© 2026 IEEE. This is the author's version of the article that has been published in IEEE Transactions on Visualization and Computer Graphics. The final version of this record is available at: xx.xxxx/TVCG.201x.xxxxxxx/

Conveying Causality Through Visualization: A Survey of Visual Causal Inference

Arran Zeyu Wang, David Borland, and David Gotz

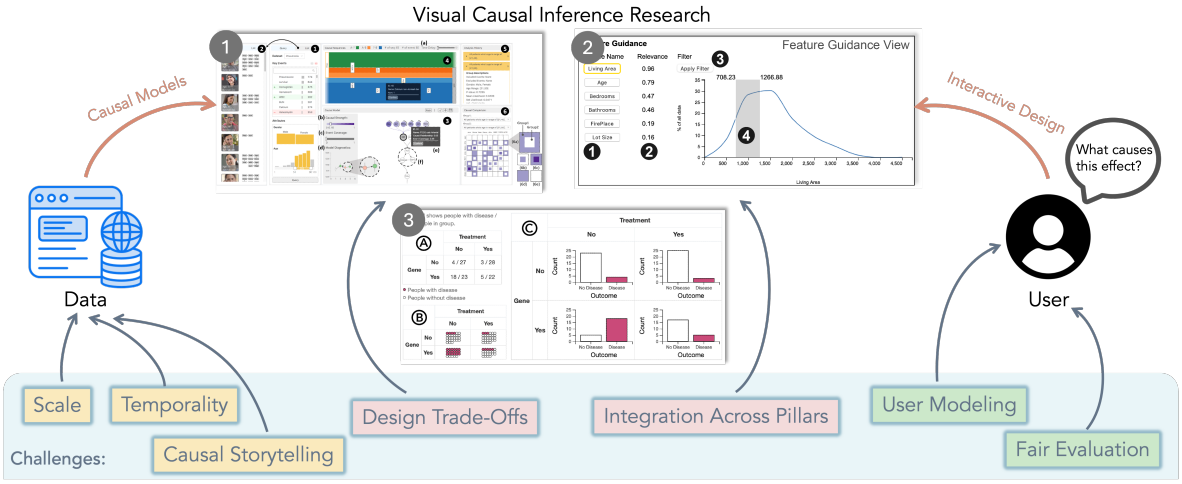

Fig. 1: An overview of visual causal inference (VCI) research, depicting a real-world pipeline of VCI research that supports and assesses users' causal inferences with visualization and visual analytics from real-world data. The middle section shows examples of VCI research aligned with three different pillars: (1) structural, e.g., *SeqCausal* [35]; (2) augmentative, e.g, *CFGuidance* [69]; and (3) cognitive, e.g., *Causal Support* [38]. Challenges to address in the future are shown at the bottom, such as scale/temporality as data grows in complexity, balancing design trade-offs for VCI tools, and user modeling and evaluation to understand VCI behaviors.

Abstract— While data visualization has historically focused on revealing frequencies, correlations, and other descriptive patterns, relying on descriptive statistics alone during exploratory data analysis may lead to spurious conclusions. Consequently, Visual Causal Inference (VCI) has emerged as a critical research topic, combining data visualization, an understanding of human perception and cognition, human-computer interaction, and the rigor of statistical causal inference and modeling. Reflecting the broad applicability of such work, the literature surrounding VCI has addressed a variety of research questions, fragmented across different contribution types, venues, and application domains. This paper provides the first systematic review of the VCI landscape and proposes a taxonomy for structuring the domain across three core categories: the *Structural Pillar*, examining interactive visual analytics techniques with a graphical structure presenting causal structures; the *Augmentative Pillar*, exploring techniques to estimate causal effects using auxiliary representations; and the *Cognitive Pillar*, studying how humans perceive and comprehend causality from visualizations. By synthesizing these areas, we map the current design space and identify critical gaps and emerging opportunities in VCI, calling for future actions in unifying the pillars, navigating design trade-offs, and bridging the evaluation gap for better visual communication of causality.

Index Terms—Visual Causal Inference, Data Visualization, Exploratory Data Analysis, Causality, Decision Making, Literature Review

1 Introduction

The field of visualization has historically defined its boundaries through the lens of descriptive statistics. The foundation of the field states that the role of information visualization is to map data to visual variables, such as position, length, color, and orientation, to reveal patterns, trends, and outliers [51, 62, 75]. Implicit in this formulation is an emergent dichotomy: the primary role of visualization is to present statistical associations (e.g., correlation), whereas the role of inference (i.e., providing causal descriptions) remains the exclusive province of the human user, via domain knowledge or rigorous statistical modeling.

For decades, the fundamental maxim **correlation does not imply causation** has been taught in every introductory statistics course, yet standard exploratory visual analytics tools are largely designed to display exactly that: *correlations*. When users rely on scatterplots, bar charts, or dashboards to make policy, medical, or business decisions, they are often implicitly treating visual associations as causal relationships [38, 66, 71, 80]. Without proper care, this translates to a risk of spurious causality judgments, where factors such as unobserved variables lead analysts to potentially incorrect and harmful conclusions (such as those hidden by Simpson's Paradox [63]).

To transition from asking "what happened?" to "why did it happen?" and "what if we act?", researchers have increasingly turned to integrating formal causal inference frameworks, such as Pearl's Structural Causal Models (SCMs) [55] and Rubin's Potential Outcomes [57], with visual analytics approaches. However, these mathematical frameworks still require human users' domain knowledge to explicitly represent causal assumptions. Algorithms cannot automatically distinguish between an observational correlation and a causal mechanism without a human expert to validate assumptions, untangle complex variable relationships, and interpret counterfactual scenarios.

Directed graphs were first used to depict causal diagrams by geneticist Sewall Wright in 1921 [78], and further work with graphical causal models was performed by Pearl in the 1970s describing SCM [55].

- *Arran Zeyu Wang and David Gotz are with the University of North Carolina at Chapel Hill (UNC). E-mail: zeyuwang@cs.unc.edu, gotz@unc.edu*
- *David Borland is with RENCI at UNC. E-mail: borland@renci.org*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

In the visualization research community, initial work on using visualizations to communicate causality began in the late 1990s [17, 76], employing animation and moving objects to present causality. However, there was little follow-up research on the topic in the following decade within the VIS community. The introduction of *Visual Causal Analyst* [72] in VIS'2015 explicitly brought causal inference to the visual analytics context, and led to increased research on using visualizations to assist causal inference in the years since. More recently, after the reported phenomenon of *Illusion of Causality* [80], there appears to be more empirical research exploring the underlying cognitive mechanisms behind users' causal inference from visualized data, uncovering ways causalities are being conveyed.

Such research has become increasingly recognized by the community, with an increasing number of publications, as well as awards [38, 43, 69]. This increase may suggest we are witnessing a paradigm shift: from describing what happened, to explaining the causality behind why it happened and predicting what would happen due to interventions. We define the use of visualization to convey and explore causal relationships as **Visual Causal Inference** (VCI). To better understand this trend, we report a literature review of VCI research spanning the last two and a half decades across major visualization venues. In this work, we define VCI as an approach that (1) *employs data visualization and visual analytics* to (2) *communicate causal information to help interpret data* and (3) *use this information to help analyze data to support human reasoning and decision-making*.

Despite the rapid growth of VCI papers in visualization research, the field currently lacks a unifying taxonomy. Research is heavily distributed across distinct smaller communities; some focus on novel visual analytics systems for uncovering graphical causal structures [35, 42, 73], some focus on the underlying visual causal discovery to allow users to estimate causal effects [63, 69, 83], and some emphasize user communication and the cognitive comprehension of causality behind data [38, 71, 82].

This paper addresses this fragmentation by systematically reviewing the literature across major visualization venues. We abstract the distinct contributions of VCI research into a cohesive taxonomy that focuses on three aspects of why a tool is built or a study is conducted (*The Causal Problem*), how it visually encodes causal data (*The Visual Solution*), and who it is designed to support (*The Human Factor*). This framing is intended to ground the survey more explicitly in causal inference theory/framework while still preserving the visualization/human-centric perspective of the field. By synthesizing empirical evidence, system designs, and theoretical models from 57 research artifacts, we identified major achievements and critical research gaps for VCI.

Based on the taxonomy categorization, we conducted a structured analysis of the VCI landscape, and organized VCI research outputs into the following three major components:

The Structural Pillar: Research that builds systems to externalize SCM-based techniques under task scenarios. These systems are designed to help users observe and refine graphical structures, focusing on the identification of causal relations and effective communication of visual causal representations.

The Augmentative Pillar: Research that leverages auxiliary information and alternative representations without relying solely on causal graph drawings to assist and improve human causal inference. This includes new types of visualization designs, ways of communicating inference scenarios, and guided visual causal discovery for users.

The Cognitive Pillar: Research that examines the human element of causality, focusing on human perception of causality, exploring questions such as "how do users' prior beliefs impact their causal interpretation of data?" and "how do visual designs mitigate or exacerbate cognitive biases and illusions of causality?"

In addition, we synthesize implications and limitations across the literature to inform a research agenda, such as: making deeper connections for empirically grounded tools; the potential of system design diversity beyond graph drawing; and missing widely-accepted evaluation methods and benchmarking of human factors.

2 Background

Research exploring VCI as a distinct sub-field of visualization builds upon early mathematical work formalizing the definitions of causality, and the intersections of these concepts with visualization and visual analytics. In particular, the field inherits a number of statistical foundations, such as potential outcomes and structural causal models, that help shape how visual causal inference systems define, encode, and validate causal assumptions.

2.1 Statistical Foundations

For much of the early 20th century, the field of statistics was dominated by the study of correlation [56]. The determination of causality and causal inference was largely restricted to the domain of Experimental Design, specifically through Randomized Controlled Trials (RCTs) introduced by Ronald Fisher [22]. In an RCT, random assignment isolates the treatment variable, neutralizing the effect of external confounders.

However, in many critical domains, such as welfare, economics, urban planning, and in some cases of healthcare and medicine, conducting RCTs is unethical, physically impossible, or prohibitively expensive [14, 28, 34]. These limitations necessitated the development of frameworks to infer causality from purely observational data. Two dominant mathematical frameworks emerged to enable this approach.

