INFORMING DESIGN OF VISUAL ANALYTICS SYSTEMS FOR INTELLIGENCE ANALYSIS: UNDERSTANDING USERS, USER TASKS, AND TOOL USAGE

A Dissertation
Presented to
The Academic Faculty

by

Youn Ah Kang

In Partial Fulfillment
of the Requirements for the Degree
Human-Centered Computing in the
School of Interactive Computing

Georgia Institute of Technology
August 2012

INFORMING DESIGN OF VISUAL ANALYTICS SYSTEMS FOR INTELLIGENCE ANALYSIS: UNDERSTANDING USERS, USER TASKS, AND TOOL USAGE

Approved by:

Dr. John Stasko, Advisor
School of Interactive Computing
Georgia Institute of Technology

Dr. Jim Foley
School of Interactive Computing
Georgia Institute of Technology

Dr. Beki Grinter
School of Interactive Computing
Georgia Institute of Technology

Dr. Jean Scholtz
Pacific Northwest National Laboratory

Dr. Chris North
Department of Computer Science
Virginia Tech

Date Approved: 06 / 29 /2012

／# ACKNOWLEDGEMENTS

I am heartily thankful to my advisor, John Stasko, for persevering with me throughout my PhD study. His patience, encouragement, and guidance enabled me to complete this research and write this dissertation. I could not have imagined having a better advisor and mentor for my PhD study.

Besides my advisor, I would like to thank the rest of my committee, Jim Foley, Beki Grinter, Jean Scholtz, and Christ North, for their encouragement and insightful comments. They have generously given their time and expertise to improve my work.

I thank my fellow lab mates in Information Interfaces Group at Georgia Tech. In particular, I am grateful Zhicheng Liu for his support, who has been my colleague and friend for the last 6 years, and will stay on to be.

My sincere thanks must also go to those who shared their memories and experiences with me in Atlanta, especially the Sugarloaf Korean Baptist Church family. Their friendship and encouragement sustained me throughout my PhD study.

Last but not the least, I would like to thank my husband Giwan for his personal support and great patience at all times. I would also like to thank my parents and brother, for the infinite trust they have in me and support along all my life. Thanks for filling my life with expectations.

As high as the heavens are above the earth

As far as the waters cover all the seas

So high, so great is the love You have for me

As far as the east is from the west

You've taken away my worthlessness

You've turned my shame into joyful praise

Forever I will thank You for the grace You've given me

I once was lost, but now I'm found, You're everything I need

I worship You, I live for You, You're all that I desire

And forever I will live to tell of Your unfailing love

"For God so loved the world that He gave His one and only Son,

that whoever believes in Him shall not perish but have eternal life." - John 3:16

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	iii
LIST OF TABLES	ix
LIST OF FIGURES	x
SUMMARY	xii
<u>CHAPTER</u>	
1 Introduction	1
1.1 Problem and Motivation	2
1.2 Purpose and Research Questions	5
1.3 Research Method	8
1.4 Research Contributions	8
1.5 Organization	9
2 Related Work	11
2.1 Models of the Intelligence Analysis Process	11
2.2 Models of Sensemaking	17
2.2.1. Sensemaking and Intelligence Analysis	17
2.2.2. Models of Sensemaking	18
2.3 Systems for Investigative Analysis and Sensemaking	24
2.4 Understanding Users through Formal User Studies in Visual Analytics	30
3 Characterizing the Intelligence Analysis Process: Informing Visual Analytics Design through a Longitudinal Field Study	33
3.1 Introduction	33
3.2 Methods and Procedures	34

3.2.1 Participants 34

3.2.2 Task 35

3.2.3 Study Protocol and Procedures 37

3.2.4 Data Collection and Analysis 38

3.3 Overall Analysis Process 38

3.3.1 Intuitive Analysis – Team A 41

3.3.2 Structured Analysis – Teams B and C 41

3.4 Tools and Methods Used 43

3.5 Understanding the Intelligence Analysis Process 44

3.6 Rethinking the Intelligence Analysis Process 47

3.6.1 Linear vs. Parallel 47

3.6.2 Pirolli and Card's Sensemaking Model 49

3.7 Where and How Collaboration Occurred 50

3.7.1 Collaboration throughout the Process 50

3.7.2 Sharing vs. Content vs. Function 52

3.8 How VA Can Help: Design Implications 53

4 Evaluating Benefits of VA Tools in an Investigative Scenario: Deriving Design Implications from a Comparative Study 62

4.1 Introduction 62

4.2 Study Design and Analysis Techniques 63

4.3 Four Investigative Strategies 65

4.3.1 Strategy 1: Overview, Filter, and Detail (OFD) 66

4.3.2 Strategy 2: Build from Detail (BFD) 67

4.3.3 Strategy 3: Hit the Keyword (HTK) 67

4.3.4 Strategy 4: Find a Clue, Follow the Trail (FCFT) 67

4.4 Jigsaw's Influence on Investigative Analysis 68

4.5 Observations on Sensemaking 69
4.5.1 Diversity in Sensemaking Processes 69
4.5.2 Power of Schematizing 70
4.5.3 Insight Acquisition 70
4.6 Implications for Design of Investigative Analysis Tools 71
4.7 Implications for Evaluation of Investigative Analysis Tools 72
5 Examining the Use of Visual Analytics Systems for Sensemaking Tasks: Case Studies with Domain Experts 74
5.1 Introduction 74
5.2 Recruitment and Study Protocol 75
5.3 Case Studies 77
5.3.1 P1: Aerospace Engineering Researcher 77
5.3.2 P2: Business Analyst at Management Services 82
5.3.3 P3: PhD Candidate in Industrial and Systems Engineering 83
5.3.4 P4: Intelligence Analyst at a Police Department 86
5.3.5 P5: Intelligence Analyst at a National Lab 88
5.3.6 P6: Intelligence Analyst at Air Force 90
5.4 Findings and Discussion – How They Used Jigsaw 93
5.4.1 Types of tasks 93
5.4.2 Learning the system 94
5.4.3 Unexpected use of tool 95
5.4.4 Issues and problems 97
5.5 Design Implications 100
6 Discussion 106
6.1 Reflecting on Research Questions 106
6.2 Revisiting Design Implications 108

6.2.1 Commonalities and Uniqueness 109

6.2.2 Existing Work on Design/Evaluation Implications 115

6.3 Reconsidering Sensemaking Models 117

6.4 Limitations of the Study 119

7 Conclusion 123

REFERENCES 126

LIST OF TABLES

Page

Table 1: Summary of research questions and studies 10

Table 2: Study results and statistics, grouped by setting 66

Table 3: Example of a mapping between NextGen and SESAR improvements 78

Table 4: Modified document: description and entities associated with the NextGen Operational Improvement OI-0320 79

LIST OF FIGURES

Page

Figure 1: Traditional intelligence cycle 12

Figure 2: Process of intelligence creation and use 13

Figure 3: Treverton's real intelligence cycle 14

Figure 4: Mark Lowenthal's model of the intelligence process 14

Figure 5: Systems model of the intelligence cycle 15

Figure 6: Wheaton's multi-phasic model of the intelligence process 16

Figure 7: Learning loop complex theory of sensemaking 20

Figure 8: The data/frame theory of sensemaking 21

Figure 9: Notional model of sensemaking loop for intelligence analysis 22

Figure 10: Iterative sensemaking model 23

Figure 11: Views in Jigsaw 25

Figure 12: An emerging analysis using the Sandbox interface 26

Figure 13: Link analysis using Analyst's Notebook 27

Figure 14: Documents, highlighting, entities, and spatial relationships in Analyst's Workspace 28

Figure 15: Compendium for literature analysis 29

Figure 16: Conceptual model. Printed from Mindmeister 39

Figure 17: Case study matrix of crimes 42

Figure 18: Wheaton's multi-phasic model of the intelligence process 48

Figure 19: Original document: Operational Improvements-0320 in NextGen initiative 78

Figure 20: A document and its relevant entities in the Graph View 80

Figure 21: The relationships between NextGen OI-0320 and relevant SESAR Operational Improvements 81

Figure 22: List of words created for “Business Expansion” entity 84

Figure 23: List View displaying event types, companies, and dates related to “strategy”85

Figure 24: Sample document of a budget summary on Air Traffic Control by Army 90

Figure 25: Critical areas identified by the circular layout in the Graph View 92

Figure 26: A mapping created manually (top) and by Jigsaw (bottom) 96

SUMMARY

Visual analytics, defined as "the science of analytical reasoning facilitated by interactive visual interfaces," emerged several years ago as a new research field. While it has seen rapid growth for its first five years of existence, the main focus of visual analytics research has been on developing new techniques and systems rather than identifying how people conduct analysis and how visual analytics tools can help the process and the product of sensemaking. The intelligence analysis community in particular has not been fully examined in visual analytics research even though intelligence analysts are one of the major target users for which visual analytics systems are built. The lack of understanding about how analysts work and how they can benefit from visual analytics systems has created a gap between tools being developed and real world practices.

This dissertation is motivated by the observation that existing models of sensemaking/intelligence analysis do not adequately characterize the analysis process and that many visual analytics tools do not truly meet user needs and are not being used effectively by intelligence analysts. I argue that visual analytics research needs to adopt successful HCI practices to better support user tasks and add utility to current work practices. As the first step, my research aims (1) to understand work processes and practices of intelligence analysts and (2) to evaluate a visual analytics system in order to identify where and how visual analytics tools can assist. By characterizing the analysis process and identifying leverage points for future visual analytics tools through empirical

studies, I suggest a set of design guidelines and implications that can be used for both designing and evaluating future visual analytics systems.

CHAPTER 1

INTRODUCTION

Visual analytics, which emerged several years ago, is a research field that combines knowledge from multiple disciplines [1,81] such as Data Mining, Databases, Cognitive Science, Information Visualization, Knowledge Management, and Decision Science. The basic idea is the integration of the outstanding capabilities of humans in terms of visual information exploration and the enormous processing power of computers in order to form a powerful knowledge discovery environment. Many consider the formal beginning of visual analytics as a field to be the publication of Illuminating the Path: The R&D Agenda for Visual Analytics [81]. In the book, multidisciplinary researchers defined visual analytics as "the science of analytical reasoning facilitated by interactive visual interfaces" and identified major challenges and a research agenda.

In its first five years of existence, visual analytics has seen rapid growth and great progress has been made in a short time, yet research challenges remain. Visual analytics is a relatively new research field, and the community has focused mostly on developing new tools and techniques. While a number of visual analytics systems have been built so far, few empirical studies have examined how such tools can help the process and the product of sensemaking [6, 39, 40, 56]. More fundamentally, the field has not yet accumulated sufficient knowledge about users and their work practices. The lack of research studies that yield critical, direct implications for the design and evaluation of visual analytics systems is a challenge that should be addressed for the further growth of the field.

One approach to address the challenge is to adopt practices that have proved successful in the human computer interaction (HCI) field. While HCI research includes developing innovative user interfaces and systems, much HCI work has put emphasis on understanding users, their behaviors and perceptions, their work practices, and evaluating user interaction with tools. Design implications from those studies have positively influenced the development and refinement of information systems in turn, increasing the adoption and use of interfaces and systems in the real world.

While visual analytics and HCI are two different fields, they do share some characteristics in common in that considering user needs and user interaction with tools is critical. In order to provide new visual analytics technologies that will be useful and widely accepted throughout the user community, I argue that visual analytics research needs to focus on enumerating design guidelines and implications by understanding the users, their work processes, and tool usage.

1. Problem and Motivation

While visual analytics tools are targeted for a certain community, mostly analysts, little research that seeks to understand their tasks and work practices has been conducted. Until recently, the main focus of visual analytics research has been on developing new techniques and systems rather than identifying how people work and what characteristics we should support. Although some researchers try to identify user tasks and design requirements through interaction with target users before implementing a system, visual analytics research still lacks a fundamental understanding of what kind of tasks analysts do, how they work, what unique characteristics they exhibit, and what they want from visual analytics tools. Without understanding those aspects deeply enough, however, it is difficult to build a system that truly helps analysts in their tasks.

The intelligence analysis community in particular has not been fully examined in the visual analytics research. While intelligence analysts are one of the major target users for which visual analytics systems are built, little research starts with understanding them and we still know relatively little about their work processes and practices. Unfortunately, it is not easy to conduct such studies due to limited time, effort, and availability of resources. Intelligence analysts are not easily accessible, and even when they are available, it is difficult to contact and study them for a long period of time. For these reasons, researchers often develop visual analytics systems based on existing models and frameworks about sensemaking and/or intelligence analysis, instead of conducting their own user studies.

One of the most widely used models in the visual analytics community is Pirolli and Card's sensemaking model [60] for intelligence analysis. While the model broadly characterizes processes used in analysis activities and has guided design processes of visual analytics tools, the model was developed based on findings from a cognitive task analysis and interviews independently from a real work context. Consequently, the model still does not provide rich details of how intelligence analysts work in the real world. For example, the model does not explain how analysts collaborate throughout their work cycle and how each phase can be omitted/modified depending on different task types. More empirical, descriptive models of the intelligence analysis process are required to better understand the sensemaking process. Another question is whether it has been validated that the model truly fits in intelligence analysis processes. No studies have proved the adequacy of the model in describing the intelligence analysis process, but researchers simply presume it as "the" model. However, a single model cannot always capture intelligence analysis processes, and we need more empirical findings and implications about intelligence analysts' work processes and practices.

Little attention about users and their work practices is also related to research in the evaluation of visual analytics systems. Although many systems have been developed to facilitate analysis, little research has examined the potential benefits of such systems in practice and how analysts are using them. One obvious reason is the difficulty of evaluating the utility of visual analytics systems [61]. Assessing whether a visual analytics system adds real value to users' current work process is challenging because the nature of work is exploratory and the quality of outcome is difficult to measure.

Nevertheless, investigating how such systems foster insight and sensemaking is important for continued growth in the field because those systems will greatly change the current processes of the end users [71].

The lack of understanding about how analysts work and how they can benefit from visual analytics systems created a gap between tools being developed and real world practices. This gap makes it difficult to know whether we are on the right track and building the right tools for analysts. A discussion with intelligence analysts revealed that they feel that researchers who build tools for intelligence analysis do not well understand how analysts work, and that such tools often do not fit in their analysis cycle. This incongruity seemed to be one of the reasons that intelligence analysts do not significantly benefit from visual analytics tools.

For better use and appropriation of tools, the visual analytics community needs to adopt successful HCI practices – understanding users and their tasks to derive design implications, integrating user requirements into a design process, and evaluating the adoption and usage of tools to further refine them and find more leverage points. Before building a system, developers need to understand what users need and how they could benefit from a tool through an exploratory study phase. Once they acquire user requirements and design implications, they can proceed based on the findings, building a system that helps the users more effectively. Then the next step would be to evaluate the potential benefit of the tool through a scientific study, validating the utility of the tool and further identifying room for improvement.

Undertaking all these processes for every single system could be daunting. Short of that, design heuristics or guidelines that can be used for tool development could save time and effort. Thus, research studies investigating work processes and practices of intelligence analysts and suggesting design implications will be extremely beneficial to researchers who build visual analytics systems. While those design guidelines might not be always applicable, simply enumerating lessons and implications through multiple studies will help researchers gain a better sense of who the users are, how they work, and what things to consider when developing tools.

2. Purpose and Research Questions

As I suggested previously, I believe that visual analytics research needs to accept good HCI practices to better support user tasks and add utility to their current work practices. As the first step, my research aimed to understand work processes and practices of intelligence analysts from a broader point of view and to identify where and how visual analytics tools can assist their tasks. By identifying unique characteristics of intelligence analysis practices, pain points with current tools, and leverage points for future visual analytics tools, I sought to suggest a set of design guidelines and implications that can be used for both designing and evaluating future visual analytics systems. My hope is that those guidelines and implications will benefit the research community by helping researchers make appropriate design choices before investing too much development effort. And more importantly, developers can build the "right" tool that truly meets user needs and smoothly integrates with their current practices.

While intelligence analysis involves various information types such as text, images, video, and numbers, most analyses involve some kind of investigation with text documents. Intelligence analysts often encounter long and complicated documents such as intelligence reports, news articles, and research papers. They need to collect, process, evaluate, understand, interpret, and integrate documents into a new form of knowledge to make an actionable decision. In my dissertation, I intentionally limited the scope of my studies to such textual document analysis.

With this purpose in mind, the overarching thesis statement of this dissertation is therefore:

Current visual analytics tools do not sufficiently support intelligence analysts throughout their work process. By examining how analysts work on projects and *how they exploit technological aids during analysis, we can better understand the intelligence process model and identify how visual analytics systems can help. By studying how investigators use an existing visual analytics tool for analytical tasks, we can better understand what characteristics of the tool add benefit or not, deriving design implications for future visual analytics systems.*

The first part of this dissertation thus explores current work processes and practices of intelligence analysts as they are. Although a few studies examined how analysts work [13,64], there is a large space to explore their work process and practices. It seems that the visual analytics community still does not have enough understanding of the intelligence analysis domain and its needs and that existing models do not accurately characterize the analysis process. Consequently, current research in visual analytics tools tends to focus on a part(s) of the analysis process, rather than the entire cycle of intelligence analysis. Particularly, I sought to address two research questions:

RQ1. Do current models used by developers of visual analytics tools adequately characterize the process of intelligence analysis? What aspects of intelligence analysis are particularly misunderstood?

RQ2. Where in the analysis process and for what kind of tasks can visual analytics tools best benefit intelligence analysts without intruding on their work practices?

To answer these questions, I conducted a field study that consisted of a series of in-depth interviews and observations with intelligence analysts, which are described more in detail in the following section.

Secondly, given an available visual analytics system for document analysis, I wanted to examine how such a system could benefit analysis processes and products.

RQ3. How do existing visual analytics systems such as Jigsaw support or fail to support investigative analysis?

Through this research question, I sought to observe people using a visual analytics system and to gain insights for analytical processes and derive design implications for investigative analysis tools. To address this question, I conducted two different evaluation studies. In the first, I compared within a laboratory setting the usage of a visual analytics system to more traditional methods of analysis. The study demonstrated how a visual analytics system adds analytical benefits and helps people perform an investigative analysis. Along with the laboratory evaluation study, I further examined the use of the tool with domain experts who had been using it with their own data in real world settings. The importance of case studies for information visualization [69, 75] has been emphasized, and I believe that this study yielded interesting findings and implications from a different perspective, complementing the findings from the laboratory study.

By answering all the research questions listed above, my research provides a detailed view of the analysis process and design implications for next-generation visual analytics systems for document analysis, thereby addressing RQ4.

RQ4: What design implications for visual analytics systems for intelligence analysis emerge from the studies of the analysis process and the use of a visual analytics system?

3. Research Method

To understand users and obtain design guidelines for intelligence analysis tools, I took two approaches in conducting research studies. The first approach is to observe intelligence analysts and their current work practices regardless of use of any specific systems. By conducting a longitudinal case study with students majoring in Intelligence Studies, I characterized the intelligence analysis process, discussed several misunderstandings about the intelligence analysis process, and identified leverage points and design implications for intelligence analysis tools.

Secondly, I examined how people perform an analysis using a visual analytics tool and identified what kind of features and characteristics such tools need to support. The main focus of this approach is to refine an existing tool and also derive design implications for future systems. Two studies are pertinent–one comparative lab experiment and one case study. The comparative lab study has been conducted with Jigsaw, a visual analytics system for document analysis, and examined how people would perform an investigative analysis with Jigsaw and how it could provide benefit compared to other traditional methods. The second study has been conducted with six professional from a variety domains who have used Jigsaw on their own data. The study identified real-world cases of how an interactive visual system for investigative analysis assisted document sensemaking in various domains and tasks. It also discussed issues and findings that emerged upon the use of the visual analytics system.

Ultimately, from the results and the findings from a series of studies, I provide a better understanding of users and suggest design implications for future visual analytics systems development.

4. Research Contributions

Understanding users and their requirements are essential parts in designing systems. Especially for visual analytics which often has specified target users, the importance of satisfying user needs and integrating with their current work practice is even more amplified. I hope that my research provides useful implications for researchers to design and evaluate their systems for investigative analysis. The expected contributions of this research include:

- Provides a deeper understanding of intelligence analysts' processes, practices, and tool usage
- Provides empirical knowledge of how an interactive visual system can assist document sensemaking
- Identifies design requirements and suggests implications for future visual analytics tools for investigative analysis

5. Organization

The remainder of this document is organized as follows. In Chapter 2, I discuss related work focusing on intelligence cycles, sensemaking models, current visual analytics tools and systems, and user studies in the visual analytics field. In Chapter 3, I describe a field study in which I investigate analysts' work processes and practices, pain points with current tools, and design implications for visual analytics systems for intelligence analysis. Chapter 4 presents an evaluation study examining how an interactive visual interface can benefit sensemaking on text documents and what evaluation methodologies we can use for the utility evaluation. In Chapter 5, I present a case study of a visual analytics tool with experts from different domains. Chapter 6 discusses all the implications from the three studies, highlighting mutually reinforcing principles, as well as limitations involved in the research. Finally I provide a summary and conclusion in Chapter 7.

Table 1. Summary of research questions and studies

Research Questions	How Addressed
RQ1. Do current models used by developers of visual analytics tools adequately characterize the process of intelligence analysis? What aspects of intelligence analysis are particularly misunderstood?	A field study of the intelligence analysis process (Chapter 3 | VAST'11, submitted to IVS)
RQ2. Where in the analysis process and for what kind of tasks can visual analytics tools best benefit intelligence analysts without intruding on their work practices?	Qualitative interviews of intelligence analysts on their work processes and pain points with current tools (Chapter 3 | VAST'11, submitted to IVS)
RQ3. How do existing visual analytics systems such as Jigsaw support or fail to support investigative analysis?	Comparative lab study of people performing an investigative analysis task (Chapter 4 | VAST'09, TVCG17(5)) Cast study of domain experts using Jigsaw (Chapter 5 | VAST'12)
RQ4: What design guidelines and evaluation implications for visual analytics systems for intelligence analysis emerge from the studies of the analysis process and the use of a visual analytics system?	Analysis and summarization of previous study results

CHAPTER 2

RELATED WORK

In this chapter, I discuss related work focusing on intelligence cycle models in Intelligence Studies, sensemaking models in HCI, visual analytics systems developed for intelligence sensemaking, and studies that seek to understand users and their usage of visual analytics tools.

2.1 Models of the Intelligence Analysis Process

Researchers and practitioners have put effort to understand how intelligence analysts work and disentangle the intelligence process. Many of them have attempted to resolve problems identified with the traditional intelligence cycle [9]. The intelligence cycle is normally illustrated as a repeating process consisting of five steps, starting from the planning and direction stage. The notion of the cycle assumes that the steps will proceed in the prescribed order and that the process will repeat itself continuously with reliable results. This type of representation gives the impression that all inputs are constant and flow automatically. Although the traditional intelligence cycle has been used widely, the model is criticized by intelligence professionals in that it does not accurately represent the way intelligence is produced.

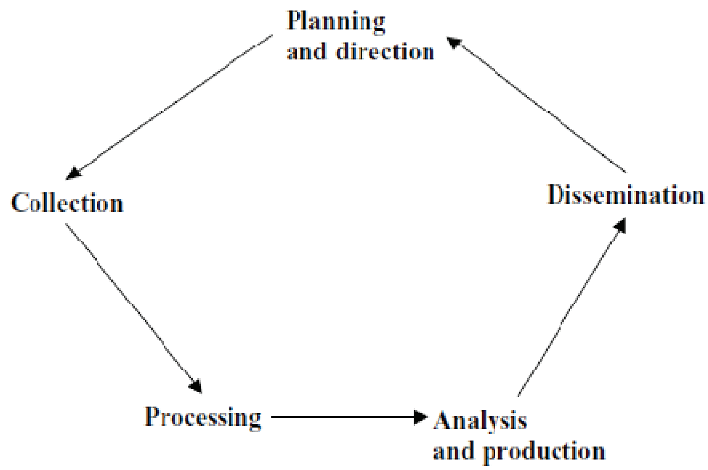

Figure 1. Traditional intelligence cycle, taken from [9]

Krizan [43], in *Intelligence Essentials for Everyone,* provides a slightly revised version of the intelligence cycle (Figure 2), which contains several component functions distinguished from the complex and dynamic cycle. In this cycle, components are identified as Intelligence Needs, Collection Activities, Processing of Collected Information, Analysis and Production. Quoting Douglas Dearth, she states “These labels, and the illustration below, should not be interpreted to mean that intelligence is a uni-dimensional and unidirectional process. In fact, ‘the process is multidimensional, multi-directional, and - most importantly - interactive and iterative [20].’”

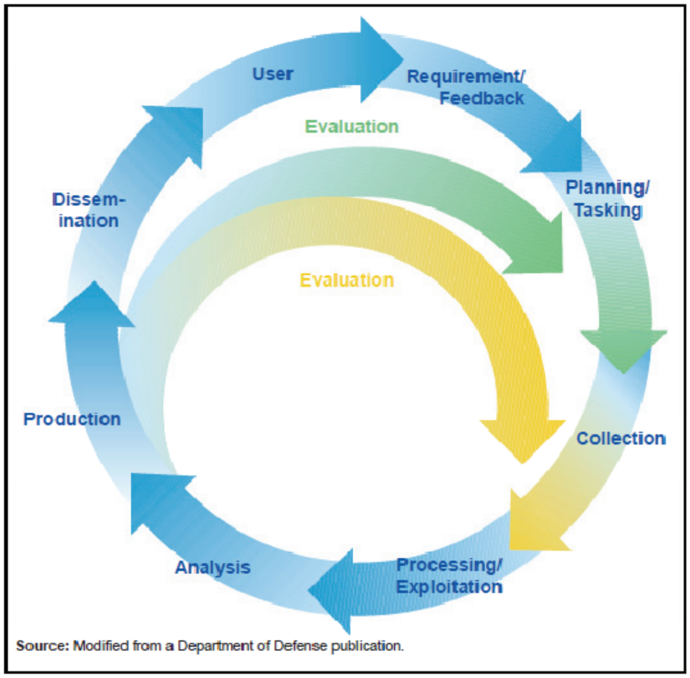

Figure 2. Process of intelligence creation and use, take from [43]

Still, this model follows the traditional depiction of intelligence cycle, which implies a sequential process, and does not provide for iterations between steps. In reality, however, there is repeated refinement in the collection and production steps. A more accurate picture of the steps in the process and their iterative tendencies may be seen in Gregory Treverton's model [83], which he terms the "Real" Intelligence Cycle (Figure 3).

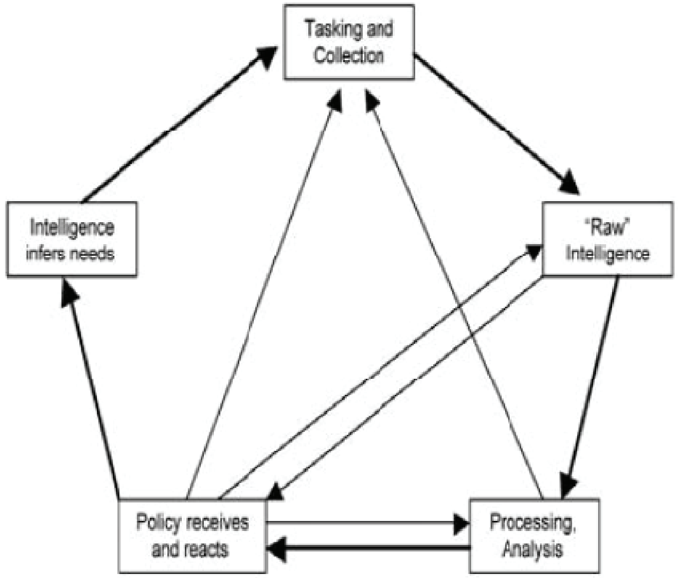

Figure 3. Treverton's real intelligence cycle, taken from [83]

Mark Lowenthal's model [45], although presented in a more linear fashion than Treverton's, focuses on the areas where revisions and reconsiderations take place, representing iteration in a slightly different light. Both models provide a more realistic view of the entire process.

