Putting digital artifacts to work

Scott Carter
scott@madpickle.net

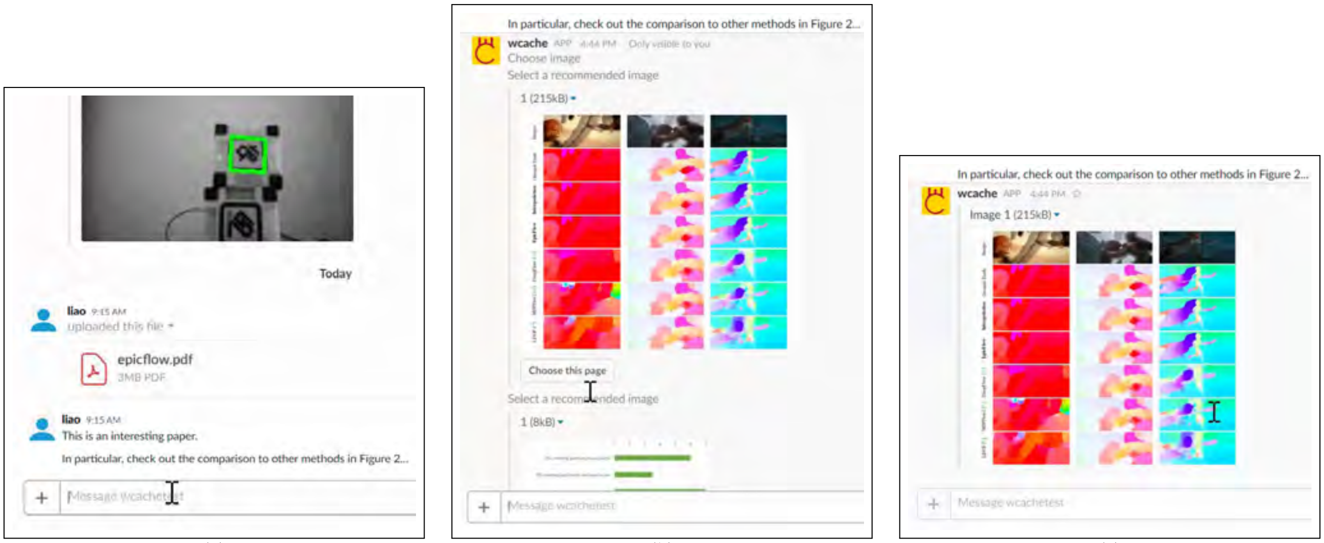

(a) (b) (c)

Figure 1: (a) Sharing documents in multimedia chat tools tends to hide critical content. Using online document analysis tools, users can more easily call out digital artifacts that can serve as conversational props. (b) An automated tool automatically identifies two candidates matching the reference "Figure 2" in the user's comment. (c) The user clicks a button to manually disambiguate the reference and the figure is injected into the chat.

ABSTRACT

Physical, collocated work environments are rife with artifacts that are ready-to-hand to serve as tools, conversation props, and content that can be remixed and reused. In digital, distributed environments, on the other hand, work artifacts tend to be siloed and difficult to extricate, making it difficult for remote workers to share interests and knowledge. In this work, we describe set of challenges for developers of systems that support digital artifact reuse. We furthermore show how these challenges are embodied in tools we built that help make digital artifacts reusable.

KEYWORDS

Future of work

Reference Format:
Scott Carter. 2020. Putting digital artifacts to work. In *New Future of Work Symposium.* 7 pages.

New Future of Work Symposium 2020, August 3-5, 2020, Online

1 INTRODUCTION

As the pandemic crisis lingers, hopes for a "V-shaped" rapid rebound of activities fades. Quarantines, shelter-in-place regulations, social distancing, and face masks seem more likely to alter the landscape of daily life permanently. For system designers and developers, the impacts of this shift are potentially profound. In particular, the crisis shifts remote, distributed work from an elective perk granted to knowledge workers on a small scale to a *mandatory and critical* aspect of worklife. This move puts even more pressure on designers of distributed systems to leverage technology to recover the nuanced experience of face-to-face interaction.

The problem: digital media lacks presence implicitly. This fact has been and remains one of the most important barriers to computational work, both synchronous and asynchronous, collocated and remote. The problem has inspired a broad array of research programs in HCI, from glanceable and peripheral displays, semi-automated information visualizations, and more recently deep learning-based visual and textual summarizations. Still, imbuing digital artifacts with presence enough to anchor conversations remains a challenge. In collocated work environments, this deficiency is easier to overcome with workaday meetings or informal chats. Furthermore, even work that is largely computational will generate some physical artifacts visible to collocated participants—printed papers, presentations, or agendas; whiteboard sketches; etc.

Klemmer et al. elucidate the value of materiality in depth [25], writing that, "the production and manipulation of visible artifacts in the workplace facilitate coordination...The visibility of a work practice manifests itself in the artifacts that the practice creates." They review an extensive body of work (including investigations of medical records [21], air traffic control [31], office work [39], publishing [4], and other domains) that find again and again that the physicality of work artifacts helps people coordinate and carry out daily work.