First is the potential outcomes framework, also known as the Rubin Causal Model (RCM) [57], developed in the 1970s. This framework conceptualizes causality as the difference between two potential states for an individual unit: one where they receive a treatment, and a counterfactual state where they do not (e.g., $Y_i(1) - Y_i(0)$). Because we can never observe both states simultaneously for the same unit (the fundamental problem of causal inference), this approach relies heavily on statistical matching and propensity score weighting to simulate a controlled environment.

Second is structural causal models (SCMs) and do-calculus [55]. In the 1990s, Judea Pearl largely changed the field by formalizing causality through graphical methods using Directed Acyclic Graphs (DAGs). Pearl introduced the $do(X)$ operator to mathematically distinguish an intervention (forcing variable X to a specific value) from an observation (passively seeing X). SCMs provided a visual and mathematical language to explicitly encode assumptions, identify backdoor paths, and calculate effects.

Currently, most techniques to explore causality from data can be regarded as stemming from these two frameworks. However, they also suggest that inferring causality requires untestable assumptions about the real world—assumptions that algorithms alone cannot verify without human domain knowledge [52]. This dependency on human judgment explains why visual analytics has become a natural partner to causal inference: visualization can expose assumptions, support comparison, and help analysts reason about evidence, but it cannot replace the need for domain expertise or methodological rigor.

2.2 The Tension of Correlation vs. Causality in EDA

Parallel to the development of these causal frameworks, another area of the statistics community heavily adopted John Tukey's approach of Exploratory Data Analysis (EDA) [64], which later contributed to the emergence of visualization research. EDA systems, with visual analytics as their core, are designed to foster rapid hypothesis generation by enabling users to slice, filter, and visually project observational data.

However, visualization can often lead to the misattribution of correlative relationships as causal. When a user explores a scatterplot or a dashboard, they are highly susceptible to identifying visual trends that are entirely driven by *unobserved* variables [32], either confounded, mediated, or collided. This creates a dangerous paradigm: highly interactive visual analytics tools may inadvertently facilitate the discovery of spurious associations, leading users to make incorrect downstream decisions. Mitigating these effects to avoid spurious analysis results during EDA has become one of the most pressing challenges, and a common research topic, in visualization [1, 46, 47]. Accordingly, VCI should be understood not as a replacement for EDA, but as an attempt to extend exploratory workflows with explicit support for graphical inference. In this way, users can move from surface-level association seeking toward more defensible inference [32].

© 2026 IEEE. This is the author's version of the article that has been published in IEEE Transactions on Visualization and Computer Graphics. The final version of this record is available at: xx.xxxx/TVCG.201x.xxxxxxx/

2.3 The Evolution of Visual Causal Inference

Foundational statistical causal inference frameworks have long applied node-link graphs to represent causal relations, such as the use of DAGs as graphical causal models [55]. Visualization researchers, recognizing the the potential of using these statistical causal theories to improve visual reasoning, began exploring the adoption of these statistical approaches in combination with interactive visualization to support visual causal inference (VCI). The development of VCI, an approach to resolving the tension between the often correlational nature of visualized data and the need for robust decision-making, can be categorized into several distinct phases within the literature (see Figure 2 for a timeline).

Early work, from 1999 to around 2010, focused almost exclusively on animation-based techniques [16–18, 37, 76]. Much of this work was influenced by an early study on the *theory of ampliation* by psychologist Albert Michotte in 1946 [48]. Michotte employed visual sequences to study the appearance of causality and introduced a famous example in which people perceived causality when a moving ball strikes another and makes it move, indicating that animation may be a useful visualization technique for conveying causality. This era established the early foundation that causality could be visually encoded, but treated causal inference as primarily a visual metaphor design problem.

Later on, after 2015, VCI research shifted toward building systems targeting domain experts for causal discovery, utilizing causal algorithms and visualizing the causal graphs. Visual analytics tools such as the *Visual Causality Analyst* [72, 73] enabled experts to constructively conduct causal inference using their domain knowledge. A number of complex causality models were integrated with visual analytics, such as Granger causality applied to event-sequence data [35]. They utilized case studies and task-based expert evaluation instead of relying purely on technical evaluations. Relevantly, causal inference mechanisms and ideas have also inspired a closely adjacent thread of work in the field: model explanation and probing. For example, counterfactual methods have been widely applied in visual analytics for ML models [9, 24, 77].

Since 2020, much VCI research has focused on more empirical factors, examining how visualizations convey causality, assist users to estimate effects, and communicate outcomes to users [38, 66, 71, 80]. Assessing and modeling the causal reasoning process with laypersons has become a focus. These studies focus on empirical validation, exploring how individuals interpret causality from visualizations. They aim to answer questions beyond just how to build effective tools, but also how and why our tools work for humans to reason causally.

However, despite the rapid maturation of VCI, the literature remains deeply fragmented. This fragmentation spans not only applications and visual forms, but also differing assumptions about what counts as causal evidence, what role visualization should play, and how causal claims should be evaluated. By systematically mapping the literature across different perspectives, this paper aims to standardize the vocabulary of VCI and provide a framework to understand the current states and inspire future VCI research.

3 Methodology

We followed the PRISMA guidelines for systematic reviews [49]. We searched full and short papers from the major venues that collectively define the state-of-the-art in visualization research: IEEE TVCG / VIS, CGF / EuroVis, the Information Visualization journal, PacificVis, and ACM CHI, from the years 1990 to 2026. The full PRISMA diagram can be found in OSF Appendix.

We initially searched through the digital libraries of these venues and then applied eligibility criteria to include papers only if they aim to use visualizations to assist or assess human users in addressing causality. This includes inferring causality, identifying causal structure, understanding factor-outcome relations, or reasoning about causality and counterfactuals. A total of 57 papers were included in our final corpus. Among them, 31 papers are from IEEE VIS / TVCG, 7 from the Information Visualization journal, 4 from EuroVis / CGF, 3 from IEEE CG&A, 2 from ACM CHI, 2 from PacificVis, 2 from CSCW, and the remainder from other venues such as the Journal of Visualization. We then categorized the papers into four types: Technique, System, Empirical Study, and Theory, following the paper type definitions used in prior visualization research [50]. See Figure 2 for a summary of the paper distribution. The collection of papers and taxonomy is available in the OSF.

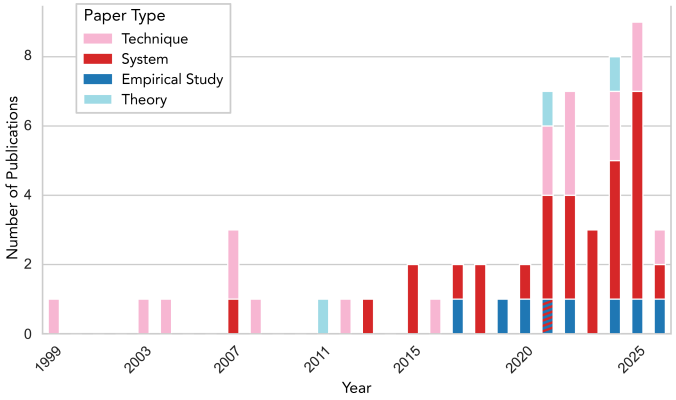

Fig. 2: The distribution of reviewed papers from 1999 to 2026, including publication year, number of papers, and their paper types.

4 A Taxonomy for Visual Causal Inference Research

Given the diverse landscape of VCI research, we coded and categorized the papers in more detail beyond the initial set of four types. We first reviewed all papers to identify a preliminary list of terms that could be used to describe, group, and differentiate them. Then all authors met weekly to discuss and iteratively arrive at a final set of 44 dimensions. These were then used to code all included papers. We categorize the literature along three primary aspects: *Why* (tasks, framework, and data), *How* (visualization techniques and designs), and *Who* (human factors and evaluation), from 44 detailed dimensions identified via the process outlined above. In addition, each dimension is grouped into a thematic area within each aspect. This framework enables us to dissect not only the statistical capabilities of a given system, but also its visual design space, and its human-centered paradigms.

The taxonomy was initially inspired by previous taxonomies in visualization that have summarized paper dimensions based on their unique aspects [30, 41, 60]. It was developed through weekly meetings and discussions based on readings of the surveyed papers by the authors.

4.1 The Causal Problem: Why

The *Why* aspect describes the underlying statistical challenge(s) a system addresses and the nature of the data the system is required to process to achieve those goals. This aspect, therefore, reflects the motivation and statistical foundations behind a given research contribution. Different underlying mathematical objectives (e.g., suggesting that a relationship exists across datasets vs. calculating a specific causal effect magnitude) will result in different designs of visualization tools and experiments. We therefore distinguish between tasks, causal frameworks, and data characteristics to make explicit how a system's analytic goal constrains both its interface and its evaluation.

4.1.1 Causal Task

People often infer causality from data visualizations even when they are not designed to convey it [38, 71, 80]. This, in part, motivates the fundamental objective of VCI: enabling people to more reliably conduct causal reasoning tasks. Such tasks take many forms in the literature.

Discovery (33 papers): This task focuses on learning the structural causal relationships between variables, including identifying whether evidence for a specific directed causal link (e.g., $A \rightarrow B$) exists in a dataset, and supporting verification of those potential effects. This is a critical task because algorithms alone cannot generally differentiate complex causal structures without human domain knowledge. For example, Causal Structure Investigator [73] is designed specifically to let humans generate or constrain structural hypotheses that algorithms cannot resolve independently.

Table 1: The taxonomy of VCI research presented in this paper.

