CauseWorks: A Framework for Transforming User Hypotheses into a Computational Causal Model

Thomas Kapler[a], Derek Gray[b], Holland Vasquez[c] and William Wright[d]
Uncharted Software, Toronto, Canada

Keywords: Causality Analysis, User-driven Modelling.

Abstract: Causal Model building for complex problems has typically been completed manually by domain experts and is a time-consuming, cumbersome process. Operational Design defines a process of rapid, structured discourse for teams to envision systems and relationships about complex, "wicked" problems, however, the resulting models are simple diagrams produced on whiteboards or slides, and as such, do not support computational analytics, thus limiting usefulness. We introduce CauseWorks, an application that helps operators "sketch" complex systems and transforms sketches into computational causal models using automatic and semi-automatic causal model construction from knowledge extracted from unstructured and structured documents. CauseWorks then provides computational analytics to assist users in understanding and influencing the system. We walk through human-machine collaborative model-building with CauseWorks and describe its application to regional conflict scenarios. We discuss feedback from subject matter experts as well as lessons learned.

1 INTRODUCTION

Causal reasoning forms the basis for most complex forms of reasoning, facilitating hypotheses, inferences, explanations, and problem-solving (Jonassen & Ionas, 2006). This is true for virtually all domains: causal reasoning permeates science, engineering, public health, finance, medicine, and military planning and decision making (Keim et al., 2010; Sedig et al., 2012; Schmitt, 2017). Indeed, understanding causality, the influence by which one factor or cause contributes to the production of another factor has become an important topic in visualization within the last decade (Pearl, 2009.) Causality provides a way of understanding systems, subsystems, how they operate dynamically, their underlying characteristics, and forces that drive change. As such, there is an increasingly important role for visual analytics, to support the user's understanding of causality through interactive visual interfaces.

The application of causality visualization methods to complex *wicked problems*, while emerging, is still limited (Proulx et al., 2019). Wicked problems can be difficult to define, are not discreet, and potential solutions are intertwined and complex (Rittel & Webber, 1973). Government instability, gray zone conflicts (Wirtz, 2017), food security (Zhang & Vesselinov 2017), and climate change (Gil et al., 2018) are examples of ill-structured, dynamic situations which are poorly understood and where solutions are neither readily available nor have consensus. Most causality visualization systems require a pre-existing model for the user to explore (Sedlmair et al., 2012). However, with wicked problems, key knowledge may reside within the mind of the domain expert. As such, in addressing wicked problems modelling efforts are often completed by hand, involving the knowledge of many domain experts, who may be novices in causal modeling. Inherently, this process can be time consuming, requiring access to domain and modelling experts, which are costly and hard to procure (McPherson et al., 2007).

[a] https://orcid.org/0000-0003-1957-9131
[b] https://orcid.org/0000-0003-4009-2125
[c] https://orcid.org/0000-0001-9048-3445
[d] https://orcid.org/0000-0002-5685-3999

Kapler, T., Gray, D., Vasquez, H. and Wright, W.
CauseWorks: A Framework for Transforming User Hypotheses into a Computational Causal Model.
DOI: 10.5220/0010194300500063
In *Proceedings of the 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2021) - Volume 3: IVAPP*, pages 50-63
ISBN: 978-989-758-488-6

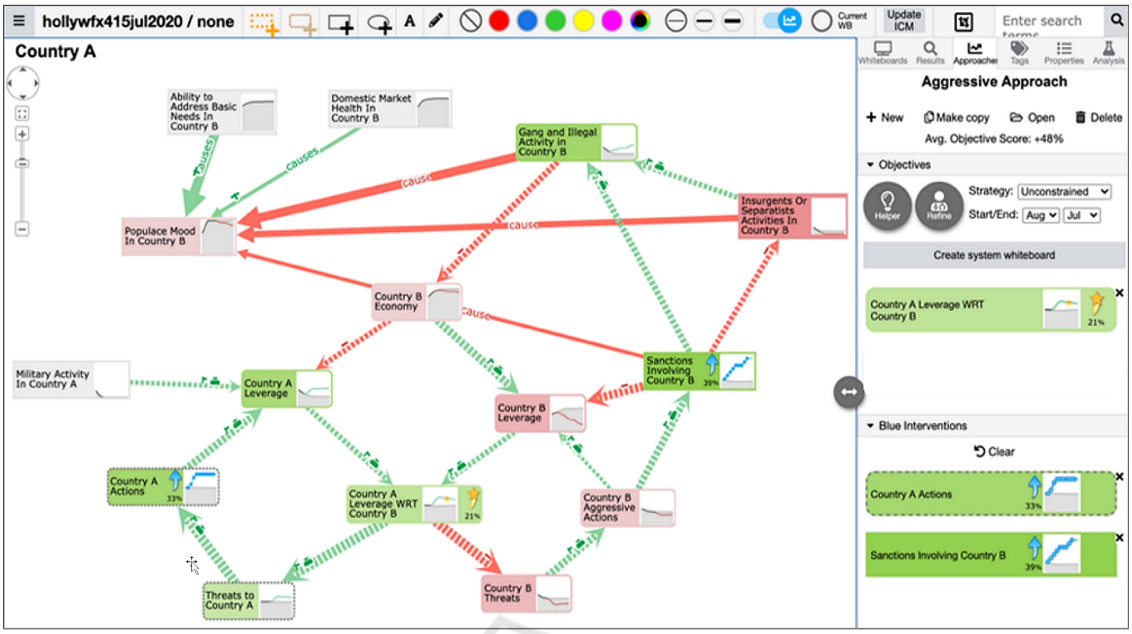

Figure 1: CauseWorks visual analytics interface for rapid creation of causal models showing a simple causal system of relationships between two countries. Side panel summarizes user objectives and the interventions made to achieve them. Green and red nodes indicate projected changes to factors resulting from interventions (in blue).

In developing solutions for wicked problems within the military, military planners will combine systems thinking, causal thinking and the Operational Design (OD) process (described in the subsequent section) to frame complex problems and arrive at potential solutions. The Defence Advanced Research Projects Agency's (DARPA, 2018) Causal Exploration (CX) Program, is focused on developing a causal modeling platform to aid expert military planners in decision-making in conflicts complicated by political, economic, social and other non-military factors where there is significant uncertainty about the problem and appropriate objective. In this paper, we present CauseWorks, a visual analytics interface developed by Uncharted Software for teams to apply causal modelling to complex problems. The problem domain of focus is regional conflict analysis; however, we note that the application of CauseWorks could extend to other domains. The key contribution in this paper is a framework to support OD experts who are novice modelers in building computational causal models for complex, wicked problems. The specific contributions include: 1) a method for capturing user-driven ideas and hypotheses, and connecting them with a knowledge base spanning thousands of documents, 2) a method for rapidly transforming hypotheses into a computational causal model backed by data, and 3) a method for interacting with and visualizing casual analytics within the context of the model to reveal system behaviors and assist in solution development. In the subsequent section, we provide additional background on OD to set the stage for the application of CauseWorks.

