Report from Dagstuhl Seminar 25511

From Psychology Enabled Visualization to Visualization Enabled Psychology

Rita Borgo*[1], Min Chen*[2], David S. Ebert*[3], Brian D. Fisher*[4], and Alvitta Ottley*[5]

1 **King's College London, GB.** rita.borgo@kcl.ac.uk
2 **University of Oxford, GB.** min.chen@oerc.ox.ac.uk
3 **University of Arizona – Tucson, US.** ebert@arizona.edu
4 **Simon Fraser University – Surrey, CA.** fisher@sfu.ca
5 **Washington University in St. Louis, US.** alvitta@wustl.edu

Abstract

This report of the Dagstuhl Seminar "From Psychology Enabled Visualization to Visualization Enabled Psychology" (25511) advances Visualization Psychology as an emerging interdisciplinary field that integrates behavioral science with visualization research and practice, while also positioning visualization as a powerful environment for studying cognition in action. Organized around six complementary themes – *foundational theory-building, learning and friction, decision-making, storytelling, perceptual effects, and human – AI interaction* – the report outlines key challenges and opportunities for understanding visualization as a human-centered process. We propose a research agenda that emphasizes integrative theory, systematic empirical methods, and interdisciplinary collaboration. Altogether, we bridge psychology-enabled visualization with visualization-enabled psychology, to establish Visualization Psychology as a cohesive field and to support the design of more effective, inclusive, and human-centered visualization systems.

Seminar December 14–19, 2025 – https://www.dagstuhl.de/25511

2012 ACM Subject Classification Human-centered computing → Visualization; Computing methodologies → Cognitive science

Keywords and phrases Visualization Psychology, Cognitive Science, Interdisciplinary Collaboration, Dagstuhl Seminar

Digital Object Identifier 10.4230/DagRep.15.12.1

1 Executive Summary

Rita Borgo (King's College London, GB, rita.borgo@kcl.ac.uk)
Min Chen (University of Oxford, GB, min.chen@ox.ac.uk)
David S. Ebert (University of Arizona – Tucson, US, ebert@arizona.edu)
Brian D. Fisher (Simon Fraser University – Surrey, CA, fisher@sfu.ca)
Alvitta Ottley (Washington University in St. Louis, US, alvitta@wustl.edu)

The Dagstuhl Seminar "From Psychology Enabled Visualization to Visualization Enabled Psychology" (25511) positions Visualization Psychology as a necessary and emerging interdisciplinary field. Visualization Psychology seeks to systematically integrate theories, methods, and empirical findings from the behavioral sciences into visualization research and practice, while also leveraging visualization as a powerful environment for studying cognition in action.

* Editor / Organizer

Dagstuhl Seminar 25511: From Psychology Enabled Visualization to Visualization Enabled Psychology,
Dagstuhl Reports, Vol. 15, Issue 12, pp. 1–78
Editors: Rita Borgo and Min Chen and David S. Ebert and Brian D. Fisher and Alvitta Ottley
DAGSTUHL REPORTS Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

This bidirectional framing – *from psychology-enabled visualization to visualization-enabled psychology* – expands the scope of both fields and opens new avenues for discovery.

On one hand, **psychology-enabled visualization** emphasizes how insights from perception, cognition, learning, and social behavior can inform the design and evaluation of visualizations. This includes understanding perceptual mechanisms that shape how visual encodings are processed, cognitive processes involved in reasoning with data (e.g., uncertainty, trust, and decision-making), and how interaction supports or hinders sensemaking. It also calls for grounding design guidance in theory, moving beyond heuristic "best practices" toward explanations of why certain designs work, for whom, and under what conditions.

On the other hand, **visualization-enabled psychology** recognizes visualization as a rich and underutilized testbed for studying complex cognitive phenomena. Interactive visualizations provide controlled yet expressive environments in which researchers can examine high-level processes such as hypothesis generation, causal reasoning, confidence, and collaborative sensemaking. These settings extend beyond traditional laboratory paradigms, enabling the study of cognition as it unfolds in realistic, task-driven contexts.

To structure this emerging field, the seminar brought together six complementary working groups, each addressing a critical dimension of Visualization Psychology:

- **Theories for Visualization Psychology** establishes a foundational structure through three efforts: a problem-driven framework grounded in psychological theory, a theory-inspired taxonomy for organizing visualization knowledge, and a conceptual alignment matrix linking psychological constructs to visualization concepts. Together, these efforts aim to create the intellectual scaffolding needed for cumulative progress.
- **Learning and Friction in Visualization** examines how difficulty and processing cost (conceptualized as friction) can both support and hinder learning. Drawing on learning sciences and HCI, this work introduces a Person – Environment – Behavior framework to explain when friction enhances understanding (e.g., through desirable difficulty) versus when it impedes it.
- **The Psychology of Decision Making with Visualization** focuses on real-world decision contexts, where interpretation unfolds under constraints such as time pressure, uncertainty, and coordination. By connecting visualization design to theories and methods from decision science, this work aims to better support situated, high-stakes reasoning.
- **The Psychology of Storytelling with Visualizations** explores how visual communication influences public understanding and behavior, with a focus on domains such as climate change. This work highlights how visualization can shape engagement, emotion, and action, bridging cognitive and social dimensions of interpretation.
- **Perceptual Illusions and Visualization** investigates systematic divergences between what is encoded and what is perceived or interpreted. This line of work connects visualization to perceptual and cognitive mechanisms, offering insight into when such effects are misleading, negligible, or even beneficial.
- **Artificial Intelligence, Psychology, and Visualization** examines how emerging AI systems, including LLMs and visual models, interpret visualizations, and how their responses compare to human cognition. This raises new questions about shared biases, differences in interpretation, and the design of visualization systems for human – AI collaboration.

Throughout the week-long seminar, several unifying challenges emerged: linking perceptual and cognitive theory to visualization design; expanding measures of effectiveness beyond accuracy to include trust, confidence, and behavior; accounting for variability across users and contexts; and explaining real-world phenomena that are not adequately captured by existing theories in either discipline. These challenges also intersect with broader concerns around inclusivity, ethics, enactive cognition in interactive systems, and behavioral prediction.

To address these challenges, we outlined a research agenda that emphasizes integrative theory-building, systematic empirical studies, and the development of shared methods and resources. We highlighted opportunities for interdisciplinary collaboration across visualization, psychology, HCI, design, and related fields, as well as the need to translate insights into practice through tools, educational materials, and design frameworks.

2 Table of Contents

Part I

Foundations

3 Theories for Visualization Psychology

3.1 Organizing Knowledge, Aligning Concepts, and Charting a Research Program

Alfie Abdul-Rahman (King's College London, GB, alfie.abdulrahman@kcl.ac.uk)
Leslie M. Blaha (Air Force Research Laboratory – Arlington, US, leslie.blaha@us.af.mil)
Min Chen (University of Oxford, GB, min.chen@oerc.ox.ac.uk)
Amy Rae Fox (University of California, San Diego, US, amyraefox@ucsd.edu)
Petra Isenberg (Université Paris-Saclay, CNRS, Inria, LISN, FR, petra.isenberg@inria.fr)
David H. Laidlaw (Brown University – Providence, US, dhl@cs.brown.edu)
Andrew I. Meso (King's College London, GB, andrew.meso@kcl.ac.uk)
Florian Windhager (Donau-Universität Krems, AT, florian.windhager@donau-uni.ac.at)

Visualization Psychology (VisPsych) has emerged as a scientific research program at the intersection of computing, behavioral science, and design [1]. Yet the field remains conceptually diverse, theoretically pluralistic, and organizationally fragmented. In this section, we propose three complementary efforts to structure and advance VisPsych: (1) a visualization research problem – inspired framework grounded in psychological theory, (2) a theories-inspired taxonomy for organizing visualization knowledge, and (3) a VisPsych conceptual alignment matrix mapping psychological constructs to visualization concepts.

Rather than offering a comprehensive review, we articulate shared foundations, highlight productive tensions, and identify gaps where theory is underdeveloped. We argue that theory enables formalization, prediction, systematic comparison, and principled evaluation [2], yet must remain sensitive to visualization's interpretive, distributed, and socio-cultural dimensions [3, 4]. We conclude with open challenges and a roadmap for coordinated research and interdisciplinary collaboration.

3.1.1 Introduction & Motivation

Visualization research has long drawn on psychological theories of perception, cognition, and decision-making. Classic work on graphical perception [5], working memory [6], attention [7], and heuristics and biases [8] inform everyday visualization design questions.

Recent VisPsych work emphasizes that visualization is not merely a perceptual artifact but part of a distributed cognitive system [3, 9]. Meaning is not transmitted but constructed through interaction with external representations [10]. Fox and Hollan characterize Visualization Psychology as a scientific research program linking visualization phenomena and behavioral theory [1].

However, as noted in the Dagstuhl discussions, "Visualization Psychology is enormously diverse" [11]. Its theoretical landscape spans perceptual psychophysics, cognitive load theory [12], dual coding [13], ecological perception [14], affect and persuasion [15], semiotics [16], and distributed cognition [3]. This diversity motivates the need for clearer organization and theoretical articulation.

3.1.2 Scope & Definitions

Visualization is construed as a form of *external representation* [10, 9, 17], encompassing both the artifact (a constructed visual representation of data) and the process through which it is created, interpreted, and used. Interaction with visualization is therefore not a passive transmission or extraction of meaning, but an active, interpretive, and semiotic process in which viewers construct meaning in relation to tasks, goals, and context. From a distributed cognition perspective [3], intelligent action with visualization emerges from systems comprising human actors, material artifacts, and their spatio-temporal environment, rather than from isolated perceptual events. This frame of reference positions visualization as part of a broader cognitive ecology that includes social, institutional, and technological dimensions.

We adopt a broad conception of theory as a structured explanatory framework that clarifies concepts, generates predictions, and organizes empirical findings [2]. In the context of Visualization Psychology, theories may be formal, semi-formal, or conceptual, but they must enable systematic comparison, operationalization of constructs, and cumulative knowledge building. We distinguish between importing established psychological theories, adapting them to visualization contexts, and articulating candidate theories that emerge from the interdisciplinary space itself. This scope intentionally includes perceptual, cognitive, affective, and social perspectives, while recognizing that some explanatory questions may require engagement with neighboring disciplines beyond psychology.

3.1.3 Relevant Theories & Empirical Foundations

3.1.3.1 Perception

Perceptual theories provide foundational constraints and opportunities for visualization design. The Gestalt principles [18] explain how viewers spontaneously group elements, informing layout, enclosure, and panel organization. Research on pre-pretentive processing and Feature Integration Theory [7] distinguishes rapid single-feature "pop-out" from slower conjunction search, guiding effective highlighting and attention management. Psychophysical principles such as just-noticeable differences (JND) and the Weber – Fechner law constrain discriminability in color, size, and position, with implications for colormap design and multi-attribute encodings [19]. Ecological perception [14] emphasizes task and context, suggesting that perceptual effectiveness depends on situational goals and affordances rather than isolated visual properties. Finally, Signal Detection Theory [20] formalizes sensitivity, threshold and classification trade-offs, quantifying how visualizations communicate uncertainty, alarms, and decision boundaries. Together, these theories provide mechanistic explanations for encoding choice, grouping, salience, and threshold design in Visualization Psychology.

3.1.3.2 Cognition

Cognitive theories explain how viewers interpret, integrate, and retain information from visualizations. Working memory models [6] and Cognitive Load Theory [12] highlight strict limits on how much visual and verbal information can be processed simultaneously, motivating simplification, chunking, and progressive disclosure in design. Mental models shape how viewers map visual structure onto prior knowledge, influencing metaphor choice and representational conventions [17]. Dual-coding theory [13] suggests that combining visual and verbal explanations can enhance comprehension and recall when properly aligned. Attention theory and research on change blindness [21] further demonstrate that viewers may miss important transitions or relationships unless cues explicitly guide focus. Together, these perspectives elucidate why comprehension, comparison, and interaction depend not only on perceptual salience but also on cognitive capacity, prior knowledge, and task structure.

3.1.3.3 Judgment and Decision-Making

Theories of judgment and decision-making explain how viewers draw conclusions from visual evidence under uncertainty. Research on heuristics and biases [8] shows that anchoring, framing, availability, and representativeness can systematically shape interpretation, meaning that axis ranges, baselines, normalization choices, and visual emphasis may subtly influence conclusions. Framing effects highlight how equivalent data can lead to different judgments depending on presentation (e.g., gains vs. losses). Risk perception research further informs how uncertainty encodings – such as intervals, distributions, or ensemble displays – are understood or ignored. Together, these perspectives underscore that visualizations do not merely present information but actively structure decision environments, making principled design of scales, thresholds, and uncertainty representations essential in Visualization Psychology.

3.1.3.4 Affect and Social Cognition

Affective and social cognitive theories explain how visualization influences engagement, trust, and persuasion beyond perceptual accuracy. Research on emotion-cognition interaction shows that affect can guide attention, shape interpretation, and influence risk perception, meaning that color, imagery, and narrative framing can alter not only engagement but judgment. Persuasion models such as the Elaboration Likelihood Model [15] distinguish between central (analytic) and peripheral (heuristic) processing routes, suggesting that visualization effectiveness depends on audience motivation, expertise, and context. In high-engagement settings, detailed evidence and uncertainty may support careful reasoning, whereas in low-attention contexts, visual salience and narrative cues may dominate interpretation. These perspectives highlight that visualization operates within social and communicative environments, where credibility, framing, and audience characteristics shape meaning construction.

3.1.3.5 Visualization-Specific Foundations

Visualization-specific research bridges general psychological theory and design practice. Graphical perception studies [5] empirically compare encodings (e.g., position, length, angle, area), demonstrating systematic differences in perceptual accuracy and informing canonical design recommendations. Task-based evaluation paradigms further align visualization effectiveness with specific judgment goals, such as comparison, trend detection, clustering, or anomaly identification, rather than with abstract notions of clarity. More recent work extends beyond precision to consider interaction, sensemaking, and insight, linking evaluation to cognitive processes and user goals. Together, these foundations ground Visualization Psychology in empirically validated design principles while highlighting the need to connect controlled perceptual findings with richer, real-world analytic and communicative contexts.

3.1.4 Theories and Practical Case Studies

In psychology, theoretical discourse is commonplace, and many theories have been proposed, such as:

- Conceptual Mapping Theory [22],
- Cognitive Load Theory [12],

- Distributed Cognition [23],
- Dual Coding Theory [24],
- Dual-Systems Theory [25],
- Emotional Design Theory [26]
- Gestalt Principles of Perception [27],
- Human Information Processing Theory [28],
- Mental Model Theory [29, 30],
- *and many more.*

Meanwhile, there are also a few theories proposed in the field of visualization, such as:

- Algebraic Process Theory [31],
- Data-ink Ratio [32],
- Information-theoretic Cost-Benefit Analysis [33],
- Semantic Discriminability Theory [34].

There is no doubt that many of these theories can explain some phenomena in various visualization processes. However, visualization researchers and practitioners often encounter phenomena that they do not know what existing theory can explain. For example, given a bipartite graph representing the relations between two sets of vertices, it can be displayed as a node-link graph, as shown in the upper part of Figure 1, or as a matrix, as shown in the lower-right part of the figure. In some applications, node-link graphs are found to be better, but in other applications, matrices are better. It would be highly desirable to identify all theories that can explain such phenomena and thus guide visualization practitioners in making appropriate design choices.

It will also be interesting to identify those theories that would lead to conclusions that are inconsistent with practical experience or findings obtained from empirical studies. In this way, visualization case studies can be used to validate or falsify existing theories, helping further theoretical advancement.

Figure 1 In visualization applications, bipartite graphs are sometimes visualized as node-link graphs, and sometimes as matrices. In the example shown here [35], the two sets of vertices represent words in two passages, and each edge indicates that the two corresponding words (one in each passage) are similar. In this case, the matrix representation proved more effective at revealing similarities between two passages.

Consider a list of N existing theories in psychology and visualization, and M visualization case studies, we can organize them into a matrix with N columns and M rows. We can examine each case study as a row, and examine if those N theories can explain it one-by-one across the row. Collectively, the matrix allows us to align proposed theories with practical case studies.

3.1.5 Key Challenges & Open Questions

C1: What constitutes a VisPsych phenomenon? A foundational challenge concerns the unit of explanation. Is Visualization Psychology primarily concerned with perceptual discrimination, cognitive insight, decision outcomes, interaction episodes, or distributed analytic activity across people and artifacts? Because visualization operates across a vast landscape spanning perception, cognition, and social interaction, singling out "units of explanation" might be insufficient. Instead, the field must establish conceptual bridges that facilitate transitions and translations among these focal areas of theorization and their respective lenses. Only through such an interphenomenal and intertheoretical discourse can we build a cumulative framework capable of supporting the full breadth of visualization practice without flattening its inherent complexity.

C2: Do we need distinct VisPsych theories? Visualization research already draws heavily on established psychological theories. However, visualization systems introduce interactive, external, and distributed representations that may not be fully captured by existing frameworks. Should VisPsych synthesize and adapt established theories, or develop genuinely interdisciplinary theories specific to visualization contexts? This question directly relates to the field's intellectual identity and its potential contribution to psychology.

C3: How do we organize knowledge? The rapid growth of VisPsych research demands a structured organization. We propose three complementary efforts:

1. A research problem – inspired framework linking practical design questions (e.g., encoding, highlighting, uncertainty) to underlying psychological mechanisms.
2. A theories-inspired taxonomy organizing visualization methods, systems, and evaluations through high-level theoretical lenses.
3. A conceptual alignment matrix mapping psychological constructs (e.g., attention, working memory, risk perception) to visualization constructs (e.g., salience, layout, threshold design).

Together, these structures are designed to support the synthesis, comparison, and identification of theoretical gaps.

C4: Where are the limits of psychological explanation? Psychological theories explain individual perception and cognition, but visualization operates within broader social, cultural, and historical contexts. When do interpretive practices, disciplinary conventions, institutional norms, or collective meaning-making require engagement with sociological or philosophical perspectives [36, 4, 37]? Identifying these limits is crucial for avoiding reductionism while maintaining explanatory rigor.

C5: How do we formalize without oversimplifying? Formal models enable clarity, prediction, and systematic comparison [2]. Yet visualization practice is often contextual, interpretive, and design-driven. How can VisPsych develop formal or semi-formal frameworks that capture essential mechanisms without stripping away contextual richness? Balancing formal rigor with ecological validity remains a central methodological and theoretical tension.

3.1.6 Roadmap & Next Steps

1 Year. Develop a proof-of-concept VisPsych alignment matrix as a shared, web-based resource that maps psychological constructs to visualization problems and design choices. Conduct a structured expert elicitation study with Dagstuhl participants and broader VisPsych researchers to identify shared assumptions, tensions, and perceived blind spots. Consolidate findings into a short position or workshop paper (e.g., BELIV or CG&A Viewpoints) that articulates the scope, gaps, and initial structure of the field.

3 Years. Extend the alignment matrix into a richer, community-informed knowledge infrastructure that integrates literature analysis, taxonomic structures, and example use cases. Develop a systematic account of VisPsych as a research program, distinguishing core commitments from auxiliary assumptions, and identifying recurring fault lines. Produce shared terminology, operational definitions, and teaching materials to support training at the intersection of visualization and psychology, helping to stabilize the field's conceptual foundations.

5 Years. Articulate and evaluate candidate VisPsych theories that integrate perceptual, cognitive, affective, and distributed perspectives into coherent explanatory frameworks. Establish theory-driven evaluation benchmarks, shared datasets, and comparative study protocols to enable cumulative empirical progress. Foster sustained interdisciplinary collaboration across visualization, psychology, HCI, and related fields through workshops, joint doctoral training, and coordinated publication venues, positioning VisPsych as a mature and self-reflective research program.

3.1.7 Conclusions

Psychological theory already permeates visualization research, shaping encoding choices, evaluation paradigms, and interpretive claims, yet it remains unevenly integrated and rarely organized into a coherent structure or intertheoretical discourse. The diversity of perceptual, cognitive, affective, and distributed perspectives reflects both the richness and fragmentation of the emerging field. By developing structured taxonomies, research problem – inspired frameworks, and conceptual alignment matrices, Visualization Psychology can make its theoretical commitments explicit, clarify what existing theories explain (and where they fall short), and support cumulative, theory-driven knowledge building. Through conceptual consolidation and interdisciplinary collaboration, VisPsych can evolve from an implicit intersection of fields into a mature and self-reflective research program. While many empirical studies on visualization are related to findings in psychology [38], it is highly desirable to explain empirical findings about visualization phenomena using existing theories developed in psychology and, indeed, in visualization. Meanwhile, empirical studies and practical experience in visualization can be used to test existing theories, thus stimulating new theoretical advances in both visualization and psychology.

3.1.8 Acknowledgments

This working group report originated at Dagstuhl Seminar 25511, "From Psychology Enabled Visualization to Visualization Enabled Psychology" (December 14–19, 2025). We thank the seminar organizers – Rita Borgo, Min Chen, David S. Ebert, Brian D. Fisher, and Alvitta Ottley – for convening the event, and all seminar participants for their thoughtful discussions and contributions. We also thank Professor Brian Rogers, University of Oxford, who participated in some discussions in the working group.

References

1 Amy Rae Fox, James D. Hollan. Visualization Psychology: Foundations for an Interdisciplinary Research Program. In: Visualization Psychology (Danielle Albers Szafir, Rita Borgo, Min Chen, David J. Edwards, Brian Fisher, Laura M. K. Padilla, Daniel A. Szafir, ed.), pp. 217–242. Springer, Cham, Switzerland, 2023. doi: 10.1007/978-3-031-34738-2_9.
2 Robert Spence. Information Visualization: Design for Interaction. 2nd ed. Pearson Prentice Hall, Harlow, UK, 2007.
3 Edwin Hutchins. Cognition in the Wild. MIT Press, Cambridge, MA, 1995.

4 Helen E. Longino. Science as Social Knowledge: Values and Objectivity in Scientific Inquiry. Princeton University Press, Princeton, NJ, 1990.

5 William S. Cleveland, Robert McGill. Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods. Journal of the American Statistical Association, 79(387), pp. 531–554, 1984. doi: 10.1080/01621459.1984.10478080.

6 Alan D. Baddeley, Graham Hitch. Working Memory. In: The Psychology of Learning and Motivation: Advances in Research and Theory (Gordon H. Bower, ed.), pp. 47–89. Academic Press, New York, NY, 1974. doi: 10.1016/S0079-7421(08)60452-1.

7 Anne M. Treisman, Garry Gelade. A Feature-Integration Theory of Attention. Cognitive Psychology, 12(1), pp. 97–136, 1980. doi: 10.1016/0010-0285(80)90005-5.

8 Amos Tversky, Daniel Kahneman. Judgment under Uncertainty: Heuristics and Biases. Science, 185(4157), pp. 1124–1131, 1974. doi: 10.1126/science.185.4157.1124.

9 Liu, Zhicheng, Nersessian, Nancy, Stasko, John. Distributed cognition as a theoretical framework for information visualization. IEEE Transactions on Visualization & Computer Graphics, 14(06), pp. 1173–1180, 2008.

10 Donald A. Norman. Things That Make Us Smart: Defending Human Attributes in the Age of the Machine. Addison-Wesley, Reading, MA, 1993.

11 Danielle Albers Szafir, Rita Borgo, Min Chen, David J. Edwards, Brian Fisher, Laura M. K. Padilla, Daniel A. Szafir (eds.). Visualization Psychology. Springer, Cham, Switzerland, 2023. doi: 10.1007/978-3-031-34738-2. ISBN: 978-3-031-34738-2.

