data was used to better understand how online communities influence the design process through online discussions.

4.4 Concordance Analysis of Posts

Concordance is a method for displaying a search term with the words immediately surrounding it in a document. This additional context allows the searcher to determine whether the search term is found in the corpus and how it is used in context. In previous sections, we describe multiple thematic analyses. To better understand whether and how the themes were representative of the whole corpus, we used concordance as a mixed-methods analysis tool to quantify the qualitative trends that we observed. Based on our memos, we identified keywords that exemplified each theme. As an example, in one analysis, we observed that posters self-promoted their content by asking other community members for votes. In these instances, posters often used the words *vote* or *voting*. Using concordance, these search terms (vote and voting) were displayed as a list with the surrounding context of each post where they appear. For example, in a comment someone might say “I vote we move on” which is a way of saying that it is okay to continue, but it is not an idea or problem related to Canvas’s voting features. We report on two results for our concordance analysis, first how many times each search term was matched in the whole corpus across 1412 posts. Then by reading the sentence in which the keyword is contained, we manually coded whether the matched search terms exemplified the theme. Finally, we report on the percentage of posts that featured a matched search term that exemplified the theme. The concordance analysis was used across our different qualitative methods to determine the extent to which our observations extended to the larger dataset without needing to read every posts in detail.

5 RESULTS

Canvas supports a collaborative design process that starts when members post ideas to an online discussion forum. Ideas are discussed and voted on by community members to determine whether new features should be integrated into the Canvas product. This raises research questions about how community members orient themselves in the design process, how they contribute, and what impact they have on the Canvas design process. At the time of collecting our data, the Canvas platform featured 1412 active posts and 18,335 corresponding comments.

5.1 What challenges do community members face when orienting themselves to ongoing design conversations in a discussion forum?

To better understand how community members oriented themselves to the design discussion in Canvas Studio, we analyzed the tagging system to see whether it was helpful for getting oriented and as a way to quantify the number of duplicates on the platform.

5.1.1 Members repeat ideas and lack awareness of the uniqueness of their contributions. Based on the metadata provided by Canvas (10,854 posts), we identified 1266 duplicate posts (11.6% of posts). 11.6% is likely a conservative estimate, since we observed many cases of repeated topics that were not tagged as being duplicates. For example the following two posts were missed in our analysis (not tagged as both *poezd* and *duplicate*):

Edit posted comments in Grade book (P1675) Occasionally, I have the need to add comments at different times for a grade in the grade book. Or sometimes I have entered

a wrong comment for a student. There is no way to edit a comment once it is posted to the gradebook. I would like to have the capability of editing a comment so that I can add to an existing comment or remove information that is no longer needed.

Editing assignment comments (P4093) I would like to be able to edit comments that I make in the speed grader on student assignments. Often I find myself wanting to correct a type, or amending a phrase. The only possible way to do it now is to make a new comments. I keep looking for an editing feature similar to the way we can modify announcements.

These two ideas are very topically similar—about editing comments placed on assignments when grading—but they are not included in the 11.6% presented earlier. The large number of posts and comments combined with the many duplicates may explain why multiple posters explicitly expressed concern about potentially posting duplicates. For example, by saying, “Sorry if this is a repeated idea. I couldn’t find it elsewhere and I’m a newbie” (P1555). Another poster said, “This may already be posted somewhere, but I cannot find it.” (P11072).

To further understand the extent to which community members discussed the possible overlap of topics raised in different parts of the discussion forum, we conducted a concordance analysis (described in section 4.4). We used the search words *originally* (17 matches), *posted* (51), *repeat* (16), *find* (157), and *already* (85). By manually coding these 326 matched words in the context of their use, we see that only 16 unique posts matched the theme that ‘participants were unsure if their post was a duplicate.’ This suggests that at least 1.1% of posters (16/1412) explicitly expressed concern that their idea was a duplicate. Despite challenges associated with duplicate ideas and navigating to find relevant ideas, some posters appeared to be aware of other community members and relevant ideas. In a few cases, posters quoted other posts or referred to other posters by name, “I believe this request was first raised in early 2015” (P4169). One poster quoted another member’s post, “The original text provided by the one and only [name omitted] is as follows” (P4095). Another mentioned that a specific idea “has come up numerous times, with slight variations” (P9135). This indicates that at least some community members not only did the work to read the ideas of others’ before posting, but actively attempted to raise awareness of that work for others.

5.1.2 Community members and admins used the same tag system for different purposes. Based on the 1412 posts for which we had the idea content, we observed 9556 tags with an average of 6.8 tags per post. Of these 9556 tags, we found 2701 distinct tags. Some of these tags were used by administrators to track release dates and share information with developers, such as *poezd* and *dobrij*. A comparison of the tags (community tags, administrative tags) is presented in Figure 5. Community members

Browse all tags in Ideas




Fig. 4. A screenshot of the tag cloud provided in Canvas Studio to aid navigation. Community members click to select one or more tags to browse tagged ideas. Only *poezd*, *duplicate*, and *instructor* appear to stand out.

	Community Tags	%	Count	Administrative Tags	%	Count
1	instructor	11.0	571	mar18	45.4	229
2	quiz	77.4	469	radar_stage	1.4	221
3	gradebook	76.2	416	jun18	18.3	115
4	modify current	31.4	309	dobrij	0.0	113
5	student	37.9	224	apr18	7.1	112
6	new	12.4	217	may18	14.0	107
7	admin	15.2	210	jul18	17.9	78
8	assignment	76.3	207	panda-cat	3.7	54
9	canvas teacher	88.9	189	q.next_complete	0.0	48
10	course	78.7	155	EmptyTag	0.0	39
11	teacher	23.3	129	aug18	64.7	34
12	discussion	72.9	118	existing	0.0	32
13	speedgrader	32.4	111	ta	0.0	18
14	groups	66.4	110	inbox	0.0	18
15	ui	4.8	105	rce	0.0	17

Fig. 5. The top 15 tags contributed by community members and the Canvas Community Team. Tags were automatically split using heuristics into community- and administrative-oriented tags. % indicates the percentage of tags that resolved to the primary tag after applying entity resolution (e.g.: combining instructor, instructional designer, instructor role, and instructors). This indicates that the community and admins used the same tag system for different purposes.

appear to have difficulty finding how to contribute to the growing conversations around design, as evidenced by the presence of duplicates. Though tags may help, it is not clear how effective this mixed-use tagging system was for orienting individuals. This secondary purpose for tagging by moderators likely made navigating to relevant ideas challenging because community members need to sift through tags about both content and process. This suggests that in community design contexts—where there are tags about the design, development, and release of a product—multiple tag systems may be needed to orient different stakeholders. We also observed that many of the tags were semantically similar, such as *grade*, *grader*, *grades*, and *grading*. The semantic similarity of some tags and the presence of 2701 distinct tags on the platform resulted in a tag folksonomy where there were many tags with very few tags standing out from the others, shown in Figure 4.

5.2 RQ2: What design information do community members contribute to a large and evolving design discussion?

In Canvas, community members contribute by posting or commenting. These contributions can include design information such as problems, solutions, or design contexts. To understand whether and how community members expressed design information in posts, we segmented the posts and coded those segmented units. This allows us to see how much emphasis was placed on various aspects of design and to see how community members structured their design contributions. Analyzing these posts provides insights about community members' designerly behaviors.




Fig. 6. (a) Across all posts, more total words were dedicated to solutions than to any other code. (b) The percentage of posts that contained at least one instance of the code. 94.9% of the posts included a segment coded as a *solution*. This suggests that across posts, more attention was paid to solutions.

