

# Temporal Perspectives on Exploring PaperWave: a Document-to-Audio Adaptation System powered by LLMs

Yuchi Yahagi  
yahagi@hc.ic.i.u-tokyo.ac.jp  
The University of Tokyo  
Bunkyo, Tokyo, Japan  
JSPS DC  
Chiyoda, Tokyo, Japan

Rintaro Chujo  
chujo@hc.ic.i.u-tokyo.ac.jp  
The University of Tokyo  
Bunkyo, Tokyo, Japan

Yuga Harada  
yugaharada@gmail.com  
Educe Technologies  
Tokyo, Japan

Changyo Han  
hanc@hc.ic.i.u-tokyo.ac.jp  
The University of Tokyo  
Bunkyo, Tokyo, Japan

Kohei Sugiyama  
sugiyama\_kohei@iii.u-tokyo.ac.jp  
The University of Tokyo  
Bunkyo, Tokyo, Japan

Takeshi Naemura  
naemura@hc.ic.i.u-tokyo.ac.jp  
The University of Tokyo  
Bunkyo, Tokyo, Japan




Figure 1: Timeline of the PaperWave design process. This includes the key events and publication activities.

## Abstract

This paper presents a temporal reflection on designing PaperWave, a document-to-audio adaptation system that converts research papers into conversational podcasts using large language models (LLMs). Through an autobiographical design approach spanning over half a year, we explored how the design process evolved alongside shifts in lifestyles and technical developments. Our continued exploration revealed three key themes: shifts in lifestyles that changed users' relationships with PaperWave, challenges in maintaining fast tinkering, and implications of parallel development with commercial products. These findings contribute to understanding temporal dimensions in research through design, particularly with rapidly evolving AI technologies.

## 1 Introduction

This paper shares our experience designing PaperWave (Figure 2) [14]. PaperWave lets users listen to research papers while doing other tasks by turning them into audio. It uses large language models (LLMs) to create podcast-style conversations about the papers. In these podcasts, a host asks questions and an AI-generated author explains their paper. We used an autobiographical design approach to study PaperWave [10]. While we will present a two-month case

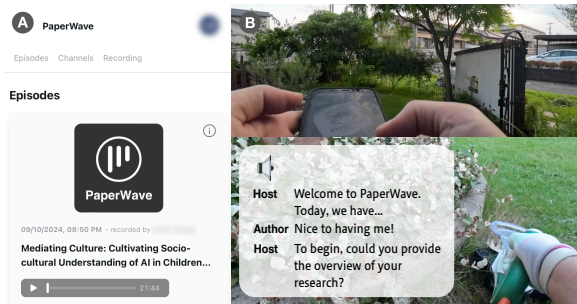
study at CHI '25 [14], here we reflect on our extended exploration over six months, as shown in Figure 1.

We believe PaperWave's design process can add to this year's Things of Design workshop discussion. Through autobiographical design, we explored how AI technology can change ways we engage with knowledge in daily lives [14]. Unlike other research about reading support systems (e.g. [1, 7, 8]), we framed listening research papers as an activity situated in everyday routine. For this reason, after two rounds of prototyping, we deployed PaperWave as a research product [12] and used it long-term. This process showed that to design a digital app like PaperWave, we need to think about how listeners interact with the physical world during listening [14].

Beyond the case study, we are keeping exploring PaperWave's design. Our exploration has now lasted over six months. In this paper, we look back at how time played a role in PaperWave's design process. We discuss three main themes that emerged: shifts in lifestyles, why quick tinkering slowed down, and what it meant to develop alongside commercial products.

## 2 Context of Our Study

Let us briefly introduce our study context. For more details, please see our case study [14].



**Figure 2: The PaperWave app and its real-world usage. (A) The PaperWave web app interface. (B) A1 listening to PaperWave while mowing the lawn on August 18.**

## 2.1 Methods

We conducted an autobiographical design to study how researchers use audio versions of papers. The study ran from June 25, 2024, to now (February 28, 2025). Our case study covered the period up to September 2, 2024, when we held design workshops. During the first study until September, five authors (A1–A5) and six collaborators (P1–P6) participated as secondary users. Everyone logged their PaperWave experiences in Slack diaries. After September, collaborators could keep using PaperWave but weren’t required to keep diaries. Since then, mostly just the authors continued logging their experiences. This paper focuses on the authors’ perspectives over time.

## 2.2 Who We Are

We are an interdisciplinary research team focused on learning and design rather than LLM technology. This lets us concentrate on user experiences. Over the past few years, A1 and A2 have grown particularly interested in design alongside their existing research work. PaperWave marks our first deep dive into research through design (RtD). A1 and A2 built PaperWave, while A3–A5 contributed to its design through weekly meetings. A3, who works as a product manager, brought valuable insights by regularly reading papers both for hobby and business needs. A5 initially joined as a potential non-user to balance our enthusiasm for PaperWave. While Japanese is our main communication language, A4’s first language is Korean, offering insights from a different linguistic background.