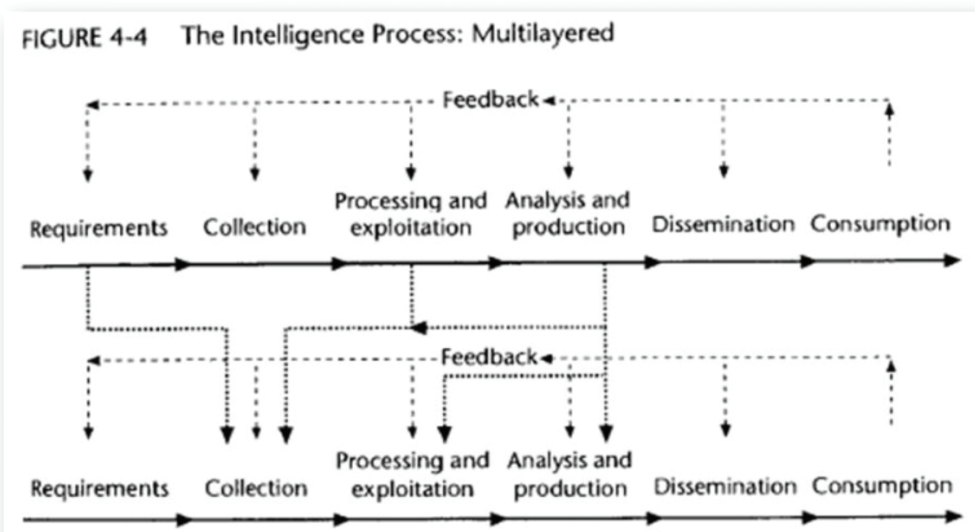

Figure 4. Mark Lowenthal's model of the intelligence process, taken from [45]

While all these models were suggested by an intelligence researcher, Rob Johnston, an anthropologist, conducted an ethnographic study of the CIA for a year and proposed a dramatically revised intelligence cycle from a systems perspective [33]. This model (Figure 5) provides a detailed representation of the process using four icons to represent actions and relationships within the system: stocks (accumulations), flows (activities), converters that change inputs to outputs, and connectors that link elements to other elements. Rather than replacing the traditional intelligence cycle, his model seeks to describe it more accurately.

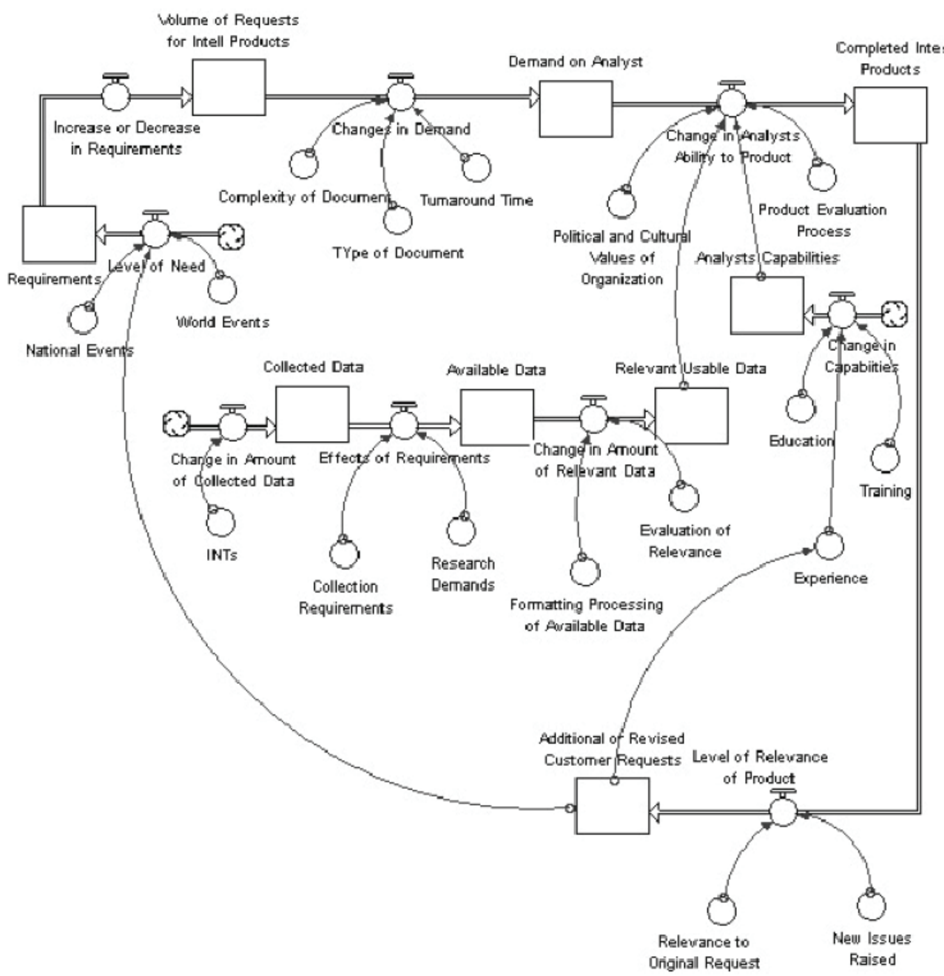

Figure 5. Systems model of the intelligence cycle, taken from [33]

Recently in his book, Wheaton [85] claims that the traditional intelligence cycle does not clearly and accurately describe what people do in intelligence analysis and even none of the alternatives proposed has yet captured the nuance of the process as practiced. Emphasizing the disconnect between theory and real-world practice, he proposes a new model that is fundamentally different way of thinking about the intelligence process. While it does not ignore the core elements of all intelligence activities such as collection and analysis, it abandons the linearity of previous models and acknowledges what intelligence professional actually do: they work on everything simultaneously. The notion of parallelism is very important in his model, in which he explains the process as a series of phases. In each phase, one of the core functions (modeling, collection, analysis, and production) would be mostly emphasized but all other functions would be operating in parallel. He argues that "All four functions begin almost immediately, but through the course of the project, the amount of time spent focused on each function will change, which each function dominating the overall process at some point."

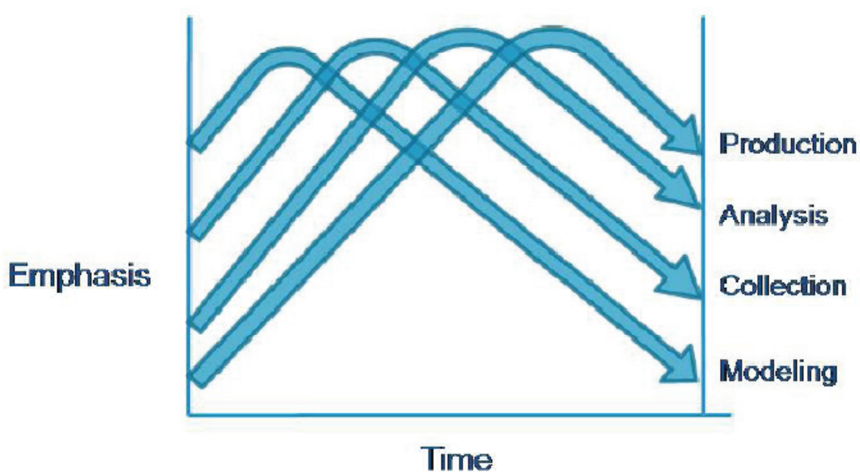

Figure 6. Wheaton's multi-phasic model of the intelligence process, taken from [85]

While all these models are hypotheses, or guesses about how intelligence process works, they are developed to describe what is occurring within a specific community (i.e., the US national security community), rather than to find leverage points for designing systems. These models by themselves do not provide a sufficient basis for developing technological support to analysis or what the end-to-end experience of analysis is like for the analyst. Considerable additional detail is required to communicate any real understanding of the process.

2.2 Models of Sensemaking

2.2.1 Sensemaking and Intelligence Analysis

The term "sensemaking" refers to the process of understanding an unfamiliar, unstructured, information-rich situation. It can be considered as the strategies and behaviors evident when people collect, evaluate, understand, interpret, and integrate new information for their own specific problem/task needs.

Sensemaking studies have been initiated by the Information Science and the HCI communities. While sensemaking is clearly different from intelligence analysis in terms of scope and subject matter, its underlying notion is quite relevant. Especially in the Visual Analytics domain, sensemaking and intelligence analysis are considered interchangeable, as stated in [60]:

> "Many forms of intelligence analysis are what we might call sensemaking tasks. Such tasks consist of information gathering, re-representation of the information in a schema that aids analysis, the development of insight through the manipulation of this representation, and the creation of some knowledge product or direct action based on the insight."

This shared meaning also can be found in intelligence literature [7], in which he distinguishes "intelligence" from "information":

> "Intelligence is more than information. It is knowledge that has been specially prepared for a customer's unique circumstances. The word *knowledge* highlights the need for human involvement. Intelligence collection systems produce... data, not intelligence; only the human mind can provide that special touch that makes sense of data for different customers' requirements. The special processing that partially defines intelligence is the continual collection, verification, and analysis of information that allows us to understand the problem or situation in actionable terms and then tailor a product in the context of the customer's circumstances. If any of these essential attributes is missing, then the product remains information rather than intelligence."

2.2.2 Models of Sensemaking

There are several influential theoretical works on sensemaking. Among sensemaking models developed from different perspectives, well-known are Dervin's sense-making methodology [17, 18, 19], Russell's learning loop model [66], the data-frame model by Klein et al. [41,42], and Pirolli's notional model of the sensemaking loop [60].

Dervin's model of sensemaking [17] sees the individual as continually making sense as s/he moves through time and space in an ongoing life journey. People continually make sense of their actions and their environment and this makes movement possible. Occasionally, people encounter a situation where movement is blocked by a discontinuity that does not fit their internal sense—there is a cognitive gap. The person defines the nature of the gap, and based on this interpretation, selects tactics to bridge the cognitive gap. People then cross the cognitive bridges in order to continue on the journey. Dervin's sensemaking approach emphasizes understanding of how individuals define a gap situation and how they attempt to bridge cognitive gaps. In order to bridge cognitive gaps, users seek, process, create, and use information.

Russell et al. [66] define sensemaking as “a process of searching for a representation and encoding data in that representation to answer task-specific questions”. His model (Figure 7) indicates the iterative nature of sense-making: processing may go through several iterations until sense-making is successful. The first process is a search for a good representation; a sensemaker creates an initial representation which he thinks could capture salient features of the information in a way that support the accomplishment of the task (the Generation Loop). Then there is an attempt to encode information in the representation, which results in the Instantiate Representations (the Data Coverage Loop). However, when the sensemaker’s understanding of the sensemaking task grows, he may find that the initial representation is not adequate to explain the sensemaking problem. When s/he finds this mismatch between his representation and the task (called “residue”), the person is motivated to adjust the representation or find a better representation so that it has better coverage, (the Representational Shift Loop). The result is a better, more compact representation of the essence of the information relative to the intended task. Thus, structural representation plays a crucial role in all sense-making processes.

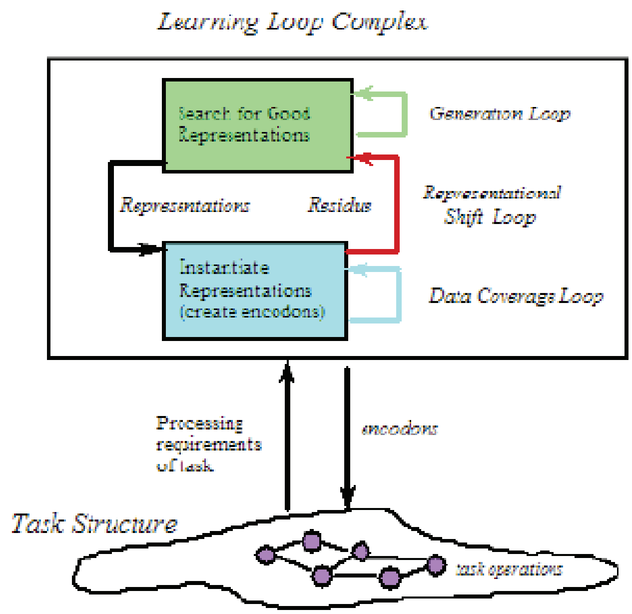

Figure 7. Learning loop complex theory of sensemaking, taken from [66]

In the data-frame model of sensemaking [41, 42], sensemaking is defined as the deliberate effort to understand events and is typically triggered by unexpected changes or surprises that make a decision maker doubt their prior understanding. The data-frame sensemaking model provides a description of how people generate an initial account to explain events and understand the current situation with new information flowing in. Situation awareness is a model of the current situation held in working memory and sensemaking is the active process of building, refining, questioning and recovering situation awareness. The process of building up situation awareness is explained by the data-frame model. The frame is the explanatory structure into which current data go. The frame defines and explains relationships, and guides the search for new data.

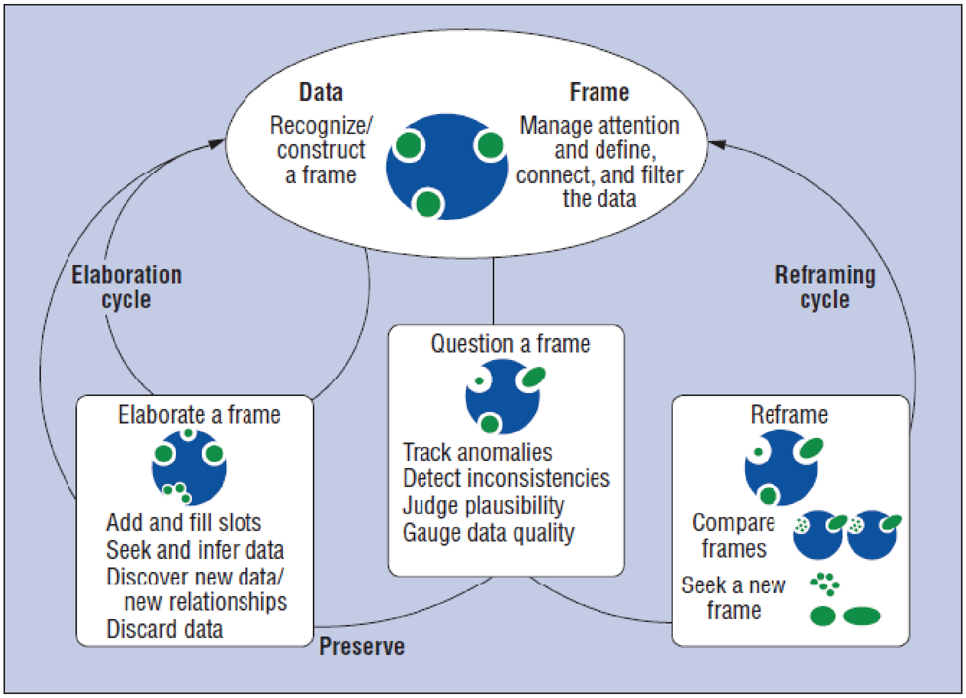

Figure 8. The data/frame theory of sensemaking, taken from [42]

Perhaps the most widely used model in the Visual Analytics domain would be Pirolli and Card's Think Look Model [60]. Pirolli and Card developed another model as a broad brush description of intelligence analysis as an example of sensemaking, based on the preliminary results from a cognitive task analysis and think aloud protocols. This model was developed specifically for intelligence analysis process, and the visual analytics community, which has few empirical studies on analysis processes, has largely applied the model to designing new systems.

In this model, they introduce the data flow that illustrates the transformation of information as it flows from raw information to reportable results. The overall process is organized into two major loops of activities: (1) a *foraging loop* that involves processes aimed at seeking information, searching and filtering it, and reading and extracting information (Pirolli & Card, 1999) possibly into some schema, and (2) a *sense making loop* [39] that involves iterative development of a mental model (a conceptualization) from the schema that best fits the evidence. According to the model, an analyst through filtering of message traffic and active search, collects raw information into an information store or "shoebox". Snippets of this evidence are collected into another store or "evidence file". Information from this evidence may be represented in some schema or conceptual form (the framework of the Learning Loop Complex model). This organization of information is used to marshal support for some story or set of hypotheses. Finally the information is cast into an output knowledge product, such as a briefing or a report.

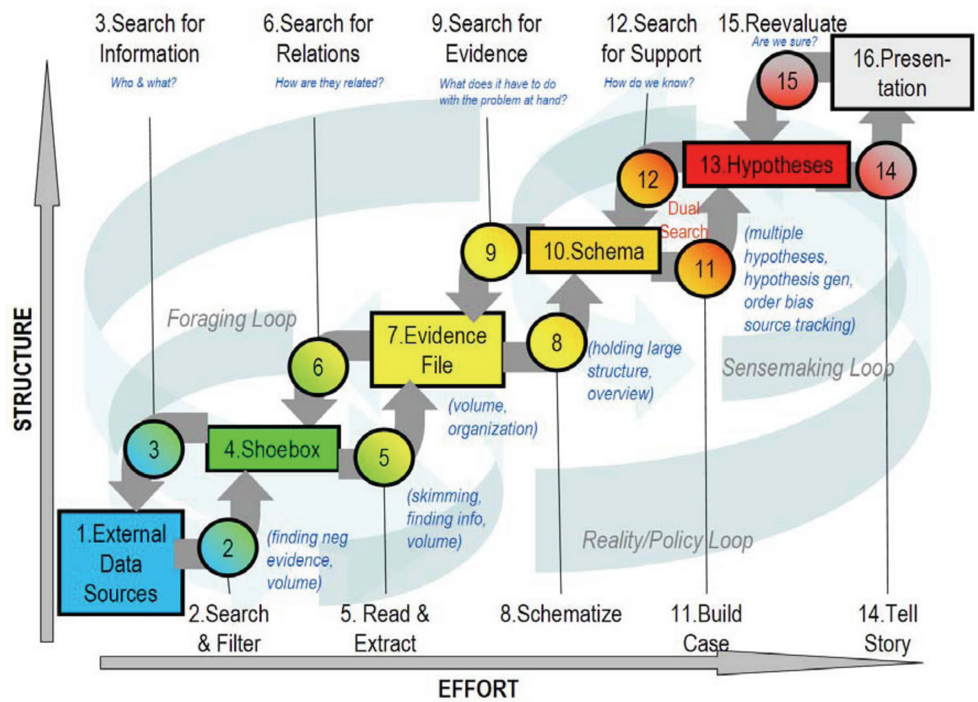

Figure 9. Notional model of sensemaking loop for intelligence analysis, taken from [60]

While this model provides new insights about the intelligence process and helps researchers find leverage points for analysis tools, it was suggested as a starting point to investigate the domain as they explicitly stated. Furthermore, this model was developed based on preliminary findings from cognitive task analysis and individual interviews, which might not be enough to describe details of the process and the domain. More empirical studies that closely investigate users' work processes and the domain are required. What have we misunderstood about the intelligence or sensemaking process?

How should our assumptions be changed to build a tool that truly helps what they are really doing? More empirical studies that seek to validate and refine existing models may be needed.

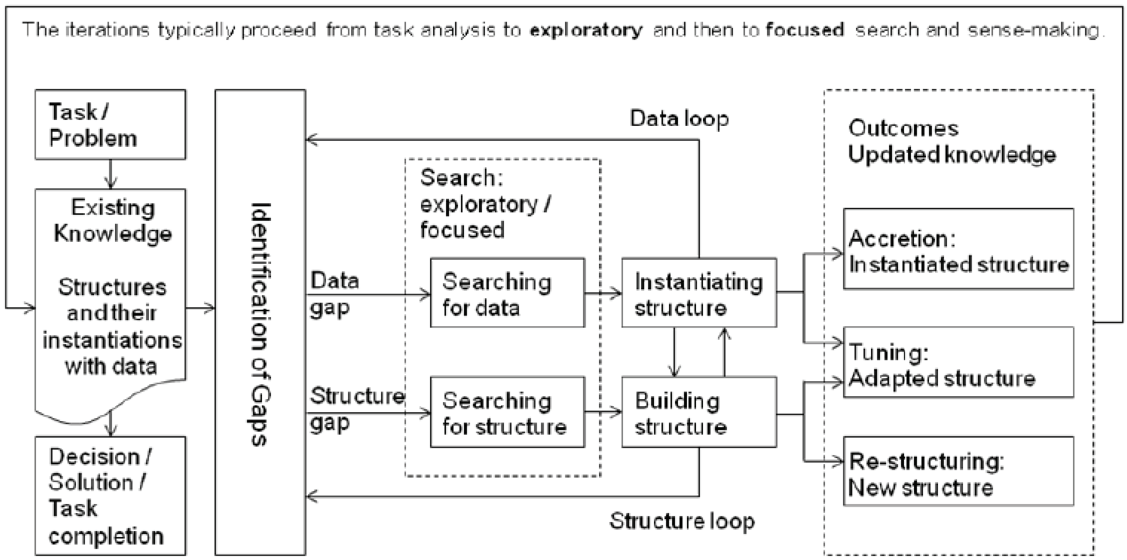

In each process, the cognitive mechanisms listed below can be used as applicable

Inductive (data-driven, bottom-up)	**Structure-driven (logic-driven, top-down)**	**Other**
Key item extraction	Definition	Comparison
Restatement	Specification	Analogy
Judgment or evaluation	Explanation-based mechanisms	Classification
Summarization	Elimination	*Stereotyping*
Schema induction	Inference	Semantic fit
Generalization		Socratic dialogues

Figure 10. Iterative Sensemaking Model, taken from [87]

Zhang [87] pointed out that while Pirolli and Card's model clearly illustrates the steps and outputs involved in a sensemaking process, sensemaking does not always have clear beginning and ending points as described in the model. He argues that "the simplified waterfall model runs counter to empirical evidence about several sensemaking tasks, for example, expert decision making." Through a qualitative study with fifteen students working with news writing and business analysis tasks, he investigated how people structure their conceptual space with the assistance of note-taking and concept mapping tools. Based on the results and previous sensemaking research, learning theories, and cognitive psychology, he proposes an iterative sensemaking model (Figure 10).

This model views the sensemaking process as several “search – sensemaking” iterations. In each iteration, the sensemaker goes through some search activities (exploratory and focused search for data or structure) followed by some sensemaking activities including gap identification, building structure, instantiating structure and creating products activities. Zhang emphasizes that sensemakers may go through the paths in the model idiosyncratically and heterogeneously.

While his model also provides cognitive mechanisms involved in the process, it was developed specifically focusing on how sensemakers create and use structured representations of sensemaking.

2.3 Systems for Investigative Analysis and Sensemaking

While there are few visual analytics systems that have been developed based on a theoretical foundation or model, most of them, especially those related to intelligence analysis, adopted Pirolli and Card’s sensemaking model as a reference. This section illustrates examples of systems for investigative analysis and/or sensemaking and how existing models and guidelines influenced the design process.

Jigsaw [35, 78, 79], a visual analytics tools for investigative analysis, was initially designed to address two leverage points identified by Pirolli and Card: (1) the cost structure of scanning and selecting items for further attention and (2) analysts' span of attention for evidence and hypotheses. By showing visualizations as separate views onto a text document collection and highlighting connections between entities across the collection, Jigsaw helps the analyst have a better understanding of the overall themes and investigate the collection by particular events and facts in the documents. To help initiate the development, researchers adopted analysis exercises created by Hughes of the Joint Military Intelligence College [30], which involve collections of fabricated reports with an embedded master plot. By using the exercises as a task, Jigsaw provides a number of features that can support the investigative process.

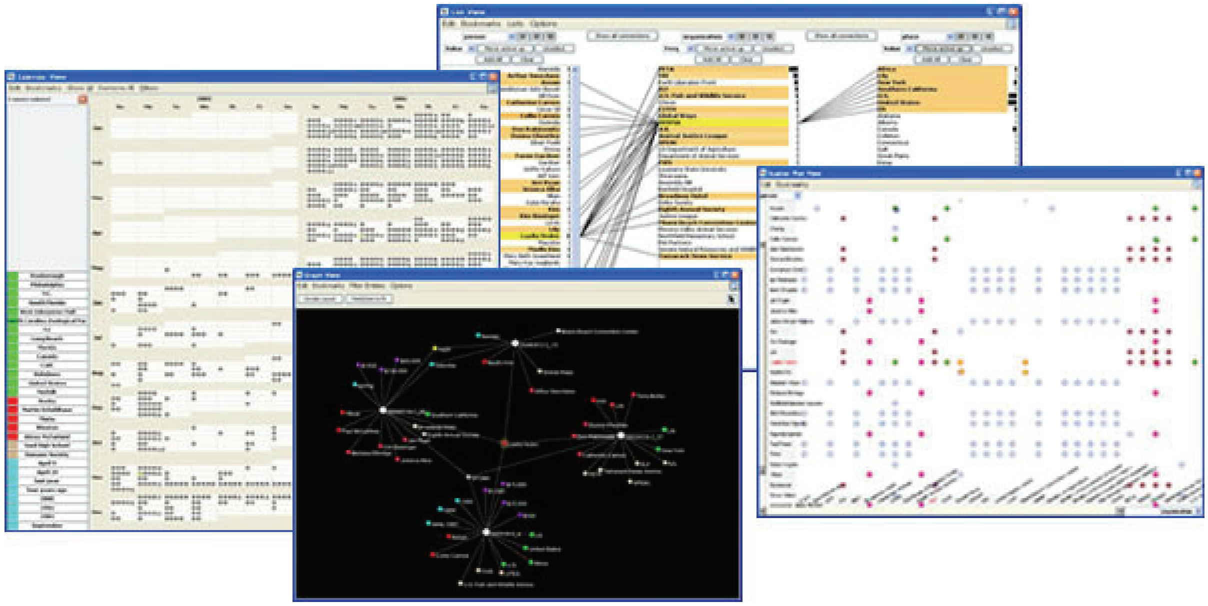

Figure 11. Views in Jigsaw

The Entity Workspace [5] is another tool to amplify the usefulness of a traditional evidence file - an electronic document into which text snippets and hand-typed notes are placed - that is widely used by analysts to keep track of facts. Entity Workspace builds up an explicit model of important entities (people, places, organizations, phone numbers, etc.) and their relationships. Using this model, it helps the analyst find and re-find facts rapidly, detect connections between entities, and identify important documents and entities to explore next. The authors explicitly states that their approach to analytical processes is based on the Pirolli and Card's sensemaking model.

While both Jigsaw and Entity Workspace put an emphasis on the foraging stage of the Think Loop Model, a few tools focus on the sensemaking stage of the model such as "schematize" and "hypothesize." The Sandbox system [62, 87] is one of such systems: "The Sandbox work presented here is focused on the 'sense-making loop' and the 'exploiting' process of the exploration-enrichment-exploitation stages of foraging." The

Sandbox is a sensemaking work that provides alternatives to paper or text editors for analysis activities such as hypothesizing, fleshing out hypotheses with evidence, corroborating, grouping, annotating and prioritizing (Figure 12). The goal of the system is to help ensure more rigorous thinking and increase an analyst's cognitive span by making the nature and structure of the analysis more explicit. The Sandbox supports authoring and organizational infrastructure by providing interactive visualization techniques and templates for building visual models of information and visual assessment of evidence.

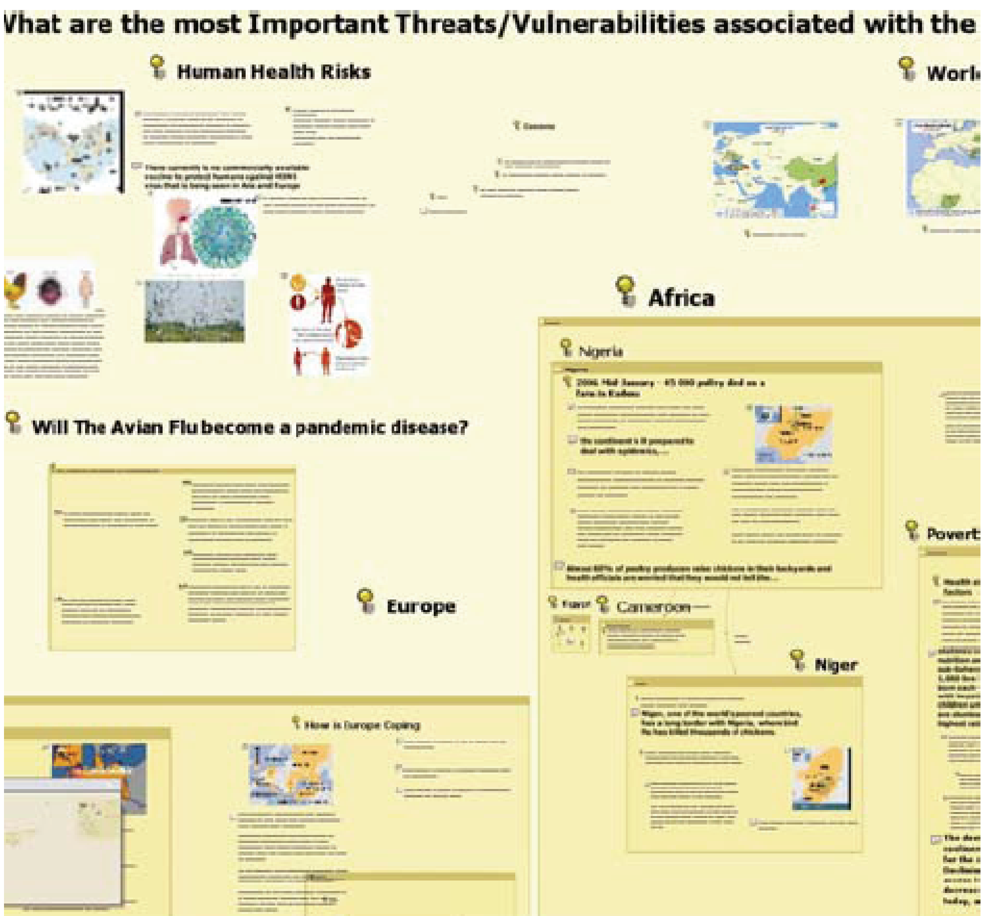

Figure 12. An emerging analysis using the Sandbox interface

Analyst’s Notebook [33], one of the most widely used visual analytics tools in practice, provides diagrammatic visual representations and is mainly used for link analysis (e.g., transactions, phone calls). While the system can import text files and do automatic layout, its primary application seems to be the creation and refinement of case charts. That is, rather than a thinking tool, it seems better suited as a presentation or report tool since it does not provide a variety of ways to visualize information.

