Supporting Collaborative Writing with Microtasks

Jaime Teevan
Microsoft Research
Redmond, WA USA
teevan@microsoft.com

Shamsi T. Iqbal
Microsoft Research
Redmond, WA USA
shamsi@microsoft.com

Curtis von Veh
Microsoft Research
Redmond, WA USA
curtisvv@microsoft.com

ABSTRACT
This paper presents the MicroWriter, a system that decomposes the task of writing into three types of microtasks to produce a single report: 1) generating ideas, 2) labeling ideas to organize them, and 3) writing paragraphs given a few related ideas. Because each microtask can be completed individually with limited awareness of what has been already done and what others are doing, this decomposition can change the experience of collaborative writing. Prior work has used microtasking to support collaborative writing with unaffiliated crowd workers. To instead study its impact on collaboration among writers with context and investment in the writing project, we asked six groups of co-workers (or 19 people in total) to use the MicroWriter in a synchronous, collocated setting to write a report about a shared work goal. Our observations suggest ways that recent advances in microtasking and crowd work can be used to support collaborative writing within preexisting groups.

Author Keywords
Collaborative writing; social computing; microtasking.

ACM Classification Keywords
H.5.3. Group and Org. Interfaces: Collaborative computing.

INTRODUCTION
The proverb, "A journey of a thousand miles begins with a single step," is attributed to the sixth century BCE philosopher Laozi. People have attempted to accomplish large personal tasks by decomposing them into manageable parts for millennia, and modern approaches like *Getting Things Done* [2] put Laozi's words into practice. Research shows that concrete plans with actionable steps enable people to complete their tasks better and faster [17], and that breaking a large macro-task down into a series of small, context-free microtasks leads to higher quality work, makes the task easier, and supports recovery from interruption [8].

While task decomposition previously had to be done by hand, it is now often possible to algorithmically break complex tasks all the way down into microtasks that can take as little as a few seconds each to complete. For example, on crowd platforms microtasks are now increasingly being composed to accomplish complex tasks that are not obviously achievable via standalone microtasks [18], such as taxonomy creation [9], itinerary planning [36], and writing [1, 3, 15, 16, 23, 24]. For crowd workers to complete large-scale collaborative tasks where each worker performs a small part of the task individually and without coordination with other workers, all of the context required to complete each microtask must be encompassed in the microtask itself.

Since microtasks are designed to minimize the need for coordination, it is likely that such task decomposition could also benefit known collaborators working towards a common goal where coordination can be an issue. Past research has identified challenges in collaborative writing, including disproportionate contributions, lack of a coherent structure, and difficulty understanding the current context of the paper [32]. In this paper we explore whether dividing writing into microtasks impacts these coordination issues. In particular we focus on understanding how writing via microtasks affects coordination when collaborators are collocated, and study whether microtask use eliminates or reduces common coordination activities for collocated collaborators.

As an example of a typical existing collaborative writing process, imagine a group writing a description of a shared project [25]. The group may meet to brainstorm ideas using the whiteboard and take notes for personal reference. These notes may later be manually organized in the form of an outline to provide a coherent structure. The outline may then be expanded collaboratively with the group gathered around a shared screen, or individually by dividing up the writing task and compiling the pieces into a coherent whole. Throughout the process the task requires tightly coupled coordination because the steps are interdependent and significant context is needed to perform each one. However, recent crowdsourcing research demonstrates that writing can be decomposed in such a way that the resulting microtasks require significantly less context to complete [1, 3, 16].

To study the impact of microtasking on collaborative writing we introduce the *MicroWriter*, a system that supports writing using microtasks. MicroWriter users perform three different types of microtasks to produce a single written report: 1) *Idea Generation*, where they enter ideas to include in the final report, 2) *Labeling*, where they provide meta-information about each idea that is used to create a coherent structure with the ideas, and 3) *Writing*, where they produce a paragraph of text in response to a small group of related ideas. At each stage in the process, individual users do not need to be aware of what has already been done or what others are doing; the microtasks themselves encompass the necessary context.

CHI '16, May 07-12, 2016, San Jose, CA, USA
 ISBN 978-1-4503-3362-7/16/05…$15.00
DOI: http://dx.doi.org/10.1145/2858036.2858108

The specific research question we address is: *How does writing via microtasks impact the writing experience for a group of collocated collaborators in an office setting?* To answer it, we asked 19 people in six groups of 2 to 5 co-workers to use the MicroWriter in a synchronous collocated setting to write a report related to a current project. Their experiences demonstrate some of the benefits and drawbacks of using microtasks to support existing collaborations. We find it encourages active participation from all group members, rapid generation of a wide variety of ideas, and the identification of emergent structure. However, it must also account for users' desire for ownership of contributions and does not remove the need for coordination in truly collaborative tasks such as discussion and brainstorming.

RELATED WORK

Previous research suggests that microtasking can be a useful way to complete a large task. Cheng et al. [8] found that breaking a large macro-task down into a series of small, context-free microtasks leads to higher quality work, makes the task easier, and supports recovery from interruption. Action plans that provide actionable steps are useful for getting hard, persistent tasks done, even when these action plans are externally created [17]. There is evidence that information workers already implicitly break larger tasks down; people perceive tasks in segments [35], and mental workload dips at task boundaries, rising during individual subtasks [34]. Additionally, common tasks such as email are accomplished in short bursts of less than five minutes [12].

Microtasking is prevalent in crowdsourcing, where a number of workflows have been developed that decompose large, seemingly complex tasks like taxonomy creation [9] and copyediting [3] into microtasks [18]. In crowdsourcing, microtasking allows crowd workers schedule flexibility [22]. The rising success of crowd work suggests that traditional information workers may stand to benefit from microwork structure [31], which can enable people to complete large tasks in many brief moments when they feel productive but do not have a long, uninterrupted period of time [6, 33].

While most crowd work consists of microtasks that require limited context, researchers have successfully identified ways to support limited generation and use of context in crowdsourcing [7, 19, 36]. Different from crowdsourcing, however, people using microtasks to complete their personal information work have significant pre-existing context, and we study that in detail with the MicroWriter.

Writing as a process resists being encapsulated in a single theory or model [13]. Of the models that have been proposed, one of the more widely accepted ones considers writing as a series of cognitive processes [11]. These processes are distinctive and hierarchical, but any process may be embedded in any other process, leading writers to construct highly personal, idiosyncratic, and complex workflows. Decomposing writing by different rhetorical categories has been shown to aid the process, especially with weak writers who spend the majority of their time and effort attending to surface-level, mechanical details rather than more complex processes of meaning-making [27]. Decomposing by a document's developmental stages, such as pre-writing, writing, and re-writing, is valuable because it scaffolds the work, prompts thoughtful engagement with tasks people would normally skip, and eases cognitive load [14].

