The Future of Microwork

What happens when we algorithmically break complex productivity tasks down into microtasks? At Microsoft Research, the author and her team are accelerating a shift toward microproductivity to make it easy for people to get big things done one small step at a time.

By Jaime Teevan

DOI: 10.1145/3019600

I found it hard to start writing this article, and put it off for as long as I could. Complex tasks like writing are difficult to do because they seem to require long, uninterrupted periods of deep engagement to make meaningful progress, and it can be hard to successfully carve out the necessary time. Research suggests it takes up to 25 minutes to get up to full productivity, but that we are interrupted, on average, every 11 minutes [1]. And even when we are not interrupted by someone else, it can be difficult to stay focused—we regularly self-interrupt to check our emails, visit Facebook, or browse the web. Getting something big done seems to require a non-stop battle with ourselves for our attention.

But we can change this! My colleagues and I are exploring "microproductivity" as a way to help people easily perform large personal information tasks by breaking the tasks all the way down into a series of self-contained "microtasks" that only take a few seconds each to complete. Rather than fighting to create uninterrupted chunks of time to get things done, we are instead changing the tasks themselves to fit our fragmented work style.

Microproductivity: *The transformation of a large information task into a series of smaller microtasks for productivity purposes.*

Microtask: *A task that is quick and easy to perform and contains within it all of the context necessary to get it done.*

As an example, writing an article like this one normally requires me to block out a chunk of time on my calendar, successfully tear myself away from Facebook when the block comes up, open up Word, stare at the blank screen for a while, and then eventually start writing the text. In contrast, a microproductivity approach might have me capturing a bunch of ideas from my mobile device as they come to me, and organizing the ideas by tagging them individually in my spare time while waiting in line at Starbucks or on the bus ride home. Each of these actions take only a few seconds each to do and can then be used to produce an outline with clusters of ideas. All I need to do to create a first draft of the article is glue together the ideas in each cluster with a bit of text—a series of microtasks that can, themselves, be done in short bursts with limited focus. Microproductivity approaches to writing like this have been used to create published research papers [2], Wikipedia-like articles [3], and news articles [4].

Although microproductivity approaches can be applied to any number of information tasks, our team at Microsoft Research is exploring microproductivity in the context of writing. We chose writing as a model productivity task because it requires fundamental, but varied, skills such as reading, analysis, reasoning, and communication. Through writing, people solidify concepts that were previously

The Microsoft Research microproductivity team.

hazy, challenge and transform existing knowledge, and construct entirely new models of reality. As such, writing tools provide a valuable lens through which to understand and explore a range of problems related to information work. Writing is also a fun domain to study because the task is not obviously decomposable, and many people balk at the idea that we might be able to create a coherent narrative in short bursts of time with limited attention. To figure out how to successfully transform the creation of written text into lightweight microtasks, we have brought together experts in the domains of attention management, crowdsourcing, natural language processing, machine learning, personal information management, computer-supported cooperative work, and artificial intelligence.

Our team believes microproductivity will have a significant impact on when and how people complete information work, enabling individuals and groups to efficiently and easily complete large tasks that currently are challenging. New frontiers for the future of microwork are opening up as the result of rapid developments in three key areas related to our ability to: 1) decompose large tasks into microtasks, 2) efficiently complete each microtasks, and 3) source microtasks to the most appropriate actor—be it the task owner, another person, or even an automated process.

TASK DECOMPOSITION

Microproductivity can only occur if we know how to successfully decompose complex information tasks into their component microtasks, and then how to recompose the output of these microtasks into the finished product. For example, TurboTax makes filing your taxes easy by decomposing the information needed into a series of questions, and then recomposing the answers you give into the tax forms you need to submit.

We call the predefined structured series of microtasks that can be used to complete a larger task a "workflow." While the workflow for doing your taxes may be obvious, it is harder to imagine a good workflow for writing a blog post, preparing a presentation, or even scheduling a meeting. Fortunately, we can take inspiration from crowd platforms, where crowdsourced microtasks are increasingly being composed to accomplish complex tasks that are not obviously achievable via standalone microtasks. A number of successful workflows for decomposing tasks into smaller microtasks already exist, ranging from organization to planning to writing.

Workflow: *A predefined series of microtasks designed such that the output of the set of microtasks can be used to successfully complete a larger task.*

While crowdsourcing provides insights into how to decompose tasks, there may be reasons to decompose personal productivity tasks differently than we currently decompose crowdsourced tasks. A person doing their own task via microtasks has more context and different motivations than a crowd worker doing someone else's microtasks. Personal productivity microtasks might, for example, require less validation because the person performing the work has a vested interest in the task being performed well. Additionally, when you do your own tasks via microtasks you lose the opportunity to take advan-

Photo courtesy of Shamsi Iqbal

tage of the many different perspectives crowd workers can provide, but intentional use of context could help people identify creative solutions and draw unexpected connections.

