# Beyond Tasks: An Activity Typology for Visual Analytics

Darren Edge, Nathalie Henry Riche, Jonathan Larson, and Christopher White

**Abstract**—As Visual Analytics (VA) research grows and diversifies to encompass new systems, techniques, and use contexts, gaining a holistic view of analytic practices is becoming ever more challenging. However, such a view is essential for researchers and practitioners seeking to develop systems for broad audiences that span multiple domains. In this paper, we interpret VA research through the lens of Activity Theory (AT)—a framework for modelling human activities that has been influential in the field of Human-Computer Interaction. We first provide an overview of Activity Theory, showing its potential for thinking beyond tasks, representations, and interactions to the broader systems of activity in which interactive tools are embedded and used. Next, we describe how Activity Theory can be used as an organizing framework in the construction of activity typologies, building and expanding upon the tradition of abstract task taxonomies in the field of Information Visualization. We then apply the resulting process to create an activity typology for Visual Analytics, synthesizing a wide range of systems and activity concepts from the literature. Finally, we use this typology as the foundation of an activity-centered design process, highlighting both tensions and opportunities in the design space of VA systems.

**Index Terms**—Activity theory, visual analytics, activity-centered design, literature review, human-computer interaction

## Introduction

Since the establishment of Visual Analytics (VA) as "the science of analytical reasoning facilitated by interactive visual interfaces" [71], the field has grown and diversified to encompass a wide variety of systems, techniques, and use contexts. Each of these advances the state of the art in a new direction, often tightly bound to a specific problem, task, or domain. However, Visual Analytics is no longer the preserve of the research community, or indeed the security and intelligence services that were the original external audience for such work.

In the wild, business intelligence (BI) systems are expanding beyond the data types and visualizations of the dashboard paradigm. IBM are advocating the mining of "dark data" for insights—data collected by organizations but not analyzed because of its unstructured nature [23]. Microsoft is also extending its Power BI platform with new visualizations [52] to support analysis of tweets [54] and text documents [53] using graph, set, list, facet, timeline, and summary representations. Across these examples and the industry as a whole, the characteristics of both data and visualizations are now closely aligned with those traditionally associated with Visual Analytics.

However, despite the mythology of big data— "the widespread belief that large data sets offer a higher form of intelligence and knowledge that can generate insights that were previously impossible" [15]—such insights are unlikely to be achieved without a fundamental transformation in the general tools (like modern BI systems) used to analyze such data. The challenge is finding the appropriate conceptual frameworks for thinking about the design of such tools, when their use in practice spans such a diversity of users, use cases, and use contexts.

Theory has the potential to help with this challenge—indeed, the InfoVis community has been advocating the need for theory for several years [74]. Both InfoVis and Visual Analytics have a theoretical tradition oriented towards the development of taxonomies for tasks and interactions (e.g., [1][16][19][30]). However, such low-level analysis omits the details of user, use case, and use context that are essential for understanding the broader activity to be supported.

Since establishing a theory from scratch is inherently challenging, we explore the potential of analyzing Visual Analytics through the lens of Activity Theory (AT) [26][27][47][48][49]—a well-established theory in the field of Human-Computer Interaction (HCI) [4][6][7][14][21][28][33][34][43][44][55]. We reformulate the theory for the analysis of high-level, general activities, applying the resulting activity-based design process to the activity of Visual Analytics.

At the core of our design approach is the concept of an *activity typology*—a structured organization of six activity elements and the various types taken by these elements across the many forms of the activity. Compared with existing task taxonomies, activity typologies encourage an expansion of focus "beyond tasks" to encompass the broader contexts of activity in which tools are embedded and used.

We begin by presenting an overview of how Activity Theory has been applied within the field of HCI, before elaborating on the concepts and development of the theory. We then describe the process of using Activity Theory to create a typology of a target activity and then using that typology as a foundation for activity-centered design. Next, we illustrate this process in detail for the activity of Visual Analytics. The outcome is a set of *design tensions* characterizing key challenges in the design of multi-purpose VA tools, and corresponding *target qualities* whose embodiment in VA tools would resolve these tensions and advance the practice of Visual Analytics in general.

## 1 Background and Motivation

In this section, we chart the use of Activity Theory within the field of Human-Computer Interaction and describe the limitations of existing HCI-AT methods when designing for general, high-level activities like Visual Analytics. We then present the fundamental concepts of Activity Theory, their application to Visual Analytics, and their relationship to the task-level taxonomies of InfoVis and VA research. In preparation for the remainder of the paper, we establish the activity system model as a framework for unifying disparate concepts arising from a broad spectrum of taxonomies, theories, systems, and studies.

### 1.1 Impact of Activity Theory within HCI

Activity Theory is well established in the HCI community and the topic of multiple influential books, including *Context and Consciousness* [55], *Acting with Technology* [44], and *Activity-Centered Design* [28]. It was also a key contributor to the "second wave" of HCI [22], described by Bannon as the shift "from human factors to human actors" [3] in which the concerns of HCI research expanded from individual users to the contexts of interaction. The influence of the emerging field of Computer-Supported Cooperative Work (CSCW) also led to a broader focus on groups, group work, and the workplace. In this second wave, proactive research methods like Participatory Design [63] and Contextual Inquiry [9] brought both users and use contexts into the design process, while Activity Theory and theories of Distributed Cognition [39] and Situated Action [69] gave researchers new frameworks for understanding interaction across individuals, tools, and contexts. Whereas third wave HCI has

- *Darren Edge, Nathalie Henry Riche, Jonathan Larson, and Christopher White are with Microsoft Research.*
- *E-mail:{darren.edge|nath|jolarso|chwh}@microsoft.com.*

*Manuscript received 31 Mar. 2017; accepted 1 Aug. 2017.*
*Date of publication 28 Aug. 2017; date of current version 1 Oct. 2017.*
*For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.*
*Digital Object Identifier no. 10.1109/TVCG.2017.2745180*

generally been defined in contrast to the second wave—"non-work, non-purposeful, non-rational, etc." [14]—there are many domains, including Visual Analytics, where the support of purposeful work remains the primary concern of both users and researchers.

Of the second wave theories, Activity Theory has generated the widest range of derived HCI methods for the design and evaluation of interactive systems. The Activity Walkthrough [7] adapts the cognitive walkthrough method for UI prototypes, illustrated through the running example of purchasing tickets from a railway ticket vending machine. The Activity Checklist [43] lists AT-inspired questions for tool design and evaluation—questions that were used in the design of Apple's Data Detector agent framework for recognizing structured data in free text and creating actions accordingly (e.g., creating an "Open URL" action for embedded URLs). This checklist has also been specialized for Visual Analytics [32], proposing that the analysts, tasks, and environments in real-world analysis contexts be profiled and evaluated against a list of criteria for tool use and design.

Developed for research in CSCL (Computer Support for Collaborative Learning), the Activity Space [33] assigns empirical findings to distinct subject groups, visualizes tensions within the categories of subject, object, outcome, and mediation, and proposes design changes to reduce inter-group conflicts. This framework was used to organize findings of a field study into why groupware designed to support collaborative learning (Lotus Notes) failed to achieve this objective, identifying critical tensions between the needs of teachers versus the needs of learners. The use of "design tensions" as springboards for design has also been explored [70], with tensions between competing design options first articulated and then resolved through trade-offs, insights, or project reformulation. The resulting Design Tensions framework was used in the design of the NetCalc project, which integrated the desktop SimCalc system for mathematics education with handheld computing. The influence of AT on HCI in the 25 years since its introduction is also captured in a recent meta-review of 109 AT-related HCI articles, spanning the spheres of health, education, science, work, and domestic, family, and social life [21].

#### 1.1.1 Limitations in Previous Uses of Activity Theory

One limitation of these methods is that they focus on the people, places, and contexts engaged in a specific instance of a more general activity. While such proactive and reflective intervention in real-world contexts is incredibly valuable (and indeed characteristic of second wave HCI), it is not possible to design customized tools for each concrete instance of activity, nor may it be desirable. In his critique "Human-Centered Design Considered Harmful" [57], Norman argues that HCD's "focus upon individual people (or groups) might improve things for them at the cost of making it worse for others", and that the target of design should instead be the activities in which users participate. The same argument applies to a methodological focus on individual forms of activity—if the goal is to design tools that support general use across a broad user base (as is typically the case in industry), then we should design to support the general activity system that unites the specific activities of these users.

The challenge in designing for a general activity system like Visual Analytics is that empirical knowledge (e.g., from lab studies or field investigations) is invariably linked with specific forms of the activity, and existing theories (e.g., task and interaction taxonomies) only contribute concepts relevant to specific parts of the activity system. However, the generality of Activity Theory makes it amenable to organizing the literature of an activity domain into a coherent whole, synthesizing activity concepts from a wide range of sources and clarifying the conceptual structure of the general activity system.