## 2.3 How PaperWave Works

PaperWave is a web app that turns research papers into conversational podcasts. Users upload a PDF file and PaperWave creates a script featuring a dialogue about the paper. This script is then converted to audio using text-to-speech (TTS) technology. To create comprehensive yet adjustable-length podcasts, we use multiple LLMs in different roles (program writer, info extractor, and script writer) working together.

## 3 Temporal Aspects of PaperWave’s Design

We identified temporal aspects in our design process by reviewing diary entries from Slack using MAXQDA. The following themes emerged from this review.

For clarity, quotes from diaries are cited as (author, date). Japanese diary entries were translated to English by the authors.

### 3.1 Shifts in Lifestyles

During our long-term use of PaperWave, we noticed changes in how we used the system. As our lives changed, so did our relationship with PaperWave.

**3.1.1 Seasonal Occurrences.** Seasonal breaks gave participants chances to try PaperWave in new ways. During Japan’s winter holidays, A3 wrote about a new use case:

*Listening to the paper while reading it at home. This is my first time trying this. (A3, December 27)*

These deviations from typical usage revealed new design insights. When A3 tried this combined reading-listening approach, he struggled to find where the podcast’s discussion matched the paper’s text (A3, December 27). This experience suggested adding an interface feature to show which part of the paper corresponds to the current audio.

However, others also tried similar combined reading-listening without seasonal triggers. When A1 couldn’t find a seat on a crowded bullet train, he tried listening while reading. “*Since my eyes were free, I tried reading along, but I kept losing track of which part they were talking about*” (A1, November 17). He faced the same challenge as A3. While seasonal breaks can spark new insights, the resulting use cases aren’t limited to those periods.

**3.1.2 Shifts in Daily Routines.** We also observed changes in how participants used PaperWave in their everyday lives. For example, A1 described changing his commute:

*From Ueno to UTokyo on foot. I’m trying to walk instead of riding a bicycle because I feel I’m not getting enough exercise, which gives me a good chance to listen. (A1, December 27)*

While a product’s usefulness depends on how it fits into a user’s lifestyle, this experience shows that users’ habits themselves can change over time. Such lifestyle changes may not always happen during a short study period, but longer-term observations increase our chances of seeing these shifts.

**3.1.3 Life events.** Our study captured major changes that rarely occur in people’s lives. A5, who initially said he would never use PaperWave [14], started using it after receiving a job offer.

*This was my first time actively using PaperWave. [...] I noticed that my need for PaperWave changes with life changes. The biggest change was getting a job teaching media studies lectures and seminars next year. Before, I didn’t need PaperWave because I couldn’t use it for reviewing papers to write papers. However, teaching and student guidance are different. For education, I need to know broad topics of media studies beyond my own research. Instead of understanding each paper deeply, knowing what topics exist in media studies is more important, and PaperWave seemed well-suited for this purpose. [...] While I’m not sure how much I fully grasped, I was able to take*

*it in as a topic for media studies, and it will probably be useful for class material. ANONYMOUS-san posted on X about studying various fields for undergraduate research guidance, and I think PaperWave is well-suited for that kind of use. (A5, December 23. The X user is anonymized.)*

Life changes help designers develop a deeper understanding of their designs. In this case, A5's career change led to a new way of using PaperWave, revealing its value for teaching and student guidance. While studying many users can provide broad insights, long-term research that captures designers' own life changes can offer unique perspectives by comparing how individuals use products in different situations.

### 3.2 Stalled Fast Tinkering

While fast tinkering is a key aspect of autobiographical design [10], our long-term exploration revealed situations that often limited it. A1, PaperWave's main developer, frequently noted this challenge in his diary:

*I have ideas for improvements, but since I haven't implemented them, I keep using the old version while thinking, "It could be better." The changes I want to make most involve platform updates, like connecting to Zotero or switching LLM to Google's Gemini, so I don't feel like doing small changes lightly, which makes it awkward to do so. (A1, October 2)*

This situation also affected how A1 used PaperWave:

*I would like to look back at the past two months, when I was writing a paper or preparing for public release, I used it less often, and updates to features did not progress. I wanted to use it more but kept thinking I should add certain functions first. It was frustrating that I couldn't implement these changes because other work took priority. (A1, November 4)*

These cases show the challenge of balancing research publication work with development work. This was especially important for PaperWave, where there was pressure to publish quickly due to parallel development with commercial products, as discussed in 3.3.