			Pillars			**The Causal Problem: Why**																**The Visual Solution: How**									**The Human Factor: Who**													
						Causal Task				Causal Framework					Data Type						Data Complexity				Causal Effect Encoding					Causal Interaction				User Role				Target Expertise			Evaluation			
No.	Paper	Year	Structural	Augmentative	Cognitive	Discovery	Communication	Estimation	Refutation	SCM (Pearl)	Counterfactual	Granger	Bayesian	Potential Outcomes	Tab./Multivar.	Synthetic	Text/Graph	Time-Series	Event-Seq.	Spatial	One-dim	Multi-dim	Large	Heterogeneous	Node-Link	Animation	Inferential Stats.	Partitioned Data	Stats. Charts/Text	Parametric	Constructive	Conditioning	Steerable	Explorer	Generator	Verifier	Interpreter	Domain Expert	Data Analysts	Layperson	Case Study	Quan. User Study	Qual. User Study	Tech. Evaluation
1	[76]	1999					✓									✓					✓				✓													✓					✓
2	[17]	2003					✓									✓					✓				✓													✓					✓
3	[18]	2004					✓									✓					✓				✓													✓					✓
4	[16]	2007					✓									✓					✓				✓													✓					✓
5	[19]	2007	✓			✓											✓						✓	✓	✓	✓												✓			✓	✓		
6	[37]	2007					✓									✓					✓				✓													✓					✓
7	[4]	2008					✓									✓					✓				✓													✓					✓
8	[8]	2011				✓																															✓						
9	[61]	2012	✓					✓		✓					✓							✓					✓				✓						✓					✓		
10	[65]	2013	✓			✓				✓								✓					✓	✓	✓												✓				✓			
11	[12]	2015	✓			✓				✓					✓							✓			✓					✓							✓				✓			
12	[84]	2015	✓			✓				✓								✓		✓				✓							✓						✓					✓		
13	[72]	2016	✓			✓		✓		✓					✓							✓			✓					✓				✓					✓			✓		
14	[2]	2017	✓		✓		✓			✓						✓						✓			✓											✓				✓				
15	[73]	2017	✓	✓		✓		✓		✓					✓							✓			✓					✓		✓		✓					✓			✓		
16	[54]	2018	✓			✓				✓					✓									✓						✓				✓					✓					
17	[45]	2018	✓	✓		✓						✓					✓		✓					✓	✓		✓					✓				✓				✓				
18	[82]	2019			✓	✓				✓						✓						✓			✓						✓					✓					✓		
19	[80]	2020			✓		✓									✓					✓							✓							✓					✓	✓		
20	[40]	2020		✓				✓				✓						✓				✓					✓	✓					✓				✓				✓			
21	[83]	2021	✓	✓			✓			✓							✓					✓			✓				✓		✓			✓					✓			✓		
22	[32]	2021				✓									✓							✓																✓					
23	[79]	2021	✓	✓		✓			✓	✓					✓							✓			✓				✓	✓			✓	✓	✓				✓				
24	[35]	2021	✓	✓		✓						✓		✓					✓			✓			✓			✓	✓	✓				✓	✓				✓					✓
25	[10]	2021		✓	✓		✓						✓			✓	✓					✓			✓				✓				✓			✓			✓			✓		
26	[26]	2021		✓		✓			✓					✓	✓							✓							✓	✓			✓					✓					
27	[33]	2021	✓	✓		✓				✓						✓						✓			✓				✓	✓			✓						✓			✓		
28	[15]	2022		✓		✓				✓								✓				✓			✓						✓				✓				✓				
29	[6]	2022	✓			✓												✓				✓													✓				✓		✓		
30	[3]	2022	✓				✓			✓					✓							✓			✓									✓					✓					
31	[38]	2022			✓	✓	✓						✓					✓			✓														✓	✓				✓		✓		
32	[13]	2022	✓	✓		✓						✓					✓			✓		✓			✓			✓	✓					✓				✓			✓			
33	[31]	2022	✓	✓				✓		✓	✓				✓							✓			✓				✓	✓		✓			✓				✓	✓	✓		
34	[39]	2022		✓		✓		✓			✓				✓	✓						✓			✓						✓		✓			✓				✓	✓		
35	[27]	2023	✓			✓		✓	✓	✓				✓	✓							✓			✓				✓	✓			✓	✓	✓	✓			✓		✓		
36	[74]	2023	✓	✓		✓						✓					✓					✓			✓			✓	✓			✓	✓					✓			✓		
37	[23]	2023	✓	✓				✓	✓	✓					✓							✓			✓			✓	✓			✓		✓	✓	✓			✓		✓		
38	[7]	2024		✓				✓			✓											✓						✓				✓				✓							
39	[81]	2024	✓			✓						✓							✓					✓	✓					✓	✓			✓	✓				✓	✓			✓
40	[59]	2024	✓				✓			✓							✓					✓			✓									✓			✓				✓			
41	[67]	2024			✓		✓				✓					✓						✓						✓	✓							✓			✓		✓			
42	[11]	2024		✓				✓			✓		✓		✓							✓				✓			✓			✓			✓				✓		✓		
43	[68]	2024		✓			✓	✓			✓											✓				✓			✓	✓	✓		✓	✓	✓				✓		✓	✓	✓
44	[63]	2024		✓			✓		✓					✓	✓		✓					✓				✓	✓	✓				✓			✓	✓			✓		✓	✓		
45	[42]	2024	✓	✓		✓				✓					✓		✓					✓		✓					✓	✓				✓	✓				✓		✓	✓		
46	[85]	2025	✓			✓						✓					✓					✓			✓					✓				✓				✓	✓		✓		
47	[71]	2025			✓	✓	✓									✓					✓															✓					✓			
48	[87]	2025	✓			✓			✓	✓							✓					✓			✓					✓				✓	✓				✓		✓	✓		
49	[36]	2025	✓	✓		✓						✓					✓			✓				✓			✓	✓					✓	✓					✓			✓		
50	[88]	2025	✓			✓						✓							✓			✓			✓					✓				✓					✓			✓	✓	✓
51	[21]	2025	✓				✓			✓					✓							✓			✓					✓	✓			✓		✓			✓	✓		✓		
52	[69]	2025		✓		✓		✓			✓					✓						✓				✓					✓								✓	✓				
53	[86]	2025		✓		✓								✓	✓							✓					✓	✓				✓								✓					
54	[53]	2025	✓				✓			✓							✓					✓			✓						✓				✓						✓			
55	[66]	2026			✓		✓									✓					✓					✓		✓							✓						✓			
56	[43]	2026	✓	✓		✓				✓	✓						✓		✓				✓	✓	✓					✓						✓				✓	✓	✓	✓	✓
57	[58]	2026	✓			✓	✓			✓					✓							✓			✓										✓				✓	✓		✓		

© 2026 IEEE. This is the author's version of the article that has been published in IEEE Transactions on Visualization and Computer Graphics. The final version of this record is available at: xx.xxxx/TVCG.201x.xxxxxxx/

Communication (21): Communication tasks focus on explaining the causal links between data variables and help decompose complex data relations to support understanding by non-experts or general audiences, which is important because people naturally interpret causality from charts, but not always in ways that are appropriate or accurate. This disconnect can lead to incorrect interpretations that, in some circumstances, can have harmful real-world consequences. A number of studies have focused on human perception of causality from static charts, including perceptual illusions [80] and the effect of priors [71], demonstrating that visual designs can either exacerbate or mitigate human biases when interpreting data in terms of causality.

Estimation (12): Once a causal structure is defined and identified, it is often necessary for analysts to estimate the impact and effect size of specific causal relationships. This is crucial for many domain applications, such as epidemiology, where it is important to translate abstract models into quantified measures and actionable insights. As one example, ADVISES transferred medical surveys and geographic information into measurable visual data representations to support inferences about the causes of obesity [61].

Refutation (6): This task involves conducting sensitivity analyses to test the robustness of an observation. The need for refutation commonly appears when there's a need to consider additional factors such as unobserved confounders and alternative model specifications. It is a vital step to ensure decision-makers do not over-trust purely algorithmic outputs. Outcome-Explorer [31] is an example of relevant approaches that enable users to intentionally perturb models and algorithms to help users better understand them.

4.1.2 Causal Framework

Causal inference has been statistically defined in various ways. Reflecting this, VCI research that relies on different underlying definitions may have very different visual forms and capabilities.

Structural Causal Models (24): SCMs, defined by Pearl, are the most widely applied approach to modeling causality [55]. Systems built on SCMs use Directed Acyclic Graphs (DAGs) to encode causal assumptions. Within a DAG, nodes represent variables and directed links represent potential pairwise causal relationships. A system that extracts causal DAGs from data can be viewed as based on the SCM framework even if they do not explicitly mention it (e.g., [21] supports ML algorithms to construct causal graphs from datasets).

Granger (9): First introduced in econometrics, Granger causality is a probabilistic account of causality based on prediction rather than structural intervention [25]. It utilizes conditions in a past event to predict a future event. Reflecting this temporal dependency, it is used for data where variables unfold dynamically over time, such as time-series data, streaming telemetry, or event-sequences. The work of Kesavan et al. leverages Granger causality for progressive causal relation analysis from streaming performance data over time [40].

Counterfactual (8): Counterfactual reasoning is a causal inference approach that involves simulating alternatives to a current condition to help reason about causal effects by comparing observed outcomes under the current condition to hypothesized outcomes from the alternative conditions. This method is highly relevant for tasks that estimate strengths between potential causes and effects while avoiding spurious relationships like confounders. For example, Co-op provides a library to help people reason about counterfactual conditions when interpreting multi-dimensional data [68].

Potential Outcomes (5): The potential outcomes (PO) framework frames causality around comparing treatment versus control groups [57]. This approach is widely applied in health science and epidemiology. VISPUR defines covariate and causality spaces to help users avoid spurious data-driven decisions from news articles [63]. XplainAct [86] and SeqCausal [35] support cohorts analysis for different type of datasets.

Bayesian (3): Unlike SCMs or Potential Outcomes, which focus on formal confounder adjustment, other methods such as Bayesian frameworks model causality as probabilistic structures. In Bayesian causal modeling, causality is treated as a cognitive or statistical prior that is updated continuously as new evidence is observed. Causal Support cognitively models the probability of causal explanations as a Bayesian benchmark and measures the performance of visualization to update a user's internal belief in causal inference [38].

We note that given the increasing complexity of data and tasks, a single system can often rely on multiple frameworks. As one example, CausalVis [27] is designed to support treatment-outcome estimate within a PO framework, while its primary DAG module facilitates interactive and collaborative causal structure modeling which falls within the SCM framework.