5.2.1 Community members focus mostly on solutions and often did not specify problems explicitly.

After segmenting and coding the posts, we analyzed how many of the posts included at least one instance of each code. We found that 94.9% of posts contained a solution, but that only 56.8% of the posts articulated a clear problem (see Figure 6). Occasionally, the problem was implied through the situation or the proposed solution, but not explicitly stated. For example, the following post included three segments coded as solution. They imply the problem, but do not clearly articulate it:

“As an account or subaccount admin, I want to set default notification settings for my users by role. [So] Like role-based permissions, these notification settings could be locked or unlocked; [So] that is, I could choose to allow users in a given role to change a default notification setting or not. [So]” (P1051)

In this example, the problem—that some people need to receive notifications while others should be able to opt out—is only implied. It is unclear whether the ‘users’ are students, teaching assistants, or instructors. Their description of ‘roles’ is similarly ambiguous. The lack of a clearly defined problem makes it harder for others to reframe the problem and makes it difficult for Canvas to determine if a solution meets the implied need.

Solutions were much more common than problems, but even when problems were included, posters still dedicated more words on average to the solutions ($avg = 23.2, \sigma = 11.8$) than they did to the problems ($avg = 21.4, \sigma = 11.8$) as measured by a t-test ($t(1179) = 2.43, p < 0.05$). Assumptions of the t-test were verified visually and via the Levene’s test. This was also higher than the number of words dedicated to acknowledgements ($avg = 13.9, \sigma = 18.3$), situations ($avg = 17.9, \sigma = 8.4$), or workarounds ($avg = 19.7, \sigma = 10.8$). The percentage of segments associated with solutions (39.1%) was also higher than it was for problems (19.4%), situations (26.7%), acknowledgements (11.1%), or workarounds (3.7%).

5.2.2 Key contextual information was often omitted. To implement an effective solution, it is often important to understand the problem, the context surrounding the problem, and the potential solutions that might be relevant for the problem. As explained above, few posters specified the problem explicitly, and even fewer posters included context through which the problem could be better understood. Only 40.4% of posts included a problem, solution, and description of the situation. Often, even when information was included, there were aspects that were left unclear. In the following example, the problem is articulated, but the ‘situations’ that justify the need are not explained:

“There are situations where instructors need to review the grading history for grade challenges and grade changes, but currently cannot without the admin changing the Teacher end date setting.[P] (P1850)”

In this example, vague information makes it challenging for designers and developers to address the problem, or to know how solving the stated problem would impact other stakeholders. This requires designer and developers to fill in the gaps on their own or to ask for more information.

5.2.3 Problems and solutions are expressed in a variety of ways. While previous sections describe the extent in which members express problems and solutions, those analyses do not give insight into how they are expressed. To further analyze the segments that were coded as *problem* or *solution*, we engaged in collaborative affinity diagramming [60] to identify these aspects within the codes. This helps us to understand how people structure information and collaboratively explore a design space. Quotes exemplifying each theme are presented in Table 2.

Themes	Definitions	Examples
Solution Talk		
Proposed Solutions	Suggest a high-level solution	"I suggest that Canvas develop a feature that would allow an admin to simply set the course settings defaults prior to the creation of courses ..." (P4922)
Goals	Define an overarching, wide-reaching objective	"We want our teachers to use Canvas to create engaging, meaningful, creative and collaborative learning experiences for students." (P4883)
Use Cases	Place the solution into a context in which it might be used	"I may want to order them alphabetically, by the way they fall in my schedule, those that require the most time and/or effort, etc." (P3450)
Feasibility Analyses	Evaluate the ease of implementation or the difficulty	"Seems like all the info necessary to create a branded mobile login is already in the Theme system." (P5891)
Alternatives	Propose alternative solutions to an ideal	"..., but if you can't do that, surely you could skip the external tool intro page when Canvas detects that the address is for Quizzes.Next." (P11907)
Design Details	Specify details about how to realize the solution	"Here are suggestions related to this idea for the development team: 1. Make an option available to save a test and/or a question bank electronically. 2. Create a database of questions that can be searched and viewed rather than only by opening each question bank separately. 3. Create a flag that ..." (P11568)
Outcomes	Describe the impact on specific users	"This would make it easy to start and stop the video while editing." (P9626)
Problem Talk		
Unwanted Behaviors	Describe a negative or unexpected behavior	"when a blueprint courses' content are updated and is synced, the order of the content is changed into a different order to what is designed from the blueprint" (P11286)
Inconvenient Workarounds	Share a non-ideal solution that was not intended by Canvas	"The only alternative (displaying one question at a time), means that the image/scenario/case study/data needs to be repeated with each question, which is cumbersome and confusing..." (P3912)
Negative Outcomes	Explain unintended outcomes	"Some faculty members feels it's incredibly concerning, as it is amplifying student competitiveness and anxiety, and marginalizing certain students." (P11650)
Limitations	Describe a need that is not being met	"...but we have no way of downloading a record of the Assignment and we have no way to download the rubric" (P1275)

Table 2. Examples from the posts for each of the themes associated with solution talk and problem talk.

Solution-talk When talking about solutions, community members contributed a variety of information to help fully convey their ideas. In previous sections, we show that community members talk about solutions; this section examines how they communicate their solutions. Our thematic analysis revealed seven themes describing how communities describe solutions: proposed solutions, goals, use cases, feasibility analyses, alternatives, design details, and outcomes. The way these contributions are communicated varies and impacts the emerging design space. The theme *proposed solutions* often implied a problem, but the proposed solutions could be ambiguous if the problem itself was not explicitly defined. *Design details* were specific to an individual idea, but it was often difficult to infer how these details relate to a corresponding problem or solution. *Goals* were mostly independent of specific solutions and often included an impacted user group (e.g.:teacher, student), adding generalizable knowledge to the design space that could be expanded beyond an individual idea. We also identified some solution-specific language that people used when talking about solutions. This language included: “it should”, “[this, it] would [be useful, be awesome, help..]”, “I [want, would like]”, and “it could.” These varied ways of describing solutions means that although most community members shared a solution, the way that contribution is communicated may introduce ambiguity or lack sufficient context.

Problem-talk Just over half of posts included a segment that was coded as *problem*. Our thematic analysis shows that there was less variation in terms of how people describe problems when compared to solutions. We identified only four themes from our thematic analysis of problems: unwanted behaviors in Canvas, inconvenient workarounds, negative outcomes, and limitations. A common idea structure we observed when communicating *unwanted behaviors* is first orienting the user to a specific workflow in Canvas, then describing how that behavior is problematic. This contribution helps establish context within the Canvas system, but typically leaves out information about the impact of these unwanted behaviors. *Inconvenient Workarounds* are often used to justify the legitimacy of the problem because the burdensome workaround emphasizes a deficiency in the Canvas platform. While sharing these workarounds can help validate the urgency of an issue, they don’t effectively frame the problem or present solutions. These various ways of describing problems shows that while around half of posts included a problem, they may have only described a single negative outcome and not engaged in deeply framing the problem. This may indicate the design novices need additional scaffolding to help guide them through the problem framing process.

5.3 RQ3: What barriers do community members face when trying to influence prioritization and decisions making?

In Canvas, the design process is mediated by social practices and technology affordances. The social dynamics associated with a discussion can be positive, where members of a community mentor each other and provide guidance; however, they can also be negative, such as trolling and bullying. Canvas facilitates these social interactions through a discussion forum. The forum presents posts sorted based on recent activity and uses voting to prioritize ideas. In this section, we explore challenges members faced when trying to influence the design process.

5.3.1 Getting views led to traction, but receiving votes did not guarantee acceptance. Discussion forums can organize content by recency, topic, author, among other factors. These factors can affect what gets seen and consequently what gets attention from the community. Because Canvas organizes posts by recent activity, recent comments push posts to the top of the discussion forum and increase the likelihood that the post is seen. We observed that the number of votes that a post received were significantly correlated with views ($\tau_b=0.704, p < 0.01$). The number of comments was also significantly correlated with the number of views ($\tau_b=0.593, p < 0.01$). This finding aligns




Fig. 7. Boxplots of idea status by view, vote, and comment counts. Having more views or comments appears to correlate with higher chances of acceptance. Having more votes, appears to have less of an effect. These differences were not tested for significance for reasons outlined in the limitations section.

with the *rich get richer* phenomenon documented in various socio-technical systems [39]. This raises questions about whether the best ideas are actually receiving the most attention.