12 John Sweller. Cognitive Load During Problem Solving: Effects on Learning. Cognitive Science, 12(2), pp. 257–285, 1988. doi: 10.1207/s15516709cog1202_4.

13 Allan Paivio. Mental Representations: A Dual Coding Approach. Oxford University Press, Oxford, UK, 1986.

14 James J. Gibson. The Ecological Approach to Visual Perception. Houghton Mifflin, Boston, MA, 1979.

15 Richard E. Petty, John T. Cacioppo. Communication and Persuasion: Central and Peripheral Routes to Attitude Change. Springer, New York, NY, 1986. doi: 10.1007/978-1-4612-4964-1.

16 Weber, Wibke. Towards a semiotics of data visualization–an inventory of graphic resources. In: 2019 23rd international conference information visualisation (IV), pp. 323–328. IEEE, 2019.

17 Liu, Zhicheng, Stasko, John. Mental models, visual reasoning and interaction in information visualization: A top-down perspective. IEEE transactions on visualization and computer graphics, 16(6), pp. 999–1008, 2010.

18 Max Wertheimer. Laws of Organization in Perceptual Forms. In: A Source Book of Gestalt Psychology (W. D. Ellis, ed.). Routledge & Kegan Paul, London, UK, 1938. Original work published 1923.

19 Taveras-Cruz, Yesenia, He, Jingyi, Eskew Jr, Rhea T. Visual psychophysics: Luminance and color. Progress in brain research, 273(1), pp. 231–256, 2022.

20 David M. Green, John A. Swets. Signal Detection Theory and Psychophysics. John Wiley & Sons, New York, NY, 1966.

21 Daniel J. Simons, Ronald A. Rensink. Change Blindness: Past, Present, and Future. Trends in Cognitive Sciences, 9(1), pp. 16–20, 2005. doi: 10.1016/j.tics.2004.11.006.

22 Alberto J. Cañas, Joseph D. Novak, Fermin M. González (eds.). Concept Maps: Theory, Methodology, Technology. University of Navarre, Navarre, Spain, 2004.

23 Gavriel Saloman (eds.). Distributed Cognitions: Psychological and Educational Considerations (Learning in Doing: Social, Cognitive, and Computational Perspectives). Cambridge University Press, New York, NY, 1993.

24 Paivio, A. Mind And Its Evolution; A Dual Coding Theoretical Interpretation. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 2006.

25 Jolicoeur, P. Identification of disoriented objects: A dual-systems theory. In: Understanding Vision: An Interdisciplinary Perspective (G. W. Humphreys, ed.), pp. 180-198. Blackwell Publishing, 1992.

26 Amic G. Ho, Kin Wai Michael G. Siu. Emotion Design, Emotional Design, Emotionalize Design: A Review on Their Relationships from a New Perspective. The Design Journal, 15, pp. 9-32, 2012.

27 Mather, George. Foundations of Perception. Psychology Press, 2006.

28 Simon, Herbert A. The information-processing theory of mind. American Psychologist, 50(7), pp. 507-508, 1995.

29 Nersessian, Nancy J. In the Theoretician's Laboratory: Thought Experimenting as Mental Modeling. In: Proceedings of the Biennial Meeting of the Philosophy of Science Association, pp. 291 – 301, 1992.

30 Staggers, Nancy, Norcio, A.F. Mental models: concepts for human-computer interaction research. International Journal of Man-Machine Studies, 38(4), pp. 587-605, 1993.

31 G. Kindlmann, C. Scheidegger. An algebraic process for visualization design. IEEE Transactions on Visualization and Computer Graphics, 20(12), pp. 2181 – 2190, 2014.

32 Edward R. Tufte. The Visual Display of Quantitative Information. Graphics Press, Cheshire Connecticut, 2001.

33 Min Chen, Amos Golan. What may visualization processes optimize?. IEEE Transactions on Visualization and Computer Graphics, 22(12), pp. 2619-2632, 2016.

34 Mukherjee, K., Yin, B., Sherman, B. E., Lessard, L., Schloss, K. B. Context matters: A theory of semantic discriminability for perceptual encoding systems. IEEE Transactions on Visualization and Computer Graphics, 28(1), pp. 697-706, 2022.

35 A. Abdul-Rahman, G. Roe, M. Olsen, C. Gladstone, R. Whaling, N. Cronk, R. Morrissey, M. Chen. Constructive visual analytics for text similarity detection. Computer Graphics Forum, 36(1), pp. 237 – 248, 2017.

36 Engebretsen, Martin, Kennedy, Helen. Data visualization in society. Amsterdam university press, 2020.

37 Meyer, Miriah, Dykes, Jason. Criteria for rigor in visualization design study. IEEE transactions on visualization and computer graphics, 26(1), pp. 87–97, 2019.

38 Alfie Abdul-Rahman, Rita Borgo, Min Chen, Darren J. Edwards, Brian Fisher. Juxtaposing Controlled Empirical Studies in Visualization with Topic Developments in Psychology. arXiv:1909.03786, 2019.

4 Learning and Friction in Visualization

4.1 Learning and Friction in Visualization

Katerina Batziakoudi (INRIA Saclay – Orsay, FR)
Anne-Flore Cabouat (INRIA Saclay – Île-de-France, FR)
Mennatallah El-Assady (ETH Zürich, CH)
Sara Irina Fabrikant (Universität Zürich, CH)
Lane T Harrison (Worcester Polytechnic Institute, US)
Kuno Kurzhals (Universität Stuttgart, DE)
Carolina Nobre (University of Toronto, CA)
Jan Rummel (Universität Heidelberg, DE)
Filip Sadlo (Universität Heidelberg, DE)
Antonia Schlieder (Universität Heidelberg, DE)
Karen Schloss (University of Wisconsin – Madison, US)
Michael Wybrow (Monash University – Clayton, AU)

This white paper synthesizes discussions from the Learning and Friction working group at the Dagstuhl Seminar on Visualization Psychology. We examine how friction – conceptualized as intentional difficulty or processing cost – can both support and hinder learning in visualization contexts. Drawing on learning sciences, particularly the concept of desirable difficulties, and frameworks from game design and HCI, we propose a Person-Environment-Behavior system for understanding when and where friction arises in visualization activities. We identify key challenges, including measuring learning processes non-intrusively, balancing information access with cognitive engagement, and designing adaptive systems that leverage friction productively. Our framework and research agenda aim to bridge visualization research and learning sciences, offering new perspectives on how visualizations can be designed to promote deeper understanding and durable knowledge construction.

4.1.1 Introduction & Motivation

Whether people create, explore, or interpret visualizations, they engage in processes through which knowledge (including declarative knowledge, skills, mental models, beliefs) is constructed, transformed, or stabilized. In this sense, visualizing data almost invariably involves learning processes; however, learning is often implicit, underspecified, or reduced to performance in data visualization studies, and visualization researchers frequently operationalize success in terms of task accuracy, efficiency, or decision quality. Conversely, the learning sciences distinguish among different learning processes, outcomes, and time scales and emphasize the role of effort, struggle, and reflection in constructing deep understanding or durable knowledge.

Our group brought together visualization researchers and learning scientists to examine "friction"; as a useful but ambiguous concept derived from the notion of desirable difficulty in learning sciences, and explore when friction in visualization activities can support learning and when it hinders productive engagement. Here we synthesize emerging perspectives from the seminar, outline a shared conceptual framing for discussing friction in visualization, and identify open questions and opportunities for future research.

4.1.2 Seminar Process & Methodology

Our working group employed a structured brainstorming and collaborative synthesis approach throughout the seminar. We began by identifying key research problems related to friction and learning in visualization and proposing an initial conceptual framing, which we describe later in Section 4.1.3.1. We then generated concrete scenarios in which friction might arise or be intentionally introduced in visualization activities, spanning different data domains, analytical tasks, and user contexts. These examples were organized through thematic mapping, which helped identify recurring categories of friction. During this process, we also developed shared vocabulary to situate friction within visualization reasoning processes and attempt to bridge learning science and visualization research. We also refined and expanded our conceptual framework by identifying factors associated with each component and reconsidering behavior as a dynamic process linking person and environment. In later discussions, we tested the framework by mapping example scenarios to these components and levels, and we explored how related concepts appear in other disciplines such as learning sciences, HCI, object design, and game design. The final stage of the seminar focused on consolidating these discussions into the present white paper and outlining future outputs, including a position paper and a subsequent theory paper developing the proposed framework in greater depth.Throughout the process, we iteratively refined our understanding of friction, its manifestations in visualization contexts, and its relationship to learning outcomes.

4.1.3 Scope & Definitions

We define learning as the structural reorganization of cognition, or the acquisition of knowledge (mental structures of information, including skills or procedural knowledge, declarative knowledge, episodic memory, attitudes, and beliefs). In this context, friction refers to the difficulty that increases processing and memory costs. While our framework has implications for educational contexts, our primary focus is on learning as it occurs during the use, creation, and interpretation of visualization, rather than on curriculum design or pedagogical interventions. As such, we also include incidental learning experiences in our scope.

4.1.3.1 Relevant Theories & Empirical Foundations

This section summarizes the main psychological theories and empirical findings relevant to understanding friction and learning in visualization, highlighting alignments, tensions, and gaps between learning sciences and visualization research.

4.1.3.2 Psychology of Learning: Desirable Difficulties

The concept of desirable difficulties, developed by Robert Bjork and Elizabeth Bjork, challenges the intuitive assumption that learning conditions that produce the best immediate performance also produce the best long-term retention and transfer. Their “new theory of disuse” [3, 5] distinguishes between two forms of memory strength: retrieval strength (the current ease of accessing information) and storage strength (the depth or durability of the memory trace). Critically, these two forms of strength are independent – conditions that rapidly increase retrieval strength often fail to build storage strength, and vice versa.

This creates a metacognitive trap: learners naturally prefer learning conditions that feel easy and produce immediate fluency, mistaking high retrieval strength for durable learning. However, the Bjorks argue that introducing certain difficulties during learning – by making retrieval more effortful – enhances storage strength and leads to superior long-term retention and transfer. Examples of desirable difficulties include: varying the conditions of practice rather than keeping them constant; interleaving different types of problems or topics rather than blocking them by type; spacing study sessions over time rather than massing them together; and using tests as learning events rather than relying solely on re-reading or re-presentation of material.

The key insight for visualization is that optimizing for immediate task performance (high retrieval strength) may inadvertently undermine deeper learning (storage strength). A visualization that makes information retrieval effortless may not support the kind of active processing that builds durable knowledge structures or transferable skills.

4.1.3.3 Friction in Interface Design

Prior research has also explored the concept of friction, both in terms of how it can shape how people experience user interfaces [8], and towards quality of games and particular game mechanics [9, 10]. Hicks et al. [9] discuss friction in the concept of game "juiciness"; a conceptual framework covering qualities like feedback (i.e. "Is feedback giving in response to game critical events or is feedback received on minor player actions that require no further action?") and game state (i.e. "Are feedback elements that highlight information in harmony with other systems?"). Crucially, these dimensions are presented as considerations rather than explicit requirements or guidelines because, depending on the game creator's goal, game elements may be intentionally designed to be less efficient to achieve a particular effect on player experience. Cox et al. [7] coined the term "microboundaries" to refer to design elements that promote beneficial friction, for example, disrupting mindless interaction patterns. As one example, they describe data entry tasks, where fully frictionless interfaces may actually lead to uncaught errors. Again, a key distinction is the proper design of friction. Cox et al. [7] describe "micro" moments of friction as balancing disruption with beneficial reflection. Modern smartphones, for example, can interrupt social media scrolling sessions with a time-use reminder to enable users to reflect – even briefly – on how they are using their time. They further connect friction to behavioral economics concepts such as System 1 and System 2 thinking (e.g., treating friction as facilitating a transition to System 2 thinking). Visualization presents many unique opportunities to build on conceptualizations and manifestations of friction across domains. We draw on these, in part, in developing our proposed framework and research opportunities. While visualization research has extensively utilized performance metrics like accuracy and speed for evaluating visualizations, and there has been some investigation of visualization memorability [6], there has been limited explicit engagement with learning outcomes or the productive role of difficulty. Recent work on guidance systems in visual analytics (e.g., interaction provenance, recommendation systems) touches on related themes but rarely frames these through a learning lens. This represents a significant gap, as visualization activities inherently involve knowledge construction – from developing literacy with visual encodings to building domain-specific insights through exploratory analysis.

4.1.4 Emerging Framework: Friction in a Person-Environment-Behavior System

The Person-Environment-Behavior system allowed us to discuss friction in visualization without having developed a perfect shared vocabulary.

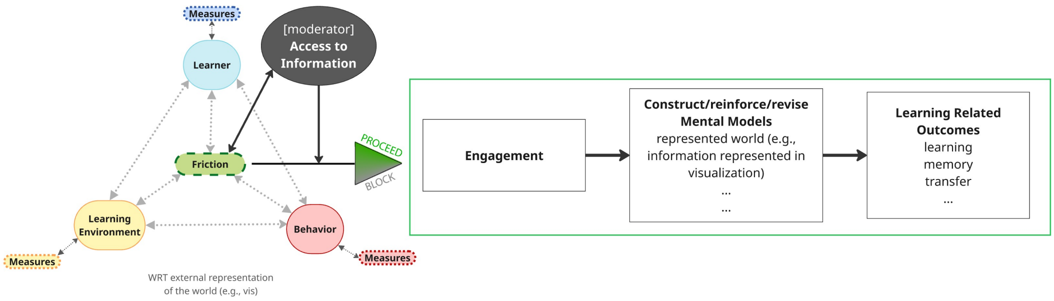

Figure 2 An emerging framework to describe friction as an emergent property of a triadic system, and possible outcomes of friction relevant to data visualization research.

4.1.4.1 Triadic Reciprocal Causation Triangle: Person, Environment, Behaviour

We conceptualize friction as emerging from the dynamic interplay between three core components: the Person (P), the Environment (E), and Behavior (B). This triadic model, inspired by Bandura's reciprocal determinism [1], allows us to examine friction as a systemic property rather than as something located solely in the visualization design or in the user. Table 1 summarizes factors within each component that we think can contribute to friction:

Table 1 Factors contributing to friction in the Person–Environment–Behavior system.

Person (P)	Behaviour (B)	Environment (E)
Visual acuity	Learning activity	Stimuli
Working memory	Cognitive effect	Distance
Language processing	Time	Lighting
Literacy	Movement	Font size/type
Situational state of a person	Frustration	
Tiredness	Performance	
	Memory	

4.1.5 Expanding the Triangle

When trying to expand the triangle by considering the behaviour as the edge between the person and the environment, we realized that behaviour is not an entity like the person and the environment. This realization led us to reconceptualize behavior not as a static node but as a dynamic process that mediates between person and environment (see figure 2). Behavior in visualization contexts encompasses the active processes of visual search, encoding interpretation, interaction, and reasoning – all of which unfold over time and are influenced by both characteristics of a person and environmental affordances.

4.1.5.1 A Procedural View of Behavior: Where and When Does Friction Live?

To better explore "when and where friction lives" in visualization, we included a time dimension and five levels of the visual analysis process, each with guidance: task-specific guidance, data exploration, analysis interactions, visual presentation, and system functionality.

These levels represent different points in the visualization reasoning process where friction can be introduced or encountered. By decomposing behavior across these levels, we can ask more precise questions about when friction supports learning versus when it becomes an unproductive obstacle. For instance, friction at the visual presentation level (e.g., requiring manual decoding of a complex encoding) might support schema development for novice users, while friction at the system functionality level (e.g., cumbersome interaction mechanics) may simply impede analysis without clear learning benefits.

4.1.6 Key Challenges & Open Questions

Through our discussions, we identified several central challenges that require interdisciplinary collaboration between visualization researchers, learning scientists, and HCI researchers. Each challenge is difficult precisely because it lies at the intersection of multiple research traditions, each with distinct methodological approaches and theoretical frameworks.

4.1.6.1 Where and When Does Friction Arise?

Friction is an emergent property of a complex system; therefore, its desirability is highly context-, person-, and task-dependent. Understanding the conditions under which friction emerges requires mapping the interactions between person characteristics (expertise, cognitive capacity, affective state), environmental factors (visualization design, interaction affordances, physical context), and behavioral processes (task goals, analytical strategies, time constraints). We need systematic frameworks for predicting when a given design choice will create productive versus unproductive friction for different user populations and use contexts.

4.1.6.2 How can we balance the tradeoff between desirable difficulty/cognitive friction and easy access to information?

In visualization research, there is a tradeoff between optimizing for as-easy-as-possible information retrieval/decision-making and triggering reflection and learning through introducing some form of friction. This raises the question of when friction is desirable and where it should occur in the visual reasoning process. While learning sciences show that certain difficulties enhance long-term retention and transfer, evaluation of visualization typically prizes efficiency and accuracy. Promoting learning is an implicit design goal of visualizations and systems for visual data analysis. However, our understanding of how learning with visualization works is still limited. We cannot design visualizations that support learning without a deeper understanding of the perceptual and cognitive processes that enable it.

4.1.6.3 How Can Guidance Employ Friction to Promote Learning?

One way to look at friction is to add cost to an easy-to-understand reasoning process. In the context of guidance for visual analytics, introducing friction into user support would only make a complex reasoning process slightly simpler. Adaptive systems represent a promising but complex approach: guidance as a form of learning-promoting friction. Such systems could dynamically adjust the level of support or difficulty based on user characteristics and behavior, potentially optimizing the learning experience. However, this raises questions about how to detect users' current knowledge states and needs in real time, which types of friction are appropriate for different users and contexts, and how to avoid creating learned helplessness or overreliance on system guidance.

4.1.6.4 Measuring Learning in Visualization

How do we measure learning activity in non-intrusive ways? From a learning science perspective: Online measures = performance, Offline measures = learning (but only if you pre-tested).

Measuring learning processes during reasoning with visualizations offers an opportunity to expand the field of visualization research by incorporating well-established measures from the learning sciences that go beyond task accuracy and completion times (e.g., metacognitive assessments). At the same time, data collected during visual data analysis can inform new measures specific to learning with visualization (e.g., provenance data, eye-tracking data). Current visualization studies predominantly measure immediate performance outcomes, which may not reflect learning. We need to develop measures that capture both learning processes and learning outcomes in ecologically valid settings, while distinguishing between different types of learning relevant to visualization: perceptual learning, conceptual understanding, procedural skill development, and metacognitive awareness.

4.1.6.5 The Role of Adaptive Systems and Metacognitive Awareness

What role do adaptive systems play in managing difficulty and metacognitive awareness?

Beyond adjusting friction levels, adaptive systems must support metacognitive awareness so users can recognize and reflect on their own learning. This involves helping users develop accurate judgments about what they know and don't know, when to seek additional information or support, and how to monitor their own understanding. The challenge is designing systems that scaffold metacognition without undermining the development of independent self-regulation skills.

4.1.7 Roadmap & Next Steps

We are currently developing a paper that addresses the discourse on friction in visualization by integrating our group discussions, prior research, and perspectives from academic and industry practitioners collected through interviews. Building on insights from this work, we plan to write a theory paper that formalizes the Person-Environment-Behavior framework for friction in visualization by systematically reviewing manifestations of friction across disciplines, including the learning sciences, game design, HCI, and design, to develop a unified theoretical model. Once the groundwork is laid to have a bigger discussion on friction in visualization, we plan to propose a workshop at EuroVis 2027 that invites visualization researchers, learning scientists, and HCI practitioners to collaboratively explore how friction can be intentionally designed into visualization systems to support learning, drawing on the frameworks and open questions identified throughout the previous work.

References

1 Albert Bandura. The self system in reciprocal determinism. *American Psychologist*, 33(4):344–358, 1978.
2 Elizabeth Ligon Bjork and Robert A. Bjork. Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In *Psychology and the Real World: Essays Illustrating Fundamental Contributions to Society*, pages 56–64. Worth Publishers, New York, NY, US, 2011.
3 Robert A. Bjork. Memory and metamemory considerations in the training of human beings. In *Metacognition: Knowing about Knowing*, pages 185–205. The MIT Press, 1994.

4 Robert A. Bjork and Elizabeth L. Bjork. Optimizing treatment and instruction: Implications of a new theory of disuse. In *Memory and Society: Psychological Perspectives*, pages 116–140. Psychology Press, 2006.
5 Robert A. Bjork and Elizabeth L. Bjork. Desirable difficulties in theory and practice. *Journal of Applied Research in Memory and Cognition*, 9(4):475–479, 2020.
6 Michelle A. Borkin, Azalea A. Vo, Zoya Bylinskii, Phillip Isola, Shashank Sunkavalli, Aude Oliva, and Hanspeter Pfister. What makes a visualization memorable? *IEEE Transactions on Visualization and Computer Graphics*, 19(12), 2013.
7 Anna L Cox, Sandy JJ Gould, Marta E Cecchinato, Ioanna Iacovides, and Ian Renfree. Design frictions for mindful interactions: The case for microboundaries. In *Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems*, pages 1389–1397, 2016.
8 Verena Distler, Gabriele Lenzini, Carine Lallemand, and Vincent Koenig. The framework of security-enhancing friction: How ux can help users behave more securely. In *Proceedings of the New Security Paradigms Workshop 2020*, pages 45–58, 2020.
9 Kieran Hicks, Patrick Dickinson, Jussi Holopainen, and Kathrin Gerling. Good game feel: an empirically grounded framework for juicy design. In *Proceedings of DiGRA 2018 Conference: The Game is the Message*, 2018.
10 Isaac Sung. *Productive and Unproductive Friction in Game Design*. The University of Wisconsin-Madison, 2021.

5 On the Psychology of Decision Making with Visualization

5.1 Dialogues on Decision Making and Visualisation Support in Real-World Settings

Nadine Malcolm (Quantellia – Denver, US, nadine.malcolm@quantellia.com)
B L William Wong (Auckland University of Technology, NZ, william.wong@aut.ac.nz)
Lorien Pratt (Quantellia – Denver, US, lorien.pratt@quantellia.com)
Peter Baudains (King's College London, GB, peter.baudains@kcl.ac.uk)
Brian Fisher (Simon Fraser University – Surrey, CA, bfisher@sfu.ca)
Evanthia Dimara (Utrecht University, NL, evanthia.dimara@gmail.com)[1]

We aim to study real-world decision making in order to inform the design of visualization tools that can meaningfully support it. Decision making has been extensively studied across multiple fields, each offering distinct methodological approaches. On the one hand, decisions unfold as situated activity within concrete contexts characterized by high stakes, constraints, time pressure, coordination, and adaptation. Such settings foreground the role of environment, tacit knowledge, and context, and motivate approaches that study decision processes as they occur in real work situations. On the other hand, each decision also involves individual cognitive processes, heterogeneous expertise, learning, mental representations, and cognitive limitations, which align with approaches that focus on internal reasoning and interpretation. At the same time, while real-world decisions resist purely data-driven or fully structured

[1] Author order chosen by random permutation.

formulations, they must remain data-informed when information and analytics are available, creating a need for interfaces that allow data to be accessed, interpreted, and integrated into ongoing decision processes. Visualization tools occupy this space, yet methodological guidance for designing and studying such tools in real-world decision contexts remains limited. Thus, it remains unclear which perspectives, or which combinations thereof, can adequately capture the processes and constraints that characterize and support real-world decision making in practice. To begin investigating this question, we adopt a dialogic approach based on four perspectives. We begin with practitioners operating in real-world decision contexts, who ground and situate the discussion in concrete decision support practices and constraints. We then draw on perspectives from cognitive engineering and naturalistic decision making research, which provide methods for studying decision making in context and under operational conditions. Next, we incorporate perspectives from psychology research on human decision making, offering complementary and contrasting accounts of cognitive processes and limitations. Finally, we present the perspective of visualization researchers, who reflect on the preceding views and examine how existing visualization paradigms, methods, and assumptions align or conflict with them. Through this dialogue, we surface points of alignment, tension, and mutual blind spots across perspectives, and identify research challenges and agenda items that must be addressed collaboratively, with practitioners positioned as active contributors throughout the research and design process rather than as downstream users.