Because workflows provide structure to open-ended goals, they can enable people to do things they may not be able to otherwise. For example, people without domain expertise can already make progress on tasks like filing taxes or producing a will simply by filling out a form. In the future, we all may be able to use workflows to write a novel or build a new mobile app.

As we develop custom workflows for complex tasks, we can take advantage of the fact that there are a number of common aspects to many information tasks that can be decomposed in common ways. For example, preparing a presentation, organizing notes, and structuring ideas into a report outline all involve organizing content into a hierarchy. Likewise, reading a paper, changing a document into a presentation, and creating an executive summary for a whitepaper all involve distilling and synthesizing key points from information. Several successful approaches to decompose these common subtasks already exist, which could be shared, reused, and composed into more complex workflows.

Moreover, while most existing microproductivity approaches currently expect task structure to be provided up front, workflows can also be developed dynamically. For example, the system TaskGenies [5] uses crowd workers to create action plans with concrete steps that help individuals successfully complete personal tasks. These workflows, once created, can then become part of the library of workflows available to others.

MICROTASK COMPLETION

Once a task is decomposed into its component microtasks, each individual microtask must be completed. Fortunately, concrete plans with actionable steps allow people to complete tasks better and faster [5]. Microproductivity has been shown to help people engage in difficult tasks, recover quickly from interruptions [6], and take advantage of small gaps in time, also called "micromoments," to be productive while mobile [7].

Micromoment: *A small gap in time that occurs between other segments of occupied time during which microtasks can be completed.*

To quantify the costs and benefits of doing a large task via a series of small microtasks, we conducted a controlled study with 110 participants and three task categories: arithmetic, sorting, and transcription [6]. In the study, we asked participants to complete a task and compare their performance on that task with an equivalent set of microtasks. For example, for the sorting task we displayed seven lines of text, each containing a list of 10 numbers, and asked participants to order the lines by the number of odd numbers in each list of numbers. We also had participants complete the same task via a series of microtasks by implementing a human-powered quicksort where individuals only compared pairs of lines and selected the line with more odd numbers. We found breaking the large tasks into microtasks resulted in longer overall task completion times, but higher quality outcomes and an easier experience.

Our findings from this study also suggested microproductivity may be more resilient to interruptions. Task resumption is easier when a person is interrupted at a breakpoint, and when the task being returned to has a clearly achievable short-term outcome. Researchers have tried to use these insights to decrease interruption costs by strategically scheduling interruptions to occur at breakpoints, helping users set goals upon interruption, and reminding users of their goals upon return. In contrast, microproductivity changes the nature of the task itself to make it more interruption-friendly by introducing breakpoints and creating clear, actionable steps to return to.

Modern information work is increasingly fragmented and mobile, and microproductivity allows us to adjust how we complete our tasks accordingly.

Recent advances in cloud and mobile technologies make it possible for people to complete microtasks from anywhere, but for microproductivity to be successful people need access to the right task at the right time. The microtask a person is asked to do should match the person's available form factor. For example, a colleague wrote the first draft of a research paper we collaborated on from his smartwatch, but he did almost no text entry from the watch. Instead, he got help from crowdsourced writers working at desktop computers. Turning his hastily dictated notes into coherent text required writing skills and a desktop environment, but occasionally providing feedback and answering questions about the underlying research ideas along the way did not. When we studied people using this watch-based writing system, we found they appreciated being able to take advantage of their mobile micromoments [7]. As one participant reported, "I don't think you can do anything productive with the watch these days. I was surprised I could do something interesting."

In addition to matching the available form factor, microtasks should match the person's available cognitive resources. Some microtasks are harder than others, and that makes them challenging to do while distracted. Fortunately, it is possible to help people build up to hard tasks. In a study we conducted looking at writing microtasks [8], we found people were able to do a better job with hard tasks on a sentence (like changing the tone of the sentence) after they had first done an easier task (like spell checking) on the same sentence. Thoughtful task ordering can help people take advantage of this by using easy microtasks to help draw people into larger and more complex microtasks, and even, potentially, into the larger complete tasks.

Microtasks not only make it possible to make incremental progress toward a larger goal, but also have the

potential to yield an early sense of accomplishment. Research shows concrete progress with frequent feedback can help reinforce a sense of self-efficacy. Completing microtasks could provide motivational benefits in addition to functional efficiency, encouraging continued productivity beyond an isolated moment.