This challenge could be viewed as part of creatively repairing complex digital objects [9]. A1 did eventually implement some changes: *"I had been putting off the implementation because it seemed troublesome, but GitHub Copilot Edit made it quick and easy to try, and I finished it in a short time."* The change fixed slow loading times caused by too many episodes (A1, February 15). This was similar to fixing wear and tear from long-term use and accumulation. Even when issues aren't obvious bugs, we can think of them as repairing gaps that appear during daily long-term use.

### 3.3 Parallel Development with Commercial Products

**3.3.1 Impact from Commercial Products.** Our exploration of PaperWave happened at the same time as commercial products, such as NotebookLM [6], Illuminate [5] and Wondercraft [13], were being developed. This timing had both advantages and disadvantages. As

mentioned earlier, commercial products sometimes create pressure. One author expressed this feeling: *"I wish Google had waited a little longer to release the product"* (A3, September 22).

However, having commercial products around also helped us. When comparing PaperWave to Google's NotebookLM, we found our unique strengths: while NotebookLM dives deep into specific topics, PaperWave explains the whole paper and gives users a sense of *having heard the paper* (A1, November 27). A4 also shared their thoughts after trying Google Illuminate:

*I also listened to Illuminate for comparison, but it focused on being "fast" and took the stance of pounding knowledge into the listener quickly, which made me feel exhausted. (A4, February 13)*

These comparisons helped us understand PaperWave's unique value. As Gaver notes, design research and practice are generative, and multiple incompatible worlds can coexist [4]. This suggests that conducting RtD alongside commercial product development is valuable. Even with the presence of commercial products, it remains valuable to conduct research focused on our specific context, such as user values and daily routines.

**3.3.2 Technology Updated Frequently.** The rapid advancement of generative AI affected PaperWave's development. For example, when we switched to gpt-4o-2024-11-20, one participant noted: *"I feel that indeed the quality of PaperWave scripts is also improving [...] (The host's reactions) used to be all interesting, but now it seems like there are more variations like, 'Enjoy,' or 'Is there anything you want to add to the discussion?'"* (A3, December 13). While quick technological changes created challenges mentioned in 3.2, they also opened doors for new insights.

**3.3.3 Impact of Public Engagement.** We are opening our PaperWave to the public. While commercial products lack Japanese language support, PaperWave has gained traction among Japanese users. From early November 2024 to mid-February 2025, 486 users signed up, and 1642 episodes were created. Conducting research at a time when commercial products were being actively developed may have had the advantage of facilitating public engagement.

Gathering insights from public users presented several challenges. We set up a feedback form that received 14 voluntary responses. While many responses echoed findings from our autobiographical design, this overlap validated that different users shared similar impressions of PaperWave. Not all user feedback came through our form—users shared their thoughts on social media instead. Determining how to incorporate this informal feedback into our research was challenging.

Notably, public user reactions provided emotional support for the authors, who shared and celebrated user feedback on Slack. For example:

*This is a pleasant feedback 😊 (A1, December 26)*  
*> Since it is specialized for papers, the parts of the papers that I need to read are well audible and useful. It will be indispensable for my graduate school life starting next April. I used to use Google's NotebookLM, but even when the prompt said something like, "Explain for researchers," the content inevitably turned out to be like a general news radio program,*

*so I thought it was not practical in terms of reading the paper. I hope that PaperWave will continue and evolve over time.* (Feedback from a user)

## 4 Discussions

Here, we discuss how time influenced PaperWave’s design process in relation to the workshop’s key questions.

### 4.1 Contributions of Long-term Deployments

The lifestyle changes and life events described in 3.1 are unpredictable and may not emerge during a short study. However, longer deployments increase the likelihood of capturing such major personal changes. Our two-month study [14] did not capture these shifts, but our continued deployment enabled us to observe them. This case highlights a different type of unlike event, distinct from major societal events like environmental catastrophes or innovations [3]. For seasonal changes, conducting research for at least a year is recommended, similar to other fieldwork. However, some artifacts may not need to capture seasonal patterns, and short-term intensive studies can also be valuable [11].

Making the research product public allows more people to use it over time. While online feedback rarely led to new insights, we found that user feedback provided emotional support for researchers (3.3.3).

### 4.2 Challenges in Long-term Deployment

During our exploration of PaperWave, we experienced a slowdown in development as noted in 3.2. While A1 called this an implementation challenge, it could be viewed as a form of repair. The need for repair is typical in autobiographical design projects that go beyond simple usage [2]. Repair creates opportunities to interact with surrounding actors [2], and in our case, AI emerged as a potential new actor (3.2).

