THE LENFEST INSTITUTE

2026 | CELEBRATING 10 YEARS

Insights from the Lenfest AI Collaborative & Fellowship Program

CAN AI HELP POWER A RENAISSANCE IN LOCAL NEWS?

FALL 2026 **AI IMPACT REPORT**

THE LENFEST INSTITUTE

2026 | CELEBRATING 10 YEARS

Contents

INSIGHTS FROM THE LENFEST AI COLLABORATIVE AND FELLOWSHIP PROGRAM

3 INTRODUCTION

6 PHILADELPHIA INQUIRER PROFILE

8 AI IN LOCAL NEWS: A PRACTICAL LIST

10 MEET THE AI FELLOWS

11 REMEMBERING DAVID CHIVERS

12 RAJESH BARADE, FELLOW

14 AARON BREZEL, FELLOW

16 DANA CHIUEH, FELLOW

18 MARK CHONOFSKY, FELLOW

20 KEVIN DELANEY, FELLOW

22 EJ FOX, FELLOW

24 KEVIN HOFFMAN, FELLOW

26 PETER RASMUSSEN, FELLOW

28 OMAR REID, FELLOW

30 RYAN THAYER, FELLOW









1. Yasser Mohammed, Applied AI Engineer at OpenAI **2.** Michelle Micone, Chief Marketing and Strategic Initiatives Officer at The Boston Globe **3.** Sonali Verma, AI Team Lead at The Lenfest Institute **4.** Adam Hirsch, Director of Product Engineering at ProPublica; Sapna Satagopan, Director of Product and Innovation at NEWSWELL; and Angie Schulke, Chief Product and Technology Officer at the Minnesota Star Tribune

Photos by Hannah Yoon

AI can strengthen the business of local news

At The Lenfest Institute for Journalism, we explore how technology can strengthen the business of local news and empower the public-service impact of journalists without undermining either. That has been core to our mission since the Institute's founding, and it is the central goal of our work with artificial intelligence.

This Impact Report details lessons from the first two years of the Collaborative and Fellowship Program. The program was supported by OpenAI and Microsoft, which provided direct funding for each newsroom to hire an engineer, along with software credits and engineering support to enable experimentation. Each participating newsroom independently set its own priorities. Editorial decisions stay with the journalists, and the program is designed to protect the human judgment, editorial integrity, and trust at the heart of local journalism.

In September, The Lenfest Institute announced a renewed commitment from OpenAI that will allow more news organizations to enter the program, more experimentation, and more resource sharing to help the broader industry implement lessons from the Collaborative.

We hope insights from the Collaborative and Fellowship along with additional programs can give independent local news organizations room to experiment, champion promising ideas, help journalists and practitioners understand and harness the technology, and develop the infrastructure they rely on — all to support strong, diverse news enterprises that can meet consumers' news and information needs over the long term.

What We Stand For

We believe AI presents an opportunity to help enable a renaissance in local news and information, but that opportunity is by no means guaranteed. It depends on the choices and commitments news organizations, funders, technology companies, and the people building and using these tools make along the way.

Journalists have a duty of care to seek the truth and report independently while minimizing harm. News organizations should apply the same principles to how they approach new technologies, including artificial intelligence.

We treat AI as a broad category of technology, not as a single product. Useful applications may involve large language models, coding assistants, data tools, machine learning, automation, or systems that do not use generative AI at all once they are built.

Here are the 5 principles behind our approach to AI in local news:

1. Start with journalism and the people it serves

News organizations should begin with their purpose: What does their community need from them? What problems are worth solving? Where could technology help?

When AI is used in service of accuracy, independence, accountability, human judgment, and community, we can leverage the opportunities it presents rather than treat it primarily as a threat.

2. Give the power of choice to news organizations

We design our programs so that our grantee partners have latitude to decide which problems are worth prioritizing and solving, and which tools are right for the job.

We believe user needs should drive news organizations' work. The Lenfest Institute and companies that co-fund the work with us do not prescribe what gets built or which technologies are used.

3. Be ambitious about what AI could make possible

Other industries are putting serious money and talent against AI. Journalism should take its own AI future just as seriously. We are interested in a broad range of experiments, from practical workflow improvements to uses of AI in journalism that haven't been invented yet.

AI makes it relatively easy for people to try things that once seemed impractical or have long been on their technology wish lists.

Having an expert guide the process, along with strong collaboration, makes that experimentation faster and safer and increases the chances of learning what works and what does not.

4. Opportunity and guardrails have to develop together

News organizations should establish clear, transparent rules for where AI may be used, how it is constrained, and how and when its use is disclosed. They should not use AI to justify demoting reporters or relaxing journalistic standards.

Responsible experimentation requires both a good offense and a strong defense: exploring what the technology can make possible while protecting journalistic integrity, sensitive information, intellectual property, and public trust.

5. Journalism should shape its own AI future

Dismissing or underestimating the change happening around us is a major risk for journalism — indeed for any business. The Internet transition of the 1990s taught the news industry some very big lessons about arriving late and underinvesting in technological change.

We will work with technology companies, foundations, publishers, researchers, users, and others who can contribute capital, expertise, and ideas. We do not advocate for any one company's model, product, or worldview. We seek — and have found — like-minded partners who share our principles and from whom we can learn. We are eager to partner and to learn more.

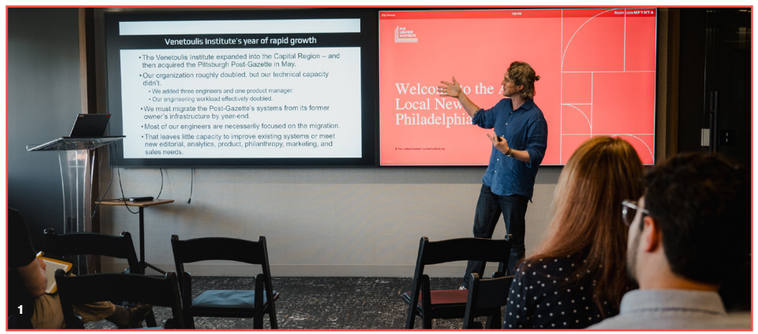
Please join us.

A front-row view of generative AI at news organizations

BY SONALI VERMA

The Lenfest AI Collaborative and Fellowship Program has given AI practitioners a front-row view of how generative AI is reshaping work, roles, and culture at news organizations. Our AI Fellows, embedded in newsrooms and product teams, have been catalysts for transformative change over the past couple of years.

When they started in late 2024, OpenAI's GPT-4o was the frontier model. Vibe coding and agents were still mostly aspirational, and conversations in journalism focused heavily on copyright and training data.



1

In the two years since the start of the fellowship, AI has become far more capable, far easier for non-technical staff to use, and increasingly woven into everyday editorial and business workflows.

How has this evolution affected news organizations – and the work that our AI Fellows do? Here are some key learnings from Fellows Rajesh Barade of The Seattle Times, Mark Chonofsky of Chicago Public Media and Kevin Hoffman of The Philadelphia Inquirer, shared at a gathering to mark the program's second anniversary.

Role shift: from engineer to coach

The "AI person," who was once seen simply as a coder, has increasingly become someone who empowers and enables colleagues to use AI.

"It started off as somebody who could help users automate something that they are doing," said Rajesh. "Now, it's shifting more toward a coach-type role, who helps people accelerate their journey, as opposed to somebody sitting in a corner building something."

Education and training were vital along the way, said Kevin. "We embarked on a teaching journey of, 'This is what AI looks like, this is how it can help you, this is how you might not want to use it.'"

Culture and trust are still stronger than technology

Trust-building and empathy are prerequisites for creating useful AI tools, each of the Fellows said.

"People are used to getting software systems that work for some startup in Silicon Valley, but don't really work for them in the news environment for their specific job, and so there's a lot of just sitting with people and fixing whatever they have," said Mark.

"It's just being genuine friends with people. Expressing real interest in what they do goes a long way."

Kevin, who initially approached the newsroom with caution, found that working with journalists was an eye-opener.

"I quickly realized that I had the best user-testers of all time," he said. "Journalists are naturally curious and ask questions. So, when I wanted to have people test out things I was building and stress-test different prototypes, it was the exact group who would want to test it. They knew exactly how to break it."

And if you build something genuinely useful, your journalists can also be your best ambassadors, he said.

"It feels like once we got some buy-in from the newsroom, they bought each other in, and there was a lot of conversation amongst each other: 'This is actually a really cool tool, I swear. Just give it a shot,' and suddenly, there magically were more people trying out some of our tools," Kevin said.

Mark agreed, citing literature that showed "if you want to get one of these interventions to stick, you have to train not one person but two people so they can talk about it. There's something deep to that."

AI is easy to work with, but the right user experience matters

The real difficulty does not lie in using AI. It is in integrating AI well into people's work.

"I think the industry likes to tell itself that doing AI is intrinsically difficult. That it's very much like computer programming, and I think that the real defining characteristic of these tools is that they are dramatically easier to use than any other computer system that we have ever encountered," said Mark.

"So the real work is around sitting with users and ensuring that they have the opportunity to do the kind of work that feels intuitive to them."

Democratization: everyone can be a builder

Unlike in previous eras of technological change, non-technical staff can now build useful tools themselves. Anyone can simply instruct an AI coding app in their native language to generate code, solving technical problems in real time.

"English is the new programming language," Rajesh said.

Kevin pointed to the rise of vibe-coding in the Inquirer newsroom. "We've really seen it lower the barrier to entry to building some of these more technical tools that might have

"I quickly realized that I had the best user-testers of all time. Journalists are naturally curious and ask questions. So, when I wanted to have people test out things I was building and stress-test different prototypes, it was the exact group who would want to test it. They knew exactly how to break it."

otherwise been launched by product teams or more technical experts," he said. In the past, it often took weeks or months before a tool could be deployed.

AI must serve real business and editorial needs

AI is only good when it solves real problems.

"We're not an AI organization, so we need to have the patience, the time, the judgment, the discernment to figure out whether it solves a business need. If so, we should do it, and if not, we should do something else," said Mark.

And even as AI proves itself useful across different parts of the news operation, human sensitivity and experience remain irreplaceable, Kevin said.


