Causeworks Collaboration: Simultaneous Causal Model Construction and Analysis

Thomas Kapler
Uncharted Software, Toronto
tkapler@uncharted.software

Derek Gray
Uncharted Software
dgray@uncharted.software

Holland Vasquez
Uncharted Software
hvasquez@uncharted.software

William Wright
Uncharted Software
bwright@unchated.software

ABSTRACT

Military planners use "Operational Design" (OD) methods to develop an understanding of systems and relationships in complex operational environments. Here, we present Causeworks, a visual analytics application for OD teams to collaboratively build causal models of environments and use analytics to understand and find solutions to affect them. Collaborative causal modelling can help teams craft better plans, but there are unique challenges in developing synchronous collaboration tools for building and using causal models. Collaboration systems typically organize information around varying degrees of synchronization between data "values" and user "views." Our contribution is in extending this collaboration framework to include analytics as *layers* that are by nature derived *from* the data values but utilized and displayed temporarily as private views. We describe how Causeworks overlays analytics inputs and outputs over a shared causal model to flexibly support multiple modeling tasks simultaneously in a collaborative environment with minimal state management burden on users.

CCS CONCEPTS

• **Human-centered computing** → Visualization; Visualization systems and tools.

KEYWORDS

Collaborative visualizations, Collaborative and social computing

ACM Reference Format:
Thomas Kapler, Derek Gray, Holland Vasquez, and William Wright. 2021. Causeworks Collaboration: Simultaneous Causal Model Construction and Analysis. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '21 Extended Abstracts), May 08–13, 2021, Yokohama, Japan.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3411763.3451748

CHI '21 Extended Abstracts, May 08–13, 2021, Yokohama, Japan

ACM ISBN 978-1-4503-8095-9/21/05…$15.00
https://doi.org/10.1145/3411763.3451748

1 INTRODUCTION

While visual analytics tools have traditionally been designed for a single user to interact with the application, complex problems call for *collaborative visualizations* [7, 13]. Collaborative visualization is "the shared use of computer-supported visual representations of data by multiple users with the goal of contributing to joint team cognition" [7]. Such applications have become an increasingly important within the visualization community, as they have the potential to enhance human information processing and increase the efficiency in which a task can be completed by bringing together many experts that can contribute towards the common goal of understanding the phenomenon under investigation [13].

The Defense Advanced Research Project Agency's Causal Exploration (CX) program aims to develop tools to aid *Operational Design* (OD) teams in constructing *causal models* of complex, *wicked* problems [5]. Wicked problems are unstructured and difficult to define and characterized by uncertainty about the problem scope and appropriate objectives [15]. Examples include climate change, food scarcity, and operational planning problems with political, social, economic, and other non-military factors [17, 18]. Causal Modelling provides a way of understanding systems, subsystems, how they operate dynamically, and the forces that drive change [12]. As part of the CX program we developed Causeworks a visualization platform for collaborative knowledge construction and solution development through causal modelling [9]. Causeworks allows OD teams to work collaboratively with computer assistance to 1) document and refine the team's knowledge of the problem space, 2) execute models and simulations to explore the consequences of particular interventions, and 3) create, discuss, and debate potential solutions. We begin by providing a brief background on OD methodology, relate Causeworks to other collaborative tools, and outline the tasks that Causeworks facilitates. We then elaborate on the collaboration architecture that allows model-building, analytics execution and solution development to be performed in parallel.

2 MILITARY OPERATIONAL DESIGN

Military operations occur in increasingly complex environments. Accordingly, there is a growing need for planners to develop a holistic understanding of the governments, population, security forces, violent non-state actors and other factors that characterize the environment. To achieve this understanding, military planners engage in OD, a method for applying critical and creative thinking

to understand, visualize, and describe unfamiliar problems that permeate an operational environment [1, 8].

During the OD process, a diverse team of military planners and domain experts will engage in extensive discussions with the goal of framing their operational environment, defining their problem systems, and developing approaches to transforming the problem area [1, 8]. To achieve this, the team will brainstorm potential environmental factors and identify key relationships between them. This is a creative process. The team will strive to generate as many ideas as possible based on their experience, augmented with supporting research to identify additional factors and relationships. The team will then identify key system factors and develop approaches for how to influence the system to achieve mission objectives.