2 OPERATIONAL DESIGN

The United States military is increasingly conducting operations in complex, ill-structured environments that are characterized by a diverse, ambiguous set of actors, enemies, and unknowns, (Joint Publication JP 5-0, 2017). In order to develop the best course of action, military planners need to develop a comprehensive understanding of interconnected, complex systems such as the governments, population, security forces, and non-state actors that make up the environment (Schmitt, 2017). This understanding is achieved through OD, a method used by military planners involving team brainstorming, system sketching, and other methods, of characterizing the intertwined relationships between entities and factors involved in current and desired states.

During the OD process, a team of six to nine operational designers and domain experts will engage in extensive discussion with three goals: 1) framing the operational environment, 2) framing the problems that permeate the environment, and 3) creating approaches to transform the problem (ATP 5-0.1, 2015). This involves rapid brainstorming of potential variables on a whiteboard and developing lists of key actors and factors, and hypotheses about the relationships between them. The design team will supplement initial hypotheses with supporting evidence obtained from researching government documents, articles, and additional relevant source material. This process is completed manually, using physical whiteboards and post-it notes that document the team's conceptualization of the environment. Written narratives are composed in Microsoft (MS) Word documents. Pre-existing templated conceptual diagrams may be filled out in MS PowerPoint. This process is completed in a short period of time (i.e. days). Figure 2 provides an example of a system sketch created by SMEs using the OD process to solve a fictional problem.

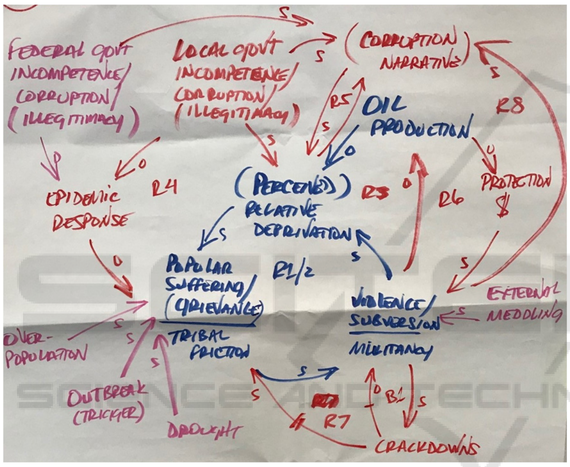

Figure 2: Example system sketch of a fictional scenario created during a user exercise with CX system designers and SMEs.

During planning exercises attended by the authors, several limitations to the current OD process were noted. First, domain experts across a range of specializations are often unavailable. Secondly, the design team identifies actors, systems, variables and proposes hypotheses manually, based on prior knowledge. As such, the quality of the resulting system concept model largely depends on the expertise, interpersonal "chemistry" and creativity of the design team's ability to identify relevant variables and connections between variables. In doing so, the team may rely on heuristics, which may be error prone and subject to human biases (Das & Teng, 1999). Additionally, in supporting hypotheses with evidence, the team manually searches through documents, which inherently is a cumbersome, time-consuming process. The lack of time reduces the opportunity for an evidence-based weighing of alternatives. Additionally, the scale and scope are limited. OD exercises over three days typically result in systems with ten to twenty factors. Current and future states are envisioned, but system dynamics and changes in factors over time are not strongly considered. Finally, because this process is completed by hand, the end product does not result in a computational causal model with which to leverage analytical tools. Validation of conceptual models and solutions are performed manually by reviewing static, slide-based products with senior staff.

3 RELATED WORK

In this section, we relate our work to relevant tools that support causal reasoning and causality visualization, including tools for mind mapping, argument mapping, and causal modelling.

Mind mapping is a tool for visualizing one's mental model of a problem. While traditionally, mind-mapping has been completed by hand on whiteboards, several digital tools have been developed to facilitate this activity (e.g., Subramanian & Krishnamurth, 2020; Shih et al., 2009). However, as Chen, and colleagues (2020) point out, few of these provide computer-based support for idea and hypotheses generation, or assist users with the learning process. With this limitation in mind, they developed *QCue*, which provides the user with system-generated ideas to assist the user in exploring topics as they develop a mind-map and user-elicited queries that allow the user to explore a given topic in depth. Wright et al., (2017) presented Argument Mapper, for developing hypotheses in the mind-mapping process through computer-based analytics, with the goal of reducing human cognitive bias. Analysts construct argument trees composed of hypotheses, sub-hypotheses, assumptions and evidence, and assess the credibility and evidence of each item. Support for the upper-level hypotheses is automatically calculated. CauseWorks provides tools for mind mapping, allowing users to sketch their hypotheses about causal relationships and factors within a given system and then automatically finds supporting or refuting evidence. Like *QCue,* users are presented with suggestions to help facilitate learning and encourage creativity, however in CauseWorks, suggestions include true model additions, thereby expanding the scope of the causal model.

In the context of causal modelling, automated algorithms exist for extracting causal relationships from factors within multivariate data sets (Wang &

Mueller, 2017). Factors are synonymous with variables and are attributes (characteristics) of an entity or their environment that influence the question of interest (McPherson et al., 2007). Causal relationships can be depicted in a directed acyclic graph (DAG), which consists of a set of nodes and links (Von Landesberger et al., 2011; Pearl, 2009). Nodes typically represent relevant factors while links represent the causal connections between factors. Link properties include direction, strength, and certainty (Bae et al., 2017), described through visual encodings including color, line width, fuzziness, and transparency. A central component of CauseWorks is a computational causal model. The model is displayed using nodes and links that visually encode different stages of model construction, as well as several static and dynamic model properties. In CauseWorks the user is not constrained to a DAG automatic graph layout, rather CauseWorks supports a user-driven layout with the option to leverage automatic graph layouts.

Causality cannot be computed from the data alone. Wang (2018) noted the need for the domain expert to interact with the causal graph. Zhang et al., (2015) developed an interactive correlation map for filtering edges with weak correlations. Wang and Meuller (2016) presented a platform in which users modify a causal graph by connecting nodes, assigning causal direction, deleting or marking edges as unknown. In Causemos (Proulx et al., 2020), expert modellers start with a knowledge base of extracted causal statements displayed in a causal graph, and then refine it. In CauseWorks, causal relationships and factors are extracted from source documents, resulting in a model database that users can search, pull content from and edit to construct a computational model.

In general, causality visualization methods are limited in expression, scale, dimensionality, and do not provide sufficient support for "what if" analysis, injecting interventions, and development of solutions to impact system behaviour (Kapler & Wright, 2018). The CauseWorks system offers advanced causal analytics to assist the team with answering sensemaking questions (e.g., "what-if" and "how-to") in addition to considering multiple approaches to solution development. In the next section we describe the implementation of CauseWorks.

4 CAUSEWORKS SYSTEM DESCRIPTION

To understand the model construction interface of CauseWorks, it is necessary to have a basic understanding of the underlying system, and how it automatically generates causal model elements for users to leverage in constructing their own models.