Following these conclusions, we discuss concrete next steps for sustaining this collaboration through joint paper writing, strategic publication efforts, and coordinated grant acquisition to extend this dialogue into longer-term research partnerships.

5.1.1 Introduction & Motivation

Real-world decision makers increasingly operate in environments where data, analytics, and algorithmic outputs are available, yet are difficult to integrate meaningfully into ongoing decision processes. Decisions in domains such as healthcare, emergency response, operations, management, and policy making are shaped by time pressure, high stakes, evolving goals, and coordination with others. In these settings, decision makers rarely pause to "analyze a dataset" in isolation. Instead, they must interpret partial information, reconcile competing constraints, and act while situations continue to unfold. Despite this reality, many existing decision support and visualization tools remain poorly aligned with how decisions are actually made in practice, offering displays of data without desired outcomes, static analyses, post-hoc explanations, or narrowly-defined optimization views that are difficult to embed into real-world workflows.

This gap is not due to a lack of research on decision making. On the contrary, decision making has been studied extensively across multiple disciplines, each bringing its own assumptions, methods, and units of analysis. Some approaches emphasize decision making as situated activity embedded in concrete contexts, where tools, practices, and organizational constraints play a central role. Others focus on decision making as an individual cognitive process, shaped by uncertainty, experience, learning, and cognitive limitations. Visualization research, in turn, has developed rich methods for externalizing information and supporting sensemaking, yet has often relied on simplified task formulations or predefined data structures. As a result, while each perspective captures important aspects of decision making, it remains unclear how they relate to one another, where they conflict, and how they might jointly inform the design of visualization tools that can be used effectively in real-world decision contexts.

A further complication is that real-world decision making does not fit neatly into a single methodological frame. Such decisions are rarely fully data-driven, yet they are not independent of data; they rely on cues, signals, and analytics when available, but these must be interpreted, weighed, and revisited over time. They are neither purely rational nor arbitrary, but adaptive, experience-based, and sensitive to context. Visualization, as a means of externalizing and interacting with data, sits at the intersection of these characteristics. However, the visualization community currently lacks clear methodological guidance on how to study, design, and evaluate visualization tools that are meant to be embedded in such complex decision processes rather than applied to well-defined analytical tasks. In this paper, we examine real-world decision making by placing four perspectives in direct conversation, each contributing a distinct way of approaching the problem of supporting decision makers and deriving solutions. We begin with the perspective of practitioners, who provide both a general account of common industry practices for decision support and concrete reflections on solutions they have personally designed, articulating how decisions unfold in practice, what currently works, and where limitations and unmet needs remain. We then introduce the perspective of cognitive engineering and naturalistic decision making research, which views decision making as situated activity embedded in socio-technical systems and attends to how decisions unfold over time under constraints, coordination, and environmental conditions, as well as to the role of tools and representations in ongoing work. Next, we introduce the perspective of psychology research on decision making, which examines decisions as individual cognitive processes shaped by perception, judgment, learning, mental representations, expertise, and cognitive limitations under uncertainty and pressure. Finally, we introduce the perspective of visualization research, which reflects on the preceding views to examine how existing visualization paradigms, design methods, and evaluation practices conceptualize decision support, and where their underlying assumptions align or conflict with the characteristics of real-world decision making.

Our dialogic approach is motivated by the observation that real-world decision making sits at the intersection of multiple domains that differ not only in methods, but also in underlying values, assumptions, and criteria of validity. Rather than attempting to reconcile these differences into a unified or ostensibly objective account, the dialogue is designed to make them explicit. By placing perspectives side by side, the approach foregrounds what each domain treats as salient, what it tends to leave implicit, and where tensions arise in how decision making is conceptualized and studied. The resulting accounts are intentionally situated and expertise-informed, reflecting the standpoints of their contributors rather than claiming comprehensiveness or neutrality. This methodological choice allows the paper to begin from non-unified views and to surface authentic points of disagreement, alignment, and blind spots that are difficult to capture through integrative summaries alone. In doing so, the dialogue serves not as a synthesis, but as a structured starting point for an interdisciplinary conversation grounded in practice, cognition, and design.

The paper concludes by discussing concrete next steps for continuing this collaboration, including joint paper writing, publication efforts, and coordinated grant acquisition to extend the dialogue beyond the seminar.

5.1.2 Perspectives from Practitioners on Real-World Decision Making

This section presents a practitioner-grounded view of real-world decision making to contextualize the theoretical perspectives that follow. We begin by situating the practitioners, Lorien Pratt and Nadine Malcolm, and the scope of their exposure (Section 5.1.2.1). We then describe how decisions typically unfold in organizational settings (Section 5.1.2.2) and how decision support is commonly provided in industry, including recurring limitations (Section 5.1.2.3). Next, we introduce the Causal Decision Diagram (CDD) as a representation that has emerged in practice to structure action-to-outcome reasoning (Section 5.1.2.4), and finally reflect on persistent challenges encountered when supporting complex decisions (Section 5.1.2.7). Together, these observations provide a grounded reference point for the cognitive engineering, psychology, and visualization perspectives discussed in the sections which follow.

5.1.2.1 Practitioner Profile and Scope of Exposure

We write from the perspective of practitioners who have spent the last two decades building and deploying decision-support methods and software in organizational settings where decisions are high stakes, distributed across teams, and accountable over time. Our work is anchored in Decision Intelligence (DI), a practitioner-led discipline first articulated in our early field-building documents [83, 84] and later operationalized through Quantellia, a decision intelligence software and advisory company working with organizations of multiple sizes on complex decisions. Quantellia was founded in 2010 to develop decision-centric methods and platforms that make action-to-outcome reasoning explicit, shareable, and monitorable in real work contexts [25]. In contrast to data-first approaches that organize work around datasets, models, or reports, Quantellia's DI work starts from concrete decisions and works outward to the data, analytics, and assumptions needed to support them. We documented these methods for practitioners in [81, 80] and have also contributed to academic discussions that frame DI as a paradigm integrating causal reasoning, human judgment, and computational analytics [82, 104].

Our perspective is shaped by repeated delivery of decision support projects across domains rather than by a single application area. Reported deployments span healthcare, finance, energy, education, agriculture, sports, and manufacturing [25], as well as commercial and marketing decisions such as planning and allocating spend across channels and time periods using interactive decision simulation (ibid.). Across these contexts, the recurring focus is on decisions involving multiple stakeholders, uncertainty, resource constraints, and explicit commitments to action. This cross-domain exposure reflects work on both operational and strategic decisions and informs how we structure decisions, representations, and supporting tools.

Our backgrounds combine academic and industrial AI. One of us previously held a faculty position in computer science at Rutgers University, a large US research university, and published in machine learning, including early work on inductive transfer, later known as transfer learning [80]. The other brings over 25 years of experience delivering enterprise software, data science, machine learning, and DI projects, including telecom decision-support programs, applied ML systems, the establishment of a DI center of excellence for a G7 central bank, and multiple national space agencies [80]. These experiences reflect long-term involvement in both the technical and organizational aspects of how decision support systems are used in practice.

The term Decision Intelligence is today used beyond our own work. Industry analysts and conference programs increasingly use DI language. For example, Gartner, a widely used global technology research and advisory firm whose reports shape enterprise IT and analytics markets, has announced a dedicated Magic Quadrant, its flagship vendor-comparison report used by large organizations for technology selection, for Decision Intelligence Platforms, announced in February 2026 [17]. Gartner defines DI as a discipline focused on explicitly understanding and engineering how decisions are made and how outcomes are evaluated and improved through feedback, and describes Decision Intelligence Platforms as software systems that support, automate, and augment human or machine decision making by integrating data, analytics, knowledge representations, and AI, with support for collaborative decision modeling, execution, and monitoring (ibid.). Gartner's decision to introduce a Magic Quadrant signals its assessment that this market has reached sufficient maturity and commercial relevance for systematic vendor comparison. DI language also appears in optimization and analytics communities, including events labeled as Decision Intelligence Summits [35, 36].

Market research firms likewise treat DI as a distinct and growing segment. Grand View Research, a commercial market intelligence firm producing industry forecasts used in investment and strategy contexts, estimates the global DI market at several billion USD with projected compound annual growth rates around 20 – 25% [33]. Precedence Research, an industry forecasting firm, projects sustained double-digit DI growth into the 2030s [85], while SNS Insider, a commercial market research provider, projects a DI market size on the order of tens of billions of USD by the early 2030s [92]. These estimates vary in methodology, but collectively indicate that DI platforms are being developed and adopted at scale across industries.

For this paper, this background provides context to origins of the decision-support approaches discussed later, including how they are used in organizational decision settings.

5.1.2.2 What Real-World Decision Making Looks Like in Practice

Across engagements, real-world decision making rarely appears as a single, discrete moment of choice. Instead, it unfolds as an extended activity in which organizations commit to actions under uncertainty, monitor their consequences, and revise their understanding as situations evolve.

The decisions encountered in DI practice typically concern interventions in complex systems. These include reallocating resources, launching or delaying initiatives, adjusting policies, or responding to emerging risks. In such settings, outcomes depend on long causal chains, heterogeneous stakeholders, and factors beyond any single actor's control.

Decision makers in these settings often work collaboratively with diverse contributors. While formal authority may reside with an executive, board, or governing body, decisions are shaped collectively by analysts, domain experts, operational leaders, stakeholders affected by the decision, and those responsible for execution. Accountability is correspondingly distributed: some participants are accountable for the integrity of data and models, others for the plausibility of assumptions or operational feasibility, and ultimate decision makers for committing the organization to action and bearing the consequences. As a result, decision making is inherently social and coordinative, requiring shared understanding of what is being decided, why particular actions are justified, and how success or failure will be evaluated.

A recurring characteristic of such decisions is that they unfold over time rather than being resolved at a single point. Initial framing decisions determine which actions are within the domain of authority of the decision maker, and for which outcomes they are responsible. These are subsequently refined through iterative review and as constraints, information, and external conditions change. Even after a decision is nominally "made", substantial effort is devoted to monitoring execution, interpreting early signals, and determining whether observed outcomes are consistent with expectations. In practice, decisions therefore persist as living structures that continue to shape attention, interpretation, and action.

Crucially, practitioners we work with rarely experience these decisions as instances of classification, regression, or prediction. While predictive models, forecasts, and analytics frequently inform decision making, they are not the decision itself. Practitioners consistently distinguish between knowing something about the world and deciding what to do about it.

The latter requires explicit reasoning about interventions: what actions are available, what outcomes are desired, what trade-offs are acceptable, and what risks must be borne. This reasoning serves two complementary functions. First, it articulates the why of a decision, the rationale that explains why a particular action is justified given the organization's goals, values, and constraints. Second, it articulates the how of an outcome, the causal mechanism by which actions are believed to produce measurable effects and ultimately achieve those goals. External representations that make this action-to-outcome logic explicit function as shared mental models, enabling decision makers to explain, defend, and remain accountable for their choices over time.

This action-to-outcome orientation overlaps with traditions such as Decision Analysis, Operations Research, and systems thinking, which also model alternatives, uncertainties, constraints, and dynamic effects. In practice, however, these methods are typically brought in as components within a broader decision effort rather than defining the decision itself. Optimization models, simulations, forecasts, and system models inform parts of a decision, but practitioners still need a coherent structure that connects actions, assumptions, responsibilities, and expected outcomes into a form that can be discussed, challenged, and revised collectively. In this sense, analytical methods often act as inputs to decision processes rather than as the decision process itself.

A concrete example illustrates these dynamics. Consider a commercial organization deciding how to allocate marketing spend across channels and time periods in advance of a major seasonal demand window. Predictive models may estimate customer response to promotions, media mix models may forecast return on investment, and dashboards may summarize historical performance. However, the decision itself concerns which actions to take: how much to spend, when to spend it, with a context including budget limits, brand considerations, operational capacity, and competitive behavior.

Decision makers must justify why a particular allocation is preferred, for example alignment with strategic priorities or risk tolerance, and how it is expected to lead to desired outcomes, for example increased demand translating into revenue without overloading downstream operations. Once the decision is made, it must be monitored. If early indicators deviate from expectations, decision makers need to diagnose whether assumptions were wrong, external conditions changed, or execution failed, and determine whether corrective action is warranted.

In practice, Decision Intelligence approaches structure this activity through a recurring decision lifecycle rather than a single analytic step. As described in the Decision Intelligence Handbook [80], this lifecycle typically involves: (1) eliciting and framing the decision; (2) identifying available actions and desired outcomes; (3) surfacing assumptions and external factors; (4) constructing a causal representation that links actions to outcomes; (5) identifying analytical or technological assets that inform parts of this causal structure; (6) simulating potential consequences of alternative actions; (7) committing to action(s); (8) monitoring outcomes and intermediate signals as those action(s) play out; and (9) diagnosing deviations and revisiting earlier assumptions as needed. These steps are not strictly linear. Rather, they form an iterative process through which decisions are designed, enacted, and learned from over time.

Taken together, these observations suggest that real-world decision making is best understood as a situated, collaborative, and temporally extended activity centered on commitments to action under uncertainty. Representations and visualizations intended to support such decisions must therefore do more than summarize data or model outputs. They must externalize both the rationale (why) and the mechanism (how) of decisions, support shared understanding among stakeholders, and remain usable throughout the decision's lifecycle.

This characterization of practice provides essential context for the remainder of this section, and motivates closer examination of the representations used to support action-to-outcome reasoning in real-world decision making, which we examine through cognitive engineering, psychology, and visualization perspectives in the sections that follow.

5.1.2.3 How Decision Support is Commonly Provided in Industry

In many organizational settings, decision support does not take the form of explicit, formalized decision representations. Instead, decisions are often made through a combination of informal reasoning, experience-based judgment, organizational power dynamics, and narrative justification. Data, if used, is commonly in the form of reports, spreadsheets, or dashboards. Senior leaders often rely on past experience, intuition, or persuasive storytelling to mentally combine these assets and then to justify courses of action, particularly when decisions are time-constrained or politically sensitive. While data and analysis may be invoked to support these narratives, they are sometimes used selectively to justify already-preferred courses of action rather than as part of a systematic process for reasoning about action-to-outcome relationships.

Historically, such informal and experience-driven approaches have been effective in relatively stable environments, where decision makers could rely on heuristics that had been reinforced over long periods. However, practitioners increasingly report that these heuristics have a diminishing "half-life" in contemporary settings characterized by rapid change, global interdependence, and complex organizational dynamics. As systems become more interconnected, and where uncertainty propagates across parts of the system, local intuition becomes less reliable: actions taken in one part of an organization may produce delayed, indirect, or counterintuitive effects elsewhere. This erosion of heuristic validity has been a significant driver behind the growing demand for more structured forms of decision support, and distinguishes "P (pattern)-decisions" (backward facing) from "C (causal)-decisions" (forward facing) [66].

When formal tools are used, they are often centered on data rather than decisions. Many organizational workflows equate "decision support" with the provision of analytics platforms, dashboards, reports, or spreadsheets that summarize historical data or generate forecasts. These tools are typically designed to support sensemaking, helping users understand what has happened or what might happen, rather than to support explicit reasoning about what actions should be taken. As a result, they stop short of representing decisions themselves. The responsibility for translating insights into action is left to the decision maker, who must mentally integrate disparate metrics, projections, and constraints without explicit computational support.

Compounding this challenge is a widespread terminological ambiguity in industry. The word "decision" is often used to refer to data-to-data transformations, such as classifications, regressions, or rule-based determinations, as might be encoded in decision trees or Decision Model and Notation (DMN) artifacts. While such models are valuable for automating routine judgments, they differ fundamentally from the action-to-outcome decisions considered here. This conflation obscures the distinction between predicting or categorizing states of the world and committing to interventions that change those states. As a result, organizations may believe they are "supporting decisions" when they are in fact only supporting inference.

Beyond analytics tools, decision support is frequently distributed across a patchwork of consulting engagements, internal strategy teams, and ad hoc working groups. Consultants may provide recommendations based on analyses conducted off-line, while internal teams assemble slide decks or business cases that frame options in qualitative or financial terms.

Spreadsheets are commonly used to explore scenarios or estimate impacts, but these artifacts are rarely shared as enduring representations of the decision; instead, they function as transient calculation tools embedded within presentations or reports. Once a decision is made, these materials are often discarded or archived, even though the assumptions and causal reasoning they contained remain relevant as outcomes unfold.

Across these approaches, a common pattern emerges: decision support focuses on producing inputs to decisions – data, metrics, forecasts, and narratives – rather than on supporting the decision as a coherent, actionable construct. Few tools provide explicit support for articulating the rationale (why) behind a decision, the mechanism (how) by which actions are expected to lead to outcomes, or for revisiting those assumptions as conditions change. For those studying decision making in controlled or formal settings, this landscape may appear fragmented or under-specified. Yet it represents prevailing norms through which many organizational decisions are currently supported in practice and forms the baseline against which alternative decision-centric approaches must be understood.

5.1.2.4 A Decision-Centric Approach to Decision Support: Causal Decision Diagrams (CDDs)

Any representation of a decision implicitly commits to a theory of how humans reason about action, uncertainty, and outcomes.

In much of the data science, analytics, and visualization literature, decisions are implicitly framed as the downstream activities from analytical processes – such as regression, classification, forecasting, or optimization – rather than as explicit, action-to-outcome constructs in their own right. Within this framing, analytical models take data as input and produce predictions, scores, or recommendations, while the decision itself is treated as an external or downstream step that lies outside the scope of representation. As a result, representational choices are typically driven by data schemas, model architectures, or computational convenience, and the structure of the decision – what actions are available, what outcomes are sought, and how those actions are believed to causally produce those outcomes within some context of external factors – remains implicit. In real-world settings, however, it is precisely this action-to-outcome reasoning that mediates commitment, accountability, and consequence: decisions are not merely inferences from data, but deliberate commitments to intervene in the world under uncertainty.

For this reason, an explicit visual representation of a decision – both as a diagram and also as an interactive exploratory and optimization framework – has proven value but is understudied.

Specifically, over more than a decade of practice, we and others in the Decision Intelligence (DI) community have converged on a particular representational form (the CDD) to make decisions explicit. CDDs externalize decisions as causal structures linking actions, external factors, intermediate states, and outcomes, and are used as a general-purpose visualization formalism across a wide range of domains. This convergence has occurred independently across multiple organizations and software platforms, suggesting that CDDs capture something fundamental about how decision makers reason about intervention and consequence. At the same time, this representational form did not arise from cognitive theory, perceptual research, or visualization science. It emerged pragmatically, shaped by repeated industrial use rather than by explicit grounding in established models of human decision making. Despite its growing adoption within Decision Intelligence (DI) – a rapidly expanding industry discipline projected to reach tens of billions of dollars in economic impact and recognized by Gartner as an emerging discipline with its own Magic Quadrant – the cognitive and epistemic assumptions embedded in this visualization remain largely unexamined.

This gap presents a significant opportunity. Because CDDs function as an abstract, domain-independent decision representation – analogous in scope to the role played by dashboards and reports in business intelligence – insights into their cognitive and perceptual properties have the potential to generalize across a very large class of decision problems. Just as foundational research in visualization and cognition has shaped the practice of business intelligence (most commonly rendered as dashboards) over decades, a deeper understanding of decision-centric representations could materially influence how organizations design, evaluate, and trust decision-support systems at scale.

The practitioner perspective presented in this section begins from this observation. It uses sustained, cross-domain experience with CDDs to motivate a research agenda in which decision representations themselves – not merely their underlying data or algorithms – become objects of cognitive and visual analysis.

A CDD is a low-friction, decision-centric visual representation designed to externalize how a specific decision maker believes their actions will lead to desired outcomes under uncertainty. In practice, CDDs can be constructed rapidly with untrained decision makers because they are grounded in questions that decision makers are already prepared to answer. Across engagements, decision makers can reliably articulate: (1) what outcomes will be used to judge success, (2) what action choices they are currently considering that they believe will influence those outcomes, and (3) what observable aspects of their environment may affect how those actions translate into results. Once these elements are identified, decision makers readily draw causal chains linking actions to outcomes via intermediate and external factors. Empirically, such causal structures are found to be natural for participants to express, suggesting that CDDs closely align with existing mental models used to reason about action and consequence.

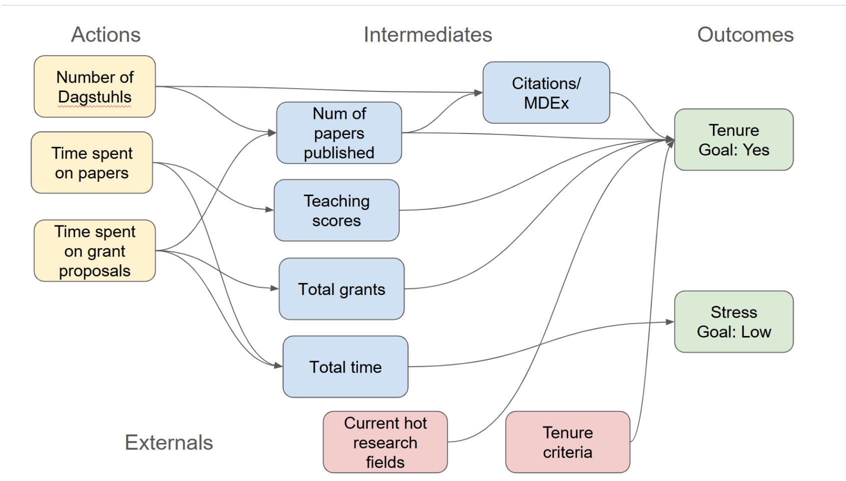

Figure 3 A simple illustrative CDD developed by our team during Dagstuhl, which captures some of the decisions faced by a tenure-track faculty member who has the goal of achieving tenure at their university.

Formally, a CDD consists of a small set of core components connected by directed causal links. These include actions (choices available to the decision maker), outcomes (measures of success), externals (factors outside the decision maker's direct control that influence outcomes), and intermediate variables that capture mechanisms linking actions and externals to outcomes. Importantly, a CDD is always situated within a specific decision context: it is constructed from the perspective of a particular role, making a particular decision, under particular circumstances. This scoping constraint is essential. Rather than attempting to model all roles, all actions, or all data relevant to an organization – as in comprehensive system models – a CDD deliberately focuses only on what matters for the decision at hand. This role- and decision-specific framing makes the representation tractable, actionable, and cognitively manageable.

The structure of the CDD also mirrors organizational accountability. Outcomes correspond directly to the responsibilities for which the decision maker is held accountable, often aligning with formal performance metrics and evaluation criteria. Actions correspond to the authorities granted to that role – what the decision maker is empowered to change or influence. This explicit alignment between outcomes and responsibilities, and between actions and authorities, makes the rationale for decisions transparent and auditable. It also clarifies what success means, who owns it, and which levers they are legitimately expected to pull to achieve it.

Within Decision Intelligence practice, CDDs are embedded in a broader decision lifecycle. As described in the Decision Intelligence Handbook, this lifecycle typically involves nine recurring steps: (1) eliciting and framing the decision, (2) identifying desired outcomes, (3) identifying available actions, (4) surfacing external factors and assumptions, (5) constructing the CDD linking actions to outcomes, (6) populating parts of this structure with analytical or empirical models where appropriate, (7) simulating alternative actions to explore potential consequences, (8) committing to a course of action, and (9) monitoring outcomes and diagnosing deviations to inform future decisions. The CDD serves as the central representational artifact throughout this process, persisting across analysis, execution, and learning. These steps are described in more detail in Section 5.1.2.6 below.

