CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information

Amy Rae Fox[1], Arvind Satyanarayan[2], Philip Guo[3], Haijun Xia[4], and James D. Hollan[5]

[1,3,4,5]Design Lab and Department of Cognitive Science, UC San Diego
[2]Visualization Group, Massachusetts Institute of Technology

Overview: 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, ranging from the omnipresent cellphone to the complex web of sociotechnical systems that envelop every sphere of personal and professional life. They connect our activities to ever-expanding information resources with previously unimaginable computational power. Yet with all the increases in capacity, speed, and connectivity, information-based activities 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 effective and convivial interaction with computer-mediated information.

The heart of the project is to rethink the nature of computer-mediated information as a basis to fully realize the potential of computers to assist information-based activities. We propose linking the existing information world of documents and applications to a multi-modal virtual workspace in which entities behave in accordance with cognitively-motivated rules of behavior sensitive to the context of activity history. Iterative development and evaluation of a series of prototypes designed to facilitate pervasive information tasks will inform design of dynamic personalized representations and the architecture of a human-centered information space.

Focusing on the domain of data analysis and visualization and the problem of activity fragmentation, this effort will (1) address the challenges of capturing and representing cross-application information activity, (2) design and evaluate activity-centered representations, (3) develop a cognitively-principled grammar of dynamic information behavior, and (4) integrate these features into a prototype information space realized as a web-based JupyterLabs extension to facilitate systematic evaluation in real-world use and encourage widespread adoption.

Intellectual Merit: The intellectual merit of this effort derives from challenging the presupposition that information is passive data, disconnected from processes, tasks, context, and histories. We propose to move beyond this legacy view and the silos of current applications to design, develop, and evaluate information entities that dynamically alter their appearance and behavior in response to task, history, and context. Developing effective principled composable general rules of behavior as the basis for a dynamic visual information space is the primary scientific focus. We refer to the rules as being *cognitively convivial* because they are specifically designed to operate in ways attuned to our perceptual and cognitive abilities.

Broader Impacts: The broader impacts of the proposed activity derive from the potential to radically improve the efficacy of computer-mediated information activities and reshape how we think with computers. In addition to research publication, all software will be open source and made available via a GitHub repository. Educational impact will result from training students in the interdisciplinary approaches required to design a new generation of information environments, providing research opportunities for graduate and undergraduate students. Student exchanges and interaction with collaborating international research labs will enrich student experience. We will continue to participate in UC San Diego outreach programs to engage students from diverse locations, including the innovative on-campus Preuss Charter School for low-income middle and high school students. Beyond publication, education, and outreach, our goal is to seed crystallization of a research community to further develop and evolve this novel approach for designing and interacting with dynamic information.

Keywords: activity history; application silos; cognitive principles and laws of cognition; cognitive tools; co-adaptive systems; distributed cognition; dynamic media; human computer interaction; information visualization; interruptions; instrumental paradigm; Vega-Lite; Webstrates.

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 a Human-Centered Information Space—a dynamic computational environment—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. Specifically, we will develop two prototype systems of this novel paradigm, tackling the engineering challenges to iteratively develop the computational environment and experimental challenges of deriving and testing rules for cognitively-driven information-behaviors. Of course this future, one in which information transcends traditional application and device boundaries, cannot be achieved by a single research project, so the motivation of the current effort is to begin to design, develop, explore, and evaluate this radical alternative. 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 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 re-arranges 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 suddently feels as though she has been transported back in time to that point in her writing, even remembering her prior train of thought. Just incase 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 begin to address them via the iterative design and evaluation of a series of prototypes, strategically selected to target the core research challenges of developing a Human-Centered Information Space. Before detailing our research plan, we describe the core concept of a Human-Centered Information Space, its empirical grounding and associated architectural requirements.

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. Specifically, we propose an environment that 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 context of activity history 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 replace the user's ecosystem of applications, but act as a home, a control center, a multi-modal but fundamentally spatial 'workshop' where information *across these applications* 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 [33], Victor [60, 58], and Berners-Lee [10]) 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* [5, 36, 6]. 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*[1] because they are derived from the empirical science of cognition and designed to operate in ways attuned to our perceptual-cognitive abilities. We employ the notion of *conviviality* 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[2], this project brings us closer to realizing the potential of computers to not only *assist*, but to *collaborate* in information-processing.

2.1 Requirements For A Human-Centered Information Space

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 argue that information itself must become a first-class citizen: imbued with behavior (2.1.1), with the context of activity (2.1.2), made available outside applications (2.1.3).

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. ***We begin our project by proposing three requirements for a Human-Centered Information Space. Our research plan centers on developing computational prototypes to evaluate the technical feasibility and cognitive efficacy of each of component.***

2.1.1 Information with Behavior: Animating Dead Bits Under Glass

In his *Brief Rant on the Future of Interaction Design* [59], 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 for active representations[3], 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 of dynamically changing their representation and interaction in accordance with empirically-grounded rules derived from human cognitive abilities.*

2.1.2 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.*

[1]For PI Hollan, the notion of a dynamic, cognitively convivial information arose from early work on simulation-based training systems [28] designed to create interactive worlds assisting people in developing mental models similar to those used by experts to reason qualitatively about complex physical systems. In this case, rules of behavior privileged conceptual over physical fidelity to create more interpretable instructional worlds. The same underlying notion motivated the design of multiscale interfaces [9] operating according to an alternate physics of an infinitely zoomable information space (see [26] for a brief history) and the beyond-being-there work [25] that questioned the efficacy of imitating face-to-face communication as a presupposition of supporting communications in electronic media.

[2]We use the terms *information* and *data* pragmatically and interchangeably, making no ontological commitment to a particular philosophical view. We do reserve the term *knowledge* to refer only to meaning actively constructed by a human.

[3]See Victor's talk entitled *Stop Drawing Dead Fish* (`https://vimeo.com/64895205`).

2.1.3 Beyond Application Silos: Integrating External 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, 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 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 harddrive (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 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 this project 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[4] [46, 48, 49], and contemporary research in Cognitive Science which emphasizes the fundamental importance of *space* to *thought*.

In her new book, *Mind in Motion* [57], our collaborator Barbara Tversky cogently describes decades of research on how we think about space—and how we use space to think[5]. 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 [44]. 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 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.

3 Research Agenda and Approach

Although the scope and complexity these research challenges are beyond what can be fully addressed in any single project, articulating them helps to characterize both the need *for* and the potential *of* the Human-Centered Information Space concept. We are fortunate to have assembled a research team ideally suited to this challenge. In addition to the PIs, with expertise in cognitive science, computer science, and psychology,

[4] *Acitivy-Enhanced Computing*, PI Hollan, NSF 1219829, 09/2013–08/2018.

[5] 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.

our team includes collaboration with the well-funded[6] European groups who designed the novel Webstrates software (described later) and continues its development (Clemens Klokmose from Aarhus University, Wendy Mackay from INRIA, and Michel Beaudouin-Lafon from University of Paris) as well as Aurélien Tabard from University of Lyon who collaborates with PI Hollan on activity capture and visualization. In addition, Barbara Tversky, an internationally-renowned expert in spatial cognition, and Jeff Zacks, a leading researcher in event cognition, will consult on the design of experiments to help ensure that our prototypes are informed by the most current empirical literature. Our Advisory Panel includes Alan Kay, an ACM Turing Award and Kyoto Prize winner who pioneered many of the ideas that led to this proposal; Gloria Mark, whose work has been central to understanding information work, interruptions, and the fragmentation of activity; Jaime Teevan, Chief Scientist at Microsoft who investigates information seeking activities; Sep Kamvar, who initiated the personalized search effort at Google; Viki Hanson, CEO of ACM, known for her work on accessibility and social impact will advise on our outreach efforts; and Fernando Pérez, who leads the Jupyter project will be our liaison for the evaluation of our software in conjunction with Jupyter Notebook. Together the PIs, collaborators, and advisors comprise a uniquely qualified team with the expertise and experience needed to successfully carry out the planned project. The preparation of this proposal further affirms to us the promise of our shared vision and the importance of the proposed research.

Although the long-term agenda for a Human-Centered Information Space is ambitious, we approach it strategically by choosing to situate our initial efforts in the context of a specific *domain activity*: data analysis and visualization in computational notebooks, with a focusing *domain problem*: mitigating activity fragmentation. In addition to constraining the scope of our initial exploration, this approach allows us to leverage the framework provided by distributed cognition [29, 24] and methods of cognitive ethnography [27]. The theory of distributed cognition seeks to understand the organization of cognitive systems. Unlike traditional theories of cognition, it extends the reach of what is considered *cognitive* beyond the individual to encompass interactions between people and with resources and materials in the environment. Methods of cognitive ethnography build on this framework, providing tools for determining what things mean to the participants in an activity and to document the means by which these meanings are constructed.

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 [56, 18, 32, 34]. 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 reported in [49][7], PI Hollan and collaborators analyzed over a million Jupyter notebooks available on GitHub, 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 [45] 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

[6] Beaudouin-Lafon and Mackay have each been awarded five year European Research Council Advanced Grants to support their work. These are prestigious long-term funding to pursue "groundbreaking high-risk projects" and will support collaboration with us on this project. Similarly, Klokmose has a three year (2018-2021) 4.1 million DKK ($638K) grant for "Literate Computing for Computational Thinking" to explore the computational notebook as a medium for thinking with the computer using prototypes built on Webstrates. He also has a three year (2017-2020) 3.9 million DKK ($607K) grant on "Shareable Dynamic Media in Design and Knowledge Work". To further support this international collaboration via student exchanges and workshops we plan to submit an NSF Accelerating Research through International Network-to-Network Collaborations (AccelNet) proposal.

[7] This paper received an honorable mention award (top 5% of submissions) at CHI2018.

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 our efforts in the realm of data analysis and visualization in computational notebooks, we 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 mean that our systems 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 [23] 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 [38, 12, 31, 13, 30]. Mark et al. [38] found that during the 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 refamiliarizing 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 proposal, 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 our selected domain of analysis and visualization, we will use this problem as a focus of evaluation for our research efforts. The success of our of 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 Research Plan

Our research plan (Figure 1) is structured by overlapping activities, each centered on developing and evaluating one of our requirements for a Human-Centered Information Space (Section 2.1). In stage one we leverage existing prototyping environments to develop a catalogue of behaviors for information entities. In stage two we leverage the web-based JupyterLabs infrastructure to develop a proxy Human-Centered Information Space and evaluate the use of system-wide Activity Context in notebook work. In stage three we extend this developing infrastructure to integrate data from external applications. This structure allows us to progressively evolve and evaluate component parts of the Human-Centered Information Space concept. Following a human-centered design approach, each stage will be informed by empirical research in applied cognition, design, and systems engineering.

