Towards a Dynamic Multiscale Personal Information Space

Beyond application and document centered views of information

Amy Rae Fox
Department of Cognitive Science and UCSD Design Lab, University of California, San Diego, USA
amyraefox@ucsd.edu

Philip Guo
Department of Cognitive Science and UCSD Design Lab, University of California, San Diego, USA

Clemens Nylandsted Klokmose
Department of Digital Design and Information Studies, Aarhus University, Denmark

Peter Dalsgaard
Department of Digital Design and Information Studies, Aarhus University, Denmark

Arvind Satyanarayan
Visualization Group, Department of Electrical Engineering and Computer Science, MIT, USA

Haijun Xia
Department of Cognitive Science and UCSD Design Lab, University of California, San Diego, USA

James D. Hollan
Department of Cognitive Science and UCSD Design Lab, University of California, San Diego, USA
hollan@ucsd.edu

ABSTRACT

The historical moment when a person worked in front of a single computer has passed. Computers are now ubiquitous and embedded in virtually every new device and system, connecting our personal and professional activities to ever-expanding information resources with previously unimaginable computational power. Yet with all the increases in capacity, speed, and connectivity, our experiences too often remain difficult, awkward, and frustrating. Even after six decades of design evolution there is little of the naturalness and contextual sensitivity required for convivial interaction with computer-mediated information.

We envision a future in which the existing world of documents and applications is linked to a multiscale personalized information space in which dynamic visual entities behave in accordance with cognitively motivated rules sensitive to tasks, personal and group interaction histories, and context. The heart of the project is to rethink the nature of computer-mediated information as a basis to begin to fully realize the potential of computers to assist information-based activities. This requires challenging fundamental presuppositions that have led to today's walled gardens and information silos. Our goal is to catalyze an international research community to rethink the nature of information as a basis for radically advancing the human-centered design of information-based work and helping to ensure the future is one of convivial, effective, and humane systems. In this paper, we propose a new view of information systems, discuss cognitive requirements for a human-centered information space, and sketch a research agenda and approach.

Programming '20, March 23–26, 2020, Porto, Portugal

ACM ISBN 978-1-4503-7507-8/20/07.
https://doi.org/10.1145/3397537.3397542

CCS CONCEPTS

• **General and reference** → *Design*; • **Human-centered computing** → *HCI theory, concepts and models*; **Interaction paradigms**.

KEYWORDS

activity history, cognitive tools, co-adaptive systems, distributed cognition, dynamic media, human computer interaction, information visualization, instrumental paradigm, Vega-Lite, Webstrates

ACM Reference Format:
Amy Rae Fox, Philip Guo, Clemens Nylandsted Klokmose, Peter Dalsgaard, Arvind Satyanarayan, Haijun Xia, and James D. Hollan. 2020. Towards a Dynamic Multiscale Personal Information Space: Beyond application and document centered views of information. In *Conference Companion of the 4th International Conference on the Art, Science, and Engineering of Programming ('20), March 23–26, 2020, Porto, Portugal.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3397537.3397542

1 INTRODUCTION

For far too long we have conceived of thinking as something that happens exclusively in the head. Thinking happens in the world as well as the head. We think with things, with our bodies, with marks on paper, and with other people. Thinking is a distributed, socially-situated activity that exploits the extraordinary facilities of language, representational media, and embodied interaction with the world.

Today we increasingly think with computers. But the computers we think with are rapidly changing. The monolithic computer of the recent past is coming apart and being reassembled in myriad new forms. Computers are now ubiquitous and intertwined with every sphere of life. This evolution is accelerated by a radically changing cost structure in which the cost to use a thousand computers for a second or day is not appreciably more than to use one computer for a thousand days or seconds. Yet with all the advances in capacity, speed, and connectivity, using computers too often remains difficult, awkward, and frustrating. Even after six decades of design evolution, there is little of the naturalness, spontaneity, and contextual sensitivity required for convivial interaction with information.

We argue that this is a result of a legacy document and application-centered design paradigm that presupposes information is static and disconnected from the context of processes, tasks, and personal histories. ***We propose a new human-centered view of information: as dynamic entities whose representation and behavior are designed in accordance with the cognitive requirements of human activity.***

Both user and activity-centered design paradigms have drawn attention to the myriad contextual factors that bear upon human interaction with computational media. Still, it is rare that modern software design reaches the true scope of situated human activity, too often supporting only simple component tasks. This is why we endlessly switch between applications, fragmenting activity across time (sessions), physical (devices) and digital spaces (documents). The job of coordinating activity falls to the user, taxing our precious time and attentional resources. We argue that to produce truly convivial interactions, designers should focus not only on the applications with which users interact, but also on the *information* that underlies, connects, and integrates complex cognitive activities.

