How We Built the Reader

By Alexia Cambon, Alex Farach, Scott Counts, Andrey Zaikin

October 8, 2026

Read the article online

THE NEW FUTURE OF WORK READER / RESEARCHER TAKE

What an AI-first publishing experiment taught us about collaboration, expertise and the changing shape of work

BY THE READER EDITORIAL TEAM

Alexia Cambon, Alex Farach, Scott Counts, and Andrey Zaikin

SECTION 01

THE ESSAY

Starting with the problem, not the platform

Researchers are trained to produce research. They know how to frame questions, design studies, analyze evidence, and explain their findings to other experts, usually in the form of an academic paper. But increasingly here at Microsoft, we have seen demand rise to publish our research in other more consumable formats, to reach a broader audience. This requires different capabilities from the research team, including editorial development, web design, data visualization, communications, and distribution. Those capabilities often sit in different teams, with different priorities and timelines.

The New Future of Work Reader began as an experiment in closing that capability gap. For those academic research papers where there was value in creating an asset for a broader audience to consume, could it be done by a small research team empowered by AI? And could we build a repeatable process for doing so, rather than treating every piece as a bespoke publishing project?

In this piece, we report our observations so far. Our initial observation is that it is feasible for a four-person editorial team to build an AI-first publishing process that can turn an existing academic paper into a published, highly consumable asset with an amplification package within one working week. The system seems to have expanded what a small team can produce. But it has also shown that making production easier relocates the hardest parts of knowledge work as opposed to eliminating them.

> *We wanted to make the research accessible to people who might not read the academic paper. Much of the analysis was already in GitHub, which was familiar to us as data scientists, so GitHub Pages seemed like a practical place to start. It gave us flexibility in how we presented the work while we figured out what the Reader should be.*
>
> Alex Farach

This reflects something we increasingly see in AI-native teams. AI does not automatically make a team creative, but it can lower the cost of exploring a previously impractical path. A research team without a dedicated web-development function could now try their hand at building a website.

The editorial team agree that one of their most important learnings was first figuring out how to converge on a process. We needed to understand how a researcher's original work could retain its rigor and voice while being reshaped for people who might never read the academic paper. That meant breaking down publishing into a sequence of decisions that could be made visible, tested, and improved.

It is important to note that we do not publish academic papers on the Reader – this process still lives in its usual form through conference submissions or sites like arXiv. Publishing on the Reader is more akin to a translation exercise: translating the more academic research into something consumable to additional audiences. This translation exercise lends itself particularly well to AI, but we also wanted to ensure researchers have the choice of publishing through us without relying heavily on AI. Some of the researchers we work with prefer to write their asset themselves, with little assistance from AI, and we happily accommodate. But there are also many researchers who do not have the time or capacity to turn their papers into something more akin to a blog post, and in these cases, the ability to leverage the Reader's agentic system reduces a significant load. This piece explores that system.

Map the work before automating it

In building the Reader's agentic system, one of the most useful things we did was map the workflow before trying to automate it. We identified the recurring stages, the judgement required for each, and the conditions that sent a piece down a different path. That map became a workflow router: an operating layer that helps the AI invoke the appropriate process, skills, and rules for the task at hand.

THE WORKFLOW ROUTER SELECT AN INPUT TO TRACE ITS PATH

Research paper
Peer-reviewed source, dense argument, expert audience
Process: Editorial development
Skills: Editorial development; Journalism; Anti-AI-slop
Rules: No blank page: start from the paper; Preserve the researcher's claim

Recorded conversation
Researcher explains what matters in their own words
Process: Context capture
Skills: Journalism; Editorial development
Rules: Voice of the researcher is load-bearing; Human bonds before async work

Dataset / chart
Quantitative claim that needs a visual
Process: Data visualization
Skills: Data visualization; Anti-AI-slop
Rules: Specialist review before publish; Never overstate the effect size

Publish request
Draft approved by editor and researcher
Process: Release
Skills: GitHub Actions; Data visualization
Rules: Staging branch first — never straight to main; Amplification package ships with the piece

This mattered because an effective agentic workflow is not one enormous prompt. It is a collection of heuristics made explicit, that applies at every instance. What makes a paper suitable for the Reader? What must be preserved from the researcher's account? When does a data visual need specialist review? Understanding those heuristics up-front made the work more reproducible and constrained the AI where consistency mattered.

The system combines several capabilities: specialist skills support editorial development, journalism, GitHub Actions, and data visualization. An anti-AI-slop skill looks for the patterns that can make AI-assisted writing feel hollow or generic. Repository rules ensure that nothing moves directly to the main publishing branch without first passing through staging. Other rules avoid unnecessary blank-page requests. For example, before asking a researcher for a biography, the system compiles public information for them to review and correct.