### 1.2 Fundamentals of Activity Theory

The theorization of activity has its roots in the work of Vygotsky and the cultural-historical school of psychology that formed in Russia in the 1920s and 1930s. His student, Leontiev, later came to establish the core tenets of what we now call Activity Theory. This theory conceptualizes activity as the purposeful actions of a living *subject* motivated by and directed towards an idealized *object* [47][48][49]. It also characterizes the subject–object relationship as being:

1. *mediated*, in that a subject uses a variety of tools (both physical and psychological) to work towards their object;
2. *cultural-historical*, in that subjects, objects, and tools carry the culture that has shaped their historic evolution;
3. *developmental*, in that activity transforms subjects, objects, tools, and cultures over time;
4. *distributed*, in that activity is a continuous reconfiguration of internal mental processes and external physical behavior; and
5. *hierarchical*, in that motive-driven "activity" is fulfilled by conscious, goal-directed "actions", which are themselves fulfilled by unconscious, skill-based "operations".

#### 1.2.1 The Activity of Visual Analytics

Using this framework, we can loosely characterize Visual Analytics activity as the purposeful actions of a human analyst motivated by and directed towards an idealized analytic product (e.g., some combination of insights, artefacts, and their intended effects). We can further say that this activity is mediated by visual analytic representations and workspaces that have evolved historically (e.g., from physical to digital media), and whose use leads to the acquisition of tradecraft. As the analyst develops their tradecraft over time, and in response to the momentary conditions of their environment, they constantly shift their attention between internal and external representations and reformulate task goals based on action feedback. This connection to user tasks touches upon the fundamental building blocks of VA activity as experienced by an acting analyst. The study of user tasks, particularly through the derivation of task taxonomies that seek to identify common structure across a range of analytic contexts, has a long history in both the InfoVis and VA literature.

#### 1.2.2 Activity Theory and Task Taxonomies

An example of AT concepts revealing themselves in task-level theories is Brehmer and Munzner's multi-level typology of abstract visualization tasks [16], which synthesizes and builds upon a wide range of prior InfoVis task taxonomies. This work aims to address the weak distinction in the literature between the means and ends of tasks—while low-level taxonomies based on primitives like Select, Filter, and Navigate cover only the means of interaction, high-level taxonomies based on primitives like Confirm Hypotheses, Present, and Explore cover only the ends. The resulting typology distinguishes the "why" of the task (corresponding to the AT motive–goal structure) from the "how" (the AT action–operation structure). It also gives a hierarchical breakdown of "why" into the high level (to consume or produce—AT motives), the mid-level (to search in various ways based on whether the target and location are known or unknown—AT goals), and the low-level (to query in various ways—elements of AT goals). This typology thus embodies the *hierarchical* and *distributed* nature of analytic activity. In contrast, taxonomies of visual variables (e.g., position, size, color, shape) [8], representational structures (e.g., maps, trees, graphs, tables) [19], and action semantics (e.g., exploratory vs insight actions) [30] focus on the *mediated* nature of interaction.

#### 1.2.3 Expanding Analysis to the Activity System

The previous review of taxonomies demonstrates that at the level of tasks, there is already much theory in the InfoVis and VA literature that can be organized and interpreted using Activity Theory. However, the literature is also replete with user profiles, system features, empirical findings, and design concepts that have not yet been assimilated into a coherent theory of VA activity.

One limitation of Leontiev's Activity Theory as presented so far is that it focuses on the activities of individual subjects rather than the socially-constructed systems of activity within which many subjects participate. This limitation has been addressed by the work of Engeström, who expanded the scope of Activity Theory to collective subjects acting in social contexts [26]. In Engeström's activity system model (Figure 1), *community* provides the primary mediation between the collective subject and their shared object. All pairwise relationships in the subject–object–community triangle are themselves mediated by further elements: the *instruments* through which the subject works towards the object; the *division of labor* through which the community contributes towards the object; and the *rules* by which the subject interacts with the community. The idealized object of activity is also extended to include the actual *outcome* of activity. The resulting triangles of activity provide a schematic framework with which to analyze the structure of activity systems.

Engeström's Activity Theory also contains a complementary framework for modelling the dynamics of activity systems, based on the concept of *contradictions*—sources of intense conceptual difficulty, or "double binds", experienced by the subject of activity. The insights and actions that lead subjects to overcome such contradictions are what lead to more advanced forms of activity over time, in the developmental process of "expansive learning" [26]. More recently, Engeström has proposed the analogous notion of "expansive design", in which the design process is explicitly oriented towards the production of tools that resolve contradictions between interacting activity systems [27]. He proposes an interventionist approach to expansive design, illustrated through a case study on transforming Helsinki-based healthcare for children with multiple chronic illnesses. Interviews with activity participants identified contradictions between the overlapping activity systems of family care, primary care, and hospital care. The expansive solution took the form of an annual, holistic care agreement negotiated by all parties, addressing the inadequacies of care relationships confined to a single institution and critical care pathways that assume only a single disease or diagnosis. Although the case study shows that an interventionist approach can be successful in transforming specific activity systems, it does not provide guidance on expansive design for general activities like VA.

## 2 Analysis of General Activity Systems

Using the structure of the activity system, we first describe how researchers and designers can develop their understanding of a target activity by building an *activity typology* using extant concepts from the literature (or indeed any source of activity-related concepts, e.g., interviews, observations, surveys, studies, interaction logs, and so on). Next, we explain how each element of a typology becomes a possible *design target* when designing new tools for the corresponding activity. Finally, we describe an expansive design process that can transform a list of design targets into a more tractable *design tension*, and a design tension into a more prescriptive *target quality*. The embodiment of such qualities in future tools is the mechanism by which such expansive design addresses the needs of a general activity audience.

### 2.1 Constructing an Activity Typology

Previous uses of the activity system model in HCI have focused on the organization of empirical observations and findings. This approach does not scale to the analysis of general activity systems, as the enumeration of concrete activity attributes soon becomes untenable. This problem is addressed in InfoVis and Visual Analytics through the notion of a typology (e.g., a task typology [16]) that connects and abstracts over many concrete instances, in a way that is descriptive, comparative, and generative of new ideas. We employ the same notion of typology in the modelling of general activity systems, resulting in what we call the *activity typology model*, shown in Figure 2. This structure highlights the distribution of knowledge and the presence of gaps, while the associated labels support the articulation of information needs and communication of research findings. We use the term typology over taxonomy to reflect our structuring of abstract concepts rather than empirical observations [2] (and as in [16]).

In line with Norman's view that it is the tools that define the activity and not the converse [57], and because the goal of activity-centered design is to design tools that transform their target activity for the better, our typology model adopts tool-centered definitions of activity system elements. We also rename some of these elements for contemporary relevance and approachability, as outlined in Table 1 (bottom left) and explained below.

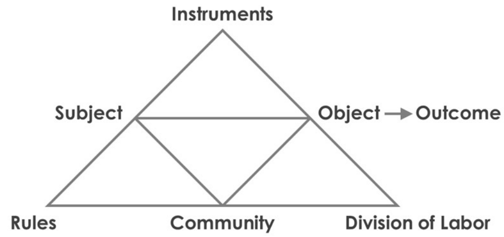


Figure 1. The activity system model [26] depicting the elements, relations, and mediations of object-oriented human activity.

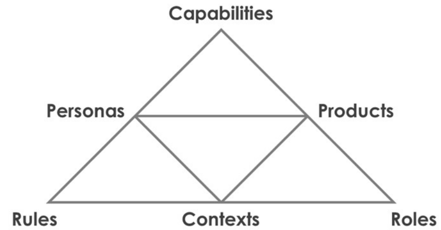


Figure 2. Our activity typology model for analysing general activities like Visual Analytics. It closely follows the activity system model [26] but with elements renamed and redefined to reflect the multiplicity of concrete activity forms that constitute any general activity system.

**Personas**—*types of people using the tools of the activity*. We map subject to personas in the sense of "kinds of people", avoiding the confusion between the subject of activity and its topic or domain. An example VA persona could be "desk analyst".

**Products**—*types of outcome that motivate the activity*. We collapse object–outcome to products in a way that avoids confusion with the everyday sense of object and encompasses all products of activity, whether tangible or intangible, and desired or actual. An example VA product could be "derive insights".

**Capabilities**—*types of task supported by the tools of the activity*. We map instruments to capabilities to emphasize the key tasks supported by the collective tools of the activity, rather than how current tools happen to bundle and distribute such capabilities. An example VA capability could be "searching sources".

**Contexts**—*types of contextual factor that shape the activity*. We map community to contexts in our most significant change to the model, reflecting both the importance and open-endedness of context [14]. This allows capture of whichever dimensions of context are most relevant to the activity system in question (inclusive of community). An example VA context could be "co-located team".

**Rules**—*types of constraint on the performance of activity*. We retain the term rules. An example VA rule could be "relevance".