4.1.3 Data Type

Input data for causal analysis can take many forms, and the nature of that data can have a significant influence on what users might intend to learn during an analysis. Common data types include **Tabular/Multivariate** (19), **Text/Graph** (16), **Synthetic** (14), **Time-Series** (6), **Event-Sequences** (5), and **Spatial** (3) data. Please see OSF Appendix for a full synthesis of data types with example papers.

4.1.4 Data Scale

VCI has been applied to data with a range of data scale, including: **one-dimensional** (10), **multi-dimensional** (36), **large** (3, hundreds of dimensions or more), and **heterogeneous** (12, e.g., network, map, and social media). We include data scale as a separate coding dimension because a single task can require different visual and computational strategies depending on the number of variables, data modalities, and relational structures.

4.2 The Visual Solution: How

This aspect captures the graphical encodings and interaction paradigms used to translate complex causal mathematics into readable visualizations and analytics. It is of significant importance because design sits at the core of visualization, and a poor design can obscure causal inferences or introduce difficulties in comprehension. In addition to direct graph drawings, we include encodings that communicate causal information through alternative representations, such as partitioning, statistics, narrative, and animation, when these are used to support causal reasoning.

4.2.1 Causal Effect Encodings

Various encodings have been used to visualize cause-and-effect relationships. Broadly, these include *explicit encodings*, which directly show a link between cause and effect, and *implicit encodings*, which use animation metaphors or clustering of objects to convey causality, or guidance values to direct users' exploration of potential causal effects.

Across the reviewed articles, the primary explicit encoding is via direct-acyclic **Node-Link** (33) graphs, with a supplemental matrix view used in one instance [35]. A majority of the visual causal inference tools, from the Visual Causality Analyst [72] to the Causality Explorer [79], use this encoding to support tracing causal pathways.

Other methods do not directly show a linkage between cause and effect. Early work uses **Animation** (7) with different metaphors to convey causality, such as Growing Polygons [16] and Growing Squares [17]. Another way to convey causality between variables is to develop inferential metrics and display visual representations of **Inferential Statistics** (7) accompanied with a data overview, such as prediction scores with heatmaps [11] or analytical guidance panels with data distributions [69]. Some methods visualize **Partitioned Data** (10), such as counterfactual visualizations, which present alternative data subsets via supplemental charts to improve visualized data understanding [67]. Lastly, **Statistical Charts/Text** (19) are also commonly used to convey causality to users, including basic statistical charts in empirical research [38, 71, 80], text narratives [10], and others like parallel coordinates [26].

4.2.2 Causal Interaction

Visual analytics tools for causal inference by their nature include interaction to help explore and convey causal relationships. Here, we consider only interactions related to causal communication within VCI tools; interactions for other stages, such as loading and initializing variables during data pre-processing, are not considered.

Parametric (25): This is the most common interaction, whereby the user manipulates sliders, thresholds, or mathematical weights to adjust analysis parameters, but without changing the underlying causal model. A large number of causal graph-based VCI tools enable such interactions to revise SCM weights. For example, Narratives + Diagrams enables users to change an *effect possible strength* value for causal links through a slider [83].
Constructive (13): This refers to interactions that enable the user to directly alter the underlying topological structure of the causal model, usually by physically drawing, adding, or removing nodes and edges in a causal graph. For example, a recent tool supports interactive graphs and node-level editing to enable visual analysis for multi-outcome causal graphs [21].
Conditioning (9): This includes interactions that are executed by slicing, filtering, or stratifying the dataset to guide and isolate causal estimations from specific subgroups, populations, or clusters. Conditioning interactions are common for tools and systems that require analyzing detailed differences between data samples, such as those visualizing counterfactual alternatives [39, 69].
Steerable (9): With steerable interactions, the user guides part of the causal output from the tool (e.g., through active learning, prompt engineering, or selecting reference samples), rather than directly drawing the model or tweaking global parameters. CausalChat [87] exemplifies steerable interactions by enabling users to refine models via chatting with large language models.

4.3 The Human Factor: Who

The final aspect involves the human-centric component, focusing on the target audience and how the system's efficacy is measured. This aspect is particularly important in VCI because causal reasoning is inseparable from prior knowledge, expertise, and cognitive bias. For this reason, human-centered evaluation is necessary to determine whether a visual design supports robust inference rather than merely increasing engagement.

4.3.1 User Role

Users are the final consumers of any VCI research, but the role of the user may differ based on various factors in VCI research.
Explorer (28): This may be the most common role for exploratory data analysis, whereby the user engages in open-ended data browsing to learn about the dataset's causal mechanisms and discover heterogeneous effects. Most VCI tools treat users as explorers, e.g., CausalMap supports user exploration from complex social media data to understand sentiment contagion [43].
Generator (14): Part of the user's objective is to generate data to aid with the causal analysis, such as providing contextualized input for causal modeling or creating a structural hypothesis. Belief Miner [59] leverages its visual design to assist users in generating and externalizing their prior beliefs on causality.
Interpreter (14): Users are required to understand a visualization via a lens of causality or how a model makes decisions, rather than conducting novel analysis from data. Empirical research in this field commonly treats users as learners to understand how they generate causal hypotheses or perceive causality from visual presentation [66].
Verifier (13): Users are required to audit algorithms, review automated causal discovery, and check for fairness or bias in machine learning models. This role is common in VCI tools aimed at algorithmic auditing and machine learning model probing, where tools like D-BIAS [23] position the user to check, validate, or audit automated models and algorithms.

4.3.2 Target Expertise

VCI research can focus on users with a variety of backgrounds, which also affects appropriate system and visualization designs.

Much VCI research targets **Domain Experts** (29), with backgrounds in professional areas, to assist in easily integrating their contextual knowledge into visual analytics, from manufacturing [6] to product design [11]. Others focus on **Data Analysts** (19) who possess general technical knowledge of analysis and visualization, enabling them to process complex data visualizations, such as techniques to generate multivariate visualization sequences [58], and utilize libraries to support visual causal inferences [27,68]. The remainder targets **Laypeople** (15), the general public and non-specialist audience, to assess or provide novel visual designs to simply but effectively convey causality via visualizations [38, 63, 69].

4.3.3 Evaluation Type

Lastly, as with many visualization-related approaches, VCI research has employed different evaluation methods to provide evidence for efficacy and investigate human factors. The most common type is the **Case Study** (33). Others include **Quantitative User Studies** (17), **Qualitative User Studies** (23), and **Technical Evaluations** (11).

Many VCI research papers contain multiple types of evaluations. 33 papers conduct only a single type of evaluation, 13 contain two evaluation types, 7 contain three evaluation types, and a recent work [43] contains all four types of evaluations. An increasing trend in the number of evaluations per paper is evident [44], with the average number of evaluations in one paper doubling post-2020 from one to two evaluations.

5 The Structural Pillar: VA for Graphical Inference

Analyzing the 44 taxonomy dimensions, we found that a majority of the work investigates the design of interactive systems relying on graphical causal structures. These systems explicitly represent causal assumptions as directed acyclic graphs (DAGs), usually based on **SCM** or **Granger** causal frameworks, and provide **Node-Link** diagrams, often with varying **Interaction** mechanisms for constructing, editing, interrogating, and refining these graphs. Therefore, we refer to this category of DAG-based visual analytics tools as the Structural Pillar, serving as the largest component of VCI research and comprising 32 papers.

5.1 System Design Space

VCI systems are typically organized around the causal inference pipeline, progressing from raw data to validated models. After reviewing their designs, we found that, due in large part to their reliance on SCMs, the system design space centers around causal graphs for both tasks and interactions.

5.1.1 Frameworks

Of the structural pillar papers, 21 explicitly adopt the Structural Causal Model framework. The remainder of the papers mainly employed the Granger causality framework because of the need to process time-dependent data. The dominance of SCM has provided conceptual coherence—a shared language of nodes, edges, confounding, and mediation—that unifies otherwise disparate systems. Tools for algorithm tracing [23], confounding and mediation analysis [82], and unstructured data analysis [42], all build on this common foundation, enabling insights across areas.

However, this dominance has also constrained the design space. Nearly all structural systems use node-link diagrams as their sole graph encoding. A few works support other representations of causal links, e.g., SeqCausal [35] explores matrix representations as complements to graphs. This representational monoculture may introduce similar problems across tools, such as incomprehensible large graphs [70]. It remains unclear whether alternative visualizations might better support graphical causal tasks in specific conditions. For example, whether matrices may scale better to hundreds of variables, compared to node-link graphs, or whether arc diagrams or other methods may better reveal hierarchical structure.

5.1.2 Task Choices

Historically, causal discovery is the major task choice in VCI systems (22). However, in recent years, research focused on communication [21] and refutation [87] has emerged. Meanwhile, more systems have started to include multiple types of tasks together. For example, CausalVis supports a wide range of causal tasks by providing a generalized library for visualizing causal graphs [27]. A multivariate visualization

© 2026 IEEE. This is the author's version of the article that has been published in IEEE Transactions on Visualization and Computer Graphics. The final version of this record is available at: xx.xxxx/TVCG.201x.xxxxxxx/

sequence technique not only allowed deeper discovery of data using automated visualization sequences, but also simplified causal communication by constructing narratives [58]. This emergence potentially reflects a more mature understanding that causal inference is not only about finding the correct model, but also about effective communication and stress-testing assumptions.

5.1.3 Node-Link Diagram Design

Node-link diagrams act as the foundational visual metaphor for this pillar. The primary design factors revolve around layout strategies and encodings of node/edge attributes. Across the literature, papers often default to using nodes to represent variables and directed arrows to indicate causal pathways. This has provided users with a clear representation of the system's structure in many interfaces from early work on Visual Causal Analysts [73] to the more recent CausalVis [27].

However, systems differ in how they handle the visual complexity of these connections in node ordering and edge encodings. For example, most visual designs utilize strict sequential (hierarchical) topological ordering to clearly delineate root causes from derived effects [2, 74, 82], whereas others rely on non-sequential layouts, such as force-directed or clustered, that prioritize spatial proximity and visibility [36, 83], or radial layouts to visualize traffic-triggering relations [84]. Furthermore, papers differ in how they encode the quantitative properties of the causal links; while most tools present uniform structural links [27, 72], others explicitly map the strength of a causal influence to edge width/color [2] or node color [88]. Other link information, such as statistical certainty, trustworthiness, or semantics, can be represented with additional indicators or fuzziness [2, 37]. Moreover, some papers slightly adjust or reorder the graph based on specific use cases, such as Compass [13].