There are several limitations to the OD process that were observed by CX system designers attending traditional OD exercises. OD teams currently complete the brainstorming and research efforts manually, and as such the generation of ideas may be limited in scope to those the team identifies through discussions and those identified through manually searching through relevant documents and articles. Additionally, the design team relies on whiteboards, post-it notes, and butcher paper to document their conceptualizations. This poses challenges to preserving design artifacts, and the current process does not leverage computational techniques to improve quality and coverage of the resulting solution. The goal of Causeworks is to inject computer assistance, causal modelling and analytics into the OD process to improve solution development.

3 RELATED WORK

Generally, causal model graphs consist of nodes representing causal factors, and links the relationships between them [12]. While causal modelling software exists that provides tools for creating and running causal systems [e.g., 16, 19], such applications do not support multi-user collaboration, and are limited in their ability to provide users with analytical tools that can assist in system and solution development [10]. Causeworks provides machine assistance for building models while encouraging human learning, analytical tools for understanding system characteristics, and finally, system simulations for projecting futures and generating solutions.

Collaborative visualization tools typically focus on shared awareness of data, views and interactions [2, 11], NOT the shared use of analytics. Analytic tools present unique challenges for coordinating interactions that affect the value (raw data), and the view (visualization) [3]. Value interactions modify the underlying data by adding, changing or deleting subsets of data, whereas view interactions only change how the data is visualized by users. In the context of tools for collaboration, view and value interactions can occur in shared and/or private views. To illustrate, we consider DARPA's Command Post of the Future (CPOF) for collaborative command, control, and decision making [4]. In CPOF, users are provided with both private and public views of public and private data. All views are dynamically updated to reflect all changes in the values data. For instance, multiple users can freely edit public data in a public shared map view at the same time. Alternatively, one can add private "view-only" data to a public workspace for all to see but not edit. As Chi and Riedl (1998) noted, the distinction between the value and view interactions is not always clearly defined, and a system such as CPOF includes a host of controls and settings for managing user privacy and sharing permissions across values and views that add significant complexity and training overhead.

Information generated by analytic functions adds a level of complexity in that analytics results are derived from shared data values, as opposed to being shared data values themselves. Nor are they a user's view of data values. And because analytics results become stale if underlying data values change, they are by nature short-lived. Value-derived analytics thus present unique challenges for collaboration. In Causeworks, we use a unique approach involving private viewing of analytics results for each user, however in one case we enable shareable *analytics inputs* to enable synchronous collaboration on analytics. Section 5 elaborates on our design for analytics usage within the Causeworks collaborative system, and how it supports various concurrent and overlapping user tasks with minimal user state management.

4 CAUSEWORKS WORKFLOW

Causeworks was developed by Uncharted Software Inc as the human machine interface component for the DARPA CX program. It allows a team of operational designers to build a computational causal model, provides machine reading to assist the team with learning about the problem space and growing the model, and incorporates analytics for solution development—all of which are new to OD methods that today are based on human verbal team discourse and sketching on paper and chalkboards. The Causeworks interface is centred around a digital whiteboard that is the canvas on which causal models are constructed, arranged and refined and also where analytics are executed and displayed. The following subsections describe user workflow and the role of Causeworks analytics in the context of the three stages of OD.

4.1 Frame the Environment

To support the OD process, Causeworks enables teams to sketch ideas for factors and relationships on the whiteboard as the first step in building a causal framework. Causeworks includes a corpus of documents related to the team's planning scenario, including intelligence reports, news stories, papers, social media posts and more. From this data, 3rd party Natural Language Processors and causal modelling engines automatically build a database of scenario-specific, pre-generated factors and relationships. These computational causal factors can be placed onto the whiteboard and tied into the users' model sketch, transforming it into a functional causal model that supports computational analysis. In addition, the Causeworks "Suggestion" system automatically identifies relationships and factors that users can add to the model, thus leveraging analytics towards the creation of a more comprehensive model. However, the automated factor generation system may not find factors that users need to represent their environment, so users can create factors from scratch, manually set their initial value and trend, and include them in the model. The team's shared causal model consists of all factors and relationships included on whiteboards, and all users synchronously view the same set of whiteboards, though each user can individually control which whiteboard they are looking at, and its zoom/pan position.