**Roles**—*types of coordinated contribution to the activity*. We map the division of labor to the roles that it creates, reflecting the everyday sense of roles as the different ways people act in different contexts. In contrast, Personas reflect more stable characteristics of individuals. An example VA role could be "decision maker".

We can reflect on the theoretical value of this model using Halverson's enumeration of the powers of theory [34]. The ability to create an activity-centered lexicon for any target activity gives this model clear *descriptive* power. The concise and memorable form of the resulting typologies, grounded in the triangles of activity, also gives it an added degree of *rhetorical* power. However, an activity typology alone does not have the *inferential* power to identify new insights, or the *application* power to generate new design ideas. To extend the model with these powers, we reinterpret the notion of expansive design [27] within the context of such activity typologies.

### 2.2 Transforming Design Targets into Target Qualities

In any activity typology, each typology element (e.g. "desk analyst") of each activity element (e.g., personas) provides a potential *design target* for future tools that aim to support the corresponding activity (e.g., future VA tools should support the work of desk analysts).

| **Activity-Centered Design using the Activity Typology Model (Section 2)** | | | |
|---|---|---|---|
| **1. Define target activity** | **2. Organize design targets** | **3. Identify design tensions** | **4. Generate target qualities** |
| *Define the activity that is to be supported through the design and introduction of new tools. This can be a general activity that takes many different forms in practice. Use the fundamentals of Activity Theory (Section 1.2) to elaborate.* | *Use each element of the activity typology model to organize related concepts from the many forms of the target activity. Each of these concepts is a potential target in the design of new tools.* | *Search for bipolar tensions that cut across the design targets of each activity element, e.g., identify attribute pairs relevant to all design targets but varying in their nature and extent.* | *Use each design tension as a springboard to generate creative design resolutions. Express these resolutions as the qualities of tool interactions that would resolve design tensions and thus support the many forms of the target activity in general.* |

| **Activity Typology Model (2.1)** | | **Visual Analytics Activity Typology (Section 3)** | | |
|---|---|---|---|---|
| **Element** | **Definition** | **Design Targets** | **Design Tension** | **Target Quality** |
| **Personas** | *types of people using the tools of the activity* | desk analyst, case investigator, field officer, domain expert, professional analyst | *Acting as:* Data collector *vs* Data analyst | **Portable Analysis** *ability to transfer analytic work across people, places, time, and devices* |
| **Products** | *types of outcome that motivate the activity* | derive insights, develop options, make arguments, present assessments, manage situations | *Acting to:* Make sense *vs* Make artifacts | **Presentable Analysis** *ability to curate presentable summaries of the analytic discovery process* |
| **Capabilities** | *types of task supported by the tools of the activity* | searching sources, visualizing data, reading reports, tagging interests, recording viewpoints | *Acting by:* Funneling data *vs* Testing data | **Perspectival Analysis** *ability to create and annotate workspaces with analytic perspectives* |
| **Contexts** | *types of contextual factor that shape the activity* | co-located teams, distributed teams, distributed communities, synchronicity, mobility | *Acting in:* Defined teams *vs* Defined areas | **Proxemic Analysis** *ability to locate oneself and others within the space of analytic views and viewpoints* |
| **Rules** | *types of constraint on the performance of activity* | relevance, confidence, provenance, access rights, time pressure | *Acting under:* Competing interpretations *vs* Competing demands | **Provisional Analysis** *ability to view and proactively reduce the uncertainty of analytic work at any time* |
| **Roles** | *types of coordinated contribution to the activity* | producers, consumers, responders, decision makers, policy makers | *Acting with:* Process partners *vs* Product partners | **Polymorphic Analysis** *ability to export analytic reports tailored by audience, format, and purpose* |

Table 1. Four stages of activity-centered design (top) using the activity typology model (bottom left) applied to Visual Analytics (bottom right).

However, a fundamental contradiction in the process of activity-centered design is that each tool must embody a single design, yet that design must serve every instance of activity in which the tool is used. We use the term *design tensions* (as in [70]) to reflect the competing design targets of each activity element. To give a simple VA example, a persona design tension could arise from the need to design for both desk analysts and field officers, and be expressed in the form "desk analysts vs field officers". This bipolar form invites the designer to consider how they might design differently for each of the two targets, and then crucially, how they could design in a way that accommodates both (mirroring the dialectical nature of contradictions [26]). A challenge of considering tensions among more than two targets, however, is that there is no longer a single relation to focus attention, but many pairwise relations whose number grows combinatorically.

Given that designers have limited time and attention, it is desirable to have them focus on a smaller number of more general tensions than a great many narrow ones. Our solution is to encourage designers to abstract a single bipolar tension that cuts across all design targets for an activity element, providing a clear and concise problem description. An abstraction strategy we have found to be effective is to look for pairs of attributes that are shared by all design targets but which vary in their nature or extent. Thus, in the VA example, we might reason that all three personas act as both a "data collector" and a "data analyst" over time, but with different levels of emphasis and switching, leading to the design tension "data collector vs data analyst". The process of abstracting bipolar tensions is akin to the process of abductive inference, in which informed guesswork is used to identify economical and plausible hypotheses for complex phenomena. Design tensions can thus be seen as hypotheses whose validity can be determined empirically (e.g., use as interview probes), complementing their role as a springboard for the ideation of design resolutions. Such resolutions may come to mind as ideas for new tools or features applicable to specific use contexts. They may also take the form of more general interaction qualities to which new tools or features should aspire. These *target qualities* for tool-mediated activity can both inspire the generation of lower-level design concepts and provide a means for evaluating and comparing such concepts. In principle, tools embodying such target qualities will better serve the diverse needs of activity participants. Table 1 (top) shows the result of codifying this approach into an activity-centered design process.

## 3 An Activity Typology for Visual Analytics

Following Table 1 (top), we now present an application of activity-centered design with the activity typology model. For stage 1 of this process, we define the target activity to be Visual Analytics (see Section 1.2.1 for our elaboration of this activity definition using the fundamentals of Activity Theory). Our presentation of the process hereafter is organized in a depth-first manner to give a more coherent reading experience, with each subsection focusing on a single typology element (e.g., personas) and proceeding directly through stages 2–4 in succession. However, in practice we freely moved between different stages and elements in the iterative and non-linear fashion typical of any design activity. The contributions of each subsection are a brief, tightly-focused literature review of relevant concepts, a single design tension that cuts across these concepts, and a single target quality for the design of new activity tools. While each of the stages leading to these contributions requires a degree of creativity, insight, or interpretation, this does not guarantee that the resulting contributions will themselves be original. Indeed, if the design tensions and target qualities are an accurate reflection of real problems and successful solutions, it is highly likely that they have been touched upon in the literature. The core value of building an activity typology is not so much the identification of original parts, but their systematic naming and assembly into a coherent whole that provides a holistic view of activity. The result of applying this process to Visual Analytics is shown in Table 1 (bottom right).

### 3.1 Visual Analytics Personas

Any practitioner of Visual Analytics is an "analyst" in the general sense. However, within this category lies a variety of analyst types conducting different kinds of analytic work in different ways.

#### 3.1.1 Organizing Personas as Design Targets

Based on interviews with law enforcement professionals, the designers of the Scalable Reasoning System (SRS) [58] identified three broad user categories. Traditional "analysts" work on longer-term strategic analyses in loosely-defined areas; "investigators" perform shorter-term tactical analyses on well-defined cases; and "field officers" providing real-time analysis of real-world situations. The SRS uses a mix of mobile and web interfaces to support distributed analytics for all users. In a similar breakdown of analyst personas, Illuminating the Path [71] differentiates "desk analysts" assessing threats and vulnerabilities from "field personnel" working on situation management and emergency response.

The foreword to Psychology of Intelligence Analysis [38] identifies an additional cultural distinction between "substantive experts" employed for their domain knowledge and "professional analysts" employed for their ability to adapt and apply effective tradecraft to diverse problem areas. The latter is viewed as significantly rarer and more valuable in the world at large, with the unique abilities of the "master analyst" attributable to mental schemata developed through many years of experience. We can combine these different kinds of analyst into the following representative personas:

**VA persona design targets**: *desk analyst, case investigator, field officer, domain expert, professional analyst*

#### 3.1.2 Identifying a Design Tension across Personas

All personas engage in the dual processes of collecting and analyzing data, but with varying emphasis and switching frequency. For a desk analyst, the emphasis is on analyzing large bodies of data assembled over long periods of time, with additional data collected to fill knowledge gaps (e.g., by searching a database or the web). For a field officer, the emphasis is on collecting new data from the world, with rapid analysis of that data used to make decisions in real time. Case investigators sit in the middle, balancing data collection and analysis, and operating over intermediate time frames. Switching between acting as a *data collector* and acting as a *data analyst* could cause problems for individuals in terms of using multiple tools, managing multiple work threads, and incorporating new data into existing analyses. Similarly, whereas domain experts focus on historical *data collection*, professional analysts focus on the tradecraft of *data analysis*. For such personas, whose emphasis is more of an enduring trait than a fluid state, additional problems arise through the need to coordinate and communicate with analysts possessing the opposing mindset and its associated values. Designers of VA tools should therefore consider how their tools would be used by personas at multiple points along the *data collector* vs *data analyst* spectrum and with differing degrees of switching between these two mindsets.