5.1.4 Causal Interaction Design

With respect to interaction, the taxonomy reveals two primary patterns in structural systems.

Parametric interaction appears in 18 structural pillar papers, making it the most common interaction type within them. It enables users to adjust parameters, enabling them to perform actions such as indirectly revising the causal graph. The slider is the most common control, which users can manipulate to adjust effect sizes, time lags, or threshold values, seeing immediate updates in the causal graph output. The Visual Causal Analyst provides sliders for filtering weak causal relations [72]. This pattern supports sensitivity analysis, exploring how conclusions change under different assumptions.

Constructive interaction appears in 10 structural pillar papers, especially in newer papers. This interaction can be used to directly edit node-link structures. It enables users to bring their own knowledge to specify and change causal relationships. For example, Xu and Sun proposed a method to analyze heterogeneous causal inference in event sequences, allowing direct modification of causal relations at three granularities [81]. These systems treat the causal graph as a design artifact that users can iteratively refine. Constructive interaction typically appears together with parametric interaction in these papers (8/10), as they complement each other in terms of direct and indirect edits.

Other interactions are all rare. Conditioning appears in 4 papers, and only as complementary to the main interactions. Steerable only appears in 2 papers that require users to steer tasks in algorithm tracking [23] and model refactoring [87].

5.1.5 Common Evaluation Patterns

Evaluating structural systems is often challenging due to unknown ground truth in real-world data. Their evaluation paradigms therefore rely primarily on case studies to capture the nuances of exploratory tasks, and sometimes include interviews with domain experts as explorers; no benchmark-based technical evaluation is reported. Experts commonly evaluate how well a system identifies causal relations in a real dataset, providing qualitative insights into usability. For example, Guo et al. utilized such mixed methods to demonstrate their utility with both example cases and expert opinions [27].

Some papers do report controlled user studies, with the majority of them relying heavily on subjective ratings or qualitative feedback. In the few studies that analyzed accuracy, the tasks are relatively simple, such as "*how many one-to-one relationships*" [88].

While these methods demonstrate feasibility and face validity, it becomes harder to establish a concrete sense of whether and how much these systems improve causal inference accuracy compared to alternatives. There is still a lack of accuracy-based causal inference quality assessments on complex tasks (e.g., [5]).

6 The Augmentative Pillar: Beyond Graphs

The Structural Pillar leverages graphical causal models to visualize causal relationships, relying primarily on DAG-based visualizations. However, analyzing the **Causal Effect Encoding** across different **Causal Tasks**, we noticed at least two groups of designs other than node-link diagrams among the papers. Under **SCM** frameworks, many systems provided auxiliary representations together with DAGs to help users make inferences. While in other frameworks, e.g., **Counterfactual** and **PO**-based methods, often provide alternative visualizations for causal inference without showing DAGs at all. We grouped these papers within the Augmentative Pillar which explores alternatives—the role of alternative and auxiliary representations in supporting causal reasoning. This smaller but rapidly growing pillar comprises 24 papers, with systems that leverage methods such as auxiliary representations, textual narratives, and counterfactual reasoning, to help users make causal judgments. There are 12 SCM-based papers that also provide auxiliary representations so they were counted in both pillars.

6.1 Auxiliary Representations of Causal Graphs

Because a causal graph primarily illustrates abstract structural hypotheses, it could potentially bias specific information. Therefore, recent VCI systems often also rely on auxiliary representations to ground the graph in evidence from underlying data.

The majority of auxiliary views consist of alternative charts that display data distributions, effect sizes, or additional measures. Common statistics charts such as bar [79], pie [23], and line [31] graphs, as well as boxplots [15] are used to show statistical information. For instance, in Causality Explorer [79], the causal graph functions as an interactive control panel where users explore connections between variables, while an adjacent visualization panel dynamically renders bar charts to reveal the comparative analyses of dimension distributions present in the dataset. Scatterplot [40], bubble charts [45], heatmaps [40, 74], and parallel coordinates [72, 73] are commonly used for relations within data subgroups and partitions. For example, Lu et al. [45] visualize a clustered view using bubble charts to show word concepts in topics. Areachart [45], dotplots [6], small multiples [74], and sequences [35] have been used to visualize time-dependent data like event or timeline. Besides, textual explanations and narratives [10, 83] are also common to assist causal inference, helping users externalize causal belief and understand complex events. Maps are used if the tool works with spatial data [13].

Other systems juxtapose alternative DAG representations, usually using Matrix-like visualizations, allowing users to look into details of the theoretical graph structure against the precise data. For example, SeqCausal [35] provided a causal comparison view using adjacency matrices showing connections and strengths between variables. Questionnaire Explorer [42] shows question combinations with set visualizations. Ultimately, these auxiliary numeric and visual statistics are critical, as they allow analysts to verify whether the observed data actually supports the causal mechanisms proposed by the node-link diagram.

6.2 Alternative Visualizations of Causal Graphs

Other VCI methods provide alternative visualizations instead of DAGs. For example, the core idea of counterfactual VCI is to allow users to explore hypothetical alternative outcomes that statistics alone cannot directly depict, providing benefits in tasks requiring careful validation, probing, and refinement. CFGuidance supported counterfactual-guided

causal discovery in EDA [69]. XPlainAct [86] allows individual-level insight interventions among subgroups.

Besides, Potential Outcome-based VCI also visualizes different cohorts instead of DAGs to compare outcomes. VISPUR [63] visualizes potential confounder variables and covariate spaces between subgroups.

This approach offers several potential advantages as an alternative to graph-based methods. Counterfactuals align with natural human reasoning patterns, perhaps making them more suitable for lay audiences [20]. Rather than abstract graph structures, users see tangible alternative datasets or outcomes, and pose intervention questions ("*what would change if I did X?*").

6.3 Task Characteristics

6.3.1 Detecting Spurious Associations

A practical challenge in causal inference is distinguishing true causal relationships from spurious correlations. VISPUR addresses this problem by providing visual aids for identifying and interpreting spurious associations in data-driven decisions [63]. It does not require users to look at a causal graph, but helps them recognize when observed associations might be misleading. CFGuidance proposed a counterfactual technique to guide users away from misleading correlations [69]. Spurious associations are a core challenge for visual analytics, but by using causal inference methods to detect and avoid spurious associations, such systems may help prevent false positive causal inferences.

6.3.2 Refutation and Effect Estimation

Several tools support effect estimation without requiring a complete graph specification. The Interact tool supports what-if analysis for virtual product design [11]. XPlainAct [86] constructs similar subgroups based on user interests and utilizes local models to compute counterfactual outcomes. These demonstrate, in many applied contexts, users can reason about specific interventions and effects without a full causal graph. Besides, CFGuidance evaluated their performance with a groundtruth causal graph [69], demonstrating that users can correctly estimate major effects purely from counterfactual guidance without a causal graph. Such methods may be preferable to causal graphs in some real-world practices, especially where causal knowledge is partial and uncertain, or users have low expertise about data analysis.

7 The Cognitive Pillar: Perception and Judgments

Mathematical causal inference models have been widely applied in visual analytics. However, their effectiveness hinges on the human user's ability to interpret visual representations accurately [52]. In the last part of our analysis, we summarized a group of **Empirical Study** papers that focused on **Communication** tasks and **Laypeople**, exploring the role of human factors under the VCI process as part of a Cognitive Pillar. This smallest but empirically grounded pillar (8 papers) shifts the focus from building tools to understanding users. In contrast to the other two pillars, the central contribution here is not a system design but empirical evidence about how visual encodings, prior beliefs, and interaction conditions shape causal judgment.

These papers conduct controlled experiments to measure how people make causal judgments from visualizations, what biases they exhibit, and how different visual designs affect causal reasoning.

7.1 The Impact of Visual Encodings

A fundamental challenge in VCI is whether certain visualizations and visual encodings can lead to more correct causal interpretations, even when the data is purely observational. Empirical research has confirmed that people naturally infer causality from visualized data, however there is still a lack of evidence to quantify how visualization choices and other factors contribute to causal conclusions.

In fact, sometimes insights do not clearly align, and even conflict, between different study conditions. For example, bar charts were measured to lead to weaker causal conclusions than scatterplots by two experiments [71, 80], and in a third study did not lead to stronger conclusions than text [38]. Higher levels of aggregations were shown to lead to perception of stronger causal relationships [80], however in a different study, when less aggregated data was combined with strong prior beliefs, people still inferred strong causal relationships [66].

These findings indicate that visual encoding influences perceived causality in charts; however, their effects are conditional on factors like task framing and prior beliefs. Quantifying and predicting the magnitude and direction of these effects requires further empirical investigation.

7.2 Bias in Causal Reasoning

Research in this pillar also provides consistent evidence that human causal judgment deviates from normative standards. The Illusion of Causality phenomenon demonstrated that visualizations can create unwarranted impressions of causal relationships [80]. In addition, the Causal Support study reported that people's causal inferences are insufficiently sensitive to sample size—they overweight small samples and underweight large ones relative to a Bayesian benchmark model [38].

More recent work on causal priors reveals that people's pre-existing beliefs about causal relationships strongly influence how they interpret visualized data. The study found that prior beliefs can sometimes override statistical evidence, leading people to see causal relationships that align with their expectations even when visualized data suggest otherwise [71]. The authors then extended causal priors to high-level comprehension tasks to examine how prior knowledge affects the interpretation of explanatory visualizations [66]. They found that the strength of causal priors shifted users' visualization interpretation behavior. Users can contextualize the data with their priors if the data matches, but rationalize away the data by assuming externalized factors if the data does not match a strong prior. These results underscore that causal reasoning in visualization is an interaction between evidence and expectation rather than a purely data-driven process.

As people are prone to causal illusions, designers should take steps to counteract these biases rather than assuming that accurate perception will follow naturally from accurate data representation. Further, these findings also highlight a fundamental challenge for visual causal inference: people do not approach visualizations as blank slates. They bring prior knowledge, expectations, and biases that interact with patterns in visualized data in complex ways.