We were interested by the effect of that last experience. Within this process, there is rarely an instance of a blank page. We begin with the researcher's paper, then schedule a recorded conversation in which the researcher explains, in their own words, what matters, what is easily misunderstood, and what they want a broader audience to take away. This recorded conversation creates essential context for the agentic operating system to lean on, and is constructed deliberately with that in mind. The agent then produces an initial expression of that material. Editors, researchers and partners then react to something concrete, rather than each starting from scratch. To be able to react to something concrete at every point of the process makes collaboration feels surprisingly easier.

THE MEETING, REFRAMED

TOGGLE TO SEE THE SHIFT

INFORMATION ERA

Information exchange

The meeting is designed primarily to update everyone on what they don't know. Each collaborator leaves with fragments, stored in their own individual sense of context.

INTELLIGENCE ERA

Context construction

The goal of the meeting is to produce the most valuable transcript for future AI grounding. Participants deliberately build a shared context layer which becomes critical input for the publication cycle.

- What matters
- What is easily misunderstood
- What the audience should take away

Recorded context

Agent draft

Same meeting. Different purpose: from hand-off to grounding.

> *There is an important point to be made here: it is possible that the editorial team could easily work asynchronously and use AI predominantly as their means for interacting. The reason this would be damaging in the long run is the need for the team to create a sense of unity as a group first, especially with the researchers we work with, who are entrusting us with some of their most important work. Ultimately, we realized that establishing human bonds first is critical for a better output.*
>
> Alexia Cambon

SECTION 02 · A WEEK IN THE LIFE

The production timeline, Monday to Friday

The resulting cadence is fast. In the most seamless instances, the editorial team can execute within a week: as a team, we aim to publish on a fortnightly basis. On Monday, the editorial team meets with the researcher and records the conversation. By Tuesday, a first web draft is ready to be reviewed and interrogated by the editorial team, who each provide their input and review. On Wednesday, the team tightens the argument while partners - from marketing, communications, or whichever functions may have a vested interest in the piece- contribute in parallel. On Thursday, final changes are incorporated and the piece is published. Overnight, the system produces a set of artifacts aimed at enabling amplification. By Friday, that package of materials is ready - and the editorial team share it with their network.

Speed is not the only point. The deeper outcome is that the process turns translation into an iterative collaboration, which is something AI is particularly suited to help with. That translation exercise may once have demanded more effort and capacity than was justified in terms of ROI. But today, with AI, researchers do not need to become web developers or communications specialists before their work can reach new audiences. Editors do not need to reconstruct the researcher's intent from the final paper alone. The system creates a shared object early and improves it through successive rounds of expertise. And each team member has now gained access to skills they never had - allowing them to do work they could not do before.

Monday · 09:00 · Recorded context

The editorial team meets with the researcher and records the conversation.

Tuesday · 14:00 · First expression

A first web draft is ready to be reviewed and interrogated.

Wednesday · 11:00 · Sharpened argument

The team tightens the argument while partners contribute in parallel.

Thursday · 16:00 · Live piece

Final changes are incorporated and the piece is published.

Friday · 10:00 · Amplification kit

The package of amplification materials is ready and shared with the network.

SECTION 03 · THE TEAM

Collaboration does not have to be hard

We often assume that collaboration becomes more difficult as more specialist perspectives enter the process. Our experience in building The Reader is that collaboration does not have to be hard when people share context, work from visible artifacts and understand how decisions will be made.

EDITORIAL DIRECTION

Alexia Cambon

Shapes editorial direction and audience relevance.

SYSTEMS

Alex Farach

Builds the technical and agentic publishing systems.

COMMISSIONING & RIGOR

INSIGHTS

Scott Counts

Commissions the right pieces and ensures scientific rigor and research integrity.

Andrey Zaikin

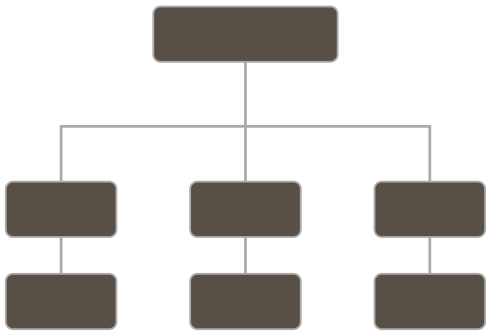
Translates evidence into accessible insights while measuring and improving the Reader's impact.

In that sense, the team resembles the "Hollywood model" of work: a small group of specialists assembles around a production, contributes distinctive expertise, and coordinates toward a shared outcome, disbanding when the "job to be done" is achieved. This model is much more achievable today, when AI allows a much smaller core team to hold more of the production system together. Skills and agents extend individual capabilities, while the team forms and reforms around the needs of a particular piece. Only Scott and Andrey share a reporting line on the team. There is no structural hierarchy within the editorial set-up.