However, while we sympathize with Marc Weiser's foundational vision for "embodied virtuality" [43], technological tools must ultimately bend to pragmatic concerns. To whit, work is increasingly dependent on remote work tools. Though they may ultimately be deficient compared to their collocated counterparts, the accelerating impacts of global environmental collapse give us as technology developers no choice but to accept a reality in which distributed work is the norm[1].

Specifically, we have found that there are four key challenges that developers face when designing for digital artifact reuse in distributed systems:

- Finding the right content
- Making content ready-to-hand
- Limiting impacts on other tasks
- Managing privacy issues

In this paper, we first review problems with digital artifact reuse in contemporary distributed systems, then review a collection of systems that we built that can help workers address these challenges by reusing and remixing digital artifacts to support digital bricolage [7], specifically in the context of remote work. We then expand on the challenges above and put forth a call to action for future work.

2 DIGITAL ARTIFACT REUSE IN THE WILD

Supporting distributed work requires a combination of systems that allow workers to collaborate *live* and *asynchronously.*

Several studies have investigated distributed live meetings [6, 13, 20, 33, 38]. These studies have documented on a wide range of issues, from social and representational problems, the communicative ability of gaze and gesture and their impacts on user privacy, as well as technical and social asymmetries and beyond. Here, we focus on those findings related to work artifacts per se.

Carter et al. [13] conducted a broad set of surveys and interviews and found that remote participants had trouble accessing key pieces of information in meetings, "a need for remote participants to be aware of physical objects, such as paper designs or sketches that were present and visible to the group in the main room. This use case often came up for groups of designers, who often printed and hung up sketches and mockups around the room for others to view and comment on. However, when one designer was out of the room, this common practice broke down: *You could just [draw a sketch] in the notebooks and then I send pictures of the notebooks and I don't have to capture what was on the wall.*"

Another example of needing to share information with remote participants came from a mobile app designer who would hold up a phone to the laptop's webcam, as this was the easiest way to share awareness of an interaction with a physical device (the phone): "*I deploy the application on my phone. Then we use the videoconferencing so I can put the phone in front of my laptop camera so you can see what it looks like.*"

The authors found that difficulties sharing work artifacts extends beyond remote participants. In live meetings, even collocated participants run into barriers sharing information both in the meeting itself and in live backchannel chats:

"Say we were in video and Julia wanted to show [a document] quickly, it's probably not her computer that we're actually connected to so and because it just wouldn't be a quick simple thing to do, we probably wouldn't do it. The person's computer, they would try to bring up the file, so it would be, in an ideal situation, you would just, Julia would do something and boom, her screen would be shared. She doesn't even connect to the meeting, usually, through her computer. It's usually just that one computer."

The fact that digital artifacts are difficult to extract and share makes live meetings far less useful, and far more frustrating, than they could be.

The problems extend beyond live meetings to *asynchronous* media sharing. Marlow et al. describe problems that people have "refinding scattered information" both within and across different distributed systems [33]:

"When I go back to history to look for something, I know I stored some information but I don't remember what it is, I know it's by Skype so I can go in there and look for it. If I'm not sure it was on Skype, it would be one of the places I'd be looking into."

Beyond meetings, most digital work is accomplished alongside shared collaboration tools, such as Slack or Trello (or, more recently, Clubhouse), that are often difficult to search and peruse, and may not provide facilities to expose work artifacts in a useful way. Past work has found that workers will repurpose such social systems to maintain awareness of other's state [36].

3 RESEARCH SYSTEMS SUPPORTING ARTIFACT REUSE

We have developed tools for mining and exposing digital content for reuse in the context of a variety of tasks, including distributed meetings, awareness systems, and multimedia chat.

3.1 Live distributed meetings

To make digital artifacts ready-to-hand for use in workaday tasks requires mining and visualizing multimedia content. Some early research tools have begun to make inroads on this problem.

The MixMeet system is a set of tools we built that analyze meeting content in real-time to make it easier to *find the right content* and *make content ready-to-hand* for sharing. The system provides a digital project room environment [15] that allows users to collaborate synchronously via web-based meetings as well as asynchronously via a shared document environment [17]. With the system, users can annotate and extract content from live meetings both to improve in situ communication as well as to guide others as they explore documents later. These online spaces allow a group of participants to create, view, annotate and edit documents in a flexible and persistent way that combines synchronous (e.g., showing slides in a meeting) and asynchronous use (e.g., perusing slides from a

[1] https://www.fastcompany.com/90449975/were-on-track-to-hit-1-5c-of-warming-in-2040-heres-how-it-will-change-our-work-lives