3


2

"At the end of the day, we're still an institution that relies on human judgment," he said. "That's our brand. Whatever we're building cannot interfere with those sorts of operations."

Shifting bottlenecks

What is hard and what is easy keep changing as the technology and user attitudes rapidly evolve.

"The bottlenecks are shifting – the hard things become easier, and easier things become harder," said Rajesh.

For example, initial hesitation from non-technical colleagues around using AI has now given way to vibe-coding, he said.

"When I look at adoption, there's a shift in perspective. It's proceeding at a much faster pace. People can totally vibe-code an entire application. That gives them huge confidence."

The new "hard" thing? Governance, in organizations where everyone has access to AI tools, and many are swiftly building their own applications.

"The level of curiosity has kind of shifted to the other side," Rajesh said. "I frequently get requests saying, 'Can we use this connector? Can we connect to this tool?' We have to think really long and hard. Being able to figure out how to structure the data and the tools that people use is going to be a big thing. I think that's also going to be a really difficult thing to do, considering the very thin layer of AI understanding across an organization."

1. Baltimore Banner Principal AI Engineer Peter Rasmussen **2.** ProPublica AI Engineering Fellow Aaron Brezel **3.** Fellows at the Lenfest Institute AI Collaborative convening in Philadelphia in August 2026



Curiosity beat technical skill

The Philadelphia Inquirer didn't appoint AI champions. It made tools easy to find and use.

BY SONALI VERMA

The Problem They Set Out to Solve

The Philadelphia Inquirer set out to harness AI in ways that genuinely helped journalists and business teams, without undermining trust, skills, or culture.

Shortly after ChatGPT's release in late 2022, leaders saw real potential in large language models (LLMs) for archives, local coverage, research, and analytics. But they also faced strong skepticism around intellectual property protection, environmental impact, quality of writing, and job security.

As Chief Product and Technology Officer Matt Boggie put it, the goal was to "support the news judgment of our teams," not to "automate all the way out to the user."

What worked

1. Start with real newsroom problems, not "AI for AI's sake"

From the beginning, the focus was on long-standing, concrete needs. One of the first projects was a retrieval-augmented generation (RAG) application over decades of archives, enabling reporters to quickly check historical coverage and context.

2. Champions self-select; curiosity beats technical skill

The team didn't appoint "AI ambassadors" from the top. Instead, they made access to AI tools visible and low-friction.

They looked for AI champions in the organization, which was "actually really easy, because they raise their hands," Matt said.

"The way that we identify those champions is to be very public and open about the tools that we have available, to say these are now things that you can use and here's how you ask for access. This is how the costs will be maintained, so that reporters and others who are experimenting know that this is something that they can do without having to ask for budget.

"And once we do that, we get a handful of people raising their hands and saying, 'I'd love to try this out. I really want to understand how this can help me.' That's the easiest place to start because you have an audience that's very committed to understanding how the technology works and how they can apply it."

Their champions were selected not for their technical expertise but for their mindset – people with curiosity and enthusiasm to try something new.

Once their ambassadors had built useful tools, peer-to-peer demos drove adoption far better than central evangelism ever could, he said.

Their ambassadors would turn to the person at the desk next to them and say, effectively, "'Look at this great thing I made. It's not theoretical, and it's not hype. It is something that literally will change the way that you do your work immediately,'" Matt said.

3. Light training plus trust in people

Rather than heavy, tool-specific training, The Inquirer focused on principles: data security, how LLMs work, and basic computational thinking.

"We talk about some of the principles behind AI, whether it's data security, whether it's LLMs and chunking and computation, just so that people have a sense of the language they're walking into, and then we give guideposts for how to get started and a place to come and ask questions once you get into it," Matt said. Tool-specific training is done in smaller groups or on one's own, as it's needed.

Crucially, they trusted staff to experiment responsibly, with human review as a non-negotiable safeguard.

Rather than heavy, tool-specific training, The Inquirer focused on principles: data security, how LLMs work, and basic computational thinking.

Success metrics

The primary success metric was deliberately simple, Matt said: utility.

This showed up in three ways:

- **Adoption across desks:** Tools like their RAG archive, local meeting scanners and court docket scrapers spread organically once editorial staff showed their colleagues how they achieved real gains in speed, quality and depth.
- **New kinds of coverage:** AI-assisted analysis helped uncover what really resonated with readers. For example, celebrity stories that combine the person and a distinctly local angle work better than celebrity stories alone.
- **Responsible cost control:** They tracked token usage closely, aiming to "spend[] appropriately for the value that we're getting back," Matt said.

Importantly, success was never just about AI use (e.g., tokens burned or number of prompts). Matt explicitly rejected that mindset.

"I have never understood the mandates, to be fully blunt about it, to use as many tokens as possible," he said.

"We want to be sure that we're being responsible, and so rather than try to push adoption for adoption's sake, our focus was on finding the places where this can be the most useful the most quickly, to show that there's real value, and to make sure that we're spending appropriately for the value that we're getting back from that product."

The key question was always: Does this tool help staff do better journalism and better work, faster?

What they'd do differently

Looking back on experimentation, Matt is clear: They would move faster from theory to practice.

"I think I would have found a way to get the tools in people's hands faster," he said, citing the impact that he has seen.

Early on, much time was spent on policies and high-level debates at a time when capabilities were still largely theoretical. With hindsight, this slowed learning and consensus-building.

"What we found then is that a lot of the concerns and objections do kind of fall away once you can identify the utility that comes with a tool," he said. Still, starting by listening to people's excitement and fear, and developing guidelines for experimentation, helped some feel more comfortable, he pointed out.

Another key learning is to resist perfectionism in product definition. Matt sees people over-engineering prompts and specifications before they've even tested whether an idea is viable.

"I work with people a lot who will spend a day or two organizing their thoughts, thinking through how to break the problem down. And the advice I give everyone is just prompt with a two- or three-sentence description of the thing you're trying to do.

"Don't let it build anything yet. Just ask some questions, see what questions it has, help it shape the product that you're looking for, and then as quickly as you can, find all the things that you forgot to ask about the first time. Because no matter how thorough you think you're being about the way the product is going to interact, as soon as you're looking at the interface, you'll go, 'Oh, that button should be over there,' or 'I really need a control here that does something.' This is an extremely iterative process."

His guidance is to avoid writing the perfect two- or three-page prompt. Instead:

- Start with a tiny, rough prompt.
- Let the model ask you questions.
- Iterate quickly, and be ready to "throw out" a weird solution and restart.

The Inquirer has also learned that human review is non-negotiable:

"No matter how many times we will run something through a system, every now and then it will get a story wrong. It will hallucinate a summary, and you'll want to be sure that you can catch those," Matt said.

"What we've found is that with some light training and a little bit of reinforcement, people know what they're doing."

This autonomy, with clear principles, produced a better culture and better tools.

Here's how you can apply this

If you're starting (or rebooting) your AI culture change, you can adapt the Inquirer's approach:

- **Anchor "official" AI efforts to two or three stubborn problems.** If you have the capacity for a more formal, funded project to jumpstart adoption, pick from long-standing pain points, such as archive search, meeting summarization, local information gathering, survey analysis, or basic analytics queries. Frame AI as a tool to fix these, not as a broad "transformation."
- **Let champions raise their hands to experiment:** "These tools are available now; here's how to get access." Prioritize curiosity and enthusiasm over technical skill. Give them light guardrails, then get out of the way. If possible, equip these champions with story archive datasets, article performance metrics, and other internal sources they can parse and search.
- **Keep training short, principled, and repeatable.** Offer brief sessions on data security, what LLMs are, and examples of good newsroom use to prevent the worst abuses and mistakes. Provide an always-open space to ask questions instead of one-off marathons that quickly go stale.
- **Measure utility, not volume**

  Track:

  - Which tools are being reused across desks
  - Time saved on specific workflows
  - New coverage or insights you couldn't realistically produce before
  - Combine that with cost monitoring and model-routing so you can say, "This is worth it," in both journalistic and financial terms.

Above all, borrow Matt's mantra: Start small, start now, and iterate with your colleagues.



1. Sonali Verma with The Philadelphia Inquirer Chief Product and Technology Officer Matt Boggie, AI Engineer Kevin Hoffman, and Chief Information Officer George Kurtas **2.** Kevin Hoffman and Seattle Times Senior AI Engineer Rajesh Barade

How local news organizations can adopt AI while maintaining trust, judgment and their sense of purpose.

AI in Local News:
A Practical List

Artificial intelligence is rapidly becoming part of every newsroom. The question for local news organizations is no longer simply whether to use AI, but how to use it well.

We asked the Lenfest AI Fellows and their managers to help us put together a playbook for local news organizations, based on their experiences and learnings over the past two years. They came up with practical suggestions for what to do – and also strong recommendations on what not to do.

One underlying theme: Local news organizations should approach AI as an organizational capability, not simply a collection of tools. The goal is to use AI where it helps the organization do better journalism, serve its community and build a sustainable business.

Do this

"Keep a human in the lead, not just in the loop." "Intentionally set human intervention points."

Set policies, guardrails and governance

- Develop clear, flexible policies and standards on AI use.
- Create an equal balance between guidance and guardrails.
- Establish the right way to get data sources as soon as possible. Create clear governance processes – and share transparently with your teams.
- Default to approving requests early on to learn about alternative platforms. Try a task with different models, since each has its own strengths and weaknesses.

"Over-specify at the beginning. An extra hour spent planning and articulating pays huge dividends."

- Set realistic goals and expectations for reader-facing AI projects. "It's not magic," one participant added.

Culture

- Work on the culture first.
- Establish clear communication with stakeholders.
- Find champions internally. Identify people in each department who would be happy to share knowledge and support their peers.
- Require top leaders to lean in.
- Meet people where they are.
- Try vibe-coding with anyone in your organization.
- Liaise with skeptics – and being kind and open to them.

Experiment and share

- Encourage experimentation and sharing. Set aside time to explore different AI tools, and "be open to creative chaos" that lets staff explore. Get leaders involved in experimentation.
- Start small, get quick wins early, and show wins. Any learning is a win. Experiment in small, iterative steps so that you can learn quickly from results.
- Create a process to share use cases regularly, like a demo camp every six weeks. "Nudge and prod people to share what they have built, especially if they were working alone."
- "Moving from prototype to a scalable tool is the hardest part. Account for that boring part."
- "Get the newsroom involved early and enthusiastically rooting for your AI project."
- Collaborate with other organizations on similar projects.