We found that this tension arose from balancing research publication needs with system development. Long-term research inevitably requires managing these competing demands. As early-career researchers with limited experience in long-term studies, we would welcome a discussion of this challenge at the workshop.

### 4.3 Ethical Considerations

Long-term research exposes researchers to public social media feedback that was not intended for research purposes, as noted in 3.3.3. This poses a greater risk than traditional lab experiments that are only revealed during research presentations. Even when researchers do not directly cite social media posts, they are inevitably influenced by this public feedback. Balancing the need to acknowledge these influences with respecting user privacy and consent presents an ethical challenge.

## 5 Conclusion

Our paper explored how time influenced PaperWave’s design process. We identified three key themes: lifestyle changes that affected system usage, challenges in maintaining rapid development, and implications of developing alongside commercial products. These themes offer insights for design research.

We hope our findings can contribute to discussions about temporal dimensions in design. At the workshop, we can provide a hands-on demonstration of PaperWave and share our experiences. We look forward to discussing how our findings can inform future design research.

## Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP22KJ1010 and Nakayama Future Factory. We would like to thank the collaborators of the field study and the design workshop for their valuable insights and feedback. The authors wrote this paper with GitHub Copilot, which provided suggestions for English expressions.

## References

- [1] Joseph Chee Chang, Amy X. Zhang, Jonathan Bragg, Andrew Head, Kyle Lo, Doug Downey, and Daniel S. Weld. 2023. CiteSee: Augmenting Citations in Scientific Papers with Persistent and Personalized Historical Context. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3544548.3580847>
- [2] Audrey Desjardins and Aubree Ball. 2018. Revealing Tensions in Autobiographical Design in HCI. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. Association for Computing Machinery, New York, NY, USA, 753–764. <https://doi.org/10.1145/3196709.3196781>
- [3] Felix Anand Epp, Anton Poikolainen Rosén, Antti Salovaara, and Camilo Sanchez. 2024. Uncertainties as Generative Resources in Research through Design: Three Dynamics for Moving in a Design Space. *ACM Trans. Comput.-Hum. Interact.* 31, 6 (Dec. 2024), 70:1–70:31. <https://doi.org/10.1145/3689041>
- [4] William Gaver. 2012. What Should We Expect from Research through Design?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 937–946. <https://doi.org/10.1145/2207676.2208538>
- [5] Google. 2024. Illuminate. <https://illuminate.google.com/>
- [6] Google. 2024. NotebookLM. <https://notebooklm.google/>
- [7] Tae Soo Kim, Matt Latzke, Jonathan Bragg, Amy X. Zhang, and Joseph Chee Chang. 2023. Papeos: Augmenting Research Papers with Talk Videos. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, San Francisco CA USA, 1–19. <https://doi.org/10.1145/3586183.3606770>
- [8] Yoonjoo Lee, Hyeonsu B Kang, Matt Latzke, Juho Kim, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliuue. 2024. PaperWeaver: Enriching Topical Paper Alerts by Contextualizing Recommended Papers with User-collected Papers. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–19. <https://doi.org/10.1145/3613904.3642196>
- [9] Leah Maestri and Ron Wakkary. 2011. Understanding Repair as a Creative Process of Everyday Design. In *Proceedings of the 8th ACM Conference on Creativity and Cognition (C&C '11)*. Association for Computing Machinery, New York, NY, USA, 81–90. <https://doi.org/10.1145/2069618.2069633>
- [10] Carman Neustaedter and Phoebe Sengers. 2012. Autobiographical Design in HCI Research: Designing and Learning through Use-It-Yourself. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. Association for Computing Machinery, New York, NY, USA, 514–523. <https://doi.org/10.1145/2317956.2318034>
- [11] Carman Neustaedter and Phoebe Sengers. 2012. Autobiographical Design: What You Can Learn from Designing for Yourself. *interactions* 19, 6 (Nov. 2012), 28–33. <https://doi.org/10.1145/2377783.2377791>
- [12] William Odom, Ron Wakkary, Youn-kyung Lim, Audrey Desjardins, Bart Hengeveld, and Richard Banks. 2016. From Research Prototype to Research Product. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. Association for Computing Machinery, New York, NY, USA, 2549–2561. <https://doi.org/10.1145/2858036.2858447>
- [13] Wondercraft, Inc. [n. d.]. AI Powered Audio Studio - Wondercraft. <https://www.wondercraft.ai/>
- [14] Yuchi Yahagi, Rintaro Chujo, Yuga Harada, Changyo Han, Kohei Sugiyama, and Takeshi Naemura. 2025. PaperWave: Listening to Research Papers as Conversational Podcasts Scripted by LLM. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Yokohama, Japan. <https://doi.org/10.1145/3706599.3706664> (to appear).