

Ensemble: Exploring Complementary Strengths of Leaders and Crowds in Creative Collaboration

Joy Kim, Justin Cheng, Michael S. Bernstein

Stanford University HCI Group
{jojo0808, jcccf, msb}@cs.stanford.edu

ABSTRACT

In story writing, the diverse perspectives of the crowd could support an author's search for the perfect character, setting, or plot. However, structuring crowd collaboration is challenging. Too little structure leads to unfocused, sprawling narratives, and too much structure stifles creativity. Motivated by the idea that individual creative leaders and the crowd have complementary creative strengths, we present an approach where a leader directs the high-level vision for a story and articulates creative constraints for the crowd. This approach is embodied in Ensemble, a novel collaborative story-writing platform. In a month-long short story competition, over one hundred volunteer users on the web started over fifty short stories using Ensemble. Leaders used the platform to direct collaborator work by establishing creative goals, and collaborators contributed meaningful, high-level ideas to stories through specific suggestions. This work suggests that asymmetric creative contributions may support a broad new class of creative collaborations.

Author Keywords

Social computing; online creative collaboration; storytelling.

ACM Classification Keywords

H.5.3 Group and Organization Interfaces: Collaborative computing

INTRODUCTION

Peer production has proven itself a powerful lever for coordinating *knowledge* [3]; now, crowds are demonstrating their ability to empower *creativity*. Creative collaboration has drawn together diverse viewpoints [22] to support goals such as songwriting [5], animation [17] and remixing [10]. Storytelling, too, can be collaborative: if the author's role is to provide an authentic emotional experience for their story's audience [19], then the diverse perspectives and expertise of the crowd could support an author's search for the perfect character, setting, or plot.

Though the Web has produced the great collaborative encyclopedia, it has not yet been able to produce a collaborative novel of similar scale and impact. Peer

production often splits collaborative work into independent, discrete parts, but stories are not so easily untangled. Unstructured attempts at collaborative creative writing, such as the experimental wkinovel A Million Penguins [18], resulted in rampant vandalism and uncertainty about the direction the story should pursue. On the other hand, highly structured approaches, such as sentence-level round-robin writing in FoldingStory [29], constrained participant contributions and resulted in patchwork, incoherent stories. In other words, while constraints may allow a story to follow a consistent creative trajectory, too many restrictions may prevent collaboration at the level of an overall plot.

In this paper, we explore the potential of *structuring collaborative roles to reflect the complementary creative strengths of the crowd and the individual*. Many successful storytelling collaborations already separate creative leadership from more general participation [15]. In television screenwriting, for example, collaborating writers do not create a script in piecemeal and glue together the pieces at the end; instead, a leader and supporting writers each take on complementary creative powers and responsibilities [30]. By adapting these strategies for the Web, we may enable the creation of stories that draw on the unique experiences of hundreds of people.

Based on these patterns, we propose a model for collaborative storytelling where a leader maintains the high-level vision and articulates creative constraints for the crowd, while the crowd generates text and contributions within those constraints. We embody this idea in a collaborative writing system called *Ensemble*¹ (Figure 1).

Lone writers view the writing process as solving a series of rhetorical problems [8]; similarly, in Ensemble, a lead author can use *prompts* to guide overall creative direction and focus collaborative work around provided goals. For example, a leader may create a prompt that asks contributors to write a short amount of text that describes how two characters meet. Collaborators can write text for any prompt, as well as vote or comment on text contributed by others. By explicitly assigning the responsibility of maintaining creative direction to a leader, collaborators can concentrate on creating contributions beneficial to the creative work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CSCW'14, February 15–19, 2014, Baltimore, Maryland, USA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2540-0/14/02...\$15.00.
<http://dx.doi.org/10.1145/2531602.2531638>

¹ <http://ensemble.stanford.edu>

a examines the 300th bicycle. Vintage goose neck. Banana seat. Pink. The memory always finds its way into his hands whether he chooses to remember or not.

Left pedal, right pedal, chain, front tire, back tire. Nox reviews the order in his mind before he begins dismantling.

It's the only way to set the world straight.

Nox spins the pedal six times to the left when he notices a shadow across his floor. "Felicity?" he murmurs as his hand slips the pedal into its seventh deadly rotation. He freezes, glances up. No, not Felicity. A girl. Six, seven years old. Plaid pajamas.

"Get out!" he bellows. He can still fix it. He can still go round six times. She steps back. but stavs.

The screenshot shows a digital writing environment. At the top left, it says "By Andrea Ellickson" with a heart icon and "1 Love". To the right, there are draft counts: "#389" and "#187" with a plus sign between them. A circular tab labeled "b" is highlighted. The main text area contains the story draft from the left page. Below the draft is a prompt: "WRITE THIS: How should Nox meet the little girl? How does the battle over the bike ensue? How can we portray the his trauma without directly showing a flashback?". To the right of the prompt is a circular tab labeled "c". Below the prompt is a comment section with "2 Comment". A comment by "Taylor Ross" is visible, dated "30 days ago". The comment text is partially visible: "Andrea- I like how his thoughts are choppy, urgent. My favorite lines are 'his hand slips the pedal into its seventh deadly rotation' and 'She tugs with the ferocity of a seven-year-old'. The world needs him. He tugs. She flies forward, head t... handlebars'. I like the juxtaposition here. I see her, feel her...". A circular tab labeled "d" is positioned over the end of the comment.

Figure 1. A story scene showing (a) the winning draft for a scene. (b) Tabs can be used to switch to different drafts, with the winning draft being shown by default. (c) The scene prompt helps focus contributors on specific writing goals. (d) Contributors can make comments discussing the scene at a high level.