Our primary objective is to develop Human-Centered Information Spaces: dynamic computational environments linked to the existing world of information and operating according to empirically-grounded principles of behavior. We envision a future in which information *itself* is dynamic, interactive, and personalized to individuals, groups, contexts, tasks, and histories of interaction. Of course this future, one in which information transcends traditional application and device boundaries, cannot be achieved by a single research project. The motivation for this paper is to serve as a brief manifesto to help to catalyze a research community to begin to explore; to design, develop, and evaluate the radical alternative we propose in which information entities operate in accordance with cognitively inspired rules of behavior sensitive to the context of our past activity, intentions, perceptual and cognitive abilities. To help convey the future we envision, we begin by sketching a brief scenario.

1.1 A Scenario

Samantha leads a research group in microbiology. After returning from a conference, she is ready to continue writing a paper she started before her trip, but is struggling to remember where she left off. Samantha is an early adopter of technology, and has been doing her writing in a new prototype system—a human-centered information space for her research activity. She thinks of the software as a kind of desktop, a virtual workspace for her information work where she can organize and easily access the resources that support it. The system offers a novel interface to her digital information, consolidating her data (e.g., email, messages, calendars, web pages, notes, sketches, and analyses and visualizations) across applications. When she interacts with the information in her workspace, it seems to be alive, aware of when and how it was last used, and sometimes even why she was using it.

To get back to her writing, Samantha browses a timeline of her past working sessions. She vaguely remembers last searching the web for an article she'd once read, but now can't remember if she successfully found and referenced it. She scrubs through the visualization of her activity to before she left for the conference. This timeline, like the workspace itself, is multiscale, enabling her to move up and down levels of abstraction. She shifts to a level where only major activities (like a session of data analysis, writing, or web browsing) are displayed, and sees a familiar view of her text editor. When she clicks it, the image centers and thumbnails of all the other applications that were open cluster around it. There are too many to deal with, so she uses a search shortcut to enter a keyword she remembers was in the article. The thumbnails are filtered and she is left with a subset of browser tabs and a pile of PDFs. This reminds her that she had found the article she was looking for, and also downloaded a few others she thought might be related. When she hovers over the pile, she sees a sort of iconic summary—a dynamic montage of images from the documents.

She moves down a level of abstraction, and the PDFs show her how they had been interacted with. She realizes that she had skimmed a few, and identified one to read more deeply. She wants to send a list of the articles to her graduate student to investigate, so she selects the pile, and from the context menu that's triggered selects 'create list'. The titles of the PDFs are extracted into a list, which she quickly gestures over to the area of her workspace reserved for email. When she opens the PDF of the article she had been reading, the workspace asks if she wants to resume her text editor as well. The space rearranges to show her the editor beside the PDF, and automatically scrolls to the places in each document where she had last been active. She appreciates that this transition is slow and animated, first zooming out to where she can see both her current location and the target, before zooming in. She likes how this gives her a sense of location in the workspace.

Taken back to the documents of her previous writing task, she triggers a movie-like replay of those moments in time. She knows this sort of visual summary would be difficult, if not impossible, for anyone else to understand, but because it is derived from her history, it is evocative. In the replay she sees her navigation between reading part of the PDF article, and writing a paragraph in her paper. She suddenly feels as though she has been transported back in time to that point in her writing, even remembering her prior train of thought. Just in case she gets interrupted again, she uses a hotkey to tag this activity, jotting down a short description, before resuming her writing flow.

1.2 Challenges of Developing a Human-Centered Information Space

This idealized scenario glosses over a host of complex issues. *How can a parallel space of digital information be linked with existing applications? What information about past activities should be captured, and how should the context of this activity influence how information is represented? What rules should govern how information behaves in different contexts?* Although the scope and complexity of these issues are clearly beyond what can be addressed in any single research project, we hope to entice others to join in developing a research program strategically targeting the fundamental challenges required to develop a Human-Centered Information Space.

2 FOUNDATIONS OF THE HUMAN-CENTERED INFORMATION SPACE

The sophisticated cognition demanded by contemporary information work has outpaced innovation in user interfaces. In modern computing systems, data is still encapsulated in application silos, leaving users to shuttle files between applications, cobbling together workflows, requiring troublesome context switching and increasing attentional demands. ***In short, we lack a cognitively convivial space for intellectual work.***

For us, a Human-Centered Information Space is both an idea, and a computational environment. It is the idea of a cognitive workspace—a desktop for intellectual activity—reified as a computational environment that actively supports the coordination of information-based work. It develops awareness of the hierarchical structure of a user's action: how she accomplishes activities through discrete tasks across devices, programs, and working sessions. Through use, information in the environment will accumulate *context*: not only who accessed it and when, but concurrent activity and semantic relationships to other data. Just as awareness of the past influences human behavior, the content and context of the history of activity will drive the *behavior of information.* To the user, her information should seem alive, have awareness, know where it came from, how it got there, what it means—and behave accordingly. These representations and interactions will in turn guide the user's future action such that the struggle of resuming interrupted work is eased, much like finding a document is simplified by power of modern search engines. Importantly, the Human-Centered Information Space will not initially replace the user's ecosystem of applications and documents, but act as a home, a control center, a multi-modal but fundamentally spatial 'workshop' where information *across applications and documents* will converge with features that support the user in not only completing her tasks, but accomplishing activities.