Voting ideally focuses the communities' attention on the most promising solutions or the most important problems. However, Figure 7 suggests that the number of comments and views tended to have a higher effect on a post being accepted or rejected than votes. It appears from the graph that unpopular ideas (few views, votes, or comments) are rejected, however there does not appear to be a clear difference between highly voted ideas that remain open for voting and those that get implemented. Canvas has a policy of removing ideas that do not remain in the top 10% of voting every six months. In this way, not receiving votes can lead to rejection, but receiving votes does not guarantee acceptance. Through these examples, we observe how the affordances of discussion forums potentially guided community attention and prioritized some ideas over others.

5.3.2 Posters' attempted to game the system by re-posting old ideas. The previous section shows how technology affordances may be prioritizing some ideas over others. These affordances can also be manipulated by the community to drive attention to their ideas. For example, the discussion forum sorts posts by recent activity. By repeatedly commenting on one's own post, a poster could pin their post to the top of the discussion forum. We observed two main gaming behaviors: 1) posters re-posted stale ideas and 2) specific posters shared disproportionately more than others.

First, posters re-posted a variation of an old idea when the original did not get enough traction. This re-framing process is often an important part of the design process [36, 107, 117], but in this case can lead to many duplicates and appears to be a strategy that some posters use to gain attention for their ideas. One poster hypothesizes that timing matters when posting saying, "I'm resubmitting this idea, which failed to get enough votes. Maybe now is a better time for this idea!" (P4224). Other posters thought that posting again might help get traction saying, "Hi all, this idea has already been requested, but did not get enough votes." (P5689). Similarly, a poster shared,

"I have suggested this in the past but it didn't seem to gain traction so I thought that I would bring the idea up again." (P5489)

We used our concordance analysis to see how frequently this theme of 'deliberately re-posting an idea to get more attention' occurred across posts using these search terms: *originally* (17 matches), *posted* (51), *repeat* (16), *again* (50), *find* (157), *already* (85), *similar* (94), and *repost* (1). Across these 471 matched instances, we observed that 51 instances exemplified the theme. This suggests that in at least 3.6% (51/1412 posts) of the posts, community members described how they gamed the system by deliberately re-posting a duplicate idea to get traction and a better result.




Fig. 8. (a) Cumulative views, votes and comments by author. Approx. 35% of attention is received by 5% of the authors. (b) Views, votes and post counts aggregated by author. Each circle indicates a single author. There is a linear relationship, with some authors receiving disproportionate attention.

5.3.3 Posters’ attempted to game the system by posting disproportionately. We observed that some posters contributed disproportionately more than others. The top 5% of posters accounted for 19.9% of the total posts, 37.4% of the total views and 34.5% of the total votes for the entire platform respectively, as summarized in Figures 8a and 8b. This suggests that superposters post more than their peers and their posts receive more attention as well. Compared to Huang’s results (44.0% of posts coming from the top 5%) we can see that superposting was less severe in our context. But it is interesting to note that while posting behavior may have been more equitable, attention received was much less balanced.

5.3.4 The lack of transparency around the implementation process frustrated community members. In addition to posters that attempted to game the system to receive additional attention, community members also championed their ideas by asking for votes. This indicates a mismatch between the mental model of the contributors and Canvas’s process for accepting ideas. Having the top 10% of votes guarantees that the idea won’t be archived, but it does not ensure that the idea will be implemented. Canvas developers consider many aspects that posters may not be aware of, such as the maintainability of the platform, the difficulty of implementation, or effects on the broader community. Regardless, Canvas community members emphasize the importance of votes. Members self-promote their ideas, saying “vote if you agree” (P11961) or “If you share this need, please upvote this feature idea” (P12005). They even advocate for related features or features suggested by other posters, saying “please also vote for [hyper-link to post] when automatically assigning peer reviews” (P8154). In our concordance analysis, we used the search terms *vote* (50 matches) and *voting* (18) and observed that 1.9% (27/1412) of posts featured a request for votes.

Despite community emphasis on votes, the Canvas policy states that “relevance to a priority, voter demographic, total votes, and comments are just some of the factors taken into consideration.” Despite this disclaimer, many community members seem to believe that highly supported ideas (as measured by receiving many votes) will be implemented, which leads members to express confusion and frustration directed at the voting process when well supported ideas are not implemented.

To better understand how this lack of transparency was interpreted by community members, we explored the comments related to a set of 50 posts that received a sufficient number of votes to remain in the open-for-voting stage for long periods of time (marked *dobrij* by Canvas). Analysis of these comments reveals frustration, as one community member expressed:

“So you only make fixes for major bugs or feature omissions that inconvenience the Canvas community on a large scale and demand community involvement? This voting process doesn’t promise much hope for improvement of the features I would like to see at all. I haven’t seen a single improvement or feature addition for any of the threads I’ve found here...Can I vote on the voting system?” (C70496)

Others reiterate their belief that many votes guarantee that their ideas will be implemented, often by highlighting the number of votes they have received to advocate for their solution being implemented. “With 300+ votes has this been tagged for progression? 378 votes... 414 votes...when will this progress?” (C97383) and “Looks like we have the votes for this one. What are the next steps?” (C83298) “...100 votes is no longer the relevant benchmark” (C83339). In each of these cases, the magnitude of votes that the commenter suggested as a goal differed by a magnitude of four.

In addition to votes, others highlighted how long the post had remained in the open-for-voting stage by consistently maintaining the top 10% of votes. In our concordance analysis, we used the search terms *vote* (50 instances) and *voting* (18) and found that 1.2% (17/1412 posts) of posts included a complaint about the lack of transparency around voting.

5.4 Despite challenges, community members remained motivated and engaged

Despite the challenges related to equity and empowerment that have been described throughout this section, we also identified many instances where community members were motivated and cared deeply about their ideas to improve Canvas. This success speaks to the challenge of maintaining participation in large-scale design initiatives mentioned in end-user innovation efforts [47]. This motivation is demonstrated by the 10,854 posts that community members contributed to the platform and by the lengthy discussions associated with many posts. Figure 9a shows an initial discussion blooming period followed by sustained discussion for all 367 posts in the 2015 cohort. In the discussions, posters expressed waiting with anticipation, “I could not agree more; it’s been over a year since I first asked about it” (C117014). Another poster expressed frustration, but demonstrated significant effort invested and a strong desire to see their idea implemented, saying,

“I used to participate in Canvas community discussions and feature requests on almost a daily basis. I had hoped to see our features developed. But I finally gave up and so stopped about 2 years ago. However, this would truly make a difference in my teaching life, so I am trying again.” (P9029)

Despite the frustration, members clearly care about their ideas and want to see them adopted by Canvas. When ideas are implemented, community members are very excited and appreciative. One commenter said, “The update is such great news, isn’t it?!” (C63164). Other reacted to the news that an idea as being implemented with celebrations, saying, “HOORAY!!!!!!!!!!!!!!!!!!!!!!!!!!!!!! This news makes me very happy!” (C65717). Members encourage each others’ teaching practices, “Bravo David!!!! I have the same basic teaching philosophy...” (C33210) and design ideas, “Yes!!! Brilliant idea. This is the number one problem our faculty have with Canvas right now.”(C34147). Many instances like these demonstrate the care, effort, and community cohesion.

This effort is also evidenced by how quickly community members commented on the ideas of others. As shown in Figure 9, many of the posts received timely feedback as measured by the first comment received. 49.6% of posters received their first comment within one day of posting, 79.7% received their first comment within 7 days, and 90.1% of the posters received their first comment within the first 24 days. These response times are reasonable; however, it is important to remember that this marks the beginning of a conversation. Additional discussion might happen for months after this first interaction before an idea is selected to be implemented (if it ever is selected). Given




Fig. 9. (a) A density plot for 9314 comments made on 367 posts over time, all from the same 2015 cohort but aligned by their creation date. (b) A graph depicts the time until a post received its first comment. Almost every post received a comment in the first five days and some conversations continued for multiple years.

the high-latency associated with ideas being implemented, this timely encouragement in the form of a comment is likely an essential motivating factor for community members.