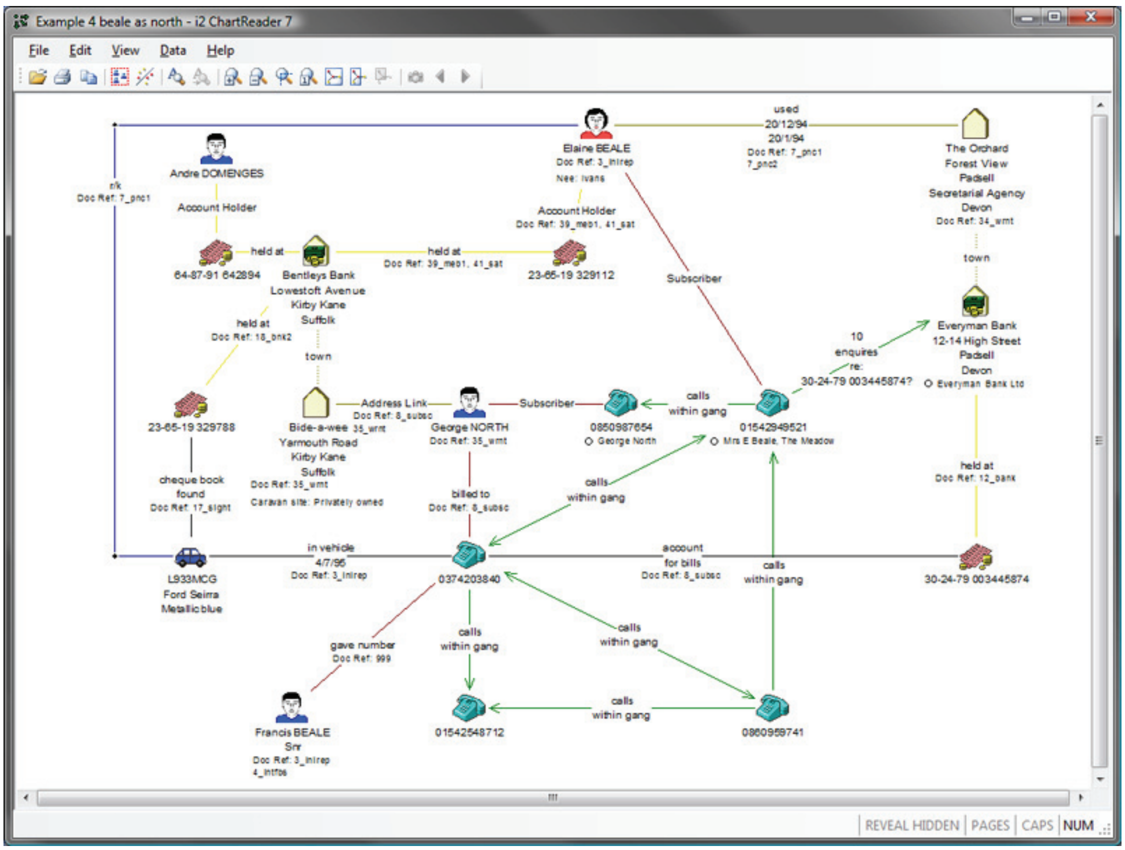

Figure 13. Link analysis using Analyst's Notebook

Analyst’s Workspace [2] is a spatial environment for document sensemaking that integrates the activities of foraging and synthesis into a single thread. AW was developed for use on a large, high-resolution display so that the analyst can use a spatial approach to manage information. The analyst can use the system to explore a collection of documents, opening them in the space and then arranging them into meaningful patterns as part of the sensemaking process. It also provides several functionalities such as notes, text highlighting, and visual links.

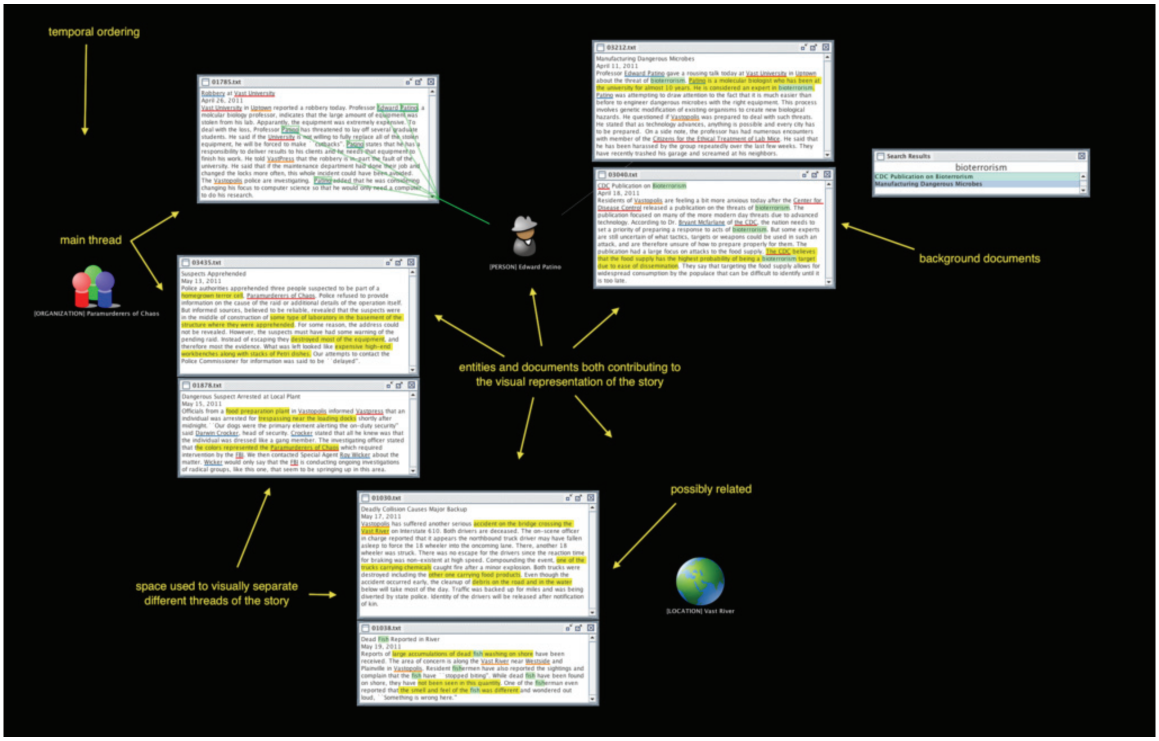

Figure 14. Documents, highlighting, entities, and spatial relationships in Analyst's Workspace [2]

Compendium [15] is hypermedia knowledge mapping software that provides a flexible visual interface for managing connections between information and ideas (Figure 15). It can be used as a personal sensemaking tool to manage one’s personal digital information resources as it allows a user to drag and drop information in any document, website, email, and image, to organize them visually, and to connecting ideas and arguments. It further supports collective sensemaking or group sensemaking in situations such as workshops and meetings [73] by allowing groups to construct graphical representations in real time.

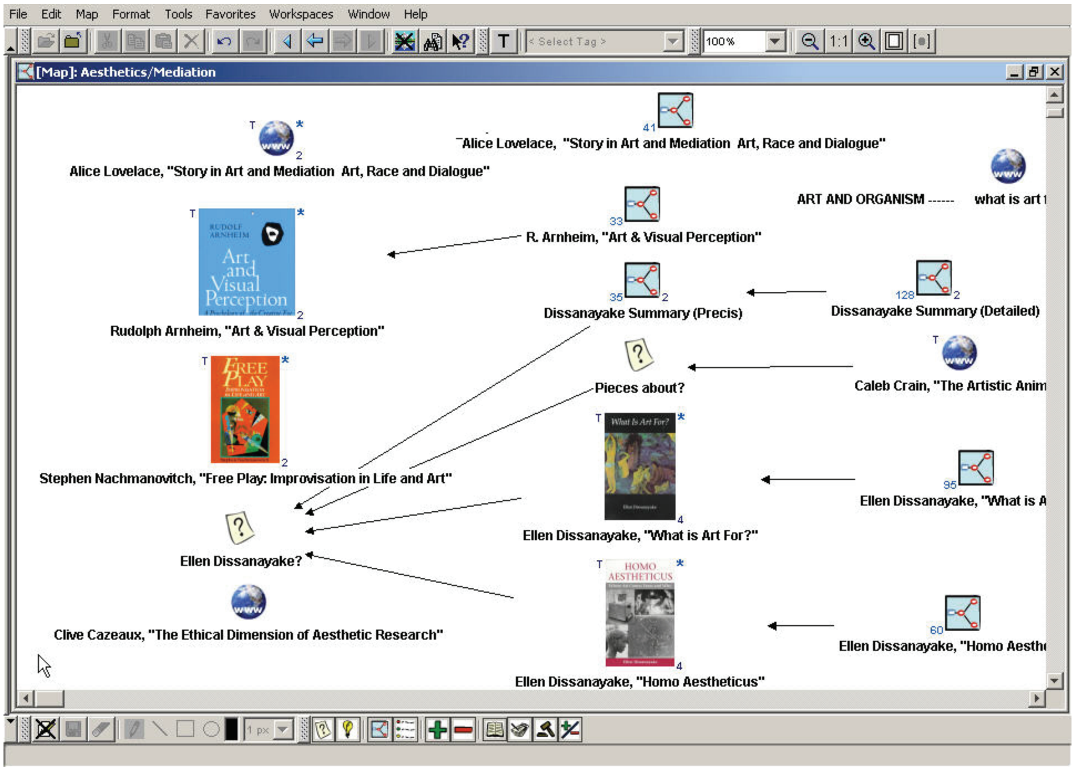

Figure 15. Compendium for literature analysis

Although not developed for sensemaking or analysis per se, there are several systems such as Zotero [89], Google Reader [24], and MediaWiki [46] used by intelligence analysts to collect and organize information. These lightweight systems seem to be preferred by professionals because of their flexibility and ease of use.

While all these systems provide a number of useful, unique functionalities, only a few of them have been developed upon existing models. Moreover, even such systems rely on a single model without validating the appropriateness of the model regarding the system. Part of the reason might be the lack of design guidelines or models of analytical processes tailored to the development of visual analytics systems. If more design implications and guidelines can be enumerated from empirical studies, it will significantly help us improve the development process and the outcome of visual analytics tools.

2.4 Understanding Users through Formal User Studies in Visual Analytics

There are several valuable resources available that examine the analytical culture in general. These include a number of books published by former government intelligence analysts [28,37,43]. These books provide insights into the complex analytical process as seen by those who practice it as well as an understanding of some critical aspects of analysis.

As Visual Analytics has rapidly grown, Infovis and visual analytics researchers have become interested in understanding the analytical culture in other communities and their work processes to derive design implications. This section describes research efforts that sought to better understand users, user tasks, and their tool usage in visual analytics research.

Several studies have captured and characterized the work practices and analytical processes of individual or collaborative analysis through a qualitative approach. Chin et al. [13] conducted an observational case study with professional intelligence analysts in which participants worked on real-world scenarios, either as an individual analyst or as an investigative team. The researchers revealed various characteristics of the analytical processes of intelligence analysts, such as the investigative methodologies they apply, how they collect and triage information, and how they identify patterns and trends. Gotz et al. [25] also recognized the lack of public studies examining analyst behavior and conducted a user study with a few analysts to explore the ways in which they gather and process information. Through interview, observations, and written notes by analysts, they report important factors surrounding analyst behavior in information gathering and results processing, such as how they keep record and what their investigative style is like.

Another study [64] examined how analysts synthesize visual analytic results by studying domain experts conducting a simulated synthesis task using analytical artifacts printed on cards on a large paper-covered workspace. Based on analysis of video coding results, he identified several characteristics in the process of synthesis such as the use of different approaches to collaborative synthesis, a variety of organizational metaphors when structuring information, and the importance of establishing common ground and role assignment. While these studies did not evaluate specific visual analytic tools or features per se, they provide valuable implications to inform design directions for future support tools.

While previous studies aimed at understanding users and their current work processes without a specific tool, researchers have also investigated the usage of visual analytics tools developed. As more visual analytics tools become available, researchers have become interested in evaluating the utility of such tools. While the reports of usability studies are helpful, there is a growing desire for alternative methods of evaluation [68] because the purpose of visual analytics systems is to assist users gain insights and new forms of knowledge by supporting their tasks [8, 76]. Demonstrating benefits provided by those systems will encourage more widespread adoption of visual analytics tools.

However, those types of evaluations involve a number of considerations, and such challenges in evaluating the utility of visual analytics systems promoted research studies in evaluation methodologies and metrics. There has been an emerging method called Multi-dimensional In-depth Long-term Case studies [75] which is adapted to study creative activities of users of information visualization systems. Encouraging information visualization researchers to study users doing their own work in the process of achieving their goals, the paper lists lessons from ethnography methods used in HCI [31,32,65] including observations and interviews and suggests evaluation methodology guidelines for information visualization researchers. Scholtz [50,70,71] emphasizes that the development of metrics and methodologies for evaluation is necessary to help researchers measure the progress of their work and understand the impact on users. She argues that the evaluation of visual analytic environments requires researchers to go beyond performance evaluations and usability evaluations, and proposes five key areas to be considered as metrics and methodologies for evaluation: situation awareness, collaboration, interaction, creativity, and utility.

Few studies have investigated the utility of visual analytic tools for investigative analysis. A study by Bier et al. [6] assessed the suitability of their Entity Workspace System in the context of design guidelines for collaborative intelligence analysis. The researchers modified their system based on five design guidelines and evaluated the system in both a laboratory study with intelligence analysts and a field study with an analysis team. Relying on analysts' subjective feedback in conjunction with quantitative logging data, they confirmed the positive effects of the tool on collaboration and the usefulness of the design guidelines for collaborative analysis. Perer and Shneiderman [56] recognized the limitations of traditional controlled experiments in examining the process of exploratory data analysis and developed an evaluation methodology for studying the effectiveness of their system, SocialAction. Consisting of a long-term case study and in-depth interviews, the evaluation confirmed the core value of SocialAction - integrating statistics with visualization – and further provided guidance for redesign of the tool.

My research is an extension of all these research efforts. In order to build useful visual analytics systems and facilitate their widespread adoption, it is crucial to understand the user community, their tasks, and their experience with visual analytics systems.

CHAPTER 3

CHARACTERIZING THE INTELLIGENCE ANALYSIS PROCESS: INFORMING VISUAL ANALYTICS DESIGN THROUGH A LONGITUDINAL FIELD STUDY

3.1 Introduction

Visual analytics applies to many domains and problem areas, but one area of particular study since the beginnings of the field has been intelligence analysis. Intelligence analysis is a cognitively demanding process, one that seems ideal for the application of visual analytics tools. Accordingly, a growing number of systems have been built for it [5, 33, 78, 87].

Research in human-computer interaction teaches us to deeply analyze and understand end-users and their problems in order to design appropriate computational solutions. I question whether visual analytics systems, including some of our own, have been based upon a deep enough understanding of the discipline. Relatively few studies of intelligence analysts, their tasks, and their work processes exist. Notable exceptions [13, 36, 60, 64] provide initial insights into the field, but I have frequently interacted with analysts who feel that their practices are misunderstood and that visual analytic systems often fail to address their important problems.

To address these concerns and to learn more about the analysis process, I conducted a longitudinal, observational field study of intelligence analysis on real world problems. Unfortunately, getting access to working, professional analysts is challenging. As an alternative, I studied analysts-in-training who are soon to become working professionals. More specifically, I studied groups of students from the Department of

Intelligence Studies at Mercyhurst College as they conducted a term-long intelligence project.

The goal was to better understand what these young analysts do, the challenges they face, and how we might be able to help them. Thus, the contributions of my research include a characterization of the processes and methods of intelligence analysis that I observed, clarification and reflection of several beliefs about intelligence analysis processes and practices, and resultant design implications for visual analytics systems for intelligence analysis.

In their recent publication, in which they suggested seven scenarios in evaluating information visualization, Lam et al. [44] quoted Munzner [48] saying, "hardly any papers devoted solely to analysis at this level [problem characterization] have been published in venues explicitly devoted to visualization" and argued for the importance and the need for this type of evaluation studies.

3.2 Methods and Procedures

In order to investigate the intelligence analysis process in-depth, I conducted an observational study of teams of analysts conducting an in-class intelligence project. In the term-long (ten-week) project, each team addressed a real intelligence problem proposed by a client. I observed three teams, monitoring their process throughout the project. At the end of the project, each team produced final deliverables and presented their findings to decision makers.

3.2.1 Participants

I recruited three groups of students, one team of four undergraduate students (Team A) and two teams of five graduate students (Teams B and C), from the Department of Intelligence Studies at Mercyhurst College [47]. Mercyhurst's Intelligence Program, started in 1992, provides education for students who want to pursue a career as an intelligence analyst. It is recognized as one of the top programs for intelligence studies in the United States, offering a broad range of classes and degrees for students seeking a career as an analyst in national security, law enforcement, or the private sector.

I recruited students who were taking the courses named “Strategic Intelligence” (undergraduate) and “Managing Strategic Intelligence” (graduate), in which teams are required to conduct an analysis project over a ten week term. The two courses are very similar with respect to the projects. The students all were close to graduation, with past internship experience, and most of whom had already received job offers.

3.2.2. Task

Different types of intelligence questions exist - I focused on one of the most common types, strategic intelligence. Strategic intelligence is “intelligence that is required for the formulation of strategy, policy, and military plans and operations at national and theater levels [27].” Strategic intelligence is exploratory and long-term in nature. The requirement for tasks within the class was that “the questions should be relevant and relatively important to the client’s success or failure but outside their control.” My colleague and I served as a client/decisionmaker for team A in order to observe the process even closer, whereas Teams B and C worked with external organizations. The specific issues each team addressed were:

Team A

The strategic assessment of potentially influential factors to the evolution of computer-mediated undergraduate and graduate distance education: What aspects of computer-mediated distance education will likely influence R1 institutions during the next 5, 10, 20 years with specific, but not exclusive, emphasis on undergraduate education and computer science?

Team B

Who are the key people, technologies and organizations that likely currently have or will develop the potential to disrupt or replace traditional US national security Intelligence Community (IC) analytic work flows and products with commercially available products available over the next 24 months?: Criteria that will be used to identify these key players are:

- *Those that are not beholden to the IC or US Government as primary sources of funding.*
- *Those that are looking at future based events or actions that are outside the control of the forecaster/predictor.*

Team C

What are the most consistent and identifiable characteristics displayed by potential insider threats to (a defense department)?

- *An insider threat will be defined as an individual or collection of individuals employed directly or indirectly by the department who violate security or access control policies with the intent of causing significant damage to the department's personnel, operations, or information.*
- *Within the broad range of insider threats, special priority will be given to violent threats and improper diversion of information or physical assets.*

The teams updated the status and the process of the project on a wiki site. At the end of the semester, they needed to produce a final report that synthesizes analytical results, and strategies of the entire analysis process.

3.2.3. Study Protocol and Procedures

The analyst teams conducted the project for ten weeks. Before the project began, the external clients formulated a draft of their initial intelligence problem. In the first week of the project, the clients conducted a conference call with the analyst team to discuss the scope and requirements of the problem. During the next two weeks, the analysts refined the problem and wrote a formal statement of the intelligence question. Upon approval from the decisionmakers, the teams began working on the problem, which took another seven weeks.

The wiki platform was used as a workspace for analysts to document their process and findings, and I was able to monitor the wiki's status throughout the project period. The final report of the projects also was placed on the wikis.

During the project period, I conducted two face-to-face meetings with each team – one in week 7 and the other in week 10. In the meetings, I interviewed each team as a group and the class instructor in order to learn more details about the project's status, process, difficulties, and future steps. Each interview took approximately an hour. While the interview was semi-structured, I followed an interview guide containing several key topics [11, 48], including:

- How do the analysts perceive their analysis process?
- What barriers and difficulties do they encounter?
- Tools and aids being used - where and why?
- Collaborative aspects in the analysis process
- Where in the process can technology help?

I also observed two team meetings firsthand, which took about 3 hours in total.

3.2.4. Data Collection and Analysis

Most of the process descriptions and produced artifacts were stored digitally. The teams reported methodologies, tools used, sources, as well as the findings on their own website (wiki). To further understand the process, I analyzed interview notes and audio recordings from the interviews. I used the artifacts produced by the analysts, such as drawings, wiki pages, tables, and slides as further data. Additionally, I had access to history logs of wiki page changes.

For analysis, I used an approach that borrowed principles from grounded theory [80]. After transcribing each interview's audio recording, I conducted open coding, in which I tried to identify and categorize phenomena found in the text. In this process, I read each sentence and paragraph, and labeled them in order to find out what it is about and what the problem is. One emergent theme focused on the analysis process, including methodology and challenges encountered. Then in the axial coding process, I began making connections between categories and themes identified through the open coding and generated a refined set of generic categories, which are described in the following sections. When making the connections between preliminary categories, I focused on similarities or causal relationships between them. Although I analyzed the data in a manner similar to that in grounded theory, I did not follow the last step of the approach, which is selective coding. Selective coding is the process of choosing one category to be the core category, and relating all other categories to that category. While the idea is to develop a single storyline–a core concept or theory, I did not intend to suggest any theory or a single concept through the study. Rather, the purpose was to suggest several core themes and communicate them with others, opening up discussions.

3.3 Overall Analysis Process

Through the project, I found four component processes essential to the overall analysis: constructing a conceptual model, collection, analysis, and production.

Phase 1: Constructing a Conceptual Model

Once the teams and clients/decisionmakers finalized the requirements of the intelligence question, the teams started to build a conceptual model, which is a map of issues and concepts that the team will be investigating to address the problem. The conceptual model illustrates the areas the analysts need to research by helping them to visualize the question at hand. The question is placed in the center, and then several high-level components of the question surround the question (Figure 16). Each component branches out and creates a bigger map, from which the team gains an idea of the areas with less/more information that they need to research. This allows the team to focus on collecting a set of data with an appropriate scope.

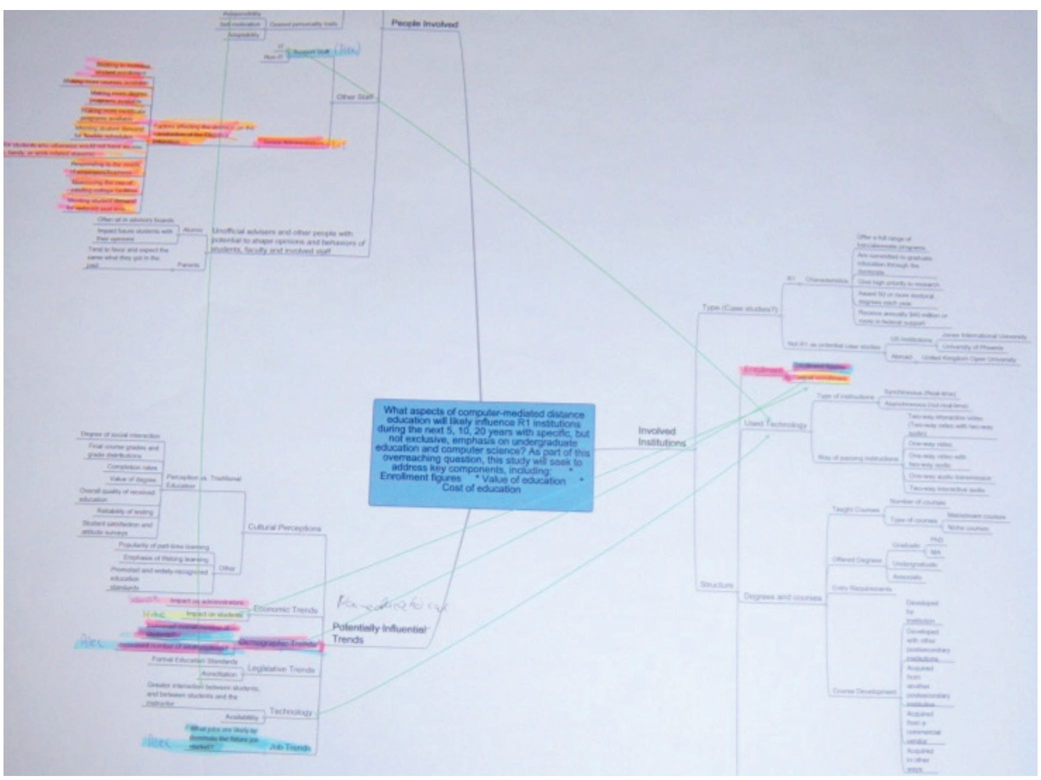

Figure 16. Conceptual model. Printed from Mindmeister

Phase 2: Collection

While working on the conceptual model, the teams also assigned areas/concepts to each member. Next, they collected information from various sources including online and offline sources (e.g., interviews with experts), which they call "all-source intelligence". While each analyst was responsible for collecting data about their assigned topic(s), the team shared their sources using Zotero, a web browser plug-in for gathering and organizing source material. This allowed teammates to view the data like a common library – other team members might already have found information that they need.

Phase 3: Analysis

The analysis phase exhibited various characteristics depending on the requirements and analytical methods used. In this phase, analysts processed data that they collected from many different sources in order to convert "information to knowledge." While team A directly began writing short format analytical reports on each topic, team B and C used a more structured format (e.g., spreadsheets) to quantify information and rank the significance of each topic or entity. No matter which method they used, the initial analysis of each topic/entity was undertaken and written by one person in accordance with the assigned topic. However, everyone on the team could review and comment on the others' work via the wiki pages. In all cases, the analysis phase was incorporated with the collection and the production phase.

Phase 4: Production

Once individual collection and analysis was almost finished, the teams met and tried to synthesize findings from each part, which led to the "key findings" – the major product of the analysis. Production was an intensive reading/writing process in which the team collaborated tightly with each other. This stage was more to prepare a presentation for the decisionmakers. Team members repeatedly checked their sources and findings to make sure that they were consistent and logical.

3.3.1. Intuitive Analysis – Team A

Team A addressed potentially influential factors to online distance education in the near future. Because the requirements were rather broad and intuitive, the team decided to take a top-down approach, investigating meta-information sources such as research that forecasts future education trends.

Instead of using a specific analytic method, this team depended considerably on the conceptual model and used it as a guide throughout the entire project. After collecting and reading information for their designated topic, each analyst wrote a short format analytical report that synthesized the information. Most of the analysis simply involved reading. For a few topics that required careful weighing of alternative explanations, the team employed analysis of competing hypothesis (ACH) [28]. While documenting results, everyone was able to review and edit the others' drafts on the wiki page, and team members frequently discussed others' analysis (short write-ups) both online and face-to-face. After working on the individual topics, the team met to write key findings together. This team invested considerable efforts in synthesizing their findings because their narrative was extremely important for their intuitive type of analysis.

3.3.2 Structured Analysis – Teams B and C

Teams B and C used structured analysis with quantified information because their research questions tended to be more specific and required rank-ordering of entities (e.g., top x indicators, key people/companies). Both teams built their conceptual model in the beginning as a base model. For these teams, however, the model was more of a collection plan rather than an actual conceptual model. Although they used the model to collect information and divide up the work, they did not refer to it for the remainder of the project. Instead, they started building a matrix in a spreadsheet to collect and analyze data from diverse sources.

The purpose of the matrix was to evaluate each entity based on criteria chosen and identify the most influential ones, those of most interest to the decisionmaker. Team C, that was asked to identify indicators displayed by potential insider threats to a defense department, analyzed data from the 117 case studies about crimes using a matrix (Figure 17). They used it to compare the relationship between crimes and motivations, as well as crimes and indicators.