Several existing systems have decomposed writing into a series of microtasks in the context of crowdsourcing. For example, Soylent [3] divides writing projects into stages and uses crowd workers to provide suggestions, shorten text, and proofread. CrowdForge [16] demonstrates that high quality text content can be written by multiple independent workers who each complete simple tasks such as preparing an outline, gathering facts, and writing simple prose. Ensemble [15] uses a team leader to direct writing projects. By taking advantage of the complementary writing skills of different crowd members, the authors find that writers are able to produce better, more creative content in less time. Similarly, WearWrite [23, 24] uses microtasks and the crowd to make it possible to write a paper from a watch. Little et al. [21] find that workers who perform writing tasks serially produce better content than workers who perform the tasks in parallel.

While the MicroWriter borrows insights from these systems to develop a microtasked approach to writing, the specific process it uses is incidental. Instead, the goal of this work is to explore how microtasking affects the user experience in collaborative writing, where writing microtasks are performed by groups of collocated people working collaboratively who are all invested in the final outcome, rather than crowd workers. There is little existing research looking at how microtasks can be used by non-crowd workers to support the completion of complex information tasks. Sadauskas et al. [29] explored how people can get started writing using existing micro-writing they have done on social media. However, the microwork is not done explicitly with the goal of creating a larger piece.

Collaborative writing is a common yet complex process that involves many discreet activities [28] and evolving roles [25]. Most existing collaborative writing relies on online synchronous collaborative authoring tools or the change tracking and version control features of modern word processors [26]. These tools explicitly facilitate coordination across authors by supporting awareness of what others are doing. For example, Tam and Greenberg [30] discuss the importance of workspace awareness drawing attention to certain changes via visual cues. Birnholtz et al. [4] argue that minimizing the visibility of some changes can facilitate social interaction. The MicroWriter explores taking this to an extreme by providing no direct awareness of other people's writing actions within the tool. In doing so, we learn how using the microtasking structure common to crowdsourcing impacts collaborative writing among work-group members.

THE MICROWRITER

The MicroWriter is implemented as a Universal Windows app, which allows it to run as a Modern Application on any

Windows device, including mobile phones, laptops, and desktops. The content generated by the MicroWriter is stored in the cloud using Azure, but all of the application logic takes place within the application itself. A new report is started by entering a title and keyword into the MicroWriter's start screen. Anyone who enters the corresponding keyword can participate in the authoring of that report, enabling the application to be used collaboratively. This collaborative writing can be done in a collocated or remote fashion, and the work can be done synchronously or asynchronously.

The MicroWriter presents the task of writing in terms of a series of smaller microtasks. Microtasks are clustered into three stages, shown in Figures 1, 2 and 3:

1. *Idea Generation*: Create the raw ideas for the report
2. *Labeling*: Cluster generated ideas into a structured outline
3. *Writing*: Given set of related ideas, produce a paragraph

These stages echo the stages used in other microtasked writing systems [1, 16]. Agapie et al. [1] create news articles for local events by using crowd workers to 'gather' content, 'curate' the content, and 'write' the article. While their stages are similar to ours, only the first is done via microtasks. Kittur et al. [16] use crowd workers to 'partition' the task of writing an encyclopedia entry into an outline, 'map' fact gathering tasks out, and 'reduce' the facts into paragraphs. While their process identifies structure in the first stage, in the MicroWriter structure emerges from the second.

A progress bar at the top of the screen shows where in the process the user is. Users can move between the stages by clicking the right arrow to get to a later stage and the left arrow to return to an earlier stage. Once the *Labeling* stage is complete and a user moves to the *Writing* stage, the outline structure for the report gets locked in and users can no longer return to the previous two stages. This is because the *Writing* stage depends on the structure that emerges during *Idea Generation* and *Labeling*; changes to the structure after writing begins invalidates the text that is already written.

Idea Generation

The purpose of *Idea Generation* (Figure 1) is to get users to input all of the content to be included in the report in the form of one or two sentences representing each idea. Short data gathering tasks have been shown to be useful at the start of the writing process: Agapie et al. [1] use local crowd workers gather ideas based on structured prompts; Kittur et al. [14] use crowd workers to collect facts for writing; Sadauskas et al. [26] gather tweets to help people get started writing.

MicroWriter users in the *Idea Generation* stage type ideas into a text box and press return to submit them. The entered ideas are then stored in a table in the cloud and reflected back to all session members to encourage users to build off of each other. Research suggests this is useful during brainstorming [10]. Ideas are not associated with a particular user when stored, and the same idea can be entered multiple times. When users can no longer think of ideas to enter or want to move to the next stage, they can move to the *Labeling* stage.

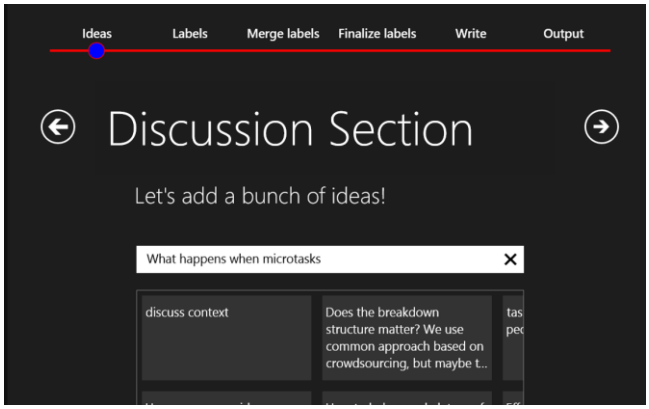

Figure 1. The MicroWriter in the *Idea Generation* stage.

Labeling

Once ideas have been entered, they need to be organized. Crowd-based microtask approaches to collaborative writing do this in a top-down fashion by either creating an initial outline (e.g., [16]) or relying on a leader with a global view (e.g., [15, 24]). However, individuals know the content they want to write even without knowing the structure, making it possible to allow structure to emerge from the content.

The process the MicroWriter uses to organize ideas is based on a modification of the Cascade taxonomy creation process [9] that uses label-based microtasks to organize information via the crowd. As such, even though the process produces a high-level organization of all of the ideas, each individual microtask is constrained in context to only a limited amount of presented information and can be done quickly and in isolation. *Labeling* is performed in three substages:

Initial Labeling

In the first substage (Figure 2a) users are presented with an idea and asked to enter one or more labels for the idea in a text box. These labels do not need to be consistent with existing labels. Labeling is constrained in context to the individual idea. During a first pass, all of the ideas are shown once in a random order, and the number of unlabeled ideas is displayed. When all ideas are labeled users may move on to the *Label Merging* substage. However, they may also choose to continue with the *Initial Labeling* and revisit ideas that have already been labeled. Users are shown all existing labels (if any) when they are asked to label an idea.