In developing this concept we join with others (e.g., Kay [22], Victor [41, 43], and Berners-Lee [8]) in questioning the prevailing view of information. Our research agenda also draws inspiration from recent work of our collaborators Wendy Mackay and Michel Beaudouin-Lafon on *co-adaptive systems* and the *instrumental paradigm* [4, 5, 24]. The innovation of our approach lies in deriving general principles for the behavior of information in computational environments. We describe these behaviors as *cognitively convivial* because they are to be derived from the empirical science of cognition and designed to operate in ways attuned to our perceptual-cognitive abilities. We employ Illich's notion of *conviviality* [18] to emphasize that information should be lively, helpful, responsive, and enjoyable to interact with. By rethinking the nature of how computers mediate interaction with information, this project brings us closer to realizing the potential of computers to not only *assist*, but to *collaborate* in information-processing.

2.1 Empirical Grounding for a Human-Centered Information Space

Much like the form and movement of matter through space is governed by the laws of physics, the representation and interaction of information should be governed by the requirements of its processors: humans, and other such intelligent agents. A primary motivation for us is to mitigate unnecessary cognitive resource expenditures during complex information activities, thus making computer-based work more efficient and enjoyable. Our approach is informed by our prior work on activity-enriched computing [33–35], and contemporary research in Cognitive Science which emphasizes the fundamental importance of *space* to *thought*.

In her book, *Mind in Motion* [40], Barbara Tversky cogently describes decades of research on how we think about space—and how we use space to think[1]. Based on decades of empirical work in spatial cognition and external representation, Tversky formulates two principles for cognitively-driven design:

- ***Principle of Correspondence:*** *The content and form of the representation should match the content and form of the targeted concepts.*
- ***Principle of Use:*** *The representation should promote efficient accomplishment of the targeted tasks.*

These principles offer useful guidance for the design of a Human-Centered Information Space, and also an explanation of why many applications fall short. While designers endeavour to craft representations conducive to their target concepts, in reality, most interfaces are driven by classic design heuristics [31]. Initially derived from empirical research on human perception, these simple heuristics are challenged by the complexity of contemporary information work. Similarly, progress in software engineering has (appropriately) trended toward encapsulation, maintenance, and agility, yielding a rich ecosystem of micro-applications with powerful offerings toward narrowly defined feature sets. The consequence for users is the need to piece together workflows across applications. Information in each application *might* be meaningfully persistent but is presented in a different encapsulated form in each application. Representations, especially those in information systems are *tools for thinking*, and so just as our thoughts transform one idea into another, so should we be able to transform one representation of digital information, into another. We argue this requirement can only be met if information transcends applications and has the flexibility to dynamically alter its representation to support the changing state of a task as it evolves. This can only be accomplished if information representations are dynamic and re-mixable outside the walled gardens of applications.

2.2 Requirements For A Human-Centered Information Space

We believe the architecture of modern personal computing systems is insufficient for achieving our vision of a convivial, human-centered computing experience. The dominant unit of personal computing is the application/program. But people do not think or organize their work in terms of apps. We operate on goals, activities, and tasks. Thus, a truly human-centered architecture must support activity at the level at which people think about their work and assist in integrating it across applications. To accomplish this, we propose three requirements for a Human-Centered Information Space. We argue that information itself must become

[1]The importance of space in human thought also motivates our choice of the term 'Information Space' to conceptualize the computational environment we aim to build.

a first-class citizen: imbued with behavior (2.2.1), with the context of activity (2.2.2), made available outside applications (2.2.3). These requirements build upon each other in an additive fashion. When information has behavior, it can support and be responsive to changing context. When information and context are available outside the silo of an application, the resulting information space can be designed to scaffold the coordination of complex cognitive activities.

2.2.1 Information with Behavior: Animating Dead Bits Under Glass.

In his *Brief Rant on the Future of Interaction Design* [42], Bret Victor describes modern digital interaction as, "Pictures Under Glass ... an interaction paradigm of permanent numbness". These pictures are lifeless; dead bits of data to be swiped and tapped, until acted upon by some program. As with Victor's call[2] for active representations, we envision a space in which information is dynamic, capable of representing itself differently depending on its surrounding context. An example of information with behavior from our sample scenario is the collection of PDF documents, represented as a pile of thumbnails, a list of titles, or montage of key images. Each representation afforded Samantha a contextually-appropriate subset of interactions, and she could navigate between them to suit the structure of her thinking at any given time. ***We propose that information entities should be imbued with behavior, capable, for example, of dynamically changing their representation and interaction in accordance with empirically-grounded rules derived from human cognitive abilities.***

When we refer to *information with behavior*, we are speaking to both the way information is represented, and its capabilities for interaction. Prior work on multiscale visual representation, mechanisms of interaction, and specification of behavior demonstrate how these constructs can be realized in computational systems.