As part of a recent project [49, 46] we have developed relationships with a number of research labs that are active users of Jupyter Notebooks (experts), as well as with instructors who require Jupyter Notebook use in their courses (novices). Over the course of our project we will perform three kinds of research activities with these groups serving as population samples. First, relying on methods of *cognitive ethnography*, we will conduct in-person observations, in-situ recording of activity, and interviews to understand how users distribute cognitive activity through the environment while analyzing and visualizing data. During the design of interfaces, we will use *participatory design* methods to ground our design-decisions in the context of how these artifacts might be used. Finally, to evaluate the efficacy of our designs and prototypes, we

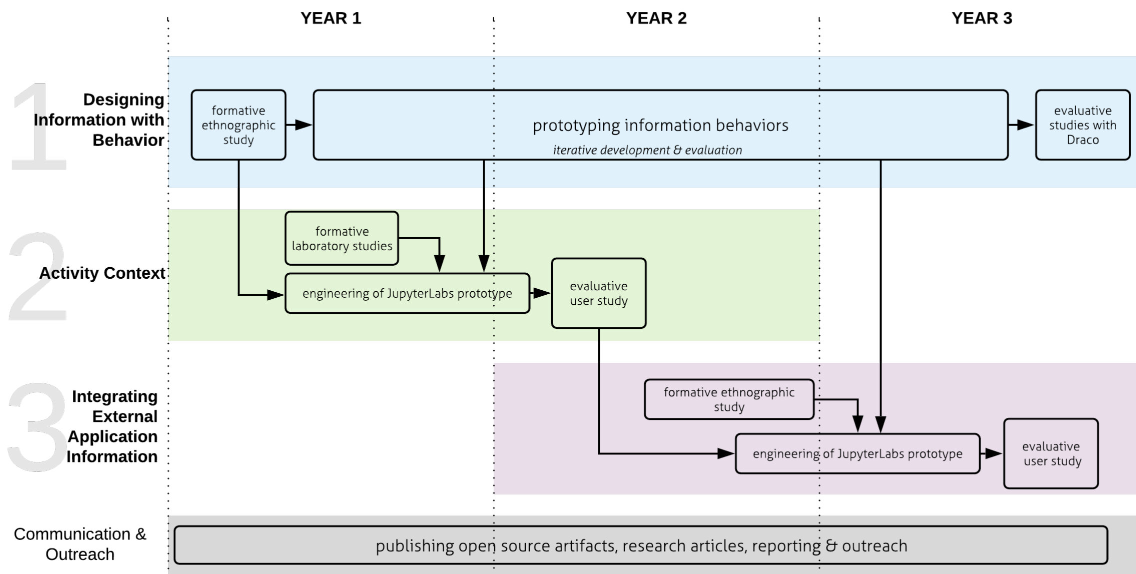

Figure 1: Timeline of Research Plan

will perform *situated user studies*. These studies are typically quasi-experimental and collect a mixture of objective measures (e.g., time on task) and subjective feedback. We will refer to these canonical research activities throughout the research plan.

4.1 Designing Information with Behavior

Developing principled composable rules of behavior for information entities is our primary scientific focus, as such will serve as an overarching activity for the duration of three-year project. The outcome of this phase will be a catalogue of primitive information behaviors. We develop this by leveraging a series of existing environments (DynaPad, WritLarge, Vega-Lite) to prototype domain-general behavioral primitives, while continuously integrating promising behaviors into our domain-specific prototypes (stages 2 and 3). Through this work we aim to evaluate our proposal that a Human-Centered Information space requires *information be imbued with behavior*.

4.1.1 Designing Information with Behavior: Prior Work

When we refer to *information with behavior*, we are speaking to both the way information is represented, and its capabilities for interaction. Prior work by PIs Hollan, Xia, and Satyanarayan on multiscale visual representation, mechanisms of interaction, and specification of behavior provide inspiration and a launchpad for our planned exploration of information with behavior.

Visual Representation. The Dynapad[8] system [3, 4], developed by PI 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 factors other 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 dynamic multiscale information environment we envision.

Interaction. The WritLarge system[9] [61], developed by Xia, exemplifies the dynamic representations

[8]Dynapad was the last version of our Pad++ [7, 8] zoomable multiscale development environment. The Pad++ software was nonexclusively licensed to Sony for $500K. It consists of a highly efficient C++ rendering core and an application development level using the Racket lanaguage [2], which supports language-oriented programming [14].

[9]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.

we propose. WritLarge provides a free-form canvas environment (on tablet and digital whiteboard) where users can flexibly transition between 'equivalent' representations of information in three ways: *semantically, structurally, and temporally.* 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.

Specifying Behavior. 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 InfoVis community, it is widely known that much more attention is paid to the nature of a representation than the nature of its interactions [51], 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 [52, 50, 40][10]—developed by co-PI Satyanarayan—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 (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.

4.1.2 Designing Information with Behavior: Planned Research

Exploring the Design Space. We will structure the exploration of this massive design space by prioritizing the problem of activity fragmentation faced by users of computational notebooks. Although many challenges were identified in our earlier work (such as the tensions between exploration and explanation), we did not examine the *low-level interactions* required to address them. Using *cognitive ethnography*, we begin here with a formative study of Jupyter notebook users recruited from labs at UCSD and MIT, designed to address the question: What are the existing behaviors of information in Jupyter Notebooks? We will interview, observe, and capture system-wide activity history of participants performing data analysis and visualization work. From these data we will identify types and sources of information entities accessed, as well as the interaction behaviors provided by Jupyter. We will map these information entities to the temporal dynamics of participants' workflows (e.g. working sessions, false starts, abandoned efforts, resumptions) and use this mapping to identify the information entities most implicated in users' breakpoints. Through this study we will prioritize entities with the greatest potential for improvement (with more dynamic behavior), as well as articulate our concept of 'information with behavior' in the context of an existing system.

Prototyping Behaviors. This work involves prototyping alternative futures: envisioning alternative representations for standard information entities, and how they might behave (with the user, and with other entities) on the basis of their context (such as goal-hierarchy, device, or history of interaction). We will leverage the existing environments from Dynapad and WritLarge to accelerate this prototyping work. The

[10] Awarded best paper at VIS 2016.

behavior specification techniques employed in Vega-Lite will serve as the foundation for our approach, as it provides a sufficiently robust vocabulary to systematically generate alternative interaction techniques for a given set of visual encodings (and vary their constituent properties). As a result, it allows us to begin to develop and empirically evaluate cognitive rules of interactive behaviors in conjunction with with the design of visual representations and mechanisms of interaction.

Evaluation. We will evaluate the conviviality of the information behaviors we design in two ways. First, we will take the promising candidate behaviors and continuously integrate them into our stage two and three prototypes, thus testing their efficacy in support of notebook activity via *user studies*. In addition, we will take the opportunity provided by working with the Vega-Lite specification language to test the relative efficacy of select behaviors in the context of Information Visualization activity. We will conduct a series of studies implementing behaviors in Draco [39], a constraint solving system developed by PI Satyanarayan's collaborators. First we will decompose the most promising behaviors into component primitives to be used to elaborate the interaction design space. Candidate points in this space will then be mapped to established task and analytic task taxonomies [21, 62, 11] to identify which ways a user might use a particular interactive behavior. Once two (or more) designs are mapped for a task, a comparative study can be run. We expect some of the studies could be conducted on crowdsourcing platforms such as Amazon Mechanical Turk, but others will benefit from our access to both expert and novice Jupyter users. We will collect quantitative metrics (e.g., task completion time, accuracy of both performing the interaction and the insights yielded) and qualitative metrics (e.g., which technique did users prefer), as well as demographic information to understand whether novice users prefer particular techniques over more advanced analysts, for example. Quantitative data will then be fed back to Draco to learn a series of weights that codify interaction effectiveness criteria. From these results, we will develop a grammar of interactive behavior for information entities. As the development of the substrate architecture, derivation of rules and development of prototypes will proceed in parallel, we will have the unique opportunity to evaluate the ecological validity of rules from the domain of visual analytics applied more generally to information tasks in our prototype information spaces. In our knowledge, this work will be the first large-scale study focused specifically on the cognitive effectiveness of behavior and interaction in data visualization, complementing the large body of work devoted to empirical assessment of visual encoding effectiveness [20, 55, 35].

4.2 Activity Context

The primary goal of our second phase of research is the development and evaluation of a prototype to support researchers performing data analysis and visualization in JupyterLabs, offering users access to the context of their system-wide activity history. Over the course of the second year, we will engineer a JupyterLabs extension to function as a proxy Information Space. We will gather field data and conduct participatory design to develop an activity-enriched interface for our human-centered space. The prototype will be evaluated based on its technical stability and effectiveness in supporting users' recovery from interruptions in activity. Through this work we aim to evaluate our proposal that a Human-Centered Information space requires *the context of activity history.*

4.2.1 Activity Context: Prior Work

PI Hollan's early Edit Wear and Read Wear [22] work pioneered capturing and visualizing activity history. 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. The approach we pursue in this phase of the project is based on operating-system-wide activity tracing, as pioneered by co-PI Guo in a series of systems: Burrito [18], Torta [41], and Porta[11] [42]. Each of these systems transparently monitor 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 [22]) and page interactions within web browsers.

4.2.2 Activity Context: Planned Research

Systems Engineering. The platforms developed by co-PI Guo [18, 41, 42] will serve as our starting point for systems engineering in this stage, enabling our prototype information space to access the full timestamped interaction history of a user working concurrently in multiple applications on their computer. A key challenge will be to determine the data structures most effective for persistence and real-time access of this expanded scope of metadata. We will leverage the JupyterLabs API to develop a web-based extension that can access this OS-level activity history, serving as a proxy Information Space for JupyterLabs-based work. The interface design will be informed by the results of our cognitive research and the functional requirement of representing cross-application history that is relevant to the notebook-based activity.

Cognition and Design. Though our own prior work has demonstrated that visual summaries of activity can help users recover from interruption [47, 48] knowing what aspects of history to present is an open empirical question. Theoretical models of event segmentation [63, 53] suggest our memory is better for the boundaries of events—the transitions—than for event contents. To apply this literature means that we need to determine what *boundary* versus *content* means when representing activities with digital information. Informed by the ongoing cognitive ethnography with our collaborative research subjects (as described in section 4), we will, as just one example, design two visualizations of personal activity history: one that gives preferential attention to information about the content of events, and a second emphasizing transitions between them. After gathering feedback on these visualizations from interviews with participants in our continuing fieldwork, we will design and run controlled between-subjects laboratory studies aimed at evaluating their efficacy in supporting resumption of an interrupted task. We expect the results, of what we anticipate will be a series of similar laboratory studies, will inform the applicability of the event segmentation literature to the design of activity history visualizations, suggesting what aspects of it we should emphasize in our stage two prototype.

Evaluation. The result of stage two research will be a prototype information space with activity history developed for JupyterLabs. We will evaluate the success of this prototype via a situated *user study* with expert notebook user based using: (1) technical stability (Does the prototype function as designed, is its performance stable, and can its architecture scale to support growing activity data?), and (2) the degree of support provided for resuming interrupted activity (as determined by user studies on both our participatory and control subject groups) as evaluation criteria. If resources allow, we may also run a laboratory user study with novice notebook users to assess the discoverability of the included information behaviors, and the extent to which they might support scaffolding of best practices in notebook use.

4.3 Integrating External Application Information

The primary goal of the third phase of research is the development and evaluation of a prototype information space that expands on the functionality developed in stage two to include an integration of information from external applications. This prototype will offer users the ability to not only access external information by launching an application through from a history of activity, but to manipulate the information *from* the external application in the Information Space itself. Over the course of the final year, we will expand the functionality of our JupyterLabs extension, leveraging the affordances of cross-domain affordances of the *Webstrates* [37] framework. Through this work we aim to evaluate our proposal that a Human-Centered Information space requires *access to data across applications.*

4.3.1 Integrating External Application Information: Prior Work

While the prior work of PI Guo [18, 41, 42] 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.

[11] Best paper award UIST 2018.

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.

A closer approximation of this integration concept is realized in the Webstrates framework. During a recent sabbatical, PI Hollan began a collaboration[12] with the developers of *Webstrates* (web + substrates) [37], a novel browser-based approach for creating sharable dynamic media. Webstrates consists of a custom web server that serves pages, called webstrates, to ordinary web browsers. Each webstrate is a shared collaborative object, and changes to the webstrate's DOM (Document Object Model), 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. An initial example [37][13] 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.