Categories of Crimes	Number of Indicators per Crime	Number of Indicators/Person	Personal	Isolationist Personality	Inflated Self Image	Mental Health Issues / Disorders	Suicidal tendencies	Hopelessness	Survivalist Mentality	Strict / Absent parents	Excessive interest in	Perce
				5	6	9	4	1	0	2	4	7
Homicide	26		9	3	2	5	4	1	0	0	4	4
Murder	26		9	3	2	5	4	1	0	0	4	4
Nidal Malik Hasan		17	4	1	0	0	1	0	0	0	1	0
John Russell		15	5	0	0	1	1	0	0	0	0	1
Hasan Akbar		15	4	1	1	1	0	0	0	0	0	1
William Kreutzer		14	6	1	1	1	1	0	0	0	1	0
Dean Mellberg		8	4	0	0	1	0	0	0	0	1	1
Dr. Bruce E. Ivirs		10	6	0	0	1	1	1	0	0	1	1
Espionage	24		7	2	3	3	0	0	0	2	0	2
Espionage	17		6	1	1	2	0	0	0	0	0	1
Brian Patrick Regan		4	2	0	0	1	0	0	0	0	0	0
Timothy Steven Smith		5	2	0	0	1	0	0	0	0	0	0
Ana Belen Montes		1	1	1	0	0	0	0	0	0	0	0
Ariel Weinmann		6	1	0	0	0	0	0	0	0	0	1
Robert Chaegun Kim		2	1	0	0	0	0	0	0	0	0	0
Kurt G Lessenthien		2	1	0	0	0	0	0	0	0	0	0

Figure 17. Case study matrix of crimes

In both teams, the matrix captured the conceptual model and how each team was thinking about the question. Filling in the cells was a time-consuming part as analysts needed to read and analyze each case/source to fill in one cell, addressing "the devil in the details." This type of analysis required additional efforts in the production phase. Initially, the teams converted qualitative information from sources into quantitative information for rank-ordering. Once they had completed the matrix, the teams needed to transform its data into a story so that it could be made useful to decisionmakers.

3.4 Tools and Methods Used

The teams used various software tools and analytical methods to develop hypotheses, arrive at analytic estimates, and create written reports and multimedia products.

Wikispaces/Google Sites: The teams used a wiki platform (Team A&B – Wikispaces, Team C- Google Sites) to exchange gathered information, aid administration, and share organizational details. The wiki sites became part of the final product, displaying the key findings, terms of reference, and all analytic reports.

Mindmeister (conceptual model): Mindmeister is an online mindmapping tool the teams used to build a conceptual model [49]. A conceptual model provides a revisable platform to view the requirements and their components. As research and facts begin to support or refute initial ideas, main ideas become more solidified and focused.

Zotero: The teams used Zotero as a source collection database [89]. Downloaded as an Add-on to Mozilla Firefox, Zotero allows the analyst to search websites and save the sites in a database that is accessible through the Zotero website. The teams used the Group Library feature to place their sources in a single database.

Website Evaluation Worksheet: To evaluate the credibility of the online sources, all the teams used the Dax Norman Trust Scale [54]. This matrix allows scores to be applied based upon criteria such as clear bias, corroboration of information, and the analyst's overall perception of the source. Based on the sum of scores, the source can score a High, Moderate, Low, or Not Credible rating.

Decision Matrix: A Decision Matrix is a decision-support tool allowing decision makers to address a problem by evaluating, rating, and comparing different alternatives on multiple criteria. Both team B and C employed a modified version of a decision matrix appropriate to address their problems.

Analytic Confidence: Each report includes an analytic confidence section that conveys to the decisionmaker the overall doubt connected with the estimative statement(s). While assessing the level of analytic confidence, the teams used Peterson's method [57]. Peterson identified seven factors that influence analytic confidence: the use of structured analytic methods, overall source reliability, source corroboration, level of expertise on subject, amount of collaboration, task complexity, and time pressure. In the analytic confidence section, the teams addressed these six factors as applicable to the particular estimate.

Social Network Analysis: Team C employed social network analysis using i2's Analyst's Notebook [33] to see relationships within industry. The team analyzed the social network analysis based on betweenness and eigenvector scores.

Analysis of Competing Hypotheses (ACH): Team A used ACH for some problems. ACH is a simple model for assessing alternatives to a complex problem. It takes analysts through a process for making a well-reasoned, analytical judgment. ACH is particularly useful for issues that require careful weighing of alternative explanations of what has happened, is happening, or is likely to happen. It also helps analysts minimize some of the cognitive limitations.

3.5 Understanding the Intelligence Analysis Process

Observing analyst teams helped me to better understand their goals and processes. In particular, the study highlighted a number of misconceptions I harbored about the intelligence process. Other visual analytics researchers may or may not share these preconceived beliefs, but I think that they have the potential for misunderstanding and are thus worth exploring.

<u>I. Intelligence analysis is about finding an answer to a problem via a sequential process.</u>

Some existing models of the intelligence analysis view it as an answer-finding process with a sequential flow, as noted in several models of the intelligence analysis process [36,43,83]. This perception presumes that the process is linear, sequential, and discrete by step. However, this model was not the intelligence process I observed. Instead, the process appeared to be more parallel and organic, as one analyst described:

> *Intelligence analysis is not about getting from point A to point B along the route, but it is better associated with basic research where you don't necessarily know where you are going to go. You're cutting a path through the jungle that's never been explored. That's what you're doing in most intelligence analysis projects. It's not a mechanical process in a sense that an assembly line is. It's a very exploratory activity by nature.*

<u>II. The key part of the intelligence process is the analysis of a specific set of data.</u>

Visual analytics systems often manipulate pre-processed data for analysis. A primary misconception about intelligence analysis is that the data analysis process, in which investigators analyze a set of collected data, is the most difficult part and takes the most time. This belief assumes that analysis occurs after investigators collect all data required for the analysis.

This view, however, needs to be changed. Although analysis is important, I observed that the process of "constructing a frame," as described in the Data-Frame theory [42] is more important. According to Klein and Hoffman, people begin sensemaking with some perspective, viewpoint, or framework—a *frame*. Frames can be in various meaningful forms such as maps and organizational diagrams. Frames define what count as data and also shape the data. Frames also change as a sensemaker acquires data. That is, frames shape and define the relevant data, and data mandate that frames change in nontrivial ways. Consequently, "constructing a frame" is not a simple process but involves a lot of thinking process and itself can be part of analysis. In other words, intelligence is about determining how to answer a question, what to research, what to collect, and what criteria to use. This process becomes part of the analysis - analysis implicitly occurs during the process of the construction.

Understanding that collection and analysis are integrated together in the process of building a frame is extremely important. Systems are not likely to be successful in supporting intelligence without acknowledging that fact.

III. Analysts do not often collaborate.

One common perception of intelligence views analysts as isolated individuals who prefer to work alone, struggling with pieces of information, rather than as collaborative teams [13]. However, a faculty member at Mercyhurst countered this perception:

> *Collaboration is almost all intelligence analysts have done in the context of the team. In the CIA or DIA, working as a team is pretty normal...Analysts are normally organized by function or geographical region. These typically operate as loose teams. Strategic projects almost always involve a team as do crisis projects (for example I am sure there are multiple Libya teams that did not exist a month ago). 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.*

During the study, I also observed many collaborative elements of intelligence analysis. Collaboration is commonplace in intelligence analysis, and understanding how that occurs is important because it influences one's whole notion of the process. The intelligence community itself has recognized the importance of improved collaboration since 9-11 [21]. Although collaborative tools have been built and they are pushing users into tighter collaboration, it is still important to understand where tighter collaboration will be beneficial and where it may not help much.

I found that multiple layers of collaboration exist in intelligence analysis and that the degree of collaboration differs depending on the type of task and the group dynamics. I observed that analysts usually do not collaborate tightly on data and content. Although the teams had meetings frequently – twice or three times per week – the main purpose was to discuss their status, issues, and the next steps.

IV. We can help intelligence analysts by developing sophisticated analytic tools that assist their thinking process.

Visual analytics researchers often seek to help intelligence analysts by developing technologically advanced analytical tools, thereby assisting their cognitive processes. The tools support specific types of analysis, specific analytical methods, and specific stages of the process. Such tools certainly can be helpful, especially to assist analysts to handle a flood of information.

However, this study revealed that analysts want something more than that. Currently, more than 50 analytic methods exist in the intelligence community [29], and analysts try many different kinds of techniques depending on the problem. Consequently, their dependency on a specific analytical technique is relatively low. Instead, the ability to manage the intelligence process effectively and employ various analytical methods and tools quickly is more important.

3.6 Rethinking the Intelligence Analysis Process

3.6.1 Linear vs. Parallel

One might believe that the way intelligence analysts work is quite simple and straightforward. First they specify requirements, build a conceptual model of what to research, then collect information, analyze data using various techniques, and finally write a report. This belief is a common misconception about intelligence as mentioned in the previous section. The reality is quite different. Rather than working linearly, analysts work on *everything* during almost the entire project. That is, analysts do not hold writing until enough information is collected; they keep revising analysis and writing as new information flows in. Analysts do not decide what to research and move on to collecting information; they start searching for information even when they are not sure what to research. Analysts do not produce final products after they are done with analysis; they already have an idea or a structure of final products in the very beginning, although it may be rough.

This "parallelism" is portrayed well in Wheaton's model of the intelligence process (Figure 18). In each phase, one of the core processes is emphasized most but all other functions operate in parallel. Wheaton argues that "All four functions begin almost immediately, but through the course of the project, the amount of time spent focused on each function will change, with each function dominating the overall process at some point [85]."

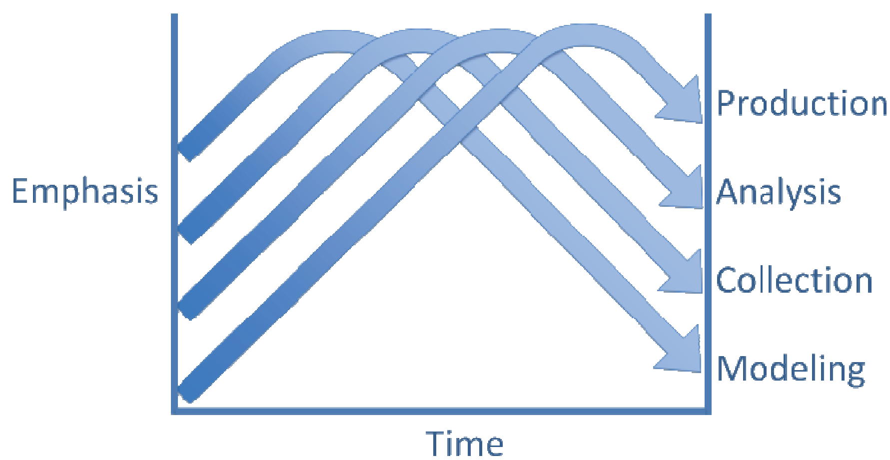

Figure 18. Wheaton's multi-phasic model of the intelligence process [85]

Although several distinct elements exist in the analysis process, all are very closely coupled and the connection is very organic. One can easily observe an analyst working on collecting new information while analyzing and checking the credibility of previously collected sources at the same time. In the study, I observed that a team's conceptual model changed drastically in the middle of the process, that a new information source was added ten days before the deadline, and that a previous analysis report was discarded and new analysis began in a late stage. While the teams were working on the matrix, they were collecting information at the same time to make sure that they were familiar with the area. Several quotes better explain this:

> *But it isn't as rigidly isolated as it's on that (traditional) cycle because you can't build a good conceptual model without knowing what's out there. So there's little bit of collection as you're building the model and we refined it.*

> *Our conceptual model is changing. It doesn't get set in phase 1 and we drive it, that's the difference between this process and an outline. An outline drives your production. But we are using it differently. As it changes, we're changing our analytic focus, we're making decisions about production, who's going to write something, who's going to do the analysis, based on how it's changing and that's being informed by new information that comes in.*

3.6.2 Pirolli and Card's Sensemaking Model

How does this new way of thinking about the intelligence process relate to Pirolli and Card's sensemaking model [60]? Because it is the most widely used model in the Visual Analytics domain, I was curious how well their model explains real-world intelligence analysis processes.

Pirolli and Card's model provides new insights about the intelligence process, suggests leverage points for analysis tools, and has guided the design of many visual analytics systems. However, I argue that the model still implies sequential, discrete stages of the intelligence process although it acknowledges that analysts can move either top-down or bottom-up or jump to different stages. For example, the model does not explain why analysts so frequently jump from one state to another state that is not adjacent. Many visual analytics tools thus support specific states only (e.g., shoebox and evidence file, evidence marshalling, foraging), and often they do not blend into the entire process of intelligence analysis.

More importantly, the model describes how information transforms and how data flows, rather than how analysts work and how they transition. It gives an illustration of how the form of information evolves from raw data to reportable results. However, it does not quite fit analysts' mental model of their work process because they do not work as information is transformed. Rather, information is transformed by how analysts proceed. Similarly, all different states of the model can exist at any point during the process. Analysts may have polished reports on certain sub-topics, drafts of analysis, structured matrices, and a collection of documents at a time.

The Pirolli-Card model identified various leverage points for visual analytics tools, but the linearity of the model could give researchers an inaccurate impression of the process. While models are inherently abstract and stage-based, it is important to understand the context and the purpose of the model. I would characterize their model as more of an information-processing process rather than intelligence analysis process. Pirolli-Card explicitly state that the model was suggested as a starting point to investigate the domain. While it has contributed to visual analytics researchers understanding of the domain, now we need to change our assumptions to build systems that better help intelligence analysts with their work.

3.7 Where and How Collaboration Occurred

3.7.1 Collaboration throughout the Process

Throughout the project, the teams worked tightly together although the degree to which they collaborated differed depending on the phase of analysis. Once the project started, the team set up weekly meetings. The first thing they had to decide was to specify requirements of the problem, and then they collaboratively worked on building the conceptual model. Whether the team kept using this model or changed to a matrix, it played a role as a representation of their "group thinking," as an analyst described:

> *You want to say that this is the way I'm thinking about this problem. These are some of things I need to think about. And what we've done by building the conceptual model is to have that sort of group interaction, which is not necessarily harmonious action. There can be disagreements about how we should be thinking about this. And if there's shifting, moving it around, that represents an evolution of the way of our thinking.*

Once the team had an idea of the areas to explore, they divided up the work and assigned concepts to each analyst. While each one worked on different concepts, they collaborated in collecting information by using a group library. Although this seems to be loose collaboration, the benefit the team gained was invaluable because it could significantly save time and effort in collection. An analyst explained how they worked in collection using Zotero:

> *Zotero is a good example of one way we collaborate. Each person creates a group library on the Zotero server. If I find a website that I think is useful, whether for my topic or someone else's topic, I add it to our group collection, and then other members can see it before they go searching the Internet for something. And if she doesn't find that in Zotero, then she might go out Google. So..try Zotero first, you might already have it.*

While working on and analyzing their own topics, team members often met with each other to check status and discuss issues. When most of the areas they had planned to explore were covered and analyzed, they collectively wrote the key findings – the crux of the analysis project. Very tight collaboration occurred in this work. They met together and spent significant time to synthesize findings from all the topics and write the key findings.

3.7.2 Sharing vs. Content vs. Function

I found that three different types of collaboration exist when analysts discuss the topic: sharing, content, and function.

Sharing is a way to collaborate by sharing information. In the study, analysts shared sources to better assist their search process and understanding of the topics. At a higher level, however, this can be the sharing of analytical products as well as information sources. This type of collaboration can be significantly supported by technology.

Collaboration also occurs at the **content** level. This type of collaboration, in which analysts work together to create analytic products, can be seen more often in a small-size team. Examples in this project include constructing a conceptual model together, dividing concepts and assigning to each analyst, commenting on each other's analysis, working on ACH together, and writing the key findings together. However, in the study, once work was divided, then each part was done individually. The degree of tightness in this type of collaboration may directly affect the quality of analysis. The more closely the team works together, the more that output is coherent and logical. However, in reality, it is difficult to collaborate on content because of efficiency. This type of collaboration is also difficult to facilitate via technology because so many subtle issues – such as social dynamics, politics, teamwork, and motivations – are involved.

Functional collaboration is needed to execute practical tasks for completing the project, such as editing, creating a matrix structure, specialized analysis on a specific topic, and polishing deliverables. Whereas analysts work on the same thing and divide up the analytic product in the content level, functional collaboration naturally emerges at the later stage of the process as the team begins to think about allocating multiple functions. In this type of collaboration, analysts reinforce their strength. For example, if one is a good editor and has a detailed eye, then that person would do the editing.

3.8 How Visual Analytics Can Help: Design Implications

How can visual analytics help intelligence analysis? Based on the study findings and reflections, I suggest several design implications for systems supporting intelligence analysis.

<u>Externalize the thinking process - Help analysts continuously build a conceptual model</u>

The analysts in the study explained that the process of making sense of a problem and building a conceptual structure is one of the most important parts of intelligence analysis as it decides the direction of analysis. They stated that they often encounter a situation in which they need to learn about new subject matter, but it takes time and effort until they become familiar with the domain. Because they cannot build a good mental model of the problem without knowing what information is available, they struggle to know more about the domain until the later stage of analysis.

Using the power of representation, visual analytics systems can help analysts build a conceptual model or a structure of the problem and domain. For example, the system can take the main question the analyst has and suggest a number of possibly related concepts and keywords based on online encyclopedias, table of contents of books, tagging services, etc. The system should allow the analyst to refine the concepts so that it can repeat the search and suggest other relevant concepts. By connecting, grouping, and organizing concepts, analysts can continuously build up their conceptual model or structure of the area throughout the process. One analyst cited experience:

Ok, I got to model something, I've got to do a report on Ghana, I don't know anything all about Ghana, where's the tool that if I hit the button, it gives me a picture of what the relationship is, the model how to think about Ghana? It gives me 60-70% of the solution. But it gives me the ability to input and tweak and change those. Because I want to have a role in that, I can't allow the computers to do all my thinking, you know.

Support for this externalization should occur throughout the analysis process because as analysts learn more about the domain, they alter their way of thinking and refine their visual model.

Externalizing the thinking process also can assist analysts when they review their analysis after the project terminates. Supporting this activity would be especially useful because it will inform how the analysts could have done better and the areas that need to be examined if they did a similar project, as the instructor said:

The other thing this model helps you 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 that impact our company, our estimates?" Because whatever reason we didn't get to it, this was important, we thought this define the space…We can sit back and go, ok, how confident are we on our estimates, knowing that our analysis is always at some level incomplete? And it's always incomplete, but how does it impact our confidence in our product? That's another way to use this representation.

<u>Support source management - enable managing both pushed and pulled information and organizing sources meaningfully</u>

One prominent characteristic of how analysts think about sources is that they have to be always vigilant of new sources. They often search for the same keywords again to see if any new materials have been added regarding the topic (pulled sources). They also receive news articles through RSS feeds everyday and check if they have received interesting information (pushed sources).

This process of searching sources takes more time than one may think, and systems should allow analysts to manage both pushed and pulled information associated with concepts they have identified. For example, a system could populate several concepts chosen by the analyst and store all the pulled sources in a database such as Zotero. Based on sources already found, the system also could recommend push resources such as blogs and news articles. For each source collected, the analyst could express if it is a useful source or not. Then the list of sources can be organized in a meaningful way – for example, by keyword queries, by tags the analyst annotated, or by date the source was added. The system also could provide several ways of representing source results such as summary and tag clouds. Further support for analysis or visualization of collected sources as a group would be extremely beneficial. Analysts commented on this functionality:

> *Sources are what we have to get, but where is the tool where I can integrate them? My RSS feeds dump into me every morning. But then I do searches as well. Where's the tool that allows me to integrate all data, the information that is useful for me?*

> *If that kind of system exists, I have the ability to go back and find all my sources. Automatically, this (keywords, phrases) gets populated. And every point, I have the ability to say no or yes, no or yes to a source. But the actual extraction or the pulling, and the organization of that is automatic from that.*

Then the list of sources can be organized in a meaningful way – for example, by keyword queries, by tags the analyst annotated, or by date the source was added. The system also could provide several ways of representing source results such as summary and tag clouds. Further support for analysis or visualization of collected sources as a group would be extremely beneficial.

All these technical capabilities currently exist in visual analytics systems. Now it is important that they be integrated together appropriately.

<u>Support analysis with constantly changing information - integrate collection and analysis in a single system and help analysts use structured methods during collection</u>

As described in the previous section, collection and analysis are not separate, but highly integrated processes. Analysts do not wait until all the data are gathered; rather, they start analysis even when they have only a few pieces of information. Through the repeated process of collection and analysis, they revise a frame and use the collected data as supporting evidence for the frame.

Currently, many systems provide analytical support assuming that processed data is available. If a system does not support a seamless transition between collection and analysis, it is likely to be less successful in assisting the analysis. Analysts collect during analysis and they analyze during collection. This differs from statistical analysis, in which a structure or a frame about how to analyze the data is clearly defined and analysis is done with clean dataset. An analyst mentioned:

> *If they had more reliable, structured data, I'd use statistical analysis. But intelligence data is unstructured and dirty. You don't know what the best way to analyze it is until the middle of the process, or even the end of the process.*

Multiple visual analytics systems provide analytical capabilities. By supporting more flexible data manipulation so that analysts can easily import and remove data from the analysis pool, these systems will be more usable, with better integration into the analysis process.

If the processes of collection and analysis are integrated in a single system, this helps analysts apply structured analytic methods such as ACH, social network analysis, geospatial mapping, and decision matrix. In the interviews, two teams mentioned that if they had more time, they would have tried other analytic techniques. Analysts always want to push their findings and triage, aggressively reshuffling their analysis. One of the most effective ways to do this is to employ multiple analytic methods and compare and contrast findings from each. The ability to try various techniques with the data can help analysts find effective ways for addressing questions and strengthening their analysis.

> *We had this time crunch. We pretty much got rid of the process of re-evaluating our hypothesis, finding what's the most important to make it perfect, and hitting on that, and going back to the stuff that we didn't deem as important. If we had time, we would fill that in.*

Help analysts create convincing production – support insight provenance and sanity checks of analytical products

Production is what differentiates intelligence analysis from general sensemaking which does not necessarily entail external representation. Even when analysts finish their analysis, they need to convert the results into a concise format so that decisionmakers can understand their findings. This can be tedious and time-consuming part of the intelligence process.

When asked about the most difficult part of their project, two teams mentioned production. Interestingly, this difficulty comes from sanity checking and insight provenance, not simply from formatting and writing issues. The sanity check, or qualitative double-check, takes time because data and findings are derived from many sources and analysts have meshed them through the process of collection and analysis.

Analysts need to return to original sources and provide a rationale by which their statements are made. They also have to add references to their statements, for which they have to revisit original sources. The following quote from an analyst illustrates those difficulties:

> *Most difficult part…basically going back through all the sources we used to grade these technologies, people, and companies, then taking basic pieces from those and making a narrative out of it. So explaining why we thought they are the keys and then relating it to the rest of the other findings.*

A system that promotes simple insight provenance during analysis could help analysts save their time in production.

<u>Support asynchronous collaboration rather than synchronous collaboration for exploratory analysis</u>

I discussed three different layers of collaboration in the intelligence process and that the degree to which technology can contribute varies. In particular, visual analytics systems seem to have the potential to help collaboration in "sharing" and "content."

From the study, I found that these types of collaboration tend to occur asynchronously, rather than synchronously. When meeting face-to-face, analysts did not work on actual tasks but spent time checking their status, coordinating next steps, and discussing issues. Even when they worked in the same lab for several hours, team members took their own computer and worked individually. Although they often talked to each other, it was for simple coordination issues or specific questions about the content. One analyst stated about his perception on collaboration:

> *We discussed how each of us interprets the data. We're very group-oriented when it comes to discussing to a consensus. Other than that, we prefer to work individually especially for the actual analysis. Of course we collaborate even when we work on our own parts, but there's no one who really knows about those concepts or entities like you do.*

In a nutshell, analysts collaborate cognitively. Rather than trying to build a system that allows analysts to work at the same time in the same workspace, providing a system that promotes individual workspaces but also provides asynchronous collaborative features - such as the ability to share sources and data, view and comment on others' work, and merge individual work together - would appear to be more beneficial.

Note that these findings are based on strategic intelligence. In other types of intelligence such as tactical and operational intelligence, which form the basis for immediate action, real-time collaboration is also important because such intelligence must be shared and used quickly.

<u>Unifying the pieces</u>

Because their typical processes of requirements gathering, collection, analysis, and production are so intertwined, and it takes considerable time to coordinate between different software systems, it appeared to us that analysts want an all-in-one system that can streamline the analysis process and save their time. When asked about their 'dream' system, a few analysts answered:

If I had to go back to the beginning and start all the way over, I should be able to jump back and forth seamlessly between all of these processes. We need a tool that compensates for that.

It should be one program. We spend more time to make it work together. Nothing's compatible with others. We want a program that syncs all the documents. Help us do our visualization with the documents. A program that is compatible with Excel spreadsheet. Don't want to open 20 different programs.

Thus, a hypothetical tool that simplifies the intelligence analysis process would function as follows:

- The analyst enters requirements into the system.
- The system suggests various concepts associated with key terms, phrases, and ideas in the requirements.
- The system automatically draws connections between concepts, but it also allows the analyst to draw connections, group, and organize them.
- The system takes the concepts and starts populating them, collecting information sources using the concepts as keywords (pull sources).
- The system uses sources the analyst identified and suggests new articles relevant to the sources (push sources).
- All of these pulled and pushed sources are integrated into a source repository.
- For documents in the database, the analyst can highlight important facts and annotate his/her thoughts. On demand, the system extracts entities requested by the analyst.
- For intuitive analysis, the analyst can write reports in a preferred format, walking through each document.

- For structured analysis, the system helps the analyst try a variety of structured methods. It takes all the information identified by the analyst and integrates it directly into the methods.
- At the end of the process, when the analyst produces final output, the system automatically links each statement to relevant sources and the process by which the statement was derived.

Thus, analysts could flexibly move between conceptual model, collection, analysis, and production. The system accompanies the analyst from requirements to product in a single platform, speeding up the process, as expressed in one analyst's comment:

> *If I had something like that, I'd be blazingly fast. I mean I would be able to do this 10-week project in three weeks.*

Interestingly, my suggestions reiterate the findings of other researchers who identified the importance of unifying disparate tools in a different domain. In an observational study of the scientific data analysis process, Springmeyer et al [77], concluded that "an effective data analysis environment should provide an integrated set of tools which supports not only visualization, but some of the additional functionality" such as capturing the context of analysis and linking materials from different stages of analysis.

CHAPTER 4

EVALUATING BENEFITS OF VA TOOLS IN AN INVESTIGATIVE SCENARIO: DERIVING DESIGN IMPLICATIONS FROM A COMPARATIVE STUDY

4.1 Introduction

Although many new visual analytics tools are being built to support investigative analysis, few empirical studies that evaluate the potential benefits of such systems have been conducted. Unfortunately, evaluating visual analytics systems for investigative analysis is very challenging and we still do not understand well how to evaluate and assess such systems. Nevertheless, determining how such systems foster insight and sensemaking is important for their continued growth and study. Furthermore, studies that identify how people use such systems and why they benefit (or not) can help inform the design of new systems in this area.

Jigsaw [35, 78] is a system developed by researchers in the Information Interfaces Lab at Georgia Tech. Jigsaw is a visual analytics system for helping analysts who deal with a large amount of documents. It reads in multiple documents in a collection and shows connections between entities and documents by using multiple visualization views. As such, Jigsaw provides a good example of the type of systems that support investigative analysis.

In this lab study, I examined use of Jigsaw in an investigative analysis scenario as compared to three other investigative methods including paper-and-pencil and simple desktop document storage and search. Each participant was given simulated intelligence case reports and asked to identify an embedded terrorist plot within allotted time in one of four conditions.

The primary goal of the study was to better understand how visualization can assist investigative analysis. I wanted to see how people would approach data analysis using a visual analytics system. What characteristics of the system lead to the main benefits? A second goal of this research was to better understand evaluation methodologies for investigative analysis systems in general. What should evaluators count, measure, and observe in order to determine the utility of systems?

Although only a single system was examined in this study, I believe that the findings and implications are still useful and applicable to those who build similar systems. Since I intended to focus on people's strategies and their sensemaking processes under an investigative analysis, rather than the system per se, I expect that developers and researchers could learn more fundamental knowledge and high-level considerations for developing such tools. Suggestions about evaluation methodologies will also benefit researchers who seek to evaluate the utility of systems and to further find design implications.