Where to use AI

- "Drive and optimize for business outcomes and team priorities – not from a 'how should I tackle this with AI' perspective." Identify a genuine business need first and then ask whether using AI is the best way forward.

"Start with the problem, not AI...Don't just sprinkle magical AI dust on a problem to say that you used AI. I call that performative innovation. It drives me crazy."

CHASE DAVIS

- "Let AI do the drudgery." Let humans do the value-added work that requires sensitivity and judgment.
- AI works best "if you point it at a tight corpus of your own reporting. When we just ask it to summarize strictly within the bounds of the article, our testing showed that it was roughly 97% accurate." LLMs can be highly accurate when supplied with high-quality context, such as reporting, documents, or other materials, as opposed to only asking them to reason over what's in their training data.

Don't do this

"Don't delay getting your AI experiments in front of real readers. The perfect is the enemy of the good."

Don't build AI tools in isolation

- Don't silo product and tech from the newsroom on AI projects. Engineering in a vacuum is a problem. Let your AI person work their way through your organization, enable and empower staff, and build relationships to ensure that their tools actually create an impact. Co-creation is ideal to develop the best solution and to get user buy-in.
- Don't give staff access to new tools without some training or guidance. "We have a responsibility of making sure that nobody is left behind in terms of AI skill. We need to be focusing on training them and getting them to some level where they are comfortable."
- "You shouldn't assume you know where and how adoption is happening."
- "Don't make employees feel unsafe by tracking and logging every single activity."

Don't make technology choices without a strategy

- Don't treat AI companies as your friend – or your enemy. Evaluate them objectively and repeatedly.
- "Don't jump to vendor point solutions. The technology is evolving too quickly."
- "Don't just use the most powerful models."
- "Don't take overwritten code as gospel. Sometimes, saying 'no, not like that,' is the best move."
- Don't overestimate the capabilities of AI or remove human judgment or let the AI think for you.

Don't wait for certainty before experimenting

- Don't expect a clear, demonstrable ROI right away. Don't build without cost and revenue metrics and impact analysis in mind, but do remember: Adoption takes time, and using the tools is often the best way to find out whether they are genuinely useful or not.
- "Don't treat every problem as one size fits all." Don't automate before you understand the problem.

Don't let governance become bureaucracy

- Don't create "iron laws," such as ensuring a human must be in the loop 100 percent of the time. Acknowledge that the level of risk differs according to the AI use case – for example, between internal newsroom assistance and producing content for public consumption. Be prepared to revisit your AI policies and guidelines as the technology evolves.
- Don't over constrain a model or apply guardrails and constraints too soon. You risk putting a freeze on experimentation without learning what the machine can actually do.
- "Do not make it harder for employees to learn, experiment, and explore AI. It is a much-needed skill for their future." As one participant said: "I think it's absolutely our responsibility as leaders to bring people along on these changes, even when it's uncomfortable. That's the job."

Meet the AI Fellows

Two years ago, The Lenfest Institute for Journalism gave a group of independent news organizations resources to hire full-time technical talent, access AI platforms and API credits, and explore what the technology could make possible.

We deliberately gave them broad latitude. Fellows could work across audience, revenue, IT, marketing, or editorial. They could use OpenAI, Microsoft, other providers, local language models, or tools they built themselves. They could change direction as they learned.

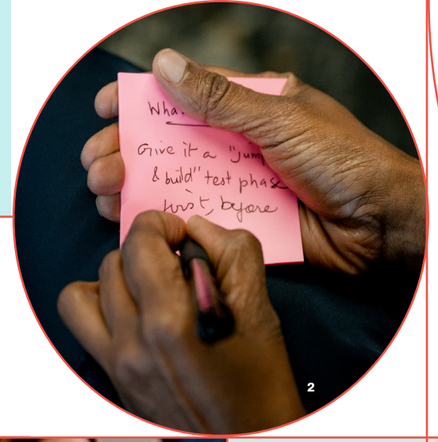
This effort has created a clear early picture of where the technology can be useful, what news organizations need to use it responsibly, and what it takes to move from a single embedded technologist to broader, responsible adoption within an organization.

Over the following pages, the Fellows will share more about their work, discussing the projects they built, the lessons they learned, and the trends they're following.

We've been inspired by the Fellows, and we hope their insights will similarly spark your imagination about what's possible.



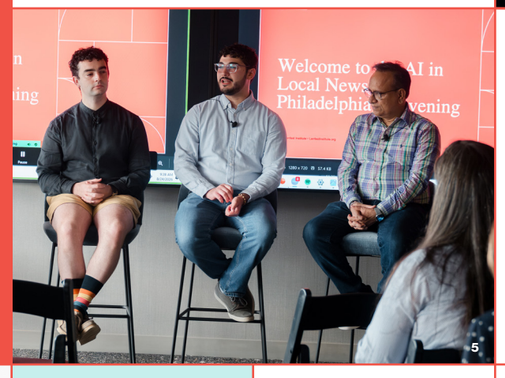
1



2


3


4



5


6


7


8

1. ProPublica's Aaron Brezel presents to the Fellows **2.** A brainstorming session at the AI Collaborative convening **3.** Microsoft Journalism Director Emily Brown **4.** OpenAI's Evan Hirsch **5.** Chicago Public Media AI Fellow Mark Chonofsky, Philadelphia Inquirer AI Engineer Kevin Hoffman, and Seattle Times Senior AI Engineer Rajesh Barade **6.** The group reviews the Inquirer's Pulitzer Prizes during a newsroom tour led by Philadelphia Inquirer Editor and Vice President Gabriel Escobar **7.** Dana Chiueh, ProPublica AI Fellow **8.** Adam Hirsch, ProPublica Director of Product Engineering

Remembering David Chivers



David Chivers speaking at a 2025 AI Collaborative and Fellowship event

I first met David in 2024, as the Lenfest AI Collaborative and Fellowship Program was taking shape. In the short time I had the chance to work with him, I was often amazed by how easily he turned new connections into friendships—and by how many people, including me, came to see him as a close friend.

David brought a range of experience to the Fellowship as a publisher, product leader, and optimistic advocate for thoughtful uses of AI in journalism. But what I remember most is how he worked with people. He was curious, generous, and warm. I never saw him get caught up in frustration or ego. Even when the work was difficult, David seemed to keep a higher value in mind: the humans involved, the mission of journalism, and what we might accomplish together.

That spirit shaped the Fellowship. David believed in bringing people together, giving them room to experiment, and sharing what they learned so others could benefit. He treated people across organizations and roles as peers, and I've rarely seen anyone with his gift for making a group feel like a community.

David helped build this program, but he also embodied its name. He practiced fellowship—in generosity, curiosity, empathy, and a belief that we could do more together than alone. Since his passing, I've believed that one way we can honor David is to carry that practice forward in our own work.

Tristan Loper
Head of National Programs
The Lenfest Institute for Journalism



FELLOW SPOTLIGHT

As an AI Fellow at The Seattle Times, Rajesh Barade is building practical AI solutions across advertising and other teams while helping staff explore and experiment with AI.

Rajesh Barade

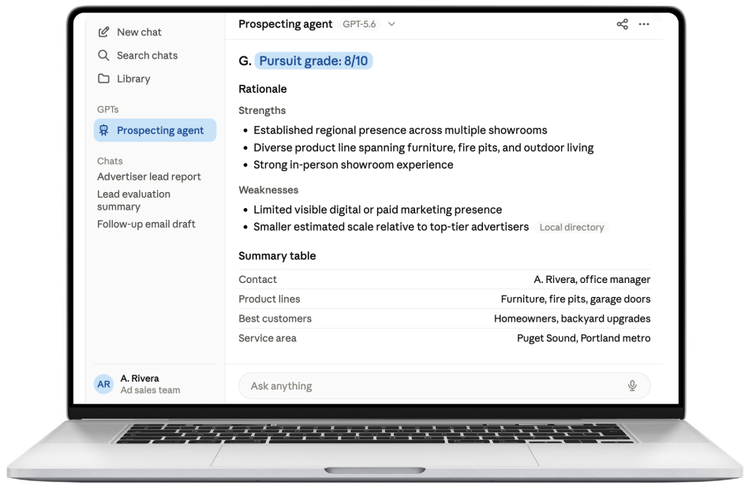
The Seattle Times

Many in the industry see AI as disrupting the work people do today. Rajesh sees it a little differently. While AI may initially automate some human tasks, he believes it will also create new roles and opportunities for people. He draws a parallel with computers: Before computers became part of our everyday lives, it would have been hard to imagine many of the roles and jobs that exist today because of them.

About his role and the problem he is trying to solve

Rajesh describes himself as "a techie who always liked to build solutions hands-on," with a career spanning over two decades in technology and financial services.

He characterizes his role at The Seattle Times as "a kind of forward-deployed engineer before that term was coined." His goal is to run useful AI experiments and be present to assist colleagues when someone needs help with AI.



Examples of the leads generated by The Seattle Times ad sales prospecting agent

Instead of solving a single, tightly defined problem, he works as an internal consultant across teams.

For the advertising team, Rajesh has built:

A lead-generation custom GPT-style agent to support prospecting, now being adapted and shared with other organizations in the Lenfest AI Collaborative cohort.

A set of copilots over SharePoint and internal data, which he envisions evolving into a collection of reusable "skills" that advertising teams across news organizations can adopt and customize.

The biggest surprise during the fellowship

Rajesh has been struck by the gap between the scale of the AI shift and its potential applications in many areas of the organization.

He hears deep concern about existential threats such as zero-click search and new AI-enabled competitors. He believes AI adoption across news media organizations will catch up with shifting industry trends, and he sees emerging chat and voice channels as new ways to reach audiences.

The advertising side of the business has been the forerunner in AI adoption, with teams that clearly understand how AI is disrupting their work and are moving quickly.

One lesson from his work that other newsrooms could apply tomorrow

Rajesh believes AI belongs in every employee's toolkit, and not only inside centralized, bespoke applications.