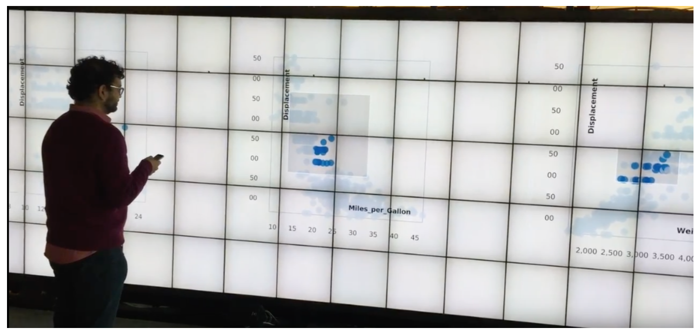

Figure 2: Webstrates. PI Satyanarayan interacts with a Webstrates prototype on a large display wall in Paris, using a cellphone to brush a section of data in a visualization.

Webstrates employs *transclusion* [43, 37] to allow one webstrate to be embedded in another. This transforms the computing environment so as to support dynamic information and web-based collaboration. Figure 2 shows an early Webstrates prototype being used for cross-device data brushing of information on a wall display using a cellphone. Importantly for our project, Webstrates provides mechanisms to enable cross-device collaboration with dynamic information entities. As a proof of concept demonstration of the type of multiscale information environment we envision, our Webstrates collaborators recently implement a Pad++-like zoomable information prototype in Webstrates that also allows pen-based collaborative sketching on an iPad.[14] In a similar proof of concept demonstration (Figure 2), PIs Satyanarayan and Hollan collaborated with Klokmose to explore integration of Vega-Lite facilities with Webstrates.

4.3.2 Integrating External Application Information: Planned Research

Cognition and Design. From a design perspective the greatest challenge of accessing information across applications is maintaining the user's understanding of provenance. To address this challenge, we will again leverage the method of *cognitive ethnography* with JupyterLabs users, this time focused on the user's workflow across the scope of an analysis and visualization project. Through the collection of observational, interview

[12]See *Facilities, Equipment, and Other Resources* in the Supplementary Documents for more information about our collaborators at Aarhus, University of Paris, and INRIA.

[13]Awarded best paper at UIST 2015.

[14]For the past several years graduate students from Aarhus have been funded by the Danish government to spend program-mandated graduate study abroad time at the DesignLab at UC San Diego, furthering this fruitful collaboration.

and activity tracking data we will determine what external applications (e.g. text editor, web browser) are most common to analysis workflows, and how users access them in concert with notebook-based work. Identification of these applications will constrain the scope of our systems engineering activities. We will also target our questioning to expose users' mental models of the provenance of notebook data, addressing questions such as: if I update information in application A, should I expect that the changes be reflected in application B?

Systems Engineering. The fundamental architectural challenge of this stage is supporting integration of information and interoperability across applications in a principled, standards-based fashion while maintaining the flexibility required for re-representation and instrumental interaction. Because our plan is to provide a unified treatment of all representations involved in an interactive behavior, this necessitates extensions to both Vega-Lite and Webstrates that we will complete in collaboration with our partners at Aarhus University and MIT. We leverage the Webstrates infrastructure to extend our stage two prototype with access to information from the external applications scoped via our cognitive ethnography, such that users can flexibly access and re-represent this information without the need to launch the external application.

Evaluation. The result of stage three research will a final prototype information space in JupyterLabs with both activity history and access to information from select external applications. As with prior stages, we will evaluate the success of this prototype via a situated *user study*, with respect to criteria for technical stability and functional support for resumption of interrupted activity.

4.4 Towards a Domain-General Web-Based Information Space

To conclude our project we will take what we have learned from the accumulation of formative and evaluative studies to critically assess our concept of a Human-Centered Information Space. We will appraise the architecture of our prototype systems, with the goal of describing the requirements to scale our approach beyond the scope of analyzing and visualization data in Jupyter notebooks. We will generalize what we have learned about supporting notebook activities to the more general case of supporting personal information management. Where appropriate, we will derive design principles and theoretical models. We will formalize our specification of information behaviors into a catalogue of primitive components that are useful for the generative design of representations and interactions. Finally, as the work we've described has been limited to the scope of supporting the activity of individuals, we will elaborate the series of challenges to be addressed in extending our work to support groups of users performing collaborative work, such as in the context of a research lab.

5 Broader Impacts

The broader impacts of the proposed activity derive from the promise to radically improve the efficacy of computer-mediated information activities and the exciting potential of cognitively convivial information entities to ultimately reshape how we think with computers. We have assembled an ideal group of collaborators (see Collaborating Institutions and Key Researchers) from active international laboratories with outstanding research and educational track records. Long-term impact will ultimately result from our students and from crystallizing an international research community to continue to evolve the Human-Centered Information Space concept and the software required to support convivial information-based activities.

Additional details of anticipated impact and plans to achieve it are provided by the Broadening Participation in Computing (BPC) Plan. We see a number of opportunities. One is to use the prototype information space we are developing to spark interest in computing for K-12 students. Another that we have been discussing with Advisor Board member Vicki Hanson, ACM CEO, is to pilot a project to support high-school and undergraduate students attending the annual ACM CHI conference.

6 Intellectual Merit

The intellectual merit of this effort derives from challenging the presupposition that information is passive data, disconnected from processes, tasks, context, and histories. We move beyond this legacy view and the silos of current applications to design, develop, and evaluate information entities that dynamically alter their appearance and behavior in response to a user's context. Developing principled composable domain-general rules for a novel dynamic information substrate is a primary focus. The heart of the project involves fundamentally rethinking the nature of computer-mediated information as a basis for fully realizing the potential of computers to assist information-based activities.

Another aspect of the project's intellectual merit is in establishing an international research collaboration spanning labs at Aarhus University, INRIA, MIT, UCSD, and University of Paris. This along with the members of the distinguished advisor board and consultants we have assembled not only strengthens the research but has potential to catalyze a broad research community to focus on developing a Human-Centered Information Space.

7 Results from Prior NSF Support

PI Hollan's recent NSF project (IIS-1319829; $499,976 9/2013-8/2018) on activity-enriched computing catalyzed our interest in the pervasive problem of interruption and fragmentation of activity as well as the challenge of restoring context. The *intellectual merit* of this research effort derived from developing and evaluating a novel activity-enriched computing framework that has become one of the bases for the proposed work. The associated publications most relevant to the current project are work on capturing activity [54] and analysis of Jupyter Notebooks [47, 48] use. The *broader impact* derives from training students in the interdisciplinary approach required to design activity-enriched computing applications and by providing research opportunities for both graduate and undergraduate students.

Hollan, Guo, and their colleagues recently began a new project (IGE-1735234; $498,751 9/2017-8/2020) mentioned earlier. The intellectual merit and broader impact center around responding to the challenge confronting virtually all graduate training programs to ensure their students have the computational skills required to function in increasingly data-intensive research domains and to be competitive in a rapidly evolving job market.

Co-PI Guo's NSF CRII award (IIS-1463864; $207,000; 9/2015–8/2018) has resulted in publications on scalable social systems that enable crowds of online learners to tutor one another in real-time on basic computer programming concepts [15, 19, 17]. Its main broader impact is expanding his open-source education platform `pythontutor.com` to reach over 5 million users across over 180 countries, many of whom do not have access to in-person tutoring resources. Co-PI Guo's NSF CAREER award (IIS-1845900; $479,860; 4/2019–3/2024) has only recently begun and not yielded publications yet.

Co-PIs Satyanarayan and Xia have not yet received NSF funding.

References

[1] Personal Communication, 2016.

[2] Racket. `http://racket-lang.org/`, 2019.

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

[4] Bauer, D. S. *The Cognitive Ecology of Dynapad, a Multiscale Workspace for Managing Personal Digital Collections*. PhD thesis, La Jolla, CA, USA.

[5] Beaudouin-Lafon, M. Instrumental interaction: An interaction model for designing post-wimp user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2000), CHI '00, ACM, pp. 446–453.

[6] Beaudouin-Lafon, M. Towards unified principles of interaction. In *Proceedings of the 12th Biannual Conference on Italian SIGCHI Chapter* (2017), CHItaly '17, ACM, pp. 1:1–1:2.

[7] Bederson, B. B., and Hollan, J. D. Pad++: A zooming graphical interface for exploring alternate interface physics. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (1994), pp. 17–26.

[8] Bederson, B. B., Hollan, J. D., Perlin, K., Meyer, J., Bacon, D., and Furnas, G. Pad++: A zoomable graphical sketchpad for exploring alternate interface physics. *Journal of Visual Languages and Computing 7* (1996), 3–31.

[9] Bederson, B. B., Stead, L., and Hollan, J. D. Pad++: Advances in multiscale interfaces. In *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems* (1994), vol. 2 of *SHORT PAPERS: Virtual and Visual Environments*, pp. 315–316.

[10] Berners-Lee, T., Hendler, J., and Lassila, O. The semantic web. *Scientific American 284*, 5 (2001), 34–43.

[11] Brehmer, M., and Munzner, T. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2376–2385.

[12] Brush, A. B., Meyers, B. R., Tan, D. S., and Czerwinski, M. Understanding memory triggers for task tracking. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (New York, NY, USA, 2007), CHI '07, ACM, pp. 947–950.

[13] Czerwinski, M., Horvitz, E., and Wilhite, S. A diary study of task switching and interruptions. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (New York, NY, USA, 2004), CHI '04, ACM, pp. 175–182.

[14] Flatt, M. Creating languages in racket. *Commun. ACM 55*, 1 (2012), 48–56.

[15] Gordon, M., and Guo, P. J. Codepourri: Creating visual coding tutorials using a volunteer crowd of learners. In *Visual Languages and Human-Centric Computing (VL/HCC), 2015 IEEE Symposium on* (Oct 2015), pp. 13–21.

[16] Guo, P. J. Online Python Tutor: Embeddable web-based program visualization for CS education. In *Proceedings of the 44th ACM Technical Symposium on Computer Science Education* (New York, NY, USA, 2013), SIGCSE '13, ACM, pp. 579–584.

[17] Guo, P. J. Codeopticon: Real-time, one-to-many human tutoring for computer programming. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology* (New York, NY, USA, 2015), UIST '15, ACM, pp. 599–608.

[18] Guo, P. J., and Seltzer, M. BURRITO: Wrapping Your Lab Notebook in Computational Infrastructure. In *Proceedings of the 4th USENIX Workshop on the Theory and Practice of Provenance* (Berkeley, CA, USA, 2012), TaPP'12, USENIX Association.

[19] Guo, P. J., White, J., and Zanelatto, R. Codechella: Multi-user program visualizations for real-time tutoring and collaborative learning. In *Visual Languages and Human-Centric Computing (VL/HCC), 2015 IEEE Symposium on* (Oct 2015), pp. 79–87.

[20] Harrison, L., Yang, F., Franconeri, S., and Chang, R. Ranking visualizations of correlation using weber's law.

[21] Heer, J., and Shneiderman, B. Interactive dynamics for visual analysis. *Queue 10*, 2 (2012), 30.

[22] Hill, W. C., Hollan, J. D., Wroblewski, D., and McCandless, T. Edit wear and read wear. In *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems* (1992), Text and Hypertext, pp. 3–9.

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

[24] Hollan, J., Hutchins, E., and Kirsh, D. Distributed cognition: Toward a new theoretical foundation for human-computer interaction research. *ACM Transactions on Human-Computer Interaction* (2000), 174–196.

[25] Hollan, J., and Stornetta, S. Beyond being there. In *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems* (1992), Perspectives on the Design of Collaborative Systems, pp. 119–125.