CauseWorks is composed of a web client and a server (Node.js), both written in the Javascript language (see Figure 3). The client application components are built on the EmberJS framework, with both D3.js and Cytoscape.js visualization libraries for graph rendering, augmented with SVG-elements for specialized visuals and interactions. The CauseWorks server marshals communication between the client and back-end analytics. These back-end analytics consist of 3rd party program performers (details available from DARPA). Firstly, three Natural Language Processing reader components extract events, causal assertions and associated locations and actors from a shared corpus of thousands of documents related to a given problem domain (including expert-authored reports, news media articles and open source material).

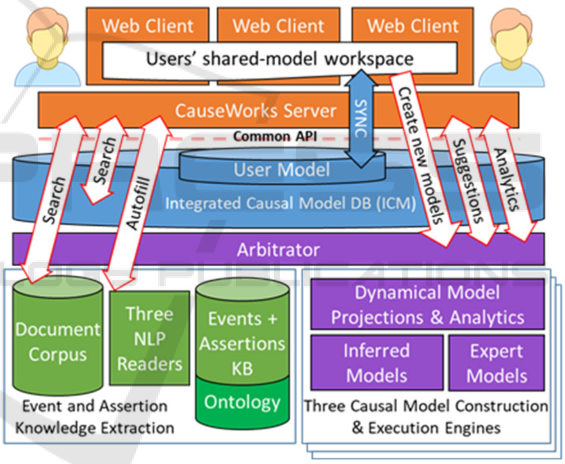

Figure 3: CauseWorks system architecture includes a federation of three NLP readers and three causal modelling framework performers. The ICM combines all models into a single database, and the Arbitrator routes work among readers and frameworks. Analytics, search, and other tools are exposed to CauseWorks as services via a common API.

A common ontology is utilized to align readers, and to associate events with "FactorTypes" curated for a specific problem domain. Secondly, reader-results are merged and then processed into a causal model database by three causal-model frameworks that provide model construction, simulation and analytic functions. Factors are created where events match FactorType definitions. Factor value and trend is inferred from event trends or structured data timeseries. Then, causal relationships between factors are generated based on the FactorTypes of the cause and effect events that comprise causal assertions. Strength is inferred from assertion count, confidence and, in some cases, correlations between factor historical trends. The process for associating extracted data with user-created Factors leverages users input values for FactorType, actors, and locations. In addition, CauseWorks includes model frameworks for pre-built subsystems (e.g. economy model for a nation-state). Figure 3 provides an overview of the CauseWorks system architecture.

5 SYSTEM OBEJECTIVES AND DESIGN REQUIRMENTS

Key human performance factors for OD are focused on human learning, thinking and creativity. Important performance objectives include: 1) increase learning and comprehension about new, unfamiliar, unknown situations, 2) engage human critical thinking in a group debate that questions and tests alternative perspectives, 3) encourage creativity by considering different perspectives and merging selected aspects, and 4) use group discourse as the catalyst to develop new ways of thinking about problems and identifying innovative solutions (Schmitt, 2006; ATP 5-0.1, 2015; Joint Publication JP 5-0; 2017)

In developing CauseWorks, we followed a user-centric approach, engaging with expert OD practitioners, trainers and students. Early design began with structured interviews and system concept sketches. Full-day, problem-focused exercises were performed for system designers to observe traditional OD teams working through regional conflict scenarios. Results included the following high-level objectives for CauseWorks functionality.

1. This system should support the rapid pace of OD team brainstorming and discourse without disruptive “care and feeding” of the software tools
2. The system should enable creation of an unconstrained, notional, causal system, i.e. a “sketch” system
3. The system should enable transformation of the hypothesized system into a computational causal model with minimal user effort or input.
4. The system should display the causal model in a manner that allows sense-making of causal structure, causal “flow”, predictions and impacts of changes.
5. The system should provide causal analytics to assist the team in understanding, validating, and improving the model, uncovering patterns and behaviours, and assist in developing a solution.

Note that one key assumption repeated by subject matter experts (SME’s) was that models are never objectively “correct”, but some *are* useful (Box, 1979). For OD of complex problems, supporting high-level thinking and introducing new factors for consideration is more important than the accuracy of a specific model. An initial CauseWorks application and analytics system was developed, followed by a

series of scenario-focused exercises conducted with SMEs to assess system performance, usability, and collect feedback to inform subsequent development cycles.

The following sections focus on CauseWorks HMI and human-machine workflow for creating and using causal models.

6 CAUSEWORKS VISUAL ANALYTICS WORKFLOW

The following sections describe how the affordances, interactions and visual encodings enable our envisioned machine assistance mantra of *“capture sketch hypotheses of a system, transform it into a causal model, and provide insights with causal analytics”*. While this workflow is described as a linear process, it is important to note that a team can work through the process iteratively and move back and forth through each step.

The main purpose of CauseWorks is working with causal models. Accordingly, the design focus is on causal model construction and analytics presentation. We begin by presenting high-level interface components and key visual encodings, and subsequently discuss how to use the system.

6.1 General Workspace & Visual Encodings

Current OD methodology relies heavily on the use of physical whiteboards to allow teams to freely capture discussion points and sketch diagrams. CauseWorks is similarly centered about a digital whiteboard workspace.

The HMI is comprised of a whiteboard workspace, side panel, and main menu bar (see Figure 1). It is within the whiteboard workspace that the team sketches and constructs their causal model. Sketching tools in the toolbar include the creation of user nodes, groups, and simple shapes. CauseWorks also includes tools for grouping nodes, and then collapsing or hiding group contents to declutter or simplify the display. Groups can also combine factor values into a single aggregate factor. A small right-side panel provides tabs to access additional functions and thus avoid floating windows that obscure the whiteboard. Tabs within the side-panel allow the team to navigate whiteboards; search the document corpus and model database; edit factors and relationships and connect them with evidence; access analytic functions; and develop approaches to achieve objectives. These functions are described in subsequent sections.

In addition, CauseWorks supports synchronous and asynchronous collaboration through shared whiteboards, shared models and shared analytics. A team can flexibly work with the system on a large touch screen display and/or multiple workstations, enabling co-located group-based OD processes as well as distribution of tasks among separate teams or individuals. Features and observations pertaining to team collaboration using CauseWorks will be described in a subsequent report (Kapler, Gray, Vasquez, and Wright in preparation).

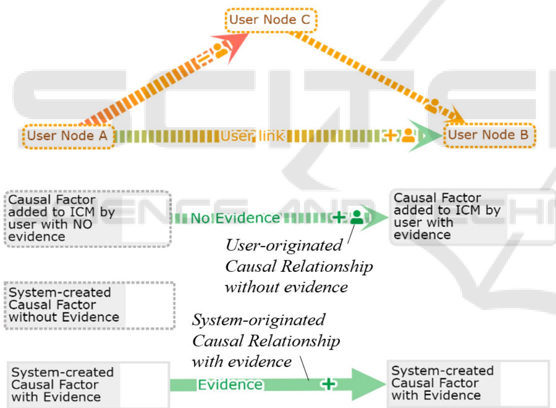

Figure 4: Visual encodings for user nodes, user links, causal factors and causal relationships.