"AI is a productivity tool that has to be on your desktop, and you must be comfortable using it. When somebody in your industry is using that to run 10 times faster, if you're not using it, you are left far behind," he said.

He has been running a pilot with Claude for some non-technical team members, who are now building tools on their own.

The key, he said, is to help people turn experiments into repeatable, ROI-positive workflows – for example, using Claude Cowork or ChatGPT to design small, reusable skills for routine tasks, and making sure the token spend is paid back as meaningful time saved.

What he is paying attention to right now

Rajesh is focused on the engineering discipline around agents – how to build robust, monitorable harnesses around LLMs so they behave reliably in production.

He has been advocating for agents for months now. For him, the frontier is not just clever prompts, but observable, governed agentic systems that organizations can trust over time.



"AI is a productivity tool that has to be on your desktop and you must be comfortable using this. When somebody in your industry is using that to run 10 times faster, and if you're not using it, you are left far behind."

The Fast Three

1. **A book, article, or newsletter you'd recommend?**
Neuron Daily, a newsletter he relies on to keep up with AI developments.

2. **What's one thing you're learning that has nothing to do with AI?**
Improving his chess.

3. **One word that describes the next year of AI in news?**
Voice, an upgrade from the chat channel – "the convergence of memory, personalization and speech interfaces."



FELLOW SPOTLIGHT

Aaron Brezel ProPublica

As an AI Engineering Fellow at ProPublica, Aaron Brezel is designing human-centered AI tools that help investigative reporters work more efficiently, without sacrificing rigor or agency. We spoke with Aaron about what he's learned so far, what the industry gets wrong about AI, and why he believes human judgment will matter more than ever in the years ahead.

> **"AI is a means to an end. You really have to understand where and how you're actually applying AI."**

About his role and the problem he is trying to solve

Aaron is one of two AI engineering fellows at ProPublica, alongside his colleague Dana Chiueh. His work is squarely focused on the investigative process, rather than on audience- or business-side tooling.

Aaron's mandate is to work in the newsroom and support reporters directly in producing journalism. With a master's degree in journalism and in engineering from Columbia University, he draws heavily on his prior experience at The Washington Post, where he worked on the Reporting Tools team.

"We were a group of engineers with engineering backgrounds," he said. "Our remit was to collaborate within the newsroom and understand the common problems that they were experiencing and build durable systems to support them in the future."

At ProPublica, he has taken the same approach: embed, observe, collaborate, and then build.

Reporters come to him with "AI-shaped problems" – massive document sets, messy data, and opaque systems, for example. He helps them break those challenges down into workable processes.

Over time, he's identified repeating patterns and is now building at least one core product to support the most common use cases – an open-source Google Sheets sidebar that lets reporters run structured, row-by-row AI operations using Gemini and related tools for tasks like document OCR and analysis, directly inside the spreadsheet environment that they already work in.

What has surprised him most during the fellowship

"What surprised me most has been the eagerness that reporters have displayed towards using AI tools," Aaron said. "There's a lot of general consternation about AI, but I think it's actually directed more towards AI mediation between a reporter and their reader. There's a lot less anxiety around reporters using the tool for their own purposes to conduct investigations."

In other words, reporters are wary of AI as a publisher, distributor or gatekeeper, but enthusiastic about AI as a backstage assistant that helps them uncover stories faster and more deeply.

"I get more requests than I can deal with, which is a good thing as far as I'm concerned," he said.



Fellows on a newsroom tour led by Philadelphia Inquirer Editor and Vice President Gabriel Escobar

> **"Our remit was to collaborate within the newsroom and understand the common problems that they were experiencing and build durable systems to support them in the future."**

What the industry gets wrong about AI

For Aaron, the biggest misconception is treating "AI" as if it were a solution in itself.

"The thing that people get wrong is that AI in itself is not a solution to anything. It's a model. It takes an input and spits out output."

He argues that the fundamentals of good product thinking and reporting support haven't changed: Understand real user problems, iterate through solutions, and focus on the minimum valuable product to build.

What has changed is the speed and volatility of the environment

"We just have to do the same things that we've always been doing. Unfortunately, now we just have to do it 10 times faster because everything is moving so quickly," he pointed out.

He's particularly critical of AI products that over-promise "magic" and encourage reporters to hand over too much of their process. "AI is a means to an end. You really have to understand where and how you're actually applying AI."

One lesson from his work that other newsrooms could apply tomorrow

Build direct, human relationships between technical staff and reporters, and treat the job as enabling reporting, not just shipping software.

"I am a big proponent of direct relationships with the people you're trying to build for. There's always value to just being a human resource in the newsroom that reporters feel comfortable talking to, and who prioritizes the same things that the reporters do."

For the first four or five months of his fellowship, Aaron simply listened carefully to what journalists needed, rather than rushing to define a product.

"I don't see my job as producing this tool. I see my job as producing reporting," he said. "I just went to reporters and said, 'I'm here to help you produce stories using AI.'"

That unstructured collaboration became the engine for the product he's now building.

What he is paying attention to right now

Aaron is closely watching how human-computer interaction is being reshaped by AI, especially the shift from code to natural-language interfaces.

His Google Sheets add-on for reporters is a concrete expression of that idea: Instead of writing code, reporters write prompts that operate over rows in a spreadsheet.

At the same time, he is focused on preserving agency and structure in reporting workflows, rather than handing everything to a black-box agent.

"We don't want a chatbot to just do everything for you. We want you to import your documents and maybe use a regular old Google Sheets function, and then the one little thing that AI needs to do, that a function can't do, that's where you apply it," he said.

What he is most optimistic about

Despite the turbulence in the news business, Aaron is optimistic about the enduring value of human judgment in journalism.

"I'm optimistic that there are a lot of really smart, really determined people working on these problems," he said. "I'm optimistic that the news industry is putting a lot of resources towards answering those questions."

He believes there will undoubtedly be a market for automated news, such as hyperlocal weather reports. But he doesn't see that eliminating the need for human reporters.

"There's still a lot of demonstrated value in a real human telling you what's important, a person who has dedicated their life to covering a subject, helping you understand these complicated systems that govern our world. That voice is still very important, and I see nothing yet that has obviated that."



The Fast Four

1
A book, article, or newsletter you'd recommend?
The writing of Geoffrey Litt, especially his essay comparing AI to teleportation.

2
What's one thing you're learning that has nothing to do with AI?
New baking techniques.

3
What local journalism challenge deserves more attention?
Direct connection and human voice in an age of information overload and AI mediation.

4
One word that describes the next year of AI in news?
"Agency—the human kind. The more you can just use AI to do anything, the more important human agency becomes in making decisions about what you spend your time on, what you think is important, what you think should be communicated to a reader."

As an AI Engineering Fellow at ProPublica, Dana Chiueh is exploring how artificial intelligence can help investigative newsrooms better manage and act on information from their communities. We spoke with Dana about the lessons she's learned so far, what excites her most about AI, and where she believes the industry is headed next.



FELLOW SPOTLIGHT

Dana Chiueh ProPublica

"I think the implementation of LLMs should only reinforce the argument for having more journalists on the ground," she said, pointing towards a misconception that LLMs can take over journalists' work.

About her role and the problem she is trying to solve

Dana's background bridges data reporting and computer science, a combination that naturally led her into AI-focused roles in newsrooms. After studying at Oxford and Stanford, she is currently an AI Engineering Fellow at ProPublica and previously held a similar role at The Minnesota Star Tribune as part of the Lenfest AI Collaborative and Fellowship.

Her core mandate at ProPublica is to explore how machine learning, AI, and other emerging technologies can be applied to engagement reporting. In practice, that means tackling a hard, very real newsroom problem: How to manage and honor the wealth of information that pours in from the public.

Newsrooms, particularly investigative ones, receive a large volume of tips, messages, and source outreach. Some of these are sensitive and need to be handled with privacy in mind. Sorting, prioritizing, and maintaining relationships with these sources at scale is difficult:

- It's hard to process and triage a high volume of incoming tips.
- It's hard to make editorial decisions about which tips to pursue when information is incomplete.
- It's hard to sustain relationships over time with so many people writing in.

"A big part of my mandate is just to figure out: Can artificial intelligence help us with this set of problems?" Dana said.

She is also working more broadly on internal reporter tools, exploring ways to thoughtfully introduce generative AI and machine learning into daily workflows inside the newsroom.

What has surprised her most during the fellowship

When Dana looks back over the past few years, what stands out is how quickly a community around AI and journalism has sprung up.

During her master's at Stanford, when she was already exploring the intersection of AI and journalism and building newsroom software, that area felt "pretty nascent." There were only a handful of people in the field, and use cases for tools like ChatGPT were still being worked out.

Since then, she has seen the emergence of:

- Pop-up communities of practice
- Events and gatherings focused on AI in journalism
- A much larger pool of people working on similar problems across newsrooms
- She notes that in the second year of the fellowship, it feels like there is "a lot more wind in our sails," with more momentum, collaboration, and funding supporting this work.

At the same time, that growth brings a new challenge: bridging the divide between editorial and product/technical practitioners working with AI. Dana is especially interested in how to foster stronger technical collaboration across newsrooms, so that people aren't solving the same problems in isolation.

What the industry gets wrong about AI

From Dana's perspective, there's a misunderstanding about what AI is actually good at and, importantly, what it is not a substitute for.

She sees a clear distinction between:

- AI's strength in "repeating the past," such as summarizing existing information, surfacing historical context, and accelerating research.
- Journalism's core mission of finding new facts and being deeply embedded in communities.

"I think the implementation of LLMs should only reinforce the argument for having more journalists on the ground," she said, pointing towards a misconception that LLMs can take over journalists' work.

What makes journalism unique is human reporting, newsgathering, knocking on doors, and discovering what's happening right now. AI can accelerate research and context-building, but it cannot replace human relationships and on-the-ground reporting, she stressed.

One lesson from her work that other newsrooms could apply tomorrow

A key lesson Dana has learned is the importance of involving stakeholders early and deeply in AI product development.

AI products, in her experience, are inherently collaborative. They often need:

- Editorial partners to help define the problem precisely.
- Journalists to label data, classify examples, or manually perform tasks once so that their judgment can be modeled.
- Ongoing feedback to refine and evaluate tools.