6 DISCUSSION

Our analysis suggests that community members face challenges orienting themselves around the many shared problems, solutions, workarounds, and relevant situations. Due to this lack of clear orientation, 11.6% of ideas were tagged as duplicates and potentially many more ideas could have been labelled as such. This challenge stems from the enormous scale of the design effort where the community posted 10,854 total ideas (474 implemented) over a 3-year period. Our deeper analysis showed that only about half of all posts described an underlying problem (56.8%) and the comprehensive information needed to make design decisions was spread across multiple posts and comments. Consequently, the Canvas Community team needed to patch together design information or make guesses to fill in the gaps. Furthermore, we observed deliberate attempts by community members to prioritize their ideas over others. For instance, some members gamed the system by posting disproportionately in an attempt to have their voices heard. 5% of posters on the platform accounted for 19.9% of posts and approximately 35% of the views, votes, and comments. Finally, we observed how a lack of transparency in Canvas about which features were implemented frustrated participants. Despite these challenges community members were highly engaged in the design process over multiple years, resulting in the implementation of 474 community-submitted ideas. Our work demonstrates challenges associated with using discussion forums to support large communities engaging in design at scale.

6.1 Design considerations for supporting large-scale collaborative design efforts

This paper explores how a large online community of end-users proposed and deliberated problems and solutions. We found that community members faced challenges getting oriented, making design contributions, and fully participating in the design process. Based on our findings, we offer three considerations grounded in Participatory Design, Computer-Supported Cooperative Work, and Social Computing research.

6.1.1 Offer training and multiple entry points into the design process. Our analysis provided indications that community members struggled to find entry points into the design process. Not only did we observe instances of duplicated effort (11.6% posts marked as duplicate), community members expressed uncertainty over whether their contributions would be unique and helpful, indicating that people struggled to navigate the many ideas posted on the platform. Duplicates are common in

large online communities, such as Reddit, where over half of links that became popular have been previously posted [48]. We also showed how moderators used the social tagging system to track releases and development stages. The presence of duplicates and irrelevant tags made it difficult for community members to orient themselves and find a place to contribute productively.

Previous research offers guidance on helping community members get oriented within a community discussion. Some advocate for socializing members to the social practices and norms of community [24, 26, 53, 126]. Another strategy involves helping community members identify relevant content and get a sense for what happened before they arrive, such as methods to summarize [129, 132, 133] and recommend [82] content. However, approaches to automatically summarize a prior discussion might be more difficult in this context where ideas emerge and decisions unfold over time. To represent an evolving process, one possibility would be to take inspiration from *Build-in-Progress*—a system that visually shows how a design process unfolds over time [127]. Sites like OpenIDEO [79] organize design discussions into design phases (e.g.: problem framing, gather information, ideation, develop solutions, evaluation [125]) which can help create entry points associated with each of these phases. Structuring discussion around the design process might help to improve transparency and more clearly convey progression while giving less experienced members a clearer indication of how they might contribute.

Finally, additional guidance could be provided to teach community members how to make high quality contributions. Currently, Canvas provides a static text form to post ideas. Additional scaffolding, such as just-in-time learning modules [106], could help community members develop design skills through mini-training tasks. Another option is to provide adaptive scaffolding to guide members to deeply reflect on their ideas. As an example, chatbots can engage users to reflect on their activity [72]. Similarly, a chatbot could coach members to make design contributions by focusing their attention on relevant design information to include. Scaffolding might help individuals develop their design skills and reflective practice which can be applied in other PD efforts.

6.1.2 Assemble design contributions using an asset-oriented perspective. Community members contributed to the design process in a variety of ways: sharing their own ideas, discussing ideas with others, and voting on how to prioritize problems and solutions. In general, the composition of posts was highly variable, with many posts omitting key information (e.g. the problem being addressed) that would be necessary to address the communities' needs. Worse, important design information was often distributed across multiple posts and comments, requiring community members and the Canvas team to stitch ideas together.

Omitting key design information likely made it difficult for the Canvas Community Team to implement solutions that meet the needs of the community. While each person brings different ideas and perspectives, it is unlikely that any one person has a full picture of the design space. An asset-based view of community development posits that skills, information, resources, and other assets are distributed among community members [97, 99]. Under this view, forum posts that offer partial bits of information can be seen as *design assets*: lived experiences and knowledge that can be collected and pieced together to create a community resource. These assets—spread across the community and shared in posts and comments—represent a large and complex design space. Similar to crisis response scenarios where residents and first responders contribute to an evolving information space on social media to coordinate aid [75, 116], we argue that these kinds of participatory design and end-user design communities contribute to an evolving design space where each individual contribution represents a design asset that further reveals and shapes the design space. By viewing assets as contributing to an emergent design space, “duplicate ideas”, which some consider a disadvantage of collaborative ideation [120], may actually provide value because each idea might introduce slightly new information and constraints. Duplicates may also

help to reframe the problem, which is a valuable part of the design process [36, 37, 107, 117]. This view of PD begins to bridge constraint-based models of design spaces [10, 19] with theories about how knowledge is socially constructed [13, 102, 122] by collaboratively uncovering the constraints that holistically describe a problem-solution space. This also aligns with past insights about how problems and solutions co-evolve [38, 91].

The idea that the community is co-constructing a design space through individual designerly behaviors introduces new opportunities for technological support. Developing micro-tasks to identify, extract, and curate design assets from posts and comments could help coordinate community members and provide system developers with a more comprehensive informational resource. For instance, Solvent [21] shows that crowd workers can annotate documents and that connections can be drawn between similar annotations across multiple documents. Similarly, Canvas could help members reflect on and highlight the most important information in a discussion. Recursive summarization (e.g.: Wikum [133]) shows that it is possible to summarize a single discussion thread, but our results identify a need to summarize and synthesize design ideas from design assets that are distributed across multiple discussion threads. It is possible that by combining these two methods, it might be possible to identify and extract design assets from multiple discussion threads and stitch them together. These micro tasks could support computational approaches that make connections across a community discussion. As one example of summarizing a design space, BlueSky uses a dimensional representation to guide contributions from crowd workers to optimize coverage of a design space [65]. These methods show how computational models of a design space can be constructed to represent a design space and used to coordinate community actions.

These annotated design assets could constitute data as the basis for design space visualization [65, 90, 120]. Such visualizations might help community members find places in the design space to make contributions. Finally, by building design ideas from information distributed across multiple threads, attribution could be shared by multiple contributors. Sharing attribution is one potential step toward mitigating the anti-social behaviors observed in our analysis.

6.1.3 Empower communities by emphasizing collaboration. Participatory design (PD) seeks to democratize the design process and empower community members to affect change [74]. Community voices and problems can be heard, discussed, and integrated into new solutions that improve community members' experience. Canvas successfully engaged a large community in a design process, but did not fully meet the democratic ideals set forth by PD. The prevalence of superposting behavior aligns with prior research on face-to-face discussions [84, 89, 121], online discussions [66], and open collaboration systems [105]. It is also one of the three consistent findings identified by Forte and Lampe in their review of open collaboration systems [45], often referred to as the "*power law of participation*". We also observed members employ strategies that attempted to prioritize their voice over others. For example, posters described deliberately posting an idea again to get more votes and attention. The affordance of the discussion forum appeared to exacerbate these problems. For example, community members could comment on their own post to pin it to the top of the discussion forum. These gaming behaviors added noise to the design process, making it harder for community members to make sense of others' ideas and to identify high-value opportunities to make their own contributions.