CLASSIC MODEL

Organized by function

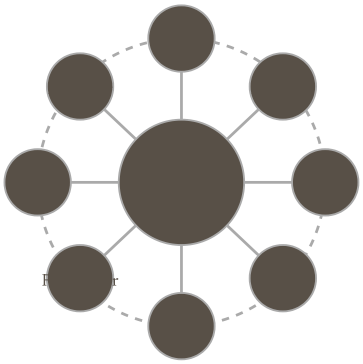
Specialized expertise separated into fixed departments, routed through hierarchy.



HOLLYWOOD MODEL

Organized by the job to be done

A small circle of specialists assembles around a production, then reforms for the next one.



The comparison with movie production has limits. In movie-making, every department influences the next: actors influence the costume department, the costume department influences the lighting department, the lighting department influences the cinematographer, etc. We observe that our collaboration is not quite as connected yet: it has often developed as a kind of parallel play, with each person building bespoke workflows around their own contribution. True team-level collaboration with AI is still difficult. We do not yet have a simple way for one person's agent to work with another person's agent using shared context, tools, standards, and decision rights. Individual augmentation is ahead of collective augmentation.

That is why improved collaboration is both a hypothesized outcome and an ongoing design challenge. The Reader provides a common workflow and shared artifacts, making human collaboration more concrete. At the same time, it exposes the next problem to solve: how to make the AI layer itself collaborative rather than a collection of individual systems. We need to be both interdependent and independent of each other in that sense. The Reader needs to be able to publish even if one of us is not here.

SECTION 04 · WHAT SHIPPED

What we think changed: our hypothesized outcomes

Note: we report on these purely from a place of qualitative observation. We are currently in the process of exploring carrying out a more empirical measurement on the longer-term impact of the Reader's system on these hypothesized outcomes.

01 Accelerated cadence

It appears that a small team can sustain a bi-weekly publishing rhythm and, in a well-defined cycle, move from the first researcher conversation to publication and rollout in five workdays. That would have been difficult to imagine if every stage required a separate hand-off and a blank-page start.

02 Increased reach

Our goal with building the Reader was to reach additional and new audiences for our scientific research. To do this, the process is designed to take an artifact created for an expert audience and iteratively adapt it for a different one without severing it from the underlying evidence or the researcher's intent. We hope to prove that the Reader has created a form of value for our research community within Microsoft that did not previously exist.

03 Improved collaboration

All four of the editorial team report positive sentiment around their collaboration, both within the team but also with the researchers who have gone through the process. Recorded grounding conversations, early drafts, and visible staging give researchers, editors, and support partners better objects around which to coordinate. Instead of spending their time describing what they might want, collaborators can evaluate and refine what is already there.

04 Organizational learning

We hypothesize that the agentic system is contributing to a bigger learning loop. The editorial team is not only publishing articles; we believe it is also discovering which parts of knowledge work can be reliably delegated, which require specialist verification, and which new capabilities are valuable enough to justify becoming part of the ongoing process. Each cycle updates the operating system.

These hypothesized outcomes should not be confused with unlimited capacity. The team can now produce more than it can always comfortably review. The supply of plausible content has increased faster than the supply of human attention. We predict that imbalance will be felt by many teams creating AI-first processes.

SECTION 05

THE NEW CONSTRAINT

When supply increases, demand and judgement become the constraint

Generative AI changes the supply-and-demand dynamics of knowledge work. Drafts, designs, summaries, and promotional artifacts that were once expensive to create can now be produced in abundance. But demand for those outputs does not automatically rise at the same rate, and the capacity to assess them remains stubbornly human. The bottleneck has moved from making to choosing.

This is not only a problem. A faster publishing cadence creates ripple effects across the organization. Communications partners, reviewers, authors, and leaders are asked to respond more frequently. If one team accelerates while the surrounding system does not, the speed can impose a cost elsewhere. Organizations may need deliberate friction, clearer thresholds, or lightweight governance to ensure that easier production does not simply create more work for everyone downstream.

Much of the solution here can be agentic: when communications partners flagged their challenges in reviewing content on a fortnightly basis, the editorial team asked them to codify their guidance and review principles in a playbook, which could then be turned into a skills file, to be incorporated into the Reader's agentic operating system. Thus, the human communications partners would only be called upon if their own skills file flagged something worthy of their involvement.

AI-FIRST WORKFLOW

Every Reader article is put before an agentic council of "judges" drawn from a library of stakeholder personas that make up our target audience, including executives, researchers, practitioners, and communicators. Each reviewer evaluates the piece through a different lens, highlighting places where claims, framing, or relevance may fall short for a particular audience.

Judges can only provide margin comments and not make changes directly. Final decisions remain with the Reader editorial team.

Illustrative review tasks, not a live status feed.