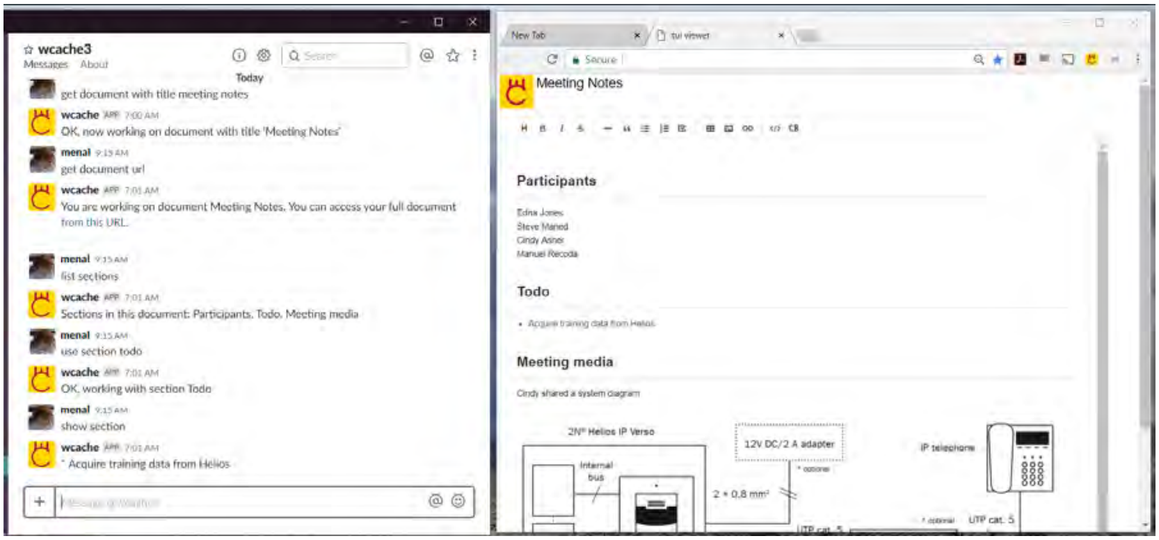

(a)

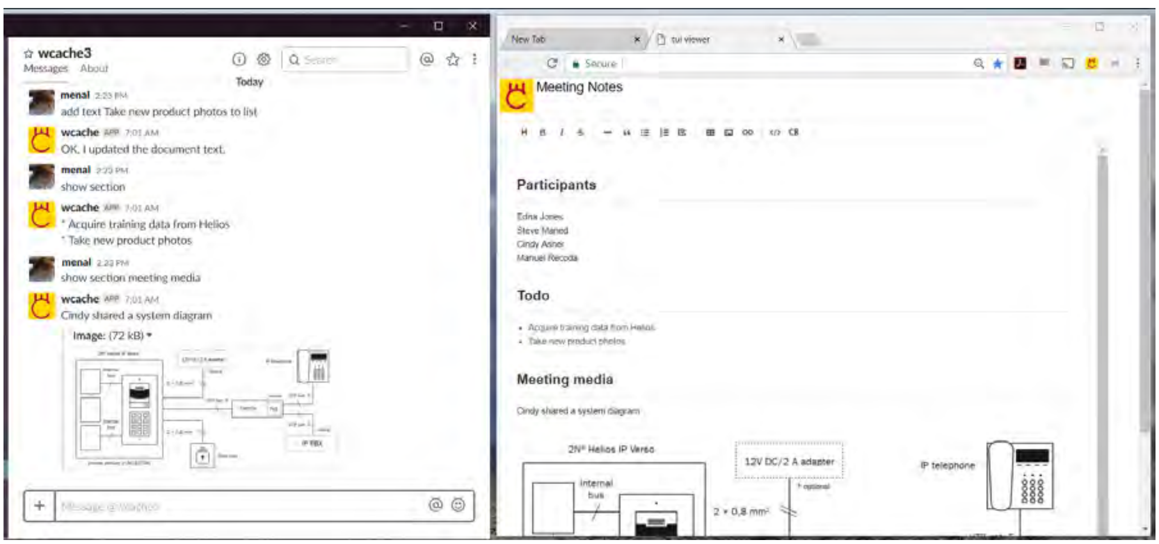

(b)

Figure 2: With a backend document analysis bot installed, users can manipulate and reuse content from live documents (note that the document and chat UIs are shown here side-by-side to help the reader understand how the system operates; normally a chat user would interact with only the chat client on a mobile device). (a) The user "loads" a document using text commands, lists sections, switches context to a section, and requests that the system show the section. The backend system will attempt to convert the document text to a view appropriate for the currently used chat client. (b) The user adds text to a section and requests that the system show the update. The user then requests to view a more complicated section.

past meeting or extracting a photo of a slide so others can refer to it later). This is accomplished using a variety of live and offline multimedia analysis tools to automatically index important keyframes, detect and extract searchable text, annotations, and user actions [16]. Coupled with such metadata, textual and image-based content was searchable and reusable for meeting notes, presentations, chats, and other asynchronous tasks [9].