These gaming behaviors also highlight a new concern about the concept of power in PD research. Where previous research has focused on institutional power and how stakeholders should be involved in decision-making [17, 18, 92], gaming behaviors show how participants can shift power to and away from other participants. This observation suggests that PD research should be similarly critical of the power relationships between participants as it has been with power relationships between participants and PD organizers.

Community votes might provide some indication of popularity, but it is not clear in this context whether this provides an effective signal of quality or priority. Our analysis finds instances where strategy and social influence may be confounding the idea that the best ideas get implemented. Although unequal participation is not new, the strong motivation that community members feel towards getting their ideas implemented may have motivated these ‘gaming’ behaviors. While voting was the main method for community members to support specific ideas, votes did not guarantee implementation. A lack of transparency associated with decision-making led to some frustration in the community. In other open collaboration systems, such as Wikipedia, researchers have described the demotivating effect of having a contribution reverted or edited [44, 57, 58], especially when the justification is missing or inadequate [44]. Reducing the presence of gaming behaviors and improving transparency around decisions would help organizers decide what design problems to prioritize and could help community members feel more empowered in this process.

Despite these barriers, community members appeared to remain highly motivated to participate. This finding departs from prior work that describes the difficulty of motivating and retaining members in collaborative innovation initiatives [47, 55, 63]. Although some members expressed frustration with the transparency, we observed many instances where people consistently advocated for ideas for over a year. Community members were quick to respond to each others’ posts, with 49.6% receiving a response within one day and over 90.1% receiving a comment within one week. To capitalize on this kind of sustained engagement, platforms could do more to support collective decision making through structured activities. As examples, summarizing and connecting design information across multiple posts into one community post could give posters shared ownership of an ‘idea.’ In that case, conflict between ideas and design criteria would need to be mediated, possibly through collaborative decision support. Previous research, like Deliberatorium, presents civic issues in structured format to make argumentative positions clearer [71]. ConsiderIt frames deliberation around the pros and cons of an issue and visually summarizes community members’ argumentative positions [77]. ConsensusUs highlights disagreements between an individuals and the community to encourage consensus [85]. These technologies demonstrate how design aspects could be highlighted while encouraging communities to reach consensus through reflecting on their individual and collective goals for the design.

6.2 Limitations

This study explores in-the-wild data from a real community. As a result, there are a number of unavoidable confounding aspects. Some ideas remained in the open-for-voting stage longer than others and could cycle in and out of the rejected status based on Canvas’s priorities and community support. Based on our data, it was not clear when or why Canvas moderators changed the status of ideas. The final status of any idea, except for those that were implemented, was not always known and could have changed after we analyzed the data. Discussions in the comments and internally at Canvas may have affected which ideas were implemented, but this data was not available for our analysis. Although case-control might have helped to pair records that were accepted compared to ones that were not by controlling for date and other variables, there are too many other confounding factors (e.g.: poster’s reputation, post quality, post length, post relevance, the quality of the discussion related to each post, etc) to attempt to establish causal relationships based on the idea status. For this reason, this paper does not analyze the factors that led ideas to be implemented or ignored. Canvas works hard to help community members design together. Despite challenges, this effort and expertise may lead to higher success rates in Canvas than would be observed in other less motivated collaborative design efforts that take place at scale. Our findings may not extend to highly unmotivated collaborative design settings.

7 ACKNOWLEDGEMENTS

We would like to thank Canvas and their community team for fostering a large and engaged design community online. We would also like to thank Brendan Jew for illustrating our process diagram in Figure 2 and Anvesh Mekala for collecting and cleaning the data from Canvas for our analysis.

8 CONCLUSION

Participatory design (PD) has the ability to bridge gaps between experience and expertise, and between problems and solutions, while surfacing the many diverse and inter-connected problems that communities face. In this paper, we collected and analyzed a subset of 1412 posts and 18,335 comments submitted by a community of end-users to a discussion forum for the education platform, Canvas. We observed that users were highly motivated and collectively offered unique and diverse assets (problems, solutions, workarounds, and relevant situations) to an evolving design space. We identified challenges associated with capturing and curating design assets, mitigating negative social influence, and more effectively engaging and empowering communities to participate in the design process. In spite of these challenges, we observed sustained, long-term community engagement in this design effort. By addressing these issues and providing additional scaffolding and coordination, our research community can create technology that better supports participatory design at scale while realizing the goals of community empowerment and authentic engagement in the entire design process.

REFERENCES

- [1] Martin A Andresen. 2009. Asynchronous discussion forums: success factors, outcomes, assessments, and limitations. *Journal of Educational Technology & Society* 12, 1 (2009), 249–257.
- [2] Cynthia J Atman, Robin S Adams, Monica E Cardella, Jennifer Turns, Susan Mosborg, and Jason Saleem. 2007. Engineering design processes: A comparison of students and expert practitioners. *Journal of engineering education* 96, 4 (2007), 359–379.
- [3] Cynthia J Atman, Monica E Cardella, Jennifer Turns, and Robin Adams. 2005. Comparing freshman and senior engineering design processes: an in-depth follow-up study. *Design studies* 26, 4 (2005), 325–357.
- [4] Brian P Bailey and Eric Horvitz. 2010. What’s your idea?: a case study of a grassroots innovation pipeline within a large software company. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2065–2074.
- [5] Liam Bannon, Jeffrey Bardzell, and Susanne Bødker. 2018. Introduction: Reimagining participatory design—Emerging voices. *ACM Transactions on Computer-Human Interaction (TOCHI)* 25, 1 (2018), 1–8.
- [6] Barry L Bayus. 2013. Crowdsourcing new product ideas over time: An analysis of the Dell IdeaStorm community. *Management science* 59, 1 (2013), 226–244.
- [7] Nicolas Bettenburg, Sascha Just, Adrian Schröter, Cathrin Weiss, Rahul Premraj, and Thomas Zimmermann. 2008. What makes a good bug report?. In *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of software engineering*. ACM, 308–318.
- [8] Mikhail Yuryevich Bilenko. 2006. *Learnable similarity functions and their application to record linkage and clustering*. Ph.D. Dissertation.
- [9] Melanie Birks, Ysanne Chapman, and Karen Francis. 2008. Memoing in qualitative research: Probing data and processes. *Journal of research in nursing* 13, 1 (2008), 68–75.
- [10] Michael Mose Biskjaer, Peter Dalsgaard, and Kim Halskov. 2014. A Constraint-based Understanding of Design Spaces. In *Proceedings of the 2014 Conference on Designing Interactive Systems (Vancouver, BC, Canada) (DIS ’14)*. ACM, New York, NY, USA, 453–462. <https://doi.org/10.1145/2598510.2598533>
- [11] Robin Blom, Serena Carpenter, Brian J Bowe, and Ryan Lange. 2014. Frequent contributors within US newspaper comment forums: An examination of their civility and information value. *American Behavioral Scientist* 58, 10 (2014), 1314–1328.
- [12] Jeanette Blomberg, Mark Burrell, and Greg Guest. 2009. An ethnographic approach to design. *Human-Computer Interaction* (2009), 71–94.
- [13] Andy Blunden. 2014. Vygotskian collaborative project of social transformation: History, politics, and practice in knowledge construction. In *Collaborative projects*. Brill, 217–238.