Label Merging

While labels are easy to produce, there can be significant variation in the labels used across people. For large group sizes, the number of labels can even equal or exceed the number ideas. For this reason, the second *Labeling* substage (Figure 2b) merges labels to produce a canonical label vocabulary. Some label merging can happen algorithmically. When a user first enters the *Label Merging* substage, an initial automatic merging is performed by normalizing for case. Stemming, stop word removal, whitespace removal, and synonym matching could also be used, but are not currently implemented. Singleton labels provide no organizational value, and thus are removed.

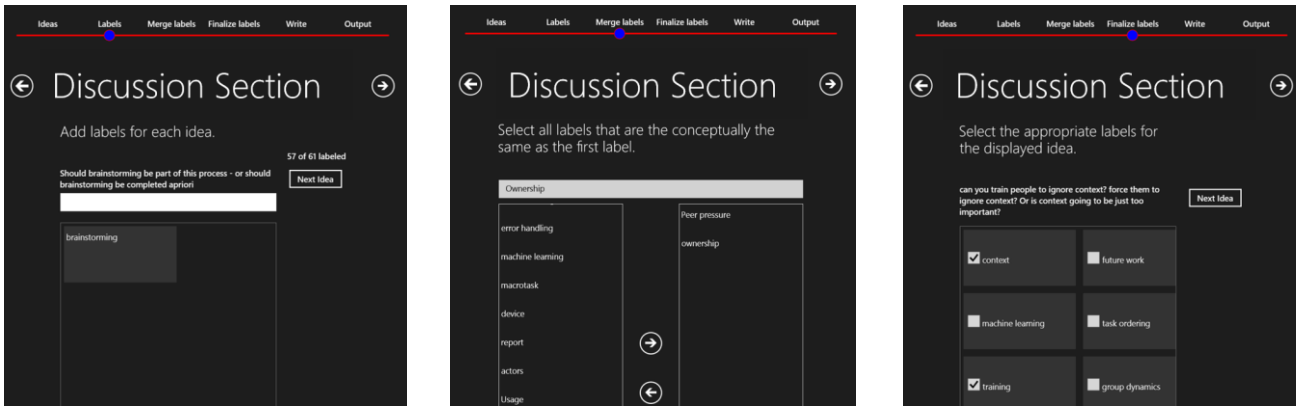

(a) (b) (c)

Figure 2. The MicroWriter in the *Labeling* stage, which consists of three substages: a) *Initial Labeling*, where labels are added to each idea, b) *Label Merge*, where a canonical label set is identified, and c) *Finalize Label*, where the canonical labels are applied.

The remaining labels are merged manually via a series of microtasks. The *Label Merging* microtask presents the user with a master label and a list of all other available labels for matching, sorted by potential relevance based on text matching and the structure of the underlying idea and label bipartite graph. The user identifies labels with similar meaning to the master label from the list by selecting them and moving them to a parallel list. The user can also optionally rename the merged label set by editing the master label. Upon completion of the microtask, the selected labels are merged and renamed, and the merged label is applied to all of the ideas associated with the set of selected labels.

Finalizing Labels

After a canonical set of labels has been identified, the system asks users to select the labels that apply to each idea (Figure 2c). Each *Finalizing Labels* microtask displays an idea and the list of canonical labels, sorted such that the labels that have been applied to it are shown first, with the remaining labels sorted using the graph structure to identify similar labels and popularity to identify likely labels. Label are accompanied by a check box with suggested labels checked. User are asked to modify this initial selection if necessary.

Writing

The final *Writing* stage in the MicroWriter involves iterating over small collections of ideas to turn them into coherent paragraphs. This is similar to what is done in CrowdForge, where workers are asked to create paragraphs by merging content together in a 'reduce' step [16], and by Little et al. [21], who show microtasks are useful for transforming rough text into better text. For each *Writing* microtask, a set of ideas related to a particular label is shown to the user (Figure 3), and the user is asked to turn the set of ideas into a coherent paragraph. They are not required to take into account any other content while writing. When the writing microtasks are complete, the paragraphs are ordered algorithmically with the goal of placing the most important paragraphs first.

Ideas are automatically grouped using the labels as a guide. A good paragraph is coherent and contains only a few related ideas. To identify a set of ideas that meets this criteria, the label with the fewest ideas associated with it is selected as the paragraph topic. The ideas associated with the label are shown to the user, along with the label, and the user is asked to write a paragraph. Since ideas can have multiple labels, any other labels associated with each idea are also displayed, and the user may choose to create subgroups within a paragraph using these auxiliary labels as they write. Once the paragraph is written, the ideas included in that paragraph are removed from the general pool to avoid redundancy in the final report. The label with the fewest number of remaing ideas is then considered, and the process continues iteratively until all of the ideas have been written about.

The final report is created by concatenating the individually written paragraphs into a report, as is done with CrowdForge [16]. Paragraphs are ordered so that those associated with the most popular labels appear first, on the assuption that these labels are particularly important and address key concepts for the report. Although popular labels are associated with many ideas, the corresponding paragraphs typically contained only a few ideas because many of the related ideas were written about in association with other labels. These leading paragraphs only contained the ideas that remained associated with the popular label after the other ideas were removed.

Figure 3. The MicroWriter during the *Writing* stage.

Group		Report	Number of			
ID	Size	Topic	Ideas	Labels	Paragr	Words
A	2	Brainstorm	30	28	12	752
B	5	Brainstorm	45	13	5	605
C	2	System	23	8	5	234
D	4	System	89	57	22	1160
E	3	System	37	15	7	801
F	3	Response	21	13	4	789

Table 1. The amount of content created by the four groups during their use of the MicroWriter.

While the process we use to organize ideas into an outline works well in practice, we suspect alternative algorithms could provide similar value; the primary benefit appears to come from grouping related ideas. We have observed that some reorganization of paragraphs and textual smoothing is necessary to ensure consistency across an entire piece. Users currently must do this outside of the system, but it could also be implemented internally using microtasked processes, such as those used by other systems to support editing [3], commenting, reviewing, or verifying other's work [15, 23].

STUDYING THE MICROWRITER

To understand how preexisting collocated groups might use microtasks for collaboration, we conducted six collaborative writing sessions with the MicroWriter and gathered insights through observation and qualitative feedback.

Participants

Six groups of 2 to 5 people used the MicroWriter, for a total of 19 participants (21% female). Exact group sizes are shown in Table 1. To target existing collaborations with shared self-directed writing needs, groups were recruited collectively via workgroup mailing lists and word-of-mouth. Half of the participants were 25 to 34 years old, and the rest were older. Each group consisted of a set of colleagues in technical roles working together on a common project within Microsoft. All participants reported being comfortable with their group members and most (12 of the 14 who replied) reported being willing to disagree with other group members. Participants did not receive any remuneration.