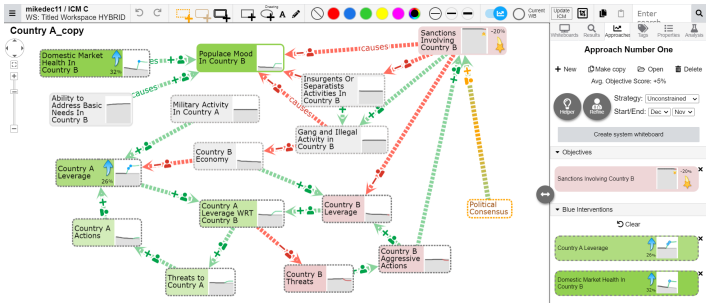

Figure 1: Causeworks human machine interface. The whiteboard displays the team's shared layout of the causal model but with individual zoom and pan state. Through the side panel each user can execute analytics on the shared team model. Analytics results are only visible to those who execute them, and do not affect other users' activities. The Approach panel (shown here) lets users view and edit shared interventions and objectives (analytics inputs), providing synchronous collaboration when viewing the same Approach.

4.2 Understand the Problem

Once a model is built, analytical tools are used to further understand and evolve it. The most important analytic is the projected value of a factor, calculated by executing simulations of the causal model over a series of time steps. System simulations automatically execute each time the model is changed, and the resulting "baseline" projections are shown in thumbnail timelines on each factor. Users can apply interventions on factors (i.e., propose an increase or decrease in its value at a point in the future, shown as blue dots in timelines) to see how downstream factor values respond. Simulations also execute each time interventions are applied, with the resulting "What-if" projections overlaid on the baseline to show the difference. An increase or decrease in a factors' projected value over time as a result of interventions is displayed on the factor using a green-red color scale.

Additionally, Causeworks integrates analytics that can be executed by users on-demand at any time. These include "Factor Sensitivity", "Most Impact", "Causal Loop Finder" and others that help users understand influences and structural properties of the model that impact how changes propagate. These analytics are executed individually by and for each user, with results overlaid privately on their whiteboard model view with magenta bar-scale indicators (see figure 2D). Each user specifies their own analytics inputs and sees results privately, but the analytics always execute over the shared model. For example, if two users run the same analysis with the same inputs at the same time, they will see the same results. However, because the model may be under constant development by other users, these analytics results are assumed to be short-lived.

4.3 Develop a Solution

The goal of OD is to develop solutions to achieve mission objectives. In Causeworks, users enter their objectives as target values for a factor at some point in the future (e.g., increase defensive capabilities by 20% in 6 months). These are displayed as orange stars in timelines. Sets of objectives and the proposed interventions to achieve them are encapsulated in a shared data object called an "Approach". Each user can select an Approach to view and edit via the Approach management list. All the objective and intervention factors for the selected Approach are displayed in the Approach tab (see Figure 1) along with a score for how well objectives are being met. Multiple users can select the same Approach at the same time.

Analytics such as "Sensitivity" can help users identify which factors to intervene on to impact their objective factors. The "Approach Helper" analytic goes one step further by automatically generating solutions to meet objectives. It uses objectives in the active Approach as inputs and allows users to set constraints on the timing and size of interventions it proposes. When executed, it creates new interventions optimized to precisely meet objectives, and adds them to the current Approach.

5 CHALLENGE: SITUATING ANALYTICS LAYERS WITHIN PUBLIC AND PRIVATE VIEWS

Teams using Causeworks can collaborate synchronously or asynchronous on building and using causal models. A key challenge however is that one user may be editing and improving the shared model, while another is simultaneously using analytics to develop Approaches for a planning problem. This can cause conflicts; for example, an Approach that appears to meet objectives may no longer work, or work as well, after relationships and factors are changed in the model by another part of the team. In fact, all analytics results may only be valid for a short time because changes could be made to the underlying model at any time.

Our novel approach to address this is to separate the shared model data from analytics execution and results display such that analytics results are temporary and private in nature within the collaborative framework. In some cases however, we allow users to share analytics *input parameters* so they can work synchronously and see the same results when desired. A key goal of our collaboration policy design is that collaboration is transparent, and modeless, without complex sharing or permissions infrastructure overhead.