MICROTASK SOURCING

Thus far we have focused on decomposing personal productivity tasks so that the resulting microtasks can be completed by the task owner, a process called "selfsourcing" [9]. However, not every piece of every information task, once broken down, requires personal knowledge to complete. For example, while it may feel like you are the only person who can write a summary of what you did this week, other people—or even automated processes—can help you out by performing some of the associated microtasks, such as providing feedback on the content you write, performing organizational microtasks, or copyediting lines of text.

Selfsourcing: *The practice of the task owner completing their own microproductivity microtasks.*

Soliciting support from peers (via "friendsourcing") or paid crowd workers (via "crowdsourcing") is easy with microproductivity, because sharing a self-contained microtask with another person requires almost no coordination overhead. Microtasks can be allocated to the most appropriate person depending on the required level of expertise, the person's availability, and their familiarity with the task. For example, when writing crowd workers could help you identify unclear sentences, your peers could rephrase these sentences in different ways, and you could select the final rephrasing for integration into the larger document.

Friendsourcing: *The practice of completing microproductivity microtasks by the task owner's friends or colleagues.*

Crowdsourcing: *The practice of completing microtasks by using a group of remote workers.*

To explore how microproductivity impacts collaboration, we conducted a study of collaborative microtasked writing with preexisting groups of colleagues creating work-related documents [2]. When we broke the task down into microtasks, the need for our participants to coordinate went down, while the rich interleaving of the content created by each individual went up. This is captured in a comment from one participant, who said: "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."

Microproductivity will have a significant impact on when and how people complete information work...

Since microtasks are small and self-contained, some common microtasks do not actually need to be done by a person. Instead, they can be done automatically. By incorporating automation as one form of task sourcing, microproductivity systems offer a pathway for integrating artificial intelligence into productivity tasks that are currently impossible to automate when considered as complete tasks. Additionally, because some microtasks can be learned from the data collected as people perform them, the system can first ask people to complete these microtasks on their own, then start suggesting solutions based on personalized or aggregate models, and finally take over the microtasks entirely. With collaboration and automation, the only aspects of a task that we will really ever have to do on our own are those for which we have unique insight.

Modern information work is increasingly fragmented and mobile, and microproductivity allows us to adjust how we complete our tasks accordingly. While microproductivity requires a focus on small details, the approach allows people to get things done in short bursts of activity, potentially with the help of other people or algorithms, so they use their focused periods of cognitive efforts on the key high-level aspects where they can contribute the most. My hope with microproductivity is one day writing an article like this will not involve my putting the task off for days and then staring at a blank screen when I finally get around to it. Instead, I want to be able to jot down ideas as they come to me and organize them from the bottom up in my spare time, so I start out with a rich populated document that the system automatically helps me engage with, edit, and develop.

References

[1] Mark, G., Gonzalez, V. M., and Harris, J. No task left behind? Examining the nature of fragmented work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '5) (Portland, April 2-7). ACM, New York, 2005, 321-330.

[2] Teevan, J., Iqbal, S. T., and von Veh, C. Supporting collaborative writing with microtasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '16) (San Jose, May 7-12). ACM, New York, 2016, 2657-2668.

[3] Hahn, N., Chang, J., Kim, J., and Kittur, A. The Knowledge Accelerator: Big picture thinking in small pieces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '16) (San Jose, May 7-12). ACM, New York, 2016, 2258-2270.

[4] Agapie, E., Teevan, J., and Monroy-Hernández, A. Crowdsourcing in the field: A case study using local crowds for event reporting. In *Proceedings of the Third AAAI Conference on Human Computation and Crowdsourcing*. 2015.

[5] Kokkalis, N., Köhn, T., Huebner, J., Lee, M., Schulze, F., and Klemmer, S. R. TaskGenies: Automatically providing action plans helps people complete tasks. *ACM Transactions on Computer-Human Interaction* (TOCHI) 20, 5 (2013)

[6] Cheng, J., Teevan, J., Iqbal, S. T., Bernstein, M. S. Break it down: A comparison of macro- and microtasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '15) (Seoul, April 18-23). ACM, New York, 2015, 4061-4064.

[7] Nebeling, M., To, A., Guo, A., de Freitas, A. A., Teevan, J., Dow, S. P., and Bigham, J. P. WearWrite: Crowd-Assisted writing from smartwatches. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '16) (San Jose, May 7-12). ACM, New York, 2016, 3843-3846.

[8] Cai, C. J., Iqbal, S. T., and Teevan, J. Chain reactions: The impact of order on microtask chains. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '16) (San Jose, May 7-12). ACM, New York, 2016, 3143-3154.

[9] Teevan, J., Libeling, D. J., and Lasecki, W. S. Selfsourcing personal tasks. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, 2014, 2527-2532.

Biography

Dr. Jaime Teevan is a principal researcher at Microsoft Research and an affiliate associate professor at the University of Washington. As a mother to four wild little boys, she is passionate about figuring out how to get big things done with limited attention.