The six groups consisted of a mix of collaborators who had been working on a project for a long time (4/6), who knew each other but recently started working on a new project (1/6), and who just met and wanted to establish a shared project (1/6). Groups were asked to use the MicroWriter to collectively produce a report related to a self-motivated work task that was important to them. For example, one group had recently changed focus, and they used their MicroWriter session to write about their group's new direction. Another wanted to produce an overview of a system they were building to share with others who wanted to learn about it.

Protocol

Because there is high variance in the writing experience as a function of topic and group composition, it was infeasible to collect sufficient data to quantitatively compare the MicroWriter with an existing collaborative writing tool. As there exists significant work on how existing tools are used, we designed the study protocol to focus on understanding participants' unique experience with the MicroWriter.

MicroWriter Stage		Coordination	Group Location	Task Enjoyment		
				Like	Neutral	Dislike
Idea generate		Indiv	Together	12	0	2
Label	Initial	Indiv	Together	6	1	7
	Merge	Group	Together	1	3	10
	Finalize	Indiv	Together	3	5	6
Write		Indiv	Alone	8	5	1

Table 2. The different MicroWriter stages. Some stages were performed individually, while others were done as a group. Group members were collocated all stages except the final writing stage. Task enjoyment by stage is also reported.

MicroWriter Writing Session

Participants met in a conference room collectively with their group members. They were provided with laptops with the MicroWriter installed and given instructions about how to use the system. In one case (Group F), a group member participated remotely via Skype. One laptop was chosen to project on a large screen to provide some group context, but each individual interacted with their own version of the MicroWriter. Figure 4 pictures one of the groups. Two of the experimenters took notes during each session, with a focus on interpersonal communication and direct feedback. The advantage of having participants be collocated during the study was that it enabled direct observation of instances where coordination outside of the tool was necessary. The experimenters' notes were analyzed qualitatively.

Writing sessions lasted approximately 90 minutes, and were structured to reflect several different ways that groups might use microtasks to collaborate. The *Idea Generation*, *Initial Labeling*, and *Finalizing Labels* microtasks were performed individually even though all group members were in the same room. Participants were told they were welcome to converse and interact outside of the tool, but were not required to do so. These tasks are represented in Table 2 as ones that were performed Together (column: Group Location) via Individual microtasks (column: Coordination). In contrast, the *Label Merging* microtasks were performed collectively by the whole Group, with the participant who was projected driving input based on feedback from the entire group. Finally, the *Writing* microtasks were performed individually (Indiv) in an asynchronous, remote fashion, with participants writing paragraphs on their own in response to small sets of grouped ideas (Alone).

Post-Session Survey

After finishing the group report each participant was asked to complete a short online survey about their experience. The survey asked participants to share their agreement (on a 5-point Likert scale) with statements related to their feelings towards writing in general (e.g., "I enjoy writing"), their

group (e.g., "I am comfortable with people in the group"), their experience writing the report (e.g., "I felt I contributed to the writing process") and the final report (e.g., "The report we produced provides a useful starting point"). Additionally, they were asked how much they enjoyed each stage of writing with the MicroWriter (again, on a 5-point Likert scale), and asked to provide optional free text feedback about their ability to contribute to the report, collaborate with group members, and create a report. The complete survey can be found at http://aka.ms/chi16-microwriter.

Fourteen of the 19 participants completed the survey, with at least one person from each group replying. All of the direct quotes presented in this paper are drawn from the free-text survey replies, and are minimally edited to support legibility.

HOW GROUPS USED THE MICROWRITER

We now give an overview of how the groups used the MicroWriter, and then look more closely at each stage.

Overview

The six groups created six written reports, ranging in length from 234 words to 1160 words (see Table 1). By way of reference, the average report contained approximately the same number of words (724) as shown on the first page of this paper. Of the participants who replied to the follow-up survey, most (8 out of 14) reported that they liked writing. However, just about everyone (all but 3) generally found it hard start writing a new document. Participants felt they produced a good starting point with the MicroWriter, with seven participants agreeing or strongly agreeing with the statement "The report we produced provides a useful starting point," and nobody disagreeing. As one participant (Group D) noted, the tool seemed to provide a good way to get going: "*It was a relatively fast way to divide the work and produce a great starting point for an essay. I think we'll actually use this as our starting point when we decide to write more about it, so it really was quite useful.*"

The process of collaborating with the MicroWriter felt very different to participants as compared with their previous collaborative writing experiences. This is illustrated by a statement from a participant in Group E: "*The collaborative aspect of the process was great. Much more interactive than web-based document authoring or worse yet coordinating via email. The interactivity probably comes with an efficiency tradeoff, but for the right project, the tradeoff is probably well worth it.*" In particular, participants noted differences with the bottom-up process of organizing content as opposed to a more typical outline-structured process, and with the fact that contributions were made in a distributed fashion rather than in clustered chunks. One participant (Group A) expressed surprise at seeing all of the different pieces come together in the end, saying, "*I enjoyed the end product; in fact, I found it unexpectedly nice, seeing as the paragraphs were formed from independent ideas/threads.*"

A dominant participation style for typical collaborative writing is 'editing and commenting' on content someone else has created [32]. In contrast, all of the content generated by the MicroWriter was created collaboratively. As a participant from Group F notes, "*Typically for this sort of writing task one of us would write a full draft and then circulate and edit over email. The tool changes this up a bit by producing an initial draft that is drawn up collaboratively.*" The general feeling was that this process supported diverse contributions. Thirteen of the survey respondents agreed that, "I felt I contributed to writing process," and nobody disagreed, and this sentiment is supported by verbatim feedback from multiple participants: "*I feel that this approach to content creation leveled the playing field for ideas and contributions from any source,*" (Group B). "*I think it .. democratized the ability for multiple team members to contribute,*" (Group D). "*I felt the process enabled us all to contribute significantly. This was a major strength as I imagine the standard process is generally dominated by a small number of people from the group doing the authoring,*" (Group E).

Figure 4. A collocated synchronous group using the MicroWriter for collaborative writing.

We now look more closely at each stage individually. Participants' enjoyment of each stage is shown in Table 2.

Idea Generation

The *Idea Generation* stage was participants' favorite. Twelve of the survey respondents (out of 14) enjoyed the associated microtasks, while only two disliked them. This stage typically took about fifteen minutes, and produced anywhere from 21 to 89 ideas (see Table 1).

Defining ideas: In general participants took some time to enter the first few ideas, needing first to establish what an "idea" meant to the group. Most groups ended up defining ideas as important points that they wanted to be included in the final report. Once the first idea was entered many others quickly followed, with participants using examples from others for fodder. Some wanted recently entered ideas to be highlighted within the interface to indicate successful entry, while others felt not having a just-entered idea prominent on the screen help them to move on quickly to new ideas.