Key visual encodings in CauseWorks delineate different stages and states in the construction of the causal model (see Figure 4). Throughout the model-building process, the team works with "user" nodes and links, and causal factors and relationships. User nodes and links have a distinct orange color with rounded corners. They have no causal function and are used to capture notes and ideas and for sketching systems diagrams. User links can however represent qualitative relationships between nodes, with width representing notional strength, a color ramp representing polarity (double-encoded with "+" or "-" icon) and color saturation representing confidence. When a user node or link is promoted to a computational causal factor or relationship, they visually change to using black text, and a grey background. A dashed line or outline indicates items that are NOT backed by evidence. User-originated causal factors retain their rounded corners while system-generated causal factors have sharp, right-angle corners.

These visual encodings were designed in conjunction with SMEs to help see at a glance the pedigree and state of the model as it evolves, indicating where there is supporting data, or where additional work is required; for example, backfilling details after an intense team discussion.

6.2 Sketch Hypotheses of Complex System

The OD process begins with team discussion about actors, factors and influences involved in a problem environment. This is typically captured as a series of point-form notes. In CauseWorks, users enter this information directly onto whiteboards using *user nodes* and *user links*, in a process we describe here as "sketching". This gives CauseWorks analytics a means to access team thinking and problem context: a critical step in providing model-building assistance.

Supporting the rapid pace of team discourse in a digital system (vs. physical markers and whiteboards) requires simple, efficient affordances. For rapid note-taking, user nodes can be created in quick succession by hitting the Tab key. Links are created by dragging a handle between nodes in one stroke, with direction, strength, polarity and confidence set simultaneously in single gesture (see Figure 5.5: Interactive Link Editor).

As the team sketches a rough causal system, they enter important concepts in the text of user nodes. User training includes guidance for labelling nodes intended to become factors in the model. For example, a label should include a measurable concept, along with locations and actors (e.g. "Economic Growth in Canada"). This specificity aligns with typical web-search syntax and enables CauseWorks to more effectively assist in evolving the model. An in-context search affordance that leverages the label is provided that operates in two ways. First it finds related factors in the model database that can be either dragged onto user nodes to substitute them in-place (see Figure 5.4), or just added to the whiteboard if they are of interest to the user. This is one way that computational causal factors are surfaced for consideration into the user's model. In many cases users may not have previously known about or considered these factors, thus potentially

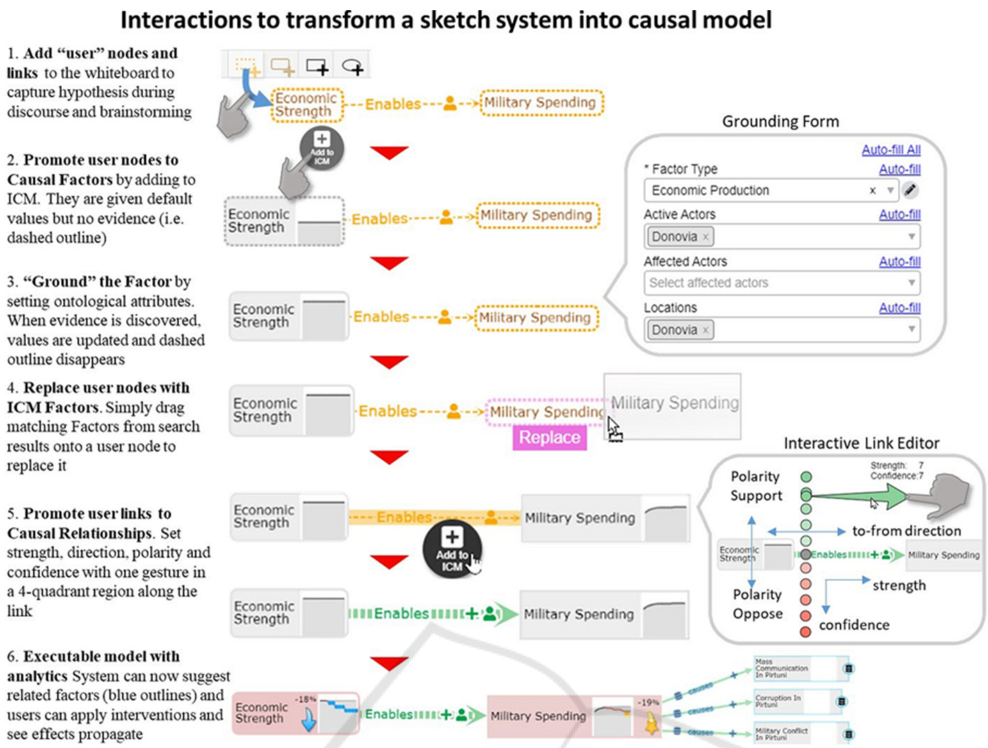

Figure 5.1-5.6: Interactions for transforming a sketch system into a computational causal model.

injecting new thinking and scope into the model. Second, the search tool performs a Lucene text search (Goetz, 2000) into the document corpus, which provides familiar web-search-like results for researching the topic represented by a node label. Results can be attached as evidence to a factor, thus allowing users to manually connect their hypotheses to data sources. This association provides additional opportunities for the machine to understand user thinking, problem context and apply analytics.

6.3 Transform Hypotheses Sketch into a Computational Causal Model

As the OD process moves into system framing, the user continues to transform their sketch into a computational causal model. User nodes and links can be instantly converted into computational *causal factor nodes* and *causal relationship links* by clicking "add to ICM" (Figure 5.2). As such, the team can construct causal factors and relationships purely based on their own knowledge or intuition without pulling from the model database. This gives CauseWorks flexibility to support unconstrained causal thinking about any domain. All causal factors have three key attributes:

1. Initial Value: value between 0-100 representing a factors current value today relative to historical norms, with 50 being "average". A 0-100 range was selected instead of discrete qualitative values because it allows for higher granularity to reveal small directional changes (e.g. a change of +5% may be insignificant, but the direction of change is perceivable)
2. Trend: The trend of the value at the current time, i.e. increasing, decreasing or staying level
3. Confidence: a scale of 0-10

User-derived causal factors are given default values; however, these can be modified at any time to reflect user's beliefs. It is important to note that in CauseWorks, the user's working model consists only of factors and relationships that are placed on a whiteboard. A dynamic model forward projection runs automatically whenever the user modifies the model. The projections can be run over a variable number of months or years and are displayed in sparklines on the nodes themselves. The grey and black timeline with grey fill shows the "baseline" projected value for a factor starting from "now" through a number of time-steps in the future, determined by executing a simulation of the model (see Figure 6).