In an initial small-scale public deployment via a month-long short story competition, over one hundred volunteer writers started over fifty short stories using Ensemble. Through analysis of user activity and interviews with competition participants, we observed that both leaders and collaborators used leader-created constraints to coordinate work and explore the space of ideas for a story, with leaders even using rejected collaborator contributions to look at the story from different perspectives. Furthermore, leaders and collaborators had complementary motivations for collaboration; while leaders sought feedback, collaborators viewed their expertise as valuable input useful for leaders. Ensemble focuses specifically on writing short stories, but by framing writing as a form of problem solving, its model may apply to many modes of creative collaboration including journalism, film, and animation. In this paper, we focus on observing in a realistic context *how* leaders and collaborators work together to complete stories; the question of whether this approach results in *better* stories is left as future work.

RELATED WORK

This paper distinguishes between *creative collaboration* and other instances of peer production. Here, creative collaboration aims to produce pieces of art such as film, music, and stories that may not have predetermined goals. A Wikipedia page, for example, has implicit but commonly understood goals that guide contributor work: contributors understand that a Wikipedia page about apples should present information about apples in a holistic way expected from an encyclopedia. However, the goals of story writing are less clear because creators must come up with their own constraints.

Writing as a cognitive process

Most previous writing process research is done in the context of pedagogy, but understanding the writing process

also has implications for the design of writing tools. Early research challenged the traditional model of teaching writing as linear, irreversible stages that moved from prewriting to writing to revision [23]. However, more recent models suggest that skilled writers create and continually revise goals to develop purpose for the work overall [8]. Ensemble adopts this more recent perspective as a theoretical base for its design, framing collaborative work around rhetorical problems defined by a leader and making it possible for writing work to progress in a non-linear, evolutionary fashion.

Collaborative creativity

Effective strategies for creative collaboration exist in many offline contexts such as television screenwriting or execution of stage plays. A primary inspiration for Ensemble is the storytelling system seen in role-playing games [31]. In many role-playing games, a single player acts as the narrator while the remaining players take on the roles of characters interacting with the world laid out by the narrator. Despite the presence of a power imbalance where the narrator ultimately decides what happens, all the players contribute ideas to the progression of the story by working within constraints agreed upon by the group (i.e., what their characters would do in the situations presented by the narrator). Luther *et al.* [16] examined how this strategy of separating a leader's role from general collaborative work is used in online animation collaborations and observed a similar distribution of power. In Ensemble, we generalize this idea of working together within agreed-upon creative constraints and apply it as a general strategy for collaborative writing.

There have been many successes in mobilizing strangers to complete complex creative projects in non-writing domains. Aaron Koblin created a number of projects where the crowd contributes to interactive artwork using a predefined

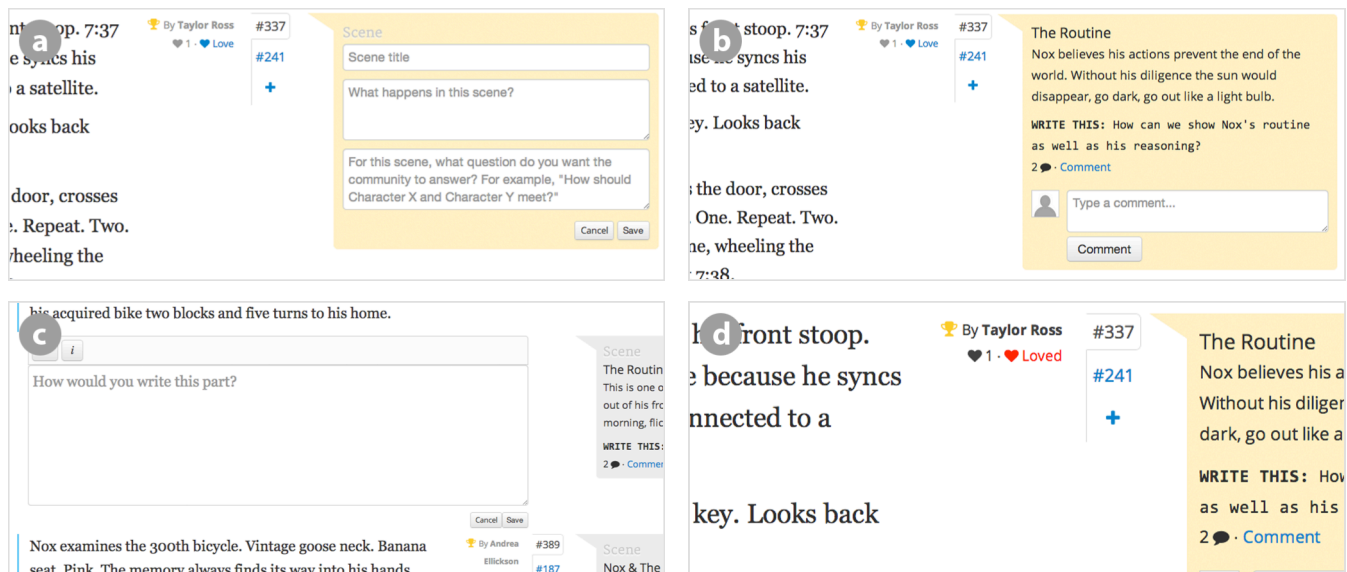


Figure 2. The collaboration process between lead authors and contributors on Ensemble. (a) The author creates a scene with a prompt to ask for specific help on a section of the story. (b) Contributors can make comments discussing the scene and possible solutions for the prompt. (c) Contributors can also submit alternative versions of the scene as solutions for the prompt. (d) Contributors can vote for alternatives they like. The author can declare a winning solution.

structure. In the Johnny Cash Project [32], for example, contributors redraw frames from a provided music video which are compiled into a redrawn version of the video. In this case, Koblin determined the project’s goals and vision before opening the project up to others; participants were not involved in the formation of the idea behind the project itself but needed only to execute the piece of work they volunteered for. While the visible uniqueness of frames in a crowdsourced video may be desired in the Johnny Cash project, a story needs a cohesive narrative arc to be understood. By structuring contributor work around high-level goals to achieve rather than low-level instructions to follow, Ensemble allows contributors the space to contribute in a way that complements work done by others.