- [14] Svenja Boberg, Tim Schatto-Eckrodt, Lena Frischlich, and Thorsten Quandt. 2018. The moral gatekeeper? Moderation and deletion of user-generated content in a leading news forum. *Media and Communication* 6, 4 (2018), 58–69.
- [15] Susanne Bødker. 2006. When second wave HCI meets third wave challenges. In *Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles*. 1–8.
- [16] Tom Boellstorff, Bonnie Nardi, Celia Pearce, and Tina L Taylor. 2012. *Ethnography and virtual worlds: A handbook of method*. Princeton University Press.
- [17] Tone Bratteteig and Ina Wagner. 2012. Disentangling power and decision-making in participatory design. In *Proceedings of the 12th Participatory Design Conference: Research Papers-Volume 1*. 41–50.
- [18] Tone Bratteteig and Ina Wagner. 2016. Unpacking the notion of participation in participatory design. *Computer Supported Cooperative Work (CSCW)* 25, 6 (2016), 425–475.
- [19] Stuart K Card, Jock D Mackinlay, and George G Robertson. 1990. The design space of input devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 117–124.
- [20] Mark Carman, Mark Koerber, Jiuyong Li, Kim-Kwang Raymond Choo, and Helen Ashman. 2018. Manipulating visibility of political and apolitical threads on Reddit via score boosting. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. IEEE, 184–190.
- [21] Joel Chan, Joseph Chee Chang, Tom Hope, Dafna Shahaf, and Aniket Kittur. 2018. Solvent: A mixed initiative system for finding analogies between research papers. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–21.
- [22] Joel Chan, Steven Dang, and Steven P Dow. 2016. IdeaGens: enabling expert facilitation of crowd brainstorming. In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion*. ACM, 13–16.
- [23] Oscar Chaparro, Jing Lu, Fiorella Zampetti, Laura Moreno, Massimiliano Di Penta, Andrian Marcus, Gabriele Bavota, and Vincent Ng. 2017. Detecting missing information in bug descriptions. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*. ACM, 396–407.
- [24] Boreum Choi, Kira Alexander, Robert E. Kraut, and John M. Levine. 2010. Socialization Tactics in Wikipedia and Their Effects. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (Savannah, Georgia, USA) (CSCW '10)*. Association for Computing Machinery, New York, NY, USA, 107–116. <https://doi.org/10.1145/1718918.1718940>
- [25] Joohee Choi and Yla Tausczik. 2018. Will Too Many Editors Spoil The Tag? Conflicts and Alignment in Q&A Categorization. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–19.
- [26] Giovanni Luca Ciampaglia and Dario Taraborelli. 2015. MoodBar: Increasing new user retention in Wikipedia through lightweight socialization. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 734–742.
- [27] Stephen Coleman and John Götze. 2001. *Bowling together: Online public engagement in policy deliberation*. Hansard Society London.
- [28] Glynis Cousin. 2010. Positioning positionality: the reflexive turn. In *New approaches to qualitative research*. Routledge, 25–34.
- [29] Justin Cranshaw and Aniket Kittur. 2011. The polymath project: lessons from a successful online collaboration in mathematics. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 1865–1874.
- [30] David P Crismond and Robin S Adams. 2012. The informed design teaching and learning matrix. *Journal of Engineering Education* 101, 4 (2012), 738–797.
- [31] Lincoln Dahlberg. 2001. The Internet and democratic discourse: Exploring the prospects of online deliberative forums extending the public sphere. *Information, communication & society* 4, 4 (2001), 615–633.
- [32] Thomas Däubler, Kenneth Benoit, Slava Mikhaylov, and Michael Laver. 2012. Natural sentences as valid units for coded political texts. *British Journal of Political Science* 42, 4 (2012), 937–951.
- [33] Steven Davies and Marc Roper. 2014. What's in a bug report?. In *Proceedings of the 8th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*. ACM, 26.
- [34] Michel M Denuit and Mhamed Mesfioui. 2017. Bounds on Kendall's tau for zero-inflated continuous variables. *Statistics & Probability Letters* 126 (2017), 173–178.
- [35] Andy Dong, Kevin Davies, and David McInnes. 2005. Exploring the relationship between lexical behavior and concept formation in design conversations. In *ASME 2005 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers Digital Collection, 61–71.
- [36] Kees Dorst. 2011. The core of 'design thinking' and its application. *Design studies* 32, 6 (2011), 521–532.
- [37] Kees Dorst. 2015. *Frame innovation: Create new thinking by design*. MIT press.

- [38] Kees Dorst and Nigel Cross. 2001. Creativity in the design process: co-evolution of problem–solution. *Design studies* 22, 5 (2001), 425–437.
- [39] David Easley and Jon Kleinberg. 2010. Power laws and rich-get-richer phenomena. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press (2010).
- [40] Pelle Ehn. 1988. *Work-oriented design of computer artifacts*. Ph.D. Dissertation. Arbetslivscentrum.
- [41] Kim VL England. 1994. Getting personal: Reflexivity, positionality, and feminist research. *The professional geographer* 46, 1 (1994), 80–89.
- [42] Dima Epstein and Gilly Leshed. 2016. The Magic Sauce: Practices of Facilitation in Online Policy Deliberation. *Journal of Public Deliberation* 12, 1 (2016).
- [43] Jenny Fan and Amy X Zhang. 2020. Digital Juries: A Civics-Oriented Approach to Platform Governance. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [44] Rosta Farzan and Robert E Kraut. 2013. Wikipedia classroom experiment: bidirectional benefits of students’ engagement in online production communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 783–792.
- [45] Andrea Forte and Cliff Lampe. 2013. Defining, understanding, and supporting open collaboration: Lessons from the literature. *American Behavioral Scientist* 57, 5 (2013), 535–547.
- [46] Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, and Peter Dalsgaard. 2019. Mapping the landscape of creativity support tools in HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [47] Johann Füller, Katja Hutter, and Rita Faullant. 2011. Why co-creation experience matters? Creative experience and its impact on the quantity and quality of creative contributions. *R&D Management* 41, 3 (2011), 259–273.
- [48] Eric Gilbert. 2013. Widespread underprovision on Reddit. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 803–808.
- [49] Barney G Glaser. 1992. *Basics of grounded theory analysis*. Mill Valley.
- [50] Vinod Goel and Peter Piroli. 1992. The structure of design problem spaces. *Cognitive science* 16, 3 (1992), 395–429.
- [51] Scott A Golder and Bernardo A Huberman. [n.d.]. Usage Patterns Of Collaborative Tagging Systems. 32, 2 ([n. d.]). <https://doi.org/10.1177/0165551506062337>
- [52] Gabriela Goldschmidt. 1990. Linkography: assessing design productivity. In *Cybernetics and System’90, Proceedings of the Tenth European Meeting on Cybernetics and Systems Research*. World Scientific, 291–298.
- [53] Timothy Gowers and Michael Nielsen. 2009. Massively collaborative mathematics. *Nature* 461, 7266 (2009), 879.
- [54] Todd Graham and Scott Wright. 2014. Discursive equality and everyday talk online: The impact of “superparticipants”. *Journal of Computer-Mediated Communication* 19, 3 (2014), 625–642.
- [55] Charles R Greer and David Lei. 2012. Collaborative innovation with customers: A review of the literature and suggestions for future research. *International Journal of Management Reviews* 14, 1 (2012), 63–84.
- [56] James Grimmelman. 2015. The Virtues of Moderation. *Yale Journal of Law and Technology* 17 (2015). <https://heinonline.org/HOL/Page?handle=hein.journals/yjolt17&id=42&div=3&collection=journals>
- [57] Aaron Halfaker, R Stuart Geiger, Jonathan T Morgan, and John Riedl. 2013. The rise and decline of an open collaboration system: How Wikipedia’s reaction to popularity is causing its decline. *American Behavioral Scientist* 57, 5 (2013), 664–688.
- [58] Aaron Halfaker, Aniket Kittur, and John Riedl. 2011. Don’t bite the newbies: how reverts affect the quantity and quality of Wikipedia work. In *Proceedings of the 7th international symposium on wikis and open collaboration*. ACM, 163–172.
- [59] Michael Hammond. 2005. A review of recent papers on online discussion in teaching and learning in higher education. *Journal of Asynchronous Learning Networks* 9, 3 (2005), 9–23.
- [60] Bruce Hanington and Bella Martin. 2012. *Universal methods of design: 100 ways to research complex problems, develop innovative ideas, and design effective solutions*. Rockport Publishers.
- [61] Christina Harrington, Sheena Erete, and Anne Marie Piper. 2019. Deconstructing Community-Based Collaborative Design: Towards More Equitable Participatory Design Engagements. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–25.
- [62] Eric B Hekler, Jennifer C Taylor, Steven P Dow, Michèle Morris, Faren J Grant, Sayali S Phatak, Don Norman, Dana M Lewis, et al. 2019. Exploring, Defining, & Advancing Community-Driven Design for Social Impact. In *Companion Publication of the 2019 on Designing Interactive Systems Conference 2019 Companion*. ACM, 373–376.
- [63] Jan Hess, Sinja Offenberger, and Volkmar Pipek. 2008. Community driven development as participation?: involving user communities in a software design process. In *Proceedings of the Tenth Anniversary Conference on Participatory Design 2008*. Indiana University, 31–40.
- [64] Kevin Hoffman, David Zage, and Cristina Nita-Rotaru. 2009. A survey of attack and defense techniques for reputation systems. *ACM Computing Surveys (CSUR)* 42, 1 (2009), 1–31.