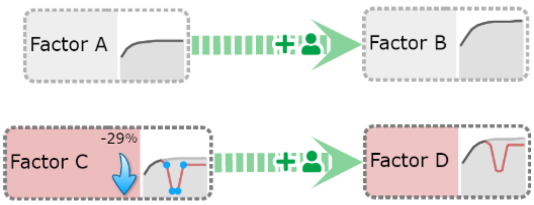

Figure 6: Baseline projections (top) and "What-if" projection overlay resulting from intervention (bottom).

The next stage in model evolution is to populate new factors and relationships with evidence from the corpus through a process referred to as "grounding". This involves defining a factor in ontological terms, and then running a process to connect it to the results of machine reading. Each factor includes properties to align it to a pre-determined domain ontology that includes FactorTypes, active and affected actors and locations (see Figure 5.3).

"Factor types" are high-level concepts detected by the reader subsystems in the document corpus (see System Description). The form can be completed manually by the user, or through the "auto-fill" feature, which uses machine reading of the user-entered node label to extract terms and complete the grounding form with suggested values.

Once grounded, a Factor's "Data" tab will list discovered events, assertions, and their associated extracts from structured and unstructured data sources, ranked by relevance. Trends in the structured data and event history are used to determine Factors' initial values and trend attributes. Explanation of value calculations are presented to users in the "Data" tab (development of explanations is ongoing and is outside of the scope of this paper).

A separate processing stage finds new causal relationships and factors related to the new factor based on information in the knowledge base and model database. These new relationships are stored in the model database, but are not added to the user's whiteboard. Rather, it is the role of the *Suggestion System* to provide in-context, discovery of new causal relationships for a selected factor, and quickly incorporate them into the model with minimal effort. Suggested relationships are displayed in context around a selected factor on the whiteboard (see Figure 4.6). They are displayed temporarily, with a blue outline to distinguish them. Suggestions are ranked based on multiple criteria including strength and context. To accept a suggestion into the model, the user clicks on it and it is added to the whiteboard. Thus suggestions surface new factors and relationships for consideration to improve and expand the users model.

6.4 Apply Causal Analytics

CauseWorks provides causal analytics services for improving the model, revealing behaviours, and achieving planning objectives. A key user experience goal is to incorporate analytics into the model building process, thus results are always presented within the whiteboard workspace to ensure continuity of thought and context.

6.4.1 "What if Projection"

The "What-if" projection is the most frequently used analytical tool in CauseWorks. "What-if" projections are generated when users create "interventions" on Factors to change the system. An intervention is defined as an applied change in the value of a factor at some point in the future. Applying interventions and assessing their effects is a key mechanism for understanding and validating model behaviour (see Figure 7). Interventions also capture the means to achieve the overall planning goals.

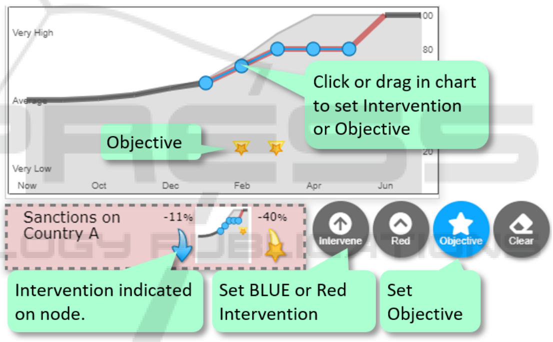

Figure 7: Intervention and Objectives Editor. Click on chart to add and remove interventions and objectives.

Fast, intuitive entry of interventions allows rapid testing and exploration of model behaviour. Users can "sketch" interventions over the baseline projection within a full-size version of the sparkline chart. When an intervention is created, CauseWorks automatically runs a model simulation and the resulting "What-If" projections are overlaid in the node sparklines for comparison, using green or red segments to emphasize values above or below the baseline. Interventions are highlighted in the whiteboard with blue arrows. Aggregate change in a factors average value over time compared to the baseline are pre-attentively indicated with a saturated scale of green or red applied to causal factor backgrounds, expressing increase or decrease, respectively (see Figure 1) making changes clearly visible within the generally monochrome graph.

6.4.2 Sensitivity

The sensitivity tool takes a factor as input and highlights other factors that it is causally sensitive to. *Degree of sensitivity* is displayed using a 5-point scale, displayed as magenta bars shown on the factor nodes. (see Figure 8.1). This helps the team answer questions such as "*which other factors are likely going to have a large impact on this factor?*" When applied to an objective factor, this tool assists the team in identifying influencing factors as targets for an intervention.

6.4.3 Most Impact

Most Impact indicates nodes with a high general impact on the system based on a series of simulations. These are also indicated with a 5-point scale, displayed as magenta bars. This helps the team answer questions such as, "*which are the central factors that all factors are most sensitive to?*"

6.4.4 "Why" Projections

"Why" Projection takes a single factor and time range as inputs and returns factors that cause changes in the value of that factor over the given time frame. Each resulting contributing factor is marked with 1 to 5 small arrows to indicate strength and direction of impact (see Figure 8.2). The "Why" function can operate on either the baseline or the what-if projection. This helps the team answer specific questions about projection results, such as, "*which factors lead to a jump at time 't' in this Factor?*"

6.4.5 Causal Loops

Causal Loops takes two factors as inputs (A and B) and outputs paths between them that include reinforcing or dampening loops occurring along each path. Loops are an important consideration in thinking about complex systems as they can significantly enhance effects of an intervention. Loops are highlighted in the whiteboard with a repeating animated cycle that is effective at emphasizing complex graph paths (Ito et al., 2016). Figure 8.3 includes a static image of a causal loop.

6.4.6 Causal Pathways

Causal Pathways takes two factors as inputs (A and B) and returns a new whiteboard that contains all significant paths in one direction from A to B. This helps the team answer questions such as "*how does a change in factor A propagate to factor B?*"

6.4.7 Approach Helper

CauseWorks provides tools for developing "approaches" to influence a system to meet planning objectives. These include tracking objectives, suggesting interventions to achieve them, finding unexpected impacts and opportunities, measuring solution performance, and managing multiple approach options (see Figure 1).

The key information that helps drive these analytics are users' *objectives.* Objectives are defined as a desired change in the value of a factor at some point in the future. Explicitly setting user objectives captures critical information about the team's goals, which in-turn helps inform machine assistance functions for suggesting factors, recommendations for interventions to achieve objectives, and generally improves system awareness of problem context and scope. Objectives are visually represented by an orange star icon on factors and in the sparklines. Objectives are entered using the same interface as interventions; users can set both time and value for an objective with a single click on the timeline (see Figure 7). An example of an objective is "increase sanctions on Country B starting in October".

In CauseWorks, an *approach* is a collection of objectives and the interventions applied to achieve them. Teams may develop multiple named approaches to a planning problem, such as "Aggressive Solution", or "Least Risk". Within the Approach Management Tab, users can name, organize, open, save and copy different approaches (see Figure 8.4). For a selected approach, the objective and intervention factors are displayed as a list of nodes. The Approach Helper is an analytic function that can automatically propose interventions to meet objectives. It uses objectives in the active Approach as the inputs and allows users to set constraints on the timing and size of interventions it proposes. When executed, this function will create interventions on factors in the user model and add them to the current Approach. The "Refine" tool is similar to the Helper, however it only adjusts existing interventions that the team has already set, optimizing them to meet objectives.