Collaborative writing on the web

Communities surrounding the creation of fan fiction and forum-based role-playing games which attest to online writers’ desires to build off of existing work, share their own work, and seek the help of others for ideas and feedback. However, existing collaborative storytelling platforms online tend to be simply implementations of round-robin storytelling games [29,33]. Ensemble, in contrast, attempts to support a larger part of the story writing process, including idea generation, revision, and discussion with others.

Our work builds on previous tools for collaborative writing online [7,21] by extending previous work’s findings concerning task-based writing to creative writing. People tend to divide collaborative writing work by identifying sections of text that are independent from each other, and then working in parallel on a single document or writing in turn [12]. This behavior often occurs through document

exchange over email or through a shared workspace such as Google Docs or a wiki. However, when collaborators all manipulate a common product, the potential for conflict and need for coordination grow [13]. In contrast to non-fiction collaborative products such as Wikipedia articles or academic papers [26], a story is not a modular presentation of ideas but a multi-layered work consisting of interdependent characters, plot elements, and settings; it is extremely difficult to divide the work of story writing into independent, parallel pieces. Furthermore, creative works measure success in terms of emotional engagement of readers and viewers [19] rather than through objective metrics people can individually assess and agree on (such as whether a piece of open-source software “works”). To manage coordination and monitor progress towards creative goals, Ensemble uses an explicitly appointed creative leader.

Hypermedia approaches, where stories are built up using links or through branching paths rather than linearly, can allow groups to exchange new ideas; in the case of groups of young children [11,24], this strategy can assist in social and emotional development. Liu *et al.* [14] compared how elementary school children create a picture book with their peers when using a branching story editing system and when using a linear story editing system. Students who used the branching story editing system reported a higher sense of appreciation for their peers’ contributions and a higher sense of authorship over the stories they contributed to. However, while a branching editing system leaves previous content intact, creates a stronger sense of ownership over work, and prevents creative conflict, it leaves unclear how one would build a cohesive, finished story out of all the generated options. In contrast, Ensemble structures story

writing around goals from the beginning in order to direct writers towards creating a final story.

Recently, the blogging platform Medium added tools that allow readers to propose changes to an author's post [34]. These new tools make it easy for others to propose changes or contributions to the original blog post written by an individual. Unlike in Ensemble, this platform does not allow authors to explicitly direct collaborator work; however, it indicates that enabling an online audience to connect with creators over ideas is viewed as a potentially valuable addition to the creative process.

For the exploratory study presented in this paper, we echo past evaluations of collaborative writing [1] by first observing the collaborative behavior of small non-collocated teams rather than large crowds. However, we focus on studying whether we can enable a certain outcome given a new collaborative *approach* rather than new technology designed to support existing behavior.

ENSEMBLE

Ensemble supports creative writing collaboration by empowering a *lead author* to guide overall creative direction, while recruiting the crowd to focus on specific writing tasks. In this section, we introduce Ensemble and its techniques for structuring asymmetric collaboration.

Scenario

Kim, a writer, has an idea for a short story. She drafts the introduction and the general framing text for the story, then posts the incomplete draft on Ensemble. Kim isn't sure how to manage some of the character encounters in the plot, so she creates *scenes* with prompts for incomplete sections, indicating that she wants feedback or even alternative rewrites to improve what she already has (Figure 2a). She posts a link to the incomplete story to her favorite writing community, hoping to attract a few contributors.

A while later, one of these prompts catches Laura's eye as she is skimming through Kim's story. She adds a quick comment with an idea for one of the scenes (Figure 2b). Miles, another Ensemble user, builds on Kim's idea by writing an alternate telling of that scene. He submits the draft text as an alternative for Kim's original (Figure 2c). Other visitors find Miles's contribution interesting, and it receives several upvotes (Figure 2d).

Having received email notifications about the new additions, Kim returns to Ensemble and looks at the newly-added contributions. She particularly likes Miles's, and marks it as a winner so that it appears by default for that scene in the story. She looks at the additions made to other parts of the story, marking some contributions as winners for a few of the scenes, and building on the ideas to write her own alternatives for others.

Contributions with constraints

The scenario above exemplifies Ensemble's approach to creative collaboration: rather than placing all participants

on the same plane of creative power, Ensemble structures online collaboration by putting a crowd of *contributors* and a single *lead author* in charge of different aspects of a story. Where the lead author is responsible for the overall creative direction, contributors focus work around leader-provided prompts. The aim is to keep the story coherent while enabling contributors to make meaningful creative contributions.

Lead authors frame contributions using scenes

Rather than allowing contributors to edit any part of the story, *scenes* direct contributors towards specific sections the lead author has chosen. In Ensemble, scenes are the basic collaborative unit of each story, reflecting the idea that a scene roughly corresponds to a turning point that reveals character development, new information, and a goal for the next scene [19]. Ensemble scenes consist of raw story text, as well as a short description and a short prompt written by the lead author. Scenes help the lead author coordinate with contributors on the intended vision for the story by providing both an outline of a story's overall frame and direction towards specific tasks to complete.

As Kim did in the scenario above, a lead author can convert any section of the story into a scene by writing a description or prompt (Figure 1c). *Prompts* allow the lead author to suggest problems that need to be solved within the bounds of the scene (e.g., "I need Character X to meet Character Y. How do they meet?"). A story may have as many or as little scenes as the lead author likes, and a scene can only have one prompt at a time. Prompts do not need to be created in any particular order — leaders may choose to create a prompt for the end of the story, then later create a separate prompt for the beginning of the story, and still later delete or modify already existing prompts. Lead authors have the flexibility to direct collaborator work to any part of the text according to the story's changing needs.

By localizing contributor participation, scenes allow contributors to focus on small sections of the story at a time [2]. In addition, descriptions of other scenes can help contributors better frame their contributions in the context of the overall story. Finally, scenes provide a quick way for contributors to identify the contribution opportunities a story contains.

Contributors participate via drafts, comments and votes

Ensemble allows for different levels of contributor involvement, an important consideration in designing successful peer production systems [3]. While contributors primarily participate in a story by writing alternative paragraphs, or *drafts* (Figure 1a), for a scene, they can also comment on scenes or vote on drafts that they like. Authors and contributors can use scene *comments* (Figure 1d) to discuss specific scenes, or discuss more general topics in a discussion forum specific to that story.