These tools are focused on primarily desktop-based use. As remote meetings and mobile devices are becoming increasingly common, it is equally important to extend support to a variety of different user contexts, including mobile devices and wearables. We developed systems for these contexts that rely on a similar backbone of live meeting analysis to support content reuse, but adapt to the needs of mobile users potentially dividing their attention between a variety of different tasks while remaining engaged with meeting content [13]. Mobile users can view and search keyframes and reuse them in a backchannel chat, their own notes, or ask to share to the main meeting thread itself.

Further extensions to this work could augment backchannel sharing by taking advantage of commonplace technologies to let users attach multimedia-indexed chats to to any meeting or presentation. For example, systems could take advantage of inaudible audio [14] to seamlessly and automatically connect collocated mobile participants to a chat-space, and use live video- and audio-analysis to ingest multimedia content for immediate reuse.

In summary, tools that make it easier to reuse digital artifacts in the context of live meetings can help users in a variety of different contexts remain more engaged in the meeting's content.

3.2 Communicating shared interests

One of the key advantages of making digital artifacts more visible is for communicating shared interests to spur fruitful conversations.

The Hebb system arose from research that found that loosely coupled work groups are less likely to discover shared interests in the way that many tightly collocated groups do, such as by overhearing conversations or noticing paraphernalia [12]. This was true even if the two groups were working on similar topics and even if they are working in the same building, or even the same floor of the same building. To help these groups discover shared interests, we developed an email-based sensor to discover shared personal interests, as well as a public peripheral display and lightweight location-tracking system to convey those interests. Peripheral displays are a useful way of conveying information while *limiting impacts on other tasks*. In this system, they were situated near common social areas encouraged people to find and discuss shared topics in a casual atmosphere. When the location-tracking system detects that two participants in the space shared a common interest, it shows the shared topic on the public display. To *manage potential privacy issues* it shares links to the documents from which it derived interests on each participant's personal device.

Tang et al. explored a related approach, mining corporate file systems for documents, apps, or libraries and visualizing their similarities to expose shared interests and encourage networking [41].

Of course, a system dependent on proverbial "water coolers" to display common interests is not workable for remote teams. This approach could be modified to show information peripheral to each user's display (using projected icons [30] or more sophisticated but subtle approaches [29]), but would have to balanced against the impact of distractions. However, as Mark et al. found, "non work-related distractions...may actually play an integral role in the connected information workplace" [32]. Therefore, notifications of shared work interests coupled with lightweight actions that workers could take to pursue them may provide a healthy balance between focus and burnout.

3.3 Reusing digital artifacts in workaday chat

Knowledge workers now have a variety of rich multimedia tools, such as Slack or Microsoft Teams, that they can use to communicate and co-edit a variety of documents. Work groups using these tools often edit papers or presentations with other applications and import them into these chats to discuss them. People might comment on specific parts of the documents (e.g., "please look at the *introduction section*") or mention aspects of their structure (e.g., "see the *table on slide 3*"). People may also explain where files are, what has been uploaded, commented on, or edited. Users will also upload different versions of the same file and describe what was differentiates them [18].

However, the representation of digital artifacts in these discussions is often impoverished. For example, Figure 1a shows users discussing aspects of a Powerpoint slide deck, but the artifact itself is represented in the chat only as a single icon. While users could manually extract and share the contents of the deck, this is cumbersome when the conversation ranges over a variety of different slides, figures, tables, texts, and topics.

We designed the DocHandles system to take a different approach [18], using automated analysis to extract and index the content of shared documents as they are shared, which both helps users *find the right content* and *makes content ready-to-hand* for reuse in a conversation (Figure 1b). For example, the user can type, "please tweak the lighting in @fig2" and the system can automatically retrieve and show the content corresponding to Figure 2 in a document. Furthermore, the system can store enough context to disambiguate different artifacts based on recency (show the Figure 2 from the most recent document shared) or optionally it can list all of the matching content and allow the user to resolve the ambiguity. This approach represents a tradeoff between *limiting impacts on the main task (the chat)* and *managing privacy issues*. For example, the system can be configured to select content automatically for public documents but require manual selection for content to which not all parties have access (e.g., documents that the current user only shared in other conversations). Other simple rules connecting text terms to content can be used to resolve ambiguities (e.g., "slide" can always resolve to the most recently shared slide deck rather than the most recently shared document).

This work shows that tools that make digital artifacts easier to reuse in the context of multimedia chat systems can help users augment their messages while preserving the ease-of-use of text-centric communication.

3.4 Conversational documents

Techniques for exposing important content buried in digital documents can have utility beyond collaboration and awareness tools. In particular, while many collaboration tools support discussions about documents, they do not focus on creating and editing underlying content. There is a disconnect between messages and the production of the higher level, aggregated, structured final document. As a result, conversation and document creation tend to live in separate siloes, only to be bridged manually when users convert suggestions from collaboration tools to content actions. This disconnect makes it difficult to *find the right content* and *make content ready-to-hand*, resulting in unnecessary copy-and-paste interactions, lost information, and repeated work. To address this, another line of work takes advantage of the natural structure in enterprise documents to bridge the gap between conversational commands and specific document edits [10].