5.1.2.5 Origins of the CDD

The development of the CDD emerged from sustained practitioner exposure to how decisions were actually made in organizations, rather than from a priori theoretical commitments. Prior to the formal articulation of Decision Intelligence, Pratt worked as a market analyst conducting interviews with senior leaders across industries and geographies. Through these engagements, she observed a striking pattern: despite widespread rhetoric about "data-driven decision making," most consequential decisions were made with little use of rigorous methods, limited integration of data, and minimal explicit reasoning about how actions were expected to lead to outcomes. Decisions were often justified post hoc through narrative or authority, rather than through shared, structured reasoning.

In the period leading up to 2008, Pratt undertook a year-long series of interviews initially intended to explore interest in a potential technology tool for pricing and profitability modeling. However, as these conversations progressed, the focus of the interviews shifted. Rather than asking what analytical tools organizations wanted, Pratt began posing a more open-ended question: If technology could solve one problem that it does not adequately address today, what should that problem be? Over time, a consistent theme emerged across respondents. Leaders were not primarily asking for better predictions or more sophisticated analytics; instead, they articulated a need for help reasoning about action-to-outcome decisions – that is, understanding how choices they controlled might plausibly lead to desired results in complex, uncertain environments.

These interview findings were synthesized and formalized in the 2008 paper by Pratt and Zangari [83], which articulated the core problem that Decision Intelligence sought to address. The paper documented two related observations. First, respondents consistently expressed frustration with the gap between analytical outputs and actionable decisions.

Second, they emphasized the need for representations that could be easily shared, discussed, and understood across stakeholders – often invoking familiar artifacts such as slide decks or whiteboard sketches. What they wanted was not an opaque model, but a visual representation of their reasoning that could function as a common reference point in discussion and decision making. Analysis of interview transcripts further revealed that, despite differences in domain and role, respondents' verbal explanations of decisions relied on a small set of recurring elements: actions under consideration, desired outcomes, external factors beyond their control, and causal mechanisms linking these together. These elements formed the conceptual basis of what later became the CDD.

Following the publication of the 2008 paper [83] and the founding of Quantellia in 2010, these ideas were tested and refined through dozens of applied projects across domains including finance, healthcare, agriculture, marketing, and large-scale organizational transformation. Across these engagements, a consistent empirical finding emerged: untrained decision makers were able to construct and reason with CDDs quickly and with minimal instruction. When asked to identify outcomes, actions, and relevant external factors, participants did so readily, and they naturally expressed causal relationships among these elements. The recurrence of this pattern across projects provided practical validation that the representation aligned well with how decision makers already conceptualized their decisions, even if that reasoning had previously remained implicit or informal. Only in retrospect did the deeper theoretical resonance of the CDD become clear. The structure of the diagram – linking actions (behaviors) to outcomes (consequences) through intervening conditions – closely mirrors the antecedent – behavior – consequence framework central to operant conditioning. From this perspective, the apparent naturalness of CDDs is less surprising: they externalize a form of causal reasoning that is deeply embedded in human learning and behavior. What began as a pragmatic response to practitioner needs thus converged with well-established principles from cognitive psychology and behavioral science, suggesting that the CDD's effectiveness may be rooted in fundamental properties of human reasoning about action and consequence.

5.1.2.6 Overview of the Decision Intelligence Process

The Decision Intelligence (DI) process begins not with modeling or analytics, but with a collaborative decision-framing exercise involving the decision customer, decision makers, analysts, and relevant subject-matter experts. Early stages focus on clarifying the decision objective, constraints, and boundaries, and on eliciting candidate actions, desired outcomes, and key external factors. This work typically culminates in the first draft of a CDD – a representation that is intentionally rough, incomplete, and sometimes messy. This initial messiness is not a flaw but a feature: it surfaces assumptions, disagreements, and gaps in understanding that would otherwise remain implicit. Research on spatial and diagrammatic reasoning shows that informal, sketch-like representations can be cognitively productive precisely because they invite interpretation, revision, and collaborative sensemaking (e.g., [98]).

In practice, untrained decision makers readily participate in this early modeling because the questions being asked – What does success look like? What choices are we facing? What external factors matter? – are ones they are already prepared to answer. Recent work also shows that large language models can assist in this phase by helping extract candidate actions, outcomes, and causal relationships from documents or meeting transcripts, accelerating early diagram construction without replacing human judgment (see `https://lorienpratt.com` for examples and discussion).

Following this initial framing, the CDD is refined into a clean, well-formed diagram. During this phase, elements are clarified, causal chains made explicit, terminology standardized, and scope aligned with the role and circumstances of the decision maker. The result is a readable

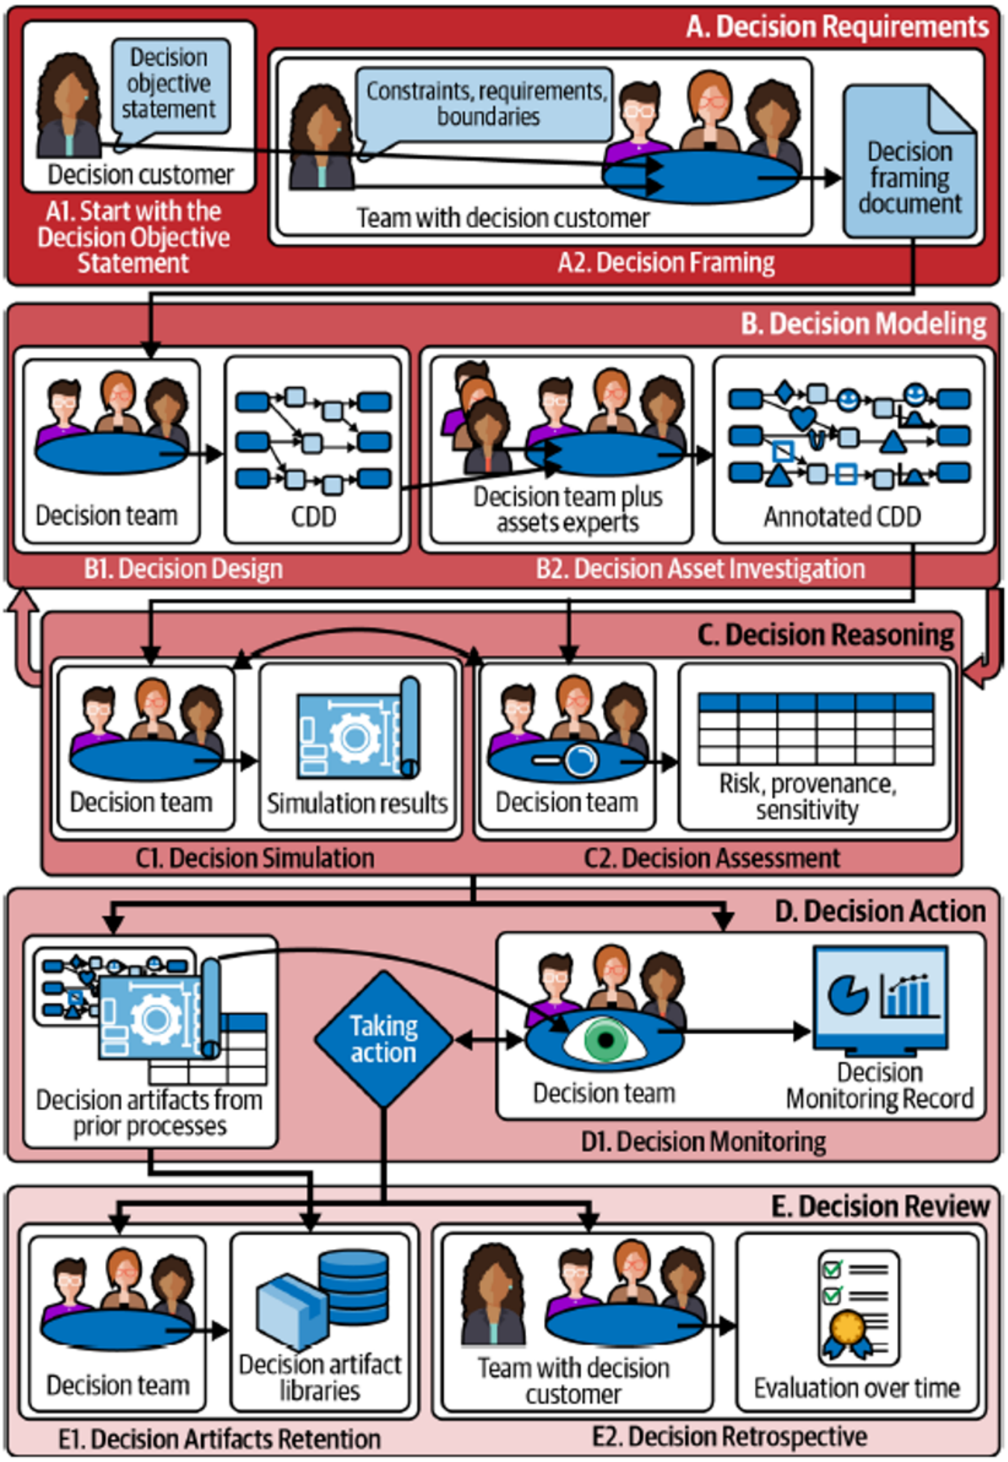

Figure 4 An overview of the DI process.

representation that makes both the rationale for the decision (why particular actions are considered) and the mechanism by which actions are expected to lead to outcomes (how effects propagate through the system) explicit and shareable.

In many projects, this visualization alone delivers substantial value. In an early Quantellia engagement, for example, the finalized CDD was printed, laminated, and distributed across teams; its use as a shared reference artifact was credited with saving several million dollars per year purely through improved alignment and coordination, even in the absence of simulation or optimization. For this reason, many DI efforts intentionally stop at the CDD stage, treating the diagram itself as the primary decision-support deliverable.

The same CDD can also serve as an architectural blueprint for decision simulation. In this mode, analytical assets from multiple disciplines – machine learning models, statistical forecasts, optimization routines, business rules, or system dynamics components – are attached to specific links or nodes in the diagram. The CDD thus functions as an integration framework that coordinates heterogeneous models around a coherent decision structure. The resulting simulator can be run forward from actions to outcomes to enhance intuitive understanding of trade-offs and sensitivities, or in optimization mode to identify action choices that best satisfy decision objectives under stated constraints.

Importantly, the role of the CDD does not end once a decision is made and action taken. The same representation can be reused as a monitoring and governance framework, organizing dashboards and metrics around the causal structure of the decision rather than around disconnected data streams. In this post-decision mode, the CDD supports early warning by highlighting deviations along causal pathways, and it facilitates root-cause analysis by helping teams reason about which assumptions or mechanisms may have failed. This continuity – from framing, through action, to monitoring and review – distinguishes the DI process from approaches that treat decision making as a one-time analytical event rather than an ongoing, learnable activity.

Taken together, these properties distinguish CDDs from dashboards, predictive models, or purely analytical artifacts. Rather than delivering insights that must be mentally translated into action, CDDs externalize both the why of a decision – its rationale and justification – and the how of outcomes – the causal mechanisms by which actions are expected to achieve goals. This makes CDDs a foundational representation for integrating analytics, machine learning, systems thinking, and human judgment within a single, decision-centric scaffold.

5.1.2.7 Real-World Challenges in Supporting Decision makers

Supporting real-world decision making presents a sequence of challenges that arise at different stages of the Decision Intelligence (DI) process. These challenges are best understood by following the lifecycle of a decision – from initial framing and diagram construction, through refinement and analysis, to simulation, action, and post-decision monitoring – rather than treating decision support as a single analytical task.

Challenges in Early CDD Construction and Decision Framing. The first set of challenges arises during initial CDD modeling, when decision makers and stakeholders collaboratively externalize their understanding of a decision. While CDDs are intentionally low-friction and align well with how people naturally reason about action and consequence, early learners often struggle to distinguish a CDD from familiar process models or workflows. Practitioners accustomed to flowcharts or business process diagrams may initially interpret nodes and links as sequences of activities to be executed.

To address this, we emphasize a strict structural distinction: actions appear only on the left of a CDD, representing commitments under the decision maker's control, while all other elements represent consequences – intermediate states, external influences, and outcomes – that explain why those actions are expected to achieve desired results. We situate this distinction within a broader taxonomy of how people reason about the future: (1) planning activities and schedules, (2) predicting events outside one's control, and (3) choosing actions by reasoning about downstream consequences. The CDD and DI approach are explicitly designed for the third mode. While planning and prediction are well studied and well supported by existing tools, action-to-outcome decision reasoning has received comparatively less systematic attention, despite its centrality to accountable decision making.

A related challenge at this stage is scope management. There is a strong tendency, particularly among analytically sophisticated teams, to attempt to model "everything that matters" from the outset, producing very large and complex CDDs early in the process. We have repeatedly observed that this "boil the ocean" approach undermines learning, alignment, and momentum. Effective practice instead begins with small, simple CDDs focused on a specific role, decision, and set of circumstances, and then extends them incrementally as understanding deepens. Starting small allows teams to validate assumptions, build shared understanding, and gain confidence before introducing additional complexity.

At the same time, we have found value in developing foundation CDDs – relatively comprehensive diagrams that capture the decision logic of an entire domain or field (e.g., pricing, supply chain resilience, patient flow). These foundation CDDs are not intended to be used directly by individual decision makers. Rather, they serve as reusable conceptual assets from which smaller, role-specific CDDs can be derived and tailored. Balancing incremental development with the existence of such foundational structures is an ongoing practical challenge.

Challenges in Refinement and Visualization. As a CDD is refined into a clean, well-formed representation, challenges related to visualization preferences and cognitive load become increasingly salient. Decision makers vary widely not only in what information they want to see, but in how they want to engage with it visually and interactively. Some prefer high-level, minimalist views that display only actions and outcomes, using large fonts, limited color, and minimal visual clutter to reduce cognitive burden and support rapid comprehension. Others seek rich, exploratory environments that resemble interactive simulations or games, in which data is presented as a heads-up display layered over an evolving, manipulable landscape. In these settings, motion, interaction, and spatial metaphor are not distractions but integral to how users build intuition and explore possible futures, and the motivation to engage with the visualization and to, ultimately, take the actions it recommends.

These divergent preferences are not merely aesthetic; they reflect different cognitive strategies for understanding complex decisions. Sparse representations can support focus, clarity, and deliberation, particularly for time-poor executives or users making infrequent, high-stakes decisions. Rich, immersive representations can support exploration, learning, and intuition-building, especially when users are comparing alternatives, experimenting with assumptions, or seeking to understand dynamic behavior over time. In practice, both modes can be valuable, often at different stages of the same decision process or for different stakeholders involved in the same decision.

When CDDs are implemented in software, these differences pose a fundamental design challenge. The same underlying decision model must often support multiple visual forms, ranging from static, summary views to highly interactive, exploratory interfaces. Treating any single visualization as canonical risks excluding users whose cognitive styles or tasks are poorly served by that representation. At the same time, attempting to accommodate all preferences within a single visual view can lead to excessive complexity and cognitive overload.

Reconciling these perspectives requires a clear separation between the CDD as an abstract decision model and its concrete visual manifestations. Rather than asking which visualization is "correct," the challenge becomes how to design systems that allow the same decision structure to be experienced through multiple, purpose-specific views, and to support transitions between them as users move from framing and exploration to commitment and monitoring. Understanding how different visual forms support reasoning, intuition, trust, and coordination – and how users shift between them over time – remains an open research problem at the intersection of visualization, cognitive psychology, and human – computer interaction.

Challenges Introduced by Simulation and Optimization. A further set of challenges emerges once data and analytical models are applied to the CDD and simulation is introduced. Simulation typically reveals not a single best action, but sets of action choices that perform well under particular assumptions about the external context. For decisions involving many actions and externals, this results in high-dimensional trade spaces rather than clear optima. Presenting these multidimensional results in a way that decision makers can grasp quickly – often under time pressure and without extensive training – remains difficult. Traditional reports, dashboards, or static charts rarely convey how trade-offs shift across contexts or assumptions.

Closely related is the challenge of intuition-building. Historically, many decision makers developed reliable intuition through long experience in relatively stable environments. In today's rapidly changing world – characterized by globalization, climate change, pandemics, regulatory shifts, and technological disruption – those intuitions can fail. Although large datasets and models provide timely signals of change, people struggle to integrate vast amounts of information into coherent mental models. Some respond by ignoring data and relying on outdated intuition; others experience tension when analytically supported choices do not feel right. Effective decision visualizations must therefore support exploration of possible futures in ways that help align cognitive understanding with intuition, so that evidence-based decisions can also be experienced as sensible and trustworthy.

Challenges After Action: Monitoring, Learning and Automatation Expectations. Finally, challenges persist after decisions are made and actions taken. Some decision makers initially assume that CDD-based systems will fully automate decisions, effectively replacing human judgment. In practice, most decision makers want to remain accountable and to exercise judgment, particularly in high-stakes or novel situations. DI systems are therefore designed to support, augment, and sometimes partially automate decision making, but not to eliminate human responsibility. Making the role of humans-in-the-loop explicit is essential for trust and appropriate use.

At the same time, using the CDD as a post-decision monitoring framework introduces its own challenges: organizing dashboards around causal structure rather than raw metrics, supporting early warning and root-cause analysis, and maintaining continuity between the original decision rationale and observed outcomes over time. While practitioners find this continuity valuable, its cognitive and organizational implications remain underexplored.

Ecosystem and Interoperability Challenges in Decision Intelligence. Decision Intelligence (DI) is developing not only as a set of methods and tools but also as an ecosystem that includes multiple vendors, research groups, and practitioner communities. As DI work expands across organizations and domains, questions of interoperability, shared vocabularies, and model portability become increasingly visible in practice.

More broadly, current AI and analytics landscapes are characterized by concerns about concentration of technical capacity within a small number of large technology firms. Such concentration can shape which tools, representations, and workflows become dominant, and may limit diversity of approaches. Within DI, similar concerns arise around fragmentation: different teams and vendors may use incompatible representations, terminologies, or software environments for modeling decisions.

From a practitioner perspective, this creates concrete difficulties. Decision models that are built for one project or platform are not always easy to transfer, compare, or reuse in another context. Even when diagrams appear similar, differences in semantics or assumptions can make them functionally incompatible. As a result, decision knowledge can remain siloed within specific tools or organizations.

Partly in response to these issues, community initiatives such as OpenDI.org have emerged to explore shared standards and interfaces for decision modeling. OpenDI positions itself as a community-driven effort to define common vocabularies, draft interoperability standards, and reference APIs for DI systems, including specifications for how causal decision models and simulation components might communicate. It also maintains draft role definitions and user stories intended to align expectations among practitioners and tool builders.

The existence of such initiatives highlights a practical tension. On one hand, shared standards may support reuse, collaboration, and continuity of decision models across tools and teams. On the other hand, decision representations are closely tied to how people naturally express reasoning about action and outcomes. Strong standardization could risk constraining that expressiveness or imposing structures that do not align with practitioners' mental models.

For researchers, this raises open questions about how decision representations should balance flexibility and standardization, how decision knowledge can be preserved and transferred over time, and how interoperability efforts influence how decisions are conceptualized and communicated. From our experience, these ecosystem-level issues increasingly affect how decision support is deployed and sustained in real organizations.

5.1.2.8 Where Industry Meets Academia

The challenges described above arise largely from practical experience with CDDs and the Decision Intelligence (DI) process. They concern how decision makers interpret representations, how visualizations are adapted to different users, how intuition is reconciled with evidence, and how human judgment is preserved alongside analytical support. A recurring practical issue is collaborative alignment: establishing and maintaining a shared understanding of a decision among diverse participants over time. This includes not only the core decision-making team but also the broader set of stakeholders who must understand, trust, or act on a decision. In real settings, misalignment among stakeholders often proves as consequential as analytical error, particularly when decisions extend over long time horizons or operate under scrutiny.

A deeper set of challenges follows from how DI methods have developed. CDDs and related practices emerged primarily through industrial use rather than through controlled empirical or theoretical programs. Many design choices that appear effective in practice, such as using shared diagrams to align teams, revisiting decision rationale during monitoring, or adapting representations for different audiences, work operationally but remain weakly explained. Practitioners can observe that these approaches help, but industry environments rarely allow for controlled comparison, longitudinal study, or systematic theory building. Delivery pressures, organizational constraints, and the uniqueness of each decision context limit how rigorously these methods can be examined from within practice alone.

At the same time, the DI landscape itself is distributed. Vendors, practitioners, researchers, and educators often work in parallel, using related ideas but without shared standards, vocabularies, or common experimental substrates. This fragmentation makes it harder to accumulate knowledge about what works, for whom, and under what conditions. Initiatives such as OpenDI (`https://opendi.org/`) reflect practical attempts to reduce this fragmentation by proposing shared conceptual frameworks, interoperability standards, and reference tools. These efforts highlight a broader challenge: decision representations increasingly function as boundary objects across teams, tools, and organizations, yet their cognitive and social properties are not well characterized.

From an industry perspective, this is precisely where academic collaboration becomes valuable. Academia is structurally better positioned to run controlled studies, compare alternative representations, test cognitive assumptions, and develop transferable theory. Industry, in contrast, provides rich, consequential, real-world contexts but limited opportunities for systematic investigation. The combination is complementary: practice exposes phenomena at scale, while academic research can help explain, validate, and generalize them.

Several issues observed in practice illustrate this need. CDDs appear to align naturally with how decision makers reason about action and consequence, yet it is not well understood why certain diagrammatic structures are easier to construct, share, or reason with than others. How stakeholders interpret causal links under uncertainty, how external representations support or distort mental simulation, and how shared decision models influence trust, accountability, and legitimacy across organizational boundaries are questions that exceed what routine industry work can answer. Similarly, DI processes emphasize iteration and post-decision learning, but the cognitive and organizational implications of carrying a decision representation forward over time, rather than treating decisions as one-off events, remain underexplored.

The remainder of this paper therefore shifts from describing practice to examining it through established research traditions. Cognitive engineering and Naturalistic Decision Making offer concepts and methods for studying sensemaking, situation awareness, mental simulation, coordination, and experience-driven reasoning under uncertainty and time pressure. Perspectives from psychology and visualization research further help analyze how representations shape reasoning, communication, and judgment. Engaging these perspectives is not about academic validation, but about building a more reliable foundation for tools and methods already in use. The goal is to better understand when decision-centric representations help, where they may mislead, and how they can be improved for the diverse settings in which real decisions are made.

5.1.3 Perspectives from a Cognitive Engineering Researcher

We approach decision making as a form of situated cognitive work performed by humans interacting with artifacts in complex socio-technical systems. From a cognitive engineering perspective, the central question is not only how people choose among alternatives, but how they perceive, interpret, coordinate, and act in dynamic environments where goals, constraints, and system states evolve over time. Decision making is therefore treated as one component of broader macro-cognitive activity that includes situation awareness, sensemaking, planning, problem detection, and adaptation [52].

A recurring commitment in cognitive engineering is that systems should be designed around the realities of human cognitive work rather than around idealized models of rational choice. This stance emerged from dissatisfaction with laboratory-only accounts of decision making and from sustained study of real-world domains such as aviation, process control, military command, and emergency response [40, 56]. In these settings, decisions are embedded in ongoing activity, shaped by time pressure, uncertainty, and coordination demands, and rarely resemble discrete choice among well-defined options.