When considering which factors to apply interventions to, it is important to recognize that not all factors represent something that can be influenced directly as part of a solution (e.g., GDP of the USA). As such, the team can toggle a flag to exclude certain factors from intervention consideration.

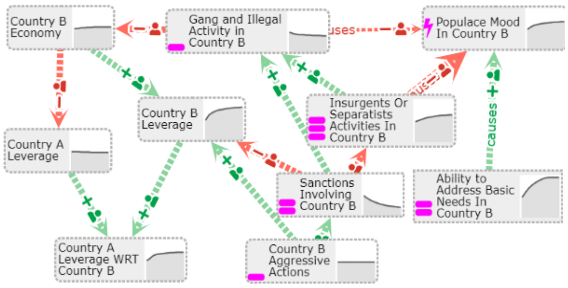

Figure 8.1: Sensitivity results for "Populace Mood In "Country B" indicated by magenta bars in the whiteboard view.

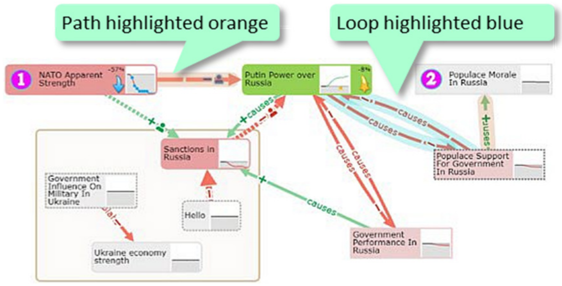

Figure 8.3: Example of Causal Loops result.

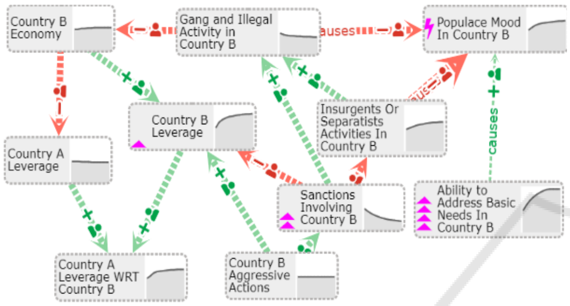

Figure 8.2: "Why" results show contributions by other factors on "Populace Mood In Country B" for a time period set by user (not shown). Up arrow-stack indicates relative amount of increasing influence.

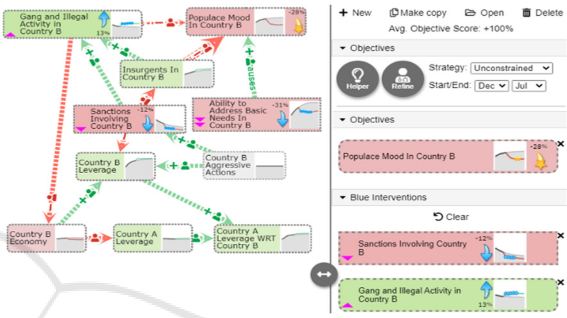

Figure 8.4: Approach Tab showing objectives in "Populace Mood in Country B" being met with interventions generated by the Approach Helper tool. Score of "100%" indicates objective target values are fully met in the What-if projection.

Each Approach can also include an optional automatically generated text summary of interventions, impacts on objectives and alternative intervention suggestions through an *Approach Narrative*. The narrative display includes embedded graphical representations of factors nodes, with linked highlighting and drag-and-drop to the whiteboard. This allows the user to connect the narrative to the whiteboard view and helps call out 2nd and 3rd order effects that users may not have noticed. As such, the narrative can surface potentially hidden or surprising information, a key benefit of combined human-machine systems (Wickens, Hollands, Banbury, & Parasuraman, 2015). The narrative engine is a separately developed plug-in component, and its design and assessment is the focus of a report by (Choudhry et al., 2020)

7 USER EVALUATION

OD SMEs were involved in CauseWorks development from the early stages of the system design through multiple hands-on evaluation exercises. Observations and interviews were conducted to determine design requirements and inform the initial system design. Multi-day exercise-based evaluations have been conducted every 6 months. Each involved a fictional scenario matched to a corpus of scenario-related documents processed into a knowledge base. Participants included experienced OD experts, government–provided problem domain experts, and OD students from the US Military. Teams were formed to address problems and present solutions over several days using CauseWorks. The goal of these exercises was to determine whether the system meets users' needs, and identify issues and functional gaps, to inform subsequent CauseWorks iterations.

Two key concerns were closely followed by the authors: 1) that users are able to quickly learn the basic interactions necessary to sketch systems, compose models, and develop approaches, and 2) users can produce planning products within normal planning timelines. HMI training took a half-day, including lower-level system background. Once users had hands-on CauseWorks, they were quickly able to use the system, possibly due to intentional similarities

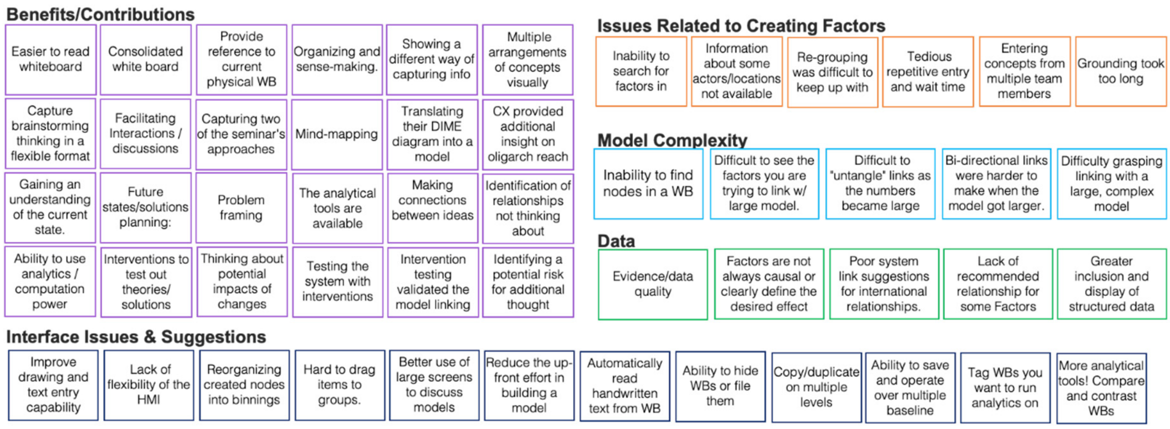

Figure 9: Affinity diagram of user feedback collected from planning exercises.

with tools such as MS PowerPoint. Within a 3-day standard planning exercise, teams using CauseWorks were able to generate models and approaches for their target problems, and then brief superiors using the live system. Further, they were able to adjust solutions on-the-fly to respond to Commanders questions and feedback. Significantly, this demonstrates that it is possible to incorporate causal modelling into the planning process with minimal negative effect on schedule or workload.