With this approach, a system with access to a multimedia enterprise chat channel can analyze documents that users have uploaded to the chat, segmenting individual regions that can be edited separately using domain-specific document segmentation methods (e.g., LexNLP[2]). A separate NLP framework can convert utterances from chat clients to edit commands using trained models. Typically, these frameworks rely on open-source NLUs (e.g., Snips[3] or Rasa[4]) to generate a large set of training data from a relatively small set of template-based document editing commands. They can also be extended and tailored to specific domains using live training data. Once the NLU is trained, commands can be connected to document editing APIs (e.g., Google Slides API[5]) to manipulate content.

With this basic framework in place, users can edit a wide variety of documents using a conversational approach, from structured forms [10] to more open-ended material (see Figure 2).

This approach shares some common-ground with later work from Iqbal et al. that analyses documents to detect microtasks that can be computed from phones [24]. In this system, tasks are extracted from a desktop application using a plugin. The tasks are also mapped to formalized workflows, making them more constrained but also allowing the system to provide more end user feedback.

A similar approach can map natural language instructions to user interface actions. For example, a system can parse a user instruction to, "enter starbucks for the SSID", to actual commands carried out in the UI [28]. In this case, instructions are first parsed into operations (e.g., "enter") and objects (e.g., "SSID") with arguments (e.g., "starbucks"); operations are then reduced to UI interactions (e.g., bringing an input field into focus and injecting text).

Other work suggests that this conversational approach may help users carry out document editing tasks in a variety of distributed attention scenarios [34].

3.4.1 Beyond chat. In many cases (e.g., listing real estate or short-term rentals, providing evidence for insurance claims, work-site inspections, or listing objects for sale) document editing involves injecting images of real-world artifacts into a template. Systems supporting this use case can couple the automatic document segmentation and editing frameworks described above with live video-based object recognition tools (e.g., using pre-trained image recognition models) to allow end users to scan an environment to catalog and capture media of objects-of-interest [8]. With this approach, users need not describe their edits to the system—they can simply load a "document" onto a mobile app, point their mobile device at a scene, and let the app automatically acquire images and inject them into the source document (see Figure 3).

Overall, this line of work shows that exposing digital artifacts can not only help users not only with activities peripheral to primary tasks, such as context awareness and chat, but can help them complete the tasks themselves.

4 CHALLENGES FOR SYSTEMS SUPPORTING DIGITAL ARTIFACT REUSE

In a woodworking shop, projects in process intermingle with each other and the tools and implements needed to carve and shape them. Communal tools are placed in areas that anyone can access them, or locked in cabinets when their access is restricted. Projects that workers want to keep more private are hidden away or keep in separate areas but can be moved easily to public view.

The goal of making components of previously monolithic digital content more accessible is to achieve a similar *done in the doing* [42] organization found in the shop. Unearthing artifacts makes them both read-to-hand for workaday tasks as well as a lightweight method of sharing interests and ideas.

In light of the systems discussed above, we here expand on the key challenges for digital artifact reuse.

4.1 Finding the right content

One of the key issues that has always kept remote work from being more widely accepted is the lack of "watercooler" chats between colleagues that lead to creative collaborations[6]. Digital artifacts can serve as props to engender these types of conversations. However, while we can develop algorithms to find overlaps in interest, determining the right prop to share among colleagues is complicated. When a user prints a paper and carries it with them to a communal space, they may expose enough of the document to allow others to get the gist of what they are reading at a casual glance, while not exposing potentially private aspects of the document. This allows anyone to strike up a conversation with that person about the general topic of the paper without overstepping. Such casual, informal shared interest awareness and privacy management is much more difficult to achieve in the digital realm.

One reason for the difficulty is the wide variety of methods people employ to organize their work. File organization strategies tend to vary widely and are "partially driven by aspects of the users' personality" [35]. Also, the effort that people put into organizing their own content tends to lead them to prefer personal file management over shared systems [5]. Furthermore, the overall number of files that people store is growing over time, along with file sizes and corresponding folder structures needed to manage them [19].

These issues make algorithmic approaches to finding the right file to share difficult for a number of reasons. How can the system determine the right content in the right document to share publicly? What types of context need to be considered in the predictive algorithm that makes such a determination? What context should actually be revealed alongside the artifact to make it understandable? How might shared context change to adjust for differences in relationships between individuals? How might methods of conveying artifacts and their contexts change for different relationships?

4.2 Making content *ready-to-hand*

It is worth considering how tools to find and expose digital artifacts might encourage (in the spirit of the Bauhaus) a "truth to [digital] materials" [40]. That is, any such tool should endeavor to expose content in a way that makes their re-use more compatible with their original design goals: exposed code can have ready-made hooks for importing into a software toolkit (including, for example, necessary libraries); images can expose themselves in editable formats where possible; video can include links to the underlying content used to create it, etc.