[26] Hollan, J. D. TIMELINES: Multiscale Zooming Interfaces: A Brief Personal Perspective on the Design of Cognitively Convivial Interaction. *Interactions 18*, 1 (2011), 71–74.

[27] Hollan, J. D., and Hutchins, E. L. Opportunities and challenges for augmented environments: A distributed cognition perspective. In *Designing User Friendly Augmented Work Environments: From Meeting Rooms to Digital Collaborative Spaces* (2009), S. Lahlou, Ed., Springer.

[28] Hollan, J. D., Hutchins, E. L., Rosenstein, M., and Weitzman, L. Steamer: An interactive simulation-based training system. *AI Magazine* (1984), 15–27.

[29] Hutchins, E. *Cognition in the Wild*. MIT Press, Cambridge, 1995.

[30] Iqbal, S. T., and Horvitz, E. Disruption and recovery of computing tasks: field study, analysis, and directions. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (New York, NY, USA, 2007), CHI '07, ACM, pp. 677–686.

[31] Jensen, C., Lonsdale, H., Wynn, E., Cao, J., Slater, M., and Dietterich, T. G. 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* (New York, NY, USA, 2010), CHI '10, ACM, pp. 767–776.

[32] Kandel, S., Paepcke, A., Hellerstein, J. M., and Heer, J. Enterprise data analysis and visualization: An interview study. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (Dec. 2012), 2917–2926.

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

[34] Kery, M. B., Horvath, A., and Myers, B. Variolite: Supporting exploratory programming by data scientists. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2017), CHI '17, ACM, pp. 1265–1276.

[35] Kim, Y., and Heer, J. Assessing effects of task and data distribution on the effectiveness of visual encodings. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 157–167.

[36] Klokmose, C. N., and Beaudouin-Lafon, M. Vigo: Instrumental interaction in multi-surface environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2009), CHI '09, ACM, pp. 869–878.

[37] Klokmose, C. N., Eagan, J. R., Baader, S., Mackay, W., and Beaudouin-Lafon, M. Webstrates: Shareable Dynamic Media. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software Technology* (2015), pp. 280–290.

[38] Mark, G., Gonzalez, V., and Harris, J. 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* (2005), pp. 321–330.

[39] Moritz, D., Wang, C., Nelson, G., Lin, H., Smith, A. M., Howe, B., and Heer, J. Formalizing visualization design knowledge as constraints: Actionable and extensible models in draco. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2019).

[40] Moritz, D., Wongsuphasawat, K., Satyanarayan, A., and Heer, J. Introducing Vega-Lite 2.0. *Medium* (2017).

[41] Mysore, A., and Guo, P. J. 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* (New York, NY, USA, 2017), UIST '17, ACM, pp. 703–714.

[42] Mysore, A., and Guo, P. J. Porta: Profiling software tutorials using operating-system-wide activity tracing. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (2018), UIST '18.

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

[44] Nielsen, J. *Usability Engineering*. Academic Press, 1993.

[45] Perez, F., and Granger, B. E. Project Jupyter: Computational Narratives as the Engine of Collaborative Data Science, July 2015.

[46] Rule, A., Drosos, I., Tabard, A., and Hollan, J. D. 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* (accepted), CSCW2018.

[47] Rule, A., Tabard, A., Boyd, K., and Hollan, J. Restoring the Context of Interrupted Work with Desktop Thumbnails. In *37th Annual Meeting of the Cognitive Science Society* (2015).

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

[49] Rule, A., Tabard, A., and Hollan, J. D. Exploration and explanation in computational notebooks. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2018), CHI '18, ACM, pp. 32:1–32:12.

[50] Satyanarayan, A. Vega-Lite: A Grammar of Interactive Graphics. `https://vega.github.io/vega-lite`.

[51] Satyanarayan, A. *The Reactive Vega Stack: Declarative Interaction Design for Data Visualization*. PhD thesis, Stanford University, 2017.

[52] Satyanarayan, A., Moritz, D., Wongsuphasawat, K., and Heer, J. Vega-Lite: A Grammar of Interactive Graphics. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2017).

[53] Shipley, T. F., and Zacks, J. M. *Understanding Events: From Perception to Action*. Oxford University Press, 2008.

[54] Stusak, S., Tabard, A., Sauka, F., Khot, R. A., and Butz, A. Activity Sculptures: Exploring the Impact of Physical Visualizations on Running Activity. *IEEE Trans. Vis. Comput. Graph. 20*, 12 (2014), 2201–2210.

[55] Szafir, D. A. Modeling color difference for visualization design. *IEEE transactions on visualization and computer graphics 24*, 1 (2018), 392–401.

[56] Tukey, J. W. *Exploratory Data Analysis*. Addison-Wesley, 1977.

[57] Tversky, B. *Mind in Motion*. Basic Books, 2019.

[58] Victor, B. Dynamicland: a communal computer, designed for agency, not apps, where people can think like whole humans. `https://dynamicland.org/`.

[59] Victor, B. A Brief Rant on the Future of Interaction Design, 2011.

[60] Victor, B. Stop Drawing Dead Fish, 2016.

[61] Xia, H., Hinckley, K., Pahud, M., Tu, X., and Buxton, B. Writlarge: Ink unleashed by unified scope, action, & zoom. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), CHI '17, pp. 3227–3240.

[62] Yi, J. S., ah Kang, Y., Stasko, J. T., Jacko, J. A., et al. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization & Computer Graphics*, 6 (2007).

[63] Zacks, J. M., Tversky, B., and Iyer, G. Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology: General 130* (2001), 29.

James D. Hollan
NSF Biographical Sketch

Department of Cognitive Science and Design Lab, UC San Diego, La Jolla, CA 92093
Email: hollan@ucsd.edu *Web:* hci.ucsd.edu/hollan *Lab:* designlab.ucsd.edu *Tel:* +1-858-534-8156

Professional Preparation

University of Florida, Gainesville, FL; Mathematics & Psychology; BS, 1969, MS, 1972, PhD, 1973

Stanford University, Palo Alto, CA; Artificial Intelligence; Postdoctoral Fellow, 1973-74

Appointments

2017–Present: Distinguished Professor of Cognitive Science, Department of Cognitive Science, UC San Diego

2016 Quarter: Visiting Professor, University of Paris and INRIA

2014–Present: Founding Co-Director, Design Lab, UC San Diego

2009 Quarter: Visiting Professor, Computer Science Division, UC Berkeley

2005 Quarter: Visiting Professor, Computer Science Department, Stanford University

1997–Present: Professor, Department of Cognitive Science & Department of Computer Science and Engineering, UC San Diego

1997–2014: Founding Co-Director, Distributed Cognition and Human-Computer Interaction Lab, UC San Diego

1993–1997: Professor and Chair, Computer Science Department, University of New Mexico

1989–1993: Director, Computer Graphics and Interactive Media Research Group, Bellcore

1987-1989: Director, Human Interface Laboratory, Microelectronics and Computer Technology Corporation (MCC)

1977–1987: Lecturer to Associate Research Cognitive Scientist, UC San Diego and Director, Future Technologies, NPRDC

1974–1977: Assistant Professor, Departments of Computer Science and Psychology, Clarkson University

Publications

Five publications related to proposed project

1. Exploration and Explanation in Computational Notebooks. A. Rule, A. Tabard, and J. Hollan. *Proceedings of CHI 2018*, 32:1–32:12, 2018.
2. Using Visual Histories to Reconstruct the Mental Context of Suspended Activities, A. Rule, A. Tabard, and J. Hollan. *Human-Computer Interaction*, 1-48, 2017.
3. Thinking with Computers. J. Hollan. *SIGCHI Lifetime Research Award, ACM Conference on Human Factors in Computing Systems*, 817-820, 2015.
4. ChronoViz: A System for Supporting Navigation and Analysis of Time-Coded Data. A. Fouse, N. Weibel, E. Hutchins, and J. Hollan. *Proceedings of CHI 2013*, 299-304, 2011.
5. Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physic. B. Bederson, J. Hollan, K. Perlin, J. Meyer, D. Bacon, and G. Furnas. *Journal of Visual Languages and Computing*, 3-31, 1996.

Five additional publications

1. Traces: A Flexible, Open-Source Activity Tracker for Workplace Studies. A. Rule, A. Tabard, and J. Hollan. *ACM Computer Supported Cooperative Work and Social Computing*, 1-6, 2016.
2. Activity-Enriched Computing: Capturing and Mining Activity Histories. J. Hollan, *Computer*, 84-87, 2012.
3. Opportunities and Challenges for Augmented Environments: A Distributed Cognition Perspective. J. Hollan and E. Hutchins. In *User Friendly Environments: From Meeting Rooms to Digital Collaborative Spaces (S. Lahlou, Editor)*, 237-259, 2009.
4. Distributed Cognition: Toward a New Theoretical Foundation for Human-Computer Interaction. J. Hollan, E, Hutchins, and D. Kirsh, In *Human-Computer Interaction in the New Millennium (J. M. Carroll, Editor)*, 75-94, 2000.
5. Beyond Being There. J. Hollan and S. Stornetta, *ACM CHI'92*, 119-125, 1992.

Synergistic Activities

Awards and Honors

ACM SIGCHI Lifetime Research Award (2015)

The SIGCHI Lifetime Research Award is presented to individuals for outstanding contributions to the study of human-computer interaction. This award recognizes the very best, most fundamental and influential research contributions. It is awarded for a lifetime of innovation and leadership. The criteria for the award are cumulative contributions to the field, influence on the work of others, and development of new research directions.

Chinese Academy of Science Invitation for Lectures in Beijing (2010)

Elected to ACM SIGCHI Academy (2003)

ACM SIGCHI Academy is an honorary group of individuals who have made extensive contributions to the study of human-computer interaction and who have led the shaping of the field.

Service to Research Community

Associate Editor, ACM Transactions on Human-Computer Interaction; Symposium Chair, ACM Information Visualization Conference; Co-Chair DARPA Information Science and Technology Study (ISAT) on Visualizing Information; Program Chair, ACM Multimedia Conference; Reviewer for CHI, UIST, CSCW, HCI journals, and NSF.

Service to University Community

Review Committee for Dean of Social Sciences, Chair of Search Committee for Chair of Department of Cognitive Science, Chair of Faculty Search Committee Department of Cognitive Science, Dean of Social Sciences Search Committee, Founding Co-Director of Design Lab, Led creation of cross-divisional undergrad Design Minor and leading development of a PhD specialization in Design, Founding Co-Director of Distributed Cognition and Human-Computer Interaction, Calit2 Divisional Council, Led creation of HCI specialization in Department of Cognitive Science, Mentor for new HCI and Design faculty in Department of Cognitive Science, Mentor for Design Lab Postdocs.

Judith Fan
Assistant Professor
9500 Gilman Dr MC 0109, La Jolla, CA 92093
jefan@ucsd.edu

(a) Professional Preparation
A list of the individual's undergraduate and graduate education and postdoctoral training as indicated below:

Harvard College	Cambridge, MA	Neurobiology	AB 2010
Princeton University	Princeton, NJ	Psychology	PhD 2016
Stanford University	Stanford, CA	Psychology	Postdoc 2017-2019

(b) Appointments

07/2019 – *present*	Assistant Professor, Department of Psychology, UC San Diego
02/2017 – 06/2019	Postdoctoral Scholar, Department of Psychology, Stanford University
06/2016 – 01/2017	Postdoctoral Research Associate, Princeton Neuroscience Institute
08/2011 – 05/2016	PhD Candidate, Department of Psychology, Princeton University

(c) Publications

(i) Publications most closely related to the proposed project

Fan J., Hawkins R., Wu M., & Goodman N. (2019). Pragmatic inference and visual abstraction enable contextual flexibility during visual communication. *Computational Brain & Behavior.* *https://doi.org/10.1007/s42113-019-00058-7*

Hawkins R., Sano, M., Goodman N., & Fan J. (2019). Graphical convention formation during visual communication. *Proceedings of the 41st Annual Meeting of the Cognitive Science Society.*

Mukherjee K., Hawkins R.., & Fan J. (2019). Communicating semantic part information in drawings. *Proceedings of the 41st Annual Meeting of the Cognitive Science Society.*

Long B., Fan J., Chai R., & Frank M. (2019). Developmental changes in the ability to draw distinctive features of object categories. *Proceedings of the 41st Annual Meeting of the Cognitive Science Society.*

Fan, J. E., Yamins, D. L., & Turk-Browne, N. B. (2018). Common object representations for visual production and recognition. *Cognitive Science*, *42*(8), 2670-2698.