Each scene can thus have multiple drafts written by many contributors, each representing a different interpretation of

that scene (Figure 1b). In editing interfaces like wikis or Google Docs where the goal is a single final version of the text, conflicts can arise through others unilaterally editing that text, which discourages further contribution [4,9]. We minimize conflict in Ensemble by asking collaborators to submit draft alternatives instead of directly editing each other's work. Collaborators receive credit for the drafts they write, supporting ownership and potentially reducing social loafing [25].

The most highly-voted draft for each scene is shown by default as part of the story, with the other drafts accessible through the use of scene-specific tabs (Figure 1b). Although contributors can communicate their vision for a story by voting for drafts, the lead author ultimately maintains creative control by choosing a winning draft for each scene. These winning contributions become part of the story text, signaling that they are aligned with the lead author's overall vision. Alternatively, the lead author can simply take inspiration from the contributed drafts and write their own.

For stories with many contributors, lead authors can optionally appoint *moderators* who can edit scene descriptions, prompts, others' drafts, and also mark drafts as winners, supporting the distribution of leadership [16].

Through these feedback channels of contributed drafts, voting, and comments, a lead author can work with contributors to continuously refine their creative vision and revise a story's text. Ensemble leverages the crowd for creativity and diversity by constraining and specifying the nature of contributions, and maintains narrative flow by granting creative control to a single author.

Limitations

Ensemble was primarily designed for an unequal power distribution, which is advantageous when there is high variance in contributor quality or motivation; alternative approaches may better suit scenarios where equitable contribution is desirable. Furthermore, because a lead author dictates the creative direction of the story, especially prolific contributors may find it difficult to make suggestions or changes beyond the bounds of a scene.

EVALUATION

Ensemble embodies the design hypothesis that the crowd can support the creative process for an author. In this section, we report on an evaluation of Ensemble that explored how users adopted Ensemble's division of creative responsibility to create short stories.

Method

We evaluated Ensemble through a public short story writing competition where teams of users were invited to write stories using the platform. Because we wanted users to behave as naturally as possible, we framed our study around a real competition so that users would focus on writing a good story rather than on the system itself. Rather than doing a controlled comparative study framed around story quality, our aim was to determine how teams would adapt

to the asymmetric structure of Ensemble and whether Ensemble allowed teams to write coherent stories. This approach allowed us to trace the ramifications of Ensemble's design on collaborative writing patterns.

The writing prompt provided to users was the following:

There are over 15,000 bikes used by students, staff, and faculty to get around campus. Over 300 bikes are stolen each year. Where do they go?

Teams were allowed to interpret the prompt in broad and unexpected ways, and stories were limited to 1,000 words. The prizes for the competition were \$500 in Amazon.com gift cards to be split among the team and publication in a major arts review publication on campus. The editor-in-chief of the arts review publication selected the grand prize and runner-up winners.

We advertised the writing challenge to several writing communities on the web, including FanFiction.net, FictionPress, and Reddit's writing communities. We also solicited a few well-known writing blogs to advertise the competition. Lastly, on-campus student arts publications and creative writing classes promoted the competition.

During the competition, Ensemble logged all participant activity, including the creation of stories, scenes, and drafts. After the competition, we emailed a survey to all Ensemble users asking for their motivations for starting stories, contributing to stories, and editing work written by others. We also asked how their experiences writing collaborative stories on other platforms compared with their experience using Ensemble.

We conducted semi-structured interviews with the seven of the most active Ensemble users (five male, two female). These participants were between the ages of 19 and 32, and most were amateur writers; their jobs ranged from student to manufacturing manager to humor columnist. We focused the interviews on how they used and adapted Ensemble to their story writing process, whether or not they perceived personal benefits of the platform on themselves as a writer, power dynamics between the head and supporting author, attitudes toward the resulting stories they participated in, and whether or not they would use Ensemble to work with others on stories again. Two interviews were conducted via Skype voice call, and five interviews were conducted via chat.

We then consolidated interview and survey responses and analyzed them according to the following predetermined themes framed by our research questions: prior experiences with collaboration on story writing, contributor activity and motivation, lead author activity and motivation, the interaction between contributors and leaders, organization of team workflow, system design feedback, and participant perceptions of the resulting story and their experiences using Ensemble. The themes that emerged from responses allowed us to form the following initial inferences about the




Figure 3. Percentage of total activity across Ensemble, by user role. Users in different roles contributed to stories in different ways.

relationship between leaders and their collaborators in online creative collaborations.

Results

The stories submitted to the short story competition displayed collaborative effort while demonstrating both coherency and creativity. Figure 1 shows one such story — this particular story was one of the runner-up winners of the competition. The team behind this story made liberal use of prompts and comments to frame the type of feedback they gave on each other’s work.

During the competition, 106 total users started 55 Ensemble stories, and 20 of these stories were submitted as entries into the competition by a total of 52 users. Over the course of the competition, 1067 unique visitors made a total of 2081 visits to the site (12165 pageviews), with stories receiving a median of 37.5 pageviews.

Visitors came from over 45 different countries, with most visitors coming from the USA (1376 visits), India (285 visits), and Canada (123 visits). Users made, in aggregate, a median of 552 comments, drafts, edits, and likes per day, with the most activity happening the day of the competition deadline (10848 actions).

Equal work despite asymmetric roles

Among the 20 submitted stories, there was a median of two members per story team, with only two stories by single authors. Team sizes ranged from one to seven people. We counted the number of comments, drafts, and edits to drafts made by each team member involved in a story and used the Gini coefficient of inequality to quantify the distribution of work across these users. The Gini coefficient ranges from 0 (work was done equally by everyone involved) to 1 (a single contributor did all of the work). The mean Gini coefficient across the 20 stories submitted to the competition was 0.286 (SD = 0.194), meaning that despite the asymmetric division of creative *responsibility* in Ensemble, the spread of work *quantity* was fairly balanced. That is, both story leaders and contributors contributed

fairly equal amounts of work, but in different ways. We describe this division of labor next.