**VA persona design tension**: Acting as *Data collector* vs *Data analyst*

#### 3.1.3 Generating a Target Quality from the Personas tension

A target quality inspired by this tension is the idea of *Portable Analysis*—that analytic work should be packaged in a way that allows it to be "put down" and "picked up" at a later time, in the same or different context, by the same or different person. Such portable work packages should retain the history of work done, the context in which the work was suspended, and the entry points required for resumption. They would help encapsulate case investigations, channel communications between desk and field, and integrate analytical reasoning from domain knowledge and data.

In terms of information sharing, this quality of portability aligns with the recommendations from Illuminating the Path that VA tools must package analysis "in a format that can be unwrapped for just-in-time use by other members of the response team without endangering security or privacy" or "during the next shift or in the next emergency" [71]. In terms of task switching, it also aligns with the visual thinking design pattern of a Task List [75] for planning and tracking work in progress. The list of incidents in the SRS is an example of such an orienting framework, and SRS remains a good example of portable analysis in general [58]. Other common techniques for supporting portable analysis are the use of notes, comments, captions and annotations. While these are typically treated as a secondary notation [13] that is not interpreted by the system, work on Annotation Graphs [77] has explored the interpretation of annotations as meta-data that can implicitly drive graph layout in mixed-initiative exploration.

**VA persona target quality**: *Portable Analysis.* The ability to transfer analytic work across people, places, time, and devices.

### 3.2 Visual Analytics Products

We can use a hierarchy to organize the products of VA activity by timescale, from immediate desires to long term outcomes.

#### 3.2.1 Organizing Products as Design Targets

The sensemaking loop for intelligence analysis [59] describes the logical flow from an unstructured evidence file, to schematizations of evidence, to articulations of hypotheses, and finally to analytic artifacts such as reports and presentations. A hierarchical product of VA activity is thus to *derive insights* (e.g., adding to an evidence file) to *make arguments* (e.g., linking evidence to hypotheses) to *present assessments* (e.g., through written reports or oral presentations) that motivate action. These products of VA activity are also evident in the following description in Illuminating the Path (emphasis added) [71]:

*"People use visual analytics tools and techniques to synthesize information and* ***derive insight*** *from massive, dynamic, ambiguous, and often conflicting data; detect the expected and discover the unexpected;* ***provide*** *timely, defensible, and understandable* ***assessments****; and* ***communicate assessment*** *effectively for action."*

When it comes to motivating action, it is necessary to *develop options* by preparing a range of reactions to possible events [71]:

*"For homeland security ... analysts may* ***develop options*** *to defend against, avert, or disrupt threats. In emergency response situations ... to understand response options and their implications".*

This also suggests the most general product—to *manage situations* by developing, communicating, and executing options for action.

A related taxonomy targeting the products of information visualization is Amar and Stasko's knowledge-task–based framework [1]. This framework delineates the "rationale tasks" of Expose Uncertainty, Concretize Relationships, and Formulate Cause and Effect from the "worldview tasks" of Determine Domain Parameters, Multivariate Explanation, and Confirm Hypotheses. In our activity typology, these are all ways in which an analyst can *derive insights*.

**VA product design targets**: *derive insights, develop options, make arguments, present assessments, manage situations*

#### 3.2.2 Identifying a Design Tension across Products

A key outcome of VA activity is for the analyst to *make sense* of data. However, almost all forms of Visual Analytics also require analysts to *make artifacts* that document and communicate this sensemaking process, e.g., reports or presentations assessing options for managing situations based on arguments and insights.

Although making such artifacts has the beneficial side-effect of forcing the analyst to clarify their thoughts and articulate their reasoning, it is challenging and potentially disruptive for an analyst to document a sensemaking process while they have yet to make sense of the data. It is also inadvisable for an analyst to wait until everything makes sense to them (or so it might seem) before beginning to document that sensemaking process for others. At each point in time, the analyst must therefore decide whether to direct their efforts towards making further sense of the data or to make artifacts that capture the sense that has been made of the data thus far. This tension adds significantly to the burden of analysis [71]:

*"Once an important piece of evidence is recognized or an inference is made, it is often exceedingly difficult to capture and record the progress directly, forcing reliance on memory, notes, or annotations. Likewise, a sudden recognition, question, or insight usually cannot be recorded without disrupting ongoing analysis"*

**VA product design tension**: Acting to *Make sense* vs *Make artifacts*

### 3.2.3 Generating a Target Quality from the Products tension

A target quality inspired by this tension is the idea of *Presentable Analysis*—that the effort invested in analysis should, as a side-effect, create artifacts that can be selected and annotated for immediate presentation within the tool itself, or exported to standard formats (e.g., as Microsoft PowerPoint slides) for archiving and distribution.

This quality can be facilitated in two complementary ways: first by increasing the standalone legibility of any visualization such that future viewers could interpret a screenshot without the benefit of seeing how it was created; and second by capturing the analyst's view history in ways that allow the analyst to retrospectively navigate, bookmark, annotate, caption, and reorder views for use in presentations. These latter concepts relate to the visual thinking design pattern of "Presentation Linking Images and Words" [75]. The Vistories system [31] also supports switching between exploratory analysis and storytelling, in which presentations of live data visualizations can be created by annotating selected states in the branching history of user interactions. Similarly, the Aruvi system [66] visualizes such branching interaction histories as an aid to navigation rather than presentation. Sampling from such visual interaction histories defers evaluations of presentation-worthiness compared with the common approach of taking explicit "snapshots". Showing users visualizations of their interaction history has also been found to help them recall their associated reasoning process [50].

**VA product target quality**: *Presentable Analysis.* The ability to curate presentable summaries of the analytic discovery process.

## 3.3 Visual Analytics Capabilities

The kind of taxonomies reviewed in Section 1.2.2 provide a detailed organization of low-level capabilities embodied by VA tools. Here we focus on high-level capabilities related to the sensemaking.

### 3.3.1 Organizing Capabilities as Design Targets

Referring again to the sensemaking loop [59], we can distinguish between tool capabilities oriented towards searching, filtering, reading, and extracting information from external data sources (so-called information foraging) and those oriented towards building schemas, hypotheses, and presentations from the foraged information (sensemaking). Data/frame theory [46] presents an alternative model of sensemaking in which the user always has a perspective, viewpoint, or framework in mind, however minimal, that focuses attention, defines what counts as data, and shapes the collection of subsequent data. Rather than information foraging progressing smoothly into sensemaking, frames are constantly elaborated, compared, and replaced as they are tested against the flow of data. In observations of trainee intelligence analysts, this model was viewed as a better fit with practice [41]. Identifying actions common to both models, we can differentiate capabilities oriented towards *searching sources*, *visualizing data*, *reading reports*, *tagging interests*, and *recording viewpoints*. The mental model of sensemaking held by tool designers plays a significant role in determining which of these capabilities serve as entry points to tool use, what flows and interactions between capabilities are allowed and encouraged, and how easily new data and perspectives can be accommodated.

**VA capability design targets**: *searching sources, visualizing data, reading reports, tagging interests, recording viewpoints*

### 3.3.2 Identifying a Design Tension across Capabilities

The two models of the sensemaking process presented previously are grounded in competing metaphors for how individuals sequence and transition between various tool actions. In the sensemaking loop [59], individuals make progress by *funneling data* through ever more refined stages of analysis. In data/frame theory [46] , individuals make progress by *testing data* against provisional hypotheses that are elaborated and replaced over time. Tools designed for only one such workflow may be inappropriate for certain working styles or scenarios, causing problems for users forced to adapt their practice to the constraints of the tool.

In practice, many historic VA systems have been inspired by the sensemaking loop [59] and reify its steps into the UI and interaction workflow. For example, the Human Interface Discourse Interface (2005) [61] from PNNL combines a document space (the In-Spire document clustering and search interface [37]), an evidence marshalling space, and a hypothesis space, "each corresponding to a step in the analytic cycle". The Jigsaw system (2008) [68] from Georgia Tech builds on the same foundation, adding multiple data visualizations (list, graph, scatterplot, calendar, and document views) to a document cluster view and a "shoebox" view combining evidence-marshalling and hypothesis-building capabilities. The funneling workflow is one-way, with each view sending entities and documents to the shoebox. The PARC platform pairs two systems: CorpusView (2004) [11] for information foraging with an emphasis on *reading reports*, and Entity Workspace (2006) [12] for sensemaking with an emphasis on *tagging interests*. Knowledge artifacts are constructed by "snapping together" entity-linked text extracts in a funneling workflow. The nSpace platform from Oculus Info similarly pairs two systems: TRIST (2005) [40] for information foraging with an emphasis on *searching sources*, and Sandbox (2006) [76] for sensemaking with an emphasis on *recording viewpoints*. The design of the Sandbox canvas as a "flexible expressive thinking environment" facilitates capturing hypotheses and assumptions against which data can be tested using TRIST.