7.3 Insights and Challenges of Empirical Causal Models

Mathematical Cognitive Models: Existing research has introduced novel methodological approaches for studying causal reasoning. Causal Support designed a Bayesian model for quantifying how strongly data support alternative causal structures, enabling normative benchmarks against which to compare human judgments [38]. It can transform the evaluation question from accuracy judgments to optimal performance in statistical inference.

Counterfactuals as Cognitive Aids: Empirical research has also generated positive evidence about what works for accurate VCI. For example, studies found that exposing users to counterfactual visualizations improved their ability to infer causal structure and comprehend data [66, 67]. These findings suggest that counterfactuals do not just align with natural reasoning, but can improve people's causal inferences.

A Gap Between Perception and Action: Lastly, compared to the existing VCI systems, a notable limitation of the cognitive research is its focus on perception rather than interaction. As VCI systems move from simple discovery to complex discovery, refutation, and communication tasks, we still have little empirical evidence about how these increasingly advanced interaction designs affect causal reasoning.

8 Towards a More Complete VCI Landscape

This survey reflects the rapid development of VCI research and its growing focus within the visual analytics field. However, it also demonstrates that VCI research remains fragmented within the literature, distributed across a multidisciplinary intersection of statistical methodology, visualization design, and human cognition with a number of critical, unsolved problems.

© 2026 IEEE. This is the author's version of the article that has been published in IEEE Transactions on Visualization and Computer Graphics. The final version of this record is available at: xx.xxxx/TVCG.201x.xxxxxxx/

8.1 Navigating Visualization Design Trade-Offs

The design space of VCI research presents a classic visualization trade-off: intuitiveness versus scalability. This can be seen in several aspects among the design choices.

Node-Link vs. Matrices/Sets: As the dimensionality of observational data grows, node-link diagrams may quickly degrade into visual hairballs, making it hard to trace variables or identify paths. Adjacency matrices/sets can easily solve the visual clutter and edge-crossing problems inherent in dense causal networks as used in [35,43]. However, the trade-off could also face a reduction in path-tracing efficiency. Sometimes, causal inference requires analysts to follow a chain of variables (e.g., $A \rightarrow B \rightarrow C$), which can be highly unintuitive in a matrix.

Hybrid and Hierarchical Layouts: Moving beyond widely-used hierarchical layouts, there is opportunity to explore more hybrid representations, such as NodeTrix [29] (which uses matrices for dense local clusters and links for global connections), or semantic zooming techniques to allow intuitive navigation of complex causal relationships while preserving overall structure.

Graph-Free Causal Visualizations: A radical alternative is to abandon the explicit causal graph entirely. For audiences lacking statistical expertise, VCI tools could embed causal semantics directly into standard charts. For instance, using static charts [69], text [83], or animated transitions [18] to implicitly demonstrate Simpson's Paradox or treatment effects to communicate causality without ever forcing the user to interpret a formal DAG.

8.2 A Need for Stronger Theoretical Foundations

The rapid development of VCI techniques and systems has not been accompanied by a corresponding development of grounded theories of causal inferences that are specific to the visualization context. The majority of existing work relies on established statistical inference foundations such as structural causal modeling using DAGs. The first theoretical call to move analytics to causal discovery in visualization appears as early as the 2010s [8]. One decade later, researchers argued for the need to build VCI theories behind exploratory systems [32], in some cases, exploratory visual analytics could mislead users' causal inferences unintentionally, thus hurting decision-making performance. More recently, the theory of counterfactual reasoning was suggested as a framework for more reliable visual causal inferences in analytics without the use of formal causal graphs [7]. However, these contributions focus on calls for action or specific methodological domains and do not provide a unified, grounded theory for VCI to help explain when, where, why, and how causal inferences are made in visualization and exploratory data analysis.

8.3 Deeper Integration Across Pillars

Structural systems have nearly universally adopted node-link diagrams, yet the cognitive pillar has not systematically tested whether node-link diagrams are actually effective for causal understanding. We assume they work because they match the mathematical formalism, but we have little empirical evidence that they support human reasoning.

The cognitive pillar has documented that people are prone to causal illusions and insufficiently sensitive to sample size. However, existing VCI systems provide few safeguards against these biases. Users can often manipulate graphs and explore parameters without any awareness that their judgments may be potentially distorted.

The augmentative pillar has embraced auxiliary information based on cognitive evidence that they reduce the complexity of causal inference. However, methods such as counterfactuals without explicit causal models have a risk of being logically inconsistency in certain scenarios, e.g., they may yield representations that are not causally valid due to statistical anomalies, and there would be no underlying causal model against which the representations could be validated.

A key open question, therefore, is how do we create systems that simultaneously leverage graphical causal models, simple yet effective visual representations, and empirical knowledge about human biases in causality? We argue that deeper integration across the three pillars in this survey can help realize more effective VCI systems.

8.4 Minding the Evaluation Gap

A critical examination of existing evaluations reveals a gap between what we typically measure and what we ultimately care about. Most evaluations of VCI tools focus on usability metrics, preferences, intuitivity, responsivity, etc. While these questions matter for adoption, they may not address whether the tool actually improves inference accuracy. Moreover, we still lack true measurements and benchmarks needed to make such evaluations of VCI performance possible.

Most typically in the literature, case studies or expert interviews are used to demonstrate that systems can be used in realistic contexts and that experts find them valuable. However, these methods cannot establish quantitatively whether systems actually improve causal inference accuracy, nor can they quantitatively compare competing designs. Laboratory experiments with laypersons have provided measurements of causal judgment under controlled conditions. However, they often use simplified stimuli (typically bivariate relationships) and artificial tasks that are hard to generalize to real-world practices.

Given these concerns, there is an urgent need for common evaluation measures and benchmarks with shared tasks and datasets. These will enable us to compare performance between systems and measure progress. This includes the need for shared ground-truth datasets with known causal relationships, and which contain the complexities of real-world applications. Similarly, we must develop improved ways to measure and cognitively model people's causal inferences.

8.5 Grand Challenges and Future Directions

As VCI matures beyond basic structural modeling and effect estimation, the research community must diversify its approaches to address the complexities of real-world decision-making tasks. Broadening the scope of VCI opens several alternative research pathways.

Visual Causal Storytelling: Integrating GenAI models with VCI could bridge the gap between complex statistical outputs and layperson comprehension. Future systems could automatically generate narrative data stories that explain why an effect occurs, visually guiding the user step-by-step through the causal chain and potential confounders.

Spatiotemporal Causality: For domains like climate science, epidemiology, and urban planning, causality is deeply tied to physical space and time. Transitioning VCI into Immersive Analytics (AR/VR) could allow experts to explore 3D spatial confounding and delayed temporal effects in a much more natural, embodied environment.

Scalability: Existing VCI systems have been demonstrated on datasets with modest numbers of variables (typically a few to 20) for observational data. Yet real-world causal questions can involve thousands of or even more variables, e.g., EHR healthcare data. The field has not seriously grappled with such scalability problems.

Temporal Dynamics: Time introduces fundamental complications for causal inference, such as feedback loops and requires special handling. Preliminary work [13, 35] has started focusing on event-sequence or time-series data. But most treat time as an additional dimension rather than rethinking causal representation for dynamic systems.

Trust and Confidence: Past work has documented that visualizations can create causal illusions from representations of the data [71, 80]. If visualizations can unintentionally induce causal impressions, then a system that presents data visually risks misleading users into seeing causes where only correlations exist. Exploring the trust and skepticism behind such inferences remains underexplored but important.

9 Conclusion

This work reviewed the emergence of visual causal inference (VCI) as a research field. We proposed a taxonomy that reveals that VCI research has developed distinct but complementary approaches to supporting causal reasoning by synthesizing insights within three pillars. The field's future lies in integration, building systems that draw on all three pillars, evaluating them with rigorous measures, and cumulating knowledge across studies to develop evidence-based design principles for VCI. The next decade of VCI research should be defined not by more systems doing more things but by deeper integration and understanding of what works, for whom, and under what conditions, including greater focus on underexplored challenges for future.

Acknowledgments

We thank the reviewers for their insightful comments. This material is based upon work supported by the National Science Foundation under Grant No. 2211845.

References

[1] Z. Armstrong and M. Wattenberg. Visualizing statistical mix effects and simpson's paradox. *IEEE TVCG*, 20(12):2132–2141, 2014. doi: 10.1109/TVCG.2014.2346297 2

[2] J. Bae, T. Helldin, and M. Riveiro. Understanding indirect causal relationships in node-link graphs. *CGF*, 36(3):411–421, June 2017. doi: 10.1111/cgf.13198 4, 7

[3] T. Banitz, T. Hertz, L.-G. Johansson, E. Lindkvist, R. Martínez-Peña, S. Radosavljevic et al. Visualization of causation in social-ecological systems. *Ecology and Society*, 27(1), 2022. doi: 10.5751/es-13030-270131 4

[4] L. Bartram and M. Yao. Animating causal overlays. *CGF*, 27(3):751–758, May 2008. doi: 10.1111/j.1467-8659.2008.01204.x 4

[5] L. Battle and J. Heer. Characterizing exploratory visual analysis: A literature review and evaluation of analytic provenance in tableau. *CGF*, 38(3):145–159, 2019. doi: 10.1111/cgf.13678 7

[6] M. Becher, D. Herr, C. Muller, K. Kurzhals, G. Reina, L. Wagner et al. Situated visual analysis and live monitoring for manufacturing. *IEEE Computer Graphics and Applications*, 42(2):33–44, Mar. 2022. doi: 10.1109/mcg.2022.3157961 4, 6, 7

[7] D. Borland, A. Z. Wang, and D. Gotz. Using counterfactuals to improve causal inferences from visualizations. *IEEE Computer Graphics and Applications*, 44(1):95–104, Jan. 2024. doi: 10.1109/mcg.2023.3338788 4, 9

[8] M. Chen, A. Trefethen, R. Banares-Alcantara, M. Jirotka, B. Coecke, T. Ertl et al. From data analysis and visualization to causality discovery. *Computer*, 44(10):84–87, Oct. 2011. doi: 10.1109/mc.2011.313 4, 9