She suggests that newsrooms:

- Use AI to quickly build very lo-fi or wireframe prototypes.
- Put these early versions in front of stakeholders to spark imagination and concrete feedback.
- Make it clear from the outset that AI tools will require significant editorial involvement, not just a handoff to product and engineering for a few months.

In other words, treat AI tools as co-created editorial products, not standalone technical projects.

What she is paying attention to right now

One emerging area that fascinates Dana is adaptive interviewing.

She's been following research on using adaptive interviewing for areas like UX research, user testing, market research, and polling, and has also seen it applied in the context of journalism.

It involves question flows that change based on a person's previous answers, making forms or bots more personalized, dynamic, and responsive to user preferences.

She is intrigued by what this might mean for reporting and engagement at scale, while also recognizing the existential questions it raises about what it means to interview and connect with someone in journalism.

"When I think about how we manage relationships at scale, it's certainly a compelling idea that piques my curiosity," she said.

What she is most optimistic about

Dana is optimistic about the future of shared, collaborative AI infrastructure for newsrooms.



The Fast Four

1. **A book, article, or newsletter you'd recommend?**
The Reboot newsletter.

2. **What's one thing you're learning that has nothing to do with AI?**
Training for a triathlon.

3. **What local journalism challenge deserves more attention?**
News deserts.

4. **One word that describes the next year of AI in news?**
Collaboration.

She sees growing recognition that:

- AI can significantly help with newsroom workflows.
- There is a need for technology that outlives specific time-limited engagements.
- Third-party organizations, such as foundations or industry associations, could help maintain and distribute shared tools across multiple newsrooms.

She believes there is a real opportunity for industry-wide collaboration, where tools developed in one newsroom are sustained and adapted for others, moving beyond one-off open-sourcing toward durable, shared systems that benefit the whole ecosystem.



FELLOW SPOTLIGHT

As a Lenfest Institute AI Fellow at Chicago Public Media, Mark Chonofsky is exploring how AI can expand access to local news, improve archival audio, and surface journalism more effectively. We spoke with Mark about the lessons he's learned so far, what excites him most about AI, and where he believes the industry is headed next.

Mark Chonofsky

Chicago Public Media

> **"Incredibly powerful capabilities that are here or are coming ... we need somebody who knows about them, and then can help the organization make informed choices."**

About his role and the problem he is trying to solve

Mark started at Chicago Public Media almost two years ago, after working on Kamala Harris's presidential campaign and going to the U.S. Army Officer Candidate School in preparation for commission as Signal Officer in the Illinois Army National Guard. He has a DPhil from Oxford, along with two degrees in physics from Cambridge.



At Chicago Public Media, his first major focus area was translation. His goal: To expand access to news in Spanish and dramatically improve turnaround times for Spanish-language coverage.

He and his colleagues built workflows that allow Spanish-language journalists and editors to edit in Spanish, instead of starting translations from scratch. They have also used AI and code-generation tools to translate entire interactive features by modularizing text into structures like JSON, translating those, and then dropping them back into the interactive.

Mark is particularly proud of one example of AI translation: Chicago Public Media's coverage of the Chicago-born Pope being appointed.

"Chicago Pope was a big thing for our Spanish audiences, who are largely Catholic," Mark said. "Previously, we would have taken 12-plus hours to translate that news out into Spanish, but we had a Pope headline out in two hours, because we have this much improved process for doing it."

Mark Chonofsky speaks at the AI Collaborative convening in August 2026 alongside The Inquirer's Kevin Hoffman

The site's traffic was five times higher than that on a normal day, as a result.

Alongside translation, Mark is also working on archival-quality transcription and smarter content discovery tools that can surface both articles and transcripts to staff, building on the Philadelphia Inquirer team's fellowship product, Dewey. "We have been exploring different ways of surfacing our content internally."

He also worked on building a propensity-to-donate model that Chicago Public Media hopes to share with other not-for-profit news organizations, once it has been thoroughly tested.

The biggest surprise during the fellowship

"People are rightly worried about what kinds of things are possible and what kinds of things may happen as a result of overuse or poor use of AI, and I think there are all sorts of news organizations around the world that are doing a bad job of this or an exploitative job of this," Mark said.

"It has taken me some time to really understand how to conceptualize the role... I'm not here to tell anybody what to do, and I'm certainly not here to replace anybody's job. I think organizationally, the issue is, we know that there are some incredibly powerful capabilities that are here or are coming, and we need somebody who knows about them, and then can help the organization make informed choices about to what extent they are useful for our work or deleterious to our mission."

He also points out that AI itself is rarely the hardest part of any project. The real challenge lies in teaching, trust, and organizational adoption.

This emphasis on teaching and judgment, rather than hype, has been a defining lesson of his fellowship.

What the industry gets wrong about AI

"We put a lot of thought into the concerns about AI replacing the journalists, and I understand why people are scared by that, but I think that anyone who spends 15 seconds with a journalist understands why that's not just unfeasible now, but unhelpful, always," Mark said.

"I really think that what this moment spurs us towards is to consider our essential humanity as news storytellers. I think that that's actually a good thing for us. I think we've lost too much of that, and we should do more of it."

Lessons that other newsrooms could apply tomorrow

Mark is skeptical of the notion that AI is moving so fast that newsrooms can't possibly keep up. He thinks that fear of being left behind is overstated.

His lesson for other newsrooms: Treat AI as a technology domain you can learn, like any other, and recognize that journalists already possess many of the skills – such as critical thinking and erudition – that are needed to understand it.

What he is paying attention to right now

Mark is particularly interested in the rise of powerful local models, systems that can be run and controlled more closely by news organizations.

What he is most optimistic about

Mark is most optimistic about AI's ability to tackle long-standing resource-intensive problems.

"AI used well has the possibility to straighten out a number of product engineering problems that have just bedeviled us for some time because we haven't been able to devote the engineering resources to fixing them," he said.

From centralizing data to fixing messy datasets and enabling smoother transitions between systems, he sees AI as a way to do crucial but unglamorous work "much faster, much cheaper."



The Fast Four

1. **A book, article, or newsletter you'd recommend?**
If you let AI do your writing, I will come to your house and kill you by Sam Kriss.

2. **What's one thing you're learning that has nothing to do with AI?**
Working with a medium-format camera and enjoying the fully mechanical, film-based process.

3. **What local journalism challenge deserves more attention?**
Consolidation.

4. **One word that describes the next year of AI in news?**
"Similar... I think there's a lot of cool stuff that people are doing, but I am not forecasting a huge amount of disruption."



FELLOW SPOTLIGHT

As a Lenfest Institute AI Fellow, Kevin Delaney, editor-in-chief and co-CEO of The San Francisco Standard, is exploring what an AI-first news interface might look like and how it can deepen reader engagement while supporting a sustainable subscription business. We spoke with Kevin about what he's learned so far, what the industry gets wrong about AI, and where he believes journalism is headed next.

Kevin Delaney

The San Francisco Standard

About his role and the problem he is trying to solve

Kevin, who previously co-founded business publication Quartz, sits at the intersection of newsroom leadership and product innovation. The Lenfest fellowship project gave him the mandate to explore a fundamental question: What would a truly AI-first news experience look like for subscribers?

The hypothesis: If The Standard could build an app that uses AI to surface the right journalism to the right people at the right time, it would both serve readers better and support subscription growth.



"Looking ahead, human journalists continue to have an important role and I see news organizations like ours wanting to hire even more of them."

The Standard's business model is built around subscriptions and sponsorship. It did not have a mobile app before this project.

"We believe that if we could create a good mobile app, it would support the regular engagement that allows you to attract and retain subscribers, and so our focus on the project quickly became: Could we use the mobile app as a way to imagine an AI-first news interface that would support our business goal of growing subscriptions?" Kevin said.

The Standard already used AI in the newsroom for help with research, monitoring news and public information sources, audio and video editing, and other workflows. This project deliberately shifted the focus to the reader interface. It is considered a Mode 2 application of AI, emphasizing learning, experimentation outside of existing workflows, and a shift toward AI-native journalism that can be fine-tuned to user needs. That contrasts with Mode 1, which focuses on optimizing existing workflows and efficiency in core operational areas to drive near-term results.

"We've seen really high engagement among subscribers," he said. "I've not been this excited about news product development in a long time."

The biggest surprise during the fellowship

Kevin's biggest surprise has been just how quickly you can now build products with AI.

The Standard received the $150,000 Lenfest Institute grant in late January 2026 and launched the first version of the app to subscribers in early April, taking roughly three months to get the live product off the ground, in partnership with Bit Bit Press, an external product and AI CMS startup.

The second surprise is how developers are using prompts to rapidly develop products.

"AI prompts are the way that we create new features in the app," Kevin said. "That's super exciting, to be able to create new product features in minutes and refine them, whereas until recently, the ability to launch a new product feature was measured in weeks and months."

For example, a single prompt can instruct the system to:

- Identify a person frequently covered by The Standard
- Create a card with their photo, name, short bio
- Add a button to explore more coverage

That ability to "ship with prompts" makes experimentation much faster and more continuous than traditional feature development cycles.

What the industry gets wrong about AI

Kevin sees two opposite but equally mistaken instincts in the industry's thinking about AI.

On one side, he believes people underestimate how much commodity news AI will be able to handle, such as the rewriting of press releases at wire services. "They've been doing this for years. AI effectively repackages them, and that trend can continue deeper into the commodity news area."

On the other side, he argues that people also underestimate the resilience and value of distinctly human journalism, such as:

- Opinion and analysis
- Narrative features
- Investigative reporting
- Local news rooted in proximity and community

"I do think that there are many areas of journalism that are pretty well protected from AI. Looking ahead, human journalists continue to have an important role, and I see news organizations like ours wanting to hire even more of them."

One lesson from his work that other newsrooms could apply tomorrow

If there's one structural lesson Kevin would recommend, it's to formally separate Mode 1 and Mode 2 work:

- Mode 1: Daily operations, such as publishing the site, running the newsroom, and hitting today's deadlines.
- Mode 2: Longer-term, strategic projects, focused on what the organization needs to look like a year or more out, especially in relation to AI.

"The ability to separate the two is really important in this moment, and to be able to think clearly about the opportunities that AI brings, about which parts of the business are positioned to thrive and which are more vulnerable," he said. "I do believe in the power of the Mode 2 project."