[2]https://contraxsuite.com/lexnlp/
[3]https://snips-nlu.readthedocs.io/
[4]https://rasa.com/
[5]https://developers.google.com/slides/
[6]https://remote.co/why-you-need-a-virtual-water-cooler/

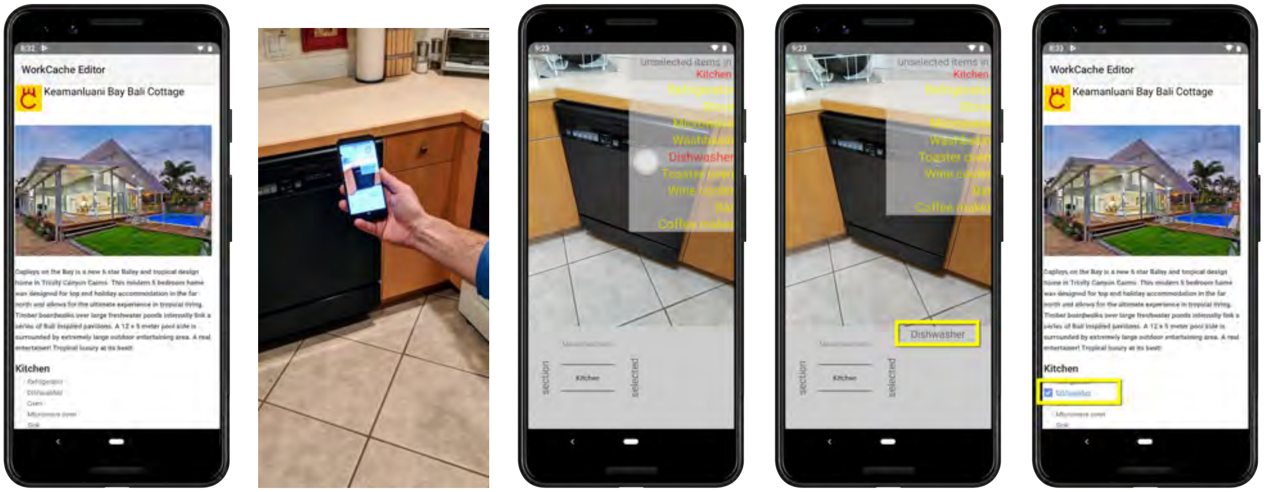

(a) (b) (c) (d) (e)

Figure 3: Editing documents using document structure analysis and live vision-based object detection. (a) A document requires a user to document different items in a rental listing. When the user first opens the mobile app, the system extracts all of the requested items for the entire document. (b) The user walks to the kitchen and selects "kitchen" in the section filter list to see only those items in the app. (c) The app detects the dishwasher in view and highlights its label in red. (d) The user selects this label and the system automatically saves a high quality image and moves the item label to the selected list. (e) The system automatically updates the linked template to indicate that the item is documented and inserts a link to the photo of the item.

While past work has investigated how to adapt interfaces for different contexts or user needs [1, 11], tools for reusing digital artifacts can also take into consideration how a file might best integrate into another workflow or visualization component.

4.3 Limiting impacts on other tasks

Exposing printed content or physical projects to colleagues is straightforward to control manually—people can physically move or adjust artifacts to make them more or less visible.

This is harder to do with digital artifacts. Tools that ask users to organize files differently for the purposes of making them available to others risks interfering with potentially highly rigorous personal file organization methods [35]. Other approaches might ask users to copy links to publicly sharable content, but this imposes obvious impacts on user's ability to remain focused on primary tasks [32]. Furthermore, systems that require changes to workaday habits can increase worker frustration and reduce their effectiveness.

This tradeoff implies that tools for extracting and sharing digital artifacts needs to maintain a model of user attention, focus, and distractability [23]. Systems should take care not to overburden users when making sharing decisions as well as mitigate distractions during times of focus, which can increase stress and decrease productivity across many types of workplaces [2].

4.4 Managing privacy issues

The problems with manual content sharing leads to the inevitable conclusion that systems for unearthing useful digital artifacts should be automated. While this may work fine for many private tools (such as editing personal documents) it has obvious privacy drawbacks in a collaborative context.

One of the core problems is that awareness and privacy are often presented to users as a tradeoff [37]. Past work has found that often users will often reduce their use of awareness systems, or turn them off entirely, when they wanted more privacy [22, 27]. Adding to the difficulty, privacy is deeply personal and context-dependent [26].

Semi-automated approaches may help find a balance between privacy and shared awareness. User modelling approaches to find task-breaks can influence visualization systems [3]. Publishing information is perhaps the thorniest problem. Returning to the woodworking example, we need methods to allow users to make decisions to expose or hide artifacts-in-progress that are as simple as moving a figure to the top of a table or a lower shelf while preserving their ease-of-access for the artifact's author. Future work might investigate semi-automated methods that embed such decisions within user's personal folder organization context [35].