4.2 Study Design and Analysis Techniques

I evaluated four settings for analysis with one of these using Jigsaw. Sixteen graduated students from Georgia Tech performed an investigation in one of the settings. Each participant was given the same data collection containing 50 plain text documents that simulated intelligence case reports and participants needed to identify an embedded terrorist plot within the allotted 90 minutes.

I told participants that they would be taking on the role of a government intelligence analyst. I gave them 50 documents, described as intelligence reports, and asked the participants to identify a hidden terrorist plot. For this task, I adapted documents from an exercise I had learned about from a military intelligence college. Embedded across some of the documents are hints to a fictional terrorist plot with four sub-stories that support the plot. Each document was a few sentences long. 23 of the documents contained information useful to identifying the threat. The other 27 documents described other suspicious activities but were not relevant to the main plot. Ultimately, participants needed to identify the plot and write a short narrative describing the potential threat. In addition, I gave participants task sheets adapted from the VAST Symposium Contest [34], which contained tables for them to list key players, events, and locations relevant to the plot.

I created four settings in the experiment and assigned each participant to one of the conditions. In setting 1 (Paper), I gave participants the reports as paper documents and asked them to perform the task without any technological aid. In setting 2 (Desktop), I gave participants the documents as separate text files on a computer and made Microsoft Desktop Search available to search for keyword(s) in the documents. In setting 3 (Entity), participants used a limited version of Jigsaw, in which only a modified version of the Document View (tag cloud removed) and text search capability were available. In setting 4 (Jigsaw), participants performed the task using the Jigsaw system. I provided participants in this setting with a short training video of the system three days before the session and gave them an additional 30 minutes of training at the beginning of the session.

In all settings, participants could take notes using pen and paper. I gave each participant 90 minutes to work on the problem and conducted a semi-structured interview after each session. I also video-taped all the sessions. For measuring performance, I created a solution to the exercise and described it in a short text narrative. In addition, I completed the task sheets (relevant people, events, places). Two external raters used this material to grade the anonymized task sheets and debriefings.

Throughout the study, I primarily focused on collecting qualitative data such as observations, follow-up interview notes, and video recordings of all sessions. Wherever possible, I also collected quantifiable data such as the number of documents viewed and the number of queries performed. In order to more closely look at usage patterns of Jigsaw, I logged user interactions and view operations.

During the early exploratory phase of analysis, I used an inductive approach to examine qualitative data. As I read through notes from observations and interviews, potential concepts and categories emerged, including participants' strategies, Jigsaw's influence to the analysis process, and characteristics of sensemaking. These broad categories were refined with specific incidents, anecdotes, and examples from detailed analysis. Detailed info such as video and log data were used as supplements to observation notes when I wanted to further examine and clarify findings. I scrutinized the videos and the log visualization after I identified investigative strategies, in order to verify each participant's process. As the analysis evolved, inductive and deductive analysis were used concurrently because some concepts developed further than others.

4.3 Four Investigative Strategies

Table 2 summarizes the results of the participants by setting. The first block indicates performance results, in which participants in the Jigsaw setting earned excellent, excellent, very good and good ratings. If we average the final scores of the four participants in each setting, those using Jigsaw clearly outdistanced those in the other three settings that produced similar average final scores. The rest blocks explain other activity patterns such as how many of the documents were viewed in total, which document was viewed most, and how many times each document was viewed. I also determined how many search queries a participant performed and when the first query was performed. For those participants who took notes on paper, I identified when they first started note-taking, as well as how many and what kind of notes they took. Additionally, I identified when each participant first began completing the task sheets.

Table 2. Study results and statistics, grouped by setting

	Paper				Desktop				Entity				Jigsaw			
	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14	P15	P16
Grading Task Sheet	-1.75	17	4	13	13	10.5	10.5	-3.5	5.5	-8.25	4	7.5	14.5	13.5	7	17
Grading Debriefing	2	2.5	1	5.5	3	4	1.5	3	3.5	2.5	1.5	6	6	2.5	5.5	5
Final Score	22.87	65.00	24.26	87.08	62.08	67.13	42.13	29.41	52.23	15.00	29.26	81.19	95.05	58.07	75.20	90.00
Performance	fair	very good	fair	excellent	very good	very good	good	fair	good	poor	fair	excellent	excellent	good	very good	excellent
Avg. Score/Setting		49.80				50.19				44.42				79.59		
Documents Viewed	50	50	50	50	50	50	50	50	49	31	45	50	31	50	46	23
Number of Queries					19	18	48	8	23	61	59	91	44	4	26	8
First Query					40:49	19:55	2:47	12:41	1:31	0:29	0:59	3:12	0:18	5:35	25:37	4:18
Amount of Notes	many	none	many	some	many	some	few	some	some	none	none	few	some	few	few	few
First Note Taking	0:07	–	0:05	0:16	1:53	19:57	2:47	8:20	2:37	–	–	3:14	0:48	0:32	5:15	78:48
First Task Sheet	43:20	32:53	70:13	3:25	61:35	20:26	7:33	64:11	28:09	0:52	2:55	7:20	48:26	41:48	43:00	5:33
Strategy Used	OFD	OFD	BFD	OFD	OFD	OFD	FCFT	BFD	BFD	HTK	HTK	FCFT	FCFT	HTK	OFD	FCFT

In addition to these descriptive results, I was interested in more details of how and by what process people performed the analysis because participants exhibited huge individual differences in performance and activity patterns. Although we could simply say that it might be individual differences in analytical capability, I wondered if there exist any patterns regarding investigative strategies used by each participant.

After examining each participant's process and strategy by videos and interview scripts, I identified four general investigative strategies being used by the participants, independent of the setting they were in. The following subsections describe each strategy more in detail.

4.3.1 Strategy 1: Overview, Filter, and Detail (OFD)

The most commonly used strategy was "Overview first, filter and select, and elaborate on details," a strategy quite similar to Shneiderman's InfoVis mantra [74]. Six participants out of 16 performed analysis using this strategy. They began by quickly scanning all documents and building up initial ideas of the plot. After scanning all documents, they revisited relevant documents selectively - either by directly looking up the document or by searching for a keyword that stood out. Then they read each one carefully, extracting key information for the task sheets. I speculate that this strategy worked well in this task because the dataset was relatively small. Participants were able to gain an initial idea of the important documents or keywords by simply scanning all documents although they sometimes missed important details.

4.3.2 Strategy 2: Build from Detail (BFD)

The strategy, “Build from detail”, contrasts the previous one. Three participants used this strategy. They started the analysis from details of each document by carefully reading it. Even though they used the search function when important phrases or words arose (where applicable), it was more of an auxiliary use than a main focus. They issued relatively few queries. Instead, they focused on every sentence of the documents, in the fear of missing any relevant information.

Because they paid attention to every detail, it was difficult for them to see the “big picture” of the plot, and therefore this strategy turned out to be less effective than other strategies.

4.3.3 Strategy 3: Hit the Keyword (HTK)

Some participants used an unexpected strategy - an intensive keyword-based exploration. They did not begin the analysis by reading a specific document, but directly looked for a few specific keywords such as “terrorist” or “Al-Qaeda”. They read only the related documents and then searched for other terms that emerged during that time. Since the effectiveness of this strategy depended on the appropriateness of the terms chosen in the initial stage, performance varied across participants using this strategy. In fact, people who used this strategy gained the poorest score in each group where it was used, and I assume that it is because they were too much focusing on specific terms, rather than trying to connect the dots by reading documents.

4.3.4 Strategy 4: Find a Clue, Follow the Trail (FCFT)

The “Find a clue, follow the trail” strategy is a hybrid approach of the previous strategies, and four participants followed it. They invested some time in reading the first few documents to understand the context and find a clue, then followed the trail rigorously using search or other functionalities provided by the tool. In theory, this may be a good strategy because the analyst's attention is focused on relevant documents only. The initial investment in reading a few documents pays off because it increases the possibility of finding the right clue. The performance of participants who used this strategy is notably good.

4.4 Jigsaw's Influence on Investigative Analysis

Among the four study conditions, the group using Jigsaw generally outperformed the other groups on the whole. Based on observations, interviews, videos, and log analyses, I identified several benefits Jigsaw seemingly provided to users. While these benefits are based on people's interaction with Jigsaw, some of them are potentially generalizable and useful for other similar systems.

- *Supporting Different Strategies.* Examining each participant's analysis process, I note that the four Jigsaw setting individuals used three different strategies. This suggests that Jigsaw supported different analysis strategies well.
- *Showing Connections between Entities*. Showing connections between entities such as people, organizations, and places was one of the benefits. While participants in the non-Jigsaw settings wanted to see comprehensive connections between entities and tried to draw connections on paper, people in the Jigsaw setting focused on the challenges in organizing and keeping track of relevant information.
- *Helping Users Find a Right Clue*. Finding an appropriate clue early in the analysis is crucial and sometimes even determines the entire performance. Even though the dataset used in this study was relatively small, participants still benefited from Jigsaw's functionality in finding a good starting point.
- *Helping Users Focus on Essential Information*. Even though analysts may find appropriate initial clues, it is still important to follow the trails in an efficient manner. If relatively unimportant information diverts their attention, the

investigative process may suffer no matter how quickly a good clue was discovered. I found that Jigsaw helped participants to follow the right trail and ignore irrelevant documents, thereby saving the participant's attention for important information.

4.5 Observations on Sensemaking

I did observations on sensemaking during the study and identified several findings in relation to Pirolli and Card's Think Loop Model of Sensemaking [60].

4.5.1 Diversity in Sensemaking Processes

While the model is not linear and can proceed top-down or bottom-up with many loops, I found that the sequence of analysis significantly differed across individuals even in the same task with the same dataset. Some participants followed the sequence linearly with iteration while some participants skipped certain processes. Other participants immediately started from a hypothesis without the schema stage, and then worked on organizing to confirm the hypothesis.

Individual differences also existed in each stage of the model. For example, the "read & extract" stage, in which evidence files are collected from the shoebox, exhibited individual differences. When encountering much unfamiliar information, it is not easy to extract nuggets of evidence simply by reading documents; the analyst usually needs some criteria to decide what to pull out. In the study, some participants started from a specific set of people and extracted information related to those people. Those who used location as a criterion gathered all information related to specific cities or countries. Participants also extracted evidence files based on specific events such as arms thefts or truck rentals.

4.5.2 Power of Schematizing

It was the schematize (organizing) stage that showed the most significant variance between individuals. Schematizing is the re-representation or organized marshalling of the evidence file so that it can be used more easily to draw conclusions. During this stage, it seemed that each person had his/her own preferred organizational scheme such as a timeline, map, or diagram. For example, while most people wanted a timeline, the representations they envisioned were all different. Some people wanted a timeline organized by person and event; some wanted a timeline by location; others wanted a timeline categorized by story. The variances in this stage seemed to affect the entire analysis performance.

The time at which a participant first reached the schematize stage and how much effort the participant invested in this stage significantly affected the performance. When I further examined those who performed well independent of the setting, I found a commonality that all of these people spent considerable time and effort in organizing information. Most people used the task sheet as a tool for gathering their thoughts since the task sheet was structured by certain schemes (e.g., people, events, and locations).

4.5.3 Insight Acquisition

It is still difficult for us to identify exactly when people gained a key insight during the investigative process. When I asked the participants how they knew they were progressing towards the goals, the common answer was "when the pieces of a puzzle started being connected and put together." Rather than a spontaneous insight occurring (the "light bulb going on"), insight seemed to form continuously throughout the investigation, not unlike that described by Chang at al. [10]. Participants had difficulties identifying when they "got" the plot.

4.6 Implications for Design of Investigative Analysis Tools

The study and its results suggest several design implications for visual analytics systems for investigative analysis.

- *Facilitate Clue-Finding.* Study participants who employed the "find a clue, follow the trail" analysis strategy generally performed well overall. Thus, investigative analysis tools that support analysts in finding appropriate starting points or clues, and then, following the trail of these clues efficiently could be beneficial. Further, the performance of those participants who were able to focus only on relevant documents was outstanding. Investigative analysis tools should help direct the analyst's attention to the most critical information.
- *Support for the "schematize" Stage.* The study demonstrated that people do frequently move between stages of the Think Loop Model, particularly in the middle parts of the model. Investigative analysis tools should allow smooth transitions between the "shoebox," "evidence file," and "schema" stages so that different sequences of the sensemaking process can be supported. Currently, the focus of Jigsaw is on the "shoebox" and the "evidence file" stages, but it lacks powerful support for the "schematize" stage. While Jigsaw does appear to help analysts finding nuggets of information effectively, it does not really support putting those pieces of evidence together. In other words, analysts may easily discover the pieces to be put in a puzzle and have a sense of which piece goes where, but they should also receive help in putting the pieces together. The ability to work on extracting evidence files and organizing them into a schema will significantly help the sensemaking process.
- *Support Evidence Marshalling.* For Jigsaw to be a comprehensive investigative analysis tool, it is crucial for the system to include a workspace in which the analyst can simply drop/paste entities, draw connections between them, and add annotations, capabilities found in systems such as Analyst's Notebook [33], the Sandbox [87], and Entity Workspace [4].
- *Allow Flexibility in Organizing.* When supporting the "schematize" stage, developers of investigative analysis tools should consider that individuals will

choose different organizational metaphors or schemes. For example, even for a timeline, individuals imagined many different types of timelines and they were quite insistent about this approach. Rather than providing one fixed schema, allowing flexibility and room for customization will be beneficial. Tool developers may consider having a system suggest a few organizational schemes when the analyst has created a significant evidence file but still does not have a schema, particularly for novice analysts. Staying too long at the "evidence file" stage appears to impede the analysis process, so suggestions of organizational schemes may be beneficial.

- *Suggest Alternative Paths but Support Task Resumption.* It is not uncommon for an analyst to confront a dead-end or find evidence that refutes an existing hypothesis. Investigative analysis tools need to support the analyst to find appropriate next steps or alternatives by making the milestones of the investigative process explicit. In this way, the analyst can come back to the point where she/he was earlier and start over from that point. This also ensures that the analyst can proceed further without being too concerned about keeping track of past states.

4.7 Implications for Evaluation of Investigative Analysis Tools

The study also suggested a number of ways to help evaluate investigative analysis systems. By comparing system usage to more traditional methods but otherwise giving participants freedom to perform as they wished, I feel that the findings are both realistic and provide ample grounds for contextual analysis and comparison.

I also suggest that the evaluation of investigative analysis tools focuses on collecting more qualitative data. While quantitative data are useful when a solution is well defined and measurable, the nature of investigative analysis is exploratory and flexible. It may be too limiting to assess the value of a system solely based on statistical results. Identifying best practices supported, particular pain points, and future design requirements can be better achieved through interviews and observations. When possible, I suggest using quantitative data such as usage log files and analysis scores to help understand qualitative results.

Findings from the study suggest potential questions to be answered in the evaluation of investigative analysis tools:

- Does the tool help to provide information scent appropriately, thus helping to find initial clues?
- Does it guide the analyst to follow a trail, without distraction?
- Does it support different strategies (sequences) for the sensemaking process? That is, does it support smooth transitions between different stages of the model?
- Does it help to find appropriate next steps when encountering a dead-end?
- Does it facilitate further exploration?

In this study, I identified and used several metrics, which are broadly applicable to evaluation of investigative analysis tools:

- The number of important documents viewed, relative to the entire collection
- When the analyst first started creating representations such as notes and drawings
- The quantity of representations created.

I also suggest two possible metrics for evaluating investigative analysis tools:

- amount of time and effort in organizing and
- amount of time the analyst spent in reading/processing essential information

CHAPTER 5

EXAMINING THE USE OF VISUAL ANALYTICS SYSTEMS FOR SENSEMAKING TASKS: CASE STUDIES WITH DOMAIN EXPERTS

While the previous evaluation study provided useful implications for visual analytics tools, I believe that an in-depth case study with domain experts would provide another valuable perspective on designing such tools. In the study, I aim to examine how domain experts use a visual analytics system for their own tasks in real world settings. In this ongoing evaluation, I explore the practical applications of Jigsaw, the visual analytics system developed by my lab colleagues. It is my hope that the anecdotal findings from this study will open up meaningful discussions for visual analytics researchers and inform design decisions, and in this way complement the findings from the laboratory study.

5.1 Introduction

In the field of visual analytics, the evaluation of systems is important but rare, probably because it is so challenging [61]. Particularly rare are actual case studies of prolonged visual analytics system use by analysts working in their domain with their own data. Case studies can provide valuable findings and insights for visual analytics researchers. By detailing the use of a system, case studies yield a description of how a tool was used and where the users had problems. Until their particular challenges are understood, it also remains difficult to know how a visual analytic system helps expert users attain their goals. These findings are difficult to achieve through controlled lab studies.

Conducting case studies is challenging, however. First of all, it can be difficult to recruit appropriate people who are willing and able to use a particular system for their task on a regular basis. Case studies also often involve issues in the reliability, validity, and generalizability of results although these issues can be mitigated by scaling up the number of users. Nevertheless, it seems valuable to study domain experts working on complex problems over long time periods and learn how they employ systems.

In this study, I profile six investigators who have been using Jigsaw in their own work. The goals of this research include the following:

- To evaluate whether Jigsaw is helping analysts with their tasks and problems
- To understand its applicability to different types of documents and analyses
- To identify particularly useful features and capabilities of the system as well as missing or problematic ones
- To reflect on usage to inform the design of next generation tools for investigative analysis.

5.2 Recruitment and Study Protocol

Jigsaw is not publicly available in general, but we distribute the system upon request. Approximately 150 people from a variety of domains including academics, government, law enforcement, intelligence, reporting, and fraud investigation have downloaded the system. However, I believe far fewer have used it extensively. I selected six analysts that we knew were using the system based on questions that they had sent to the team about it in email. I asked if they would agree to tell about their use of the system, and all agreed to conduct an interview and share their experiences.

The professionals include three intelligence analysts, two academic researchers, and one business analyst. They sought out Jigsaw after facing challenges in their own work and have been using Jigsaw for a range of 2 – 14 months. I conducted semi-structured interviews with each; two interviews were conducted face-to-face, and the other four were conducted over the phone.

Each interview lasted for about 45-60 minutes. While each was a semi-structured interview, I had a set of planned questions to make sure to cover particular important topics such as (1) What kind of tasks, data, and documents they used Jigsaw for, (2) To what extent and how Jigsaw helped their work compared to existing ways and methods, (3) What features were most/least useful, and (4) What barriers they encountered while using Jigsaw and how the tool can be improved. Sample questions are included in the following:

- For which tasks have you used Jigsaw? What kinds of documents are involved?
- What is the main purpose of using Jigsaw in analyzing those documents? What do you want to accomplish?
- Before using Jigsaw, how did you perform the tasks? What are advantages and disadvantages of the method?
- How you typically work with Jigsaw and the documents?
- Which features do you use most? How does each of those features assist your task?
- What barriers did you encounter while using Jigsaw? How did you address them?
- How your usage has changed/evolved over time?
- What kind of features do you want to see in Jigsaw in the future?

While I took some notes during the interview sessions, all conversations were audio-recorded for further analysis. I also collected several screenshots whenever possible. The interviews were transcribed and analyzed using a general qualitative analysis technique, borrowing principles from grounded theory [80]. After skimming through the transcribed texts, I conducted an open coding, in which I tried to identify and categorize phenomena. I read each sentence and paragraph, and lable them in order to find out core themes and categories. Then I began relate categories and themes to each other in order to generate a refined set of categories, which is axial coding. Then I carefully re-examined the transcript to find data that fits each categories. Since the interview guide already had core concepts and themes, I focused more on disentangling phenomena and relationships behind users' experience with the tool. Although I analyzed the data in a manner similar to that in grounded theory, I did not follow the last step of the approach because I did not intend to identify a single storyline from the findings. During the analysis, I also exchanged emails with participants as pertinent follow-up questions arose.

5.3 Case Studies

Throughout these studies, I found that professionals have unique goals and consequently, different use cases of Jigsaw. This section describes each individual's particular background, objectives, and how they used the system.

5.3.1 P1: Aerospace Engineering Researcher

P1 is an Aerospace Engineering researcher at Georgia Tech working on aerospace systems design. She was examining two air traffic control-related initiatives - the Next Generation Air Transportation System (NextGen) by the United States and The Single European Sky ATM Research (SESAR) by the European Union (SESAR). The two initiatives consist of new concepts, capabilities, and implementation plans over the next decade, pursuing more efficient air traffic management.

While the objectives of SESAR and NextGen are similar, a number of differences exist between the two initiatives. In her field, the need for harmonization between the two has been recognized, and she wanted to analyze to what extent the two initiatives are compatible with each other. Particularly, she wanted to compare similarities and differences between the two initiatives – if a concept or capability suggested in one initiative also appears in the other initiative, and if so, how each initiative describes the same concept. In order to do that, she needed to examine components, roadmaps, terminologies, and definitions in each initiative thoroughly. Each initiative has seven huge volumes of documentation, and each volume has dozens of documents in it. Each initiative has hundreds of Operational Improvements (OI) and enablers such as policy, technology, and procedures. Figure 19 illustrates one example of the many Operational Improvements in NextGen.

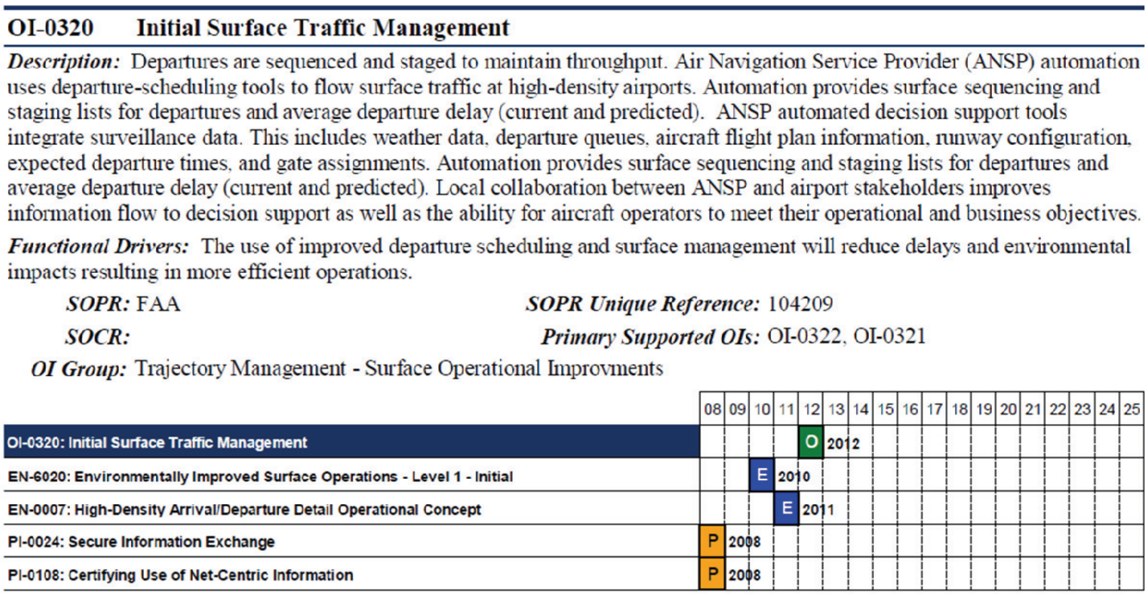

OI-0320 Initial Surface Traffic Management

Description: Departures are sequenced and staged to maintain throughput. Air Navigation Service Provider (ANSP) automation uses departure-scheduling tools to flow surface traffic at high-density airports. Automation provides surface sequencing and staging lists for departures and average departure delay (current and predicted). ANSP automated decision support tools integrate surveillance data. This includes weather data, departure queues, aircraft flight plan information, runway configuration, expected departure times, and gate assignments. Automation provides surface sequencing and staging lists for departures and average departure delay (current and predicted). Local collaboration between ANSP and airport stakeholders improves information flow to decision support as well as the ability for aircraft operators to meet their operational and business objectives.

Functional Drivers: The use of improved departure scheduling and surface management will reduce delays and environmental impacts resulting in more efficient operations.

SOPR: FAA

SOPR Unique Reference: 104209

SOCR:

Primary Supported OIs: OI-0322, OI-0321

OI Group: Trajectory Management - Surface Operational Improvments

	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
OI-0320: Initial Surface Traffic Management					O	2012												
EN-6020: Environmentally Improved Surface Operations - Level 1 - Initial			E	2010														
EN-0007: High-Density Arrival/Departure Detail Operational Concept				E	2011													
PI-0024: Secure Information Exchange	P	2008																
PI-0108: Certifying Use of Net-Centric Information	P	2008																

Figure 19. Original document: Operational Improvements-0320 in NextGen initiative

Her goal in this project was to create a mapping between the two initiatives by identifying similarities and relationships between operational improvements of each (Table 3).

Table 3. Example of a mapping between NextGen and SESAR improvements

NextGen Operational Improvements	SESAR Operational Improvements and Steps
OI-0320: Initial Surface Traffic	**L07-02 TS-0201:** Basic Departure Management (DMAN)
	L07-02 TS-0202: Departure Management Synchronized with Pre-Departure Sequencing
	L07-02 TS-0203: Integration of Surface Management Constraint into Departure Management
	L07-02 TS-0306: Optimized Departure Management in the Queue Management Process
	L10-02 AO-0205: Automated Assistance to Controller for Surface Movement Planning and Routing
	L10-03 AO-0501: Improved Operations in Adverse Conditions through Airport Collaborative Decision Making
	L10-03 AO-0602: Collaborative Pre Departure Sequencing

Originally, the comparison was done manually using Microsoft Word and a search function. That is, she searched for descriptions of NextGen and identified keywords, and then she reviewed descriptions of SESAR containing matching keywords one by one, which was lengthy and cumbersome. Given the high number of descriptions and concepts, it became increasingly difficult to form a clear understanding of the underlying relationships and similarities between the two initiatives. At that point, she searched for a more analytically efficient way of reviewing the information and found Jigsaw.

In order to import documents into Jigsaw, she modified the original document (Figure 19) so that it could be readable. After importing documents, she created entity types including 'Title," "Initial Operational Capability (IOC) indicators," "focuses," "benefits," as well as "the procedures, concepts and systems," relevant to each operational improvement (Table 4).

Table 4. Modified document: description and entities associated with the NextGen Operational Improvement OI-0320

Description	Departures are sequenced and staged to maintain throughput. Air Navigation Service Provider (ANSP) automation uses departure-scheduling tools to flow surface traffic at high-density airports. Automation provides surface sequencing and staging lists for departures and average departure delay (current and predicted). ANSP automated decision support tools integrate surveillance data. This includes weather data, departure queues, aircraft flight plan information, runway configuration, expected departure times, and gate assignments. Automation provides surface sequencing and staging lists for departures and average departure delay (current and predicted). Local collaboration between ANSP and airport stakeholders improves information flow to decision support as well as the ability for aircraft operators to meet their operational and business objectives.
Title	Initial Surface Traffic
IOC Indicators	2012
Focuses	Surface Surface sequencing Departure sequencing High-density airports
Phases	Surface Departure
Procedures, Concepts and Systems	Departure-scheduling tools Decision support tools Surveillance data Automation

She performed the analysis mainly using the Graph View and the List View. In the Graph view, she searched for any OI of interest and the document associated with the OI appeared as a node. She then further expanded the node to reveal the different entities relevant to the document of interest. After filtering out all but the "Focus" entities and expanding all nodes, all connections between relevant documents are represented. An example of a Graph View representation resulting from querying one of the NextGen OI is shown in Figure 20.