(ii) Other Significant Publications

Cullen, S., Fan, J., van der Brugge, E., & Elga, A. (2018). Improving analytical reasoning and argument understanding: a quasi-experimental field study of argument visualization. *NPJ science of Learning*, *3*(1), 21.

Achlioptas, P., Fan, J., Hawkins, R. X., Goodman, N. D., & Guibas, L. J. (2019). ShapeGlot: Learning Language for Shape Differentiation. *International Conference on Computer Vision (ICCV)*. https://arxiv.org/abs/1905.02925.

Fan J., Dinculescu M., & Ha D. (2019). Collabdraw: An environment for collaborative sketching with an artificial agent. *Proceedings of the 2019 ACM SIGCHI Conference on Creativity and Cognition.* https://dl.acm.org/citation.cfm?id=3326578

Fan J., Hutchinson, J., and Turk-Browne, N. (2016) When past is present: Substitutions of long-term memory for sensory evidence in perceptual judgments. *Journal of Vision*. 16(8), 1-12.

Fan J. and Turk-Browne, N. (2016) Incidental biasing of attention from long-term memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. 42(6), 970-977.

(d) Synergistic Activities

1. Mentor, Stanford Center for the Study of Language & Information (2017, 2018)
2. Resident Graduate Advisor, Wilson College, Princeton University (2012-16)
3. Graduate Fellow, Princeton Program in Cognitive Science (2015-16)
4. Psychology Princeton Senior Thesis Writing Group Leader (2013-14)
5. Member, Princeton Psychology Colloquium Committee (2012-14)

Biographical Sketch

Philip J. Guo
Assistant Professor of Cognitive Science
Affiliate Assistant Professor of Computer Science
UC San Diego
9500 Gilman Drive, La Jolla, CA 92093-0515

A. PROFESSIONAL PREPARATION

University	Location	Major	Degree	& Year
MIT	Cambridge, MA	Electrical Eng. & Computer Science	S.B.	2005
MIT	Cambridge, MA	Electrical Eng. & Computer Science	MEng.	2006
Stanford	Stanford, CA	Computer Science	Ph.D.	2012
MIT	Cambridge, MA	CSAIL (postdoc institution)	Oct 2013 – Jun 2014	

B. APPOINTMENTS

July 2016 – present	Assistant Professor of Cognitive Science	UC San Diego
July 2016 – present	Affiliate Assistant Professor of Computer Science	UC San Diego
July 2014 – June 2016	Assistant Professor of Computer Science	University of Rochester

C. PRODUCTS

(i.) Products most closely related to the proposed project

Philip J. Guo and Margo Seltzer. Burrito: Wrapping Your Lab Notebook in Computational Infrastructure. USENIX Workshop on the Theory and Practice of Provenance (TaPP), 2012.

Alok Mysore and Philip J. Guo. Torta: Generating Mixed-Media GUI and Command-Line App Tutorials Using Operating-System-Wide Activity Tracing. *ACM Symposium on User Interface Software and Technology* (UIST), Oct 2017.

Alok Mysore and Philip J. Guo. Porta: Profiling Software Tutorials Using Operating-System-Wide Activity Tracing. *ACM Symposium on User Interface Software and Technology* (UIST), Oct 2018.

Xiong Zhang and Philip J. Guo. DS.js: Turn Any Webpage into an Example-Centric Live Programming Environment for Learning Data Science. ACM Symposium on User Interface Software and Technology (UIST), Oct 2017. ***(Honorable Mention Paper Award)***

Philip J. Guo. Older Adults Learning Computer Programming: Motivations, Frustrations, and Design Opportunities. *ACM Conference on Human Factors in Computing Systems (CHI)*, May 2017. ***(Honorable Mention Paper Award)***

(ii.) Other significant products

April Y. Wang, Ryan Mitts, Philip J. Guo, Parmit K. Chilana. Mismatch of Expectations: How Modern Learning Resources Fail Conversational Programmers. *ACM Conference on Human*

Factors in Computing Systems (CHI), Apr 2018. ***(Honorable Mention Paper Award)***
Philip J. Guo. Non-Native English Speakers Learning Computer Programming: Barriers, Desires, and Design Opportunities. *ACM Conference on Human Factors in Computing Systems* (CHI), Apr 2018.

Jeremy Warner and Philip J. Guo. CodePilot: Scaffolding End-to-End Collaborative Software Development for Novice Programmers. ACM Conference on Human Factors in Computing Systems (CHI), May 2017.

Philip J. Guo. Codeopticon: Real-Time, One-To-Many Human Tutoring for Computer Programming. *ACM Symposium on User Interface Software and Technology (UIST)*, Nov 2015.

Hyeonsu Kang and Philip J. Guo. Omnicode: A Novice-Oriented Live Programming Environment with Always-On Run-Time Value Visualizations. *ACM Symposium on User Interface Software and Technology* (UIST), Oct 2017.

D. SYNERGISTIC ACTIVITIES

1. *Software*: I created Python Tutor, which is now the most widely-used program visualization system, with over 3.5 million total users in over 180 countries. This open-source system has been used worldwide in dozens of university courses and MOOCs (Massive Open Online Courses).

2. *Education/Outreach*: My personal website contains over 300 articles, videos, and podcast episodes on topics ranging from research to education, and receives over 750,000 page views per year. My most popular articles include advice for Ph.D. program applicants (~150,000 total views) and graduate research fellowship applicants (~180,000 views).

3. *Education/Outreach*: I have created over 500 research, education, and scientific outreach videos on my personal YouTube channel, which now has 4,000+ subscribers and 500,000+ total video views. My Twitter social media account has 9,000+ followers and has been officially verified by the company as a notable account of public interest.

4. *Outreach*: In July 2016, I started a podcast interview series on YouTube and iTunes where I talk to scientists about their research, work habits, and advice for students. I have recorded 38 podcast interview episodes so far, 20 of which are with female scientists and engineers.

5. *Diversity*: In Jan 2014, I wrote an article for *Slate* called "Silent Technical Privilege" about implicit bias and discrimination in STEM fields, which has over 370,000 page views so far. This article led to interviews on NPR and the BBC, and a workshop on STEM diversity at MIT.

Haijun Xia
Assistant Professor
Department of Cognitive Science, University of California, San Diego
haijunxia@ucsd.edu

(a) Professional Preparation

Tsinghua University	Beijing, China	Computer Science	Bachelor of Engineering, 2013
University of Toronto	Toronto, Canada	Computer Science	Master of Science, 2015
University of Toronto	Toronto, Canada	Computer Science	Doctor of Philosophy, 2019

(b) Appointments

University of Toronto, Teaching Assistant 2018
Microsoft Research, Research Intern, 2017
Microsoft Research, Research Intern, 2016
Autodesk Research, Research Intern, 2015

(c) *Publications*

(i) Closely related to the proposed project

Xia, H., Herscher, S., Perlin, K., and Wigdor, D. 2018 Spacetime: Enabling Fluid Individual and Collaborative Editing in Virtual Reality. *In Proceedings of the ACM Symposium on User Interface Software and Technology.* UIST 2018. ACM, New York, NY, 853-866.

Xia, H., Riche, N., Chevalier, F. Araujo, B., and Wigdor, D. 2018 DataInk: Enabling Direct and Creative Data-Oriented Drawing. *In Proceedings of the ACM annual conference on Human Factors in Computing Systems.* CHI 2018. ACM, New York, NY. 223-236. **Best Paper Honorable Mention**

Xia, H., Hinckley, K, Pahud, M., Tu, X., and Buxton, B. 2017 WritLarge: Ink Unleashed by Unified Scope, Action, & Zoom. *In Proceedings of the ACM annual conference on Human Factors in Computing Systems.* CHI 2017. ACM, New York, NY. 3227-3240. **Best Paper Honorable Mention**

Xia, H., Araujo, B., and Wigdor, D. 2017. Collection Objects: Enabling Fluid Formation and Manipulation of Aggregate Selections. *In Proceedings of the ACM annual conference on Human Factors in Computing Systems.* CHI 2017. ACM, New York, NY. 5592-5604. **Best Paper Honorable Mention**

Xia, H., Araujo, B., Grossman, T., and Wigdor, D. 2016. Object-Oriented Drawing. *In Proceedings of the ACM annual conference on Human Factors in Computing Systems*. CHI 2016. 4610-4621. **Best Paper Award**

(ii) Other significant publications/products

Hayatpur, D., Heo, S., **Xia, H.,** Stuerzlinger, W. and Wigdor, D. Plane, Ray, and Point: Interaction Methods for ad hoc Creation of Alignment and Manipulation Constraints in VR. 2019. ACM Symposium on User Interface Software and Technology. UIST 2019. ACM, New York, NY. Accepted. To appear.

Zhang, Y., Pahud, M., Holz, C., **Xia, H.**, Laput, G., McGuffin, M., Tu, X., Mittereder, A., Su, F., Buxton, W., Hinckley, K. 2019. Sensing Posture-Aware Pen+Touch Interaction on Tablets. In *Proceedings of the ACM annual conference on Human Factors in Computing Systems*. CHI 2019. ACM, New York, NY, USA, Paper 55, 14 pages. **Best Paper Honorable Mention**

Xia, H., Herscher, S., Perlin, K., and Wigdor, D. 2018 Spacetime: Enabling Fluid Individual and Collaborative Editing in Virtual Reality. *In Proceedings of the ACM Symposium on User Interface Software and Technology.* UIST 2018. ACM, New York, NY, 853-866.

Xia, H., Grossman, T., and Fitzmaurice. G. 2015. NanoStylus: Enhancing Input on Ultra-Small Displays with a Finger-Mounted Stylus. *In Proceedings of the ACM symposium on user interface software and technology.* UIST 2015. ACM, New York, NY, 447-456.

Xia, H., Jota, R., McCanny, B., Yu, Z., Forlines, C., Singh, K., and Wigdor, D. 2014. Zero-Latency Tapping: Using Hover Information to Predict Touch Locations and Eliminate Touchdown Latency. *In Proceedings of the ACM symposium on user interface software and technology*. UIST 2014. ACM, New York, NY, 205-214.