5 CONCLUSION

Over time distributed work has grown in acceptance and prevalence, but only slowly. Many of the keys to productive, healthy work, including awareness of shared interests, informal sharing, and interpersonal communication, come relatively easy in collocated spaces but are difficult to recreate for remote workers. Even after decades of research, many of these issues remain unsolved. Regardless, global ecological issues are pushing knowledge workers into a future in which collocated work will be necessarily curtailed, or in some cases, elided entirely. Research into tools to support distributed workers is paramount, now more than ever. Research into tools for making digital artifacts accessible and reusable in particular are critical to help replace some of the rich interactions we now stand to lose.

ACKNOWLEDGMENTS

Thanks to Steve Whittaker for his feedback.

REFERENCES

[1] Mariam Kimiaei Asadi and J-C Dufourd. 2005. Context-aware semantic adaptation of multimedia presentations. In *2005 IEEE International Conference on Multimedia and Expo*. IEEE, 362–365.

[2] Anja Baethge and Thomas Rigotti. 2013. Interruptions to workflow: Their relationship with irritation and satisfaction with performance, and the mediating roles of time pressure and mental demands. *Work & Stress* 27, 1 (2013), 43–63.

[3] James "Bo" Begole, John C Tang, Randall B Smith, and Nicole Yankelovich. 2002. Work rhythms: analyzing visualizations of awareness histories of distributed groups. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work*. 334–343.

[4] Victoria Bellotti and Yvonne Rogers. 1997. From Web press to Web pressure: multimedia representations and multimedia publishing. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. 279–286.

[5] Ofer Bergman, Steve Whittaker, and Noa Falk. 2014. Shared files: The retrieval perspective. *Journal of the Association for Information Science and Technology* 65, 10 (2014), 1949–1963.

[6] Jed R Brubaker, Gina Venolia, and John C Tang. 2012. Focusing on shared experiences: moving beyond the camera in video communication. In *Proceedings of the Designing Interactive Systems Conference*. 96–105.

[7] Scott Carter, Matthew Cooper, Laurent Denoue, John Doherty, and Vikash Rugoobur. 2014. Supporting media bricoleurs. *interactions* 21, 4 (2014), 46–49.

[8] Scott Carter, Laurent Denoue, and Daniel Avrahami. 2019. Documenting Physical Objects with Live Video and Object Detection. In *Proceedings of the 27th ACM International Conference on Multimedia*. 1032–1034.

[9] Scott Carter, Laurent Denoue, and Matthew Cooper. 2015. Searching and browsing live, web-based meetings. In *Proceedings of the 23rd ACM international conference on Multimedia*. 791–792.

[10] Scott Carter, Laurent Denoue, Matthew Cooper, and Jennifer Marlow. 2018. FormYak: Converting forms to conversations. In *Proceedings of the ACM Symposium on Document Engineering 2018*. 1–4.

[11] Scott Carter, Amy Hurst, Jennifer Mankoff, and Jack Li. 2006. Dynamically adapting GUIs to diverse input devices. In *Proceedings of the 8th international ACM SIGACCESS conference on Computers and accessibility*. 63–70.

[12] Scott Carter, Jennifer Mankoff, and Patrick Goddi. 2004. Building connections among loosely coupled groups: Hebb's rule at work. *Computer Supported Cooperative Work (CSCW)* 13, 3-4 (2004), 305–327.

[13] Scott Carter, Jennifer Marlow, Aki Komori, and Ville Mäkelä. 2016. Bringing mobile into meetings: enhancing distributed meeting participation on smartwatches and mobile phones. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services*. 407–417.

[14] Scott A. Carter, Daniel Avrahami, and Nami Tokunaga. 2020. Using Inaudible Audio to Improve Indoor-Localization- and Proximity-Aware Intelligent Applications. arXiv:cs.HC/2002.00091

[15] Lisa M Covi, Judith S Olson, Elena Rocco, William J Miller, and Paul Allie. 1998. A room of your own: what do we learn about support of teamwork from assessing teams in dedicated project rooms?. In *International Workshop on Cooperative Buildings*. Springer, 53–65.

[16] Laurent Denoue, Scott Carter, and Matthew Cooper. 2015. Searching Live Meeting Documents" Show me the Action". In *Proceedings of the 2015 ACM Symposium on Document Engineering*. 195–198.

[17] Laurent Denoue, Scott Carter, Andreas Girgensohn, and Matthew Cooper. 2014. Building digital project rooms for web meetings. In *Proceedings of the 2014 ACM symposium on Document engineering*. 135–138.

[18] Laurent Denoue, Scott Carter, Jennifer Marlow, and Matthew Cooper. 2017. DocHandles: Linking Document Fragments in Messaging Apps. In *Proceedings of the 2017 ACM Symposium on Document Engineering*. 81–84.