From this viewpoint, the practitioner perspective presented earlier resonates strongly with core concerns in cognitive engineering: decisions as temporally extended, action-oriented, and socially distributed. At the same time, cognitive engineering treats such richness not only as context but as an object of systematic analysis. The goal is to understand how cognitive work is structured by the work domain, by available representations, and by the coupling between humans and technology [41]. In this sense, cognitive engineering asks not only what decisions look like in practice, but how systems and representations can be designed so that cognitive work remains effective under complexity and surprise.

5.1.3.1 Conceptual and Methodological Reflections from Cognitive Engineering on Real-World Decision Making

Cognitive engineering draws heavily on the Naturalistic Decision Making (NDM) tradition, which studies how people make decisions in real-world settings characterized by ill-structured problems, shifting goals, uncertainty, feedback loops, time pressure, high stakes, and multiple stakeholders [56, 55]. Rather than treating these factors as noise to be controlled away, NDM treats them as defining features of decision environments.

One influential model within NDM is the Recognition-Primed Decision (RPD) model [53, 54]. RPD describes how experienced decision makers rapidly recognize situations as familiar patterns, retrieve plausible actions, and mentally simulate their consequences. Instead of comparing multiple options, they evaluate one workable option at a time, accepting it if it satisfies constraints. This account emphasizes the role of experience, pattern recognition, and mental simulation, and helps explain why experts often report that they "just know" what to do.

These ideas connect to broader work on expertise and adaptive performance, which shows that expert decision making relies on richly structured experience, domain-specific knowledge, and the ability to adapt under variability [27, 26]. From a cognitive engineering standpoint, effective decision support must therefore be compatible with how expertise is actually exercised, rather than forcing experts into unfamiliar analytic routines.

Methodologically, cognitive engineering has developed a family of approaches for analyzing cognitive work. Cognitive Task Analysis (CTA) focuses on eliciting how experts notice cues, interpret situations, and manage uncertainty, often revealing tacit strategies that are not captured in formal procedures [16]. Work Domain Analysis (WDA), within the broader Cognitive Work Analysis framework, models the structure of a domain in terms of goals, constraints, functional relationships, and physical resources, with the aim of making deep system structure visible and usable [99].

Together, these approaches shift attention from isolated decisions to the ecology of cognitive work. Decisions are seen as emerging from interactions among actors, artifacts, and environments, rather than as purely internal computations. This orientation aligns with the concept of joint cognitive systems, where cognition is distributed across humans and tools and where the unit of analysis is the coupled system rather than the individual alone [41].

Design methods in cognitive engineering follow from this analysis. Ecological Interface Design (EID) aims to make system constraints, functional relationships, and invariants perceptually available so that users can detect problems, diagnose causes, and anticipate consequences [100]. The Semantic Mapping Principle, central to EID, argues that the meaning of information should map as directly as possible onto its visual form, reducing the need for mental translation. The Law of Requisite Variety further emphasizes that a support system must be able to accommodate the range of variability present in its environment if it is to remain effective [4].

From this perspective, decision support is not merely about presenting data, but about shaping the cognitive ecology so that meaningful relationships, constraints, and trade-offs become perceptually and conceptually available for action.

5.1.3.2 Cognitive Engineer's Reflections on the CDD Method

From a cognitive engineering perspective, Causal Decision Diagrams (CDDs) are immediately recognizable as artifacts that structure cognitive work. They externalize relationships among actions, conditions, and outcomes, and thus function as shared representations around which reasoning and coordination can occur. In this sense, they align with the broader goal of making important relationships visible so that they can be inspected and discussed [40].

At the same time, CDDs reflect a particular choice of unit of analysis. They center on a focal decision and its action-to-outcome pathways, whereas cognitive engineering often begins from the structure of the work domain as a whole. A WDA-informed analysis would ask whether the causal links in a CDD adequately reflect domain constraints, system invariants, and means – ends relationships [99]. From this viewpoint, a CDD is a useful but partial slice of a larger constraint space.

CDDs also resonate with NDM and RPD ideas. Their emphasis on action – consequence chains and scenario exploration is compatible with mental simulation, a core mechanism in RPD [54]. As boundary objects, they may support sensemaking and shared understanding in teams, functions that are central to macrocognition [52, 51].

However, cognitive engineering would raise several questions. First, how well do CDDs capture the constraints and invariants of the work domain, rather than only stakeholders' current beliefs? Second, how do they scale when the variety of the environment increases, as highlighted by the Law of Requisite Variety [4]? Third, how do they interact with expertise? Representations that are helpful for novices may be redundant or even constraining for experts, and vice versa [26].

In this reading, CDDs are promising as cognitive artifacts, but their effectiveness depends on how well they are grounded in domain structure and how they fit into existing patterns of cognitive work.

5.1.3.3 Challenges for Cognitive Engineering to Support Real-World Decision Making

From our perspective as cognitive engineering researchers, the practitioner account also exposes important limits in how our field currently studies and supports decision making.

A first challenge concerns generalization. Much of the empirical base of cognitive engineering comes from high-risk domains such as aviation, military operations, and process control. Corporate and strategic decision contexts, as described by practitioners, differ in time scale, accountability structures, and degrees of formalization. While NDM principles likely transfer in part, the boundary conditions of that transfer remain underexplored [55].

A second challenge concerns evaluation. Cognitive engineering evaluation emphasizes macro-cognitive functions such as situation awareness, sensemaking, and problem solving under normal and abnormal conditions [40]. However, many practitioner success criteria include legitimacy, defensibility, and organizational alignment, which are not easily captured by standard performance metrics. Assessing decision support in such contexts may require broader evaluative frameworks that consider social and organizational outcomes.

A third challenge concerns resilience and adaptation. Resilience engineering argues that successful performance often involves flexible adaptation to variability rather than strict adherence to plans [42, 39]. Decision support tools, including CDD-based systems, must therefore be examined not only for how they support planned decisions, but for how they help users adapt when conditions change in unexpected ways.

Finally, there is a translational challenge. Cognitive engineering offers rich theories and methods, but these do not always translate into design practices that are accessible to industry teams working under time and resource constraints. Practitioner-developed methods like CDDs move quickly because they are usable and adaptable. For cognitive engineering to remain relevant, it must find ways to engage with such methods, not only to critique them but to help refine and evaluate them systematically.

In sum, the practitioner perspective does not simply provide another application domain. It challenges cognitive engineering to examine how its concepts apply in strategic, organizational, and decision-centric contexts, and to treat decision representations themselves as primary sites of cognitive work. This challenge also represents an opportunity to extend cognitive engineering beyond its traditional domains and to develop more general accounts of how humans and artifacts jointly support consequential decisions.

5.1.4 Perspectives from a Psychology Researcher

Cognitive Psychologists typically approach decision making as an assembly of cognitive and perceptual mechanisms that enable us to interpret sensory information, allocate attention to important elements of the environments, form beliefs in the form of narrative structures, and translate those beliefs into plans for action. A recurring methodological commitment across much of psychology research is that robust explanations require isolating component processes under controlled conditions, so that stable regularities can be identified and theory can cumulate [77, 31, 67, 73]. From this viewpoint, the practitioner perspective offers a compelling characterization of decision making as temporally extended, socially distributed, and anchored in action-to-outcome reasoning. At the same time, we treat this richness as both a feature and a risk: real-world settings introduce substantial variability in attention, interpretation, and learning, and the fact that groups converge on a shared representation does not guarantee that the representation is cognitively tractable or epistemically reliable [78, 96, 86]. We therefore see cognitive engineering and naturalistic decision making as essential complements rather than alternatives. These traditions focus on decision making as situated activity in real-world contexts, emphasizing uncertainty, coordination, and evolving constraints. Psychological research, by contrast, typically seeks explanatory leverage on underlying cognitive and perceptual mechanisms, such as attentional selectivity, limited working representations, and learning dynamics, which are often easiest to diagnose in simplified paradigms [79, 15, 91]. In this sense, psychology asks not only what decisions look like in practice, but why systematic strengths and failures recur across contexts.

5.1.4.1 Conceptual and Methodological Reflections from Psychology on Real-World Decision Making

We can organize the psychological lens on the practitioner account into several partially overlapping traditions, each with characteristic assumptions about what decision making is and how it should be studied. First, within a classical information-processing tradition, we treat decision making as an internal sequence of representational and computational steps, and we operationalize progress in terms of mechanistic decomposition and controlled experimentation [67, 73, 2]. This tradition resonates with the practitioner emphasis on making structure explicit, but we would ask which cognitive operations are actually supported by external artifacts, and under what complexity limits those supports break down. We also note that perception and action are not mediated by a single unified pathway. Work on separate visual pathways for perception and action suggests that some aspects of behavior can be guided by visual information without explicit, reportable percepts [32], which matters for how we interpret practitioner reports about what decision makers "reason about" versus what they do effectively in action.

Second, in inferential and constructivist accounts of perception and cognition, we emphasize that people do not passively read information from the environment but actively construct interpretations and hypotheses under uncertainty [88, 34, 57, 37]. This framing aligns naturally with practitioner descriptions of action-to-outcome reasoning and with the need to surface assumptions. At the same time, it motivates a concrete caution: shared causal stories can become compelling even when they are weakly constrained by evidence. From this perspective, the central methodological question is not only whether stakeholders agree on a representation, but whether the representation supports revision when evidence conflicts with prior beliefs [34, 57].

Third, Brunswikian and probabilistic functionalist traditions explicitly treat cognition as adaptation to the statistical structure of environments, and they foreground representative design as a methodological response to the ecological validity critique [8, 19, 7]. This line of work shares naturalistic decision making's concern with real contexts, but it pushes us toward measurable questions about which cues people rely on, whether those cues are valid in the environment, and how cue use should change across contexts. It also connects naturally to the heuristics-and-ecological-rationality view that "good enough" strategies can be effective when matched to the environment, even if they deviate from normative ideals [30].

Fourth, Gibsonian ecological psychology emphasizes direct perception of actionable possibilities in the environment, often downplaying internal representations in favor of perception-action coupling [29, 97]. This tradition is sympathetic to practitioner and cognitive engineering emphases on situated action and real constraints, but it also invites skepticism toward representations that require decision makers to externalize elaborate internal causal models. From this viewpoint, a key question is whether decision-support artifacts preserve the action-relevance of information or instead push users toward detached, meeting-room explanations that do not map onto embodied practice [29].

Finally, two cross-cutting bodies of work constrain all of the above: attention and higher-level reasoning. We know that attention is selective, limited, and shaped by both goals and salience [78, 79, 96, 15, 103, 95, 101]. In complex scenes, what users believe they have perceived can diverge sharply from what they actually encoded, as shown by classic change blindness findings [86, 91, 38]. This matters directly for real-world decision making and for naturalistic settings: failures may stem less from "bad judgment" and more from what information was never attended to or retained. Separately, higher-level cognition emphasizes that people use diagrams, analogies, and structured representations to support reasoning [61, 43, 28, 44, 47]. We also know that reasoning in groups can be deeply shaped by argumentative dynamics, where the ability to justify a position can dominate truth-seeking [69]. Together, these results motivate a psychologically grounded reading of practitioner accounts: decision support must manage not only information and uncertainty, but also attention, memory, and social reasoning.

5.1.4.2 Psychologist's Reflections on the CDD Method

We interpret CDDs as psychologically plausible because they externalize a form of causal, action-to-outcome thinking that people readily use when explaining and justifying behavior. As diagrams, they can reduce search and make some inferences easier by placing related information in a single spatial structure, a well-known advantage of diagrammatic representations [61]. They also appear well aligned with constructivist and Brunswikian emphases on making hypotheses and cue structures explicit [88, 34, 19], and they plausibly support organizational coordination by providing a shared object around which stakeholders can align.

At the same time, psychology points to specific risks that we should treat as design and evaluation requirements rather than as abstract critiques. First, attentional selectivity implies that CDD layout, visual salience, and complexity will strongly shape what stakeholders notice and discuss, raising the possibility of structured blind spots even when a diagram feels comprehensive [78, 96, 101]. Evidence on change blindness further suggests that users can remain unaware of meaningful differences unless attention is explicitly guided [86, 91, 38]. Second, inferential accounts predict that users will treat the diagram as a hypothesis and may fill in missing links with prior beliefs [34, 57]. This motivates explicit representation of uncertainty, alternative causal accounts, and mechanisms for revision as evidence accumulates. Third, argumentative theories of reasoning suggest that shared causal diagrams can become tools for persuasion and post-hoc justification rather than for truth-tracking, especially in organizational contexts with power asymmetries [69]. Finally, from an ecological psychology stance, we should ask whether CDDs preserve action relevance and affordance information or whether they over-privilege explicit verbalized reasoning at the expense of perception-action coupling [29].

Taken together, we view CDDs as a promising representational scaffold that aligns with several psychological principles, but we also see a clear research agenda: characterizing cognitive load and scalability limits, measuring attentional guidance and blind spots, testing how diagrams interact with priors and belief updating, understanding how "sensemaking" enables the construction of narrative accounts that are consistent with evidence, of and studying how group reasoning and accountability dynamics build consensus on situation understanding and action plans in practice with the support of diagrammatic visualizations.

5.1.4.3 Challenges for Psychology to Support Real-World Decision Making

From our perspective as psychology researchers, the practitioner account also exposes important limits in how our field currently studies and supports decision making. Much of psychological science has been built on tightly controlled tasks with well-defined stimuli, short time horizons, and individual participants. These designs are powerful for isolating mechanisms, but they often abstract away precisely the features that define practitioner decision making: long causal chains, organizational accountability, evolving goals, and coordination among multiple stakeholders [67, 73, 19]. As a result, our theories can be mechanistically precise yet under-informative about how decisions unfold across weeks, months, or organizational cycles.

A first challenge is temporal scale. Many classic paradigms examine decisions as discrete trials, whereas practitioner decisions are extended processes that are revisited, monitored, and revised. Psychological models of learning and memory do speak to change over time [2], but we rarely study how representations persist as living artifacts that shape future attention, interpretation, and action. We therefore risk missing how decision representations function as long-term cognitive scaffolds rather than momentary aids.

A second challenge is social and organizational embedding. Standard experiments treat decision makers as isolated individuals, yet practitioner accounts emphasize distributed responsibility and shared rationale. While there is work on group reasoning and argumentation [69], we still lack mature paradigms for studying how diagrams and decision models operate as boundary objects across roles, hierarchies, and expertise levels. In practice, a representation may succeed less because it improves individual cognition and more because it stabilizes coordination and accountability, a function our individual-centric methods often miss. One approach to this is paired/group analytics, a Visual Analytics methodology that builds upon Clark's Joint Action Theory from Pragmatic Psycholinguistics by considering the role of interactive visualization displays in building common ground and advancing analysis and decision-making [14, 48, 3].

A third challenge concerns ecological validity and representative design. Brunswikian traditions have long warned that cognition must be studied in relation to real environments [7, 8, 19], yet it remains methodologically difficult to construct studies that are both representative and experimentally tractable. Practitioner settings include rare events, high stakes, and non-repeatable contexts, which resist classic replication logic. We therefore face a trade-off between control and realism that is especially acute for decision-centric representations.

A fourth challenge is attentional and informational overload. Real decisions involve far more variables than typical laboratory tasks. We know that attention is severely limited and guided by both goals and salience [78, 96, 15, 101], and that people routinely miss large changes in complex scenes [86, 91]. However, we do not yet have strong predictive models for how these limits play out in large, semantically rich decision diagrams. Practitioners may experience a representation as clarifying, while cognitively it may still induce systematic blind spots.

A fifth challenge is that practitioners care about action, not just judgment. Many psychological tasks evaluate perceptual accuracy or choice quality, but practitioner success criteria include implementability, legitimacy, and defensibility. A decision representation may be valuable because it supports explanation and justification, even if it does not optimize normative choice. This complicates evaluation: we must ask not only whether a representation leads to "better" decisions, but better for whom, by which metric, and over what time horizon.

Finally, we must acknowledge a translational gap. Psychological findings are often reported at a level of abstraction that is difficult to convert into design guidance for decision-support systems. Practitioner methods like CDDs move quickly because they are actionable and adaptable, even if theoretically under-specified. For psychology to contribute meaningfully, we need frameworks that connect attentional limits, learning dynamics, and reasoning biases to concrete representational design choices, and we need empirical methods that scale beyond micro-tasks.

In sum, the practitioner perspective does not simply provide a new application domain for existing theories. It challenges us to expand our methodological toolbox, to study decisions as temporally extended and socially distributed processes, and to treat external decision representations as primary objects of cognitive analysis rather than as peripheral aids. These challenges also define an opportunity: by engaging seriously with practitioner settings, we can test the scope of our theories and develop more ecologically grounded accounts of human decision making.

5.1.5 Perspectives from a Visualization Researcher

From a visualization perspective, decision making often appears as an intended context of application rather than as the primary theoretical object. The field has traditionally focused on how data is represented, transformed, and interacted with, and on how these representations support tasks such as comparison, pattern finding, or sensemaking. Several reflections within visualization research have noted that decisions typically enter the picture as downstream activities that visualizations are assumed to inform [23, 6, 11, 1]. In psychology and cognitive engineering, a comparable pattern can be observed: decisions are rarely examined in isolation, with greater emphasis placed on processes such as perception, attention, learning, expertise, or coordination in real work settings. Decisions are studied in areas such as judgment and decision making or naturalistic decision making, but they are often treated as one manifestation of broader cognitive or work processes rather than as a standalone endpoint. The distinction, therefore, lies less in which fields study decisions and more in what each field treats as its primary analytic lens. Visualization research has historically evaluated representations and interactions on their own terms, frequently treating improved understanding or insight as proxies for downstream impact, a tendency that has been explicitly discussed within the field [71, 89].

This leads to a subtle epistemic distinction. A substantial portion of visualization research begins from available data or from a data-abstraction problem [63, 90] and asks how that data can be made interpretable and explorable. Many real-world decision contexts instead begin from a problem, a commitment, or a need for action, and only then seek relevant data. This difference in starting point is consequential. Beginning from data or data abstractions can unintentionally prioritize what is measurable or already structured over what is decision-relevant, and may privilege open-ended exploration even when the central challenge lies in framing a decision. In such cases, visualization may support analyses that are informative yet peripheral to the actual commitment at stake. Recognizing this limitation helps clarify where visualization support aligns well with decision needs and where closer coupling to decision framing may be necessary [24].

5.1.5.1 Conceptual and Methodological Reflections from Visualization Research on Real-World Decision Making

Methodologically, when visualization research engages with decision making, it draws largely on traditions from cognitive psychology, behavioral decision research, and decision theory, where decisions are typically operationalized as judgments or choices among predefined alternatives under uncertainty [74, 9, 22, 45, 76]. Many visualization studies therefore adopt controlled experimental paradigms that measure estimation error, bias, preference, or decision performance in well-defined task settings [49, 74, 50, 45]. Reviews in this area explicitly ground visualization-based decision studies in cognitive theories of judgment and choice, including dual-process accounts [76]. Within this framing, decisions are treated as observable responses that reveal how visual encodings shape perception, reasoning, and evaluation.

A second line of visualization research addresses decisions in multi-attribute and trade-off contexts, often drawing from multi-criteria decision analysis and engineering optimization [10, 21, 13, 18, 105]. Here, the decision problem is formalized as a set of alternatives characterized by multiple attributes or criteria, and visualizations are designed to support comparison, trade-off inspection, and preference articulation. Pareto-front and multi-criteria visualizations, for example, support the selection of compromise solutions among many candidates, where preferences may evolve during exploration but the alternative space and criteria are predefined. Empirical studies in this area evaluate how multidimensional visualizations support structured choice tasks using measures such as accuracy, time, and subjective assessments when users compare alternatives under fixed criteria. Compared to psychology-grounded studies, this strand accommodates richer decision structures and trade-offs at the cost of less insight into how preferences and judgments are formed or influenced by cognitive factors, yet it still relies on formalized and bounded problem formulations where options and criteria are specified in advance.

Visualization research has also engaged with situated and practice-oriented inquiry through design study and qualitative methodologies. Design study frameworks emphasize sustained collaboration with domain experts, evolving problem definitions, and reflective accounts of visualization use in real contexts [70]. Epistemically, this aligns with ethnographic and cognitive-engineering traditions in grounding knowledge in real work practices and contextual constraints. At the same time, these approaches typically remain data-centric: they begin from existing datasets and examine how data representations are integrated into practice. Decision making is therefore often treated as part of the surrounding context rather than as the primary object of modeling. Related work has drawn on concepts such as sensemaking, situation awareness, and decision elicitation methods to inform design [5, 59, 12].

Visualization research has nevertheless made increasing efforts to treat decision making more explicitly, with empirical studies showing that accurate judgments or perceptual performance do not necessarily translate into effective decisions [50, 20, 74], characterizations and typologies of decision making tasks that were previously omitted from visualization taxonomies [11, 6, 23], literature reviews pointing to limitations in visualization tools for decision support [75], and decision making research gaps relevant to visualization [94]. Taken together, these developments reflect a gradual shift toward recognizing decisions as distinct analytic targets rather than implicit endpoints of analysis.

5.1.5.2 Visualization Researcher's Reflections on the CDD Method

Several forms of visualization are required to support the CDD method across its lifecycle and the group discussions spent some time reflecting on these visualisation challenges.

First, the construction of the CDD is grounded in participatory design and collaborative sensemaking. Visualization is used to externalize mental models and support the joint elicitation of goals, decision variables, causal relationships, intermediate states, and external factors. In this stage, node – link representations such as causal loop diagrams serve as exploratory visual abstractions that facilitate discussion, hypothesis generation, and alignment among stakeholders. There has been some recent interest in visualisation research in regard to causal loop diagrams (e.g. [72, 102]) and there are opportunities to consider the psychological implications of these approaches. For some of the problems discussed, causal loop diagrams can quickly become very large and complex, with up to several hundred inputs and outputs and a large number of links connecting them with intermediate nodes. These inputs are likely to have hierarchies and relationships which provides opportunities for effective visualisation. Indeed, network visualisation techniques such as edge bundling [62] might be used to provide clarity of network layout, making it easier for decision-makers to identify the most important causal links. Finally, one consideration for the use of causal loop diagrams is capturing design provenance, such as the rationale for including or excluding causal links, assumptions made by participants, and areas of disagreement. Interactive visualisation tools to track this provenance are likely to form an important design requirement. One such prototype is provided by Bou Nassar et al., 2025 [72].

Second, the exploration of simulation outputs under varying initial conditions and decision parameters introduces challenges related to scale, dimensionality, and uncertainty. Supporting these analytic tasks requires interactive visualization techniques, including filtering, evaluation of parameter selections, and dimensionality-reduction methods. Visual representations of decision boundaries in optimization and design spaces may support comparative analysis and trade-off reasoning [93]. Given the inherent uncertainty in model structure and parameter estimates, uncertainty visualization techniques such as ensemble views, probability distributions, or confidence encodings will be essential for conveying the robustness and sensitivity of simulation results.

One important goal discussed in our group was the need to provide aesthetic and immersive visualisation solutions to promote engagement with executives who may be time poor. Immersive analytics and virtual reality environments may also be explored as interaction paradigms. The ability for such environments to provide effective ways of interacting with high volume data in collaborative ways is a key research goal for Immersive Analytics [68] but there are also opportunities for considering rendering of natural scenes with information overlays with the goal of promoting engagement and time spent exploring different parameter configurations and simulation outputs.