(d) Synergistic Activities

- Program Committee, Associate Chair, ACM SIGCHI 2019
- Served as ACM SIGCHI Peer Reviewer (2015-2018)
- Served as ACM UIST Peer Reviewer (2016-2019)
- Program Committee, Associate Chair, Late Breaking Work, ACM SIGCHI, 2018
- Program Committee, Associate Chair, Interactivity, ACM SIGCHI, 2017

Biographical Sketch: Arvind Satyanarayan

Assistant Professor of Electrical Engineering and Computer Science (EECS)
MIT CSAIL
32 Vassar St., 32-G706
Cambridge, MA 02139
arvindsatya@mit.edu | http://arvindsatya.com

Professional Preparation

BS 2011	Computer Science	**University of California, San Diego**, *San Diego, CA*
MS 2014	Computer Science	**Stanford University**, *Stanford, CA*
PhD 2017	Computer Science	**Stanford University**, *Stanford, CA*
Postdoc 2017–2018	Google Brain	**Google, Inc.**, *San Francisco, CA*

Appointments

2018–	**Massachusetts Institute of Technology**, *Cambridge, MA*
	Assistant Professor, Electrical Engineering & Computer Science (EECS)
2013–2016	**Apropose, Inc.**, *Mountain View, CA*
	Co-Founder, Chief Architect (2013–2014), Advisor (2014–2016)

Products

(i) Related Products

1. Vega-Lite: A Grammar of Interactive Graphics. Arvind Satyanarayan, Dominik Moritz, Kanit Wongsuphasawat, Jeffrey Heer. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis'16)*, 2016. **Best Paper.**
2. VizNet: Towards a Large-Scale Visualization Learning and Benchmarking Repository. Kevin Hu, Neil Gaikwad, Madelon Hulsebos, Michiel Bakker, Emanuel Zgraggen, César Hidalgo, Tim Kraska, Guoliang Li, Arvind Satyanarayan, Çağatay Demiralp. *Proc. ACM Human Factors in Computing Systems (CHI)*, 2018.
3. Lyra: A Visualization Design Environment. Arvind Satyanarayan, Jeffrey Heer. *Computer Graphics Forum (Proc. Eurovis'14)*, 2014.
4. Reactive Vega: A Streaming Dataflow Architecture for Declarative Interactive Visualization. Arvind Satyanarayan, Ryan Russell, Jane Hoffswell, Jeffrey Heer. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis'15)*, 2015.
5. Visual Debugging Techniques for Reactive Data Visualization. Jane Hoffswell, Arvind Satyanarayan, Jeffrey Heer. *Computer Graphics Forum (Proc. EuroVis)*, 2016.

(ii) Other Significant Products

1. Webzeitgeist: Design Mining the Web. Ranjitha Kumar, Arvind Satyanarayan, Cesar Torres, Maxine Lim, Salman Ahmad, Scott R. Klemmer, Jerry O. Talton. *Proc. ACM Human Factors in Computing Systems (CHI)*, 2013. **Best Paper**.
2. Authoring Narrative Visualizations with Ellipsis. Arvind Satyanarayan, Jeffrey Heer. *Computer Graphics Forum (Proc. EuroVis)*, 2014.
3. The Building Blocks of Interpretability. Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine Ye, Alexander Mordvintsev. *Distill* 3(3), 2018.

Synergistic Activities

1. Developed *Vega*, *Vega-Lite*, and *Lyra* visualization systems. *Vega* has been deployed on Wikipedia, *Vega-Lite* is widely used in the Python data science community (Jupyter, Observable), and both are in use in industry (Apple, Google, Netflix, Fitbit, etc.).
2. Program Co-Chair for OpenVis Conf 2018.
3. Co-Editor of Distill.pub, an interactive academic journal for machine learning research.
4. Advisory Board member for the Harvard Data Science Review (HDSR).
5. Invited speaker at NICAR 2013, OpenVis Conf 2013, 2015, 2016, and Keynote speaker at the U.N. DataViz Camp 2016 and IEEE Visualization In Practice (VIP) workshop 2016.

CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information

Amy Rae Fox[1], Arvind Satyanarayan[2], Philip Guo[3], Haijun Xia[4], and James D. Hollan[5]

[1,3,4,5]Design Lab and Department of Cognitive Science, UC San Diego
[2]Visualization Group, Massachusetts Institute of Technology

Project Personnel and Institutions

1. James Hollan; UC San Diego; PI
2. Philip Guo; UC San Diego ;co-PI
3. Haijun Xia; UC San Diego; co-PI
4. Arvid Satyanarayan; Lead MIT Subaward
5. Barbara Tversky; Stanford University and Columbia University; Consultant
6. Jeff Zacks; Washington University; Consultant
7. Alan Kay; Viewpoints Research Institute; Advisory Panel
8. Fernando Pérez; UC Berkeley; Advisory Panel
9. Gloria Mark; UC Irvine; Advisory Panel
10. Jaime Teevan; Microsoft; Advisory Panel
11. Sep Kamvar; Mosaic; Advisory Panel
12. Vicki Hanson; Association for Computing Machinery; Advisory Panel
13. Clemens Nylandsted Klokmose; Aarhus University; Unpaid Collaborator
14. Michel Beaudouin-Lafon; Université Paris-Sud; Unpaid Collaborator
15. Wendy Mackay; INRIA; Unpaid Collaborator
16. Aurélien Tabard; Université Lyon; Unpaid Collaborator

Viewpoints Research Institute
1880 Century Park East, Los Angeles, 90067 (310) 471-4241

September 19, 2019

Professor James D. Hollan
Department of Cognitive Science
University of California, San Diego
La Jolla, CA 92093

If the proposal submitted by Dr. James D. Hollan, entitled CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information, is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Sincerely,

Dr. Alan Kay
President
Viewpoints Research Institute

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

Tel 425 882 8080
Fax 425 706 7329
www.microsoft.com

Microsoft

September 16, 2019

Professor James D. Hollan
Department of Cognitive Science
University of California, San Diego
La Jolla, CA 92093

If the proposal submitted by Dr. James D. Hollan, entitled CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Sincerely,

Jaime Teevan
Chief Scientist
Microsoft
Corporation

Jaime Teevan is Chief Scientist for Microsoft's Experiences and Devices, where she is helping Microsoft create the future of productivity. Previously she was the Technical Advisor to Microsoft's CEO, Satya Nadella, and a Principal Researcher at Microsoft Research AI, where she led the Productivity team. Dr. Teevan has published hundreds of award-winning research papers, technical articles, books, and patents, and given keynotes around the world. Her groundbreaking research earned her the Technology Review TR35 Young Innovator, Borg Early Career, Karen Spärck Jones, and SIGIR Test of Time awards. She holds a Ph.D. from MIT and a B.S. from Yale, and is an affiliate professor at the University of Washington.

TEACHERS COLLEGE

COLUMBIA UNIVERSITY

DEPARTMENT OF HUMAN DEVELOPMENT

September 19, 2019

To: NSF:

If the proposal submitted by Dr. James D. Hollan, entitled CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information, is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Kind regards,

Barbara Tversky

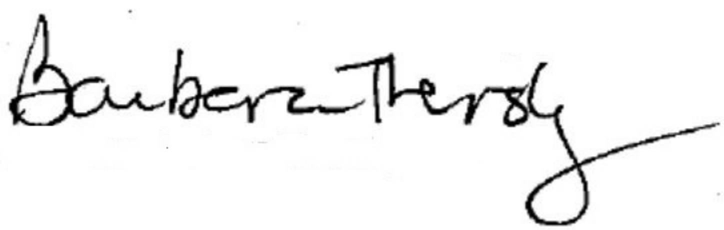

Professor of Psychology and Education

Columbia Teachers College

525 WEST 120th STREET, NEW YORK, NY 10027-6696

AARHUS UNIVERSITY
Aarhus University

To whom it may concern,

If the proposal submitted by Dr. James D. Hollan entitled "CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information" is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Clemens N. Klokmose

Department of Digital Design and Information Studies

Clemens N. Klokmose

Associate Professor

Date: 17 September 2019

Mobile Tel.: +45 28556442
Email: clemens@cavi.au.dk

Web: klokmose.net/clemens

Sender's CVR no.: 31119103

Aarhus University
Nordre Ringgade 1
DK-8000 Aarhus C
Denmark

Tel.: +45 89421111
Fax: +45 89421109
Email: au@au.dk
www.au.dk/en

UNIVERSITY OF CALIFORNIA, BERKELEY

BERKELEY • DAVIS • IRVINE • LOS ANGELES • MERCED • RIVERSIDE • SAN DIEGO • SAN FRANCISCO SANTA BARBARA • SANTA CRUZ

Fernando Perez, Ph.D.
Assistant Professor
Department of Statistics
UC Berkeley
fernando.perez@berkeley.edu

September 19, 2019

NSF Program Officers

Dear érogram éfficers,

If the proposal submitted by Dr. James D. Hollan, entitled *CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information*, is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Yours sincerely,

Fernando Pérez

UNIVERSITY OF CALIFORNIA, IRVINE UCI

Berkeley - Davis - Irvine - Los Angeles - Riverside - San Diego - San Francisco Santa Barbara - Santa Cruz

INFORMATION AND COMPUTER SCIENCE

IRVINE, CALIFORNIA 92697-3425
GLORIA MARK
TEL: (949) 824-5955
FAX: (949) 824-1715
EMAIL: gmark@uci.edu

Sept. 19, 2019

Professor James D. Hollan
Department of Cognitive Science
University of California, San Diego
La Jolla, CA 92093

Dear Jim:

Here is my statement to include in your NSF proposal submission:

If the proposal submitted by Dr. James D. Hollan, entitled *CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information*, is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Sincerely yours,

Gloria Mark
Professor
Department of Informatics

Laboratoire de Recherche en Informatique
Université Paris-Sud, Centre d'Orsay
U.M.R. 8623 du C.N.R.S.
Bâtiment 650
F - 91 405 Orsay Cedex
fax : +33 (0)1 69 15 65 86
Michel Beaudouin-Lafon
Professeur, membre de l'Institut Universitaire de France
tél.: +33 (0)1 69 15 69 10
e-mail: mbl@lri.fr

Orsay, 12 September 2019

Professor James D. Hollan
Department of Cognitive Science
University of California, San Diego
La Jolla, CA 92093

If the proposal submitted by Dr. James D. Hollan, entitled *"CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information"* is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Prof. Michel Beaudouin-Lafon

Association for Computing Machinery

Advancing Computing as a Science & Profession

19 September 2019

Re: *CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information*

NSF Awards Panel:

If the proposal submitted by Dr. James D. Hollan, entitled *CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information*, is selected for funding by NSF, it is my intent to collaborate as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

This proposal has my full support.

Sincerely,

Vicki L. Hanson

Vicki L. Hanson, PhD, FACM, FRSE, FBCS CITP
ACM Chief Executive Officer

2 Penn Plaza, Suite 701
New York, NY 10121-0701, USA

Tel: +1-212-626-0550
Fax: +1-212-944-1318

Vicki.hanson@hq.acm.org
https://vickihanson.org/
https://acm.org

Monday, 16 September, 2019

To: Professor James D. Hollan
Department of Cognitive Science
University of California, San Diego
La Jolla, CA 92093:

If the proposal submitted by Dr. James D. Hollan, entitled *CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information*, is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Wendy E. Mackay
Research Director (DR0), Inria
Responsable for the ex)**situ** joint lab
with Inria and Université Paris-Saclay

Phone : *+33 (1) 69 15 69 08*
Fax : *+33 (1) 69 15 42 16*
Email : *wendy.mackay@inria.fr*

LIRIS

CNRS
UMR 5205

Aurélien Tabard
Bâtiment Blaise Pascal
7, avenue Jean Capelle
69622 Villeurbanne cedex - France
Tél : 06 98 54 35 33
Mél : aurelien.tabard@liris.cnrs.fr

September 18th, 2019

If the proposal submitted by Dr. James D. Hollan, entitled CHS: Medium: A Human-Centered Information Space: CHS: Medium: A Human-Centered Information Space:Designing Dynamic Personalized Visual Information, is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Aurélien Tabard,
Associate Professor of Computer Science
Université de Lyon & CNRS

Washington University in St. Louis

Department of Psychological and Brain Sciences

September 26, 2019

Dear Colleagues,

If the proposal submitted by Dr. James D. Hollan, entitled *CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information*, is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Yours Sincerely,

Jeffrey M. Zacks, PhD
Professor of Psychological and Brain Sciences
Professor of Radiology
Associate Chair of Psychological and Brain Sciences
Director, Dynamic Cognition Laboratory

Washington University
Campus Box 1125
One Brookings Drive
Saint Louis, Missouri 63130-4899
(314) 935-8454

Mosaic

December 5, 2018

To whom it may concern,

If the proposal submitted by Dr. James D. Hollan, entitled *CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information*, is selected for funding by NSF, it is my intent to collaborate and/or commit resources as detailed in the Project Description or the Facilities, Equipment or Other Resources section of the proposal.