2.2.1.1 Visuospatial Representation. The Dynapad[3] system [2, 3], developed by Hollan and colleagues, realized a zoomable, multiscale virtual space with innovative user interactions that made information objects active and reactive, inspiring our vision of *information with behavior*. Chief among the facilities is *semantic zooming*—in which representations of information objects at different levels of granularity are determined by semantic factors rather than simple geometric scaling. The implementation of *lenses* enable filtered views of portions of the space, such that users have a sense of viewing the same information in different contexts. *Portals* allow connected views to other portions of the space that are independently pannable and zoomable, and *hyperlinks* afforded rapid movement to specific virtual locations (while maintaining object permanence and supporting wayfinding, unlike the experience of hyperlinks in web browsers). Though developed in 2006, Dynapad remains the best approximation of the spatiality of the dynamic, multiscale information environment we envision.

2.2.1.2 Interaction and Re-representation. The WritLarge system[4] [44], developed by Xia and colleagues, exemplifies the capacity for dynamic representation we argue are fundamental to convivial computing experiences. WritLarge provides a free-form canvas environment (on tablet and digital whiteboard) where users can flexibly transition between 'equivalent' representations of information along three structured axes: *semantic, structural ,and temporal*. For example, if a user scribbles a note to themselves on the canvas, they produce a series of vector-based strokes. Along the semantic axis, they can transition 'up' a layer of abstraction, and have the system recognize the text they have written, or 'down' a layer, and edit the strokes as pixels. Along the structural axis, the user can alter the organizational structure of the representation, while the temporal axis offers the ability to 'scrub' forward and backward in time. These features are equally powerful should the user scribble text on the canvas, or import an image. This eliminates the need for the user to fragment her activity (and therefore her thinking) from one application to another, perhaps typing her scribbled notes into a word processor, or exporting a vector-based image to a raster-editing application. The movement along these axes enables flexible transformation of representations, empowering her to express her thoughts at a natural level, rather than being confined to the fixed level applications are typically designed to support. Although WritLarge is a discrete application, it demonstrates the representational flexibility we seek in a separate space of information representations linked to existing information and applications.

2.2.1.3 Specification. One of the key insights in our concept of *information with behavior* is the simultaneous consideration of a representation and its afforded interactions. In the Information Visualization community, it is widely known that much more attention is paid to the nature of a representation than the nature of its interactions [37], in part because the languages we have for specifying such behaviors confound the two aspects. We plan to address this challenge by leveraging the high-level declarative grammar approach of Vega-Lite, [27, 36, 38][5] a widely used state-of-the-art substrate for developing dynamic interactive behavior for data. Vega-Lite employs a concise JSON syntax for rapidly generating visualizations by describing mappings between data fields and the properties of graphical marks. The Vega-Lite compiler automatically produces additional necessary components, such as scales, axes, and legends, and determines their properties (e.g., default color palettes) based on a set of carefully crafted rules for perceptual effectiveness. This approach allows specifications to be concise yet expressive. Vega-Lite enables authoring a wide range of interaction techniques including tooltips, brushing & linking, panning & zooming, focus+context, and interactive filtering. Critically, these techniques are not instantiated through top-down templates but rather with a set of bottom-up composable language primitives called "selections" and "selection transformations." These primitives allow simply describing the high-level intent of an interaction

[2]See Victor's talk entitled *Stop Drawing Dead Fish* (https://vimeo.com/64895205).
[3]Dynapad was the last version of our Pad++ [6, 7] zoomable multiscale development environment. The Pad++ software was non-exclusively licensed to Sony for $500K. It consists of a highly efficient C++ rendering core and an application development level using the Racket language [1], which supports language-oriented programming [11].

[4]WritLarge received a Best Paper Honorable Mention Award at the ACM CHI Conference in 2017. It is challenging to describe dynamic representations with text. A video of WritLarge (https://www.youtube.com/watch?v=6lWe9PvabAo) is available.
[5]Awarded best paper at VIS 2016.

(e.g., "highlight points on click") and the Vega-Lite compiler synthesizes the lower-level details such as registering event-handling callbacks and updating visual encoding rules. Vega-Lite provides a layered stack of declarative representations (all expressed as JSON). Having multiple levels of abstraction, with a correspondence mapping between them, is critical for enabling the rich cognitively convivial behaviors we envision, allowing end-users to work with the representation most suited to the task at hand.