By treating the AI-first app as a Mode 2 initiative built with an external partner, The Standard was able to protect it from daily operational firefighting while it was getting the app off the ground. (Aspects of the app today rely on the core Standard team, which will likely take increased responsibility for it over time.)

What he is paying attention to right now

Looking ahead, Kevin is particularly focused on the emerging world of AI agents talking to AI agents, and what that might mean for journalism as a business.

He's watching:

- How news content is being made available to agents
- Whether there is a viable agent-to-agent marketplace for journalism
- If models like micropayments or microtransactions could fairly compensate publishers when agents draw on their reporting

What he is most optimistic about

Kevin is most excited about AI's ability to turn a city's journalism into practical, personalized utility for residents.

- By combining reporting with signals like location and interest profiles, the app can:
- Surface news that's hyper-locally relevant to each user
- Match people to the topics and beats they care about (e.g., Golden State Warriors fans vs. non-fans)
- Help residents navigate different aspects of their city, such as schools, events, and restaurant recommendations from trusted critics

Many of these features were theoretically possible pre-AI, but for a smaller organization, AI is a powerful enabler that dramatically reduces the resources and complexity required to deliver them.



The Fast Four

1

A book, article, or newsletter you'd recommend?
"Seeing Like a State" by James C. Scott, which examines how individuals and communities interface with large, centralized institutions. "I think it has real relevance now relative to tech companies and how to preserve the specificities of your culture and lifestyle in the face of a homogenizing global force."

2

What's one thing you're learning that has nothing to do with AI?
Learning how to play the accordion.

3

What local journalism challenge deserves more attention?
Subscription top-of-funnel.

4

One word that describes the next year of AI in news?
Deep experimentation.



FELLOW SPOTLIGHT

EJ Fox

Arizona State University
NEWSWELL

As an AI Fellow at Arizona State University's NEWSWELL, EJ Fox is rethinking how small, overstretched local newsrooms can use AI to make better newspapers with fewer resources. We spoke with EJ about what he's learned so far and what the industry gets wrong about AI.

About his role and the problem he is trying to solve

EJ's work sits squarely at the intersection of code and journalism. He describes himself as "the computer person" in the newsroom — the one who helps reporters tell their stories with data, graphics, and now, AI.

He is the lead developer on Newskick, a tool that automates print production, built for small local newsrooms that regularly run on skeletal staff.

"Often it's one person doing many jobs, being editor and managing the ads and doing print," he said, referring to news operations that NEWSWELL has acquired.

"The main thing we're trying to do is make it so that the process of creating a local newspaper is a lot easier for a smaller staff. In what ways can we make that person's life easier so that we can have more, better local newspapers?"

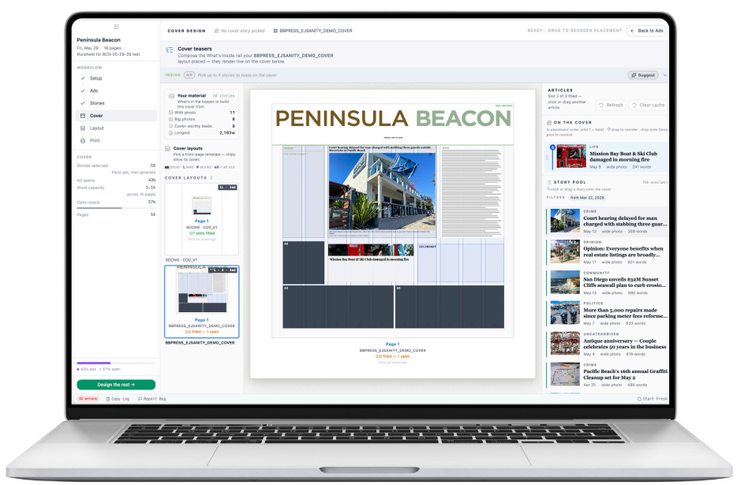
> **"Hey, it looks like you're doing a chore that you hate,' and they're like, 'Yeah, I do hate this chore.' And I'm like, 'The robot can do that chore.'**

What has surprised him most during the fellowship

EJ came into this work optimistic about what large language models could do for newspapers — especially because news is fundamentally about text.

"So, you'd say okay, this is a pretty good application for this technology, so let's move around thoughts and concepts in a newspaper," he said.

What surprised him was how limited the models still are for some of the tasks that matter most.



Newskick, a NEWSWELL tool that automates print production for small newsrooms

> **"Even with a human in the loop, we're not seeing results that even come close to what a human page designer would do. The model makes choices that experienced editors "would never do in a million years."**

On the plus side, they're reasonably good at:

- Summarizing articles
- Shrinking an 800-word story into a 500-word print slot

But when he tried to let the AI control the entire layout of a newspaper — deciding which article goes on the front page, where feature images should sit, or where ads should land — the results fell short both technically and editorially.

"Even with a human in the loop, we're not seeing results that even come close to what a human page designer would do," he said. The model makes choices that experienced editors "would never do in a million years."

What the industry gets wrong about AI

Journalists are often too pessimistic about AI's potential because they see the technology through the lens of automation, restructuring, and job security, he said.

EJ finds that working closely with journalists on their pain points is the solution.

"I say, 'Hey, it looks like you're doing a chore that you hate,' and they're like, 'Yeah, I do hate this chore.' And I'm like, 'The robot can do that chore.' Everyone loves having a task taken out for them, and when they see good results, it's good."

One lesson other newsrooms could apply tomorrow

Start extremely small. It helps both technically and with culture change.

Instead of creating a big, formal AI team or trying to overhaul workflows overnight, he advocates for:

- Tiny, minimum-viable product experiments
- Removing one hated chore at a time
- Earning incremental buy-in from reporters and editors

What he's paying attention to right now

EJ is closely watching the economics of AI, especially falling token prices and what they reveal about the true cost of these systems.

He likens today's AI ecosystem to the early days of Uber, when rides were heavily subsidized.

For him, the strategic question is: Do you build tools that will always require tokens, or use tokens now to build durable systems you own?

What he's most optimistic about

EJ's optimism comes from a sober view of the industry: In many ways, journalists have already been "pre-fired." Newsrooms are so thinly staffed that the best use of AI isn't in automating people out of jobs so much as filling in gaps.

He sees huge upside in freeing journalists from logistics and drudgery so they can do more of what humans excel at.

"We can just use it to make better journalism – you'll have time to make more calls, the robot can handle the scheduling, just go talk and do the thing that only a human can do."



The Fast Four

1
A book, article, or newsletter you'd recommend?
"Inventing the Future: Postcapitalism and a World Without Work," by Nick Srnicek and Alex Williams.

2
What's one thing you're learning that has nothing to do with AI?
Becoming an amateur ham radio operator.

3
What local journalism challenge deserves more attention?
Existence.

4
One word that describes the next year of AI in news?
Reckoning. EJ anticipates a period when the industry confronts both the real capabilities and the real costs of AI — and when an increasingly skeptical public forces news organizations to be more thoughtful about how they deploy it.



FELLOW SPOTLIGHT

Kevin Hoffman

The Philadelphia Inquirer

As an AI Fellow at The Philadelphia Inquirer, Kevin Hoffman is helping a 197-year-old newsroom navigate a fast-moving AI landscape while staying grounded in editorial values and real-world workflows. In conversation, he reflected on how his role has evolved and how another Lenfest AI Fellow helped him build trust with reporters.

About his role and the problem he is trying to solve

Kevin joined The Inquirer at the end of 2024, initially with one clear goal: To build a digital librarian. He worked closely with the Microsoft AI Development Acceleration Program (MAIDAP) team, including a two-week hackathon with their engineers.

"We quickly realized that we were able to prototype a lot of these more complex applications far quicker than anticipated, just through the resources made available to us in the grant through the MAIDAP team at Microsoft," he said.

In the early months, the fellowship was highly engineering-focused. Kevin spent his time with his head down, building prototypes and showing tools to stakeholders to get them to buy in.

"The role has largely transformed throughout the year and a half," he said. "I have definitely taken on a more considerable role with also being like an AI liaison or consultant with the newsroom."

That bridge role includes working with leadership on internal AI policies, talking with reporters about their experiences with the tools, and increasingly, helping reporters build their own tools as they become more comfortable with AI systems.

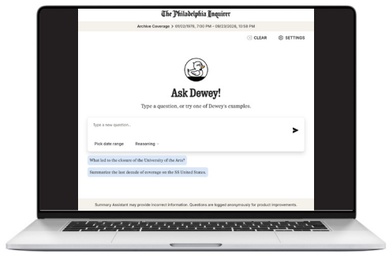
"That's been the big topic recently," Kevin said. "I hear from lots of different staff in the newsroom, like, 'Hey, I vibe coded this project, I don't really know much about coding, but it really helps me, and I'd like to share it with my coworkers. Is there any way we could build this into production and have multiple people collaborate on this project?'"

His flagship project

Kevin's original digital librarian project, Dewey, is now moving into its next phase. He is focused on two expansions:

- Image search over the Inquirer's photography archives, so reporters can search as easily for visuals as text.
- Exploring digitizing historical newspaper issues that pre-date the existing digital archives, with the goal of introducing much older material into a searchable system.

He is also supporting another tool that The Inquirer developed, a local media aggregator called Scrape, designed for journalists working on hyper-local newsletters for townships around Philadelphia.



"Dewey," a chatbot Inquirer reporters are using to retrieve relevant reporting archives

Reporters would previously spend 12 hours a week checking hundreds of sources, such as government and school district sites, for potential stories. "It was a real bottleneck," Kevin said.

Now, reporters give Scrape a list of sources to monitor and receive a daily report summarizing updates from the last 24 hours.

"We've really tried to simplify the news aggregation process from a 12-hour task a week to something that is more like opening your inbox and gathering tips from Scrape's report," Kevin said.



"We'd gotten really close with several different investigative reporters that strongly evangelized the use of some of these tools within the newsroom. And building that level of trust with people with deep institutional knowledge at The Inquirer accelerated that pace."

The biggest surprise during the fellowship

Kevin went into the role expecting trust-building to be a major challenge. What surprised him most is how quickly newsroom needs drove ideas for AI tools, once the usefulness of AI — and trust — had been established.