**VA capability design tension**: Acting by *Funneling data* vs *Testing data*

### 3.3.3 Generating a Target Quality from the Capabilities tension

A target quality inspired by this tension is the idea of *Perspectival Analysis*—that all analysis occurs in its own workspace, and each workspace can be labelled with the perspective that motivated the analysis and the perspective resulting from the final view of the analysis (if complete). Workspaces would encapsulate their internal interaction history, and creating new workspaces by branching from existing ones would create a legible activity history of what was investigated when, why, and to what end. The labels used when "funneling data" would thus describe the data sources examined and the resulting discoveries, while the labels used when "testing data" would describe the hypotheses tested and the resulting conclusions. In other words, the same semantic structure could support a blend of workflows through alternative labelling strategies.

While activity-labelled desktop workspaces are a foundation of activity-based computing (ABC) [5], they have yet to be explored at this lower level for Visual Analytics. Instead, VA tools rely on the user to manually record such details inside dedicated views or data annotations. Perspectival Analysis inverts this structure in ways that support "Cognitive Reconstruction" [75] of work context, both by the constructing analyst and by others.

**VA capability target quality**: *Perspectival Analysis.* The ability to create and annotate workspaces with analytic perspectives.

## 3.4 Visual Analytics Contexts

One of the most significant aspects of context for Visual Analytics is the social and group structures within which the activity is embedded. Environmental factors also play an important role in shaping context.

### 3.4.1 Organizing Contexts as Design Targets

The coordination of individuals is a fundamental dimension of activity context. Research on design considerations for collaborative VA [35] notes how groups of peers can be defined explicitly (e.g., work teams dividing, allocating, performing, and integrating work towards a shared goal) or emerge implicitly (e.g., visitors to social data analysis sites such as Many Eyes [73] or sense.us [36] who share a common interest), and that groups can act synchronously or asynchronously. Work groups can also be co-located or distributed, and operate at different scales within, across, or beyond organizational boundaries, as described by a trainer of intelligence analysts [41]:

*"While working on a particular topic within an agency is typical, also typical is working on an interagency team... Strategic projects almost always involve a team as do crisis projects... In short, teamwork is the norm although the teams differ in the degree of formality and to the degree that there is a designated leader"*

Early collaborative systems for Visual Analytics typically supported asynchronous collaboration among distributed teams. For example, a modification [10] of Entity Workspace [12] supports shared document collections and can offer "quiet collaborative entity recommendations" based on the commonalities and differences in the "entities of interest" to different users and their relationships (a pair of entities are typically viewed as related if they are both referenced in the same document, with the strength of the relationship a function of their co-occurrence frequency). Another multi-analyst tool [17] performs the same kind of entity comparison more explicitly, with an analyst able to translate a node-link graph of tracked entity relations between their own private perspective, the perspective of a collaborating analyst, or into a fused graph representing the combined perspective of all analysts. Asynchronous interaction is also especially important in the context of mobile VA because of the need to maintain situation awareness, as explored by the mobile interface to SRS [58].

More recent systems have begun to offer additional facilities for synchronous collaboration. For example, VizCept [20] is based on a global concept map shared by all analysts, and incorporates real-time presence indicators and chat channels. In contrast, CLIP [51] retains the use of private workspaces but extends the representation of entity graph nodes to reveal linked notes, images, and evidence contributed by other analysts. It also provides a tab to view each collaborator's private workspace in real-time, with node attachments and workspace tabs color-coded per analyst. Finally, through use of interactive tabletop displays, Hugin [45] and Lark [72] support co-located and mixed-presence collaboration in synchronous visualization tasks.

**VA context design targets**: *co-located teams, distributed teams, distributed communities, synchronicity, mobility*

#### 3.4.2 Identifying a Design Tension across Contexts

The nature of joint activity can be classified according to the extent to which there is a *defined team* of individuals collectively working towards shared goals with mutual coordination, versus a *defined area* in which individuals are conducting independent but overlapping analyses. Area can be interpreted both conceptually (as in a subject area or area of interest) and spatially (as in the same physical space) Tools designed for defined teams of users will make assumptions about identity and communication needs that are inappropriate for users whose only relationship is their shared interest in some item or area. Similarly, tools designed for single or independent users may not adequately support the coordination and collaboration needs of teams. Key characteristics of collaborative tools that make them suitable for *defined teams* vs *defined areas* include how interaction is organized asynchronously (e.g., as conversations with particular people versus on particular topics), how it is initiated synchronously (e.g., in which modality and with what signaling of availability), and how the tools of activity adapt to or facilitate mobile interaction.

**VA context design tension**: Acting in *Defined teams* vs *Defined areas*

#### 3.4.3 Generating a Target Quality from the Contexts tension

A target quality inspired by this tension is the idea of *Proxemic Analysis*—that the analyst should be able to locate both themselves and others within the space of shared data and be made aware when others enter their personal space or previous territory (by analogy to human interactions in physical space, which are the subject of proxemics proper). Defined teams of users could create persistent links to one another, viewing real time status with the option to view both location, orientation, and trajectory within the shared data space. Users with defined areas of interest, whether stated explicitly or mined from interactions, could also elect to view other interactions and individuals touching upon those areas, as a bridge to opportunistic synchronous interaction. While these ideas are only just beginning to appear in VA systems, they have a long history of use in support of Computer Supported Cooperative Work (e.g., [24]), especially for collaborative tabletop interaction (e.g., [64]). The quality applies literally to users working in the same physical spaces, with proximity indicating the potential for face-to-face interaction.

**VA rule target quality**: *Proxemic Analysis.* The ability to locate oneself and others within the space of analytic views and viewpoints.

### 3.5 Visual Analytics Rules

The rules of VA activity serve to guide both the attention and interactions of the analyst over time, placing constraints on such things as what is possible, allowable, and desirable.

#### 3.5.1 Organizing Rules as Design Targets

In VA activity, rules act at each stage of information foraging and sensemaking to determine what counts as data, as information, and as insight. Considerations of *relevance* apply both to the selection of data sources and the filtering of data from those sources. A relevance-oriented cost metric for data foraging [59] is the extent to which the analyst can get more of the relevant information in a shorter amount of time (whether by obtaining the same information in less time, by obtaining more information in the same time, or a combination of both). One way in which tools can help is to represent relevance explicitly as a means of creating information scent. For example, Entity Workspace [12] allows entities to be tagged with their "degree of interest" to analysts. This is then propagated through the entity graph using spreading activation, suggesting additional entities of interest and highlighting discrepancies across analysts.

A related determination made by analysts, captured in tools such as SRS [58], is the degree of confidence in the accuracy of information and its evidential support for stated hypotheses. In SRS, evidence and hypotheses combine to form a belief network within which confidence scores are propagated using Dempster–Shafer reasoning [62]. This method generalizes the Analysis of Competing Hypotheses (ACH) approach popularized by Psychology of Intelligence Analysis [38], in which each evidence source is compared against each hypothesis in search of diagnostic evidence that best differentiates between the competing set of hypotheses.

A third class of rules concerns the *provenance* of insights arrived at through analysis [60]. While manual recording of how insights came about (e.g., using notes or diagrams) requires substantial effort, automatic approaches (e.g., event-based interaction logs) fail to capture the semantics of interaction. An extension [30] of the HARVEST system [29] resolves this problem by logging user interactions using a high-level action typology that is independent of both the domain and visualizations used. It defines an "insight trail" as a series of exploration actions followed by a series of insight actions, and extracts these automatically using regular expressions. The SensePath system [56] also focus on the capture of analytic provenance, in the context of browser-based online sensemaking.

A fourth, and large class of rules govern *access rights* to data and analytic products, arising from concerns including privacy, security, auditability, and legal compliance [71]:

*"Analysts work with information drawn from multiple sources, each of which has associated security and privacy constraints. Laws exist that govern how information may be used and combined, and those laws must underpin visual analytics approaches. Tools must proactively adopt and support approaches such as data anonymization, data minimization, audit trails, and access controls to both protect privacy and ensure information security"*

Finally, a significant driver of VA activity is *time pressure*. Especially in situations of crisis or emergency response, the timing of actions it critical. Intelligence analysts also operate under strong pressure for "premature closure" [38], in which "demand for interpretive analysis is greatest within two or three days after an event occurs". The pressurizing effects of time are at odds with other rules that demand comprehensive, rigorous, and reasoned analysis.