2.2.2 Activity Context: Realizing the Potential of Activity History. As we move through the world, we leave rich traces of activity throughout our environment. These traces serve as the *context* for what, when, how, and potentially even indicators of why we do the things we do. Computationally, we record and make use of only a fraction of this context, storing it as metadata. In a document-centered paradigm, the user has easy access to administrative metadata: such as who created a document, of what type, and when it was last accessed. But imagine you could recall *all* of the times you accessed a particular document? Better yet, what if you knew what searches you performed while the document was open, what applications were in concurrent use, and how you developed the document's structure? In our sample scenario, Samantha's history of interaction was explicitly represented in a feature-rich timeline, and used to guide her interaction with information resources. The representation of some entities (PDF and browser search results) were enhanced with a history of her interaction (scrolling and click input). ***We argue that information should be responsive to the context afforded by a user's personal history of interaction.***

2.2.2.1 Capturing Activity History. The early Edit Wear and Read Wear [13] project pioneered the capture and visualization of activity history, allowing users to drive future interaction with a document from a context-appropriate representation of their past. As with subsequent research, however, the focus was on capturing history within a specific application. Today it is common for applications to include similar facilities (e.g., track-changes in Microsoft Word) to provide access to the modification history of a document. Of course, most modern computational workflows span multiple independent applications. A data scientist might search for open data sets in a web browser, write Python scripts within an IDE to scrape and wrangle that data, connect those scripts to black-box Unix command-line applications to run proprietary machine learning algorithms, and then feed the resulting models into a Jupyter Notebook with embedded Vega-Lite widgets to interactively visualize the results. The functions available in each application fail to support the user's higher level activity; the complete history of her interactions across all of these applications operate independently without awareness of one another.

To provide cognitively convivial interactions that support semantically meaningful higher level activities it is essential to first be able to capture cross-application interaction histories in a generic, application-independent manner. We call for operating-system-wide activity tracking, as exemplified by Guo and colleagues in the Burrito [12], Torta [28], and Porta[6] [29] projects. Each of these systems transparently monitors application activity at the OS level, creating a timestamped trace of activities such as which files were opened and/or modified, which system calls were executed, which GUI windows were opened/closed, and which sub-processes were launched. They also provide a layered architecture to connect this generic trace with application-specific tracers such as those that track editing/navigation actions within text editors (similar to Edit Wear and Read Wear [13]) and page interactions within web browsers.

2.2.3 Beyond Application Silos: Integrating Information. It is as difficult to conceptualize a computing paradigm *not* centered around documents and applications as it is to envision an interface paradigm not centered around windows, icons, menus, and pointers. Nonetheless, we believe it is time to move beyond aging metaphors and software structures convenient for the design and maintenance of machines, to those conducive to the thoughts and actions of users. A fundamental aspect of our vision is that the nature of representation of an information entity should be flexible, integral to the structure of the entity itself rather than a function of a specific application. The complementary design challenge lies in ensuring that the behavior of a representation provides the cross-task generality, consistency, and learnability that is too often missing from today's applications. To accomplish this, a distinct but connected space for representations is required. In our sample scenario, Samantha could retrieve PDFs from her search activity from both her hard drive (those she had downloaded) and web browser (those she was perusing), from a single point of access. ***We argue that a Human-Centered Information space needs access to data across applications. Information must become a first-class citizen in such a computational environment, owned by the user, available for re-representation and instrumental interaction.***

2.2.3.1 Application Integration. While the prior work of Guo [12, 28, 29] enabled the collection and visualization of data across web and native applications, these projects did not involve the integration of information into an independent space. Of course, in an application-centric paradigm, one might construe such consolidation as an instance of yet another application. The closest approximation of the cross-application integration we envision is realized in the realm of enterprise computing. To manage the integration of data and business processes across an ecosystem of enterprise-level applications requires standard data interchange formats, service oriented architectures, middleware infrastructure and business logic engines. One construal of the Human-Centered Information Space concept is as the personal-computing analog to the enterprise integration engine, plus a user interface. Such a solution does not exist in the world of personal computing, owing to the rapid pace, prolific number and democratic nature of application development. In recent years, an alternative solution to the problem of application silos has emerged via point to point integration services like *If-This-Then-That*, and Zapier. These (primarily cloud-based) task automation systems allow end-users to construct cohesive workflows without programming by mapping API triggers and data objects. Although this automation can ease the user's experience of manually porting information from one application to another, she is still limited to the features and representations of each application, and constrained by the expressiveness of each application's API.

[6] Best paper award UIST 2018.

2.2.3.2 Beyond Applications. A closer approximation of the integration concept is realized in the Webstrates framework [25], a novel browser-based approach for creating shareable dynamic media. Webstrates consists of a custom web server that serves pages, called webstrates—(web + substrates)—to ordinary web browsers. Each webstrate is a shared collaborative object, and changes to the webstrate's DOM, as well as changes to its embedded JavaScript code and CSS styles, are transparently made persistent on the server and synchronized with all clients sharing that webstrate. By sharing embedded code—behavior typically associated with browser-based software—can be collaboratively manipulated across devices. Webstrates employs *transclusion* [25, 30] to allow one webstrate to be embedded in another. This transforms the computing environment so as to support dynamic information and web-based collaboration. An initial example [25][7] was collaborative editing, enabling authors to interact with the same document via functionally and visually different editors. By making the DOM of web pages persistent and collaboratively editable, content and functionality become re-programmable and extensible. This is achieved through a conceptually simple change to the web stack that effectively blurs the distinction between applications and documents. The removal of the traditional hard distinction between applications and documents is crucial for the dynamic information environment we propose.