Post-exercise surveys were conducted by government staff and CX program performers to collect feedback about CauseWorks. We compiled survey answers from one exercise and constructed an affinity diagram (Harboe & Huang, 2015) comprising five themes in the data (see Figure 9). Many participants noted benefits and contributions of the system, suggesting that the CauseWorks provides advantages over traditional operational design methods. As with most prototypes, opportunities to improve the system were noted. Some users observed that creating and grounding factors was a cumbersome multi-step process. Indeed, the grounding process is not part of the traditional OD process, however it is necessary to connect factors to extracted data from the corpus. Additionally, SMEs observed that as models increased in scale, there were challenges with untangling links and grasping complex causal flows. SMEs also questioned the accuracy of event associations to Factors, though this can be attributed to limitations in automatic event and assertion extraction technologies. Finally, SMEs provided recommendations and suggestions to improve the interface, several of which have since been addressed.

8 LESSONS LEARNED & LIMITATIONS

In this section we discuss limitations and potential future research suggested by this work. This includes the efficacy of user-driven layouts, fixed vs. zoomable whiteboards, the distribution of models across multiple whiteboards, the construction of causal models by novice modellers, and the use of color to convey causal attributes and effects.

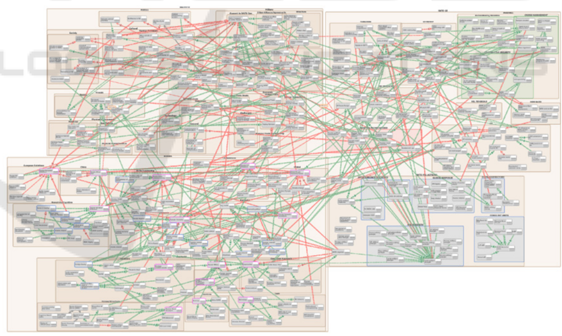

Figure 10: Example system sketch from problem-solving exercise with CauseWorks. This model consists of approximately 1000 nodes and edges created during a multi-day exercise. Image intentionally blurred to obfuscate confidential details.

An early design decision was to emphasize user-driven layouts, rather than applying automated graph layouts to the evolving causal model. This allows teams to arrange information according to their own criteria and logic and provides an additional contextual dimension from which machine analytics could extract meaning (e.g. users often cluster related items). User layouts also support long-term shared team cognition, as users become familiar with groupings and arrangements and remember where to find things. We observed that this contributed to maintaining comprehension of models containing nearly 1000 nodes and edges, well beyond our scalability expectations (see Figure 10).

At this scale, maximizing use of screen space and label readability has a significant impact on a teams' ability to view a model together, even on an 80-inch 4k display. Independent virtual collaboration spaces are a technical alternative to wall-size displays, however the impact of virtual meetings on OD discourse is unknown. A large touchscreen display was available at one exercise, and we observed stronger team engagement, discussion and model development compared to a single-laptop-per-team configuration. Users did ask for layout tools to automatically arrange factors based on topical groupings, such as geography or affiliation.

Early implementations of CauseWorks provided a fixed-scale whiteboard, as most virtual whiteboards take this approach (e.g. Google Jamboard) and there are documented challenges associated with zoomable workspaces. Zooming is a weak method for subgraph comparison and in some instances, users prefer an overview context compared to zoom interaction due to simpler navigation (Büring, Gerken, & Reitere, 2006). However, as models increased in size and complexity, SMEs demanded zoom-able whiteboards to fit growing models into a single view.

A related challenge is how to effectively work with large models that are distributed across multiple whiteboards. Initial thinking was to place major subsystems on separate whiteboards, however, observations from exercises revealed difficulties in recalling relationships between nodes on different whiteboards. Indeed, as the distance between sources of information increases, whether across multiple displays or whiteboards, there is a cost to the user in maintaining information in working memory (Wickens et al., 2015) and users tend to perform fewer information seeking behaviours (Fu and Gray, 2006). Eventually users elected to place large models on a single whiteboard to avoid this effect. Further research is required to effectively allow comprehension of connections in systems that span multiple pages.

Functions are needed to support model building by domain experts and planning experts who are not model building experts. Determining the appropriate scope and level of detail to model impacts the speed of model development, validity, speed of experimentation and confidence in model results. Building more complex models, because it is possible to do so, does not necessarily equate to effective modelling (Robinson, 2004). It is important to avoid modelling every aspect of a system. Simpler models can be developed faster, are more flexible, require less data and are easier to interpret, validate and maintain because the structure of the model is clearer (Chwif, 2000). The ease with which CauseWorks users can construct models impacts the size of the systems they create. This can enable inclusion and consideration of many more factors, causes and solutions than might otherwise occur using traditional methods, however one should acknowledge the possibility of overwhelming human comprehension. A cycle of model simplification or filtering could help reduce a large model to make it more comprehensible. In general, however, we observed that users find value in building larger systems to reflect real-world complexities, and also that the system can help with sense-making of larger models through use of analytics, and by enabling distributed work through collaboration.

CauseWorks use of green and red for indicating both link polarity *and* factor value change, merits explanation. Use of color is important when a user must attend to changing patterns on an interface (Brewer,1996). Our hypothesis is that the shared color schemes reinforce each other (increase/support vs. decrease/oppose; Wickens et al., 2015), and thus makes recalling combinations of link-factor effects easier than with two separate, two-color schemes (4 colors total; Brewer, 1996). A 4-color scheme also limits unique color use for other information. User feedback on color usage includes mention that red is often used to represent "enemy" in military convention, which may cause confusion in interpretation, however despite providing alternative color options in CauseWorks, users continued to use the green-red color scheme without issue. We suggest that further research into optimal color schemes for military use of causal model representations may be useful.

9 CONCLUSIONS

Causal Model building for complex problems has typically been completed manually by domain experts and is a time-consuming, cumbersome process. OD is a process of rapid, structured discourse for teams to envision systems and relationships about complex, "wicked" problems, however, the resulting models are simple diagrams produced on whiteboards or slides, and as such, do not support computational analytics, thus limiting usefulness. In this paper we introduced the HMI and workflow for CauseWorks, a tool for expert planners (but novice model builders) to create computational causal models of complex problems. We presented how users can sketch hypotheses about a system on a digital whiteboard and connect it to automatically extracted information that suggests system behaviour, thereby transforming the sketch into a computational causal model. CauseWorks also helps expand OD team thinking and model development by suggesting new factors to add to the model, and by providing analytics to support sense-making and solution development. In applied planning exercises, military operational planners with no prior modelling experience were able to use CauseWorks to construct and use computational causal models to develop approaches for realistic complex planning scenarios, within typical planning time constraints. Planners thought CauseWorks supported the OD process and helped them consider new ideas.

Future work should investigate the following: ways to present connections between models spanning multiple whiteboards; assessment of model characteristics built by novice modellers; deeper investigation into causal symbology and color-use for military applications. Formal experiments should be performed to assess impact of using CauseWorks modelling tools in operational design vs traditional methods.