**VA rule design targets**: *relevance, confidence, provenance, access rights, time pressure*

### 3.5.2 Identifying a Design Tension across Rules

VA activity is fundamentally constrained by the limitations of human perception, memory, and reasoning. Analysis of multiple *competing interpretations* is one way of ensuring that confidence is not placed prematurely on the first reasonable interpretation that comes to mind. However, the process of generating multiple plausible interpretations is mentally demanding, their documentation for the purposes of provenance are time consuming, and it is unclear at which point additional time and effort might be better spent searching for additional information of relevance. Searching and documenting, all while respecting access rights, therefore comprise *competing demands* to the generation and analysis of competing interpretations. The balancing of these concerns to satisfy requirements in all areas, especially under time pressure, presents a considerable challenge.

**VA rule design tension**: Acting under *Competing interpretations* vs *Competing demands*

### 3.5.3 Generating a Target Quality from the Rules tension

A target quality inspired by this tension is the idea of *Provisional Analysis*—that the fundamental uncertainty of analytic work is acknowledged and represented, that explicit consideration is given to competing interpretations of data, that support for interpretations could change in the face of new or revised data, and that the frontier of work that would be done to strengthen the analysis (time permitting) is clearly visible and available for direct continuation. Prior work on classifying the various sources of uncertainty in information visualization [67] could be helpful in revealing both data and analytic uncertainty at the interface level. The importance of externalizing all such concerns as a conceptual model has also been highlighted in studies of trainee analysts [41]:

*"The other thing this model helps you to do is at the end of the project you can look back and go, 'What did we not have time to do? And how does this impact our company, our estimates?'"*

This quality could be supported through such tool capabilities as proactive "what-if" analysis to identify the lynchpin evidence on which an argument rests, and the use of arguments and their constituent entities as standing queries to public and private data streams that would alert the analyst to events that may affect the strength of these arguments. Both relate to the visual thinking design patterns of "Query By Example" and "Visual Monitoring" [75].

**VA context target quality**: *Provisional Analysis.* The ability to view and proactively reduce the uncertainty of analytic work at any time.

## 3.6 Visual Analytics Roles

The final element of the typology is the roles assumed by individuals who contribute to products of the activity. These roles may be transient, e.g., played by the core personas of the activity in their interactions with one another, or more enduring, e.g., played by additional activity participants who provide an established kind of contribution to the activity.

### 3.6.1 Organizing Roles as Design Targets

Considering the personas described previously, each is sometimes a *producer* of analytic products: desk analysts produce analytic assessments, case investigators produce leads and case resolutions, field officers produce field reports, domain experts produce domain precedents, and professional analysts produce rigorous analyses. Each is also sometimes a *consumer* of analytic products, e.g., a case investigator might consume both assessments from desk analysts and reports from field officers. When new events or information require immediate action, any person taking that action is also a *responder*, with crisis and emergency response personnel being examples of this role. A crucial role in formal contexts is that of *decision maker*. Decisions to be made include which groups or individuals are tasked with the production of what analytic products, which groups or individuals are permitted or required to consume the resulting products, and what actions should be taken as a result. Both analysts and decision makers must also recognize the role of *policy makers* who establish the rules by which analysis is conducted and decisions are made, and who use analytic products to inform policy evolution.

**VA role design targets**: *producers, consumers, responders, decision makers, policy makers*

### 3.6.2 Identifying a Design Tension across Roles

We can distil roles into *process partners*—those contributing to the process of analysis and the construction of analytic products—and *product partners*—those relating to the appraisal, distribution, and consumption of completed analytic products. These partner types are likely to speak different languages and hold different concerns. They are also likely to differ in terms of their desired media formats and thresholds of detail, confidence, and structure. The need for analytic products differentiated by audience poses significant challenges [71]:

*"information presentations must be scaled or adapted to the audience. For example, an analyst's presentation to other analysts will contain far more detail than the summary analysis presented to the President. Current techniques require that this be done manually in an ad hoc fashion"*

**VA role design tension**: Acting with *Process partners* vs *Product partners*

### 3.6.3 Generating a Target Quality from the Roles tension

A target quality inspired by this tension is the idea of *Polymorphic Analysis*—that a single analysis can produce artifacts of multiple forms for different intended audiences, from casual annotations and discussions around data points to more formal reports and presentations. The creation of such artifacts should not be a separate exercise using a different set of tools. Rather, the component discussion comments, report paragraphs, presentations slides, etc. of the ultimate analytic products should be tightly bound to the data and views that inspired them, such that they may be compared, translated, and cross-referenced during authoring. Previewing artifacts for export would be achieved by assembling and ordering these components, with edits to the assembled artifact propagated back to their source.

An example VA tool supporting polymorphic analysis is GeoTime Stories [25], in which the user can capture snapshots of the XYT geo-temporal trajectories that GeoTime [42] renders in 3D. These snapshots can be given a title and color-coded descriptive sections. All sections of a given color can be included or filtered for presentation, and different story sequences can be composed from the same set of underlying snapshots. Additional uses for such color-coding included feedback, versioning, change tracking, details, back story, alternatives, and comments.

**VA role target quality**: *Polymorphic Analysis.* The ability to export analytic reports tailored by audience, format, and purpose.

# 4 Discussion

The preceding construction of an activity typology provides a focused yet holistic view of VA activity. Central to such potential transformation is our set of six "target qualities" for the activity—qualities that are realized in practice through the use of tools designed to support these qualities in principle. Note that these qualities should not be assessed in isolation. In the following paragraph, we make their interdependency clearer by linking them in narrative form:

> In VA activity, analysts should be able to record their *perspectives* on work as they are doing it, including what prompted and resulted from each analysis, with gaps indicating both the *provisional* nature of analytic work and opportunities for continuation. The additional semantic structure provides a medium for *proxemic* interactions between individuals and makes analysis more *portable* across people, devices, and contexts.

> Finally, the linking of annotated workspaces to curatable interaction histories enables a more *presentable* discovery process, with opportunities for *polymorphic* export of multiple tailored analytic products.

Such a description of interdependent qualities provides a framework for considering the cross-activity implications of design decisions. While such framework remains removed from concrete design concepts and implementation details, it articulates a coherent vision for a cross-domain and cross-task analytics tool. High-level examples of how our VA typology could be applied and extended are given next.

### 4.1 Application of VA Target Qualities to BI Tools

Consider attempting to extend a commercial Business Intelligence (BI) tool like Tableau or Microsoft Power BI with features that embody the target qualities for Visual Analytics described above. We might start by replacing the fixed pages of the dashboard paradigm with a more flexible notion of workspaces that can be labelled and forked in ways that retain their provenance (supporting Perspectival Analysis). Aggregating workspace labels in a scrollable list (rather than as page tabs) and encouraging the use of workspaces to represent work "to do" would allow case or investigation-level work to be conducted through workspace creation and management (supporting Provisional Analysis). The presence and actions of other users could be revealed through highlight, filter, and sort options on this list, which would thus play a dual role as a collaborative notification center (supporting Proxemic Analysis). This compact meta-interface to planned and asynchronous collaborative analytic activity would be particularly amenable to interaction on mobile devices, unlike the core tasks of data preparation, visual binding, and cross-visual data exploration (supporting Portable Interaction). The workspace labels used to plan and record the progress and outcomes of work would also provide natural captions of visual states suitable for presentation (supporting Presentable Analysis). Finally, creating presentation sequences by sampling from workspace-level interaction history supporting semantic zoom between workspace labels visual and content could allow a wide range of presentations of the same underlying discovery process (supporting Polymorphic Analysis).

This use of target interaction qualities to expand the scope of a design concept from a minor feature into a qualitatively new way of working shows the benefits of having a holistic model of VA activity.

### 4.2 Integration of "Guidance" into VA Typology

A complementary example type is the integration of new concepts into our VA activity typology. Sometimes new tools or techniques grow in influence over time in ways that have a transformative effect on the practice of an activity. One candidate in the field of Visual Analytics is the idea of automated guidance that assists the analyst as they conduct their analysis. A recent article surveys examples of guidance in the literature and synthesizes them into a conceptual model [18]. Here, we show how the activity-level aspects of guidance discussed in that paper could be incorporated into our VA typology.

The users described as being most in need of guidance are analytic "novices" (a new persona) working to "resolve a knowledge gap" (a new product) by "choosing the correct target" (a new capability). Guidance can be presented with varying levels of "prescription" (a new rule) and be offered by a virtual "assistant" (a new role). It should also account for environmental variability (a new context): "Guidance has to be … adaptive to the particular context, as the type of assistance a user requires varies and depends on many factors" [18].

While new typology elements may trigger a reformulation of associated design tensions, in many cases they can still be rationalized in the context of the existing tensions. For example, a novice is a learner "Data analyst", resolving a knowledge gap is one way to "Make sense" of data, choosing the correct target is enabled by the system proactively "Funneling data" on the user's behalf, prescription is a way to deal with "Competing demands", an assistant is a kind of "Process partner", and one form of environmental variability arises from not having a "Defined area" in the physical sense of the term.