Thank you,

Sep Kamvar

CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information

Amy Rae Fox[1], Arvind Satyanarayan[2], Philip Guo[3], Haijun Xia[4], and James D. Hollan[5]

[1,3,4,5]Design Lab and Department of Cognitive Science, UC San Diego
[2]Visualization Group, Massachusetts Institute of Technology

Facilities, Equipment, and Other Resources

The Design Lab at UC San Diego and PI's Satyanarayan's lab at MIT EECS/CSAIL will be the primary locations for the proposed work. CSAIL is committed to pioneering new approaches to computation that will bring about positive changes in the way people around the globe live, play, and work. It involves more than 60 research groups. The Design Lab is located in the Qualcomm Institute which is the UCSD campus of the California Institute of Telecommunications and Information Technology (Calit2). Calit2 is a major state/industry-funded research institute at UCSD and UCI, whose mission is to facilitate large-scale interdisciplinary research. Calit2 is dedicated to information technology research and its primary mission is to bring researchers, industry, and the public together for the benefit of society.

Calit2 Facilities

Calit2 researchers have access to a 127,000 ASF facility consisting of research neighborhoods, specialized laboratories, general and public spaces, and administrative spaces. This provides numerous adjustable spaces for the manipulation of various multimedia technologies. The research neighborhoods consist of office spaces, related support space and open project areas. Additionally, Calit2 researchers have access to campus (OptIPuter, ABC), regional (CENIC), national (NLR) and international (STAR TAP) wireless and optical networking test beds and visualization centers. Calit2 also offers collaborators access to a variety of high definition video teleconferencing systems ranging from room to room collaborations using the LifeSize HD system to more one-on-one communication via a variety of telepresence facilities. These systems allow research teams, like ours, to meet on a regular basis even though they involve a group at UC San Diego, MIT CSAIL, and two in Europe (Aarhus and Paris). High-end visualization facilities are also available for viewing data and for facilitating policy and decision support. These specialized facilities include the Calit2 StarCAVE and one of the world's largest scaled tiled display walls known as the Highly Interactive Parallelized Display Wall project (HIPerWall) providing high-capacity visualization capabilities to the project. The StarCAVE is a 5-sided immersive visualization environment that projects stereo imagery on the 5 walls surrounding the user as well as to the floor using a total of 34 rear projectors. The HIPerWall is composed of a 70 LCD tile display, which offers over 491 million pixels in resolution in its current configuration.

Collaborating Institutions and Key Researchers

We have assembled an exciting group of collaborators from active internationally respected laboratories with strong research and educational track records. The proposed research will benefit from collaboration with labs in France and Denmark. They also have excellent access to modern visualization and research facilities. Professor Beaudouin-Lafon, for example, leads the large (22M €) French Digiscope project. Although our European collaborators and students from their labs will receive no funding from this project, these labs are currently funded for work on Webstrates and related projects. The collaboration will have immense benefits for the proposed effort.

- *Michel Beaudouin-Lafon, Professor Computer Science, Université Paris-Sud, France.* He was director of LRI, the Laboratory for Computer Science, joint between Université Paris-Sud and CNRS, from 2002 to 2009, elected in 2006 to the ACM SIGCHI Academy, received the ACM SIGCHI Lifetime Service Award in 2015, is a senior member of the Institut Universitaire de France (Only 2% of French university professors have been currently distinguished by the Institut Universitaire de France), and recently was awarded a prestigious European Research Council Advanced grant.

- *Clemens Nylandsted Klokmose, Associate Professor, Department of Digital Design and Information Studies, School of Communication and Culture, Aarhus University, Denmark.* He is also a member of Centre for Advanced Visualization and Interaction (CAVI), affiliated with the Centre for Participatory Information Technology, a collaborator with INRIA Saclay, and the primary designer of Webstrates.

- *Wendy Mackay, Research Director (Tenured Professor), INRIA Saclay, France.* She heads the ExSitu Human-Computer Interaction Research Group, was elected to the ACM SIGCHI Academy (an honorary group of individuals who have made extensive contributions to the study of human-computer interaction and who have led the shaping of the field) in 2009, received the ACM SIGCHI Lifetime Service Award in 2014, and is a recipient of a prestigious European Research Council Advanced grant. She recently served as Vice President of Research for the Computer Science Department at Université Paris-Sud.

- *Aurélien Tabard, Associate Professor, Université Lyon, France.* His research focuses on consequences of the digital revolution: the pervasiveness of digital traces we leave behind us, how we can control them, and how they can enrich our lives.

The team we have assembled and the laboratories they represent are uniquely positioned to catalyze international research and education collaborations. In addition, connections to Project Jupyter will expand the impact of the proposed research and the development of next generation web-based information design environments.

Each lab has additional faculty who will also assist with training. The UC San Diego Design Lab that PI Hollan co-directs along with Professors Scott Klemmer and Don Norman focuses on designing complex sociotechnical systems. PIs Guo and Xia are also members of the Design Lab. PI Satyanarayan just joined MIT CSAIL as an Assistant Professor. Our European collaborators are Professor Beaudouin-Lafon, Klokmose, Mackay, and Tabard. The ExSitu research group is joint between Université Paris Sud and INRIA and focuses on interaction design and human-computer interaction. Professor Klokmose works on design and visualization at Aarhus University. He collaborates with the ExSitu research group on Webstrates. For the last several year Professor Tabard of Université Lyon has collaborated with PI Hollan. The Digiscope project that Professor Beaudouin-Lafon leads involves a unique network of nine sites in France that have large wall displays and are interconnected with high-end video-conferencing facilities. Similarly, the Design Lab is housed in the Qualcomm Institute (QI) of the California Institute of Telecommunications and Information Technology (Calit2) with access to advanced visualization and videoconferencing facilities to support the proposed effort. Access to these advanced networking and visualization facilities is of particular significance not only for the proposed research but also for educational activities, weekly research meetings, and informal interactions as needed.

Advisory Panel and Consultants

The proposed research will benefit from access to a stellar advisory panel and consultants. We are extremely fortunate to have assembled a set of exceptionally distinguished researchers.

Alan Kay

Viewpoints Research Institute. Dr Kay's experience spans across Xerox PARC, Stanford University, Atari, Apple, Walt Disney Imagineering, and Hewlett-Packard. Among his many awards are the ACM Turing Award and Kyoto Prize. The Turing Award is for pioneering many of the ideas at the root of contemporary object-oriented programming languages, leading the team that developed Smalltalk, and for fundamental contributions to personal computing. The Kyoto prize recognizes him for creation of the concept of modern personal computing and outstanding contribution to its realization. Dr. Kay is a fellow of the American Academy of Arts and Science, National Academy of Engineering, Royal Society of Arts, Computer History Museum, and the Association for Computing Machinery. He has taught at New York University, UCLA, Kyoto University and MIT.

Barbara Tversky

Professor of Psychology and Education and Emerita Professor of Psychology at Stanford. Professor Tversky specializes in cognitive psychology. She is a leading authority in the areas of visual-spatial reasoning, language and communication, comprehension of events and narratives, and the mapping and modeling of cognitive processes. She was elected to the American Academy of Arts and Sciences, is a Fellow of American Psychological Society, Cognitive Science Society, Russell-Sage Foundation, and Eastern Psychological

Association. She has received a Phi Beta Kappa Excellence in Teaching Award and Distinguished Software Award from EDUCOM/NCRIPTAL.

Jeff Zacks
Associate Chair and Professor of Psychological & Brain Sciences and Professor of Radiology at Washington University Professor Zacks studies perception and cognition using behavioral experiments, functional MRI, computational modeling, and testing of neurological patients. One line of Zacks' research examines how people parse the continuous stream of behavior into meaningful events, and how this affects memory and cognition. Another line examines how mental imagery contributes to reasoning about spatial relations, especially how mental representations of one's body are updated during imagery and reasoning.

Fernando Pérez
Assistant Professor, Statistics Department, UC Berkeley, Faculty Scientist at the Data Science and Technology Division of Lawrence Berkeley National Laboratory, and founding Co-Investigator of the Berkeley Institute for Data Science (BIDS). Professor Pérez created the iPython project (precursor of Jupyter) and is the founder and leader of Project Jupyter. The Jupyter Notebook is the flagship open-source software product of Project Jupyter. For his work on Jupyter is received the Award for the Advancement of Free Software.

Gloria Mark
Professor, Department of Informatics, University of California, Irvine. Visiting Senior Researcher, Microsoft and; Visiting Professor at MIT Media Lab, National University of Singapore, University of Haifa; Groupware Laboratory Director, Electronic Data Systems Center for Advanced Research. Professor Mark was elected to the ACM CHI Academy, and has received Donald Bren School Outstanding Graduate Student Mentor Award, an NSF Career Grant, and multiple Google Research Awards.

Jaime Teevan
Chief Scientist for Microsoft's Experiences and Devices. Previously she was the Technical Advisor to Microsoft's CEO, Satya Nadella, and a Principal Researcher at Microsoft Research AI, where she led the Productivity team. Dr. Teevan has published hundreds of award-winning research papers, technical articles, books, and patents, and given keynotes around the world. Her research earned her the Technology Review TR35 Young Innovator, Borg Early Career, Karen Spärck Jones, and SIGIR Test of Time awards. She holds a Ph.D. from MIT and a B.S. from Yale, and is an affiliate professor at the University of Washington.

Sep Kamvar
Founder of Wildflower Schools and multiple companies, including Mosaic, Celo, and Kaltix. Until recently the LG Career Development Professor of Media Arts and Sciences and Director of the Social Computing Group at the MIT Media Lab. He was head of personalization at Google. He has been a consulting professor at the Institute for Computational and Mathematical Engineering at Stanford University. Dr. Kamvar is a computer scientist, artist, and entrepreneur. He is the author of three books and his artwork has been exhibited in museums worldwide, including The Museum of Modern Art in New York, the Victoria and Albert Museum in London, and the National Museum of Contemporary Art in Athens.

Viki Hanson
Executive Director and CEO of the Association for Computing Machinery (ACM). Previously Distinguished Professor at Rochester Institute of Technology, Professor and Chair of Inclusive Technologies at the University of Dundee, and Past President of ACM. Among her many awards, Professor Hanson has received the ACM SIGCHI Social Impact Award, ACM SIGACCESS Award for Outstanding Contributions to Computing and Accessibility, Anita Borg Institute Woman of Vision Award for Social Impact, and is an ACM Fellow.

CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information

Amy Rae Fox[1], Arvind Satyanarayan[2], Philip Guo[3], Haijun Xia[4], and James D. Hollan[5]

[1,3,4,5]Design Lab and Department of Cognitive Science, UC San Diego
[2]Visualization Group, Massachusetts Institute of Technology

Collaboration Plan

As a group of cognitive and computer scientists, we have come together jointly convinced of the importance of questioning the presupposition that information is fundamentally passive data disconnected from processes, tasks, context, and personal histories. We are excited by the opportunity to rethink the nature of computationally-based information. The current effort builds on a long history of collaboration on human-computer interaction research among the participants, recent work on capturing activity history, the maturing of the Vega visualization framework, and a new collaboration on Webstrates, a novel approach for creating sharable dynamic media.