Lead authors maintained creative authority

Lead authors worked to maintain creative ownership over the story, often by managing privileges and permissions. In Ensemble, anyone can participate in a story as a *contributor* (only able to comment and write alternatives for scene prompts), and lead authors can designate some users as *moderators* (able to edit anybody’s work). Of the 34 stories that had multiple team members, 25 stories had teams with a leader and contributors only: that is, only the leader was able to directly edit the work done by others. The remaining nine stories had all-moderator teams, where all members had the power to edit work by anybody else. There were no teams with a mix of contributors and moderators, perhaps due to small team sizes.

50% of lead authors’ activity was spent revising drafts they wrote themselves (Figure 3). In practice, this means they manually integrated ideas from others’ contributions, rather than directly including contributions as written. Lead authors would interpret the contributed draft for the high-level idea being suggested by the contributor, then manually tweak their existing work while keeping the suggestion in mind:

Participant 3: for the one that another person submitted for the story i submitted, i didn't end up using [their contribution], but it helped me pinpoint things that i liked and didn't like in both my version and theirs.

Interestingly, moderators did not add many new drafts to the story even though they had the power to do so: instead, they directly edited drafts created by the lead author, with this making up 45% of their activity on Ensemble (Figure 3). Moderators reported making edits mainly for the purpose of fixing typos, grammatical errors, and adjusting sentence structure, engaging in behavior similar to wiki-gardening [28].

Though they tended to reserve creative control for themselves, lead authors went through pains to portray themselves as benevolent dictators. Instead of simply accepting or rejecting contributions, lead authors would use comments to express their rationale and give the contributor the chance to respond:

Participant 1: Actually she did make quite a few changes at one point, but at that point in the story I had already plotted it out... And I let her know in the comments that, 'Look, I like the suggestions, but at this point I kind of know what I'm doing with the prose, so I would rather get suggestions with the plot elements, character motivations, things like that.'

Prompts reflected the variety of issues leaders thought about during the high-level writing process. Two researchers independently coded all prompts according to

Prompt Type	Count	Percentage
Story problems	29	32%
Outline	25	28%
Writing problems	11	12%
Request for review	11	12%
Placeholder	7	8%
Call for ideas	6	7%
Commentary	1	1%

Table 1. The types of prompts created by lead authors.

the categories in Table 1, and disagreements were resolved through discussion. We ignored 41 prompts that either had not been changed from the default text or were spam.

We found that 32% of all prompts created presented story-related problems that needed to be solved (“Is this the first time the bully and the main character have met?” or “When do Dean and Melanie let the others know the truth?”). 12% of prompts asked about writing technique or stylistic decisions (“How thorough should the descriptions of the character’s research be?” or “Is the ending too ambiguous?”). A smaller percentage of prompts called for brainstorming of ideas, sometimes while giving broad constraints (“Some kind of snappy ending here.” or “Any suggestions?”), or asked readers to review the written text (“Is there a certain imagery invoked when reading?” or “Does it sound alright?”). 28% of prompts did not present questions at all, and instead acted as points in an outline representing the overall story structure (“In which our protagonist borrows something of ours.” or “Intro! The first clue is found, and our narrator is intrigued.”). 8% of prompts simply marked the existence of a scene, but did not contain meaningful information about the type of scene the author was envisioning (“Bike #1”). One prompt was not related to collaboration, but rather served as a comment on the author’s creative choice for a scene (“yay for jokes”). In other words, prompts reflected how story writing goals ranged from very specific to very broad at different points in the story at different times.

Contributors offered high-level ideas through low-level work Unsurprisingly, creating comments and drafts made up 60% of contributor activity on Ensemble (Figure 3). Limiting contributors to these actions resulted in collaborative work that was directed and detailed, facilitating discussion between leader authors and their collaborators. Furthermore, because new drafts never overwrote work written by others, contributors felt free to express ideas.

Contributor efforts were directed by scenes. Contributors used scenes provided by the leader to make comments and write story text. The number of comments per story ranged widely, with some stories getting as many as 33 comments.

Out of the 20 submitted stories, the median number of comments was 5.5 (Q1 = 1.75, Q3 = 14). There was no correlation between comment count and team size (Pearson’s $r(18) = 0.268$, $p = 0.2537$). Comment areas essentially served as localized discussion boards surrounding individual scenes, and served three main purposes in Ensemble stories. First, contributors would make comments directly answering the prompts posed by lead authors in scene descriptions. Second, contributors would make unsolicited comments concerning grammar, sentence structure, and other technical issues within the scene, as opposed to unsolicited ideas for plot and character. Lastly, in the case of moderators, comments served as a place to notify others about direct changes that had been made to drafts for that scene (e.g., “I just edited the opening paragraph”).

Because they created drafts in the context of scenes rather than the story as a whole, contributors felt that they were able to make focused and substantial contributions to the story. Drafts were not so small as to feel incremental, as with the sentence-sized contributions one would make in a round-robin game such as FoldingStory; neither were they so big as to be daunting, as with a single document shared by all team members.

Participant 3: it still was really rewarding... because the way the website's set up helps that whole first draft, peer critiques, second draft process go smoothly and in little, manageable sections.

Scene descriptions acted as planning notes that kept track of overall story composition so that contributors did not have to. By structuring the story into scenes, contributors were able to focus on local creative issues without worrying about how their work would affect the story as a whole.

Contributors used drafts to safely share ideas. Users wrote an average of 1.2 drafts per scene (SD = 0.48). Without the ability to edit work by others, contributors used drafts (Figure 4) for two separate purposes. First, contributors would copy and paste existing drafts, make desired changes, then save the revised text as a new draft. Second, contributors and leaders alike would use drafts as a form of version control, creating drafts as true alternatives for writing a particular scene. Users viewed these alternative drafts as a safe way to show rather than describe ideas without changing existing work:

Participant 6: It's harder to describe what you want to do than just do it. So you could do it and have them look at it and say, that's what I was intending, [instead of them saying] I think we should go here, and I think that's maybe what should happen, and then they start criticizing their picture of what it is rather than your picture, because it's impossible to talk around it.