"Once we had crossed that bridge, just how fast the pace would move and how much conversations would shift surprised me," he said. Adoption went from journalists asking whether they should use AI to complaining that a tool was down for even 30 minutes. "I didn't realize it was being depended on like that."

The real accelerator came when reporters became vocal advocates for AI tools.

"We'd gotten really close with several different investigative reporters that had strongly evangelized the use of some of these tools within the newsroom," he said. "And building that level of trust with people with deep institutional knowledge at The Inquirer really accelerated that pace."

What the industry gets wrong about AI

For Kevin, one widespread mistake is the way AI has been rolled out.

"The way that AI has been deployed across industries, not just the news industry, has been kind of in everybody's face," he said. "It's in Zoom, it's in Slack, it's in Gmail, it's kind of everywhere you go, and that can feel really exhausting."

AI was often sold as a tool that can do anything, leading to requests that don't match what large language models are actually good at.

"We were finding that a lot of people would ask questions like, you know, 'Give me every single time we've ever mentioned this one person in our reporting, and I would like you to figure out how many articles we've written about them that's positive,'" he said. "That's a really hard thing for AI to do... it's not a tool that's meant to be used as an exhaustive feature."

Instead, he emphasizes the need for AI literacy and clear communication about the limitations and appropriate uses of these tools.

"Once we are able to arrange them in a way where people understand exactly where the sort of pitfalls are, we get a lot further," he said.

Lessons other newsrooms could apply tomorrow

One of Kevin's most practical lessons is about showing up and listening.

He credits ProPublica AI Fellow Dana Chiueh, a reporter-turned-engineer, with pushing him to spend more time simply talking with journalists about how they felt about AI before pitching solutions.

"Have you made an effort to go out there and connect with the journalists?" he recalls her asking.

That advice helped him see that building effective AI tools for journalism meant immersing himself in editorial workflows. He believes publishers should embed people at the intersection of engineering and the newsroom, and then give them the time and support to deeply understand pain points, constraints, and values.

"Building AI as an AI engineer doesn't mean much unless you have a subject matter expert in the area you're trying to build AI for," he said. As a loyal Inquirer reader but a news industry newbie, he needed to learn "what it looks like for high-quality journalism to be produced" and where tools could genuinely help.

What he is paying attention to right now

Explainable AI, a field in which Kevin worked extensively before this fellowship.

"We're increasingly becoming reliant on a lot of these tools that are AI-focused in so many areas, but we still feel like they're a black box," he said.

What he is most optimistic about

Kevin marvels at how he is witnessing the cultural evolution underway in the newsroom, with reporters growing more comfortable with AI, becoming builders themselves, and using AI tools to support serious investigative work.

"We're just starting to discuss how to make it easier for reporters to deploy these mini products on their own."

The Fast Four

1. **A book, article, or newsletter you'd recommend?**
   The AlphaSignal newsletter.
2. **What's one thing you're learning that has nothing to do with AI?**
   Urban sketching.
3. **What local journalism challenge deserves more attention?**
   News deserts.
4. **One word that describes the next year of AI in news?**
   Re-strategizing.



FELLOW SPOTLIGHT

Peter Rasmussen

The Baltimore Banner

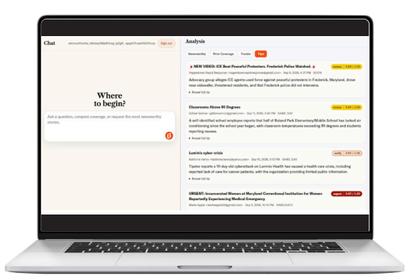
As an AI Fellow at The Baltimore Banner, Peter Rasmussen is building tools that directly support both impactful journalism and the Banner's business operations. We spoke with Peter about the lessons he's learned so far, what excites him most about AI, and where he believes the industry is headed next.

About his role and the problem he is trying to solve

Peter Rasmussen is the principal AI engineer at The Baltimore Banner. His work focuses on building practical tools that help journalists go faster and deeper, and on helping the business solve real problems.

His flagship project is News Detector, which has several applications.

The first is to scrape news from publications across the internet and assign them a newsworthiness score based on criteria



A look at how reporters can search through relevant news articles through News Detector

> **"We have human beings who are experts – that's the truly differentiating factor."**

suggested by editors. The goal: To save editors from having to manually scan several different news sites every morning. To do this, Peter developed a system where the news is provided to editors, with a layer of news judgment already applied on top.

News Detector also contains a tool that lets journalists use natural language to query what the Banner and other publications have written about a particular topic, for research. The tool was partly inspired by Lenfest AI Fellow Kevin Hoffman and The Philadelphia Inquirer's projects, he said.

"It's nice having that collaboration because that made things for us easier, as they shared their code with us."

Also useful: A "foodie bot" that scrapes Instagram posts and turns them into a weekly digest, integrated into Slack so an editor can more easily scale a concept that was originally vibe coded but too expensive to run at scale.

Peter is also working on a tool to sift through audience tips, separating genuine leads from self-promotional pitches, and to power an events calendar that surfaces "top things to do" in a user-friendly way.

Peter came to AI in local news with two degrees in engineering and a career that led him through environmental consulting in New Orleans and a year of work in Kenya, where he worked for a startup focused on energy and agricultural initiatives. He also worked as a management consultant in Washington, D.C., where he helped build an infrastructure mapping and navigation platform, a project that sparked his interest in data science.

He then completed a master's degree in AI at Johns Hopkins University, focusing on computer vision and geospatial analytics, including projects that used satellite imagery to detect heavy construction and other changes.

He was a Banner subscriber when he saw the job posting, "and I was like, 'Oh my god, I gotta apply to this.'" The chance to combine his technical skills with mission-driven local journalism was too compelling to pass up.

The biggest surprise during the fellowship

One of Peter's biggest ongoing lessons is how hard it is to estimate effort accurately. For example, he thought he would be able to finish building the News Detector scraper in a week, over his Christmas holidays.

"It took five weeks. I wasn't done until the end of January," he said.

Another major surprise has been just how powerful agentic coding can be. Early on, faced with a long list of requirements for the fellowship and for his role, he realized he couldn't just hand-code everything himself. Instead, he focused on building a harness to govern the behavior of LLMs, creating a structure that constrains and regulates the agents.

The combination of better frameworks and rapidly improving models has delivered a "step-change in productivity."

What the industry gets wrong about AI

Peter is clear that he does not see AI as a replacement for journalists and believes some organizations misuse AI by treating it as a substitute for human reporting.

"We have human beings who are experts – that's the truly differentiating factor."

At the Banner, he sees a healthier approach: Use AI as an assistant for summarizing, surfacing information from outside the newsroom, and understanding customer behavior, while keeping editorial judgment firmly human.

"Everything else has to be vetted by humans," he said, pointing to the rigorous daily editorial conversations he overhears from his desk.

Lessons that other newsrooms could apply tomorrow

Peter offers two tightly linked lessons: build scalable frameworks and start from real use cases.

"Build frameworks that allow you to scale your productivity," he said. His agent harness has now been adopted by other engineers, including colleagues at the recently acquired Pittsburgh Post-Gazette. One of their engineers, he notes, now "has a better setup than I do," and is tackling a massive migration project with impressive productivity.

The second lesson is about requirements discipline: "Really understand what the use case is, what the needs are, and elicit the use cases, because otherwise you're building something that isn't going to be used."

In other words, don't start with the model. Start with clearly defined newsroom and business problems.

What he is paying attention to right now

Right now, Peter is closely tracking rapid advances in large language models, both proprietary and open source.

What he is most optimistic about

Deepening the Banner's understanding of its audiences and using that insight to build better products and more sustainable revenue.

His recent work clustering users by engagement and connecting that to donor propensity is already showing that "people who are more engaged tend to... donate more." He is developing a blueprint for donor modeling that leverages those insights as well as learnings from Chicago Public Media AI Fellow Mark Chonofsky.

Peter is also excited about using AI to take repetitive, time-consuming tasks off people's plates, for example, building structured rubrics to help HR sift through hundreds of resumes for a single position, "basic but onerous tasks that take forever" and keep staff from doing higher-impact work.

"I'm so optimistic that we'll be able to take care of a lot of the stuff that's just a pain... that people don't like, that's time-consuming and onerous, and we can take that off people's plate."



> **"I'm so optimistic that we'll be able to take care of a lot of the stuff that's just a pain... that people don't like, that's time consuming and onerous, and we can take that off people's plate."**

The Fast Four

1
A book, article, or newsletter you'd recommend?
"Doppelganger" by Naomi Klein.

2
What's one thing you're learning that has nothing to do with AI?
How to create a "lush garden," digging up weeds and understanding wildflowers, in his front and back yards.

3
What local journalism challenge deserves more attention?
Resourcing and workload. "People wear multiple hats and are very busy. Workflow automation is crucial to relieving that pressure."

4
One word that describes the next year of AI in news?
Adoption.

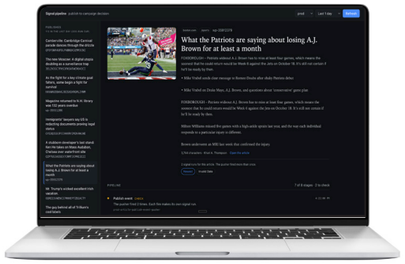
As an AI Fellow at the Boston Globe, Omar Reid is tackling brand safety, personalization, and data silos to help the newsroom unlock the full value of its journalism.



FELLOW SPOTLIGHT

Omar Reid

The Boston Globe



The Boston Globe is using AI to help advertisers place more targeted ads.

About his role and the problem he is trying to solve

Omar's path into journalism is unconventional but tightly focused on AI.

"All of my career, I've basically been an AI engineer," he explains. He worked briefly at a bank, where he realized that he was merely a cog in a huge machine and unlikely to have an impact. So, when he had the opportunity to do a one-year work placement, he went to a startup in Tokyo, Cogent Labs, which had an English website and had hired many foreign engineers.

At The Boston Globe, Omar's role sits at the intersection of data, AI, and organizational change.

The core challenge he's working on is what he calls a "family of problems" rooted in silos:

- Data silos: "All the data sets [are] in separate systems."
- Organizational silos

To tackle this, he centralized archives and built an internal AI chatbot in Slack that lets colleagues search across the Globe's data.