Some participants used the notion of an outline to influence the *Idea Generation* process, and their rendition of ideas resembled the structure of an outline. One participant (Group D) suggested a similar process to that used for *Idea Generation* to create structure within the document, "*I think the* Idea Generation *stage encouraged a diverse range of contributions (it was a really efficient way to cover the space*

of topics). I think it would have been great if there was also a collaborative phase where we could decide together what are the important topics. That way we could flush out ideas that the group as a whole thinks are important."

Group interaction: Although all of the MicroWriter microtasks were designed to not require coordination, we observed considerable interaction among group members. During ideation some of this interaction happened verbally, and some took place within the tool. We observed that group members often took a break to read the ideas that had been entered so far and then return to enter additional ideas. However, some deliberately refrained from checking ideas from others while entering their own to maintain their flow.

Interaction varied across groups based on how developed the report topic was. For projects in a brainstorming phase, entering ideas was interleaved with verbal discussion. For more mature projects discussion was more limited, and we sometimes observed participants transferring ideas from external sources such as notebooks or pieces of paper. Group size also appeared to impact the group's interaction. For larger groups, discussions took place both with the entire group and within smaller subsets while others continued to enter ideas. Groups with only two members coordinated the most; typically both group members discussed what to enter and then one entered the idea to avoid duplication.

Some participants felt that switching between discussion and entering ideas negatively impacted brainstorming. These participants requested affordances to reduce the transition cost, such as automatic transcription of discussion points. Highlighting a desire to take advantage of being collocated, one participant (Group B) commented: "*We were focused on screens as opposed to be talking to each other. When we are already in the same room, I would like to spend as much time as possible in discussion and save individual work to when I am in my office.*" Another participant (Group E) desired more support for discussion: "*Finding a way to encourage discussion while using the MicroWriter would be good.*" Prior work has highlighted the importance for supporting back-channel discussions during collaborative writing [32].

Anonymity: Participants valued anonymity while generating ideas. In groups with more than two people, the *Idea Generation* phase allowed them to enter ideas rapidly without concern for quality, syntactic or semantic issues, or how the idea may be perceived by the group. For awareness, many participants wanted the system to show which ideas belonged to them while not revealing the source to others.

Even though an individual's work was anonymous, three participants agreed with the statement, "I felt self-conscious while using the MicroWriter." Some felt peer pressure when they saw lots of ideas being generated by other people, and several noted that it felt like others were producing more ideas than they were. Participants from Groups A and C were particularly self-conscious, as each contained only two group members and this made it obvious who contributed what. A member of Group C reported, "*I feel I perform worse at these tasks with an audience, and so when there is an audience I prefer to be prepared or rehearsed, at least enough to boost my confidence.*"

Labeling

During the *Labeling* phase groups produced between 8 and 57 labels (Table 1). The bottom-up approach to organization enabled the groups to view their ideas in new ways. As one participant (Group C) said, "*There was value in the process, with regard to brainstorming, identifying salient themes (labels), and then seeing what emerged when asking how the themes applied to topics we hadn't explicitly thought to apply them to previously.*" However, it was hard for participants to separate the process of generating ideas from organizing them. A participant from Group F reported, "*Perhaps the most difficult part of the process was the discussion about how to construct a reasonable argument. During this early phase the tool was not helpful. So it's best to have at least a good idea of what you're writing before using the tool.*"

The *Labeling* process is divided into three substages, and participants responded differently to each, appearing neutral towards the *Initial Labeling* and *Finalizing Labels* substages and reacting negatively to *Merging Labels* (see Table 2).

Initial Labeling

During the *Initial Labeling* substage each group member was presented with an idea, one at a time, and asked to provide one or more relevant label. Even though this substage was designed to support independent contributions, in practice we observed that some coordination was required.

Defining labels: Similar to what we observed during the *Idea Generation* stage, at the start of the *Initial Labeling* substage participants had to establish what a "label" meant to the group. They often found it difficult to distinguish between a label and an idea, with many ideas receiving a single label with similar content to the idea. Generating a label that aptly represented an idea required thought, and participants often wanted to revisit a previously labeled idea to correct it.

The goal of the *Initial Labeling* substage is to produce a large set of diverse and independently generated labels. For this reason we chose not make label suggestions accessible to participants for reuse through the interface, though they could reuse them from memory. However, most participants wanted to converge on a fixed vocabulary first. They asked to see automatic label suggestions, both in the form of label-completion when they started to type and of label-suggestion based on analysis of the text or the labels that other people had generated. When participants were shown ideas that had already been labeled by other group members, we observed this helped convergence and inspired new label ideas.

Understanding ideas: Participants did not always understand the ideas others had entered, and sometimes had to ask for clarification. This stimulated discussions among group members, and occasionally resulted in the entry of new ideas. Overall, the difficulty of the *Initial Labeling* substage appeared to relate to the level of coordination during the *Idea*

Generation stage; when there was a lot of discussion during *Idea Generation*, the acquired awareness seemed to benefit the labeling process and reduce the need for clarification.

Progress: The MicroWriter provides a progress update during the *Initial Labeling* substage reflecting the total number of ideas labeled by the group (e.g., "57 of 61 labeled" in Figure 2a). Many participants found the fact that this represented group progress confusing, and wanted instead to see how many ideas they themselves had labeled.

Merging Labels

After providing an *Initial Labeling,* the group moved on to *Merging Labels*. For this substage everyone worked together using a shared screen to create a single set of canonical labels. As one participant (Group D) reported, "*The most collaborative task during the process was in categorizing the tags. This was also the most frustrating task*." Overall, the *Merging Labels* substage appeared particularly challenging, and most (10 out of 14) participants did not enjoy it. Even though the stage was fully collaborative and participants were able to quickly identify similar labels, the process was thwarted by the volume and the complexity of the task. This was particularly true for large groups, perhaps due to the larger number of ideas and labels. In contrast, for smaller groups merging was simpler. This could be because the *Idea Generation* stage was collaboratively performed, resulting in a smaller set of ideas and more cohesive labels.

Understanding labels: Just as participants needed to understand each other's ideas to perform the *Initial Labeling* microtasks, participants needed to understand each other's labels to merge them. Participants sometimes meant different things by the same label. In accordance with our design decision to provide minimal context for each microtask, participants were asked to merge labels using only the label text. However, participants sometimes felt that the context necessary to understand even the labels that they themselves had produced got lost, and this discouraged people from merging ambiguous labels. Our observations suggest that including a few representative ideas with each label as context might help address this issue.

Finalizing Labels

The *Merging Labels* substage reduces the space of possible labels into a manageable canonical set. The final *Labeling* substage, *Finalizing Labels*, once again involved participants working on their own, this time to apply the canonical labels to each individual idea. The task of selecting, rather than producing labels appeared to be a good exercise for many participants in applying labels that they had not thought about before to an idea. While this substage was designed to be performed either independently or collaboratively, we found that some participants preferred to collaborate.