5.1 Public Shared Model

At a high-level, changes affecting the causal model are always viewed publicly, while tasks not involving changes to the model are

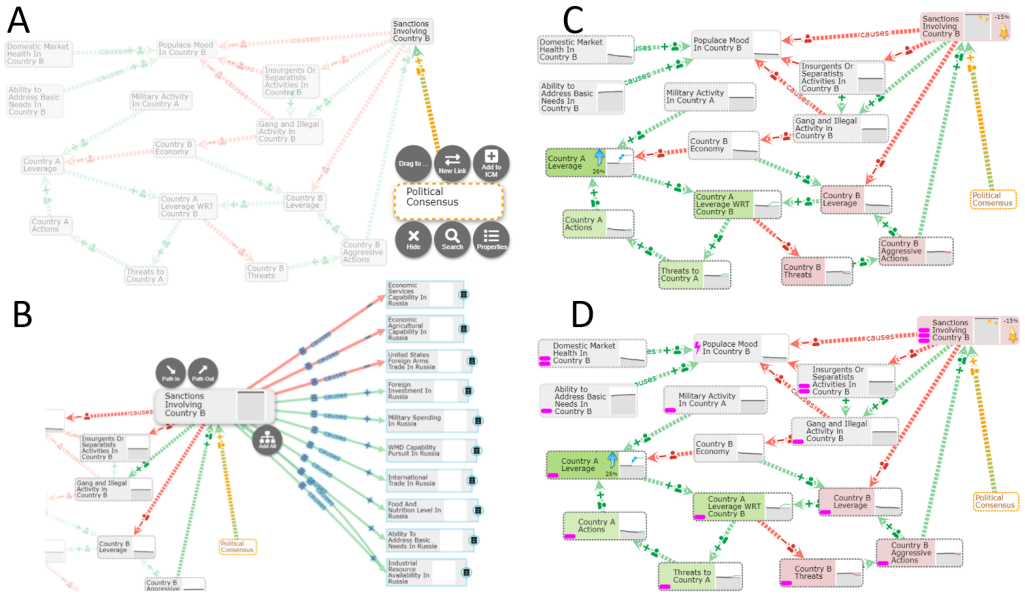

Figure 2: Four views of the same Causeworks whiteboard at the same time from 4 different users' perspectives: User A is adding a new "political consensus" factor that all other users immediately see. User B is exploring suggested relationships for a selected factor that only they can see (blue highlighted items). User C and D are looking at the same Approach, and therefore see the same "what-if" projected effects in green and red, but user D is also running sensitivity analysis on factor "Populace Mood in Country B" that only they can see (magenta annotations).

viewed privately. This policy sets the stage for how collaborating teams perform parallel efforts in editing and using model analytics. Model composition and whiteboards are public and any change to the model or a whiteboard is visible to all. For example, when an individual user adds, changes or remove factors or relationships in a whiteboard, updates are immediately sent and visualized for all users. This allows the team to distribute model construction among members, ensuring all users are working with the most up-to-date model.

5.2 Private Views

Certain tools in Causeworks operate privately. For example, the use of the property viewers, search tool and model suggestions are viewed privately for each user, allowing them to explore potential changes to the model privately before committing them. In addition, each user has their own private view of the shared whiteboards, (i.e., similar to Google Docs, multiple users can simultaneously edit the model, with each user being able to zoom the scale and scroll to different areas independently). The need for synchronized zoom/pan views across users is not considered critical, and significantly simplifies window state management across users. (The need for synchronized views, along with voice collaboration, can be best addressed via screenshare capabilities in Zoom, Google Hangouts, WebEx or Microsoft Teams.)

5.3 Analytics Layers: Between Public and Private

Causeworks analytics help users make immediate decisions about aspects of the model or Approach they are working on. Analytics inputs include the shared casual model and various user-entered parameters. Outputs are presented visually as privately-viewed overlays on users' whiteboards. Analytics are implemented as independent services executed by each client. Baseline and What-if projection analysis results update automatically for all users each time the model is updated. Other analytics, such as factor Suggestions, Sensitivity, Approach Helper, and Causal Loop Finder are executed *on-demand only*, with results viewable as a layer on the model view until dismissed. The reason that these do *not* update automatically is to allow users time to assess results in a stable state. Users are trained that analytics results are, by nature, short-lived due to potential model changes. In exercises, we found users comfortable with a low risk of short-term stale analytics results, as they are primarily used to gain rapid insights, and do not involve changes to the underlying model.