3 A PRAGMATIC APPROACH

A long-term research agenda for developing *human-centered information spaces* is necessarily ambitious. We argue it should be pursued incrementally and strategically by situating initial efforts in the context of specific domain activities and focusing on central problems of those domains. This approach allows one to narrow the enormous design space of information behaviors, while chipping away at the cognitive design and systems engineering challenges therein. It should also be informed by theory and data. In our work, we leverage the framework provided by distributed cognition [15, 17] and methods of cognitive ethnography [16]. Distributed cognition seeks to understand the organization of cognitive systems. Unlike explanatory theories of cognition, it extends the reach of what is considered 'cognitive' beyond the individual to encompass interactions between people and with resources and the material environment. Methods of cognitive ethnography build on this framework, providing tools for determining what things mean to the participants in an activity and for documenting the means by which these meanings are constructed.

Based on our prior work and ecosystem of current tools, we suggest that the activity of *data analysis and visualization* and the pervasive problem of *activity fragmentation* are an ideal testing ground for exploring these concepts.

3.1 Domain Activity: Analysis and Visualization in Computational Notebooks

One cannot study complex cognitive activity in the abstract. Data analysis and visualization are exploratory processes of extracting insights from data and communicating those insights to others [12, 21, 23, 39]. These processes have become more visible due to the widespread use of computational notebooks. The fact that analysis and visualization tasks typically cross application boundaries and require a characteristic mix of formal and informal information make computational notebook use an ideal domain of activity to focus our initial research efforts.

In a series of studies [35], Hollan and colleagues analyzed over a million Jupyter notebooks from a GitHub repository, selected 200 notebooks associated with academic publications for more detailed analysis, and interviewed 15 academic data analysts. A major finding was that many of the problems with notebooks result from a tension between exploration and explanation. Although notebooks provide tools for users to write rich computational narratives, analysts do not necessarily use them to great effect, as they continuously face a tension between exploring their data, or pausing to explain their process. Because users are torn between between using notebooks as a sandbox for exploratory work and as a repository for publication-ready analyses, they often delete intermediate and "failed" analyses. As a result, analysts lose the ability to retrace their steps at a later time, a task often crucial to reinstating the context of an interrupted analysis. In addition, the linear structure of current notebooks is constraining, failing to match the iterative and complex flow of most data analyses.

A primary feature of notebook environments like the increasingly popular JupyterLab [32] is the ability to combine code, commentary, and visualizations in a single document, rather than being scattered across multiple files. This unification attempts to reduce the time and effort needed to manage information and enable quickly retracing complex analyses or succinctly communicating them to colleagues. However, the scale, complexity, and exploratory nature of analysis means that notebooks quickly becoming "messy" and "too long" to understand. As a consequence users separate phases of their analyses into separate notebooks, fragmenting activity and reintroducing the issue of finding information across multiple files.

The challenges faced by users of computational notebooks are common to activities across many application domains. By situating efforts in the realm of data analysis and visualization in computational notebooks, one can address the management and navigation of multiple information resources (e.g., papers, notes, sketches, and the complex evolution of analyses and visualizations). Additionally, the infrastructure provided by Project Jupyter and the web-based JupyterLab environment means that engineering efforts will be concentrated on a platform that can be readily extended to domain-general web applications.

3.2 Domain Problem: Mitigating Activity Fragmentation

Research on activity-enriched computing [14] reveals that the need to coordinate activity over time and distributed media is a primary source of frustration and lost productivity in information work. Just as the need to employ multiple applications leads to increased complexity, rapidly expanding network connectivity brings a growing number and variety of *interruptions*—increasingly accepted as normal components of modern life. Observational studies of office workers reveal that real-life work is highly fragmented [9, 10, 19, 20, 26]. Mark et al. [26] found that during the

[7] Awarded best paper at UIST 2015.

course of a typical day information workers spend an average of only 12 minutes on any given task and most uninterrupted "events" average about 3 minutes in duration. Often, interruptions present serious challenges for resuming a task and re-familiarizing ourselves with the context of the interrupted activity. Many of the most challenging issues we face involve the disconnection of related information from our tasks and the associated problem of recreating the context required to resume activities that have been interrupted. Even when there aren't external interruptions, requirements of collaboration, time limitations, and the frequent requirement to switch back and forth between applications make interruptions unavoidable and fragment information resources. The preparation of this paper, for example, was distributed across email, sketches on whiteboards, text messages, recordings of video-conferences, annotated drafts, notes on paper, and the invisible histories of our individual activities.

In our view, a solution to fragmentation is unlikely to arise from a consolidation of functionality: a step back to the time of limited choice between feature-bloated programs. Rather, we argue that the sophistication of modern computing environments should be leveraged to combat this problem rather than exacerbate it. Because the problem of fragmentation is pervasive in the domain of analysis and visualization, we suggest this problem as a focus of evaluation for research efforts. The success of any prototype Human-Centered Information Space will be largely determined by its ability to mitigate the problem of fragmented activity, helping users recover from interruption and reinstate mental context.