- [65] Gaoping Huang and Alexander J Quinn. 2017. BlueSky: crowd-powered uniform sampling of idea spaces. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. 119–130.
- [66] Jonathan Huang, Anirban Dasgupta, Arpita Ghosh, Jane Manning, and Marc Sanders. 2014. Superposter behavior in MOOC forums. In *Proceedings of the first ACM conference on Learning@ scale conference*. ACM, 117–126.
- [67] Inc. Instructure. 2019. *About Canvas | Edtech Learning Platform | Instructure*. <https://www.instructure.com/canvas/about?newhome=canvas>
- [68] Jiwan Jeong, Jeong-han Kang, and Sue Moon. 2020. Identifying and Quantifying Coordinated Manipulation of Upvotes and Downvotes in Naver News Comments. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 303–314.
- [69] Shagun Jhaver, Iris Birman, Eric Gilbert, and Amy Bruckman. 2019. Human-machine collaboration for content regulation: The case of reddit automoderator. *ACM Transactions on Computer-Human Interaction* 26, 5 (7 2019). <https://doi.org/10.1145/3338243>
- [70] Joy Kim, Sarah Sterman, Allegra Argent Beal Cohen, and Michael S Bernstein. 2017. Mechanical novel: Crowdsourcing complex work through reflection and revision. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 233–245.
- [71] Mark Klein. 2011. How to harvest collective wisdom on complex problems: An introduction to the mit deliberatorium. *Center for Collective Intelligence working paper* (2011).
- [72] Rafal Kocielnik, Lillian Xiao, Daniel Avrahami, and Gary Hsieh. 2018. Reflection companion: a conversational system for engaging users in reflection on physical activity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (2018), 1–26.
- [73] Yubo Kou and Colin M Gray. 2017. Supporting distributed critique through interpretation and sense-making in an online creative community. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–18.
- [74] Laura Kozak, Charlotte Falk, and C Jean Chisholm. 2020. Sticks, Ropes, Land: Confronting Colonial Practices in Public Space Design. In *Proceedings of the 16th Participatory Design Conference 2020-Participation (s) Otherwise-Volume 2*. 63–67.
- [75] Peter Krafft, Kaitlyn Zhou, Isabelle Edwards, Kate Starbird, and Emma S Spiro. 2017. Centralized, parallel, and distributed information processing during collective sensemaking. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2976–2987.
- [76] Beate Krause, Christoph Schmitz, Andreas Hotho, and Gerd Stumme. 2008. The anti-social tagger: detecting spam in social bookmarking systems. In *Proceedings of the 4th international workshop on Adversarial information retrieval on the web*. 61–68.
- [77] Travis Kriplean, Jonathan Morgan, Deen Freelon, Alan Borning, and Lance Bennett. 2012. Supporting reflective public thought with considerit. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. 265–274.
- [78] Klaus Krippendorff, Yann Mathet, Stéphane Bouvry, and Antoine Widlöcher. 2016. On the reliability of unitizing textual continua: Further developments. *Quality & Quantity* 50, 6 (2016), 2347–2364.
- [79] Karim R Lakhani, Anne-Laure Fayard, Natalia Levina, and Stephanie Healy Pokrywa. 2012. OpenIDEO. *Harvard Business School Technology & Operations Mgt. Unit Case* 612-066 (2012).
- [80] Karim R Lakhani and Eric Lonstein. 2011. InnoCentive. com (B). *Com (B)(August 17, 2011). Harvard Business School General Management Unit Case* 612-026 (2011).
- [81] Anne Lamb, Jascha Smilack, Andrew Ho, and Justin Reich. 2015. Addressing Common Analytic Challenges to Randomized Experiments in MOOCs: Attrition and Zero-Inflation. In *Proceedings of the Second (2015) ACM Conference on Learning @ Scale (Vancouver, BC, Canada) (L@S '15)*. ACM, New York, NY, USA, 21–30. <https://doi.org/10.1145/2724660.2724669>
- [82] Andrew S Lan, Jonathan C Spencer, Ziqi Chen, Christopher G Brinton, and Mung Chiang. 2018. Personalized thread recommendation for MOOC discussion forums. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 725–740.
- [83] J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics* (1977), 159–174.
- [84] Colleen M Lewis and Nirali Shah. 2015. How equity and inequity can emerge in pair programming. In *Proceedings of the eleventh annual International Conference on International Computing Education Research*. ACM, 41–50.
- [85] Weichen Liu, Sijia Xiao, Jacob T. Browne, Ming Yang, and Steven P. Dow. 2018. ConsensUs: Supporting Multi-Criteria Group Decisions by Visualizing Points of Disagreement. *Trans. Soc. Comput.* 1, 1, Article 4 (Jan. 2018), 26 pages. <https://doi.org/10.1145/3159649>
- [86] Kurt Luther and Amy Bruckman. 2008. Leadership in online creative collaboration. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*. 343–352.