Third, once a decision is enacted, visualization supports ongoing monitoring, model validation, and accountability. Measures identified during the participatory design phase are tracked over time using dashboard-style visualizations commonly found in business intelligence systems. These time-oriented views enable the detection of deviations, threshold violations, and emerging trends, while uncertainty-aware encodings can help distinguish meaningful change from noise. Provenance visualization plays a key role at this stage by linking observed outcomes back to prior decisions, assumptions, and model versions, thereby supporting retrospective analysis and learning. When significant deviations occur, they may trigger renewed sensemaking activities and revisions to the underlying CDD. In this monitoring phase, established visualization principles – such as perceptually effective visual encodings, appropriate use of color, and faithful representation of quantitative values and uncertainty – are critical to supporting accurate interpretation and decision-making.

5.1.5.3 Challenges for Visualization Research to Support Real-World Decision Making

Across visualization research, there is increasing evidence that many domains routinely visualize processes, workflows, and conceptual structures, yet these representations are rarely treated as first-class objects of visualization research. In process mining and visual process analytics, the primary visual objects are process models derived from event logs, such as control-flow graphs and directly-follows relations, which support inspection of execution order, variants, and deviations but largely assume a fixed underlying process structure [60, 58]. Recent work acknowledges that such representations become insufficient when analysts must reason across multiple perspectives, including temporal, organizational, and compliance-related views, and calls for deeper integration of visual analytics techniques to support sensemaking rather than mere inspection [64]. Ribarsky et al. (2014) similarly argue that complex business processes should be treated as analytical entities composed of events and states, motivating visual analytics environments that go beyond monitoring raw data streams [87]; however, their framework remains grounded in event-centric data analysis rather than the interactive construction or negotiation of process meaning. In applied visualization research on Building Information Modeling, Ivson et al. (2020) document extensive use of visual representations such as schedules, networks, and schematic diagrams to coordinate workflows and information flow, but these practices are primarily surveyed as domain-specific solutions rather than articulated as general visualization paradigms [46]. From a complementary visualization perspective, Lohfink et al. (2022) introduce knowledge-assisted visualization and explicitly distinguish data from knowledge, emphasizing externalization and guidance [65]; nonetheless, knowledge structures are typically assumed to be predefined rather than dynamically constructed, revised, or contested through interaction. Taken together, these works reveal a consistent gap: while processes, workflows, and conceptual structures are widely visualized across domains, visualization research has only partially addressed how such non-dataset objects can be supported as interactive, revisable, and epistemic artifacts for reasoning and decision making.

5.1.6 Conclusion

Across the perspectives discussed in this report, a consistent pattern emerges: decision making is rarely absent from research agendas, yet it is seldom treated as the central unit of analysis. Psychology, cognitive engineering, and visualization each illuminate important facets of decision processes, but they often do so through lenses optimized for their own methodological traditions. Psychology provides models of cognition, bias, and judgment, often under controlled conditions. Cognitive engineering brings attention to real work, context, and distributed cognition. Visualization research contributes powerful methods for externalizing information and supporting sensemaking. However, these contributions do not automatically converge on a shared understanding of decisions as temporally extended, revisitable, and socially embedded commitments to action.

A recurring tension concerns the gap between judgment and decision, between analysis and commitment, and between individual cognition and organizational practice. Many methods evaluate how well people interpret information, but practitioners ultimately care about how representations support action, coordination, justification, and accountability over time. This does not invalidate existing approaches; rather, it clarifies their scope. Controlled experiments, formal decision models, and design studies each capture different slices of decision phenomena. The challenge is not to replace one with another, but to articulate how they relate and where their assumptions diverge.

One implication is that decision support cannot be reduced to improving perceptual accuracy or analytic efficiency alone. Decisions unfold across episodes, actors, and constraints, often under evolving goals and incomplete information. External representations, including visualizations, function not only as cognitive aids but also as social and organizational artifacts. Studying them as such requires methods that tolerate complexity, partial observability, and context-dependence.

Looking forward, progress may lie less in proposing a single unified theory and more in developing bridges between levels of analysis: linking micro-level cognitive findings to meso-level representational practices and macro-level organizational decisions. Visualization research is well positioned to contribute here because it already operates at the interface between data, cognition, and action. Treating decisions themselves as first-class analytic objects, rather than implicit endpoints of analysis, could help align research more closely with how decisions are actually made in practice.

In this sense, the dialogue between disciplines is not only descriptive but generative. By making assumptions explicit and comparing methodological lenses, we gain a clearer view of what each field can and cannot explain. This clarity is a prerequisite for building decision-support approaches that are both scientifically grounded and practically meaningful.

5.1.7 Next Steps

Our immediate goal is to submit a dialogic paper containing the four perspectives provided above. During the seminar, we have identified the importance of sharing disciplinary perspectives concerning decision making, acknowledging that different fields bring associated methodological frameworks that can be complementary across decision-making contexts. The visualization community can play a pivotal role in supporting the integration of different approaches, by providing technical capability and by playing a brokering role between pure research performed in psychology and industry applications. Possible venues include CG&A, the Asian Conference on Ergonomic Design taking place in September 2026 or the International Conference on Engineering Design. Other options include Cognition, Technology and Work and Topics in Cognitive Science. We have also identified two funding opportunities which could support an interdisciplinary project on visualisation supported decision making with industry applications.

A broader goal of our working group is to consider whether decision making deserve its own "subfield" or community within vis research. IEEE VIS has an area for "analytics and decisions" yet the focus here is on human-machine workflows and less on human-centric approaches. There may be scope for decision making research in visualisation which incorporates the practitioner, cognitive engineering and psychology perspectives presented here.

5.1.7.1 Acknowledgments

This working group report was conceived at Dagstuhl Seminar 25511 "From Psychology Enabled Visualization to Visualization Enabled Psychology" (December 14–19, 2025). We thank the seminar organizers – Rita Borgo, Min Chen, David S. Ebert, Brian D. Fisher, and Alvitta Ottley – and seminar participants for insightful contributions and discussion. We also acknowledge Schloss Dagstuhl – Leibniz Center for Informatics where this collaboration took place.

References

1 Tomás Alves, Tiago Delgado, Joana Henriques-Calado, Daniel Gonçalves, and Sandra Gama. Exploring the role of conscientiousness on visualization-supported decision-making. *Computers & Graphics*, 111:47–62, April 2023.
2 John Robert Anderson. *Learning and memory: An integrated approach, 2nd ed.* Learning and memory: An integrated approach, 2nd ed. John Wiley & Sons Inc, Hoboken, NJ, US, 2000.
3 Richard Arias-Hernandez, Linda T. Kaastra, Tera M. Green, and Brian Fisher. Pair analytics: Capturing reasoning processes in collaborative visual analytics. In *2011 44th Hawaii International Conference on System Sciences*, pages 1–10, 2011.
4 W. Ross Ashby. *An introduction to cybernetics.* Martino Publishing, Mansfield Centre, CT, 2015.
5 Simon J. Attfield, Sukhvinder K. Hara, and B. L. William Wong. Sensemaking in Visual Analytics: Processes and Challenges, 2010. EuroVAST 2010: International Symposium on Visual Analytics Science and Technology.
6 Camelia D. Brumar, Sam Molnar, Gabriel Appleby, Kristi Potter, and Remco Chang. A Typology of Decision-Making Tasks for Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 31(10):8536–8551, October 2025.
7 Egon Brunswik. *The conceptual framework of psychology.* University of Chicago Press, Chicago, 1979.
8 Egon Brunswik. Perception and the Representative Design of Psychological Experiments [1956]. In Kenneth R Hammond and Thomas R Stewart, editors, *The Essential Brunswik*, pages 260–264. Oxford University PressNew York, NY, September 2001.
9 Mengyu Chen, Andrew Yang, Seungchan Min, Kristy A Hamilton, and Emily Wall. A Novel Lens on Metacognition in Visualization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–16, Yokohama Japan, April 2025. ACM.
10 Lena Cibulski. *Visual Analysis for Multi-Attribute Choice.* PhD thesis, Darmstadt, Technischen Universität and Informatik, Fachbereich, 2024.
11 Lena Cibulski and Stefan Bruckner. Towards Understanding Decision Problems As a Goal of Visualization Design, July 2025. arXiv:2507.18428 [cs].

12 Lena Cibulski, Evanthia Dimara, Setia Hermawati, and Jorn Kohlhammer. Supporting Domain Characterization in Visualization Design Studies With the Critical Decision Method. In *2022 IEEE 4th Workshop on Visualization Guidelines in Research, Design, and Education (VisGuides)*, pages 8–15, Oklahoma City, OK, USA, October 2022. IEEE.

13 Lena Cibulski, Hubert Mitterhofer, Thorsten May, and Jörn Kohlhammer. PAVED: Pareto Front Visualization for Engineering Design. *Computer Graphics Forum*, 39(3):405–416, June 2020.

14 Herbert H. Clark. *Using Language*. "Using" Linguistic Books. Cambridge University Press, 1996.

15 Maurizio Corbetta and Gordon L. Shulman. Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3):201–215, March 2002.

16 Beth Crandall, Gary A. Klein, and Robert R. Hoffman. *Working minds: a practitioner's guide to cognitive task analysis*. MIT Press, Cambridge, Mass, 2006.

17 David Pidsley. Gartner To Publish New Magic Quadrant for Decision Intelligence Platforms, December 2025. Available online at: https://www.linkedin.com/pulse/gartner-publish-new-magic-quadrant-decision-platforms-david-pidsley-xanhe (Accessed: 2026-01-05).

18 Tiago Delgado, Tomás Alves, and Sandra Gama. How neuroticism and locus of control affect user performance in high-dimensional data visualization. *Computers & Graphics*, 109:88–99, December 2022.

19 Mandeep K. Dhami, Ralph Hertwig, and Ulrich Hoffrage. The Role of Representative Design in an Ecological Approach to Cognition. *Psychological Bulletin*, 130(6):959–988, 2004.

20 Evanthia Dimara, Anastasia Bezerianos, and Pierre Dragicevic. Narratives in Crowdsourced Evaluation of Visualizations: A Double-Edged Sword? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 5475–5484, Denver Colorado USA, May 2017. ACM.

21 Evanthia Dimara, Evanthia Dimara, Anastasia Bezerianos, Pierre Dragicevic, and Pierre Dragicevic. Conceptual and Methodological Issues in Evaluating Multidimensional Visualizations for Decision Support. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):749–759, January 2018.

22 Evanthia Dimara, Evanthia Dimara, Steven Franconeri, Steven Franconeri, Catherine Plaisant, Catherine Plaisant, Anastasia Bezerianos, Pierre Dragicevic, and Pierre Dragicevic. A Task-Based Taxonomy of Cognitive Biases for Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 26(2):1413–1432, February 2020.

23 Evanthia Dimara and John Stasko. A Critical Reflection on Visualization Research: Where Do Decision Making Tasks Hide? *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1128–1138, January 2022.

24 Evanthia Dimara, Harry Zhang, Melanie Tory, and Steven Franconeri. The Unmet Data Visualization Needs of Decision Makers Within Organizations. *IEEE Transactions on Visualization and Computer Graphics*, 28(12):4101–4112, December 2022.

25 Eric Hill. Quantellia: Take Charge of Outcomes with Decision Intelligence, December 2025. Available online at: https://ciotechworld.com/quantellia-take-charge-of-outcomes-with-decision-intelligence (Accessed 2025-01-06).

26 K. Anders Ericsson, Robert R. Hoffman, Aaron Kozbelt, and A. Mark Williams, editors. *The Cambridge Handbook of Expertise and Expert Performance*. Cambridge University Press, 2 edition, April 2018.

27 K. Anders Ericsson and Jacqui Smith. *Toward a general theory of expertise : prospects and limits*. Cambridge University Press, Cambridge, UK, 1991.

28 D Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, June 1983.

29 James J. Gibson. *The Ecological Approach to Visual Perception: Classic Edition*. Psychology Press, 1 edition, November 2014.

30 Gerd Gigerenzer and Peter M. Todd. *Simple heuristics that make us smart*. Evolution and cognition. Oxford University Press, New York, 1999.

31 E. Bruce Goldstein. *Sensation and perception*. Wadsworth Cengage Learning, Australia Brazil Japan Korea Mexico Singapore Spain United Kingdom United States, ninth edition, student edition edition, 2014.

32 Melvyn A. Goodale and A.David Milner. Separate visual pathways for perception and action. *Trends in Neurosciences*, 15(1):20–25, January 1992.

33 Grand View Research. Decision Intelligence Market To Reach $36.34 Billion By 2030, June 2025. Available online at: https://www.grandviewresearch.com/press-release/global-decision-intelligence-market (Accessed: 2026-01-05).

34 Richard Langton Gregory. Perceptions as hypotheses. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 290(1038):181–197, July 1980.

35 Gurobi Optimization. The Decision Intelligence Summit, 2025. Available online at: https://www.gurobi.com/lp/all/the-decision-intelligence-summit (Accessed 2026-01-05).

36 Gurobi Optimization. Gurobi to Bring Decision Intelligence Summit to Vienna, October 2025. Available online at: https://www.gurobi.com/news/gurobi-to-bring-decision-intelligence-summit-to-vienna (Accessed 2026-01-05).

37 Hermann von Helmholtz and James P. C. Southall. *Treatise on physiological optics*. Dover phoenix editions. Dover Publications, Mineola, NY, dover ed edition, 2005.

38 John M. Henderson and Andrew Hollingworth. The Role of Fixation Position in Detecting Scene Changes Across Saccades. *Psychological Science*, 10(5):438–443, September 1999.

39 Erik Hollnagel, Robert L Wears, Jeffrey Braithwaite, and others. From Safety-I to Safety-II: a white paper. *The resilient health care net: published simultaneously by the University of Southern Denmark, University of Florida, USA, and Macquarie University, Australia*, 2015.

40 Erik Hollnagel and David D. Woods. Cognitive Systems Engineering: New wine in new bottles. *International Journal of Man-Machine Studies*, 18(6):583–600, June 1983.

41 Erik Hollnagel and David D. Woods. *Joint Cognitive Systems: Foundations of Cognitive Systems Engineering*. CRC Press, 0 edition, February 2005.

42 Erik Hollnagel, David D. Woods, and Nancy Leveson, editors. *Resilience engineering: concepts and precepts*. Ashgate, Aldershot, England ; Burlington, VT, 2006.

43 Keith James Holyoak and Robert G. Morrison. *The Cambridge handbook of thinking and reasoning*. Cambridge university press, Cambridge, 2005.

44 Keith James Holyoak and Paul Thagard. *Mental leaps: analogy in creative thought*. A Bradford book. MIT Press, Cambridge, Mass., 1996.

45 Jessica Hullman, Alex Kale, and Jason Hartline. Underspecified Human Decision Experiments Considered Harmful. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, pages 1–15, New York, NY, USA, April 2025. Association for Computing Machinery.

46 Paulo Ivson, André Moreira, Francisco Queiroz, Wallas Santos, and Waldemar Celes. A Systematic Review of Visualization in Building Information Modeling. *IEEE Transactions on Visualization and Computer Graphics*, 26(10):3109–3127, October 2020.

47 Philip Nicholas Johnson-Laird. *Mental models: towards a cognitive science of language, inference, and consciousness*. Number 6 in Cognitive science series. Harvard Univ. Press, Cambridge, Mass, 6. print edition, 1995.

48 Linda T. Kaastra and Brian Fisher. Field experiment methodology for pair analytics. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, BELIV '14, pages 152–159, New York, NY, USA, 2014. Association for Computing Machinery.

49 Alex Kale, Matthew Kay, and Jessica Hullman. Visual Reasoning Strategies for Effect Size Judgments and Decisions. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):272–282, February 2021.

50 Jade Kandel, Jiayi Liu, Arran Zeyu Wang, Chin Tseng, and Danielle Albers Szafir. Graphical Perception of Icon Arrays versus Bar Charts for Value Comparisons in Health Risk Communication. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–11, 2025.

51 G. Klein, B. Moon, and R.R. Hoffman. Making Sense of Sensemaking 1: Alternative Perspectives. *IEEE Intelligent Systems*, 21(4):70–73, July 2006.

52 G. Klein, K.G. Ross, B.M. Moon, D.E. Klein, R.R. Hoffman, and E. Hollnagel. Macrocognition. *IEEE Intelligent Systems*, 18(3):81–85, May 2003.

53 Gary Klein. A Recognition Primed Decision (RPD) Model of Rapid Decision Making. In *Decision making in action: models and methods*. Ablex Publishing Group, January 1993.

54 Gary Klein. *Sources of power: how people make decisions*. MIT Press, Cambridge, Mass., 7th print edition, 2001.

55 Gary Klein. Naturalistic Decision Making. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 50(3):456–460, June 2008.

56 Gary A. Klein, editor. *Decision making in action: models and methods*. Ablex, Norwood, N.J, 2. print edition, 1995.

57 David C. Knill and Whitman Richards, editors. *Perception as Bayesian inference*. Cambridge University Press, Cambridge, 2008.

58 Steven Knoblich, Jan Mendling, and Helena Jambor. Review of Visual Encodings in Common Process Mining Tools. *VIPRA 2024 – Visual Process Analytics Workshop*, 2024. The Eurographics Association.

59 Jörn Kohlhammer, Thorsten May, and Marcus Hoffmann. Visual Analytics for the Strategic Decision Making Process. In Raffaele De Amicis, Radovan Stojanovic, and Giuseppe Conti, editors, *GeoSpatial Visual Analytics*, pages 299–310, Dordrecht, 2009. Springer Netherlands.

60 Simone Kriglstein, Margit Pohl, Stefanie Rinderle-Ma, and Magdalena Stallinger. Visual Analytics in Process Mining: Classification of Process Mining Techniques, 2016. EuroVis Workshop on Visual Analytics (EuroVA).

61 Jill H. Larkin and Herbert A. Simon. Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science*, 11(1):65–100, January 1987.

62 A. Lhuillier, C. Hurter, and A. Telea. State of the Art in Edge and Trail Bundling Techniques. *Computer Graphics Forum*, 36(3):619–645, June 2017.

63 Haihan Lin, Derya Akbaba, Miriah Meyer, and Alexander Lex. Data Hunches: Incorporating Personal Knowledge into Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):504–514, January 2023.

64 Sanne Van Der Linden, Velitchko Andreev Filipov, Luise Pufahl, Silvia Miksch, and Stef Van Den Elzen. Towards Integrating Visual Analytics in Multi-Perspective Conformance Checking: A Call to Action. *EuroVis Workshop on Visual Analytics (EuroVA) 2025*, 2025. The Eurographics Association.

65 Anna-Pia Lohfink, Simon D. Duque Anton, Heike Leitte, and Christoph Garth. Knowledge Rocks: Adding Knowledge Assistance to Visualization Systems. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1117–1127, January 2022.

66 Nadine Malcolm and Lorien Pratt. P-decisions, c-decisions, and why the difference matters more than ever, February 2026. Accessed: 2026-02-12.

67 David Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. MIT Press, July 2010.

68 Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H. Thomas, editors. *Immersive Analytics*, volume 11190 of *Lecture Notes in Computer Science*. Springer International Publishing, Cham, 2018.

69 Hugo Mercier and Dan Sperber. Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, 34(2):57–74, April 2011.

70 Miriah Meyer and Jason Dykes. Criteria for Rigor in Visualization Design Study. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2019.

71 Tamara Munzner. Process and Pitfalls in Writing Information Visualization Research Papers. In Andreas Kerren, John T. Stasko, Jean-Daniel Fekete, and Chris North, editors, *Information Visualization*, volume 4950, pages 134–153. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. Series Title: Lecture Notes in Computer Science.

72 Jessica Bou Nassar, Yu Xuan Yio, Nethara Athukorala, Simran, Songhai Fan, Cynthia A. Huang, Lyn Bartram, Tim Dwyer, and Sarah Goodwin. 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)*, pages 111–115, Vienna, Austria, November 2025. IEEE.

73 Ulric Neisser. *Cognitive Psychology*. Psychology Press, November 2014.

74 Başak Oral, Pierre Dragicevic, Alexandru Telea, and Evanthia Dimara. Decoupling Judgment and Decision Making: A Tale of Two Tails. *IEEE Transactions on Visualization and Computer Graphics*, 30(10):6928–6940, October 2024.

75 Emre Oral, Ria Chawla, Michel Wijkstra, Narges Mahyar, and Evanthia Dimara. From Information to Choice: A Critical Inquiry Into Visualization Tools for Decision Making. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–11, 2023.

76 Lace M. Padilla, Sarah H. Creem-Regehr, Mary Hegarty, and Jeanine K. Stefanucci. Decision making with visualizations: a cognitive framework across disciplines. *Cognitive Research: Principles and Implications*, 3(1):29, December 2018.

77 Stephen E. Palmer. *Vision science: photons to phenomenology*. MIT Press, Cambridge, Mass, 1999.

78 Michael I. Posner. Orienting of Attention. *Quarterly Journal of Experimental Psychology*, 32(1):3–25, February 1980.

79 Michael I. Posner, Charles R. Snyder, and Brian J. Davidson. Attention and the detection of signals. *Journal of Experimental Psychology: General*, 109(2):160–174, June 1980.

80 L. Y. Pratt and N. E. Malcolm. *The Decision Intelligence Handbook*. O'Reilly Media, Incorporated, Sebastopol, 1st ed edition, 2023.

81 Lorien Pratt. *Link: How Decision Intelligence Connects Data, Actions, and Outcomes for a Better World*. Emerald Publishing Limited, Bingley, 2019.

82 Lorien Pratt, Christophe Bisson, and Thierry Warin. Bringing advanced technology to strategic decision-making: The Decision Intelligence/Data Science (DI/DS) Integration framework. *Futures*, 152:103217, September 2023.

83 Lorien Pratt and Mark Zangari. Decision Engineering White Paper, 2008. Available online at: https://www.quantellia.com/Data/WP-OvercomingComplexityv1_1.pdf (Accessed 2026-01-06).

84 Lorien Pratt and Mark Zangari. High Performance Decision Making: A Global Study, 2009. Available online at: https://www.quantellia.com/Data/HighPerformanceDecisionMaking.pdf (Accessed 2026-01-06).

85 Precedence Research. Decision Intelligence Market Size, Share, and Trends 2024 to 2034, October 2024. Available online at: https://www.precedenceresearch.com/decision-intelligence-market (Accessed 2026-01-05).

86 Ronald A. Rensink, J. Kevin O'Regan, and James J. Clark. To See or not to See: The Need for Attention to Perceive Changes in Scenes. *Psychological Science*, 8(5):368–373, September 1997.

87 William Ribarsky, Derek Xiaoyu Wang, Wenwen Dou, and William J. Tolone. Towards a Visual Analytics Framework for Handling Complex Business Processes. In *2014 47th Hawaii International Conference on System Sciences*, pages 1374–1383, Waikoloa, HI, January 2014. IEEE.

88 Irvin Rock. *The logic of perception.* MIT Press, Cambridge, Mass, 1983.

89 Michael Sedlmair. Design Study Contributions Come in Different Guises: Seven Guiding Scenarios. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization*, pages 152–161, Baltimore MD USA, October 2016. ACM.

90 Michael Sedlmair, Miriah Meyer, and Tamara Munzner. Design Study Methodology: Reflections from the Trenches and the Stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, December 2012.

91 Daniel J. Simons and Daniel T. Levin. Change blindness. *Trends in Cognitive Sciences*, 1(7):261–267, October 1997.

92 SNS Insider pvt ltd. Decision Intelligence Market to Reach USD 74.23 Billion by 2033, Owing to Rising Adoption of AI-Enabled Data-Driven Decision-Making, December 2025. Available online at: https://www.globenewswire.com/news-release/2025/12/07/3201138/0/en/Decision-Intelligence-Market-to-Reach-USD-74-23-Billion-by-2033-Owing-to-Rising-Adoption-of-AI-Enabled-Data-Driven-Decision-Making-SNS-Insider.html (Accessed 2026-01-05).