The proposed effort will benefit immensely from collaboration with ongoing European-funded efforts[1] at Aarhus University and the University of Paris and INRIA. The leaders of this effort are: Clemens Nylandsted Klokmose (Aarhus University, Denmark), Michel Beaudouin-Lafon (Université Paris-Sud, France), and Wendy Mackay (INRIA Saclay, France). Our European collaborators and students from their labs will receive no funding from this project. Existing EU funding will support their collaboration with our effort.

As part of the collaboration, we will also benefit from support from a stellar Advisory Panel: Alan Kay, an ACM Turing Award and Kyoto Prize winner who pioneered many of the ideas that led to this proposal; Gloria Mark, who's work has been central to understanding information work, interruptions, and the fragmentation of activity; Jaime Teevan, Chief Scientist at Microsoft who focuses on people's information seeking activities; Sep Kamvar, who initiated personalized search at Google; Viki Hanson, CEO of ACM, who is known for her work on accessibility and social impact and will advise on our outreach efforts; and Fernando Pérez, who will act as a liaison to the Jupyter community. In addition, we will have consulting support from international-known psychologists Barabara Tversky and Jeff Zacks. They will advise on behavioral experiments.

It is extremely valuable that the faculty involved have a shared vision and a history of working together. PI Hollan spent a sabbatical recently at INRIA and collaborated on visualization and Webstrates with Klokmose, Beaudouin-Lafon, and Mackay. PI Satyanarayan visited during this period. Various subsets of us and our students currently meet weekly and we have already seen the benefits of this not only for research but also for the education of our students. This interdisciplinary interaction encourages students to appreciate cultures and disciplines outside their own. The type of research we propose necessitates interdisciplinary collaboration. In the design phase, cognitive scientists and computer scientist need to work closely to design innovative yet feasible solutions. In collecting data in the field, cross-disciplinary collaboration is just as useful, as the team seeks to gather, analyze, and interpret data to evolve the software systems being developed.

In the following we outline the responsibilities and specific roles that each team member will play in the proposed research activities. Then we discuss the collaboration mechanisms that will be employed to ensure that each role contributes to the overall effort.

Specific Roles

PI Jim Hollan will provide overall direction for the project. Working closely with Satyanarayan, Guo, Fan, and our European collaborators, he will carry the major responsibility for coordinating project activities. He will lead the effort on designing and evaluating the dynamic personal information space we propose with assistance from consultants Barbara Tversky and Jeff Zacks. In addition, he will have responsibility for

[1]See Synergistic Collaboration (Section 3, page 5).

assuring the quality of the scientific output of the project. He has over three decades of experience leading collaborative research projects. He is a member of the ACM SIGCHI Academy and in 2015 received the ACM SIGCHI Lifetime Research Award.

Arvind Satyanarayan will lead the MIT portion of the effort and focus on extension of Vega-Lite to support visualizations with the types of behavioral physics we propose. As a principal designer of the Vega stack he is perfectly positioned to fill this role.

co-PI Philip Guo will lead the activity capture effort. His expertise and extensive experience developing web-based software are especially valuable for the project.

co-PI Haijun Xia, who will be arriving as a new UCSD faculty member in November 2019, will lead development of a activity-centered representations and assist in designing a grammar for dynamic behavior and interaction.

Graduate Students

We propose two graduate students for the first two academic years and three graduate students working full-time each of the three summers. As detailed in the MIT subcontract, graduate students at MIT will also be supported. Students will be heavily involved in all aspects of the proposed research and will particularly benefit from the interdisciplinary focus and the involvement of faculty from multiple laboratories. We also expect additional graduate students and researchers from collaborating labs at UC San Diego, MIT, Aarhus, INRIA, and University of Paris to be actively involved given their interests in the new dynamic visual information substrate we propose.

Collaboration Mechanisms

Bi-Weekly Meetings: The collaboration began during PI Hollan's recent sabbatical and following his return to UC San Diego it has continued via Skype meetings. The current proposal has benefited from this. To establish and ensure a successful collaboration, we will continue to have weekly Skype project meetings to assess progress, remove roadblocks, and make plans.

Meetings at Conferences: We plan to schedule an added day at the major conferences we all attend (e.g., CHI, UIST, and CSCW) to get together for face-to-face meetings. This will provide added opportunities to address progress and to discuss project plans, collaboration, future publications, and research dissemination.

Slack Channels: To ensure that all are "on the same page" we will create Slack channels to support team conversations and interactions on specific topics, including general project information, software development, and evaluations. We find such online channels to be indispensable for supporting collaborative research.

CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information

Amy Rae Fox[1], Arvind Satyanarayan[2], Philip Guo[3], Haijun Xia[4], and James D. Hollan[5]

[1,3,4,5]Design Lab and Department of Cognitive Science, UC San Diego
[2]Visualization Group, Massachusetts Institute of Technology

Broadening Participation in Computing

We have assembled an exciting group of collaborators from active international laboratories with strong research and educational track records. The ability to interact with students from international labs will augment the educational experience for all students. The broader impacts of the proposed activity derive from the promise to radically improve the efficacy of computer-mediated information activities and the exciting potential of cognitively convivial information entities to ultimately reshape how we think with computers.

Target Populations and Plans

One direction for broadening participation is to exploit the personalized information environment we propose to better serve individuals with disabilities. Fundamental to the Human-Centered Information Space concept is the idea that information entities and methods of interaction should be tailored at an individual level. Although we've describe these entities as 'dynamic personalized *visual* information', the representations themselves can be multi-modal. Vicki Hanson, a member of our Advisory Panel, is a recognized leader in addressing issues of accessibility and diversity. She will advise our activities in this area. We will also capitalize on the connections and outreach advice of our colleague Professor Beth Simon. She has been involved for a decade in efforts to bring quality computing education into K-12 education to support the new AP CS Principles (CSP) course. She is a co-PI on another NSF project with PI Hollan and we see interesting opportunities for the prototypes we will build to be used in her efforts to spark the interest of K-12 students. As part of our efforts to broaden participation in computing, we also mention PI Guo's development of online learning platforms. Over the past 7 years his Python Tutor [16] has been used by at estimated 5 million users and has been integrated into MOOCs from edX, Coursera, and Udacity.

Another project, currently in an early stage of planning, is to catalyze broader participation in computing by a pilot program to invite high-school and undergraduate students to attend the annual ACM CHI conference. We have discussed this with Vicki Hanson, former President of ACM, current ACM Chief Executive Officer, and a member of our Advisory Panel, and she notes that "the idea of bringing students to CHI is really intriguing and I think something that could be a model for others in the future." [1] There are numerous issues that would need to be address, such as age of students and appropriate permissions. If those are resolved, we are confident funding from other sources could be secured to support a pilot effort and we plan to continue to pursue it.

Prior Experience

The lack of appropriate socioeconomic, race, and gender diversity in computing research is well documented. Only 13-16% of technical roles are filled by women; proportional representation by African Americans and other underrepresented groups is even worse. Our lab has long been committed to attracting a diverse[1] group of students and involving them in research. We have participated in the UCSD STARS[2] outreach program. As in the past, we plan to employ the NSF REU mechanism to further increase the diversity of student participation. We also actively participate in a number of outreach activities hosted by Calit2 involving the innovative Preuss School [3] on campus.

[1]50% of our postdocs and 38% of our graduate students are women. The PIs have appointments in the Cognitive Science Department where women comprise 40% of the faculty.

[2]The University of California, San Diego Summer Training Academy for Research Success (STARS) program is an eight week summer research academy for community college students, undergraduate students, recent college graduates, and masters students.

[3]Preuss has been recognized by Newsweek as the top transformative high school in the nation for three years in a row. The Preuss School is a unique charter middle and high school for low income students who strive to become the first in their families to graduate from college. Nearly 100 percent of graduates go on to higher education.

CHS: Medium: A Human-Centered Information Space: Designing Dynamic Personalized Visual Information

Amy Rae Fox[1], Arvind Satyanarayan[2], Philip Guo[3], Haijun Xia[4], and James D. Hollan[5]

[1,3,4,5]Design Lab and Department of Cognitive Science, UC San Diego
[2]Visualization Group, Massachusetts Institute of Technology

Data Management Plan

Data coupled with the proposed research will be systematically managed by the research team. As part of this project we expect to produce four types of data: (a) recording of screen-capture activity data, as well as associated metadata resulting from analysis, (b) software to enable the collection, analysis and visualization of interaction, (c) software prototypes developed to explore designing a human-centered information space with dynamic personalized visual information entities, and (d) interview and data from experiments to iteratively design the prototypes. We will enforce following policies for data produced during this project:

1. Access and Sharing of Research Resources

We will ensure that intellectual property and data generated under this project will be administered in accordance with both University and NSF policies, including the NSF Data Sharing Policy and AAG Chapter VI.D.4. Ownership of sole or joint inventions developed under the project will be owned by the institution(s) employing the inventor(s). Inventors shall be determined by U.S. Patent law, Title 35 USC. University and participating investigators will disclose any inventions developed under the project and such inventions will be reported and managed as provided by NSF policies. Sole inventions will be administered by the institution employing the inventor. Joint inventions shall be administered based on mutual consultation between the parties. Similar procedures will be followed for copyrights.

2. Re-use and Re-Distribution of Research Resources

Access to data and associated software tools generated under the project will be available for educational, research and non-profit purposes. Such access will be provided using web-based applications, as appropriate. Materials generated under the project will be disseminated in accordance with university and NSF policies. Depending on such policies, materials may be transferred to others under the terms of a material transfer agreement. Publication of data shall occur during the project, if appropriate, or at the end of the project, consistent with normal scientific practices. Research data that documents, supports and validates research findings will be made available after the main findings from the final research data set have been accepted for publication. Such research data will be redacted to prevent the disclosure of personal identifiers.

3. Source of Data, Standards, Privacy and Archiving

Given the heterogeneous and sometimes confidential nature of the data that will be produced during this project, we specifically address standards, archiving, privacy and additional sharing and distribution policies for each data separately.

(a) Recording of interaction of participants in experiments and in use of prototypes.

Recorded data will be stored using well-recognized standards such as JPEG (for video frames), MP3 (for audio recordings), and comma-separated-value (CSV) files.

Maintaining participant confidentiality and research integrity through the proposed project will require significant data management capabilities. For data that may require it we have access to data management state-of-the-art information technologies include an extensive set of secure functions for participant tracking, distributed data entry, and protocol management complimented by real-time dynamic reporting tools. This approach to research information security and confidentiality is specifically designed to comply with the requirements of relevant federal regulations and guidelines.

(b) Software to analyze and visualize human-computer interaction and communication

We anticipate the development of a range of software tools that will allow researchers to investigate the nature of the collected data. As we already do with our current software (e.g. ChronoViz is freely available at http://chronoviz.com), we will make the developed software available to other researchers on dedicated servers or on servers made available by the university. The source code of the developed software will be available on open source repositories such as GitHub (http://github.com). The distributed nature of this version control system will also support sharing, backup and archiving of the developed code. We believe that making analysis software available to the community is an important part of researchers responsibility.

(c) Software prototypes

One of the results of the proposed project will be a range of prototypes. We plan to make those prototypes available to the larger research community, and as live demonstration during conferences and research meetings. Software at the basis of the developed prototypes will be available online. Source code will also be made available through open source repositories such as GitHub (http://github.com), which automatically provides sharing facilities, backup and archiving possibilities