[9] F. Cheng, Y. Ming, and H. Qu. Dece: Decision explorer with counterfactual explanations for machine learning models. *IEEE TVCG*, 27(2):1438–1447, Feb. 2021. doi: 10.1109/tvcg.2020.3030342 3

[10] A. Choudhry, M. Sharma, P. Chundury, T. Kapler, D. W. S. Gray, N. Ramakrishnan et al. Once upon a time in visualization: Understanding the use of textual narratives for causality. *IEEE TVCG*, 27(2):1332–1342, Feb. 2021. doi: 10.1109/tvcg.2020.3030358 4, 5, 7

[11] V. Ciorna, G. Melançon, F. Petry, and M. Ghoniem. Interact: A visual what-if analysis tool for virtual product design. *Inf. Vis*, 23(2):123–141, Dec. 2023. doi: 10.1177/14738716231216030 4, 5, 6, 8

[12] T. N. Dang, P. Murray, J. Aurisano, and A. G. Forbes. Reactionflow: an interactive visualization tool for causality analysis in biological pathways. *BMC Proceedings*, 9(S6), Aug. 2015. doi: 10.1186/1753-6561-9-s6-s6 4

[13] Z. Deng, D. Weng, X. Xie, J. Bao, Y. Zheng, M. Xu et al. Compass: Towards better causal analysis of urban time series. *IEEE TVCG*, 28(1):1051–1061, Jan. 2022. doi: 10.1109/tvcg.2021.3114875 4, 7, 9

[14] S. Edwards, R. Lilford, D. Braunholtz, J. Hewison, J. Jackson, and T. Thornton. Ethical issues in the design and conduct of randomised controlled trials. *Health Technology Assessment*, 2(15):1–132, 1999. doi: 10.4135/9781848608344.n1 2

[15] J. Eirich, G. Koutroulis, B. Mutlu, D. Jackle, R. Kern, T. Schreck et al. Manex: The visual analysis of measurements for the assessment of errors in electrical engines. *IEEE Computer Graphics and Applications*, 42(2):68–80, Mar. 2022. doi: 10.1109/mcg.2022.3155306 4, 7

[16] N. Elmqvist and P. Tsigas. Causality visualization using animated growing polygons. In *IEEE Symposium on Information Visualization 2003*, INFVIS-03, pp. 189–196. IEEE. doi: 10.1109/infvis.2003.1249025 3, 4, 5

[17] N. Elmqvist and P. Tsigas. Growing squares: animated visualization of causal relations. In *Proceedings of the 2003 ACM symposium on Software visualization*, SOFTVIS03, p. 17. ACM, June 2003. doi: 10.1145/774833.774836 2, 3, 4, 5

[18] N. Elmqvist and P. Tsigas. Animated visualization of causal relations through growing 2d geometry. *Inf. Vis*, 3(3):154–172, Sept. 2004. doi: 10.1057/palgrave.ivs.9500074 3, 4, 9

[19] N. Elmqvist and P. Tsigas. Citewiz: A tool for the visualization of scientific citation networks. *Inf. Vis*, 6(3):215–232, Sept. 2007. doi: 10.1057/palgrave.ivs.9500156 4

[20] K. Epstude and N. J. Roese. The functional theory of counterfactual thinking. *Personality and social psychology review*, 12(2):168–192, 2008. doi: 10.1177/1088868308316091 8

[21] M. Fan, J. Yu, D. Weiskopf, N. Cao, H.-Y. Wang, and L. Zhou. Visual analysis of multi-outcome causal graphs. *IEEE TVCG*, 31(01):656–666, 2025. doi: 10.1109/TVCG.2024.3456346 4, 5, 6

[22] R. A. Fisher. On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London*, 222(594-604):309–368, 1922. doi: 10.1007/978-1-4612-0919-5_2 2

[23] B. Ghai and K. Mueller. D-bias: A causality-based human-in-the-loop system for tackling algorithmic bias. *IEEE TVCG*, pp. 1–10, 2022. doi: 10.1109/tvcg.2022.3209484 4, 6, 7

[24] O. Gomez, S. Holter, J. Yuan, and E. Bertini. Vice: Visual counterfactual explanations for machine learning models. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*, pp. 531–535, 2020. doi: 10.1145/3377325.3377536 3

[25] C. W. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, pp. 424–438, 1969. doi: 10.2307/1912791 5

[26] G. Guo, M. Glenski, Z. Shaw, E. Saldanha, A. Endert, S. Volkova et al. Vaine: Visualization and ai for natural experiments. In *2021 IEEE VIS*, pp. 21–25, 2021. doi: 10.1109/VIS49827.2021.9623285 4, 5

[27] G. Guo, E. Karavani, A. Endert, and B. C. Kwon. Causalvis: Visualizations for causal inference. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pp. 1–20. ACM, Apr. 2023. doi: 10.1145/3544548.3581236 4, 5, 6, 7

[28] S. D. Halpern, J. H. Karlawish, and J. A. Berlin. The continuing unethical conduct of underpowered clinical trials. *JAMA*, 288(3):358–362, 2002. doi: 10.1001/jama.288.3.358 2

[29] N. Henry, J.-D. Fekete, and M. J. McGuffin. Nodetrix: a hybrid visualization of social networks. *IEEE TVCG*, 13(6):1302–1309, 2007. doi: 10.1109/TVCG.2007.70582 9

[30] F. Hohman, M. Kahng, R. Pienta, and D. H. Chau. Visual analytics in deep learning: An interrogative survey for the next frontiers. *IEEE TVCG*, 25(8):2674–2693, 2018. doi: 10.1109/TVCG.2018.2843369 3

[31] M. N. Hoque and K. Mueller. Outcome-explorer: A causality guided interactive visual interface for interpretable algorithmic decision making. *IEEE TVCG*, 28(12):4728–4740, Dec. 2022. doi: 10.1109/tvcg.2021.3102051 4, 5, 7

[32] J. Hullman and A. Gelman. Designing for interactive exploratory data analysis requires theories of graphical inference. *Harvard Data Science Review*, July 2021. doi: 10.1162/99608f92.3ab8a587 2, 4, 9

[33] F. Husain, P. Proulx, M.-W. Chang, R. Romero-Gómez, and H. Vasquez. A mixed-initiative visual analytics approach for qualitative causal modeling. In *2021 IEEE VIS*, pp. 121–125, 2021. doi: 10.1109/VIS49827.2021.9623318 4

[34] A. R. Jadad and M. W. Enkin. *Randomized controlled trials: questions, answers and musings*. John Wiley & Sons, 2007. doi: 10.1002/9780470691922 2

[35] Z. Jin, S. Guo, N. Chen, D. Weiskopf, D. Gotz, and N. Cao. Visual causality analysis of event sequence data. *IEEE TVCG*, 27(2):1343–1352, Feb. 2021. doi: 10.1109/tvcg.2020.3030465 1, 2, 3, 4, 5, 6, 7, 9

[36] C. Jung, S. Yim, G. Park, S. Oh, and Y. Jang. Catom: Causal topology map for spatiotemporal traffic analysis with granger causality in urban areas. *IEEE TVCG*, 31(9):6036–6052, Sept. 2025. doi: 10.1109/tvcg.2024.3489676 4, 7

[37] N. R. Kadaba, P. P. Irani, and J. Leboe. Visualizing causal semantics using animations. *IEEE TVCG*, 13(6):1254–1261, Nov. 2007. doi: 10.1109/tvcg.2007.70528 3, 4, 7

[38] A. Kale, Y. Wu, and J. Hullman. Causal support: Modeling causal inferences with visualizations. *IEEE TVCG*, 28(1):1150–1160, Jan. 2022. doi: 10.1109/tvcg.2021.3114824 1, 2, 3, 4, 5, 6, 8

[39] S. Kaul, D. Borland, N. Cao, and D. Gotz. Improving visualization interpretation using counterfactuals. *IEEE TVCG*, 28(1):998–1008, Jan. 2022. doi: 10.1109/tvcg.2021.3114779 4, 6

[40] S. P. Kesavan, T. Fujiwara, J. K. Li, C. Ross, M. Mubarak, C. D. Carothers et al. A visual analytics framework for reviewing streaming performance data. In *2020 IEEE PacificVis*, pp. 206–215. IEEE, June 2020. doi: 10.1109/pacificvis48177.2020.9280 4, 5, 7

[41] X. Lan, Y. Wu, and N. Cao. Affective visualization design: Leveraging the emotional impact of data. *IEEE TVCG*, 30(1):1–11, 2023. doi: 10.1109/TVCG.2023.3327385 3

[42] R. Li, W. Cui, T. Song, X. Xie, R. Ding, Y. Wang et al. Causality-based visual analysis of questionnaire responses. *IEEE TVCG*, 30(1):638–648, Jan. 2024. doi: 10.1109/tvcg.2023.3327376 2, 4, 6, 7

© 2026 IEEE. This is the author's version of the article that has been published in IEEE Transactions on Visualization and Computer Graphics. The final version of this record is available at: xx.xxxx/TVCG.201x.xxxxxxx/

[43] R. Li, S. Ye, Y. Lin, B. Zhou, Z. Kang, T.-Q. Peng et al. Causality-based visual analytics of sentiment contagion in social media topics. *IEEE TVCG*, 32(1):35–45, 2026. doi: 10.1109/TVCG.2025.3633839 2, 4, 6, 9

[44] F. Lin, A. Z. Wang, M. D. Rahman, D. A. Szafir, and G. J. Quadri. Striking the right balance: Systematic assessment of evaluation method distribution across contribution types. *2024 IEEE Evaluation and Beyond-Methodological Approaches for Visualization (BELIV)*, pp. 129–135, 2024. doi: 10.1109/BELIV64461.2024.00020 6

[45] Y. Lu, H. Wang, S. Landis, and R. Maciejewski. A visual analytics framework for identifying topic drivers in media events. *IEEE TVCG*, 24(9):2501–2515, Sept. 2018. doi: 10.1109/tvcg.2017.2752166 4, 7

[46] A. McNutt, G. Kindlmann, and M. Correll. Surfacing visualization mirages. In *ACM CHI 2020*, pp. 1–16, 2020. doi: 10.1145/3313831.3376420 2