We think of analysis results as analogous to toggleable layers situated conceptually between private and public views (see Figure 3): Private in that they are executed and displayed for only one user, but public in that users will see the same results *if* they run the same analytic under the same conditions (i.e., at the same time). Analytic inputs, such as Approaches (which contain interventions and objectives), are treated more like model content, and are persisted as shared data across users. When multiple users choose to activate the same Approach, they see and edit the same objectives and interventions, and therefore view the same results from projections analytics. In this case, collaboration appears synchronized even though analytics execute independently. Figure 3 provides a schematic overview of collaboration categories and policy. This structure clearly defines the boundaries of collaboration, allowing

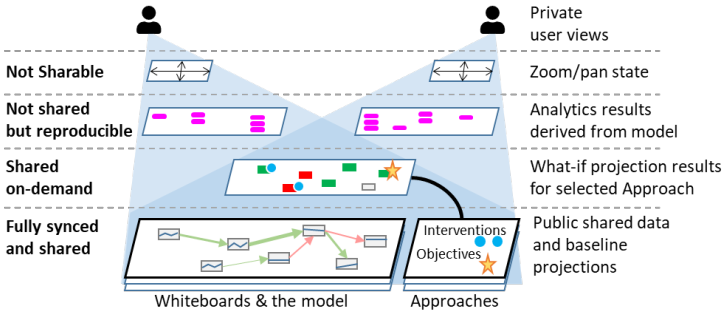

Figure 3: Schematic representation of analytics as "layers" between public and private data and the extent to which they are sharable.

utilization of analytics without users having to manage synchronization and sharing states with each other. It also eliminates the need for storage and redistribution of dated analytics results. Most important, it allows each team member to work independently, yet always with the most current version of the model.

5.4 Other Collaboration Support Tools

Causeworks provides additional features to support collaboration, including popup notifications of model changes, and display of most recent edits to properties and the users who make them. Users can also mark factors with colors or add custom tags that are shared across users. These are typically used to categorize elements according to user-defined criteria. In addition, text comments can be made on objects or directly on whiteboards. Causeworks also allows users to make copies of the model in cases where total separation from changes is desired, however users noted that it is generally preferable to work with a continuously evolving version of the model rather than a stable but potentially out-of-date version.

5.5 Workflow Observations

Multi-day evaluation exercises were conducted every 6 months with teams of five to ten people to inform the system design and ensure it meets users' needs. Each exercise involved fictional planning scenarios and the participation of OD experts, government–provided domain experts, and OD students from the US Military. Teams using Causeworks worked over several days and then presented their planning solutions.

When using Causeworks, operational designers self-organized to balance individual work and collaborative work according to physical environment and roles, skills, expertise and tasks. Collaboration occurred synchronously and asynchronously in both collocated and distributed settings. The ability to switch between tightly and loosely coupled work was leveraged when allocating tasks [6]. For instance, some team members worked synchronously around a single touchscreen display to identify factors, while individuals or smaller groups worked on laptop computers in different rooms to explore evidence and define relationships. Remotely located domain-experts were able to contribute to model and solution development. A key team decision concerned division of labor, and how to most efficiently organize more resources to contribute to time-sensitive tasks. For example, one group focused on building a causal model of political factors and another on the economy. In other cases, effort was divided by nation states, and after completing their separate models on different whiteboards, the team reconvened to combine all state models into a single regional whiteboard. They then proceeded to build relationships *between* countries. Teams also split into groups to game competing approaches (e.g., "red team", "best case", or "most likely"), and then compared their solutions before presenting to their commander. In all cases, analytics were used to help answer questions related to causal forces among the factors, and generally understand and validate that the systems behaved as expected.

Effects on the resulting model and solution quality have yet to be evaluated in a formal experiment, however additional exercises and evaluations to compare collaborative causal modelling using Causeworks against traditional operational design methods are planned. These evaluations are expected to inform future work.

6 CONCLUSIONS

In this paper we introduced Causeworks, a tool that allows distributed teams to collaboratively combine effort in building and using causal models for complex military planning problems. We described how an OD team defines the operational environment as a causal system, applies causal analytics to refine the model, and conceives interventions to meet planning objectives. Performing all these activities simultaneously presents unique challenges for coordinating collaboration. In our collaboration design, the model structure and layout are always shared and synchronized across all users, whereas analytics execution and results are privately viewable by individual users because they do not change the model data. We also described how separating analytics inputs as shared data objects enables effective opt-in synchronous collaboration and discussed observations about how OD teams organize collaborations by various tasks and roles.