### 4.3 Use of Activity Typology Models with Other Methods

Use of the activity typology model to ground an activity-centered design process does not preclude use of other design approaches. For example, the Design Study Methodology [65] provides practical guidance on the nine stages of a visualization design study: learn, winnow, cast, discover, design, implement, deploy, reflect, and write. Although grounded in theory, the design targets, design tensions, and target qualities of an activity typology contribute directly to the learn, discover, and design stages of this process. The remaining stages provide additional and complementary guidance on the real-world project concerns of collaboration, evaluation, and presentation.

### 4.4 Limitations and Future Work

The application of our activity-centered design process to Visual Analytics has several limitations. First, it only draws on concepts from the VA literature rather than from original empirical studies or data collection. Expanding the scope of analysis to other sources may lead to conceptual insights that are not represented in the literature. Second, it does not provide exhaustive coverage of the literature, only a representative selection of papers with relevance to VA activity and influence on the field. Because Visual Analytics is rapidly growing and diversifying, such works are biased towards the dominant concerns during the early development of the field. Our activity typology thus reflects security, intelligence, and investigation more so than the many subsequent domain specializations. Third, we only generate a single design tension and a single target quality per activity element, when more of each may be desirable and where design tensions between activity elements may also be possible. Fourth, our design tensions remain hypotheses pending empirical validation.

Expanding this analysis to address these limitations requires considerable effort (and space) that is beyond the scope of this paper. We do not suppose that we have the definitive activity typology for Visual Analytics, but simply a blueprint for the development of such a typology by the community at large (or indeed, multiple competing visions of what it should be). Such holistic reflection on the field may take years of work, and will need to evolve with the state of the art. We view our activity typology as just the first step in this process.

## 5 Conclusion

Activity Theory is well-established in the field of HCI and played a major role in shaping the second wave of evolution in HCI research. In this paper, we have reformulated AT for the analysis and design-led transformation of high-level, general activities, and applied the resulting design process to the activity of Visual Analytics.

At the core of our design approach is the concept of an activity typology—a structured organization of six activity elements and the types taken by each element across many forms of the activity. Activity typologies build upon and extend the taxonomic tradition in the InfoVis and VA literature, in ways that expand the focus "beyond tasks" to the broader systems of activity in which tools are embedded and used. We invite the community to refine our VA activity typology, to reapply our design process in specific domains, and to drive a new wave of activity-centered research and design for Visual Analytics.

## Acknowledgements

We thank James Landay and Sergio Paolantonio for their contributions to the activity-centered design process used in this work.

## References

[1] Robert Amar, James Eagan, and John Stasko. 2005. Low-level components of analytic activity in information visualization. IEEE Symposium on Information Visualization (INFOVIS) 2005,111-117.

[2] Kenneth D. Bailey. 1994. Typologies and taxonomies: an introduction to classification techniques. Sage.

[3] Liam Bannon. 1986. From human factors to human actors: the role of psychology and human-computer interaction studies in system design. In Greenbaum, J. & Kyng, M. (eds). Design at work: cooperative design of computer systems, Erlbaum, 25–44.
[4] Jakob Bardram. 1998. Designing for the Dynamics of Cooperative Work Activities. ACM Conference on Computer-Supported Cooperative Work (CSCW) 1998, 89-98.
[5] Jakob Bardram, Jonathan Bunde-Pedersen, and Mads Soegaard. 2006. Support for activity-based computing in a personal computing operating system. ACM Conference on Human Factors in Computing Systems (CHI) 2006, 211-220.
[6] Paulo Barthelmess and Kenneth M. Anderson. 2002. A View of Software Development Environments Based on Activity Theory. ACM Conference on Computer Supported Cooperative Work (CSCW) 2002, 13-37.
[7] Olav W. Bertelsen. 2004. The activity walkthrough: an expert review method based on activity theory. ACM Nordic Conference on Human-Computer Interaction (NordiCHI) 2004, 251-254.
[8] Jacques Bertin. 1977. Semiology of Graphics: Diagrams, Networks, Maps.
[9] Hugh Beyer and Karen Holtzblatt. 1997. Contextual Design: A Customer-Centered Approach to Systems Designs. Morgan Kaufmann Series in Interactive Technologies.
[10] Eric A. Bier, Stuart K. Card, and John W. Bodnar. 2008. Entity-based collaboration tools for intelligence analysis. IEEE Symposium on Visual Analytics Science and Technology (VAST) 2008, 99-106.
[11] Eric A. Bier, Lance Good, Kris Popat, and Alan Newberger. 2004. A document corpus browser for in-depth reading. ACM/IEEE Digital Libraries 2004, 87-96.
[12] Eric A. Bier, Edward W. Ishak, and Ed Chi. 2006. Entity Workspace: an evidence file that aids memory, inference, and reading. Intelligence and Security Informatics, 466-472.
[13] Alan F. Blackwell and Thomas R.G. Green. 2003. Notational systems – the Cognitive Dimensions of Notations framework. In J.M. Carroll (ed.) HCI Models, Theories and Frameworks: Toward a multidisciplinary science. San Francisco. Morgan Kaufmann, 103-134.
[14] Susanne Bødker. 2006. When second wave HCI meets third wave challenges. ACM Nordic Conference on Human-Computer Interaction (NordiCHI) 2006, 1-8.
[15] danah boyd and Kate Crawford. 2012. Critical questions for Big Data: Provocations for a cultural, technological, and scholarly phenomenon. Information, Communication & Society, 15(5), 662-679.
[16] Matthew Brehmer and Tamara Munzner. 2013. A multi-level typology of abstract visualization tasks. IEEE Transactions on Visualization and Computer Graphics, 19(12), 2376 – 2385.
[17] Susan E. Brennan, Klaus Mueller, Greg Zelinsky, IV Ramakrishnan, David S. Warren, and Arie Kaufman. 2006. Toward a multi-analyst, collaborative framework for visual analytics. IEEE Symposium on Visual Analytics Science and Technology (VAST) 2006, 129-136.
[18] Davide Ceneda, Theresia Gschwandtner, Thorsten May, Silvia Miksch, Hans-Jorg Schulz, Marc Streit, and Christian Tominski. 2016. Characterizing Guidance in Visual Analytics. IEEE Symposium on Visual Analytics Science and Technology (VAST) 2016, 111-120.
[19] Ed Chi. 2000. A taxonomy of visualization techniques using the data state reference model. IEEE Symposium on Information Visualization (INFOVIS) 2000, 69-75.
[20] Haeyong Chung, Seungwon Yang, Naveed Massjouni, Christopher Andrews, Rahul Kanna, and Chris North. 2010. VizCept: Supporting Synchronous Collaboration for Constructing Visualizations in Intelligence Analysis. IEEE Symposium on Visual Analytics Science and Technology (VAST) 2010, 107-114.
[21] Torkil Clemmensen, Bonnie Nardi, and Victor Kaptelinin. 2016. Making HCI Theory Work: An Analysis of the Use of Activity Theory in HCI Research. Behaviour & Information Technology, 35(8), 608-627.
[22] Geoff Cooper and John Bowers. 1995. Representing the user: Notes on the disciplinary rhetoric of human-computer interaction. Cambridge Series on Human Computer Interaction, 1995, 48–66.
[23] Dark data. Retrieved 6 Sept. 2016 from https://en.wikipedia.org/wiki/Dark_data
[24] Paul Dourish and Sara Bly. 1992. Portholes: supporting awareness in a distributed work group. ACM Conference on Human Factors in Computing Systems (CHI) 1992, 541-547.
[25] Ryan Eccles, Thomas Kapler, Robert Harper, and William Wright. 2007. Stories in GeoTime. IEEE Symposium on Visual Analytics Science and Technology (VAST) 2007, 19-26.
[26] Yrjö Engeström. 1987. Learning by expanding: An activity theoretical approach to developmental research. Helsinki.
[27] Yrjö Engeström. 2006. Activity theory and expansive design. Theories and practice of interaction design, 3-23.
[28] Geraldine Gay and Helene Hembrooke. 2003. Activity-Centered Design: An ecological approach to designing smart tools and usable systems. MIT Press.
[29] David Gotz, Michelle X. Zhou, and Vikram Aggarwal 2006. Interactive visual synthesis of analytic knowledge. IEEE Symposium on Visual Analytics Science and Technology (VAST) 2006, 51-58.
[30] David Gotz and Michelle X. Zhou. 2008. Characterizing Users' Visual Analytic Activity for Insight Provenance. IEEE Symposium on Visual Analytics Science and Technology (VAST) 2008, 123-130.
[31] Samuel Gratzl, Alexander Lex, Nils Gehlenborg, Nicola Cosgrove, and Marc Streit. 2016. From Visual Exploration to Storytelling and Back Again. Computer Graphics Forum (EUROVIS), 2016, 35(3), 491-500.
[32] Tera Marie Green, Ron Wakkary, and Richard Arias-Hernandez. 2011. Expanding the scope: Interaction design perspectives for visual analytics. IEEE Hawaii International Conference on System Sciences (HICSS).
[33] John Halloran, Yvonne Rogers, and Mike Scaife. 2002. Taking the 'No' out of Lotus Notes: Activity Theory, Groupware and student work projects. International Society of the Learning Sciences Conference on Computer Support for Collaborative Learning (CSCL) 2002, 169–178.
[34] Christine A. Halverson. 2002. Activity theory and distributed cognition: Or what does CSCW need to DO with theories?. ACM Computer Supported Cooperative Work (CSCW) 2002, 243–267.
[35] Jeffrey Heer and Maneesh Agrawala. 2008. Design considerations for collaborative visual analytics. Information Visualization, 7(1), 49-62.
[36] Jeffrey Heer, Fernanda B. Viégas, and Martin Wattenberg. 2007. Voyagers and voyeurs: supporting asynchronous collaborative information visualization. ACM Conference on Human Factors in Computing Systems (CHI) 2007, 1029-1038.
[37] Elizabeth Hetzler and Alan Turner. 2004. Analysis experiences using information visualization. IEEE Transactions on Computer Graphics and Applications (TVCG) 24(5), 22-26.
[38] Richard J. Heuer. 1999. Psychology of Intelligence Analysis. Center for the Study of Intelligence.
[39] Edwin Hutchins. 1995. Cognition in the Wild. MIT press.
[40] David Jonker, William Wright, David Schroh, Pascale Proulx, and Brian Cort. 2005. Information triage with TRIST. In 2005 International Conference on Intelligence Analysis, 2–4.
[41] Youn-ah Kang and John Stasko. 2011. Characterizing the intelligence analysis process: Informing visual analytics design through a longitudinal field study. IEEE Symposium on Visual Analytics Science and Technology (VAST) 2011, 19-28.
[42] Thomas Kapler and William Wright. 2004. GeoTime Information Visualization. IEEE Symposium on Information Visualization (INFOVIS) 2004, 25-32.
[43] Victor Kaptelinin, Bonnie A. Nardi, and Catriona Macaulay. 1999. The Activity Checklist: a tool for representing the "space" of context. ACM Interactions, 6(4), 27–39.
[44] Victor Kaptelinin, Bonnie A. Nardi. 2006. Acting with Technology: Activity Theory and Interaction Design. MIT Press.
[45] KyungTae Kim, Waqas Javed, Cary Williams, Niklas Elmqvist, Pourang Irani. 2010. Hugin: A Framework for Awareness and Coordination in Mixed-Presence Collaborative Information Visualization. ACM Conference on Interactive Tabletops and Surfaces (ITS) 2010, 231-240.
[46] Gary Klein, Brian Moon and Robert R. Hoffman 2006. Making sense of sensemaking 2: A macrocognitive model. IEEE Intelligent Systems, 21(5), 88-92.
[47] Aleksei Leontiev. 1974. The Problem of Activity in Psychology. Soviet Psychology 13, 4-33.