4 TOWARD A MORE CONVIVIAL FUTURE

Often in science, it is the case that paradigm shifts are accompanied by changes in the prevailing metaphor. As you imagine the kind of computational environments with which you want to interact, we invite you to think beyond applications and documents, to envision alternative futures.

As a group of cognitive and computer scientists, we have come together jointly committed to this vision and convinced of the crucial importance of questioning the presupposition that information is fundamentally passive data disconnected from processes, tasks, context, and personal histories. As an underlying substrate, it is information that is processed, represented, acted upon and processed again—across applications, digital and physical media—to support human activity.

Although this may appear to be a problem for software engineering, it in reality depends equally on models of cognition. We argue an integration of methods from computing, behavioral and cognitive sciences is required to transform information into an active participant in distributed cognitive systems. In this way, we aim to re-designate the role that computers play in human life from tools with which we interact to (convivial) partners with whom we collaborate.

REFERENCES

[1] 2019. Racket. http://racket-lang.org/. http://racket-lang.org/

[2] Dan Bauer, Pierre Fastrez, and Jim Hollan. 2004. Computationally-Enriched 'Piles' for Managing Digital Photo Collections. In *Proceedings of the 2004 IEEE Symposium on Visual Languages - Human Centric Computing (VLHCC '04)*. 193–195.

[3] Daniel S. Bauer. [n.d.]. *The Cognitive Ecology of Dynapad, a Multiscale Workspace for Managing Personal Digital Collections*. Ph.D. Dissertation. La Jolla, CA, USA. Advisor(s) Hollan, Jim and Kirsh, David.

[4] Michel Beaudouin-Lafon. 2000. Instrumental Interaction: An Interaction Model for Designing post-WIMP User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (The Hague, The Netherlands) *(CHI '00)*. ACM, New York, NY, USA, 446–453. https://doi.org/10.1145/332040.332473

[5] Michel Beaudouin-Lafon. 2017. Towards Unified Principles of Interaction. In *Proceedings of the 12th Biannual Conference on Italian SIGCHI Chapter (CHItaly '17)*. ACM, 1:1–1:2.

[6] Benjamin B. Bederson and James D. Hollan. 1994. Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. 17–26.

[7] Ben B. Bederson, James. D. Hollan, Ken Perlin, Jon Meyer, David Bacon, and George Furnas. 1996. Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics. *Journal of Visual Languages and Computing* 7 (1996), 3–31.

[8] Tim Berners-Lee, James Hendler, and Ora Lassila. 2001. The Semantic Web. *Scientific American* 284, 5 (2001), 34–43. http://www.sciam.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21

[9] A.J. Bernheim Brush, Brian R. Meyers, Desney S. Tan, and Mary Czerwinski. 2007. Understanding memory triggers for task tracking. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (San Jose, California, USA) *(CHI '07)*. ACM, New York, NY, USA, 947–950. https://doi.org/10.1145/1240624.1240768

[10] Mary Czerwinski, Eric Horvitz, and Susan Wilhite. 2004. A diary study of task switching and interruptions. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (Vienna, Austria) *(CHI '04)*. ACM, New York, NY, USA, 175–182. https://doi.org/10.1145/985692.985715

[11] Matthew Flatt. 2012. Creating Languages in Racket. *Commun. ACM* 55, 1 (2012), 48–56.

[12] Philip J. Guo and Margo Seltzer. 2012. BURRITO: Wrapping Your Lab Notebook in Computational Infrastructure. In *Proceedings of the 4th USENIX Workshop on the Theory and Practice of Provenance* (Boston, MA) *(TaPP'12)*. USENIX Association, Berkeley, CA, USA. http://dl.acm.org/citation.cfm?id=2342875.2342882

[13] William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. 1992. Edit Wear and Read Wear. In *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems (Text and Hypertext)*. 3–9.

[14] Jim Hollan. 2012. Activity-Enriched Computing: Capturing and Mining Activity Histories. *IEEE Computer* 45, 10 (2012), 84–87.

[15] J. Hollan, E. Hutchins, and D. Kirsh. 2000. Distributed Cognition: Toward a New Theoretical Foundation for Human-Computer Interaction Research. *ACM Transactions on Human-Computer Interaction* (2000), 174–196.

[16] James D. Hollan and Edwin L. Hutchins. 2009. Opportunities and Challenges for Augmented Environments: A Distributed Cognition Perspective. In *Designing User Friendly Augmented Work Environments: From Meeting Rooms to Digital Collaborative Spaces*, Saadi Lahlou (Ed.). Springer.

[17] Edwin Hutchins. 1995. *Cognition in the Wild*. MIT Press, Cambridge.

[18] Ivan Illich and Anne Lang. 1973. Tools for conviviality. (1973).