- [87] Kurt Luther, Casey Fiesler, and Amy Bruckman. 2013. Redistributing Leadership in Online Creative Collaboration. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work* (San Antonio, Texas, USA) (CSCW '13). Association for Computing Machinery, New York, NY, USA, 1007–1022. <https://doi.org/10.1145/2441776.2441891>
- [88] Stephen MacNeil, Sarah Abdellahi, Mary Lou Maher, Jin Goog Kim, Mohammad Mahzoon, and Kazjon Grace. 2018. Designing with and for the Crowd: A Cognitive Study of Design Processes in NatureNet. In *International Conference on Design Computing and Cognition*. Springer, 55–73.
- [89] Stephen MacNeil, Kyle Kiefer, Brian Thompson, Dev Takle, and Celine Latulipe. 2019. IneqDetect: A Visual Analytics System to Detect Conversational Inequality and Support Reflection during Active Learning. In *Proceedings of the ACM Conference on Global Computing Education*. ACM, 85–91.
- [90] Stephen MacNeil, Johanna Okerlund, and Celine Latulipe. 2017. Dimensional reasoning and research design spaces. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. 367–379.
- [91] Mary Lou Maher, Josiah Poon, and Sylvie Boulanger. 1996. Formalising design exploration as co-evolution. In *Advances in formal design methods for CAD*. Springer, 3–30.
- [92] Henry Mainsah and Andrew Morrison. 2014. Participatory design through a cultural lens: insights from postcolonial theory. In *Proceedings of the 13th Participatory Design Conference: Short Papers, Industry Cases, Workshop Descriptions, Doctoral Consortium papers, and Keynote abstracts-Volume 2*. 83–86.
- [93] Juliane Manitz, Mark Hempelmann, Goeran Kauermann, Helmut Kuechenhoff, Shuai Shao, Cornelia Oberhauser, Nina Westerheide, and Manuel Wiesenfarth. 2013. R Package 'samplingbook'.
- [94] Kent Marett and Kshiti D Joshi. 2009. The decision to share information and rumors: Examining the role of motivation in an online discussion forum. *Communications of the Association for Information Systems* 24, 1 (2009), 4.
- [95] Jane Margolis. 2010. *Stuck in the shallow end: Education, race, and computing*. MIT press.
- [96] Yann Mathet, Antoine Widlöcher, and Jean-Philippe Métivier. 2015. The unified and holistic method gamma (γ) for inter-annotator agreement measure and alignment. *Computational Linguistics* 41, 3 (2015), 437–479.
- [97] Alison Mathie and Gord Cunningham. 2003. From clients to citizens: Asset-based community development as a strategy for community-driven development. *Development in practice* 13, 5 (2003), 474–486.
- [98] Alice McIntyre. 2007. *Participatory action research*. Vol. 52. Sage Publications.
- [99] John McKnight and John Kretzmann. 1993. Building communities from the inside out. *A path toward finding and mobilizing a community's assets* (1993).
- [100] Axel Menning, Andrea Scheer, Benjamin Heinz Meier, and Claudia Nicolai. 2015. Designing as weaving topics: coding topic threads in design conversations. *Proceedings of the international association of societies of design research. IASDR, Brisbane* (2015).
- [101] Michael J Muller, Kate Ehrlich, and Stephen Farrell. 2006. Social tagging and self-tagging for impression management. *Technical Report No. 06-02* (2006).
- [102] Aileen Oeberst, Iassen Halatchliyski, Joachim Kimmerle, and Ulrike Cress. 2014. Knowledge construction in Wikipedia: A systemic-constructivist analysis. *Journal of the Learning Sciences* 23, 2 (2014), 149–176.
- [103] Daniel FO Onah, Jane E Sinclair, and Russell Boyatt. 2014. Exploring the use of MOOC discussion forums. In *Proceedings of London International Conference on Education*. 1–4.
- [104] Anne-Marie Oostveen and Peter van den Besselaar. 2004. From Small Scale to Large Scale User Participation: A Case Study of Participatory Design in e-Government Systems. In *Proceedings of the Eighth Conference on Participatory Design: Artful Integration: Interweaving Media, Materials and Practices - Volume 1* (Toronto, Ontario, Canada) (PDC 04). Association for Computing Machinery, New York, NY, USA, 173–182. <https://doi.org/10.1145/1011870.1011891>
- [105] Felipe Ortega, Jesus M Gonzalez-Barahona, and Gregorio Robles. 2008. On the inequality of contributions to Wikipedia. In *Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS 2008)*. IEEE, 304–304.
- [106] Vineet Pandey, Justine Debelius, Embriette R Hyde, Tomasz Kosciolk, Rob Knight, and Scott Klemmer. 2018. Docent: transforming personal intuitions to scientific hypotheses through content learning and process training. In *Proceedings of the Fifth Annual ACM Conference on Learning at Scale*. 1–10.
- [107] Bec Paton and Kees Dorst. 2011. Briefing and reframing: A situated practice. *Design Studies* 32, 6 (2011), 573–587.
- [108] Mercedes Paulini, Paul Murty, and Mary Lou Maher. 2011. Understanding collective design communication in open innovation communities. *Journal of CoCreation in Design and Arts* (2011).
- [109] Ronald Silva Pimentel. 2009. Kendall's Tau and Spearman's Rho for Zero-Inflated Data. (2009).
- [110] Sheizaf Rafaeli and Yaron Ariel. 2008. Online motivational factors: Incentives for participation and contribution in Wikipedia. *Psychological aspects of cyberspace: Theory, research, applications* 2, 08 (2008), 243–267.
- [111] Jennifer A Rode. 2011. Reflexivity in digital anthropology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 123–132.
- [112] Lars Kristian Roland, Terje Aksel Sanner, Johan Ivar Sæbø, and Eric Monteiro. 2017. P for Platform. Architectures of large-scale participatory design. (2017).
- [113] Johnny Saldaña. 2015. *The coding manual for qualitative researchers*. Sage.

- [114] Elizabeth B-N Sanders. 2002. From user-centered to participatory design approaches. In *Design and the social sciences*. CRC Press, 18–25.
- [115] Elizabeth B-N Sanders and Pieter Jan Stappers. 2008. Co-creation and the new landscapes of design. *Co-design* 4, 1 (2008), 5–18.
- [116] Aleksandra Sarcevic, Leysia Palen, Joanne White, Kate Starbird, Mossaab Bagdouri, and Kenneth Anderson. 2012. “Beacons of Hope” in Decentralized Coordination: Learning from on-the-Ground Medical Twitterers during the 2010 Haiti Earthquake. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (Seattle, Washington, USA) (CSCW ’12). Association for Computing Machinery, New York, NY, USA, 47–56. <https://doi.org/10.1145/2145204.2145217>
- [117] Donald A Schön. 1984. *The reflective practitioner: How professionals think in action*. Routledge.
- [118] Douglas Schuler and Aki Namioka. 1993. *Participatory design: Principles and practices*. CRC Press.
- [119] Ulrike Schultz and Betty Vandenbosch. 1998. Information Overload in a Groupware Environment: Now you see it, now you don’t. *Journal of Organizational Computing and Electronic Commerce* 8, 2 (1998), 127–148.
- [120] Pao Siangliulue, Joel Chan, Steven P Dow, and Krzysztof Z Gajos. 2016. IdeaHound: improving large-scale collaborative ideation with crowd-powered real-time semantic modeling. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 609–624.
- [121] Adam Stankiewicz and Chinmay Kulkarni. 2016. \$1 Conversational Turn Detector: Measuring How Video Conversations Affect Student Learning in Online Classes. In *Proceedings of the Third (2016) ACM Conference on Learning@ Scale*. ACM, 81–88.
- [122] Anna Stetsenko and Igor Arieivitch. 1997. Constructing and deconstructing the self: Comparing post-Vygotskian and discourse-based versions of social constructivism. *Mind, Culture, and Activity* 4, 3 (1997), 159–172.
- [123] Osamuyimen Stewart, David Lubensky, and Juan M Huerta. 2010. Crowdsourcing participation inequality: a SCOUT model for the enterprise domain. In *Proceedings of the ACM SIGKDD Workshop on Human Computation*. 30–33.
- [124] Jennifer Stromer-Galley. 2007. Measuring deliberation’s content: A coding scheme. *Journal of public deliberation* 3, 1 (2007), 12.
- [125] Vanessa Svihla and Richard Reeve. 2016. Facilitating problem framing in project-based learning. *Interdisciplinary Journal of Problem-Based Learning* 10, 2 (2016).
- [126] Yla Tausczik, Rosta Farzan, John Levine, and Robert Kraut. 2018. Effects of collective socialization on newcomers’ response to feedback in online communities. *ACM Transactions on Social Computing* 1, 2 (2018), 8.
- [127] Tiffany Tseng. 2016. Build in progress: Building process-oriented documentation. *Makeology: Makerspaces as learning environments*. New York: Routledge (2016).
- [128] Alexander Van Deursen and Jan Van Dijk. 2011. Internet skills and the digital divide. *New media & society* 13, 6 (2011), 893–911.
- [129] Vasilis Verroios and Michael S Bernstein. 2014. Context trees: Crowdsourcing global understanding from local views. In *Second AAAI Conference on Human Computation and Crowdsourcing*.
- [130] Marisol Wong-Villacres, Carl DiSalvo, Neha Kumar, and Betsy DiSalvo. 2020. Culture in Action: Unpacking Capacities to Inform Assets-Based Design. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [131] Rodrigo Zamith and Seth C Lewis. 2014. From public spaces to public sphere: Rethinking systems for reader comments on online news sites. *Digital journalism* 2, 4 (2014), 558–574.
- [132] Amy X Zhang and Justin Cranshaw. 2018. Making sense of group chat through collaborative tagging and summarization. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–27.
- [133] Amy X Zhang, Lea Verou, and David Karger. 2017. Wikum: Bridging discussion forums and wikis using recursive summarization. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 2082–2096.