ACKNOWLEDGMENTS

This work was funded by the Defense Advanced Research Projects Agency with contract # FA8650-17-C-7720. The views, opinions, and findings in this report are those of the authors and should not be construed as an official Department of Defense position, policy,

or decision, unless so designated by other official documentation. The authors wish to thank everyone in the broader team for their support and encouragement.

REFERENCES

[1] ATP 5-0.1, Army Design Methodology, HQ Dept of the Army, July 2015. Retrieved from: https://armypubs.army.mil/epubs/DR_pubs/DR_a/pdf/web/atp5_0x1.pdf

[2] Aruna D. Balakrishnan, Susan R. Fussell, and Sara Kiesler. 2008. Do visualizations improve synchronous remote collaboration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (April 2008). pp 1227-1236.

[3] Ed Huai-hsin, Chi, and John T. Riedl. An operator interaction framework for visualization systems. 1998. *In Proceedings IEEE Symposium on Information Visualization* (October, 1998) pp. 63-70. IEEE.

[4] Mei C. Chuah, and Steven F. Roth. 2003. Visualizing common ground. *In Proceedings on Seventh International Conference on Information Visualization* (July 2003) pp. 365-372.

[5] DARPA, Causal Exploration of Complex Operational Environments (Causal Exploration). 2018. Retrieved from: https://www.darpa.mil/program/causal- exploration

[6] Carl Gutwin, and Saul Greenberg. 1998. Design for individuals, design for groups: tradeoffs between power and workspace awareness. *In Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work*, (Nov. 1998) pp. 207-216.

[7] Petra Isenberg, Niklas Elmqvist, Jean Scholtz, Daniel Cernea, Kwan-Liu Ma, and Hans Hagen. 2011. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization* 10, 4 (October 2011), 310-326.

[8] Joint Publication JP 5-0, Joint Planning, CJCS, June, 2017. Retrieved from: https://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/jp5_0_20171606.pdf

[9] Thomas Kapler, Derek Gray, Holland Vasquez, and William Wright. 2021. CauseWorks: A framework for transforming user hypotheses into a computational causal model. *In IVAPP 2021: 12th International Conference on Information Visualization Theory and Applications* (in press).

[10] Thomas Kapler and William Wright. 2018. Challenges in visualizing complex causality characteristics. *IEEE Pacific Visualization.*

[11] Narges Mayar and Melanie Troy. 2014. Supporting communication and coordination in collaborative sensemaking. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (November 2014) 1633-1642.

[12] Judea Pearl. 2009. Causal inference in statistics: An overview. *Statistical Survey*, 3 (October 2009), 96-146.

[13] James J. Thomas, and Kristin A. Cook. 2006. A visual analytics agenda. *IEEE Computer Graphics and Applications* 26, 1 (January 2007) 10-13.

[14] Rajeev R. Raje, Michael Boyles, and Shiaofen Fang. 1998. CEV: collaborative environment for visualization using Java RMI. *Concurrency: Practice and Experience* 10, 11-13 (September 1998), 1079-1085.

[15] Horst W.J. Rittel, and Melvin M. Webber. Dilemmas in a general theory of planning.1973. *Policy Sciences* 4,2 (June 1973), 155-169.

[16] Jun Wang, and Klaus Mueller. 2017. Visual causality analysis made practical. In *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)* (pp. 151-161).

[17] James J. Wirtz. 2017. Life in the "Gray Zone": observations for contemporary strategists. *Defense & Security Analysis* 33, 2 (April 2017), 106-114.

[18] Xiaodong Zhang, and Velimir V. Vesselinov. 2017. Integrated modeling approach for optimal management of water, energy and food security nexus. *Advances in Water Resources* 101 (March 2017), 1-10.

[19] Zhiyuan Zhang, Keven T. McDonnell, Erez Zadok, and Klaus Mueller. 2014. Visual correlation analysis of numerical and categorical data on the correlation map. *IEEE Transactions on Visualization and Computer Graphics, 21*(2), 289-303.