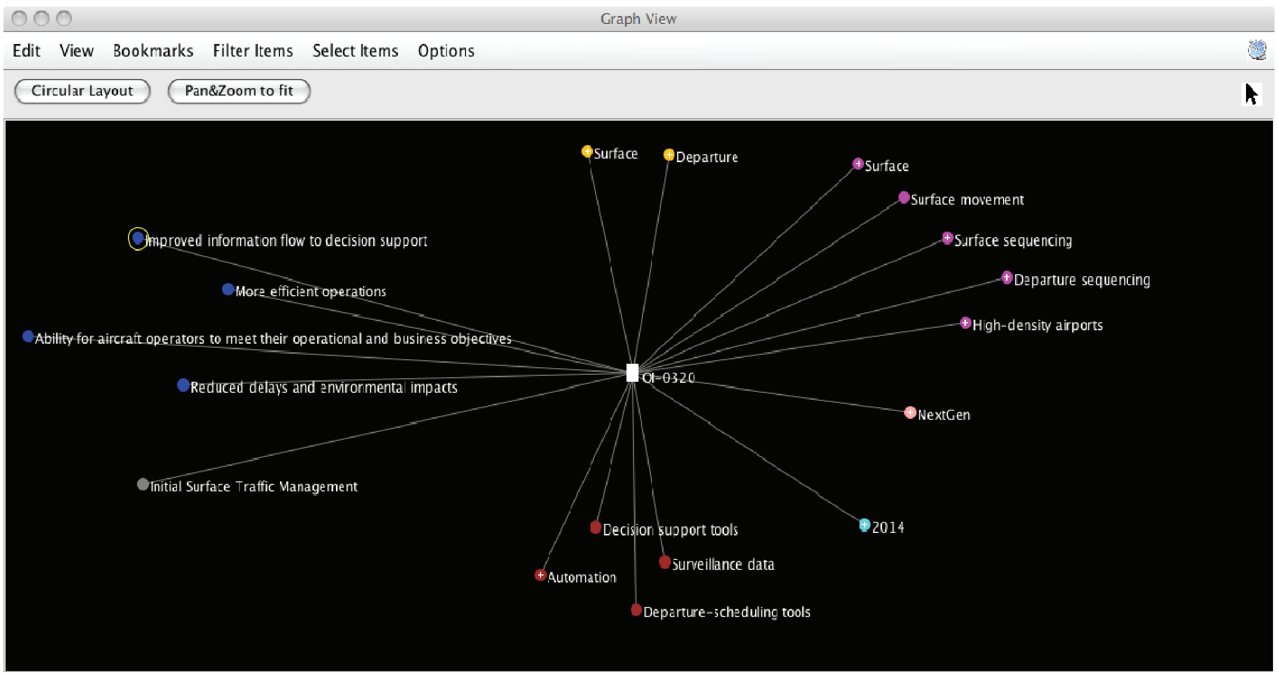

Figure 20. A document and its relevant entities in the Graph View

She used the List View to obtain similar connections, as illustrated in Figure 21. She set the first, second, and third columns to display the document's title, focus, and ID number, respectively. Then she selected the title corresponding to one of the NextGen OIs so that Jigsaw can provide a list of focuses associated with the document in the second column. By further selecting the different focuses of that particular OI, she could also see the ID number of all other relevant documents in SESAR.

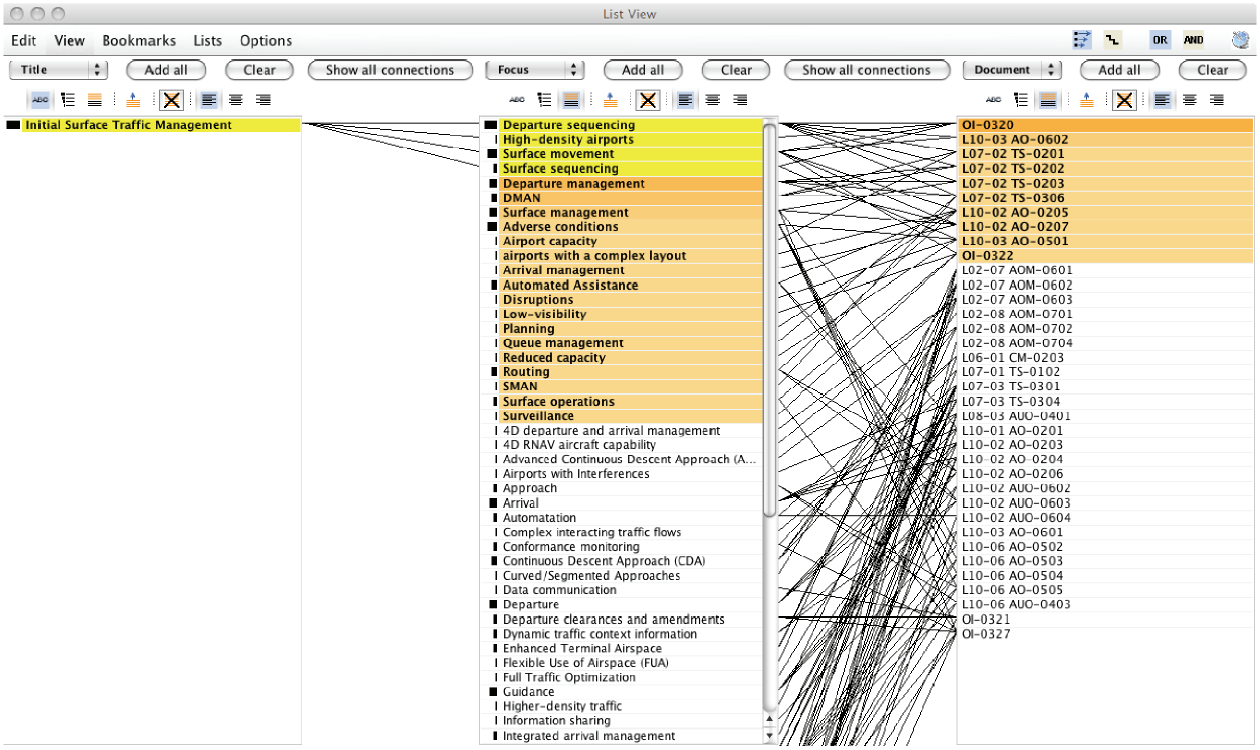

Figure 21. The relationships between NextGen OI-0320 and relevant SESAR Operational Improvements

Following this process, she was able to map all NextGen OIs to focus-related SESAR OIs more efficiently, which would have been impossible with the manual approach. She pointed out that the complexity of each initiative's structure made it difficult to rigorously investigate similarities and differences between the two, but using visual analytics, she was able to review and analyze the information in an efficient way. Her work using Jigsaw led to a publication in American Institute of Aeronautics and Astronautics [58]. After the research was done, she and her team continued to using Jigsaw for another research project.

In addition to current functionalities in Jigsaw, she wished compatibility with other query databases so that she could import documents directly from other databases and statistical capabilities that can count and measure connection strength.

5.3.2 P2: Business Analyst at Management Services

P2 is an analyst of an accounting firm in Malaysia. While his company provides a variety of services related to business management, he specializes in financial fraud and forensic investigation. He usually receives large amounts of both structured and unstructured data from his client. While his team has several tools that can effectively analyze structured data such as transactional data, they did not have an appropriate tool that can help analyze unstructured data such as emails or text files.

His main task is to examine unstructured data from financial databases of clients and to identify any linkages between people or companies relevant to financial fraud such as fictitious suppliers' invoices (i.e., bloated expenses to minimise tax), systematic deletion of suppliers' invoices, or fictitious customers' invoices to boost revenue. Before using Jigsaw, he would put all of the text documents into an Excel database, search for specific keywords within the database, and start investigation by reading all the returned documents containing that keyword. Obviously, this process required manpower to make the database and make it searchable with appropriate keywords.

He had been using Jigsaw for about 14 months. First he converts all documents into text files and imports them into Jigsaw. Then he identifies entities such as organizations, people, dates, locations, description, and zip codes. He starts an investigation with the Wordtree View, in which he searches for names of interest, for example "ABC" company, simply to see the context of the person or company. Next he examines the connections more carefully in the List View to observe what documents link the two people/companies together and who is connected most. Once he sees a potential connection between an entity and a company, he searches for the company and further investigates if other entities are linked to the company. He sometimes uses the Document Cluster or Timeline View to check the amount of documents within a certain topic or time frame.

Through this repetitive process, he can reveal connections between entities and use it as evidence for financial fraud. Thus, Jigsaw provides support for his task by making it easier to find linkages between entities in emails. In one case involving 4.5 million transactions, his team identified approximately 100,000 transactions as fictitious supplier invoices over a period of 10 years using data mining software. They suspected "John Doe" as being the prime culprit, but they needed evidence for that. They asked the HR personnel to seize his notebook and cloned his hard disk drive. After indexing all the documents on his notebook, they imported about 100,000 emails from the past 10 years into Jigsaw to find the motivation for the fraud. After analyzing the documents, they finally found that the theft of funds occurred because the suspect needed to support his children's education costs overseas.

Because his data comes as different formats such as pdfs, docs, and emails, he wanted to be able to import documents directly into Jigsaw instead of having to convert them to .txt format manually (Jigsaw's import of pdf and MS word files is sometimes problematic). Since he mainly looks for evidence, he also seeks the ability to statistically compute closeness or correlation between connections.

5.3.3 P3: PhD Candidate in Industrial and Systems Engineering

P3 is a PhD student at the School of Industrial and Systems Engineering at Georgia Tech. Her research is about enterprise transformation, in which she tries to build mathematical models of how firms would evolve over the years. In her earlier research, she formulated mathematical models about company transformation, and now, she wants to validate the models by combining them with historical data of several companies. The company data, which includes 5,000+ company announcements and news articles of nine IT companies for 10 years, contains critical information about firms such as new product releases, executive/board changes, business expansion, strategic alliances, etc. By measuring how often those events have occurred in the past, she seeks to combine this quantitative information with her model to see if the model is valid or not. That is, she is ultimately trying to transform qualitative information about the IT companies into a quantitative form that can be incorporated into her model. But she was in her initial stage of the research, and she first wanted to understand the documents and generate keywords based on the understanding of the documents for the following step – data mining. After actively searching for software, she decided to use Jigsaw for her research.

Again, her goal in using Jigsaw is to obtain an overview of the huge document collection and extract keywords from those documents. Her documents were stored in Excel spreadsheets, which is an appropriate format for Jigsaw. She added entity types such as event type, company name, capitalIQ, and date, so that she could understand key events of each company. She also created libraries for some entities. For example, she created a list of words for "Business Expansion" entity (Figure 22). While she tried all the system views, she ended up using two views: The List View (most frequently, Figure 23) and the Calendar View when she was focusing on a specific time period, e.g., if something is occurring in a certain period.

```
launch
launches
launched
unveil
unveils
unveiled
test
tests
tested
testing
open
opens
opened
opening
close
closed
announce
announces
announced
expand
expands
expanded
expanding
expansion
```

Figure 22. List of words created for "Business Expansion" entity

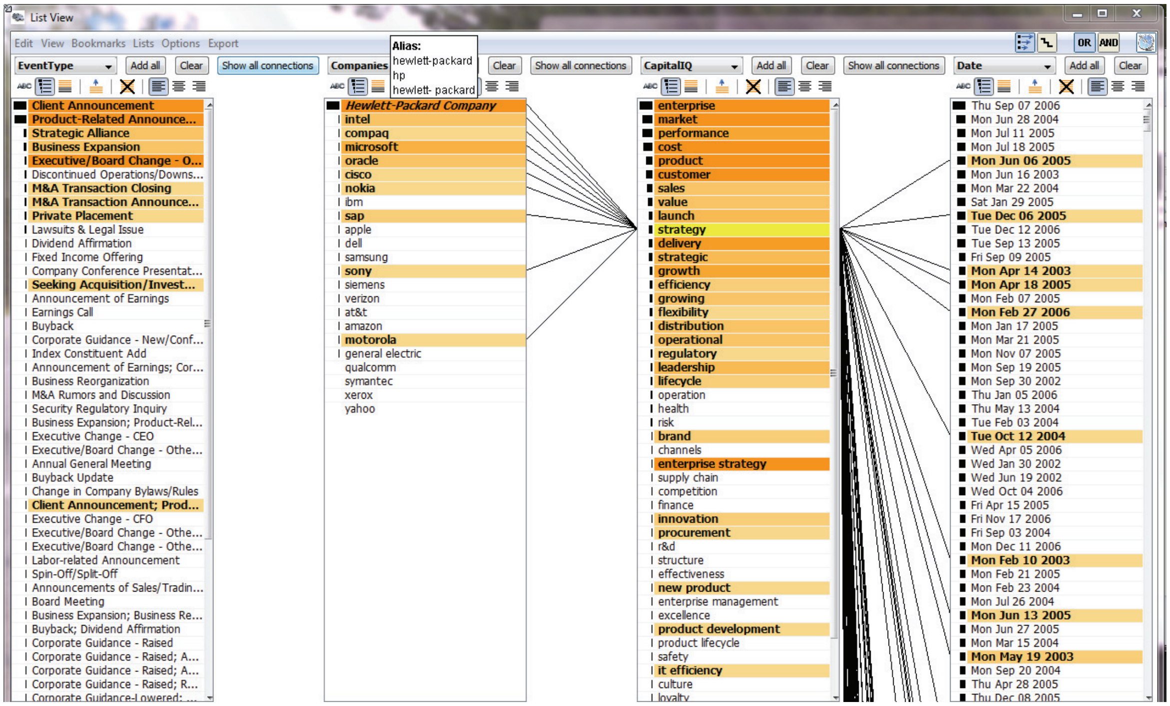

Figure 23. List View displaying event types, companies, and dates related to "strategy"

While Jigsaw helps her research primarily at the initial stage, she thinks that it is very helpful in making sense of the documents in a relatively short amount of time. For a more detailed analysis and final output, she is using other software such as Northern Light and statistical tools in conjunction with Jigsaw:

> *Jigsaw is to me for understanding. So if I need to talk to my advisor about something, I'd go back to Jigsaw and then import some documents and then I can talk what's really going on in this company. For more output formats, I need to do statistical analysis and use other software that have better output format. For this particular project, I'm using more than 5 software (tools) including Jigsaw.*

One difficulty she encountered was working with entities because the system did not identify those she really wanted. She did not find the people and organizations identified by the system very helpful and had to create her own lists. Because her purpose of using text analytics software was to finally create a statistical analysis instead of getting to know about the data in detail, she sought more functionality in terms of output such as a timeline table or word count results.

5.3.4 P4: Intelligence Analyst at a Police Department

P4 is an intelligence analyst at a police department in a city of close to 70,000 population. His work includes making sense of incident/crime reports everyday and discovering patterns, trends, and any top issues in the city. Particularly, he seeks to make better connections between individuals and other information collected in the incident reports. Because the amount of reports increases day by day, he has been trying to find ways to better analyze the narrative text data from the incident reports in their records management system.

Before using Jigsaw, he did not have any ways to systematically work with the information. Basically he could not do anything but read and remember. He read all the reports individually and tried to remember different connections between people, and then recognized names and locations that were outstanding. In order to know who is connected to whom, and in what documents, he printed a copy of the documents, put all the printed reports together, and tried to see the relationships.

When he discovered Jigsaw, he found it very helpful because he wanted to connect people, narrative text, subjects, and concepts in the same system. His goal in using Jigsaw was therefore to make sense of the crime reports and to find connections, patterns, trends, and associated names/places/other incidents.

In order to work with Jigsaw, he reads in crime reports and puts them into an Excel spreadsheet, in which he adds labels for each column such as “Case number,” “Name,” “Person involved,” “Incident address,” “Home address,” “Report date,” and “Description.” Then he imports the spreadsheet into Jigsaw and starts the investigation.

He mainly uses the List View, Graph View, and Document View. In the List view, he normally employs several lists such as persons, addresses and crime types, and conducts a search on a person and examines what addresses and crimes they are connected with. He sorts the lists by connection strength to get a quick sense of relationships between persons, addresses, and/or crimes. He also likes the visual aspect of Graph View in that he can look at connections through link analysis. He generally starts with one individual and then expands out from that person to see what documents, individuals, and addresses that person is connected to. While many times he starts with one person, after expansion of entities, he starts looking at other individuals and their relationships. In the Document View, he takes the information from the List View – such as a suspect or victim, selects the person, and then reads all the crime reports that person has been involved in and looks for any patterns or trends related to that person, involving crimes. Sometimes he uses the Calendar View by selecting an individual in the List View. Then he finds a strong connection with another individual and proceeds to look at those two individuals together in the Calendar View, in order to identify when they are associated with each other, on what dates.

Although he is relatively familiar with the features and functionalities of Jigsaw, his way of using it is still inconsistent. He is still experimenting how to more effectively analyze his data using Jigsaw.

He has already experienced the utility of Jigsaw in his work by helping the police to arrest a criminal. The police were trying to find a criminal, and he searched for the name of another related person in the document collection and examined connections between the two, finally identifying an address where the criminal might be. He liked both visual and investigative support by Jigsaw:

> *I think Jigsaw's strength is its visual support, and investigative support. It would have been impossible without it... When I showed the results and connections to other colleagues, it was easy for them to understand how a certain person is connected to others, that is, providing the context.*

One of the issues he has encountered is figuring out how much data he should import. If he imports documents from the last two years, it would be easier for him to see long-term trends and links between associates. However, this will take significant time to import the documents and clean up the entities. If he imports documents of only several months, it will be faster to import and handle, but he will be able to see short-term trends only. Considering the tradeoff, he has to spend significant time finding out the optimal point.

5.3.5 P5: Intelligence Analyst at a National Lab

P5 is an intelligence analyst at a national laboratory. His department receives a number of resumes for post-docs and researchers who are applying to the lab throughout a year. Among the applicants, he is interested in finding someone who has expertise in a specific area, and being an intelligence analyst, he utilizes his analysis skills in finding candidates. To identify who has the specialty the laboratory requires, he looks at not only the technology/specialization an applicant explicitly expressed, but also publications, co-authors and collaborators, and previous institutions of an applicant.

Before using Jigsaw, he performed the task using Analyst's Notebook, which he felt was limited because he had to manually type in all the data in resumes to the Analyst's Notebook and create connections:

> *I had to do it one at a time and tie them together manually, really. I mean I was using Analyst's Notebook, but pretty much you have to put the data by yourself. There's not a lot of ways to pull in data, so it's really a lot of work, especially when there's a lot of resumes in our system.*

When he was introduced to Jigsaw, he found that it might be a good fit for his task – finding connections between people and technologies (specialties). In Jigsaw, he looks for entities such as institutions, organizations, technologies (specific types of technologies), publications, co-publications, employment history, dates, and emails, and he always creates those entity types. Especially, he tries to find who is connected with whom within a community. By investigating the connections, he ultimately seeks to find an expert in a specialized area, for example, an energy expert.

Working with resumes, he found the Document View really helpful. Interestingly, he uses the view for “identifying what views to use,” as well as for simply reading the documents. He first reads a couple of documents in Document View and determines which other views would be appropriate and effective for analyzing those documents. That is, by getting a brief overview of what each document looks like, he decides which views to utilize for investigation. Among other views, the List View helps him clearly visualize who is connected to what technology or organization. Particularly, the view is useful when an applicant does not explicitly mention a certain technology as specialization but still has background or experience relevant to the technology in the past. Using the List View, P5 could see possible connections and find a good candidate who is knowledgeable about a technology, which would have been much more difficult otherwise. He also often uses the Document Cluster View when he wants to see how the documents can be categorized. He then would select a specific document cluster to read some of the documents in that category.

He mentioned that entity identification and being able to focus on the inter-connectedness of ideas between people and technologies were especially beneficial. Due to these features of Jigsaw, the process of investigating resumes has become more efficient and effective, as it helps him bring connected people together that he might not have been able to see otherwise.

While entity identification is a benefit, it also seemed to be a barrier to him. Because Jigsaw does not always recognize all the entities as he wants, he has to go through the documents and clean up entities after the initial import.

5.3.6 P6: Intelligence Analyst in the Air Force

P6 is an intelligence analyst in the Air Force, in which his team examined the Research and Development Descriptive Summaries, which are budget documents for R&D programs in the Department of Defense [63]. It is a large document collection (>10,000) from 20+ agencies such as Airforce, Navy, DARPA, etc., each of which contains one-page budget summary including description and justification (Figure 24).

ARMY RDT&E BUDGET ITEM JUSTIFICATION (R2 Exhibit) — **February 2008**

BUDGET ACTIVITY	PE NUMBER AND TITLE	PROJECT
5 - System Development and Demonstration	**0604633A - AIR TRAFFIC CONTROL**	**586**

COST (In Thousands)	FY 2007 Estimate	FY 2008 Estimate	FY 2009 Estimate	FY 2010 Estimate	FY 2011 Estimate	FY 2012 Estimate	FY 2013 Estimate	Cost to Complete	Total Cost
586 AIR TRAFFIC CONTROL	7877	8899	14214	2717	4844	7102	6610	Continuing	Continuing

A. Mission Description and Budget Item Justification: This program element funds continuous efforts in the development of modernized tactical and fixed base Air Traffic Control (ATC) systems that will significantly enhance aviation safety in both the tactical and strategic ATC domains. ATC systems are required to achieve or maintain compliance with civil, military, domestic, and international air traffic control and combat identification requirements and mandates. Funding will be utilized to develop, evaluate and integrate candidate systems in each key technology area. Funded in this program element is the development of the Mobile Tower System (MOTS). The MOTS is a tactical mobile tower designed to meet the deployability and communication requirements of the current to future force. The MOTS will be equipped with modernized and secure avionics to ensure highly reliable and consistent tactical aircraft communications across all frequency bands and ranges to ensure compatibility with all Army, Joint, and Allied aircraft. MOTS will provide modern digital, secure, anti-jam communications, a digital recorder, basic weather information, a precision location capability, and full compatibility with all military and civilian airfields as well as tactical landing zones in an armored, survivable vehicle.

Funded product improvements include the Air Traffic Control (ATC) Communications and Networking efforts integration and the Tactical Airspace Integration System (TAIS). Voice radios currently integrated into ATC systems will begin migration to accommodate both voice and high bandwidth data throughput. In a networked battlefield, joint service systems and radars can provide data beneficial to ATC missions assuming a communications infrastructure and data processing capability is embedded in ATC systems. As the Federal Aviation Agency (FAA) and Department of Defense (DoD) transition to aircraft self-reporting technologies such as Automatic Dependent Surveillance-Broadcast (ADS-B)/Combat Identification (CID) and Mode 5, PM ATC will equip tactical and fixed base ATC units with ground receivers and networks to process the aircraft positional data. TAIS, as a Battlefield Automated System (BAS) of the Army Battle Command System (ABCS), requires the development and testing of web-based services for both Army Airspace Command and Control (A2C2) and Air Traffic Services (ATS), and integration of these new web-based services into a Battle Command Service Oriented Architecture (SOA) under the provisions of the Army's Battle Command Migration Plan. TAIS RDTE efforts also include Pre-Planned Product Improvements (P3I). TAIS P3I include, but are not limited to, developing and testing Combat Identification (CID) technologies, and autonomous, embedded Blue Force Tracking (BFT) solutions for third dimension BFT situational awareness with minimal latency.

FY 2007 funding total includes no funding received in GWOT supplemental.
FY 2008 funding total includes no funding received in the Bridge Supplemental.
FY 2008 funding totals do not include any previously requested funding for current FY 2008 GWOT requirements, and no FY 2008 GWOT funds have been previously requested in the RDTE Project of 586.

Accomplishments/Planned Program:	FY 2007	FY 2008	FY 2009
MOTS System Development, Demonstration & Testing	7498	4829	4052
Communications			850
Networking			724

0604633A
AIR TRAFFIC CONTROL
Item No. 86 Page 1 of 7
379
Exhibit R-2
Budget Item Justification

Figure 24. Sample document of a budget summary on Air Traffic Control by Army

By analyzing these documents, he sought to identify common themes, what programs are similar, what makes them similar, and who are working on similar topics. Because it was a large document collection, he had no idea of how they are related in the beginning. So he searched for a visual analytics tool that can help his analysis, and finally found Jigsaw.

His goal in using Jigsaw was to find related tools, topics, technology, and people working on a similar topic in the documents and to discover clusters of data that he might not notice. Instead of deeply analyzing the document collection, he wanted to highlight similarities and connections among the documents so that he could narrow down to specific entities to further investigate.

For this task, he first wanted to find entities that had a similar function. That is, he used Jigsaw for a similar tool search and a synonym search. For example, if tool A forecasts certain type of data, then he tried to find other similar tools and examine their functionalities. Whenever he found a tool of interest, he queried it in the Jigsaw control panel and read returned documents that contained the tool. He also did the same process for a verb that expresses specific functionality such as "predict."

He imported all 10,000 documents into the system and added entity types such as agency, name of tool (technology), and text (description). Usually he started with the circular layout in the Graph View, in which entities appearing within multiple documents are shown inside a circle. The stronger the connection is, the closer to the center the entity is shown. That is, he sought to learn what the most common theme (Figure 25) is among the document collection. From there, he searched for interesting terms and looked for the documents that came up. Then he opened the List View to further explore the connections. Sometimes he would look at immediate clusters – a group of documents towards the center – in the circular layout in the Graph View and highlight those entities so that he could explore them in the List View. By undertaking this process repeatedly, he was able to find what he wanted.

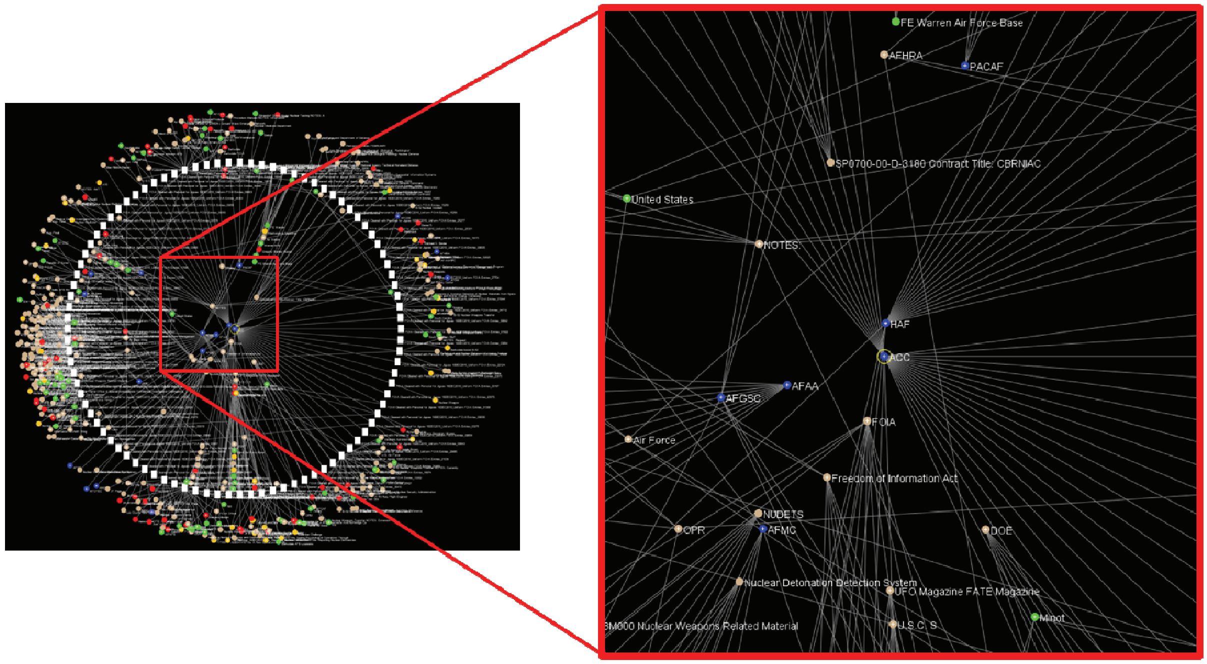

Figure 25. Critical areas identified by the circular layout in the Graph View

With Jigsaw, he was able to effectively search for similar tools and technologies that required further investigation. Through the circular layout, he was able to easily identify where to start his investigation when he did not have a clue where to begin. Even when he had some idea of what he would investigate, Jigsaw helped by showing other interesting documents and keywords so that he could investigate further, which led to a better set of documents instantly. Once he got a set of documents of interest, then he could see important connections such as what are related topics, what kind of programs are related to the topics, and who are working on the programs. He emphasized that Jigsaw was particularly helpful when convincing people because the visualization itself helped draw other people's attention to his work:

> *..and it was pretty, people who received the briefs with that picture (vis) in there, they loved it. They said that the coolest picture was the graph view in Jigsaw. That's a sign, it's a good analytic tool, but having that graphic that you are able to show the most central themes in this set of documents and say that's because of this and this, it's definitely nice to look at that kind of stuff.*

Visualization helps convince people. People pay attention a lot more than if I just told them. It proves itself.

5.4 Findings and Discussion – How They Used Jigsaw

Reflecting on the interviews and discussions with analysts, a number of common themes emerged. Ahead of time, I cared about how an interactive visual system for investigative analysis assisted document sensemaking in various domains, and what kind of issues emerged upon the use of a system. I also hoped to see if professionals used the tool in unexpected ways. I characterize four dimensions in this section.

5.4.1 Types of tasks

While all individuals in the study were from different domains and had unique problems, I could classify their tasks into a few categories, described below.

- *Relationship / connection between entities*: P2, P4, and P5 searched for a tool that could help them make the connections and find complex relationships between entities that were not apparent simply by reading documents. They were investigating emails to detect financial fraud, crime reports to make linkages, and resumes to find a candidate with specific expertise, respectively. Rather than seeing the big picture and understanding the entire story, they did a more targeted investigation. For this type of task, it seemed that Jigsaw's model of connection was sufficient and actually many professionals felt that it is highly useful and beneficial to their task.
- *Search / comparison*: P1 and P6 used the visual analytics system to compare documents and search if the documents contain specific keywords. P1 explicitly compared two sets of documents, examining whether a set of documents contain similar concepts identified in the other set of documents. P6 tried to find if certain

tools or technologies have similar functionalities within the document collection, using the system for a similar tool search and synonym search.