ACKNOWLEDGEMENTS

This work was supported by the Defense Advanced Research Projects Agency (DARPA) under Contract Number FA8650-17-C-7720. The views, opinions, and findings contained in this report are those of the authors and should not be construed as an official Department of Defense position, policy, or decision. This work has been approved for Public Release, Distribution Unlimited. The authors wish to thank all Causal Exploration collaborators for support and encouragement.

REFERENCES

ATP 5-0.1, Army Design Methodology, HQ Dept of the Army, July 2015.

Bae, J., Helldin, T., & Riveiro, M. (2017). Understanding indirect causal relationships in node-link graphs. *Computer Graphics Forum, 36*(3), 411-421.

Box, G. E. P. (1979), Robustness in the strategy of scientific model building, in Launer, R. L.; Wilkinson, G. N. (eds.), Robustness in Statistics, Academic Press, pp. 201–236.

Brewer, C. A. (1996). Guidelines for selecting colors for diverging schemes on maps. *Cartographic Journal, 33*(2), 79-86

Büring, T., Gerken, J., & Reiterer, H. (2006). User interaction with scatterplots on small screens - A comparative evaluation of geometric-semantic zoom and fisheye distortion. *IEEE Transactions on Visualization and Computer Graphics, 12*(5), 829-836.

Chen, T. J., Subramanian, S. G., & Krishnamurthy, V. R. (2020). QCue: Queries and Cues for Computer-Facilitated Mind-Mapping. Paper presented at the *Proceedings – Graphics Interface*

Choudhry, A., Sharma, M., et al. (2020). Once upon a time in visualization: Understanding of textual narratives for causality, *IEEE Transactions on Visualization and Computer Graphics*

Chwif L, Barretto M.R.P., & Paul R.J. (2000). On simulation model complexity. In: Joines JA, Barton RR, Kang K and Fishwick PA (eds). *Proceedings of the 2000 Winter Simulation Conference*. IEEE: Piscataway: NJ, pp 449–455.

Dang, T. N., Murray, P., Aurisano, J., & Forbes, A. G. (2015). ReactionFlow: An interactive visualization tool for causality analysis in biological pathways. *BMC Proceedings, 9*

DARPA, (2018). https://www.darpa.mil/program/causal-exploration.

Das, T. K., & Teng, B. S. (1999). Cognitive biases and strategic decision processes: An integrative perspective. *Journal of Management Studies, 36*(6), 757-778.

Fu, W. T., & Gray, W. D. (2006). Suboptimal tradeoffs in information seeking. *Cognitive Psychology, 52*(3), 195-242.

Gil, Y., Pierce, S. A., Babaie, H., Banerjee, A., Borne, K., Bust, G., . . . Shekhar, S. (2019). Intelligent systems for geosciences: An essential research agenda. *Communications of the ACM, 62*(1), 76-84.

Goetz, B. (2000). The Lucene search engine: Powerful, flexible, and free. JavaWorld. http://www. javaworld. com/javaworld/jw-09-2000/jw-0915-lucene. html.

Gray, W. D., & Fu, W. T. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science, 28*(3), 359-382.

Harboe, G., & Huang, E. M. (2015). Real-world affinity diagramming practices: Bridging the paper-digital gap. *In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 95-104).

Ito, T., & Misue, K. (2016). A visualization technique using loop animations. *In International Conference on Human Interface and the Management of Information* (pp. 136-147).

Joint Publication JP 5-0, Joint Planning, Joint Chiefs of Staff, June, 2017.

Jonassen, D. H., & Ionas, I. G. (2008). Designing effective supports for causal reasoning. *Educational Technology Research and Development, 56*(3), 287-308.

Keim, D. A., Kohlhammer, J., Ellis, G., & Mansmann, F. (2010). *Mastering the information age; solving problems with visual analytics*. Eurographics.

McPherson, K., Marsh, T., & Brown, M. (2007). Foresight report on obesity. *The Lancet, 370*(9601), 1755.

Pearl, J. (2009). Causal inference in statistics: An overview. *Statistics Surveys, 3*, 96-146.

Proulx, P., Husain, F., Gyori, B. M., & Pyarelal, A. (2019). MOIRE: MOdeling and Inference from REading. Demonstration presented in *Modeling the World's Systems Conference.*

Rittel, H. W. J., Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences, 4*(2), 155–169.

Robinson, S. (2004). *Simulation: The Practice of Model Development and Use*, John Wiley and Sons Ltd: England.

Schmitt, J., A Systemic Concept for Operational Design, (2006) US MC Warfighting Laboratory. http://www. mcwl. usmc. mil/file_download.

Sedig, K., Parsons, P., & Babanski, A. (2012). Towards a characterization of interactivity in visual analytics. *Journal of Multimedia Processing Technologies, 3*(1), 12-28.

Sedlmair, M., Meyer, M., & Munzner, T. (2012). Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics, 18*(12), 2431-2440.

Shih, P. C., Nguyen, D. H., Hirano, S. H., Redmiles, D. F., & Hayes, G. R. (2009). GroupMind: Supporting idea generation through a collaborative mind-mapping tool. Paper presented at the *GROUP'09 - Proceedings of the 2009 ACM SIGCHI International Conference on Supporting Group Work*, (pp.139-148).

Von Landesberger, T., Kuijper, A., Schreck, T., Kohlhammer, J., van Wijk, J. J., Fekete, J. D., & Fellner, D. W. (2011). Visual analysis of large graphs: state-of-the-art and future research challenges. *Computer Graphics Forum, 30*(6), 1719-1749

Wang, X. (2018). A Comparison of Causal Inference Methods and Their Application in Big Data Analytics (Doctoral Dissertation), Louisiana State University, Louisiana, USA.

Wang, J., & Mueller, K. (2018). Visual causality analysis made practical. Paper presented at the *2017 IEEE Conference on Visual Analytics Science and Technology, VAST 2017 Proceedings*, 151-161.

Wickens, C. D., Hollands, J. G., Banbury, S., & Parasuraman, R. (2015). *Engineering psychology and human performance*. Psychology Press.

Wright, W., & Kapler, T. (2018). Challenges in Visualizing Complex Causality Characteristics. *IEEE Pacific Visualization.* 2018. Proceedings

Wright, W., Sheffield, D., & Santosa, S. (2017). Argument mapper: Countering cognitive biases in analysis with critical (visual) thinking. *In Information Visualisation (IV), 2017 21st International Conferences,* IEEE, (pp. 250-255).

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

Zhang, Z., McDonnell, K. T., Zadok, E., & Mueller, K. (2015). Visual correlation analysis of numerical and categorical data on the correlation map. *IEEE Transactions on Visualization and Computer Graphics, 21*(2), 289-303.

Zhang, X., & Vesselinov, V. V. (2017). Integrated modeling approach for optimal management of water, energy and food security nexus. *Advances in Water Resources, 101*, 1-10.