[48] Aleksei Leontiev. 1978. Activity, Consciousness, and Personality. Prentice Hall.
[49] Aleksei Leontiev. 1981. Problems in the Development of the Mind. Progress Publishers.
[50] Heather Richter Lipford, Felesia Stukes, Wenwen Dou, and Remco Chang. 2010. Helping Users Recall Their Reasoning Process. IEEE Symposium on Visual Analytics Science and Technology (VAST) 2010, 187-194.
[51] Narges Mahyar and Melanie Tory. 2014. Supporting communication and coordination in collaborative sensemaking. IEEE Transactions on Visualization and Computer Graphics (TVCG), 20(12), 1633-1642.
[52] Microsoft Power BI custom visuals, posted 9 May 2016. Retrieved 6 Sept. 2016 from https://powerbi.microsoft.com/en-us/blog/custom-visualizations-visual-awesomeness-your-way/
[53] Microsoft Power BI visuals for text analysis, posted 11 July 2016. Retrieved 6 Sept. 2016 from https://powerbi.microsoft.com/en-us/blog/new-power-bi-custom-visuals-for-browsing-and-analyzing-collections-of-text/
[54] Microsoft Power BI Twitter solution template, posted 15 August 2016. Retrieved 6 Sept. 2016 from https://powerbi.microsoft. com/en-us/blog/twitter-solution-template/
[55] Bonnie A. Nardi. 1996. Context and consciousness: activity theory and human-computer interaction. MIT.
[56] Phong H. Nguyen, Kai Xu, Ashley Wheat, B.L. William Wong, Simon Attfield, and Bob Fields. 2016. SensePath: Understanding the Sensemaking Process through Analytic Provenance. IEEE Transactions on Visualization and Computer Graphics (TVCG), 22(1), 41-50.
[57] Donald A. Norman. 2005. Human-centered design considered harmful. ACM Interactions, 12(4), 14-19.
[58] William Pike, Joe Bruce, Bob Baddeley, Daniel Best, Lyndsey Franklin, Richard May, Douglas Rice, Rick Riensche and Katarina Younkin. 2009. The scalable reasoning system: lightweight visualization for distributed analytics. Information Visualization, 8(1), 71-84.
[59] Peter Pirolli and Stuart Card (2005). The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. International Conference on Intelligence Analysis, 5, 2-4.
[60] Eric D. Ragan, Alex Endert, Jibonananda Sanyal and Jian Chen. 2016. Characterizing Provenance in Visualization and Data Analysis: An Organizational Framework of Provenance Types and Purposes. IEEE Transactions on Visualization and Computer Graphics (TVCG), 22(1), 31-40.
[61] Antonio Sanfilippo, Bob Baddeley, Andrew J. Cowell, Michelle L. Gregory, Ryan Hohimer, and Stephen Tratz. 2005. Building a human information discourse interface to uncover scenario content. International Conference on Intelligence Analysis 2005.
[62] Antonio Sanfilippo, Bob Baddeley, Christian Posse, and Paul Whitney. 2007. A Layered Dempster-Shafer Approach to Scenario Construction and Analysis. IEEE Intelligence and Security Informatics 2007, 95-102.
[63] Douglas Schuler and Aki Namioka (eds). 1993. Participatory design: Principles and practices. CRC Press.
[64] Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori M. Inkpen. 2004. Territoriality in collaborative tabletop workspaces. ACM Conference on Computer Supported Cooperative Work (CSCW) 2004, 294-303.
[65] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. 2012. Design study methodology: Reflections from the trenches and the stacks. IEEE Transactions on Visualization and Computer Graphics (TVCG), 18(12), 2431-2440.
[66] Yedendra Babu Shrinivasan and Jarke J. van Wijk. 2008. Supporting the analytical reasoning process in information visualization. ACM Conference on Human Factors in Computing Systems (CHI), 1237-1246.
[67] Meredith Skeels, Bongshin Lee, Greg Smith, and George G. Robertson. 2010. Revealing uncertainty for information visualization. Information Visualization, 9(1), 70-81.
[68] John Stasko, Carsten Görg, and Zhicheng Liu. 2008. Jigsaw: supporting investigative analysis through interactive visualization. Information Visualization, 7(1), 118-132.
[69] Lucy A. Suchman. 1987. Plans and situated actions: The problem of human-machine communication. Cambridge University Press.
[70] Deborah Tatar. 2007. The Design Tensions Framework. Human-Computer Interaction, 22(4), 413-151.
[71] James J. Thomas and Kristin A. Cook (eds). 2005. Illuminating the Path: The Research and Development Agenda for Visual Analytics. IEEE Computer Society Press.
[72] Matthew Tobiasz, Petra Isenberg, and Sheelagh Carpendale. 2009. Lark: Coordinating Co-located Collaboration with Information Visualization. IEEE Transactions on Visualization and Computer Graphics (TVCG), 15(6), 1065–1072.
[73] Fernanda B. Viegas, Martin Wattenberg, Frank Van Ham, Jesse Kriss, and Matt McKeon. 2007. Many Eyes: A Site for Visualization at Internet Scale. IEEE Transactions on Visualization and Computer Graphics (TVCG), 13(6), 1121–1128.
[74] VIS 2011 panel. Theories of Visualization – Are There Any?
[75] Colin Ware, William Wright, and Nicholas J. Pioch. 2013. Visual Thinking Design Patterns. International Conference on Distributed Multimedia Systems (DMS) 2013, 150-155.
[76] William Wright, David Schroh, Pascale Proulx, Alex Skaburskis, and Brian Cort. 2006. The Sandbox for analysis: concepts and methods. ACM CHI 2006, 801–810.
[77] Jian Zhao, Michael Glueck, Simon Beslav, Fanny Chevalier, Azam Khan. 2017. Annotation Graphs: A Graph-Based Visualization for Meta-Analysis of Data based on User-Authored Annotations. IEEE Transactions on Visualization and Computer Graphics (TVCG), 23(1), 261-270.