- *Understanding*: P3 actively looked for a tool that can help her understand the huge collection of documents, and thus she used the system to attain a better, clear understanding of the documents. By "understanding," I mean getting an overview of the documents such as "what kind of information the documents contain," "what are important keywords and terms," and "what is happening here." She did not conduct a detailed analysis using the system. Instead, based on the overall understanding she gained from the system, she set the basis for a further analysis, which she performed using other software.

In addition to these three types of tasks, some of the analysts found the system useful as a communication aid as well.

- *As a communication aid / shared understanding of data*: P2, P4, and P6 commented that through the visualization created by the system, they were able to effectively share findings and connections with colleagues. While they did not initially expect that effect, it seemed clear that the visualization system had a persuasive power and added value in communicating with others.

5.4.2 Learning the system

Jigsaw is a relatively complex system and has a number of features that may not be intuitive at first. All the professionals I interviewed had technical knowledge enough to learn and utilize the system. To learn about the system, every person watched the video tutorials available on the web [35] and gained a general idea of how the system works before they started using it. While most of the users also read the tutorial document and found the tutorial very helpful, they admitted that mostly they went through by themselves and interacted with the principal researcher to ask questions and solve issues that arose.

A few professionals explained that they did not have any problems in learning to use Jigsaw, and the system was pretty intuitive and easy to use. But still, many of the users seemed to have a learning curve. The learning curve was more about making sense of "how to better analyze my data using this tool," rather than about learning how to use the system itself. Even after they got familiar with the system and its features, they tried to find the best way to analyze their own data among a number of views and ways to display the data in each view, thus "constructing a frame" [38]. They had questions such as "which views are most appropriate for my data and task?" or "what entity types do I want to put in this column?" Once they found the optimal approach in their own way, they seemed to settle down with it; their usage pattern did not change much.

5.4.3 Unexpected use of the system

In the study, I recognized that the professionals sometimes used the system in unexpected ways, which may provide some insights for design.

The first one is *using the views for evidence/output generation*, rather than for exploration. Originally, Jigsaw was designed for investigative analysis; it tells you which document you should be reading next based on the ones you have read. But often, people used it as a search tool with a visual aid; after they found specific connections by searching a keyword, they created a representation of these connections. For example, in the case of P1, she wanted to create a mapping between two documents and used the List View to more effectively generate the mapping, which was originally done manually (Figure 26). In these cases, it seemed that people missed the investigative power but instead used the view as a presentation aid.

NextGen Operational Improvements	SESAR Operational Improvements and Steps
OI-0320: Initial Surface Traffic	**L07-02 TS-0201:** Basic Departure Management (DMAN)
	L07-02 TS-0202: Departure Management Synchronized with Pre-Departure Sequencing
	L07-02 TS-0203: Integration of Surface Management Constraint into Departure Management
	L07-02 TS-0306: Optimized Departure Management in the Queue Management Process
	L10-02 AO-0205: Automated Assistance to Controller for Surface Movement Planning and Routing
	L10-03 AO-0501: Improved Operations in Adverse Conditions through Airport Collaborative Decision Making
	L10-03 AO-0602: Collaborative Pre Departure Sequencing

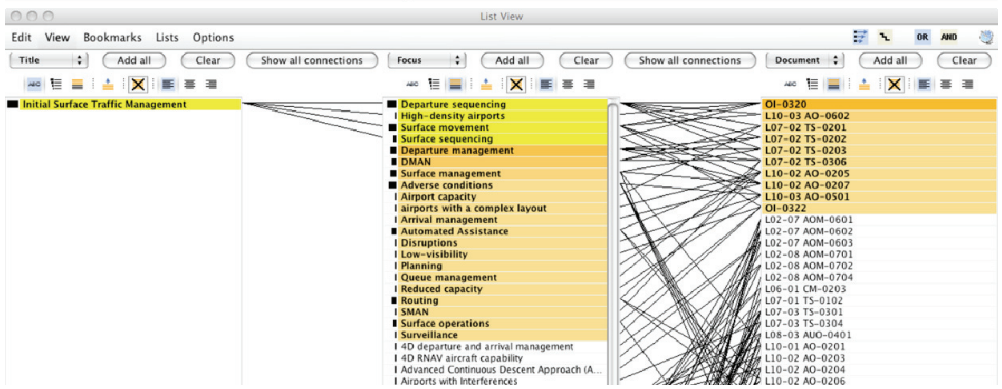

Figure 26. A mapping created manually (top) and by Jigsaw (bottom)

One of Jigsaw's goals is to help analysts with a large number of documents. In the study, however, I found a few experts using Jigsaw for a relatively small number of documents. In those cases, they worked with information-dense documents and did not want to be overwhelmed by the information shown in the system. Thus, they separated documents into several projects, making each project manageable. P5 said that he usually imported only about 10 documents into the system for his analysis:

Sometimes it's as low as five. I tried a couple of hundreds at first, but it was really too much information. Now I try stick to under 10 (documents). I don't like putting too much (documents) into one project because it becomes too complicated.

Some might say it's too few (documents) to use Jigsaw, but it's not that easy. Resumes have condensed information in a few pages. In addition to everywhere *the person's been, you're looking at people who they've worked with. Typically you have a list of publications that have 5 or 6 names, and a couple publishers per line. Using those ten documents to compare to another ten documents, it begins to become more complex. Ten documents...doesn't sound a lot, but it is quite a bit of information.*

In contrast, I found some people using Jigsaw itself as a database. Those people wanted to merge new incoming documents with an existing Jigsaw project and build a historical dataset so that later they can look up to it in a single project file. Three professionals emphasized that they wanted to accumulate new reports to the existing project so that they do not need to re-run all the computations and start over the entity clean-up process. Several users commented:

Analysis is ongoing, it's never done. I want to build on previous data.

I'm trying to figure out how much data should I import. The more data I import, I can see long-term trends and make long-term connections between associates better. But the issues would be time to import and clean up entities on a bigger set of data.

I have about 30 Jigsaw projects. An issue is that 2000-3000 is the maximum for Jigsaw to handle. I mean, the processing time is acceptable for that amount. This is HP documents from 2009 to 2011. I can't do it from 2002 to 2011 because it's gonna be more than 10,000. So I just do it like from 2002 to 2004, something like that. If it was doable, I'd definitely import them all at once.

This notion of "file management" or "project management" could have an important design implication for analytics systems. I will discuss this matter more in detail in the next section.

5.4.4 Issues and problems

Some issues and problems in using the system have been identified through the study, at various levels. Here, I want to highlight a few prominent issues.

One of the initial barriers in working with Jigsaw was technical issues in the preparation stage such as importing data into Jigsaw and identifying entities. Technically, Jigsaw can read in documents in a variety of file formats including text, html, pdft, Word, and Excel files. However, plain text files or Excel files are the most reliable type of file to import, and users are required to transform their documents into text or Excel files if possible. Because people often have documents as pdf or Word files with complex formatting and images, importing these files directly into Jigsaw is less reliable, and therefore, users need to put extra effort to convert their documents into plain text or Excel files. Identifying and working with entities is another similar issue. While Jigsaw provides the automated entity identification feature, which attracts many users, it is not perfect and many false positives and negatives can occur. In order to fix incorrect entity identification, users have to manually choose each word to add, remove, or modify. Creating a new entity type is common because users have their own interest when working with documents. Actually, all of the professionals in the study created their own entity types applicable to all documents in the collection. Users also have to go through the process of entity aliasing, which create aliases for entities that are identical but worded differently.

While this grounding process – both importing documents and cleaning up entities - does not seem to be a serious issue in terms of the analysis process, it turned out that most people considered it as one of the biggest difficulties in using Jigsaw. Without addressing these issues in the beginning, they are not even able to see their data properly displayed in the views. When they encounter any problem in the process, they need to contact the researcher and follow the instructions, which could be cumbersome and even daunting to someone without technical background. All the professionals in the study mentioned that the initial processing required a lot of time and efforts.

> *One of the biggest difficulties that I encountered was Entity Identification. When importing data, because resume is not a type of data that Jigsaw is designed to*

read, I still have issues with entity identification. Well, I have to go through it one by one. Have to do a lot of cleaning after the initial importing of data.

Once they undergo this stage, however, they became easily engaged in working with the views.

Another issue that the professionals faced was that Jigsaw has very limited filtering options and users are not able to easily select a subset of data in the views. Currently, once Jigsaw reads in documents, all the operations and computations are run upon the entire set of documents. That is, once users have ingested a collection of documents into Jigsaw, they will notice that all the document and entities are active. If they want to temporarily exclude some documents and explore connections only for another set of documents, the only way to do it is to start over from the data importing process. They have to decide what documents in the collection to look at, create another collection of those selected documents, and import the documents into the system. In other words, there is no easy way to select a subset of documents while working with views. Users wished that they had a better, flexible way to have a certain set of documents, as expressed in the quote below. He compared IN-SPIRE [86], a visual analytics system for text analysis which provides an overview of the key themes and trends across a document collection, to Jigsaw when discussing this feature. In IN-SPIRE, he was able to make a selection of documents even after all the dataset was displayed:

I started IN-SPIRE *about at the same time I used Jigsaw. The thing I liked best in Inspire, which Jigsaw doesn't have, was that you have all the dataset up there on the screen (the galaxy view), and I could easily select across all data, make the selection and make the rest of them outlier, and have the just ones I have selected. In Jigsaw, I have to have all set of data.*

I assume that if selecting and working with a subset of data was easy enough, some of the professionals might not have had to segment their data into several Jigsaw project files, since they would not have had the information overload issue.

Finally, there was an issue of trust on the system. While people favored the automatic power of the visual analytics system, they did not seem to solely rely on the system as in the quote, which is a common behavioral pattern of analysts [38].

I'm the only one who's using it in our team. They don't think it's reliable enough.

It seems that this mistrust is raised when the process does not flow smoothly. When the system fails to import documents or identify entities that they want to see, they tend to attribute it to the lack of system reliability. This tendency is more likely to appear to people with less technical capability, those who are not willing to put extra efforts in troubleshooting. Or simply, some experts think that the system assists part of their work more efficiently, but ultimately, they believe that they can do the job more accurately. For example, after working with Jigsaw, P1 double-checked its findings with those from a manual process in order to validate her analysis:

Finally, we carefully reviewed descriptions of OIs for which one or many counterparts were identified with experts, in order to ensure that the themes and ideas behind these concepts were indeed analogous. It was found that the mappings obtained through Jigsaw were similar to the ones obtained manually, and thus we could say that Jigsaw offers a valuable alternative to our manual approach.

5.5 Design Implications

While I interviewed only a small number of professionals, the study still suggests several design implications for visual analytics systems as well as Jigsaw. Although not all these suggestions are applicable to all systems because different systems serve different tasks and purposes, my hope is that these implications would be helpful for designing such systems in a larger sense.

Supplement automatic entity identification

While there exist a number of entity identification systems [4,14,23] and visual analytics systems that incorporate entity identification [4,11,67], the process typically is not perfect. Some entities may not be identified at all, some may have an incorrect entity type assigned, and some identified ones may not be entities. Systems should provide ways to correct such errors, and the process needs to be intuitive and efficient. While Jigsaw allows users to modify, remove, add, and alias entities, professionals pointed out that it is still not a simple, easy process, as mentioned in the previous section. For capabilities such as entity aliasing, the process is not automatically going forward. That is, when new documents are imported, the analyst must manually create the aliases again.

Another issue is that although Jigsaw allows users to create a new entity type and specify the instances of that entity, they seemed to be unaware of the feature. For example, once they create a new entity type "Company name," they could create a text file that has each different possible entity value such as "HP," "Apple," "IBM," etc. While every user in the study created their own entity types, most of them did not know about this feature but specified each entity every time they opened a new project, which took significant time. Four users suggested a feature that Jigsaw already provides:

> *I suggest an entity library you can draw on for every project. Then you wouldn't have to keep creating new entities.*
>
> *It will be nice to have an entity list that you can apply to each project you do and not have to recreate them. For example, list of universities that would be identified every time, list of technologies that would be identified every time, so you only have to make the list once.*

I suspect that the way to create a new entity type was not intuitive or salient enough to users. The feature could have been more nicely incorporated with the entity identification work flow, for example, by asking them to type a list of entities instead of importing a text file, so that users do not need to create an extra file outside the system.

Allow flexible data (document) management

Previously, I discussed that some people worked with multiple files of a small number of documents while some people wanted to accumulate documents into one file and build a database. In most cases, they had to try multiple imports to find an appropriate number of documents they need to import at a time. There are two reasons for this issue of data size – technical capacity and information overload (e.g., not wanting to be inundated by information shown). While technical issue needs to be solved at a lower level, information overload can be addressed by providing a flexible data management. Currently, once a user imports a document collection into a Jigsaw file, all the documents and entities are "active," which is often overwhelming. If a user wants to investigate only part of the document collection, there is no easy way to do it but create another subset of the collection and import it. This is inefficient especially when users want to try different subsets of documents in a single document collection. Ultimately, users desire to be able to flexibly work with documents within a single database, and a system should provide an ability to easily select a subset of documents to investigate once users import a document collection. For example, a system could provide a way to choose a subset of documents and run analysis only for the selected documents. Or a system could allow users to temporarily exclude a set of documents so that they can work with the remaining documents only. I assume that if selecting and working with a subset of data was easy enough, some of the professionals did not have to segment their data into several Jigsaw project files since they would not have had the information overload issue.

Systems also need to provide a way to easily accumulate documents into the existing file. In many cases, users may want to build a database over time, especially when they receive documents incoming regularly and the analysis is ongoing. Currently in Jigsaw, if users want to add only one or two documents to the existing Jigsaw file, they have to repeat the process of computational analysis on the document collection. This is very inefficient because not only the system has to be re-run, but also users have to go through the entity clean-up process again. Often, users do not have the complete documents prior to investigation or they receive new documents continuously. They would want to simply “merge” new documents into the existing file, upon which entities are already cleaned up and computational analysis is done.

Empower with numbers

Jigsaw was developed for unstructured text data and does not provide statistical analysis per se. For example, in order to show connection strength, the system uses colors (darkness) or list order. However, most of the analysts in the study strongly expressed that statistical functionality would be really desirable. Depending on the domain and task, analysts often need to convert results from investigative analysis into evidence, which is better supported with quantified information such as descriptive statistics or counts. In the study, several users wished to have statistical importance metrics such as degree centrality, betweenness, closeness, or others so that they could have more accurate metrics of the connections between entities and documents. Even for investigative analysis systems that deal with unstructured data such as text, it seems important to have simple statistics and measures, which is consistent with findings from Perer and Shneiderman’s study [56].

Consider interaction paradigm

The professionals in this study wanted to have more control and flexibility over the visualizations. They sometimes wanted to be able to annotate, mark, and change the representations. Such changes may not be feasible or desirable from the point of view of the system, however. For example, the visualizations presented by a system may communicate analysis metrics or results computed about the data. Allowing the user to modify the visualization would be, in this case, inappropriate because it could make the visualization present the analysis data inaccurately. Conversely, allowing the analyst to simply highlight or augment the visualizations would not violate the fundamental data-to-representation mapping. Presently, Jigsaw allows no view augmentation. Should it? It is important that system designers and developers carefully consider the style of changes, if any, that viewers can make to a system's visualizations.

Invest in tutorial

Usually, visual analytic systems for investigative analysis tend to have a number of features and interaction techniques, which makes it hard to get familiar with a tool without any external aids such as one-on-one training or written instructions. In many cases, tutorials seem to be quite important and helpful for learning visual analytics systems. While some people may argue that users do not pay much attention to tutorials, all of the professionals I interviewed said that they put considerable time and effort in reading the tutorial document and watching video tutorials.

Another reason for the importance of tutorial is the intermittent use of a system. Many professionals pointed that they do not use the system on a regular basis. Instead, they used the system when they have enough time, when they new data, or when they need to prepare a brief. Consequently, they often forgot about some functions and operations and had to revisit tutorials in early stages. Thus, it is desired to provide an intensive but still easy-to-understand tutorial.

For example, breaking down the tutorial into subtopics with use-cases and examples would be really helpful, as the users commented:

> *For learning, I mainly used the video tutorial. It was very useful actually. They are good because they're broken down into topics and you can pick what you need help with. I like it a lot.*
>
> *I wished a better tutorial though. I want to see more examples about each view so that I can find the best way to analyze my own data.*

Jigsaw-specific recommendations

The study helped identify issues and future work for Jigsaw:

- *Focus on useful views*: While different users have different preferences of views, it was clear that the List View was most useful. I suggest that future development focus on improving the features and interface in the List View, as it will definitely benefit real world users. The Document View, the Document Cluster View, and the Graph View were also used by several analysts. Multiple analysts mentioned that they did not find the Timeline useful, and the Scatterplot View was not even used at all. Those views may need significant changes or be removed from the system.
- *Give them power to control*: When working with their own data, users want to actively interact with the system because they have their own goals and expectations from the system. While Jigsaw is very good at "showing" documents and entities in different ways, professionals wanted to be able to annotate and manually alter visual representations. The professionals said:

  In Jigsaw, you cannot change anything as a user. You cannot annotate, draw lines, add connections, etc. it's just there.

CHAPTER 6

DISCUSSION

6.1 Reflecting on Research Questions

This work was driven by four research questions:

- Do current models used by developers of visual analytics tools adequately characterize the process of intelligence analysis? What aspects of intelligence analysis are particularly misunderstood?
- Where in the analysis process and for what kind of tasks can visual analytics tools best benefit intelligence analysts without intruding on their work practices?
- How do existing visual analytics systems such as Jigsaw support or fail to support investigative analysis?
- What design implications for visual analytics systems for intelligence analysis emerge from the studies of the analysis process and the use of a visual analytics system?

The first and second questions were raised because I felt that we do not have enough understanding about intelligence analysts and their work process. I doubted whether visual analytics systems being built are well aligned with analysts' work cycle and are truly beneficial to them. One of the best ways to answer the questions was to simply observe and ask about what they do during analysis. For this, I conducted a long-term field study that consisted of a series of in-depth interviews and observations with student analysts working on a real intelligence project. By examining how they normally perform an analysis project and what kind of problems and issues they often encounter, the study identified several misunderstandings we might have about their process and analysts' specific needs for visual analytics systems, thereby answering the two questions. Another possible study considered was having analysts use existing visual analytics systems and observe how they perform analysis. In that way, I might have been able to identify specific issues regarding existing visual analytics systems during the entire analysis cycle. However, in practice, analysts have very limited time for a project, and they likely would not be willing to learn a complex system that is not vital to their work. Thus, that approach would have not been as fruitful as the one I did.

While the questions about analysts' process and its implications were answered, I believed that the community still needs more systematic evaluation studies of existing systems. By observing people using visual analytics systems, I sought to gain insights for analytical processes and design implications for analysis tools. Are visual analytics systems truly helping analysts? If so, what characteristics are relevant? If not, what do we need to improve? While a variety of evaluation methods exist, I conducted two different evaluation studies. In the first, I compared within a laboratory setting the usage of a visual analytics system to more traditional methods of analysis. The study demonstrated how a visual analytics system adds analytical benefits and helps people perform an investigative analysis, along with a list of design implications. If there were enough time and resources, it would have been also useful to compare the usage of multiple visual analytics systems in a laboratory setting so that we can investigate various features and capabilities provided by different systems.

Although the study provided interesting findings about analytic strategies and useful implications for visual analytics tools, it would be particularly beneficial if we could identify the utility of visual analytics systems in practice and further derive implications for design through case studies. Especially in the field of visual analytics, which has a relatively short history, such case studies are rare. Fortunately, I was able to recruit domain experts who had been using Jigsaw with their own data and conduct interviews with them. From the study, I found issues and problems that would have been difficult to identify otherwise. Lessons and implications from anecdotal findings were also valuable, complementing the findings from the laboratory study. I chose case studies because they yield a description of how a tool was used and where the users had problems by detailing the use of a system. Until their particular challenges are understood, it remains difficult to know how a visual analytic system helps expert users attain their goals. These findings are difficult to achieve through controlled lab studies or other types of studies.

The last question has been addressed by combining results from the first three questions. Through a series of studies, a number of design implications for visual analytics systems for document analysis have emerged, as discussed more in detail in the following section.

6.2 Revisiting Design Implications

From the three studies, I derived a set of design implications that can be useful for designing visual analytics systems. While the ultimate goal of the studies was the same—"to inform the design of visual analytics systems for document analysis," the three studies had very different approaches and settings. In the first study, I observed teams of analysts conducting a real intelligence project for 10 weeks. In the second, I did a detailed analysis of how people with different technological aids analyzed a given set of documents. In the third study, I interviewed Jigsaw users and asked them how they used the system, what helped and what did not, and what they wanted for the system. Accordingly, implications derived from each of these are best explained under the different conditions and settings that were being examined. However, I found some implications sharing similar principles, triangulating the findings from each study. In this section, I will synthesize all the design implications and discuss commonalities and differences between them. Then I will discuss how this work is different from existing work that also provides design implications.

6.2.1 Commonalities and Uniqueness

Among all the design implications obtained, I identified four principles that are mutually reinforced across the studies.

1. Support for structuring representations in analysts' conceptual space

In the study with intelligent analysts, one of the implications derived is "externalize the thinking process - help analysts continuously build a conceptual model" because it seems to be highly relevant and important in the sensemaking process. Using the power of representation, visual analytics systems can help analysts build a conceptual structure of their problem. By connecting, grouping, and organizing concepts, analysts should be able to continuously build up their conceptual model or structure of the problem throughout the process. This principle is also suggested in the second study, as "support for the schematize stage." In the study, it seems that Jigsaw does not really support putting the pieces of evidence together. Analysts may discover the pieces to be put in a puzzle and have a sense of which piece goes where, but they should receive help in putting the pieces together. That is, they need to be able to visually structure the data and organize them into their own conceptual structure, and this will significantly help the sensemaking process.

This principle is also highlighted in Zhang's dissertation [88]. He argues that "Research in visualization has put much emphasis on visualizing the collection that users search rather than the conceptual structures that users create through sensemaking tasks" and provides the basis for designing "tools that help structure the representations in a sense-maker's conceptual space to provide better sensemaking support to information system users."

2. Support for flexible inclusion/exclusion of data during analysis

The first study recommends a system to "support analysis with constantly changing information - integrate collection and analysis in a single system" because collection and analysis are highly integrated processes. Analysts start analysis even when they have only a few pieces of information and repeat the process of collecting and analyzing throughout the cycle. That is, analysts collect during analysis and they analyze during collection. In this case, analysts may not have all data in the beginning of the process and want to include more documents later. Or, they may decide to exclude part of the data as analysis proceeds. Therefore, it is crucial for an analyst to be able to flexibly add and remove part of documents during the entire process.

The third study also emphasized "allowing flexible data (document) management." In the case studies, domain experts often wanted to try different subsets of documents in a single document collection, which indicates that users desire to be able to flexibly work with documents within a single database. Thus, a system should provide an ability to easily select a subset of documents to investigate once users import a document collection or allow users to temporarily exclude a set of documents. Systems also need to provide a way to easily accumulate documents into the existing project. Often, users do not have the complete documents prior to investigation or they receive new documents incoming regularly. In these cases, users want to build a database over time, by simply "merging" new documents into the existing file. Overall, it will highly improve the analysis process if a system supports flexible inclusion and exclusion of data.

3. Support for various organization structures

The first study suggested "helping analysts use structured methods during data collection" such as ACH, social network analysis, geospatial mapping, or decision matrix because analysts always want to triage their findings by trying multiple analytic techniques. The ability to apply various techniques with the same underlying data can help analysts address questions and strengthen their analysis. In many cases, analysis techniques imply different organization forms of the same data. For example, it could be a table, outline, network graph, or maps, depending on how the analyst wants to restructure/represent his data.

A similar principle, "allow flexibility in organizing" is identified in the second study. When supporting the "schematize" stage, developers of investigative analysis tools should consider that individuals will choose different organizational metaphors or schemes. Rather than providing one fixed schema, tool developers should consider providing a number of different organizational schemes so that analysts can try various analytic approaches.

Again, Zhang's study [88] suggests how multiple representations may help sensemakers in accomplishing sensemaking tasks. He argues that multiple representations of the same underlying structure such as network representations (e.g., maps), concept hierarchies (e.g., an outline or directory), and text representations offer different contributions to users' sensemaking.

4. <u>Support for marshalling data</u>

"Support evidence marshalling" was identified as one of design implications in the comparative lab study of Jigsaw. For a sensemaking system, it is crucial to include a workspace in which analysts can drop/paste entities, draw connections between them, and add annotations so that it can be organized in a way the analyst wants. These capabilities found in systems such as Analyst's Notebook [33], the Sandbox [87], and Entity Workspace [4]. The case study with domain experts also revealed that people want to more actively interact with their data such as annotating and manually changing visual representations (e.g., drawing connections between entities). It is important to provide support for marshalling their data during analysis.

While the previous implications are emphasized in multiple studies, I found several implications that are unique to one study due to different settings and contexts across the three studies.

5. <u>Help analysts create convincing production</u>

From the first study, I suggest it would help analysts create convincing production if insight provenance and sanity checks of analytical products were supported. This implication could have been identified in the context of the study because it was the real intelligence project in which the analysts had to produce deliverables for their own client. While intelligence analysis and sensemaking are similar, production is more emphasized in the intelligence analysis community because it is the ultimate goal of the analysis. That is, they do the sensemaking/intelligence tasks for other decisionmakers, not for themselves and even when analysts finish their analysis, they need to convert the results into a concise format so that decisionmakers can understand their findings. This production part turned out to be challenging and time-consuming because of sanity checking and insight provenance. The sanity check, or qualitative double-check, refers to a process of ensuring that statements, sources, and references are correct. This process takes time because data and findings are derived from many sources and analysts have meshed them through the process of collection and analysis. A system that promotes simple insight provenance during analysis could help analysts save time in production. For example, a system could implement an architecture that links data to sources so that the reference can be automatically made without errors.

6. <u>Support asynchronous collaboration for exploratory analysis</u>

A collaboration aspect was examined only in the first study, and the study suggests supporting asynchronous collaboration rather than synchronous collaboration for exploratory analysis. I observed that analysts collaborated cognitively and that collaboration at the "sharing" or "content" level tended to occur asynchronously. Rather than trying to build a system that allows analysts to work at the same time in the same workspace, providing a system that promotes individual workspaces but also provides asynchronous collaborative features - such as the ability to share sources and data, view and comment on others' work, and merge individual work together - would appear to be more beneficial.

7. Facilitate clue-finding

Among the three studies, a detailed observation on the analysis process was possible in the comparative lab study. From the study, participants who employed the "find a clue, follow the trail" analysis strategy generally performed better than those who used other strategies. It seems that the strategy is efficient and effective in an investigative analysis, and a system could support analysis by promoting the approach. Thus, investigative analysis tools may want to support analysts in finding appropriate starting points or clues, and then, following the trail of these clues efficiently. Further, the performance of those participants who were able to focus only on relevant documents was outstanding. Investigative analysis tools should help direct the analyst's attention to relevant information based on initial clues.

8. Suggest alternative paths but support task resumption

The lab study also found that an analyst may confront a dead-end or find evidence that refutes an existing hypothesis. When an initial clue was not a really essential piece of information or the analyst deviated from a desired path, the analyst needs to start over or revert to previous steps. Thus, investigative analysis tools need to support the analyst to find appropriate next steps or alternatives by making the milestones of the investigative process explicit. In this way, the analyst can come back to the point where she/he was earlier and start over from that point. This also ensures that the analyst can proceed further without being too concerned about keeping track of past states.

9. Supplement automatic entity identification

The case study with professionals highlighted practical issues in using a visual analytics system, which would have been difficult to identify using other types of studies. One issue is relevant to working with automatic entity identification. Because the entity identification process typically is not perfect, some entities may not be identified at all, some may have an incorrect entity type assigned, and some identified ones may not be entities. Participants in the study spent most time working with correcting such errors and wished a better support for that. Systems should provide ways to correct such errors in more intuitive and efficient ways.

10. Empower with numbers

Most of the analysts in the case study strongly expressed that statistical functionality would be really desirable. Depending on the domain and task, analysts often need to convert results from investigative analysis into evidence, which is better supported with quantified information such as descriptive statistics or counts. In the study, several users wished to have statistical importance metrics such as degree centrality, betweenness, closeness, or others so that they could have more accurate metrics of the connections between entities and documents. Even for investigative analysis systems that deal with unstructured data such as text, it seems important to have simple statistics and measures, which is consistent with findings from Perer and Shneiderman's study [56].