"Once they saw that, they were like, 'Wait, hold on, we have all of this useful data, and now we have... a way to access it,'" he said. That proof of concept helped build trust.

Now, he's applying that same approach to a high-stakes problem: brand safety and the long-term opportunity of personalization.

His project at The Boston Globe is built around extracting rich signals from every article — topics, people, places, organizations, and sentiment — so the ad system can understand what a piece is actually about, not just which section it sits in. It is designed to give advertisers fine-grained control to avoid sensitive content while targeting context that aligns with their brand.

Longer term, the same signals that power brand safety can also underpin personalized products such as tailored newsletters, letting the Globe surface stories that match what individual readers actually care about while protecting advertisers from risky contexts.

The biggest surprise during the fellowship

Omar said his biggest surprise has been the disconnect between the influence of news organizations and the money they make, especially compared to tech giants.

"I didn't realize... news used to run things, have a lot of power," he reflects. When he joined the Globe, he even told them in his interview, "Oh, the Globe still exists?"

Over time, he came to understand that the Globe is "still around and kicking," with real cultural presence and prestige, but operating at a very different financial scale from tech platforms.

"You expect it to be making Meta money, but Meta is like a trillion-dollar company... It's like a really small group of people in this industry."

His own work is driven by his understanding of how much value is locked up in news archives. "I'm less worried about the views and clicks, because I know that they will happen if I put my attention on surfacing the value from all of this knowledge that's been accumulated."

What the industry gets wrong about AI

For Omar, the problem is skepticism amplified by misconceptions.

"Reporters seem to be very skeptical, which I think is a good thing," he said. But when that skepticism is "married with the misconception," it can become paralyzing.

He describes speaking with people who think "AI started in 2023 when ChatGPT was released," and who rate the danger level of AI on a scale of zero to 10 as "10 – Terminator."

"Whereas to me, AI is like a clever washing machine," he said. Through conversations, he has tried to bring people back from science-fiction fears to a more realistic view.

"It just makes it hard to do useful stuff" when the conversation is dominated by apocalyptic narratives, he pointed out.

Lessons that other newsrooms could apply tomorrow

Omar's advice is simple but difficult: Slow down enough to learn.

He sees many journalists and newsrooms as "conditioned to react," driven by constant urgency.

That same mindset, he argues, undermines meaningful experimentation with AI:

"If I tell you now to sit down on a computer, spend 30 to 40 minutes playing with something... this is just discovery now. This is not you reacting to anything. That becomes very hard if you're very urgent and reactive."

Newsrooms need to consciously drop the urgency when approaching AI.

"Be patient with yourself, be patient with the process, be patient with the computer," he said, to build real intuition about where AI can help and what's just hype.

> **"Be patient with yourself, be patient with the process, be patient with the computer to build real intuition about where AI can help and what's just hype."**

What he is paying attention to right now

Omar is focused less on shiny new objects and more on building the right things.

He's been reading "Continuous Discovery Habits" by Teresa Torres, a book that "outlined the process which people can use to know what is worth building." With AI making it "really easy to build things," he worries that teams will "start building things that in the end no one wants, because they didn't think to check first."

He's also spending time with the revenue team to deeply understand the Globe's business model. "What are the core levers here? And they tell me what it is, and I'm like, that is perfect. Now I know what we should target and why it will move the needle and produce revenue for the company."

This lens helps him avoid "solutions in search of problems" and instead align AI projects with both user value and business value.

What he is most optimistic about

Omar believes we are at a turning point for language-driven organizations, after years of watching numbers-driven organizations succeed.

"We deal with language, and now we have a tool that makes it really easy for us to deal with language," he said, framing it as journalism's "Moneyball moment" — a time when the tools finally match the domain, just as data analytics transformed baseball.



The Fast Three

1. **A book, article, or newsletter you'd recommend?**
"Continuous Discovery Habits" by Teresa Torres — "the number one thing people need to discern what's worth building and not worth building."

2. **What's one thing you're learning that has nothing to do with AI?**
How to deal more effectively with issues of identity.

3. **One word that describes the next year of AI in news?**
"Consequences. If you successfully understood the hype, you know what was hype, what was real, and you put the work in, looking at the problems in your company, trying to solve those, and you use AI to help you do that, you're going to reap what you sow."



FELLOW SPOTLIGHT

As an AI Fellow at The Minnesota Star Tribune, Ryan Thayer is exploring how artificial intelligence can unlock the value of decades of archived journalism and reshape how journalists search for, discover, and use news. We spoke with Ryan about the lessons he's learned so far, what excites him most about AI, and where he believes the industry is headed next.

Ryan Thayer

The Minnesota Star Tribune

> **"Just connecting with other people and hearing their progress and how they go about things has really helped me out a lot."**

About his role and the problem he is trying to solve

Ryan's path to the Star Tribune runs through software engineering, enterprise architecture, data science, and machine learning. When large language models and tools like ChatGPT arrived, he was well positioned to jump in.

At the Star Tribune, Ryan describes his role as part builder, part enabler. "My team and I have built AI tools from scratch. I help people from across the Star Tribune get started on their projects. I've helped others deploy their projects and apply engineering best practices if they're not familiar with that."

One of the projects he has focused on as an AI Fellow is archive search — helping journalists make sense of the vast trove of past coverage in ways that go far beyond simple keyword queries.

Rather than just answering narrow factual questions, the team is experimenting with richer, more conceptual queries.

"Instead of just a simple question-answer, it can be like a summary type of thing. Like, 'Tell me what the Star Tribune has published about this wide-ranging topic.'"

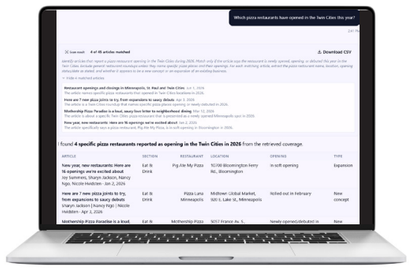
The vision is to support both reporters and, perhaps eventually, the public. The system is currently in the proof-of-concept stage.

What has surprised him most during the fellowship

Ryan came into the fellowship used to working independently or within a single organization's silo. One of the biggest surprises has been the level of cross-industry collaboration among Lenfest fellows.

He points to conversations with peers at other organizations in the cohort, where they are working on similar projects. "They have different perspectives. They are able to share things they've learned because they are at a different stage of a project, and I'm able to share feedback I've received because we went about it a little differently."

"Just connecting with other people and hearing their progress and how they go about things has really helped me out a lot."



An example of how the Star Tribune's archive search tool responds to complex questions with targeted answers

What the industry gets wrong about AI

Ryan considers it a mistake to let AI do writing that requires a human touch.

"I read a lot of articles written by AI that I feel are not good. To me, there are a lot of words but not a lot of insight."

At the Strib, there is a clear line: AI is not there to replace reporters or generate stories for publication.

"The reporters and journalists are still going to write their articles, but we're looking to see what we can do to make things easier for them."

The focus is on using AI to automate tedious, time-consuming work, such as background research, summarization, and other repetitive tasks, so journalists can focus on reporting and storytelling.

He's noticed that once people learn more about practical ways AI can be useful, their attitude toward AI experimentation & adoption changes.

One lesson from his work that other newsrooms could apply tomorrow

Ryan's key learning is product teams need to have discipline around the impact of their work.

"There's a lot of cool stuff we can do, but for the work we're doing, how does it move metrics? What are our main KPIs, and then how does the work move those metrics? And if it doesn't, then we should probably be focusing on something else," he said.

Archive search is a good example. The team constantly asks: What difference does this really make? If they can't articulate the impact, it's a sign to rethink priorities.

Right now, they're still experimenting with which KPIs archive search can meaningfully affect: deeper engagement, new ways of using legacy content, and surfacing stories that people can't easily find via search.

What he is paying attention to right now

Ryan and his team are particularly focused on agentic development, using AI agents to help build and maintain software faster.

"How can we build things faster? How can we use agents to do the coding for us? And then again, where's that line – what do people still have to do, what do experts still have to do, versus what can an autonomous agent take care of?"

The team is thinking about accelerating how they build everything, while still keeping humans in the loop for oversight, judgment, and deployment decisions.

What he is most optimistic about

How AI is reshaping roles and lowering barriers to contribution across the organization.

Ryan is seeing a world where non-engineers can meaningfully shape products.

"Once we get an agentic development system set up, then someone who isn't a software engineer can say, 'Here's an issue I see' or 'Here's an improvement I need to make,' and then the agents can make that update or at least do an evaluation and suggest updates."

Similarly, product managers and designers can move faster.

"A product manager can mock up a design pretty quickly and hand that to the design team, and the design team is not starting from scratch."

Or someone without deep coding experience can build a functioning prototype, then pass it to engineers to harden and deploy properly. "Again, that's saving weeks from the requirements gathering process."

For Ryan, this blurring of traditional role boundaries is one of the most exciting shifts underway.

> **"How can we build things faster? How can we use agents to do the coding for us? And then again, where's that line – what do people still have to do, what do experts still have to do, versus what can an autonomous agent take care of?"**



> **"There's a lot of cool stuff we can do, but for the work we're doing, how does it move metrics? What are our main KPIs, and then how does the work move those metrics? And if it doesn't, then we should probably be focusing on something else."**

The Fast Four

1. **A book, article, or newsletter you'd recommend?**
   "Thinking, Fast and Slow" by Daniel Kahneman.

2. **What's one thing you're learning that has nothing to do with AI?**
   Ryan points to curiosity as his ongoing personal focus – intentionally staying open to new ideas and digging into new topics as a mindset in daily life.

3. **What local journalism challenge deserves more attention?**
   Hyperlocal coverage and truly distinctive niches: "How do we fit in versus what everyone else is doing? What niche can we really focus on that no one else is doing?"

4. **One word that describes the next year of AI in news?**
   Fast-paced.

THE LENFEST INSTITUTE

2026 | CELEBRATING 10 YEARS

LENFESTINSTITUTE.ORG

100 South Independence Mall West
Suite 600
Philadelphia, PA 19106
support@lenfestinstitute.org
215.854.5600

Design: Piper R. Lyman