General Purpose:	Test CEO readability	Working
General Purpose:	Test COO usability	Working
General Purpose:	Audit scientific rigor	Working
General Purpose:	Test people relevance	Working
General Purpose:	Test manager scanning	Working
General Purpose:	Test provenance trust	Working

General Purpose: Test talent signal Working

SECTION 06

THE CEILING

Expertise still sets the ceiling

AI has allowed the Reader team to work across areas that sit outside traditional research roles, including publishing, visualization, web development, and social media. But we believe that access to a capability is not the same as mastery. “Good enough” can create meaningful new value where no alternative existed. It is not sufficient for every task.

> *AI can bridge an expertise gap for a time; it cannot eliminate the need for someone who knows how to verify, qualify, and improve the work.*
>
> Scott Counts

This is also why AI clichés are not confined to prose. There are recognizable patterns of AI writing, but also of AI design and visualization. Domain experts see the conventions, omissions, and shortcuts that a capable generalist might miss. The more consequential the output, the more important that judgement becomes.

AI also changes the boundaries of roles. If researchers can create websites and communications assets, will those tasks become expected parts of research work? For some people, expanded capability will reveal new interests and career paths. For others, it may add obligations they never chose. Teams will need to decide deliberately which adjacent capabilities should become core, which should remain optional, and when specialists should enter the process.

SECTION 07

THE RECIPE

A process for learning, not a finished recipe

The models, tools, and interfaces underpinning the Reader continue to change. That means the workflow cannot be treated as a machine that was built once and switched on. A model update can improve one stage, alter the quality of another or make an existing heuristic obsolete. The team must keep revisiting how tasks are routed, what the system can be trusted to do and where human review is most valuable. In short, the Reader’s agentic system requires maintenance – which is a net-new kind of work for everyone on the team.

This constant adjustment is not a temporary inconvenience on the way to a stable system. It may be a permanent feature of AI-first work. Teams need operating practices that can absorb capability changes without losing their quality standards or organizational memory. Explicit heuristics help because they make it possible to distinguish the enduring intent of the workflow from the particular model used to execute it.

The Reader's operating system is still evolving. We continue to add skills, refine rules, identify missing expertise, and adjust to new models. We are also learning where the underlying platform reaches its limits. GitHub Pages made the experiment possible, but a purpose-built publishing platform could eventually provide capabilities that creative improvisation cannot.

> *Every research team does not need to copy our exact stack. But all teams can make AI-first work more effective by making the work itself legible.*
>
> Andrey Zaikin

SECTION 08 · THE FIELD GUIDE

Make the work legible

A process for learning, not a finished recipe. Seven moves any team can apply to make AI-first work legible.

1. Map the basic process.
2. Identify the heuristics – where and how decisions are made.
3. Create grounding that captures expert intent.
4. Give collaborators something concrete to react to.
5. Separate areas where good enough creates value from those where only deep expertise will do.
6. Watch where increased supply creates new demands on attention.
7. Design collaboration rather than assuming it will emerge.

This last one may be the most important. A small team can now create outputs that once required a much larger organization. Sustainable scale, however, depends on clear roles, shared systems, feedback loops, deliberate constraints, continuous enhancement, and respect for specialist judgement.

The Reader shows what becomes possible when a team builds those conditions while it builds the product. The most interesting outcome may be that the process offers a glimpse of how future teams will work: assembling expertise around a shared purpose, extending it with AI, and continuously transforming both the system and themselves.

Final note: This piece itself was created using the very process we describe above.

AUTHORS

The team behind this Researcher Take



RESEARCH LEADER AT MICROSOFT

Alexia Cambon

Alexia Cambon studies AI adoption and the future of work at Microsoft. Her research examines how people move from using AI individually to collaborating with it in teams, including the role of workplace experiments and organizational practices.

LinkedIn



PRINCIPAL DATA AND APPLIED SCIENTIST AT MICROSOFT

Alex Farach

Alex Farach is a Principal Data and Applied Scientist at Microsoft and a former economist at the U.S. Bureau of Labor Statistics. He studies how AI changes work and productivity. He built the New Future of Work Reader and has co-authored Microsoft's Work Trend Index and New Future of Work reports.

Workforce Futures

Microsoft Research profile



SENIOR PRINCIPAL RESEARCHER AND RESEARCH MANAGER AT MICROSOFT

Scott Counts

Scott Counts is a Senior Principal Researcher and Research Manager at Microsoft, where he leads the IDEAS Research team. He uses large-scale behavioral data to study how people use AI, how it changes knowledge work, and its economic effects.

Personal website



SENIOR APPLIED SCIENTIST AND ECONOMIST AT MICROSOFT

Andrey Zaikin

Andrey Zaikin is a Senior Applied Scientist and Economist on Microsoft's IDEAS Research team. He studies AI adoption and human-AI collaboration, combining workplace data and labor-market evidence to understand changes in tasks and productivity.

LinkedIn