Most participants felt neutral about this substage (Table 2), and it generally appeared to be fairly straightforward and perhaps a little boring. The perceived complexity of this substage seemed to depend on how much interaction between the participants had taken place during the *Labeling* substages; the more the discussion, the less the conflict there was in finalizing what labels apply.

Writing

The final stage in the process was expanding a set of related ideas into text. In contrast to the preceding stages, the *Writing* stage happened remotely and asynchronously. Members of each group received a set of labels and associated ideas via email and were asked to produce text corresponding to those ideas. As intended, we found that group members produced the text in short bursts and at different times. This demonstrates that while the initial ideation and clustering process happened in sync with collocated collaborators, the final text generation could happen independently and without a huge time commitment. Participants generally reacted positively to this stage, with 8 of our 14 survey respondents enjoying it and only one not.

Summary

By studying six workgroup groups using the MicroWriter to produce personally meaningful writing, we found that the process of generating written content collaboratively using microtasks produced coherent, useful text in a reasonably short amount of time. The MicroWriter seemed to help our participants overcome the reported challenge of initiating writing and produce a useful starting point.

The fact that the MicroWriter breaks the writing process into microtasks allowed every group member to feel like they contributed to the writing process in a more integrated manner than would typically happen with collaborative writing. Participants particularly liked the initial ideation phase, and found that the bottom up approach to idea organization required by microtasking, while unfamiliar, enabled them to view their ideas in new ways.

Participants were often able to perform the MicroWriter microtasks in bursts without requiring a lot of context, similar to how crowd workers operate. However, the process did not completely eliminate the need for coordination and participants sometimes needed to clarify tasks with the group. Some participants felt that having individual rather than group tasks limited their ability to collaborate as a group, and many disliked the shared task of merging labels which required considerable coordination overhead.

DISCUSSION

We now dive into how the successes and failures of the MicroWriter can provide insight into how to design a collaborative writing system that leverages microtasks.

Managing the Need for Context

The MicroWriter's task breakdown was designed to require very little context to complete each microtask. Nonetheless, we still found that participants sometimes needed more context than they were given. For example, while labeling ideas that were contributed by other group members, participants occasionally had to ask the group to clarify the underlying message so that they could provide a meaningful label. Decoupling the label from the idea during the *Labeling*

stage also made it difficult for participants to recall the context in which they had produced the label. This suggests it may be important to carry over context across stages. Different task breakdowns or task allocation strategies than those employed by the MicroWriter may help users bypass the need for other types of context. For example, if a user were only asked to label ideas they themselves produced, they would not need to clarify the idea's meaning.

Collaboration Styles

Based on the observations of the sessions, we found that groups leveraged microtasking differently depending on the size of the group, existing group dynamics, maturity of the project (e.g., newer projects involved more brainstorming and idea generation) and individual preferences for collaboration. Some groups communicated outside of the MicroWriter as they performed the microtasks, and this ensured convergence, reduced redundancy, and resolved conflict earlier in the process. Verbal discussions appeared particularly prominent with smaller groups. Being collocated also appeared to have impacted the collaboration style, as it made participants conscious about their contributions to the process relative to others. Previous research [25] highlights the need for writing tools to support synchronous and asynchronous collaboration, and the MicroWriter does this seamlessly. While we had only one remote participant and therefore have little evidence on how this process may affect remote collaboration, we believe that being remotely located during this process would likely resulted in more independent contributions without the added pressure of what the collaborator is producing.

Improving Efficiency

Since the microtasks were designed to be completed independently, participants sometimes completed one stage before others. Rather than having them wait, it is almost certainly more efficient for users to work on different stages at different times according to preference, skill, availability, and need. It may also make sense to allocate related microtasks to the same individual to minimize the need for internal coordination and clarification. For example, the best person to label an idea may be the one who generated it.

Anonymity

Participants liked the fact that their actions were anonymous as it allowed them to freely enter content without concerns about judgment. Prior work has shown that during collaborative writing, people are concerned about how their behaviors will be perceived by others and take steps to reduce the opportunity for negative perceptions [5]. The MicroWriter makes it impossible to identify the individual user who completes any given microtask. This frees users from concerns about how they appear to others, but presents problems for larger groups when conflict resolution or clarification might require coordination. There may also be an additional issue of understanding how group members are contributing in the process, and it is possible that without explicit ownership some people may become freeloaders.

Additional Usage Scenarios

The MicroWriter implements one possible breakdown for collaborative writing, based on the bottom-up approaches used in crowdsourcing. However, the optimal task breakdown may depend on many factors, including who will perform the task, the context they have about the task, their available time, and device form factor. For example, the microtask structure employed by the MicroWriter enabled collaboration across group members with limited coordination, but a task structure providing larger subtasks might be more useful when an individual sets aside an entire day for focused writing. Smaller tasks may be needed when using the MicroWriter on a mobile phone or smartwatch.

The MicroWriter microtasks are not limited to being completed by the groups who have some shared context. Some subset of the tasks could be performed by just one individual who has complete context [31], or by many unknown individuals (i.e., the crowd) who have no context. Complex personal information tasks that require deep personal knowledge or contain private information [20] can be accomplished in smaller chunks by an individual using the same processes [31]. The structure has benefits beyond coordination, such as enabling individuals to complete tasks in short bursts using different form factors (e.g., from mobile phones where it is difficult to view a lot of content) and with varying time constraints (e.g., during micromoments that can be used for quick tasks [6]). Additional studies are necessary to understand how what we observed with the MicroWriter generalizes to these other usage scenarios.

Automation

While in this paper we focused on writing microtasks that are done manually, some of the tasks we explored could be automated. There are existing approaches to group relevant ideas or apply labels that the MicroWriter could implement. Because microtasks are small and structured, it may be possible to use machine learning algorithms with the user's input as training data and slowly take over the task.

Training

The process of using microtasks for writing was unfamiliar to all of our participants, and it is likely that frequent users would use the capabilities provided by the MicroWriter differently. To help participants understand the underlying concepts we walked participants through each stage to describe how the microtasks should be performed. However, a working microtasking system would need to be extended to teach users how to perform each type of microtask the first time it is encountered. Some aspects of training are likely to be difficult. Participants consistently found it hard to focus the task they were performing without considering other tasks. We plan to explore whether people can be trained over time to ignore context or tasks structured so as to minimize the desire for context, but it may be that context is just too important for people to ever be able to disregard it.