[19] Jesse David Dinneen, Charles-Antoine Julien, and Ilja Frissen. 2019. The scale and structure of personal file collections. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.

[20] Christian Heath and Paul Luff. 1992. Media space and communicative asymmetries: Preliminary observations of video-mediated interaction. *Human–Computer Interaction* 7, 3 (1992), 315–346.

[21] Christian Heath and Paul Luff. 1996. Documents and professional practice: "bad" organisational reasons for "good" clinical records. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work*. 354–363.

[22] James D Herbsleb, David L Atkins, David G Boyer, Mark Handel, and Thomas A Finholt. 2002. Introducing instant messaging and chat in the workplace. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 171–178.

[23] Eric Horvitz, Carl Kadie, Tim Paek, and David Hovel. 2003. Models of attention in computing and communication: from principles to applications. *Commun. ACM* 46, 3 (2003), 52–59.

[24] Shamsi T. Iqbal, Jaime Teevan, Dan Liebling, and Anne Loomis Thompson. 2018. Multitasking with Play Write, a Mobile Microproductivity Writing Tool. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology (UIST '18)*. 411–422.

[25] Scott R Klemmer, Björn Hartmann, and Leila Takayama. 2006. How bodies matter: five themes for interaction design. In *Proceedings of the 6th conference on Designing Interactive systems*. 140–149.

[26] Scott Lederer, Jason I Hong, Anind K Dey, and James A Landay. 2004. Personal privacy through understanding and action: five pitfalls for designers. *Personal and ubiquitous computing* 8, 6 (2004), 440–454.

[27] Alison Lee, Andreas Girgensohn, and Kevin Schlueter. 1997. NYNEX portholes: initial user reactions and redesign implications. In *Proceedings of the international ACM SIGGROUP conference on Supporting group work: the integration challenge*. 385–394.

[28] Yang Li, Jiacong He, Xin Zhou, Yuan Zhang, and Jason Baldridge. 2020. Mapping Natural Language Instructions to Mobile UI Action Sequences. arXiv:cs.CL/2005.03776

[29] Kris Luyten, Donald Degraen, Gustavo Rovelo Ruiz, Sven Coppers, and Davy Vanacken. 2016. Hidden in plain sight: an exploration of a visual language for near-eye out-of-focus displays in the peripheral view. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 487–497.

[30] Blair MacIntyre, Elizabeth D Mynatt, Stephen Voida, Klaus M Hansen, Joe Tullio, and Gregory M Corso. 2001. Support for multitasking and background awareness using interactive peripheral displays. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*. 41–50.

[31] Wendy E MacKay. 1999. Is paper safer? The role of paper flight strips in air traffic control. *ACM Transactions on Computer-Human Interaction (TOCHI)* 6, 4 (1999), 311–340.

[32] Gloria Mark, Mary Czerwinski, and Shamsi T Iqbal. 2018. Effects of individual differences in blocking workplace distractions. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.

[33] Jennifer Marlow, Scott Carter, Nathaniel Good, and Jung-Wei Chen. 2016. Beyond talking heads: multimedia artifact creation, use, and sharing in distributed meetings. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 1703–1715.

[34] Nikolas Martelaro, Jaime Teevan, and Shamsi T. Iqbal. 2019. An Exploration of Speech-Based Productivity Support in the Car. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. 1–12.

[35] Charlotte Massey, Sean TenBrook, Chaconne Tatum, and Steve Whittaker. 2014. PIM and personality: what do our personal file systems say about us?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 3695–3704.

[36] Bonnie A. Nardi, Steve Whittaker, and Erin Bradner. 2000. Interaction and Outeraction: Instant Messaging in Action. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (CSCW '00)*. 79–88.

[37] Sameer Patil and Alfred Kobsa. 2009. Privacy considerations in awareness systems: designing with privacy in mind. In *Awareness Systems*. Springer, 187–206.

[38] Banu Saatçi, Roman Rädle, Sean Rintel, Kenton O'Hara, and Clemens Nylandsted Klokmose. 2019. Hybrid Meetings in the Modern Workplace: Stories of Success and Failure. In *International Conference on Collaboration and Technology*. Springer, 45–61.

[39] Abigail J Sellen and Richard HR Harper. 2003. *The myth of the paperless office*. MIT press.

[40] William Smock. 2009. *The Bauhaus ideal then and now: an illustrated guide to modern design*. Chicago Review Press.

[41] John C Tang, Clemens Drews, Mark Smith, Fei Wu, Alison Sue, and Tessa Lau. 2007. Exploring patterns of social commonality among file directories at work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 951–960.

[42] Peter Tolmie, James Pycock, Tim Diggins, Allan MacLean, and Alain Karsenty. 2002. Unremarkable computing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 399–406.

[43] Mark Weiser. 1999. The computer for the 21st century. *ACM SIGMOBILE mobile computing and communications review* 3, 3 (1999), 3–11.