In other words, drafts were used as a way to clearly communicate ideas among team members rather than as a way to merge work written by multiple people. This fork-

and-merge behavior was unexpected; future versions of Ensemble could help leaders explicitly give credit where due when merging ideas into the story's text.

Prior relationships affected creative process

An unexpected difference in collaboration workflows emerged depending on whether or not the lead author knew contributors previously. Through interviews, we asked participants whether they knew their teammates prior to collaborating with them. We later sent out a survey to gather this information from authors we did not interview. Out of the 20 submitted stories, seven stories were written by teams made up of friends, four stories were written by teams made up of strangers, and two stories were written by single authors. For the remaining seven submitted stories, we were unable to gather information about the relationships among team members from the authors involved. Though we cannot quantitatively test this difference through the current study, we present observed trends regarding this difference.

Friends preferred an equal division of power. Teams that knew each other seemed to prefer a balanced power dynamic. In these cases, the entire team was often made moderators so that they could freely edit text written by anyone else; team members sometimes saw restrictions on editing ability as frustrating. Scene descriptions acted as notes from an earlier story planning session, usually carried out by the team in real-time over phone, email, or in person. The lead author would then create an outline for the story using Ensemble scenes based on the planning session. Team members would assign themselves to scenes, after which the team would begin loops of feedback and iteration.

Though this type of dynamic was not the core design scenario for Ensemble, participants explained that the system facilitated collaboration in a way that wikis and Google docs could not; scene descriptions provided a structure that aided in completing the collaboration. As one participant said, *"It was kind of nice just to have an explicit place to work on collaborative stories. I wouldn't have gone and emailed [my partners] and said, 'Hey, want to start a Google Doc and write a collaborative story?'"* (Participant 6)

Lead authors preferred creative authority over strangers. In teams made up of strangers, lead authors often took on a much more central and active role during story writing. After outlining a new story through scene descriptions and skeleton paragraphs, the lead author would then leave the door open for input before returning to the story. Surprisingly, contributions made by strangers, in part because of how out of sync they sometimes were with the leader's vision, ended up being serendipitous and influential in guiding a story to its finished state. As an example, one lead author reported that a contributor wrote a draft about how the bikes in the story could talk — an idea he had not considered, but eventually chose to include as a central

element in the story. Leaders viewed limiting the direct editing ability of collaborators as desirable, and felt that alternative drafts were sufficient for enabling collaborators' creativity.

Quality of submitted stories

Though this study focused on the relationship between leaders and collaborators, we wanted to conduct a preliminary investigation into whether or not Ensemble's collaboration strategy demonstrated an effect on creative outcomes. Did Ensemble's leader-directed collaboration strategy at least not worsen creative outcomes?

Two researchers independently rated the 20 submitted stories based on creativity, flow, and technical merit based on 7-point Likert scales, where 1 was worst and 7 was best for each category. Highly divergent ratings were discussed and re-scored by the researchers. The final rating for each story was the sum of the points linearly converted to a 1 to 7 scale to represent overall quality. The measured weighted Cohen's Kappa (squared weights) for the ratings by the two raters was 0.8, indicating substantial agreement.

Though we would need a controlled experiment to determine whether these factors truly affect story quality, it seems that for this initial study there is no relationship between quality and how work was split (Gini coefficient) amongst team members (Pearson's $r(18) = 0.385$, $p = 0.093$) or the number of collaborators per story (Pearson's $r(18) = 0.069$, $p = 0.7715$).

From the perspective of the competition participants, the written stories seemed to be of adequate quality. All seven interviewees expressed satisfaction with the story they had participated in, with some of the interviewees further stating that they had read through a few stories by other teams and "liked what they saw".

Expectations for collaboration

Lead authors and contributors had complementary motivations for participating in Ensemble stories. When asked what they had expected from collaboration through Ensemble, lead authors overwhelmingly cited that they hoped for outside perspectives on their work. Notably,

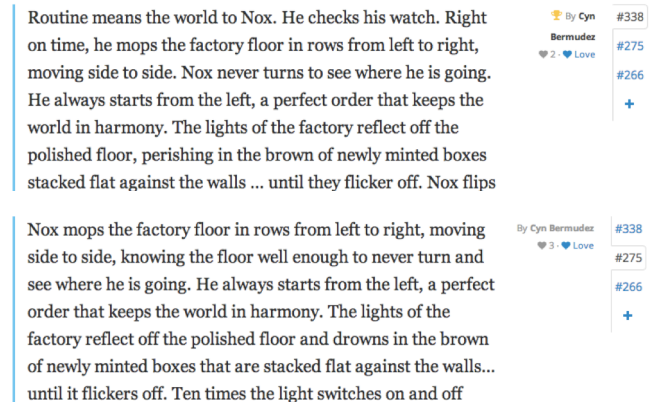


Figure 4. Two alternatives for the opening scene of a story.

contributors reported that they *wanted* to give input to others' work, describing the collaboration process as a way to hone their feedback skills as a method for improving their own writing. In turn, leaders felt more motivated to actively work on their stories because others were paying attention to their work.

Overall, the asymmetric structure of Ensemble did not just sync the complementary creative strengths of leaders and more general collaborators; it also revealed complementary motivations that served to spur work from both sides.

DISCUSSION

Through Ensemble, lead authors were able to successfully direct contributors' work through the use of scene descriptions and specific prompts, and contributors were able to offer high-level ideas by writing drafts and comments within provided constraints. Contributors did not feel creatively restricted, and the stories that were written were coherent and readable by non-writer audiences. In this section, we discuss what we might see in a large-scale deployment of Ensemble and propose a framework for collaboration by division of creative responsibility.