93 Jan-Tobias Sohns, Christoph Garth, and Heike Leitte. Decision Boundary Visualization for Counterfactual Reasoning. *Computer Graphics Forum*, 42(1):7–20, February 2023.

94 Midori Sugihara, Shuhei Takakai, Kazutaka Takamatsu, and Kazuo Misue. Contribution of Data Visualization to Decision-Making: A Classification of Data Visualization Research Based on the Characteristics of Decision Problems. In *2025 IEEE 18th Pacific Visualization Conference (PacificVis)*, pages 269–278, Taipei City, Taiwan, April 2025. IEEE.

95 Jan Theeuwes. Perceptual selectivity for color and form. *Perception & Psychophysics*, 51(6):599–606, November 1992.

96 Anne M. Treisman and Garry Gelade. A feature-integration theory of attention. *Cognitive Psychology*, 12(1):97–136, January 1980.

97 M.T. Turvey. Action and perception at the level of synergies. *Human Movement Science*, 26(4):657–697, August 2007.

98 Barbara Tversky. Spatial thought, social thought. In Thomas W. Schubert and Anne Maass, editors, *Spatial Dimensions of Social Thought*, pages 17–38. DE GRUYTER, October 2011.

99 Kim J. Vicente. *Cognitive Work Analysis.* CRC Press, April 1999.

100 K.J. Vicente and J. Rasmussen. Ecological interface design: theoretical foundations. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(4):589–606, August 1992.

101 Jeremy M. Wolfe and Todd S. Horowitz. What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, 5(6):495–501, June 2004.

102 Xiao Xie, Fan Du, and Yingcai Wu. A Visual Analytics Approach for Exploratory Causal Analysis: Exploration, Validation, and Applications. *IEEE Transactions on Visualization and Computer Graphics*, 27(2), February 2021.

103 Steven Yantis and John Jonides. Abrupt visual onsets and selective attention: Evidence from visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 10(5):601–621, 1984.

104 Abdullah Kaan Zaimoglu, Lorien Pratt, and Brian Fisher. Epistemological role of human reasoning in data-informed decision-making. *Frontiers in Communication*, 8:1250301, November 2023.

105 Ying Zhao, Feng Luo, Minghui Chen, Yingchao Wang, Jiazhi Xia, Fangfang Zhou, Yunhai Wang, Yi Chen, and Wei Chen. Evaluating Multi-Dimensional Visualizations for Understanding Fuzzy Clusters. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):12–21, January 2019.

Part II

Application

6 On the Psychology of Storytelling with Visualizations

6.1 Psychology-Enabled Visual Communication Related to Climate Change to Encourage Climate-Friendly Behavior

Rita Borgo (King's College London, GB, rita.borgo@kcl.ac.uk)
Yuri Engelhardt (University of Twente – Enschede, NL, yuri.engelhardt@gmail.com)
Eva Mayr (Donau-Universität Krems, AT, eva.mayr@donau-uni.ac.at)
Mauro Martino (IBM Research – Cambridge, US, mauro.martino@mamartino.com)
Margit Pohl (TU Wien, AT, margit.pohl@tuwien.ac.at)[2]

6.1.1 Introduction & Motivation

Visualizations are frequently used in public communication settings [19], for example, in data journalism [11], science communication [13], or cultural exhibitions [30], to support a story with facts, to better ground it in data and to make it more trustworthy. These applications share principles like storytelling [28, 23, 2], the use of metaphors [4], onboarding [25], and interactivity to strengthen their impact.

However, when taking a deeper look at the effects of these public communication projects, they frequently fail to reach a wider impact: People oftentimes do not interact with the visualizations [31] and they engage more superficially, let themselves be entertained, and want to feel informed. That is, why we consider it worth looking at the underlying mechanisms of public communication with visualizations from a psychological perspective. **How can psychological theories, models, and findings inform the design and evaluation of visualizations for public communication for a wider impact**?

6.1.2 Scope & Definitions

In this paper, with "visual communication to the public" we mean communicating information visually to general and diverse audiences – which may concern a variety of domains and include a variety of aims. We have chosen to focus on the domain of the climate crisis, which includes both the consequences of climate change and the possibilities for climate-friendly action. Visual communication to the public to encourage climate-friendly behaviour is in line with calls to turn eco-anxiety into eco-action [7].

Encouraging climate-friendly behaviour can be looked at the level of the individual, the level of the peer group (or family), and the level of the larger community. We discuss relevant theories from psychology for each of these levels.

[2] Authors listed in alphabetical order.

Table 2 Design recommendations for climate change data visualisation derived from cognitive and psychological science [14]. (Placeholder: to be populated.)

Psychological Insight	**Design Recommendation**
1. Intuitions about effective graphics do not always correspond to evidence-informed best practice for increasing accessibility	Use cognitive and psychological principles to inform the design of graphics. Test graphics during development to understand viewers' comprehension.
Direct visual attention	
2. Visual attention is limited and selective.	Present only visual information required for the communication goal. Direct viewers' attention to features supporting inferences.
3. Salient visual features attract attention.	Make important features perceptually salient (size, shape, colour, motion).
4. Prior experience and knowledge guide attention.	Design graphics informed by viewers' familiarity and domain knowledge. Provide guidance via text positioned close to the graphic.
Reduce complexity	
5. Excess visual information creates clutter.	Include only necessary information. Break graphics into meaningful visual chunks.
Support inference-making	
6. Some inferences require spatial reasoning skills.	Reduce need for spatial transformations by showing inferences directly. Provide textual guidance where necessary.
7. Visual structure influences inference.	Identify key relationships and structure data to make them salient.
8. Animation may help or hinder comprehension.	Base animation decisions on cognitive principles. Provide user control over playback and speed.
9. Conceptual thought uses cultural metaphors.	Align representations with helpful metaphors (e.g., up = good, down = bad).
Integrate text with graphics	
10. Spatial separation splits attention.	Keep graphics and explanatory text close together.
11. Language influences thought.	Use text to guide comprehension and interpretation.

Psychological theories are examined in the context of gamification, storytelling, social factors, etc., and applied to the design of a prototype example. While focussing on climate-friendly behaviour, the protoype will enable the browsing of relevant data to a limited degree – however, it will not allow for extensive data analysis.

6.1.3 Relevant Theories & Empirical Foundations

A lot of research from a psychological perspective exists on climate-related communication. When looking at papers addressing visual communication, many focus on the selection of appropriate images, less on data visualizations [14, 30, 21]. For data visualizations, Harold et al. [14] derived a set of useful design guidelines for climate visualizations from a cognitive and psychological perspective (see Table 2). These suggestions are useful for making climate visualizations easier to comprehend, but do not focus on the question of how to design visualizations going beyond comprehension – towards attitude change and climate action [30]. Table 2 summarizes design recommendations from [14].

To fill this gap, we reviewed a selection of potentially relevant psychological theories and derived potentially relevant design recommendations.

6.1.3.1 Theory of Planned Behavior / Theory of Reasoned Action

The Theory of Reasoned Action (TRA) and its extension, the Theory of Planned Behavior (TPB) [1], explain behaviour as the outcome of a reasoned process grounded in underlying beliefs. The central determinant of behaviour is behavioural intention, defined as a person's readiness to perform a specific action. Intention is shaped by three factors: (1) Attitude toward the behaviour, based on beliefs about its consequences and their evaluation; (2) Subjective norm, reflecting perceived social pressure from relevant others; and (3) Perceived behavioural control, capturing the perceived ease or difficulty of acting, based on beliefs about resources and barriers. While TRA assumes full volitional control, TPB recognises that many climate-relevant behaviours are constrained by contextual factors. Perceived behavioural control therefore influences both intention and behaviour directly when actual control is limited. Importantly, communication affects behaviour indirectly by changing the behavioural, normative, and control beliefs that underpin these determinants.

Visualization Design Recommendation. Climate visualisations should:

- Support attitudes. Clearly visualise meaningful behavioural consequences.
- Highlight norms. Show what relevant others approve of and are doing.
- Provide feasibility cues. Represent practical steps, resources, and barriers.
- Show behavioural specificity. Depict concrete, actionable behaviours rather than abstract goals.
- Align determinants. Ensure evaluative, normative, and control signals consistently support the same action.
- Target key beliefs. Focus on beliefs that differentiate actors from non-actors.

6.1.3.2 Construal Level Theory / Psychological Distance

Psychological distance is defined as the extent to which an event is perceived as being distant from one's direct experience – spatially, temporally, socially, or hypothetically [29]. According to *Construal Level Theory* (CLT), distant events are construed in a more abstract way. In turn, more abstract mental representations would make people perceive them as psychologically more distant [29].

It has been suggested that psychological distance will influence the success of climate communication. Empirical evidence on the influence of abstractness of visualizations on the perception of climate risk and climate-related behavior is mixed: Depending on risk perception [9] and visual literacy [10], people will react differently to psychologically more or less distant visualizations: "When people perceived high risk, abstract visuals strengthened this impact and made them feel psychologically proximal to the risk. When people saw themselves at less risk and perceived climate risk as less serious, concrete visuals were more effective in reducing their psychological distance" [9]. People with low visual literacy, perceive abstract visuals as more psychological distant and are less easily influenced in their behavioral intentions and climate concerns, whereas for people with high visual literacy more abstract visualizations have higher impact.

Visualization Design Recommendation. Though the empirical evidence shows an impact of psychological distance on behavioral intentions, no clear positive effect of psychological distance has been reported. Therefore, a visualization which aims to reach a diverse audience, needs to include both, psychologically more close and more distant perspectives. By combining spatially, temporally, and socially close and distant levels of abstractions, spanning from the very concrete perspective of my current daily activities to a global perspective in the future, visualizations can cater the information needs of very different people effectively.

6.1.3.3 Protection Motivation Theory (Self-efficacy)

Originally developed by Rogers [22] for health behavior, Protection Motivation Theory (PMT) explains how people respond to perceived threats. Behavior change requires two appraisals: threat appraisal (perceived severity and vulnerability) and coping appraisal (self-efficacy – believing you can take action, and response efficacy – believing your action will make a difference). For climate action, people must believe both that they can personally engage in pro-environmental behaviors AND that these behaviors will meaningfully address climate change. Research indicates that perceived self-efficacy is a relevant factor in motivating people to act on climate change [3][15]

Visualization Design Recommendation.

- Show the impact of individual actions clearly
- Provide actionable steps, not just problem descriptions; actionable steps can be elementary and feasible activities

6.1.3.4 Elaboration likelihood / narrative persuasion

The Elaboration Likelihood Model (ELM) and narrative persuasion research offer complementary accounts of how communication influences attitudes and behavior. ELM [20] posits two routes to persuasion: a central route, in which individuals engage in effortful cognitive elaboration of argument quality under conditions of high motivation and ability, and a peripheral route, in which attitudes are shaped by heuristic cues such as source credibility or aesthetic fluency. Attitudes formed via central processing are more stable, resistant, and predictive of behavior. Narrative persuasion theory extends this dual-process logic by demonstrating that stories influence beliefs through mechanisms such as transportation, identification with characters, and reduced counterarguing by embedding persuasive content within experiential structures rather than overt argumentation [12, 18]. Busselle and Bilandzic's multidimensional model of narrative engagement [5] further specifies cognitive, emotional, and attentional components of transportation, emphasizing that engagement operates through both reflective and experiential processing. Integrated, these frameworks suggest that persuasion in data-driven climate communication depends not only on argument strength and elaboration likelihood, but also on the degree to which information is embedded within coherent, engaging narrative structures that reduce defensiveness while maintaining epistemic credibility. Visualization-based storytelling interfaces are uniquely positioned to combine systematic evidence presentation (supporting central processing) with immersive, temporally structured simulations that evoke narrative engagement.

Visualization Design Recommendation.

- Support dual processing routes: Provide layered access to rigorous data, uncertainty information, and causal explanations (central route), while maintaining visual clarity, credibility cues, and aesthetic coherence (peripheral route).
- Embed data in narrative structure: Organize visualizations along narrative story arcs to foster transportation and reduce counterarguing.

- Enable identification through personalization: Allow users to simulate "someone like me" scenarios or input personal data to increase self-relevance and elaboration.
- Facilitate mental simulation: Use interactive scenario modeling to promote imagery and experiential processing, enhancing narrative engagement.
- Strengthen perceived efficacy: Visualize individual-to-collective scaling effects to support durable attitude – behavior alignment.

6.1.3.5 Persuasive technologies

Persuasive technologies have been increasingly discussed in recent years (Mintz et al.). They have been applied in areas such as health, sustainability (e.g., Ming-Chuan Chiu et al.) or in educational contexts. Such technologies are supposed to motivate users to take a specific course of action that is seen as beneficial (e.g., reducing weight, increasing sports activities, reducing energy consumption). Persuasion can result from different sources, either from careful presentation of information or from a design of the message addressing the emotion of the users without aiming at a deep processing of the message. It has been argued that presentation of information on ist own is often not sufficient to convince users (e.g., van Ham et al. 2018). Fogg (2002) has developed several design principles for persuasive technologies:

Visualization Design Recommendation.
- Reinforcement: reward desireable target behavior
- Reduction: make complex tasks simpler
- Self monitoring: users can monitor their behavior
- Suggestion: interventions should be shown at the most appropriate moment
- Surveillance: a person's behavior is monitored by other people
- Tailoring: information is adapted to the specific needs of the users
- Tunneling: tasks should be sequenced to reduce cognitive load

6.1.3.6 Social norms

People are strongly influenced by what they believe others do (descriptive norms) and what others approve of (injunctive norms). The Focus Theory of Normative Conduct [6] distinguishes these two types: descriptive norms provide information about what is typical behavior, while injunctive norms convey social approval or disapproval. Social norms function as powerful levers for behavior change because humans have deep-seated needs for belonging and social acceptance.

Visualization Design Recommendation.
- Show what "people like you" are doing (descriptive norms)
- Include social comparison feedback (e.g., neighbor comparisons in energy bills)
- Leverage group identity and belonging
- Design for social sharing and conversation
- Combine descriptive and injunctive norms for maximum effect

6.1.3.7 Gamification

Gamification is defined as the application of design principles originating from a gaming context to a nongaming context. It has been argued that gamification is especially relevant for changing attitudes and behaviors concerning climate change to a more environmentally friendly stance because it promotes motivation and engagement [8]. Several different theories from social psychology or sociology have been suggested as a foundation for research into game design and gamification (e.g., self-determination theory, flow theory, experiential learning, constructivist learning) [17]. From an overview of the literature, various design principles have been derived [17]. Some of those are especially relevant for our investigation:

It should be mentioned that some of the theories also mentioned in this paper have also been discussed as a foundation for gamification (e.g., self-efficacy).

Visualization Design Recommendation.

- Positive reinforcement. Users should get feedback on their accomplishments within the system to raise self-efficacy. Gamification uses, for example, points, levels or progress bars to achieve this.
- Social comparison. The possibility to compare one's own achievements with the achievements of others is seen as a strong incentive for learning, attitude change and behavioral change.
- Social norming. Participation in group activities and the development of common goals can help to develop normative beliefs towards behavior change.

6.1.3.8 Mental accounting

Mental Accounting Theory [26, 27] describes how individuals cognitively organize, evaluate, and track (not only) economic outcomes by assigning them to subjective "accounts" rather than treating money and value as fully fungible. Instead of optimizing globally across all resources, people compartmentalize gains and losses (e.g., "fuel budget," "vacation fund," "moral credits"), apply reference points, and evaluate outcomes relative to those mental categories. Core mechanisms include hedonic framing (segregating gains, integrating losses), loss aversion within accounts, and the influence of labeling on perceived value. Importantly, mental accounts are not restricted to financial resources; they extend to time, effort, and even moral or environmental "credits," thereby shaping how individuals perceive trade-offs and justify behavior.

In the context of climate-friendly decision-making, individuals may maintain implicit "carbon accounts" or moral balance sheets (e.g., licensing high-emission behaviors after a pro-environmental action), leading to inconsistencies between global climate goals and local behavioral choices. Consequently, how environmental impact is framed, categorized, and visually structured determines how it is encoded into these subjective accounts. Visualization-based storytelling interfaces can thus influence climate-friendly behavior not only through argument strength or narrative engagement, but by actively shaping how users cognitively partition emissions, savings, and trade-offs across behavioral domains.

Visualization Design Recommendation.

- Make carbon budgets explicit: Provide clearly bounded "carbon accounts" (e.g., monthly or annual carbon budgets) to externalize otherwise implicit accounting processes.
- Use consistent units across categories: Express impacts in standardized metrics (e.g., CO_{2e} per month) to reduce compartmentalization and improve cross-category comparability.
- Frame savings as accumulated gains: Visually segregate and highlight cumulative emission reductions to leverage preferences for gain segregation.
- Provide dynamic feedback on budget depletion: Show how individual actions "spend" from a visible carbon allowance to increase salience of marginal decisions.
- Support goal-based accounts: Allow users to set carbon-reduction targets and track progress, aligning subjective accounts with climate-consistent reference points.

6.1.3.9 Prospect theory

Prospect Theory [16] provides a descriptive account of how people make decisions under risk rather than "utility". A core insight of prospect theory is that individuals evaluate outcomes as either gains or losses relative to a reference point, rather than as final states. The associated value function is typically concave for gains and convex for losses, markedly steeper for losses than for gains, behaviour which reflects loss aversion: losses loom larger than equivalent gains. In turn probabilities are transformed into decision weights that do not correspond linearly to objective probabilities. There is a marked tendency for people to overweight small probabilities, underweight moderate to high probabilities, showing a strong certainty effect whereby outcomes framed as certain are extremely influential. These characteristics jointly explain the systematic patterns that emerge, such as risk aversion for gains, risk seeking for losses, and sensitivity to framing and representation.

Visualization Design Recommendation.

- Frame impacts relative to meaningful reference points (e.g., present-day baselines) to make gains and losses explicit rather than presenting absolute values.
- Frame impacts in terms of potential losses where appropriate, as loss aversion makes losses more behaviorally influential than equivalent gains.
- Highlight certainty versus uncertainty, recognizing that outcomes framed as certain carry disproportionate weight (certainty effect).
- Highlight low-probability but high-impact risks (e.g., extreme events), as small probabilities tend to be over-weighted in decision-making.
- Provide clear probability and uncertainty information in perceptually interpretable formats (e.g., frequencies, ranges, distributions).
- Use diverging visual encodings centered on a neutral baseline to clearly distinguish gains from losses and reinforce reference-point effects.
- Support reflective decision-making by carefully structuring visual decompositions (e.g., separating guaranteed and uncertain components), avoiding unintended framing biases while aligning with intended behavioral goals.

During the discussions within and beyond our group during the Dagstuhl Seminar, we became aware that these theories are only a selection of possibly relevant theories. When one aims to develop public visualization-based communication for a wider impact, theories from very different psychological subfields become relevant: social psychology, behavioral psychology, media psychology, cognitive psychology, … . Only by combining these different perspectives and theories, we can gain a clearer picture, why some forms of visualization-based public communication might have an impact for the wider good – or not. At the same time, inspecting these theories more closely also shows new challenges and open questions for visualization design within this context.

6.1.3.10 Key Challenges & Open Questions

We have identified the following four grand challenges, together with open questions:

1. Making an impact:
   a. How can we encourage people to change their behavior (e.g. act more climate-friendly)? This matters because: climate change is one of the biggest problems that humanity and our planet are currently facing. Difficult because: producing behavior change is known to be hard.

   b. Engaging emotion: rather than merely focusing on the cognitive aspects of communicating information, how can we influence affective responses? *This matters because:* Behavior change interventions are more likely to be successful when involving emotion [24]. *Difficult because:* the climate crisis is an overwhelming topic, with a tendency for numbness, and also involves societal polarization.
   c. Social impact – how can we go beyond individual behavior change? *This matters because:* we strive for broader adoption of climate-friendly behavior, and would like to explore and take advantage of phenomena related to peer group and (perceived) social norms.
2. What is the role of interindividual differences in visual communication for the masses (e.g. visualization literacy, attitudes)? *This matters because:* people are different in their responses to communication material. *Difficult because:* we would like communication projects to cater to the members of not only one, but ideally to all of these different groups.
3. Communicating data:
   a. How can we help people better understand and reason about large and unfamiliar magnitudes (e.g. carbon emissions)? *This matters because:* understanding such magnitudes play an important role in understanding climate-crisis-related issues. *Difficult because:* numeracy, and quantitative reasoning skills, especially regarding large and unfamiliar magnitudes, are challenging for many people.
   b. How can we balance simplicity and complexity? *This matters because:* we want to neither oversimplify nor overwhelm. *Difficult because:* these are each other's opposites, and finding a balance involves difficult choices.
   c. How can we make people interact with data? *This matters because:* interacting with data engages people and enhances understanding. *Difficult because:* people tend to prefer to just watch and browse.
4. How can we evaluate whether a visualization (project) accomplishes its communication goal? *This matters because:* this is the only way of finding out what works, what works better, and what doesn't work. *Difficult because:* communication goals may include attitude change, behavioral intentions, and behavioral change, and because actual behavioral change is very hard to find out about and measure.

The above challenges require interdisciplinary work, because they involve aspects of psychology, sociology, visualization design, and domain knowledge regarding the climate crisis.

6.1.4 Methods, Measures, and Infrastructure

This section presents the visualization prototype developed during the Dagstuhl Seminar and articulates the design rationale grounded in the psychological theories discussed in Section 6.1.3. The prototype serves as a proof-of-concept for psychology-enabled climate communication visualization, demonstrating how theoretical insights can be translated into actionable design features.

Visualization Prototype and Design Rationale

The prototype visualization system was designed to address the key challenges identified in Section 6.1.3.10, particularly the need to make climate action feel achievable at the individual level while demonstrating collective impact potential. The design integrates multiple psychological principles into a cohesive user experience organized around two primary interface components: Action Cards and Detailed Action Views.

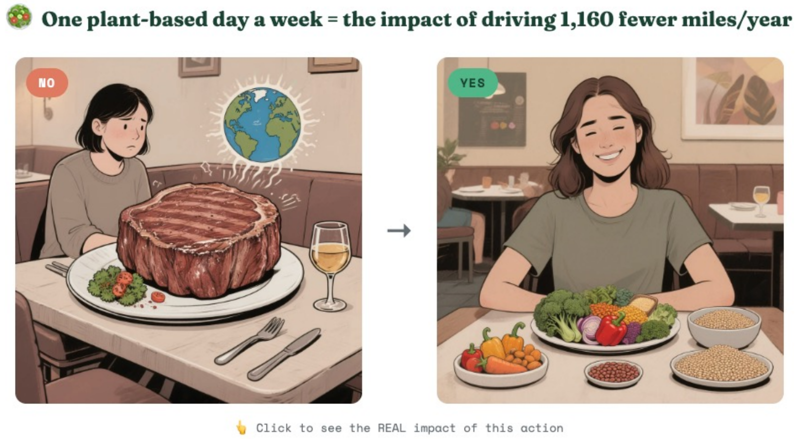

Figure 5 Food Action Card – "NO" vs "YES" framing presents the choice between meat-heavy and plant-based meals, with headline quantifying impact in relatable terms (driving miles equivalent).