11. Consider the interaction paradigm

Analysts may want to have more control and flexibility over the visualizations. They sometimes want to be able to annotate, mark, and change the representations. Such changes may not be feasible or desirable from the point of view of the system, however. For example, the visualizations presented by a system may communicate analysis metrics or results computed about the data. Allowing the user to modify the visualization would be, in this case, inappropriate because it could cause the visualization to present the analysis data inaccurately. Conversely, allowing the analyst to simply highlight or augment the visualizations would not violate the fundamental data-to-representation mapping. It is important that system designers and developers carefully consider the style of changes, if any, that viewers can make to a system's visualizations.

12. Invest in tutorials

The case study also emphasized the need for learning aids and especially, tutorials seem to be quite important and helpful for learning visual analytics systems. All of the professionals in the study put considerable time and effort in reading the tutorial document and watching video tutorials to get familiar with the system. It also turned out that many professionals use the system intermittently, rather than using it on a regular basis. Consequently, they often forget about some functions or operations and have to revisit tutorials, and it is desired to provide an intensive but still easy-to-understand tutorial. For example, breaking down the tutorial into subtopics with use-cases and examples would be really helpful.

These findings hopefully will assist the developers of future visual analytics systems as they consider different issues and factors affecting their systems.

6.2.2 Existing Work on Design/Evaluation Implications

Other researchers have also put considerable efforts to enumerate design implications in visual analytics. In their paper, Heer and Agrawala [26] emphasized the need for supporting social interaction in sensemaking and present design considerations for asynchronous collaborative visual analytics. Based on their experiences and literature survey in visual analytics, social psychology, organizational studies, and computer-supported cooperative work, they identified a set of design considerations for seven areas, including Division and allocation of work, Common ground and awareness, Incentives and engagement, and Identity, trust, and reputation. For each of these areas, they describe each topic and suggest a few mechanisms for achieving them. While the purpose is similar to the goal of my research—to inform the design of visual analytics systems, I focused more on understanding our users and their practices, thereby bridging the gap between researchers and analysts. The research also sought to derive design implications based on user behavior and feedback from the usage of a specific visual analytics system, Jigsaw.

Forsell and Johansson [22] examined Heuristic Evaluation as a useful evaluation method in Information Visualization, arguing the need for heuristics that are consistent, standardized, and well-adapted for assessing usability issues in InfoVis techniques. To identify which existing HCI heuristics are most useful for assessing interactive visual displays in InfoVis systems, they let six experts rate how well a total of 63 heuristics from 6 earlier published heuristic sets could explain a collection of 74 usability problems derived from earlier InfoVis evaluations. Based on the results, a new set of 10 heuristics for InfoVis were derived. While those heuristics are highly useful as it provides researchers with a good starting point for evaluating visualization systems, they mainly focus on usability issues in visual displays. Consequently, those guidelines are more suitable for evaluating systems than designing them. My research tried to uncover more underlying issues rather than usability issues so that implications can be made useful before developing a system. Also their study did not involve actual experiences with visualization systems.

Scholtz [72] also made a contribution to developing guidelines for the evaluation of InfoVis using another approach. By synthesizing the 2009 Visual Analytics Science and Technology (VAST) Challenge reviews from reviewers (e.g., professional analysts and visualization researchers) and results from a user study with professional intelligent analysts, they developed guidelines for evaluating visualizations in visual analytics environments. Then they incorporated the results with other heuristics developed in various domains, including Forsell and Johansson's heuristics, to provide an explanation for the issues. They also worked with analysts to understand what criteria they use in evaluating analytics reports and identified possible guidelines for evaluating the quality of analytic reports, which can be useful for visual analytics researchers to conduct an evaluation study. Similar to Forsell and Johansson's heuristics, implications in this study are derived from expert reviews, rather than from actual users or interaction with systems. The guidelines are at the UI level, making them more useful for evaluation than generating designs.

While each study has a different focus and approach, findings and implications derived from these studies can inform the design and evaluation of visual analytics systems in different ways. More research efforts to enumerate design and evaluation guidelines will certainly benefit the visual analytics community.

6.3 Reconsidering Sensemaking Models

In earlier sections, I quoted an intelligence analyst saying that traditional intelligence process models do not accurately describe how intelligence is produced. I presume that it is because the models do not capture the subtle nuance of the process as practiced, not because the models are flawed. Models are inherently abstract, and it is quite difficult to reflect the complexity of real-world practices in one simple, theoretical model. This is even harder especially for sensemaking models, which try to extricate the complex relationship between human cognition, information, and representation. That is, VA researchers may want to model analysts' processes, but analysts' work and how intelligence is produced may not be easily modeled. Their work is not discrete or procedural, but more parallel and integrated. Various sensemaking models explain the process in different ways, and their suitability depends on the type of the sensemaking task and the specific domain.

In the field of Visual Analytics, Pirolli and Card's sensemaking model has guided the design of many visual analytics systems. While the model was suggested as a starting point to investigate the domain by the authors, researchers have relied on the model without validating its accuracy or applicability, probably due to the lack of sensemaking models suited to designing systems. However, results from the study with intelligence analysts suggested that this model does not explain the parallelism and integration of the process, and that tools based on the model tend to support specific stages only and often do not blend into the entire process of intelligence analysis.

In his dissertation, Andrews [2] also discusses the importance of the integrated approach to sensemaking, especially integrating foraging and synthesis. From a user study, he found the fluidity of the process and describes "..the analysts we studied moved freely around the sensemaking loop, jumping through various levels of abstraction." Based on observations and findings from his previous studies, he tried to address the disconnect between foraging and synthesis by building a sensemaking environment that unifies the activities of both. In Andrew's work, Pirolli and Card's model was used as a useful basis, providing concepts and elements. I assume that this was possible because he carefully examined analysts' activities in conjunction with the model, trying to identify room for improvement in the process.

Pirolli and Card's model can also provide researchers with a framework for analyzing and describing user behaviors and phenomenon surrounding a visual analytics system. However, the model can be made useful only when people understand the context behind it. Through my research, I tried to provide the context by examining what aspects we might have understood regarding the intelligence/sensemaking process and how our assumptions should be changed. Based on my research, I argue that the model better describes how information transforms and how data flows, rather than how analysts work and how they transition. While it gives a nice illustration of how the form of information evolves from raw data to reportable results, it does not quite fit analysts' mental model of their work process because they work in parallel, cyclical process while information is transformed in a linear sequence. Similarly, all different states of the model can exist at any point during the process, highlighting the parallel, integrated process of sensemaking.

In sum, while Pirolli and Card's model can serve as a helpful framework to explain what is going on during analysis in terms of the state or form of data, when the model is used without enough understanding about what the model really describes, it could give researchers an inaccurate impression of how analysts really work. Furthermore, designing a system solely based on the model without a further understanding about user processes and practices can result in a fragmented system that does not really support the analysis process.

In the field of Visual Analytics, the ultimate goal of researchers is to develop a better system for visual analysis and sensemaking. Consequently, sensemaking models in the community need to provide a sufficient basis for designing technological support to analysis, as well as describing what is occurring in the process. In order to do that, VA researchers need to understand what the context of the sensemaking model is, what the model best explains, and what is not reflected in the model. That is, considerable additional detail is required to communicate the real understanding of the process in the model. For example, researchers need to conduct more empirical studies that closely investigate users' work processes in conjunction with sensemaking models such as Pirolli and Card's model. Simply relying on a sensemaking model without an enough understanding of it could lead to a misalignment between a system and users' work flow.

6.4 Limitations of the Study

Evaluation in Visual Analytics or Information Visualization is challenging. During the research, I encountered various challenges and difficulties, and some of them still remain as limitations of the study.

One of the biggest limitations is that I used student analysts in the first study instead of actual analysts. Although they were students in an intelligence analysis graduate degree program, I suspect that they are different from working professional analysts who belong to specific agencies. Students in the study might have learned specific methodologies and techniques that are prevalent only within the institution. Or their analysis might have been limited to public sources and information because student analysts have limited access to confidential information. While intelligence analysts heavily rely on and exercise their personal knowledge in a domain when defining and establishing relationships, the student analysts may have relatively little background knowledge to apply.

After the study, I wanted to better understand the difference between the student analysts and professional analysts. According to the instructor, while those student teams clearly are not practicing professional analysts, there was not a significant difference between the way the students worked and the way real analysts work. The analysis process used in the class was modeled directly after the process employed by the US National Intelligence Council to produce its strategic reports, the National Intelligence Estimates [52]. The instructor also intentionally stayed relatively detached from the students, acting as a mentor and limiting his supervision so that the teams could autonomously work on the project. The teams were diverse in expertise on the subject matter, which is common for teams in the intelligence community. One key difference from real world practice was the relative absence of administrative and bureaucratic overhead affecting the student teams, as well as issues in getting access to different levels of classified information. They operated in a much more "sanitary" environment than the real world.

Student analysts were also working on only one problem throughout the study period while intelligence analysts often work on more than one problem. Due to the limited time and resource, I was not able to take into account that issue, and this might have affected the collaboration pattern and analysis process in the study.

The second (lab) study also has several limitations that likely affected its findings. First, from the results (scores) and observations, it is clear that there was quite a bit of variability among the participants. I speculate that certain individuals simply have better innate skill at such analysis tasks.

The study compared Jigsaw to other traditional tools, but not to other visual analytics systems. Comparing the usage of the tool to other existing systems developed for investigate analysis would generate more insightful findings and implications.

A relatively small document collection was used for the study, which likely would not be the case in reality. The collection size was chosen to make the experiment feasible in a reasonable amount of time. I speculate that some of the findings would only be amplified when working with larger document collections.

For the third study, I was not able to explain the entire cycle of each participant's analysis in detail because many of their documents and analytic results were confidential and could not be shared. I was only able to provide a general description of the analysis process.

As I stated earlier, I experienced several difficulties in the process of conducting research. One of the biggest difficulties was recruiting—getting access to the intelligence community. The ideal scenario we planned was recruiting real analysts as study participants. Unfortunately, the community is largely shielded from the public and finding someone involved in that community is not easy. Even if I had a personal connection with someone in the community, conducting a formal study with a few analysts is extremely hard because of security issues. One approach I tried in order to get access to the professionals was doing an internship at an intelligence agency, but as a non-U.S. citizen, it was also not possible because I could not get security clearance. For those reasons, I had to find an alternative. Using a connection with a faculty member at the department of intelligence analysis at Mercyhurst, I recruited student analysts as study participants. Although they are not working professionals, they were more available for the study. I also broaden the scope of "analysts" to those who conduct a similar sensemaking tasks such as researchers and business analysts, which made the recruiting easier.

Limited access to sensitive information was another issue even after recruiting. Except for the second study, I tried to observe and learn analysts' processes and practices under a real world setting, rather than controlling the study environment. They conducted analyses using their own work-related data, which often involves confidential information. The analysts I worked with tended to be very sensitive about their data and the results of their analysis which made it difficult for them to share concrete examples.

CHAPTER 7

CONCLUSION

In this document, I presented my research that aimed to inform the design of visual analytics systems for investigative analysis. In order to better understand users' environments and work processes and practices, I conducted a long-term observational case study with intelligence analysts working on the real-world problems. The study documented the processes and methods they followed, clarified several misconceptions regarding the intelligence analysis process in the visual analytics field, and suggested design implications for visual analytics systems for intelligence analysis from the analysts' perspective.

While many researchers in the visual analytics community firmly believe that new visual analytics technologies can benefit analysts, showing that is the case is still a challenging proposition. In order to assess how visual analytics systems add analytic benefits, I conducted a comparative lab study and compared the use of Jigsaw to existing, more traditional methods in the context of an investigative analysis. While lacking the size and depth to identify statistically significant differences, the study nonetheless suggested how a visual analytics system such as Jigsaw can benefit investigative analysis and how its absence amplified challenges and difficulties. The study also provided a description of four analytic strategies employed by participants, as well as identifying design and evaluation suggestions to make visual analytics systems for investigative analysis more effective.

In order to evaluate long-term, field use of Jigsaw, I conducted in-depth case studies with analysts from a variety of domains. I interviewed six investigators from the intelligence, academic, and law enforcement communities who had been using the system for a period of 2-14 months. I asked them about their use of Jigsaw, the types of data they were working on, and difficulties they encountered. Analysts used Jigsaw for finding relationships, comparing documents, getting an overview, and sharing analytical products with others. Their primary difficulties included importing data into Jigsaw, identifying entities in the preparation stage, and selecting a subset of data during data exploration. The contributions of this work thus include (1) identification and review of real-world cases of how an interactive visual system for investigative analysis assisted document sensemaking in various domains and tasks; (2) discussion of issues and findings that emerged upon the use of the visual analytic system; and (3) development of design recommendations and suggestions for the system and future visual analytics tools. A growing number of visual analytics systems are being developed and used in practice. Assessing the utility and value of a system is essential for improving it, and I believe that the field needs more of case studies because it helps to understand the types of tasks and problems a system can address and to identify strengths and weaknesses of a system in real world settings.

Finally, I assembled a number of design implications identified from the studies and synthesized them into a set of implications. While those design implications may not be applicable to every visual analytics system, I believe that they can be more helpful in understanding our users, their processes, and their system usage. I hope that my research contributes to bridging the gap between users and system developers, ultimately informing the design of visual analytics systems.

My research presented in this document is the first step of an agenda of informing design of visual analytics systems for intelligence analysis. There are still a number of areas to explore in understanding our users, their tasks, and their tool usage. One important issue is that my study covered a specific environment (e.g., single organization, single problem) only, and thus more user studies with professional intelligence analysts need to be conducted. It is also desirable to validate to what extent the design implications from this research can benefit the design of visual analytics systems.

Evaluation of visual analytics systems must progress in step with new technical development for continued progress. Understanding our users, how they word, and how and why systems aid analysts will help to inform future designs and research. I believe that my research provides initial evidence and insight in this area, and sheds light on many challenging open questions.

REFERENCES

[1] W. Aigner, A. Bertone, and S. Miksch. Tutorial: Introduction to Visual Analytics. In: Holzinger, A. (ed.) USAB 2007. LNCS, vol. 4799, pp. 453–456. Springer, Heidelberg, 2007.

[2] C. Andrews. Space to think: Sensemaking and large, high-resolution displays. PhD Dissertation, 2011. http://scholar.lib.vt.edu/theses/available/etd-08232011-160350/

[3] D. C. Andrus. The wiki and the blog: toward a complex adaptive intelligence community. In *Studies in Intelligence*, 49(3), September 2005.

[4] Balie. http://balie.sourceforge.net/ (accessed 01 March 2012).

[5] E. Bier, E. Ishak, and E. Chi. Entity workspace: an evidence file that aids memory, inference, and reading.*Proceedings of the IEEE International Conference on Intelligence and Security Analysis (ISI '06)*, (San Diego, CA), Springer-Verlag: Berlin, May 2006; 466–472.

[6] E. Bier, S. Card, and J. Bodnar, Entity-based collaboration tools for intelligence analysis. In *IEEE VAST*, pages 99–106, Oct. 2008.

[7] W. S. Brei. *Getting Intelligence Right: The Power of Logical Procedure*, Occasional Paper Number Two (Washington, DC: Joint Military Intelligence College, January 1996), 4.

[8] S. Card, J. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, 1999.

[9] Central Intelligence Agency. *A Consumer's Guide to Intelligence*. Washington, DC: Central Intelligence Agency, 1993.

[10] R. Chang, C. Ziemkiewicz, T. M. Green, and W. Ribarsky. Defining insight for visual analytics. *IEEE CGA.*, 29(2):14–17, 2009.

[11] K. Charmaz. Qualitative interviewing and grounded theory analysis. In *Handbook of Interview Research: Context and Method*, (J. F. Gubrium and J. A. Holstein, eds.), CA: Sage Publications, 2002.

[12] Chappell AR, Cowell AJ, Thurman DA, Thomson JR. Supporting mutual understanding in a visual dialogue between analyst and computer. *Proceedings of Human Factors and Ergonomic Society 48th Annual Meeting* (New Orleans, LA), Human Factors and Ergonomics Society: Santa Monica, CA, September 2004; 376–380.

[13] G. Chin, O. A. Kuchar, and K. E. Wolf. Exploring the analytical processes of intelligence analysts. In *ACM CHI*, pages 11–20, April 2009.

[14] Cicero. Language Computer Corp. http://www.language computer.com/ (accessed 01 March 2012).

[15] Compendium. http://www.compendiuminstitute.org (accessed 11 September 2010).

[16] Department of Intelligence Studies at Mercyhurst College. http://intel.mercyhurst.edu/ (accessed March 15, 2011).

[17] B. Dervin. Information as a User Construct: The Relevance of Perceived Information Needs to Synthesis and Interpretation. In Qualitative Research in Information Management. J. D. Glazier and R. R. Powell. Englewood, CO, Libraries Unltd., 61-84, 1983.

[18] B, Dervin. Sense-making Theory and Practice: An Overview of User Interests in Knowledge Seeking and Use. J. of Knowledge Management 2, 2, 36-46, 1998.

[19] B, Dervin. An overview of sensemaking research: concepts, methods and results. In: *International Communications Association Annual Meeting*. Dallas, Texas, 1983 & 2002.

[20] D. H. Dearth. "National Intelligence: Profession and Process," in *Strategic Intelligence: Theory and Application*, eds. D. H. Dearth and R. T. Goodden, 2d ed. (Washington, DC: Joint Military Intelligence Training Center), 17, 1995.

[21] A. B. Eli and J. Hutchins. Intelligence after Intellipedia: improving the push pull balance with a social networking utility. Research Report in *Information Science and Technology Directoratet*. Defense Technical Information Center, February 2010.

[22] C. Forsell and J. Johansson. An Heuristic Set for Evaluation in Information Visualization, In *ACM AVI*, pages 199-206, 2010.

[23] GATE − General Architecture for Text Engineering. http://gate.ac.uk/ (accessed 01 March 2012).

[24] Google reader. www.google.com/reader (accessed 11 September 2010).

[25] D. Gotz, M. X. Zhou, and Z. Wen. A study of information gathering and result processing in intelligence analysis. In IUI 2006 Workshop on IUI for Intelligence Analysis, 2006.

[26] J. Heer and M. Agrawala. Design considerations for collaborative visual analytics. In *Information Visualization*, vol. 7, no. 1, pages 49–62, 2008.

[27] J. G. Heidenrich. The State of strategic intelligence. In *Studies in Intelligence*, 51(2), 2008.

[28] R, Heuer. *Psychology of Intelligence Analysis*. Center for the Study of Intelligence, Central Intelligence Agency: Washington, DC, 1999.

[29] R. J. Heuer and R. H. Pherson. Structured Analytic Techniques for Intelligence Analysis, CQ Press College, 2010.

[30] F, Hughes and D. Schum. Discovery-Proof-Choice, The Art and Science of the Process of Intelligence Analysis – Preparing for the Future of Intelligence Analysis, *JMIC Joint Military Intelligence College*, 2003.

[31] J. Hughes, V. King, T. Rodden, and H. Andersen. The role of ethnography in interactive systems design, *ACM interactions 2*, 2, 56-65, April 1995.

[32] J. Hughes, J. O'Brien. T. Rodden, and M. Rouncefield. Design with ethnography: A presentation framework for design, *Proc. of Design of Interaction Systems '97*, ACM Press, New York, 147–159, 1997.

[33] i2 – Analyst's Notebook. http://www.i2group.com/us (accessed 11 September 2011).

[34] IEEE VAST 2007 CONTEST. http://www.cs.umd.edu/hcil/VASTcontest07, 2007.

[35] Jigsaw website. http://www.cc.gatech.edu/gvu/ii/jigsaw (accessed 11 September 2010).

[36] R. Johnston. Analytic Culture in the United States Intelligence Community: An Ethnographic Study. Central Intelligence Agency: Washington, DC, 2005

[37] M. D. Jones. The Thinker's Toolkit: 14 Powerful Techniques for Problem Solving. Three Rivers Press, 1998.

[38] Y. Kang and J. Stasko. Characterizing the Intelligence Analysis Process: Informing Visual Analytics Design through a Longitudinal Field Study. Proceedings of IEEE VAST '11, Providence, RI, October 2011, pp. 21-30.

[39] Y. Kang, C. Görg, and J. Stasko. The evaluation of visual analytics systems for investigative analysis: Deriving design principles from a case study. In *IEEE VAST*, pages 139–146, Oct. 2009.

[40] Y. Kang, C. Görg, and J. Stasko. How can visual analytics assist investigative analysis? Design implications from an evaluation. In *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, May 2011.

[41] G. Klein, B. Moon, and R. R. Hoffman. Making sense of sensemaking 1: alternative perspectives. IEEE Intelligent Systems, 21(4), July/August 2006.

[42] G. Klein, B. Moon, and R. R. Hoffman. Making sense of sensemaking 2: A macrocognitive model. IEEE Intelligent Systems, 21(5), September/October 2006.

[43] L. Krizan. Intelligence Essentials for Everyone. Joint Military Intelligence College Occasional Paper Number 6, June 1999.

[44] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. Empirical Studies in Information Visualization: Seven Scenarios. In *IEEE Transactions on Visualization and Computer Graphics*, 2012. Appeared online: 30 November 2011.

[45] M. Lowenthal. *Intelligence: From Secrets to Policy*. Washington, DC: CQ, 2000.

[46] MediaWiki. www.mediawiki.org/ (accessed 11 September 2010).

[47] Mercyhurst. http://intel.mercyhurst.edu/ (accessed 11 September 2010).

[48] D. R. Millen. Rapid ethnography: time deepening strategies for HCI field research. Proceedings of *ACM DIS'00 (New York, New York, August 17-19, 2000)*, pages 280-286. ACM Press, August 2000.

[49] Mindmeister. http://www.mindmeister.com/ (accessed March 15, 2011).

[50] E. Morse, M. P. Steves, and J. Scholtz. Metrics and Methodologies for Evaluating Technologies for Intelligence Analysts, In Proc. *Conference on Intelligence Analysis*, 2005.

[51] T. Munzner. A nested model for visualization design and validation. In *IEEE Transactions on Visualization and Computer Graphics*, 15(6), pages 921-928. IEEE Press, 2009.

[52] National intelligence estimates. http://www.cfr.org/iraq/national-intelligence-estimates/p7758#p3 (accessed June 07, 2011)

[53] D. Norman. The Psychology of Everyday Things. Basic Books, New York, 1988.

[54] D. R. Norman. Websites you can trust. In *American Libraries*, 37(7), page 36. American Library Association, Aug 2006.

[55] G. M. Olson, J. S. Olson, M. Carter, and M. Storrøsten. Small group design meetings: An analysis of collaboration. In *Human-Computer Interaction*, 7, pges 347-374. 1992.

[56] A. Perer and B. Shneiderman, Integrating statistics and visualization: case studies of gaining clarity during exploratory data analysis. In *ACM CHI*, pages 265–274, April 2008.

[57] J. J. Peterson. Appropriate factors to consider when assessing analytic confidence in intelligence analysis (Master's Thesis). Retrieved from http://scip.cms-plus.com/files/Resources/Peterson-Appropriate-Factors-to-Consider.pdf on March 15, 2011.

[58] O. J. Pinon, D. N. Mavris, and E. Garcia. A Visual Analytics Approach to the Qualitative Comparison of the SESAR and NextGen Efforts. *9th AIAA Aviation Technology Integration and Operations ATIO Conference*, 2009.

[59] P. Pirolli and S. K. Card. Information foraging. *Psychological Review, 106*, 643-675, 1999.

[60] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *International Conference on Intelligence Analysis*, May 2005.

[61] C. Plaisant. The challenge of information visualization evaluation. In *AVI*, pages 109–116, May 2004.

[62] P. Proulx, S. Tandon, A. Bodnar, D. Schroh, R. Harper, and W. Wright. Avian Flu Case Study with nSpace and GeoTime. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology (VAST'06)*. IEEE, 2006.

[63] Research and Development Descriptive Summaries (RDDS). http://www.dtic.mil/dtic/stresources/budgetNplanning/rdds_desc.html (accessed February 8, 2012).

[64] A. Robinson. Collaborative synthesis of visual analytic results. In *IEEE VAST*, pages 67–74, Oct. 2008.

[65] A. Rose, C. Plaisant, and B. Shneiderman. Using ethnographic methods in user interface re-engineering, *Proc. DIS '95: Symposium on Designing Interactive Systems,* ACM Press, New York, 115–122, 1995.

[66] D. M. Russell, M. J. Stefik, P. Pirolli, and S. K. Card. The Cost Structure of Sensemaking. In *INTERACT '93 and CHI '93, Conference on Human Factors in Computing Systems*, ACM Press, 269-276, 1993.

[67] Sanfilippo A, Baddeley B, Cowell AJ, Gregory ML, Hohimer R, Tratz S. Building a human information discourse interface to uncover scenario content. *Proceedings of the 2005 International Conference on Intelligence Analysis* (McLean, VA), https://analysis. mitre.org/proceedings/index.html, May 2005.

[68] P. Saraiya , C. North , and K. Duca. An Insight-Based Methodology for Evaluating Bioinformatics Visualizations, IEEE Transactions on Visualization and Computer Graphics, v.11 n.4, p.443-456, July 2005.

[69] P. Saraiya , C. North , V. Lam, and K. A. Duca. An Insight-Based Longitudinal Study of Visual Analytics, IEEE Transactions on Visualization and Computer Graphics, v.12 n.6, p.1511-1522, November 2006.

[70] J. Scholtz, E. Morse, and T. Hewett. In depth observational studies of professional intelligence analysts. In Presented at Human Performance, Situation Awareness, and Automation (HPSAA), 2004.

[71] J. Scholtz. Beyond usability: Evaluation aspects of visual analytic environments. In *IEEE VAST*, pages 145–150, Oct. 2006.

[72] J. Scholtz. Developing guidelines for assessing visual analytics environments. In *Information Visualization*, vol. 10, no. 3, pages 212–231, 2011.

[73] A, Selvin and S. B. Shum. Narrative, sensemaking, and improvisation in participatory hypermedia construction. Paper for CHI 2008 Sensemaking workshop, Florence, Italy, 2008.

[74] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *IEEE Visual Languages*, pages 336–343, Sept. 1996.

[75] B. Shneiderman and C. Plaisant. Strategies for evaluating information visualization tools: multi-dimensional in-depth long-term case studies. In *BELIV*, pages 1–7, May 2006.

[76] R. Spence. Information Visualization. Addison-Wesley, 2001.

[77] R. R. Springmeyer, M. M. Blattner, and N. L. Max. A characterization of the scientific data analysis process. Proceedings of IEEE VIS'92, pages 235-242. IEEE Press, 1992.

[78] J. Stasko, C. G¨org, and Z. Liu. Jigsaw: supporting investigative analysis through interactive visualization. In *Information Visualization*, vol. 7, no. 2, pages 118–132, 2008.

[79] J. Stasko, C. G¨org, and Z. Liu. Sensemaking across Text Documents: Human-Centered Visual Exploration with Jigsaw, CHI '08 Workshop on Sensemaking, Florence, Italy, April 2008.

[80] A. Strauss and J. Corbin. Basics of qualitative research: Grounded theory procedures and techniques. Sage Publications: Newbury Park, Calif, 1990.

[81] J. J. Thomas and K. A. Cook. *Illuminating the Path*. IEEE Computer Society, 2005.

[82] Transana. http://www.transana.com (accessed 11 September 2010).

[83] G. F. Treverton. Reshaping National Intelligence in an Age of Information. Cambridge: Cambridge University Press, 2001.

[84] Typewith.me http://typewith.me (accessed March 15, 2011).

[85] K. Wheaton. Wikis in Intelligence. Unpublished manuscript, 2011.

[86] J.A. Wise, J.J. Thomas, K. Pennock, D. Lantrip, M. Pottier, A. Schur, and V. Crow. Visualizing the non-visual: spatial analysis and interaction with information from text documents. *Proceedings of IEEE Symposium on Information Visualization*, Atlanta, GA, pages 51-58, October 1995.

[87] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort. The sandbox for analysis: Concepts and Methods. *Proceedings of ACM CHI '06* (Montreal, Quebec), ACM: New York, NY, 801–810, April 2006.

[88] P. Zhang. Sensemaking: Conceptual changes, cognitive mechanisms, and structural representations. A qualitative user study. Available from ProQuest Dissertation and Theses Database (UMI No. 3409887), 2010.

[89] Zotero. http://zotero.org (accessed 11 September 2010).