[47] K. Mehta, G. Kindlmann, and A. Kale. Designing for disclosure in data visualizations. *IEEE TVCG*, 2025. doi: 10.1109/TVCG.2025.3634781 2

[48] A. Michotte. *The perception of causality*. Routledge, 1946. doi: 10.1093/bjps/XV.59.254 3

[49] D. Moher, A. Liberati, J. Tetzlaff, and D. G. Altman. Preferred reporting items for systematic reviews and meta-analyses: the prisma statement. *Bmj*, 339, 2009. doi: 10.1371/journal.pmed.1000097 3

[50] T. Munzner. Process and pitfalls in writing information visualization research papers. In *Information visualization: human-centered issues and perspectives*, pp. 134–153. Springer, 2008. doi: 10.1007/978-3-540-70956-5_6 3

[51] T. Munzner. *Visualization analysis and design*. CRC press, 2014. doi: 10.1145/3721241.3733989 1

[52] R. J. Musci and E. Stuart. Ensuring causal, not casual, inference. *Prevention Science*, 20(3):452–456, 2019. doi: 10.1007/s11121-018-0971-9 2, 8

[53] J. B. Nassar, Y. X. Yio, N. Athukorala, Simran, S. Fan, C. A. Huang et al. Out of the loop: Enhancing documentation and transparency in collaborative causal loop diagrams to capture multiple perspectives. In *2025 IEEE Visualization and Visual Analytics (VIS)*, pp. 111–115, 2025. doi: 10.1109/VIS60296.2025.00028 4

[54] Y. Onoue, K. Kyoda, M. Kioka, K. Baba, S. Onami, and K. Koyamada. Development of an integrated visualization system for phenotypic character networks. In *2018 IEEE PacificVis*, pp. 21–25. IEEE, Apr. 2018. doi: 10.1109/pacificvis.2018.00012 4

[55] J. Pearl. *Causality*. Cambridge University Press, Cambridge, 2009. doi: 10.1017/CBO9780511803161 1, 2, 3, 5

[56] K. Pearson. Notes on the history of correlation. *Biometrika*, 13(1):25–45, 1920. doi: 10.1093/biomet/13.1.25 2

[57] D. B. Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American statistical Association*, 100(469):322–331, 2005. doi: 10.1198/016214504000001880 1, 2, 5

[58] P. Ruchikachorn, D. Coelho, J. Wang, K. Striegnitz, and K. Mueller. Generating coherent visualization sequences for multivariate data by causal graph traversal. *IEEE TVCG*, 32(3):2812–2824, 2026. doi: 10.1109/TVCG.2026.3656952 4, 6, 7

[59] S. Salim, M. N. Hoque, and K. Mueller. Belief miner: A methodology for discovering causal beliefs and causal illusions from general populations. *PACMHCI*, 8(CSCW1):1–37, Apr. 2024. doi: 10.1145/3637298 4, 6

[60] A. Sarikaya and M. Gleicher. Scatterplots: Tasks, data, and designs. *IEEE TVCG*, 24(1):402–412, 2018. doi: 10.1109/TVCG.2017.2744184 3

[61] A. Sutcliffe, O. de Bruijn, S. Thew, I. Buchan, P. Jarvis, J. McNaught et al. Developing visualization-based decision support tools for epidemiology. *Inf. Vis*, 13(1):3–17, May 2012. doi: 10.1177/1473871612445832 4, 5

[62] D. A. Szafir, R. Borgo, M. Chen, D. J. Edwards, B. Fisher, and L. Padilla. *Visualization Psychology*. Springer Nature, New York, 2023. doi: 10.1007/978-3-031-34738-2 1

[63] X. Teng, Y. Ahn, and Y.-R. Lin. Vispur: Visual aids for identifying and interpreting spurious associations in data-driven decisions. *IEEE TVCG*, 30(1):219–229, 2024. doi: 10.1109/TVCG.2023.3326587 1, 2, 4, 5, 6, 8

[64] J. W. Tukey et al. *Exploratory Data Analysis*, vol. 2. Addison-Wesley, Reading, 1977. doi: 10.1007/978-3-031-20719-8_2 2

[65] T. von Landesberger, S. Diel, S. Bremm, and D. W. Fellner. Visual analysis of contagion in networks. *Inf. Vis*, 14(2):93–110, May 2013. doi: 10.1177/1473871613487087 4

[66] A. Z. Wang, D. Borland, E. Calcaterra, and D. Gotz. Contextualization or rationalization? the effect of causal priors on data visualization interpretation. *IEEE TVCG*, pp. 1–15, 2026. doi: 10.1109/tvcg.2026.3663050 1, 3, 4, 6, 8

[67] A. Z. Wang, D. Borland, and D. Gotz. An empirical study of counterfactual visualization to support visual causal inference. *Inf. Vis*, 23(2):197–214, Feb. 2024. doi: 10.1177/14738716241229437 4, 5, 8

[68] A. Z. Wang, D. Borland, and D. Gotz. A framework to improve causal inferences from visualizations using counterfactual operators. *Inf. Vis*, 24(1):24–41, Aug. 2024. doi: 10.1177/14738716241265120 4, 5, 6

[69] A. Z. Wang, D. Borland, and D. Gotz. Beyond correlation: Incorporating counterfactual guidance to better support exploratory visual analysis. *IEEE TVCG*, 31(1):776–786, Jan. 2025. doi: 10.1109/tvcg.2024.3456369 1, 2, 4, 5, 6, 8, 9

[70] A. Z. Wang, D. Borland, and D. Gotz. Visual analytics for causal reasoning from real-world health data. In *Grand Challenges Proceedings of IEEE VAHC Workshop*, 2025. doi: 10.48550/arXiv.2508.17474 6

[71] A. Z. Wang, D. Borland, T. C. Peck, W. Wang, and D. Gotz. Causal priors and their influence on judgements of causality in visualized data. *IEEE TVCG*, 31(1):765–775, Jan. 2025. doi: 10.1109/tvcg.2024.3456381 1, 2, 3, 4, 5, 8, 9

[72] J. Wang and K. Mueller. The visual causality analyst: An interactive interface for causal reasoning. *IEEE TVCG*, 22(1):230–239, Jan. 2016. doi: 10.1109/tvcg.2015.2467931 2, 3, 4, 5, 7

[73] J. Wang and K. Mueller. Visual causality analysis made practical. In *2017 IEEE VAST*, pp. 151–161. IEEE, Oct. 2017. doi: 10.1109/vast.2017.8585647 2, 3, 4, 7

[74] J. Wang and K. Mueller. Domino: Visual causal reasoning with time-dependent phenomena. *IEEE TVCG*, 29(12):5342–5356, Dec. 2023. doi: 10.1109/tvcg.2022.3207929 4, 7

[75] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers Inc., San Francisco, 2nd ed., 2004. doi: 10.1016/B978-155860819-1/50001-7 1

[76] C. Ware, E. Neufeld, and L. Bartram. Visualizing causal relations. In *IEEE InfoVis 1999*, vol. 99, pp. 39–42, 1999. doi: ccom/206/ 2, 3, 4

[77] J. Wexler, M. Pushkarna, T. Bolukbasi, M. Wattenberg, F. Viegas, and J. Wilson. The what-if tool: Interactive probing of machine learning models. *IEEE TVCG*, 26(1):56 – 65, 2020. doi: 10.1109/tvcg.2019.2934619 3

[78] S. Wright. Correlation and causation. *Journal of Agricultural Research*, 20(7):557, 1921. doi: bibtex/6cd977bb9d1ce8fb96b7f2498ba3e143 1

[79] X. Xie, F. Du, and Y. Wu. A visual analytics approach for exploratory causal analysis: Exploration, validation, and applications. *IEEE TVCG*, 27(2):1448–1458, Feb. 2021. doi: 10.1109/tvcg.2020.3028957 4, 5, 7

[80] C. Xiong, J. Shapiro, J. Hullman, and S. Franconeri. Illusion of causality in visualized data. *IEEE TVCG*, 26(1):853–862, Jan. 2020. doi: 10.1109/tvcg.2019.2934399 1, 2, 3, 4, 5, 8, 9

[81] S. Xu and M. Sun. Detecting, interpreting and modifying the heterogeneous causal network in multi-source event sequences. *CGF*, 44(1), Nov. 2024. doi: 10.1111/cgf.15267 4, 7

[82] C. E. Yen, A. Parameswaran, and W. Fu. An exploratory user study of visual causality analysis. *CGF*, 38(3):173–184, June 2019. doi: 10.1111/cgf.13680 2, 4, 6, 7

[83] C.-H. E. Yen, H. Cheng, G. Y.-C. Yen, B. P. Bailey, and Y. Huang. Narratives + diagrams: An integrated approach for externalizing and sharing people's causal beliefs. *PACMHCI*, 5(CSCW2):1–27, Oct. 2021. doi: 10.1145/3479588 2, 4, 6, 7, 9

[84] H. Zhang, M. Sun, D. D. Yao, and C. North. Visualizing traffic causality for analyzing network anomalies. In *2015 ACM IWSPA Workshop*, CODASPY'15, pp. 37–42. ACM, Mar. 2015. doi: 10.1145/2713579.2713583 4, 7

[85] X. Zhang, X. Yang, H. Hu, and H. Qin. Visual causal analysis of multivariate time series. *Journal of Visualization*, 28(2):323–339, Jan. 2025. doi: 10.1007/s12650-024-01041-6 4

[86] Y. Zhang, K. Hegde, and K. Mueller. Xplainact: Visualization for personalized intervention insights. In *2025 IEEE Visualization and Visual Analytics (VIS)*, pp. 161–165, 2025. doi: 10.1109/VIS60296.2025.00038 4, 5, 8

[87] Y. Zhang, A. Kota, E. Papenhausen, and K. Mueller. Causalchat: Interactive causal model development and refinement using large language models. *IEEE TVCG*, 31(12):10473–10487, 2025. doi: 10.1109/TVCG.2025.3602448 4, 6, 7

[88] S. Zhu, G. Sun, Y. Shen, Z. Zhu, W. Xia, B. Chang et al. Vac2: Visual analysis of combined causality in event sequences. *IEEE TVCG*, 31(9):6223–6240, Sept. 2025. doi: 10.1109/tvcg.2024.3496789 4, 7