[19] Shamsi T. Iqbal and Eric Horvitz. 2007. Disruption and recovery of computing tasks: field study, analysis, and directions. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (San Jose, California, USA) *(CHI '07)*. ACM, New York, NY, USA, 677–686. https://doi.org/10.1145/1240624.1240730

[20] Carlos Jensen, Heather Lonsdale, Eleanor Wynn, Jill Cao, Michael Slater, and Thomas G. Dietterich. 2010. The life and times of files and information: a study of desktop provenance. In *Proceedings of the 28th international conference on Human factors in computing systems* (Atlanta, Georgia, USA) *(CHI '10)*. ACM, New York, NY, USA, 767–776. https://doi.org/10.1145/1753326.1753439

[21] Sean Kandel, Andreas Paepcke, Joseph M. Hellerstein, and Jeffrey Heer. 2012. Enterprise Data Analysis and Visualization: An Interview Study. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012), 2917–2926. https://doi.org/10.1109/TVCG.2012.219

[22] Alan Kay and Adele Goldberg. 1977. Personal Dynamic Media. *Computer* 10, 3 (1977), 31–41.

[23] Mary Beth Kery, Amber Horvath, and Brad Myers. 2017. Variolite: Supporting Exploratory Programming by Data Scientists. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. ACM, New York, NY, USA, 1265–1276. https://doi.org/10.1145/3025453.3025626

[24] Clemens Nylandsted Klokmose and Michel Beaudouin-Lafon. 2009. VIGO: Instrumental Interaction in Multi-surface Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, 869–878.

[25] Clemens N. Klokmose, James R. Eagan, Siemen Baader, Wendy Mackay, and Michel Beaudouin-Lafon. 2015. Webstrates: Shareable Dynamic Media. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology*. 280–290.

[26] Gloria Mark, Victor Gonzalez, and Justin Harris. 2005. No Task Left Behind?: Examining the Nature of Fragmented Work. In *CHI '05: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*. 321–330.

[27] Dominik Moritz, Kanit Wongsuphasawat, Arvind Satyanarayan, and Jeffrey Heer. 2017. Introducing Vega-Lite 2.0. *Medium* (2017).

[28] Alok Mysore and Philip J. Guo. 2017. Torta: Generating Mixed-Media GUI and Command-Line App Tutorials Using Operating-System-Wide Activity Tracing. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (Québec City, QC, Canada) *(UIST '17)*. ACM, New York, NY, USA, 703–714. https://doi.org/10.1145/3126594.3126628

[29] Alok Mysore and Philip J. Guo. 2018. Porta: Profiling Software Tutorials Using Operating-System-Wide Activity Tracing. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology (UIST '18)*.

[30] T. H. Nelson. 1995. The Heart of Connection: Hypermedia Unified by Transclusion. *Communications of the ACM* (1995), 31–33.

[31] Jakob Nielsen. 1993. *Usability Engineering*. Academic Press.

[32] Fernando Perez and Brian E. Granger. 2015. Project Jupyter: Computational Narratives as the Engine of Collaborative Data Science. https://blog.jupyter.org/project-jupyter-computational-narratives-as-the-engine-of-collaborative-data-science-2b5fb94c3c58

[33] Adam Rule, Ian Drosos, Aurélien Tabard, and James D. Hollan. 2018. Aiding Collaborative Reuse of Computational Notebooks with Annotated Cell Folding. In *Proceedings of the 2018 ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW2018)*. 1–12.

[34] Adam Rule, Aurëlien Tabard, and Jim Hollan. In press. Stories We Tell Ourselves: Using Visual Histories to Make Sense of Interrupted Activities. *Human Computer Interaction* (In press).

[35] Adam Rule, Aurélien Tabard, and James D. Hollan. 2018. Exploration and Explanation in Computational Notebooks. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Article 32, 12 pages.

[36] Arvind Satyanarayan. [n.d.]. Vega-Lite: A Grammar of Interactive Graphics. https://vega.github.io/vega-lite.

[37] Arvind Satyanarayan. 2017. *The Reactive Vega Stack: Declarative Interaction Design for Data Visualization*. Ph.D. Dissertation. Stanford University.

[38] Arvind Satyanarayan, Dominik Moritz, Kanit Wongsuphasawat, and Jeffrey Heer. 2017. Vega-Lite: A Grammar of Interactive Graphics. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2017). http://idl.cs.washington.edu/papers/vega-lite

[39] John W. Tukey. 1977. *Exploratory Data Analysis*. Addison-Wesley.

[40] Barbara Tversky. 2019. *Mind in Motion*. Basic Books.

[41] Bret Victor. [n.d.]. Dynamicland: a communal computer, designed for agency, not apps, where people can think like whole humans. https://dynamicland.org/.

[42] Bret Victor. 2011. A Brief Rant on the Future of Interaction Design. http://worrydream.com/ABriefRantOnTheFutureOfInteractionDesign/

[43] Bret Victor. 2016. Stop Drawing Dead Fish. https://vimeo.com/64895205

[44] Haijun Xia, Ken Hinckley, Michel Pahud, Xiao Tu, and Bill Buxton. 2017. WritLarge: Ink Unleashed by Unified Scope, Action, & Zoom. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. 3227–3240.