Scale

Many people participated in the story competition, but most stories had a small number of contributors. Likewise, in many crowdsourcing systems such as open-source software, a small number of dedicated contributors do most of the work [20]. What would happen if an Ensemble story had tens or hundreds of contributors?

As we saw in Ensemble, lead authors often take it upon themselves to manually integrate contributions done made by others when producing a collaborative work. However, coordinating creative work can overburden the leader [16], especially as the *number of contributors increases*. In Ensemble, collaborative writing was effective when the story leader was able to concentrate on their role of providing creative direction for the story. Introducing a leadership hierarchy could help lead authors integrate a large number of contributions.

In this study, teams wrote short stories. As *length of a creative work* increases, it may become harder for collaborators to contribute meaningfully. Changing details about a character or setting in one section, for example, may invalidate earlier or later parts of the story. As another example, contributors may need to take the time to become familiar with what has been written so far to even be in a position to add new content. However, this may reveal other opportunities to utilize the strength of the crowd; much like in Wikipedia [27], the crowd may be able to quickly and accurately identify inconsistencies in the story as they appear, or even manually maintain summaries of story sections through scene descriptions.

Designing power balance

Lead authors were bottlenecks for the writing process. While contributors were able to contribute to prompts that

already existed, they had to request additional scenes or structural changes through comments and wait for the leader to act on their requests.

What does this imply for whether Ensemble's design strikes the right balance between lead authors and contributors? Ensemble may be able to borrow from the theory of distributed leadership [16] and shift more organizational power to contributors. For example, Ensemble could structure collaboration similarly to open source software engineering, where contributors first identify their *own* problems to solve before submitting their work for acceptance into the main project. However, this may introduce more contention over the direction of the story and further increase what the lead author must moderate. Moreover, lead authors in our study valued having strong control over the story; giving too much power to contributors may cause lead authors to feel threatened.

Framework for dividing creative responsibility

In Ensemble, lead authors controlled creative direction and merged work, moderators made detailed edits to work already present, and contributors generated ideas by contributing paragraphs. Table 2 maps how the major creative behaviors seen in Ensemble are split among different roles in other online collaborative systems. In this first attempt at mapping the space for designing collaboration systems around creative roles, we can see that there are unexplored combinations that may generate further insight into where leaders, moderators, and crowds best fit in the creative process. One possible direction for further research is to examine the types of collaborative work these unexplored combinations best support. For example, a collaboration structure where the crowd is in charge of providing story goals and an individual is in charge of merging together input from the crowd through detailed writing and editing may result in different kinds of stories than the ones created through Ensemble.

This design space may be extended to include new roles and thus generate new ideas for collaborative processes. One might imagine adding rows for the support roles discussed earlier, including "filter out spam or off-topic

Task	Folding Story	A Million Penguins	Pipeline [16]	Ensemble
Providing goals	–	–	Leader/ Mods	Leader
Merging work	Crowd	Crowd	Leader/ Mods	Leader
Generating ideas	Crowd	Crowd	Crowd	Crowd
Detailed edits	–	Crowd	Crowd	Leader/ Mods
Doing work	Crowd	Crowd	Crowd	Crowd

Table 2. A selection of online creative collaboration systems and how each system splits creative work.

contributions”. A possible collaboration system, then, could center around a story fully generated by the crowd, with a leader simply pruning contributions as they come in to keep the story following a single creative direction. Another possible collaboration system could take advantage of the fact Ensemble enables alternative drafts for each scene in a story to create a story (or several stories) that branch in several possible directions. A “navigator” support role could help readers find story paths that make sense.

Future research could also examine how the motivations of users to participate in collaborative work changes as they fulfill different types of roles. What combinations of roles are most effective at encouraging people to take part in massively participatory stories? In Ensemble, leaders were in charge of providing creative direction, and were motivated to take part in collaboration because they felt perspectives from others would benefit the writing of a story they owned. If the leader was instead responsible for generating high-level ideas according to direction provided by the crowd, they might instead view collaboration as a helpful creative exercise. How might responsibility change users’ perception of the benefits and expectations of collaboration?

Limitations

The study in this paper was not controlled; we did not compare Ensemble and other collaborative story writing platforms through metrics such as writing quality and time to completion. While we were able to provide insight on the nature of asymmetric collaboration, we cannot claim anything about the quality the stories that emerged out of collaboration. We leave this as future work: some possibilities for measures of quality may be ratings by writing experts, comparative analysis of lexical quality of Ensemble stories, or measures of reader engagement.

Another limitation is that our user evaluation was framed around a competition. It is unclear to what extent the competition affected users’ motivations for using Ensemble. Participants described round-robin storytelling games such as FoldingStory as frameworks for supporting play, but tended to describe Ensemble as a platform with the potential to support production of serious stories or to support peer learning. It may be the case that the presence of a competition prompt and a deadline impacts users’ perceptions of whether a collaborative storytelling platform is for creative play or for creative work.

CONCLUSION & FUTURE WORK

In this paper, we explored the complementary creative strengths of the crowd and an individual by framing collaborative creativity as a process where collaborator activity can be guided by leader-generated constraints. Ensemble enabled small teams of both strangers and friends to collectively write short stories by channeling the complementary creative strengths of the crowd and the individual. In a creative writing competition, leaders directed collaborator work by establishing creative goals

while collaborators assisted leaders in exploring the space of ideas for a story. Collaborators did not feel creatively restricted by a lack of edit privileges and were able to contribute meaningful, high-level ideas to the story through detailed contributions. Furthermore, leaders and collaborators encouraged work from each other through complementary motivations for collaboration; leaders sought feedback from collaborators, and collaborators felt their input would be valuable to leaders.

In future work, we would like to see if Ensemble’s asymmetric strategy results in similar behavior at large scale. Much like Neil Gaiman mobilized his fans to help write *A Calendar of Tales* [35], we envision a system that enables authors to engage their tens of thousands of fans or fellow writers in a collective creative endeavor, using participation as the medium through which story, creator and audience experience each other.