Support for Additional Aspects of Writing

We observed that reports that required brainstorming encountered unique challenges compared to reports where

Ideas (all have the label *Brainstorming*)	Corresponding Paragraph
- Brainstorming support needs different things - Should brainstorming be part of this process - or should brainstorming be completed a priori - Effects of state of project - The informal nature of the tool suggests that people might be willing to use this for a brainstorming stage - how to support that in the tool - Brainstorming projects had some trouble during ideation.	"We observed that reports that required brainstorming encountered unique challenges compared to reports where group members already knew the topic being written about. The informal nature of the tool suggests that people might want to use this for a brainstorming stage. However, groups focused on brainstorming had some trouble during the ideation stage because the group wanted to be involved in a joint conversation. All writing inherently involves some brainstorming, and an interesting direction for research is to explore whether the ideation stage could be expanded to include better brainstorming support, or if brainstorming must be completed a priori for the tool to work successfully."

Table 3. Example transition of a set of ideas into a paragraph of text. The ideas are all associated with the label *Brainstorming*. A version of this paragraph appears in the Discussion Section in a subsection titled *Support for Additional Aspects Writing*.

group members already knew the topic being written about. While the informal nature of the MicroWriter seems to lend itself to brainstorming, groups focused on brainstorming found it limited conversation during the *Idea Generation* stage. All writing involves some brainstorming [25], and it is interesting to consider whether *Idea Generation* could be extended to better support brainstorming or if brainstorming should be completed a priori. Little et al. [21] explored microtask processes for crowd brainstorming.

Understanding the Big Picture

Breaking the macro-task of writing down into small microtasks carried benefits but also presented challenges. For example, not being able to see the big picture was sometimes stressful for participants. It may be that not all aspects of writing can be done via microtasks, because some tasks may require a complete picture of the entire piece. However, microtasking may still help make these larger tasks shorter and easier by siphoning off the busy work.

Automated support for task breakdown that takes into account the user's available time and context could be useful in this regard. In future work we plan to compare people's experiences with the MicroWriter with other macro-task based approaches to collaborative authoring to build a clearer picture of the relative benefits.

Error Handling

While significant work in crowdsourcing has focused on quality assurance, collaborative microtasking requires less validation because the person performing the work has a vested interest in the task being performed well. For example, while the approach we use to organize writing builds on a workflow described by Chilton et al. [9], it employs fewer rounds and no validation. However, people were concerned about making errors, and this was particularly apparent in interactions between dyads where there was no anonymity. Error handling is important part of microtasking in crowd work, and in this particular scenario even with context and motivation in a non-adversarial setting it is still possible to make mistakes, making error handling important. There are a number of approaches from crowdsourcing that could be borrowed [3, 21].

Impact

Our results are useful when considering the impact of microtasking research on collaborative practices, where there is typically a lot of shared context and coordination is vital for successful interactions. While we demonstrate different styles of collaboration that emerged through a particular breakdown of writing, we expect the observations to generalize to other types of tasks as well. Tools built for microtasking between groups should therefore take into consideration the need for supporting some context sharing, discussion as well as awareness of group members' activity.

Demonstration of the MicroWriter in Action

To demonstrate the MicroWriter in action, we wrote this Discussion Section using the tool. The screenshots in Figures 1, 2, and 3 represent our experience with different stages during this writing process. While writing we followed exactly the same protocol that was used to study the tool. An example of how a set of ideas transitioned into a paragraph during the *Writing* stage is shown in Table 3. In all, 61 ideas were created with 22 labels, which resulted in 20 paragraph units and a 1698 word report. We tightened the text, made local edits to add references and for terminological consistency, and merged related paragraphs to form the existing Discussion Section. Our experience with the MicroWriter was similar to that of our participants. We particularly enjoyed seeing our ideas come together in individual paragraphs and felt more ownership of all of the text in the discussion section than typical.

CONCLUSION

In this paper we explored how breaking the writing process down into small, contained microtasks could be used to support collaboration with limited coordination and turn an overwhelming task into one that is less daunting and easier to coordinate. We presented the MicroWriter, a system designed for collaborative writing via small microtasks that can each be performed with limited context. By studying 19 people working in six preexisting groups of 2 to 5 people, we highlighted some of the pros and cons of using microtasks for collaborative writing. The MicroWriter provides insight into new ways to support existing collaborations using recent advances in microtasking and crowdsourcing.

REFERENCES

1. Elena Agapie, Jaime Teevan, and Andrés Monroy-Hernández. 2015. Crowdsourcing in the field: A case study using local crowds for event reporting. In *Proceedings of the Third AAAI Conference on Human Computation & Crowdsourcing* (HCOMP '15). AAAI, New York, NY, USA.
2. David Allen. 2002. *Getting Things Done: The Art of Stress-Free Productivity*. Penguin Books.
3. Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. 2010. Soylent: A word processor with a crowd inside. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology* (UIST '10). ACM, New York, NY, USA, 313-322. http://doi.acm.org/10.1145/1866029.1866078.
4. Jeremy Birnholtz and Steven Ibara. 2012. Tracking changes in collaborative writing: Edits, visibility and group maintenance. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work* (CSCW '12). ACM, New York, NY, USA, 809-818. http://doi.acm.org/10.1145/2145204.2145325.
5. Jeremy Birnholtz, Stephanie Steinhardt, and Antonella Pavese. 2013. Write here, write now!: An experimental study of group maintenance in collaborative writing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13). ACM, New York, NY, USA, 961-970. http://doi.acm.org/10.1145/2470654.2466123.
6. Carrie J. Cai, Philip J. Guo, James R. Glass, and Robert C. Miller. 2015. Wait-Learning: Leveraging Wait Time for Second Language Education. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15). ACM, New York, NY, USA, 3701-3710. http://doi.acm.org/10.1145/2702123.2702267.
7. Carrie J. Cai, Shamsi T. Iqbal, and Jaime Teevan. Chain reactions: The impact of order on microtask chains. In *Proceedings of the 34th Annual ACM Conference on Human Factors in Computing Systems* (CHI '16). ACM, New York, NY, USA. http://dx.doi.org/10.1145/2858036.2858237.
8. Justin Cheng, Jaime Teevan, Shamsi T. Iqbal, and Michael S. Bernstein. 2015. Break it down: A comparison of macro- and microtasks. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15). ACM, New York, NY, USA, 4061-4064. http://doi.acm.org/10.1145/2702123.2702146.
9. Lydia B. Chilton, Greg Little, Darren Edge, Daniel S. Weld, and James A. Landay. 2013. Cascade: Crowdsourcing taxonomy creation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13). ACM, New York, NY, USA, 1999-2008. http://doi.acm.org/10.1145/2470654.2466265.
10. Haakon Faste, Nir Rachmel, Russell Essary, and Evan Sheehan. 2013. Brainstorm, chainstorm, cheatstorm, tweetstorm: New ideation strategies for distributed HCI design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13). ACM, New York, NY, USA, 1343-1352. http://doi.acm.org/10.1145/2470654.2466177.
11. Linda Flower and John R. Hayes. A cognitive process theory of writing. *College Composition and Communication,* 1981.
12. Victor M. González and Gloria Mark. 2004. "Constant, constant, multi-tasking craziness": managing multiple working spheres. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '04). ACM, New York, NY, USA, 113-120. http://doi.acm.org/10.1145/985692.985707.
13. Nick Greer, Jaime Teevan, and Shamsi T. Iqbal. 2016. An introduction to technological support for writing. Microsoft Research Tech Report MSR-TR-2016-001, January 2016.
14. Ronald T. Kellog. 1988. Attentional overload and writing performance: Effects of rough draft and outline strategies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(2).
15. Joy Kim, Justin Cheng, and Michael S. Bernstein. 2014. Ensemble: Exploring complementary strengths of leaders and crowds in creative collaboration. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (CSCW '14). ACM, New York, NY, USA, 745-755. http://doi.acm.org/10.1145/2531602.2531638.
16. Aniket Kittur, Boris Smus, Susheel Khamkar, and Robert E. Kraut. 2011. CrowdForge: Crowdsourcing complex work. In *Proceedings of the 24th annual ACM symposium on User interface software and technology* (UIST '11). ACM, New York, NY, USA, 43-52. http://doi.acm.org/10.1145/2047196.2047202.
17. Nicolas Kokkalis, Thomas Köhn, Johannes Huebner, Moontae Lee, Florian Schulze, and Scott R. Klemmer. 2013. TaskGenies: Automatically providing action plans helps people complete tasks. *ACM Trans. Comput.-Hum. Interact.* 20, 5, Article 27 (November 2013), 25 pages. http://doi.acm.org/10.1145/2513560.
18. Anand Kulkarni, Matthew Can, and Björn Hartmann. 2012. Collaboratively crowdsourcing workflows with Turkomatic. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work* (CSCW '12). ACM, New York, NY, USA, 1003-1012. http://doi.acm.org/10.1145/2145204.2145354.
19. Walter S. Lasecki, Kyle I. Murray, Samuel White, Robert C. Miller, and Jeffrey P. Bigham. 2011. Real-