We plan to explore this dynamic through a larger scale story called Arrowhead. Tom Kealey, a lecturer in the creative writing department at Stanford University, and Chris Baty, the founder of NaNoWriMo (National Novel Writing Month), will take turns using Ensemble to write chapters of the story, using scenes and prompts to gather contributions from the NaNoWriMo community after each chapter. This may reveal to what extent the crowd can help support authors who already have significant writing experience compared to the amateur writers that participated in our earlier evaluation; at the same time, working with such well-established authors may result in a story where the crowd becomes more of a vocal audience rather than a deeply involved group of participants.

We would also like to explore how Ensemble’s asymmetric creative strategy can be applied to mediums and contexts beyond creative writing. Could we, for example, crowdsource a painting that doesn’t look crowdsourced, or create collaboratively authored newspapers? Does Ensemble have the potential to enable new ways of teaching writing in K–12 schools? The Web is a rich space for sharing experiences through many media, not just through writing; through the right collaborative strategies, we can further empower its creative potential.

ACKNOWLEDGMENTS

We would like to thank the Stanford Arts Review and the Stanford Creative Writing Department for helping make the Ensemble Story Slam possible, as well as everyone who participated in the competition and interviews. We also thank Tom Kealey, Chris Baty, and the NaNoWriMo community. This project was supported by the Brown Institute Magic Grant.

REFERENCES

1. Baecker, R.M., Nastos, D., Posner, I.R., and Mawby, K.L. The user-centered iterative design of collaborative writing software. In *Proc. CHI* (1993).

2. Beenen, G., Ling, K., Wang, X., et al. Using social psychology to motivate contributions to online communities. In *Proc. CSCW* (2004).
3. Benkler, Y. *The wealth of networks: How social production transforms markets and freedom*. 2006.
4. Birnholtz, J. and Ibara, S. Tracking changes in collaborative writing: edits, visibility and group maintenance. In *Proc. CSCW* (2012).
5. Burke, M. and Settles, B. Plugged in to the community: Social motivators in online goal-setting groups. In *Proc. ICCT* (2011).
6. Cranshaw, J. and Kittur, A. The polymath project. In *Proc. CHI* (2011).
7. Dourish, P. and Bellotti, V. Awareness and coordination in shared workspaces. In *Proc. CSCW* (1992).
8. Flower, L. and Hayes, J.R. A cognitive process theory of writing. *College composition and communication* (1981).
9. Halfaker, A., Kittur, A., and Riedl, J. Don't bite the newbies: how reverts affect the quantity and quality of Wikipedia work. In *Proc. WikiSym* (2011).
10. Hill, B.M., Monroy-Hernández, A., and Olson, K. Responses to remixing on a social media sharing website. In *Proc. ICWSM*, (2010).
11. Hourcade, J.P., Bederson, B.B., Druin, A., and Taxén, G. KidPad: collaborative storytelling for children. In *Proc. CHI Extended Abstracts* (2002).
12. Kim, H.-C.E. and Eklundh, K.S. Reviewing practices in collaborative writing. In *Proc. CSCW* (2001).
13. Kittur, A., Suh, B., Pendleton, B.A., and Chi, E.H. He says, she says: conflict and coordination in Wikipedia. In *Proc. CHI* (2007).
14. Liu, C.-C., Liu, K.-P., Chen, W.-H., Lin, C.-P., and Chen, G.-D. Collaborative storytelling experiences in social media: Influence of peer-assistance mechanisms. *Computers & Education* (2011).
15. Luther, K. and Bruckman, A. Leadership in online creative collaboration. In *Proc. CSCW* (2008).
16. Luther, K., Fiesler, C., and Bruckman, A. Redistributing leadership in online creative collaboration. In *Proc. CSCW* (2013).
17. Luther, K., Ziegler, K., Caine, K.E., and Bruckman, A. Predicting successful completion of online collaborative animation projects. In *Proc. CC* (2009).
18. Mason, B. and Thomas, S. A million penguins research report. *Institute of Creative Technologies* (2008).
19. McKee, R. *Story: Style, Structure, Substance, and the Principles of Screenwriting*. 1997.
20. Mockus, A., Fielding, R.T., and Herbsleb, J.D. Two case studies of open source software development: Apache and Mozilla. *ACM TOSEM* (2002).
21. Noël, S. and Robert, J.-M. Empirical study on collaborative writing: What do co-authors do, use, and like? In *Proc. CSCW* (2004).
22. Paulus, P.B. and Nijstad, B.A. *Group creativity: Innovation through collaboration*. 2003.
23. Perl, S. The composing processes of unskilled college writers. *Research in the Teaching of English* (1979).
24. Ryokai, K. and Cassell, J. Computer support for children's collaborative fantasy play and storytelling. In *Proc. CSCL* (1999).
25. Thom-Santelli, J., Cosley, D.R., and Gay, G. What's mine is mine: territoriality in collaborative authoring. In *Proc. CHI* (2009).
26. Tomlinson, B., Ross, J., Andre, P., et al. Massively distributed authorship of academic papers. In *Proc. CHI Extended Abstracts* (2012).
27. Viégas, F.B., Wattenberg, M., and Dave, K. Studying cooperation and conflict between authors with history flow visualizations. In *Proc. CHI* (2004).
28. Welser, H.T., Cosley, D., Kossinets, G., et al. Finding social roles in Wikipedia. In *Proc. iConference* (2011).
29. FoldingStory. <http://foldingstory.com/>.
30. Who really runs things - Los Angeles Times. <http://articles.latimes.com/2007/nov/23/entertainment/et-channel23>.
31. Dungeons & Dragons. <http://www.wizards.com/dnd/>.
32. The Johnny Cash Project. <http://www.thejohnnycashproject.com/>.
33. The Collabowriters. <http://www.thecollabowriters.com/>.
34. Why Medium Notes Are Different and How to Use Them Well. <https://medium.com/about/5972c72b18f2>.
35. A Calendar of Tales. <http://keepmoving.blackberry.com/desktop/en/us/ambassador/neil-gaiman.html>