time crowd control of existing interfaces. In *Proceedings of the 24th annual ACM symposium on User interface software and technology* (UIST '11). ACM, New York, NY, USA, 23-32. http://doi.acm.org/10.1145/2047196.2047200.

20. Walter S. Lasecki, Jaime Teevan, and Ece Kamar. 2014. Information extraction and manipulation threats in crowd-powered systems. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (CSCW '14). ACM, New York, NY, USA, 248-256. http://doi.acm.org/10.1145/2531602.2531733.
21. Greg Little, Lydia B. Chilton, Max Goldman, and Robert C. Miller. 2010. Exploring iterative and parallel human computation processes. In *Proceedings of the ACM SIGKDD Workshop on Human Computation* (HCOMP '10). ACM, New York, NY, USA, 68-76. http://doi.acm.org/10.1145/1837885.1837907.
22. David Martin, Benjamin V. Hanrahan, Jacki O'Neill, and Neha Gupta. 2014. Being a Turker. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (CSCW '14). ACM, New York, NY, USA, 224-235. http://doi.acm.org/10.1145/2531602.2531663.
23. Michael Nebeling, Anohong Guo, Alexandra To, Steven Dow, Jaime Teevan, and Jeffery Bigham. 2015. WearWrite: Orchestrating the crowd to complete complex tasks from wearables (We wrote this paper on a watch). arXiv:1509.02983.
24. Michael Nebeling, Alexandra To, Anhong Guo, Adrian de Freitas, Jaime Teevan, Steven Dow, and Jeffery Bigham. WearWrite: Crowd-assisted writing from smartwatches. In *Proceedings of the 34th Annual ACM Conference on Human Factors in Computing Systems* (CHI '16). ACM, New York, NY, USA. http://dx.doi.org/10.1145/2858036.2858169.
25. Christine M. Neuwirth, David S. Kaufer, Ravinder Chandhok and James H. Morris. 2000. Computer support for distributed collaborative writing: A coordination science perspective. *Coordination Theory and Collaboration Technology*.
26. Sylvie Noël and Jean-Marc Robert. 2004. Empirical Study on Collaborative Writing: What Do Co-authors Do, Use, and Like? *Computer Supported Cooperative Work* 13, 1 (January 2004), 63-89. http://dx.doi.org/10.1023/B:COSU.0000014876.96003.be.
27. Sondra Perl. 1979. The composing processes of unskilled college writers. *Research in the Teaching of English*.
28. Ilona Posner and Ronald Baecker. How people write together. 1992. *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences*.
29. John Sadauskas, Daragh Byrne, and Robert K. Atkinson. 2015. Mining memories: Designing a platform to support social media based writing. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15). ACM, New York, NY, USA, 3691-3700. http://doi.acm.org/10.1145/2702123.2702383.
30. James Tam and Saul Greenberg. 2006. A framework for asynchronous change awareness in collaborative documents and workspaces. *International Journal of Human-Computer Studies*, 64.
31. Jaime Teevan, Daniel J. Liebling, and Walter S. Lasecki. 2014. Selfsourcing personal tasks. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '14). ACM, New York, NY, USA, 2527-2532. http://doi.acm.org/10.1145/2559206.2581181.
32. Bill Tomlinson, Joel Ross, Paul Andre, Eric Baumer, Donald Patterson, Joseph Corneli, Martin Mahaux, Syavash Nobarany, Marco Lazzari, Birgit Penzenstadler, Andrew Torrance, David Callele, Gary Olson, Six Silberman, Marcus Stünder, Fabio Romancini Palamedi, Albert Ali Salah, Eric Morrill, Xavier Franch, Florian Floyd Mueller, Joseph 'Jofish' Kaye, Rebecca W. Black, Marisa L. Cohn, Patrick C. Shih, Johanna Brewer, Nitesh Goyal, Pirjo Näkki, Jeff Huang, Nilufar Baghaei, and Craig Saper. 2012. Massively distributed authorship of academic papers. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '12). ACM, New York, NY, USA, 11-20. http://doi.acm.org/10.1145/2212776.2212779.
33. Rajan Vaish, Keith Wyngarden, Jingshu Chen, Brandon Cheung, and Michael S. Bernstein. 2014. Twitch crowdsourcing: Crowd contributions in short bursts of time. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14). ACM, New York, NY, USA, 3645-3654. http://doi.acm.org/10.1145/2556288.2556996.
34. Christopher D. Wickens. 2002. Multiple resources and performance prediction. *Theoretical Issues in Ergonomics Science*, 3(2).
35. Jeffrey M. Zacks, Barbara Tversky, Gowri Iyer. 2001. Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology General*, 130 (1).
36. Haoqi Zhang, Edith Law, Rob Miller, Krzysztof Gajos, David Parkes, and Eric Horvitz. 2012. Human computation tasks with global constraints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '12). ACM, New York, NY, USA, 217-226. http://doi.acm.org/10.1145/2207676.2207708.