6.1.4.1 Action Cards: Behavioral Choice Framing

The Action Cards (Figures 5 and 6) present behavioral choices through relatable everyday scenarios, using a "NO" versus "YES" visual framing to contrast current behaviors with sustainable alternatives. This design choice draws on several psychological principles:

- Social Norms and Descriptive Framing: Rather than presenting information as prescriptive mandates, the cards show what "people like you" could choose to do, leveraging descriptive social norms [6]. The visual juxtaposition implies that others are already making these choices, normalizing sustainable behavior.
- Psychological Distance Reduction: By depicting familiar domestic scenes – a person at a dinner table, someone doing laundry – the visualizations reduce psychological distance [29]. The scenarios are temporally immediate (today's meal, this week's laundry), spatially proximate (home environment), and socially relatable (individual perspective).
- Construal Level Bridging: The headlines bridge concrete actions to abstract impact ("One plant-based day a week = the impact of driving 1,160 fewer miles/year"), connecting immediate behavior to meaningful environmental outcomes without overwhelming users with global-scale data.
- Positive Reinforcement: The "YES" option consistently shows positive emotional states (smiling figures, bright colors, pleasant environments), applying principles from persuasive technology about rewarding desirable target behavior visually.

6.1.4.2 Detailed Action Views: Impact Visualization and Social Scaling

When users click on an Action Card, they access a Detailed Action View (Figures 7 and 8) that provides deeper engagement with the behavioral choice. These views operationalize several additional psychological principles:

- Self-Efficacy Enhancement: Following Protection Motivation Theory (Rogers, 1983), the interface provides "Pro Tips" that offer concrete, achievable first steps ("Start with ONE meal swap per week"). This addresses coping appraisal by demonstrating that individual action is feasible and manageable.
- Response Efficacy Visualization: The "Your Impact Over Time" chart demonstrates that individual actions produce measurable outcomes, directly addressing response efficacy. The logarithmic scale allows visualization of impact from individual level ("Just You") through peer group ("You + 10 Friends") to global scale ("If Everyone Did It").

Figure 6 Energy Action Card – Visual contrast between tumble dryer use and line drying, with playful headline ("The sun is literally free") applying humor to increase engagement.

Table 3 Mapping of design elements to psychological theories and communication challenges.

Design Element	Psychological Basis	Challenge Addressed
NO/YES visual framing	Social norms, Prospect theory	Making sustainable choice salient
Relatable scenarios	Construal Level Theory	Reducing psychological distance
Impact equivalencies	Mental accounting	Understanding large magnitudes
Pro Tips	Protection Motivation Theory	Building self-efficacy
Multi-scale impact chart	Social comparison, Gamification	Demonstrating response efficacy
Quantified outcomes	Persuasive technology	Making impact tangible

- Social Comparison and Gamification: The three-tier impact visualization (individual, friends, everyone) incorporates social comparison mechanisms from gamification research [8]. Users can see their personal contribution while understanding how collective action amplifies impact.
- Concrete Impact Metrics: Bottom summary cards present specific, quantified outcomes (200kg CO_2/year for individual food choices, 400 kWh/year for individual energy choices), making abstract environmental benefits tangible and memorable.

Table 3 summarizes how each design element maps to underlying psychological theories and the specific communication challenges addressed.

6.1.5 Infrastructure Needs and Shared Resources

Advancing this subfield requires development of shared infrastructure and standardized resources. We identify the following priorities:

1. Open Visualization Testbed: A modular, web-based platform enabling researchers to deploy visualization variants, collect interaction data, and conduct A/B tests with diverse populations. The current prototype can serve as a foundation for such infrastructure.
2. Standardized Measurement Battery: Validated instruments for measuring visualization impact on climate-related attitudes, intentions, and behaviors, enabling cross-study comparison and meta-analysis.

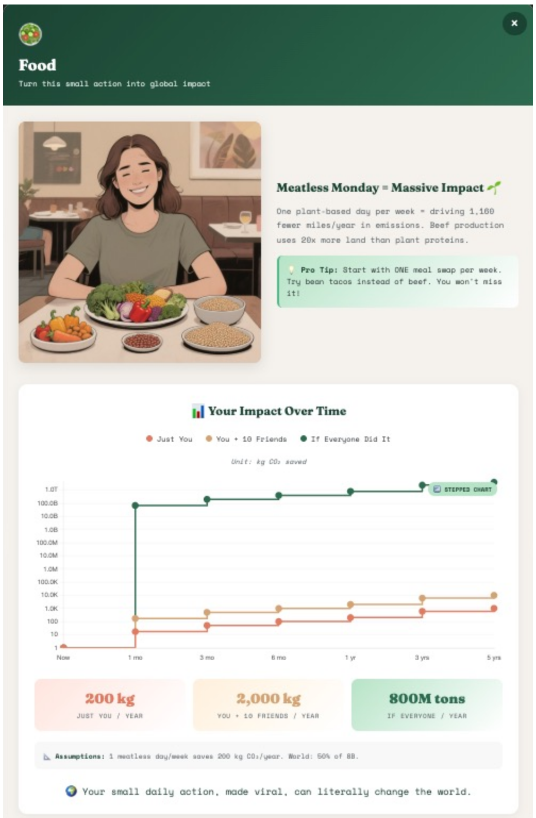

Figure 7 Food Detailed Action View – Shows "Meatless Monday = Massive Impact" with Pro Tip, impact trajectory chart, and quantified savings (200kg CO_2 individual, 800M tons global potential).

3. Behavior Change Benchmark Dataset: Longitudinal dataset linking visualization exposure to measured behavioral outcomes, providing baselines for evaluating new designs.
4. Design Pattern Library: Repository of empirically-tested visualization patterns with effect sizes and moderating conditions, organized by psychological mechanism and target behavior.

The visualization prototypes and methodological framework presented in this section provide a foundation for systematic empirical investigation of psychology-enabled climate communication. By grounding design decisions in established psychological theory and testing specific alternatives through controlled studies, we can build an evidence base that advances both visualization science and our understanding of effective climate communication.

6.1.6 Opportunities for Collaboration

Generalizing from our concrete use case of climate-related visualization to other domains, in which visual communication plays a role, we see the need of a close interdisciplinary collaboration for the successful design of visual communication for the public. At least three different disciplines should contribute for this purpose (see figure 9): (1) The application domain (in our case, climate scientists) have the data necessary and possible messages based on relevant insights. (2) Visualization designers sketch and develop visualizations and a public communication design. (3) Psychology has a large pool of theories and models, as well as methods, which can inform the design of visualizations, public communication and

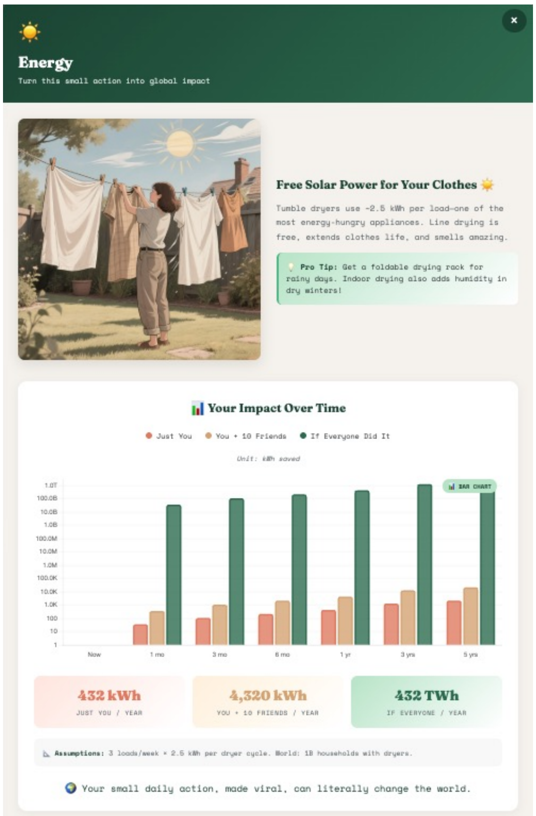

Figure 8 Energy Detailed Action View – "Free Solar Power for Your Clothes" with practical Pro Tip about drying racks, impact chart, and energy savings metrics (400 kWh individual, 1.6 TWh global).

user testing. These three parties have to team up to create the best possible communication design and have a social impact: The visualization and communication design has to be well-grounded within the domain to be trustful and create an impact. Psychological theories can inform the design to ease comprehension by the general public and increase the social impact. In addition, psychology can help to align user studies with the communication aims and go beyond standard evaluation procedures like usability and UX. At the same time, the results of user studies can inform and improve psychological theories and models, as well as visualization design guidelines.

6.1.7 Roadmap & Next Steps

The breakout group will follow up on the activities created during this seminar week in three different ways:

1. We will further develop the design recommendations towards a **full paper**, which could be submitted to a journal or conference. Within 2026, two events would especially be of interest for this topic, the Envirvis workshop at the Eurovis conference in Nottingham (Website, Deadline: March 3rd), and the Visualizing Climate Conference in Bologna (Website, Deadline April 30th).
2. There is an intention to further develop and expand the **prototype**. Ways forward that are being considered include adoption of the prototype from Higher Education Institutions as a platform to both test and tailor the tool.

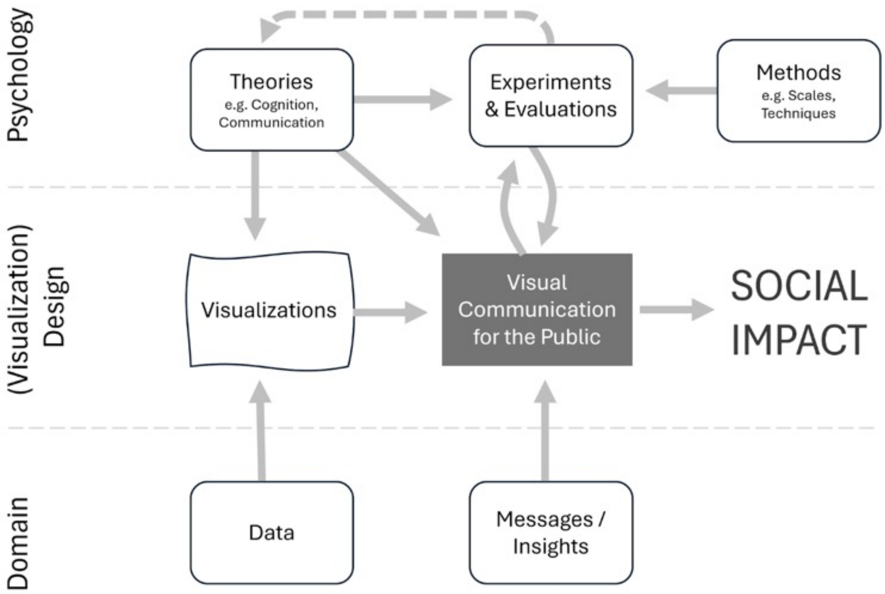

Figure 9 Interdisciplinary collaboration between domain experts, visualization design, and psychology in the design of visual communication for the general public.

3. To develop the first prototype into a wider psychology-informed platform, we envision to submit a proposal for a bigger **research project**, which would allow us to collect empirical evidence on the design recommendations grounded in psychological theories. An option within the near future could be the European call HORIZON-CL5-2026-07-D1-04: *Fighting disinformation and effectively communicating on climate change.*

We see great potential in the proposed, psychologically-grounded climate visualization approach for the expected impact, but would need to reach out to other partners to build up a competitive project consortium and proposal.

6.1.8 Acknowledgments

We would like to thank all participants of the Dagstuhl Seminar for fruitful discussions, a special thanks goes to Katerina Batziakoudi, Sara Irina Fabrikant, Kuno Kurzhals, Jan Rummel, and Barbara Tversky for their input.

References

1 Icek Ajzen. The theory of planned behavior. Organizational Behavior and Human Decision Processes. Volume 50, Issue 2, 1991, Pages 179-211, ISSN 0749-5978, doi.org: 10.1016/0749-5978(91)90020-T.
2 Benjamin Bach, Moritz Stefaner, Jeremy Boy, Steven Drucker, Lyn Bartram, Jo Wood, Paolo Ciuccarelli, Yuri Engelhardt, Ulrike Koeppen, and Barbara Tversky. Narrative design patterns for data-driven storytelling. In *Data-driven storytelling*, pages 107–133. AK Peters/CRC Press, 2018.
3 Claudia Baldwin, Gary Pickering, and Gillian Dale. Knowledge and self-efficacy of youth to take action on climate change. *Environmental Education Research*, 29(11):1597–1616, 2023.

4 Rita Borgo, Alfie Abdul-Rahman, Farhan Mohamed, Philip W. Grant, Irene Reppa, Luciano Floridi, and Min Chen. An empirical study on using visual embellishments in visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2759–2768, 2012.
5 Rick Busselle and Helena Bilandzic. Fictionality and perceived realism in experiencing stories: A model of narrative comprehension and engagement. *Communication theory*, 18(2):255–280, 2008.
6 Robert B Cialdini, Carl A Kallgren, and Raymond R Reno. A focus theory of normative conduct: A theoretical refinement and reevaluation of the role of norms in human behavior. In *Advances in experimental social psychology*, volume 24, pages 201–234. Elsevier, 1991.
7 Anastasia Denisova. *Effective Climate Communication: Turning Eco-anxiety Into Eco-action*. Springer Nature, 2025.
8 Benjamin D Douglas and Markus Brauer. Gamification to prevent climate change: a review of games and apps for sustainability. *Current opinion in psychology*, 42:89–94, 2021.
9 Ran Duan. Visualizing climate change: The interplay of construal level, threat perception, and coping efficacy. *Science Communication*, doi: 10.1177/1075547025136315, 2025.
10 Ran Duan and Christian Bombara. Visualizing climate change: the role of construal level, emotional valence, and visual literacy. *Climatic Change*, 170(1):1, 2022.
11 Yu Fu and John Stasko. More than data stories: Broadening the role of visualization in contemporary journalism. *IEEE Transactions on Visualization and Computer Graphics*, 30(8):5240–5259, 2023.
12 Melanie C Green and Timothy C Brock. In the mind's eye: Transportation-imagery model of narrative persuasion. In *Narrative impact*, pages 315–341. Psychology Press, 2003.
13 Alice Grishchenko and Mauro Martino. A network of science: 150 years of nature papers. *Nature Videos*, doi:10.1038/d41586-019-03325-6
14 Jordan Harold, Irene Lorenzoni, Thomas F Shipley, and Kenny R Coventry. Cognitive and psychological science insights to improve climate change data visualization. *Nature Climate Change*, 6(12):1080–1089, 2016.
15 Seth Heald. Climate silence, moral disengagement, and self-efficacy: How albert bandura's theories inform our climate-change predicament. *Environment: Science and Policy for Sustainable Development*, 59(6):4–15, 2017.
16 Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
17 Jeanine Krath and Harald FO Von Korflesch. Designing gamification and persuasive systems: A systematic literature review. *GamiFIN*, pages 100–109, 2021.
18 Emily Moyer-Gusé and Julia Wilson. Eudaimonic entertainment overcoming resistance: an update and expansion of narrative persuasion models. *Human Communication Research*, 50(2):208–217, 2024.
19 Raul Niño Zambrano and Yuri Engelhardt. Diagrams for the masses: Raising public awareness–from neurath to gapminder and google earth. In *International conference on Theory and application of diagrams*, pages 282–292. Springer, 2008.
20 Richard E Petty and John T Cacioppo. The elaboration likelihood model of persuasion. In *Advances in experimental social psychology*, volume 19, pages 123–205. Elsevier, 1986.
21 Margit Pohl, Markus Rester, Peter Judmaier, and K Stöckelmary. Ecodesign – design and evaluation of an e-learning system for vocational training. *e & i Elektrotechnik und Informationstechnik*, 122(12):473–476, 2005.
22 Steven Prentice-Dunn and Ronald W Rogers. Protection motivation theory and preventive health: Beyond the health belief model. *Health education research*, 1(3):153–161, 1986.
23 Nathalie Henry Riche, Christophe Hurter, Nicholas Diakopoulos, and Sheelagh Carpendale. *Data-driven storytelling*. CRC Press, 2018.

24 Ronald W Rogers, John Cacioppo, Richard Petty. Cognitive and physiological processes in fear appeals and attitude change: A revised theory of protection motivation. In *Social Psychophysiology: A Sourcebook*, 153-177. John Wiley & Sons, 1983.
25 Christina Stoiber, Markus Wagner, Florian Grassinger, Margit Pohl, Holger Stitz, Marc Streit, Benjamin Potzmann, and Wolfgang Aigner. Visualization onboarding grounded in educational theories. In *Visualization psychology*, pages 139–164. Springer, 2023.
26 Richard Thaler. Toward a positive theory of consumer choice. *Journal of economic behavior & organization*, 1(1):39–60, 1980.
27 Richard H Thaler. Mental accounting matters. *Journal of Behavioral decision making*, 12(3):183–206, 1999.
28 Chao Tong, Richard Roberts, Rita Borgo, Sean Walton, Robert S Laramee, Kodzo Wegba, Aidong Lu, Yun Wang, Huamin Qu, Qiong Luo, et al. Storytelling and visualization: An extended survey. *Information*, 9(3):65, 2018.
29 Yaacov Trope and Nira Liberman. Construal-level theory of psychological distance. *Psychological review*, 117(2):440, 2010.
30 Florian Windhager, Günther Schreder, and Eva Mayr. On inconvenient images: Exploring the design space of engaging climate change visualizations for public audiences. In *EnvirVis@EuroVis*, pages 1–8, 2019.
31 Dmytro Zagorulko, Kateryna Horska, and Nataliia Zhelikhovska. (un) necessary interaction: Audience perceptions of interactivity in digital media. *Journalism and Media*, 6(3):153, 2025.

7 Perceptual Illusions and Visualization

7.1 Illusion Delusion

Laura Matzen (Sandia National Laboratories – Albuquerque, lematze@sandia.gov)
Brian Rogers (University of Oxford, GB, brian.rogers@psy.ox.ac.uk)
Alvitta Ottley (Washington University in St. Louis, US, alvitta@wustl.edu)

Visual illusions are frequently used in visualization pedagogy to illustrate how perception and interpretation can diverge from underlying graphical encodings. Canonical examples – often drawn from Gestalt principles and related perceptual phenomena – highlight a central insight shared across visualization and vision science: what is presented is not necessarily what is perceived or understood, and design choices can both exploit and expose these discrepancies. Despite their ubiquity in teaching and discourse, "illusion" remains undefined in visualization research. This lack of conceptual clarity limits our ability to reason about when such effects might hinder interpretation, when they can be productively leveraged, and how they generalize across contexts. We propose a research agenda for studying illusions in data visualization through a psychological lens. We focus on (1) why these effects arise, grounding them in perceptual and cognitive mechanisms shaped by interaction with a three-dimensional world; (2) how they manifest in visualization contexts, including their prevalence and impact on interpretation; and (3) when they may be beneficial, misleading, or context-dependent.

7.1.1 Introduction & Motivation

Visualization research is fundamentally concerned with how people interpret visual representations of data. Yet, a persistent and often underexamined reality is that interpretation does not arise directly from the graphical encoding itself, but from the interaction between that encoding and the human perceptual and cognitive system. In this sense, every visualization is an implicit model of how people see, attend, and reason. When that model is incomplete or incorrect, systematic discrepancies can emerge between what a visualization encodes and what a viewer perceives or understands.

In visualization, examples of these discrepancies or "visual illusions" are typically used anecdotally and pedagogically to caution against misleading designs, to illustrate perceptual limits, or to motivate best practices [5, 2, 1, 4]. However, beyond these illustrative uses, illusions have not been systematically studied as a core phenomenon within visualization. They remain anecdotal, loosely defined, and disconnected from a broader theoretical framework.

This gap is consequential. Without a principled understanding of visualization illusions, we lack the vocabulary and tools to explain why certain designs consistently mislead, why some encodings appear more salient or trustworthy than they are, and why interpretation varies across contexts and populations. More critically, we risk framing these effects purely as failures or errors to be avoided, rather than as predictable outcomes of perceptual and cognitive processes that can, in some cases, be intentionally leveraged.

At the same time, this gap reflects a missed opportunity for deeper integration with vision science and psychology. While these fields have long studied perceptual illusions as a window into the mechanisms of human perception, their findings are rarely translated into the context of abstract, data-driven visual representations. Visualization operates in a space that is neither fully naturalistic nor fully artificial: it repurposes perceptual systems evolved for navigating a three-dimensional world to interpret symbolic, often compressed representations of data. This mismatch creates conditions under which illusion-like effects are not exceptions, but expected outcomes.

We argue that "illusion" should not be treated as a peripheral curiosity, but as a central lens for understanding visualization. Doing so enables us to reframe a range of persistent challenges: misleading or biased interpretations, overconfidence in graphical representations, inconsistencies in how visual encodings are perceived, and the limits of generalizable design guidelines. It also opens new opportunities to study when perceptual biases can be beneficial, for example, by enhancing salience, supporting rapid judgment, or guiding attention in complex displays.

7.1.2 Scope and Definitions

A central challenge in studying visualization illusions is definitional. Neither visualization research nor vision science offers a single, agreed-upon meaning of **illusion**. Across disciplines, the term has been used to describe a range of phenomena, including:

- **Perceptual discrepancies**, where what is seen differs from measurable stimulus properties (e.g., length, area, color, or position)
- **Interpretive discrepancies**, where visual encodings are perceived accurately but lead to incorrect or biased inferences
- **Contextual and framing effects**, where surrounding elements, comparisons, or expectations systematically influence perception or judgment

This variability reflects deeper differences in how fields conceptualize "error." In vision science, illusions are often treated as failures of perception relative to physical reality. In visualization, however, there is no single ground truth for perception; there is only an intended mapping between data, representation, and interpretation. As a result, discrepancies may arise not only from perceptual mechanisms but also from design choices, task demands, and viewer expectations.

7.1.2.1 Scope

We focus on illusion-like effects that emerge from the interaction between **visual encoding**, **human perception and cognition**, and **context of use**. Within this scope, we include:
- Effects rooted in visual perception (e.g., grouping, contrast, salience, spatial judgment)
- Effects arising during interpretation (e.g., misestimation, bias, or overgeneralization from visualized data)
- Effects shaped by context, including framing, comparison sets, and prior expectations
- Both detrimental and beneficial effects, including those that mislead as well as those that support efficient or robust reasoning

7.1.2.2 Out of Scope

We explicitly distinguish visualization illusions from several related, but separate, phenomena:
- **Data errors or distortions**: inaccuracies introduced by flawed data collection, processing, or transformation
- **Deceptive or manipulative design**: intentional attempts to mislead through selective omission, exaggeration, or rhetorical framing (though such designs may exploit illusion-like effects)
- **Arbitrary misunderstandings**: idiosyncratic or non-systematic misinterpretations that do not generalize across viewers
- **Purely semantic or domain knowledge errors**: misinterpretations arising from lack of subject-matter expertise rather than visual representation

These distinctions are not absolute. In practice, many real-world visualization failures arise from the interaction of multiple factors (e.g., a misleading design that leverages perceptual biases). However, isolating illusion-like effects allows for more precise analysis of the perceptual and cognitive mechanisms involved.

7.1.2.3 Working Glossary

To support consistency, we use the following terms throughout:
- **Encoded properties**: the visual attributes mapped from data (e.g., position, length, color, size)
- **Intended interpretation**: the meaning or inference the visualization is designed to communicate
- **Perception**: early-stage visual processing of graphical features (e.g., detecting length, grouping, contrast)
- **Interpretation**: higher-level reasoning about what the visualization implies (e.g., trends, comparisons, causality)
- **Illusion (visualization)**: a systematic divergence between encoded or intended properties and perceived or interpreted outcomes

7.1.3 Relevant Theories & Empirical Foundations

Human vision evolved to support interaction in a three-dimensional world organized around a ground plane, where depth cues such as perspective, occlusion, shading, and size constancy enable efficient navigation and object recognition [6]. These mechanisms are highly adaptive in natural settings, but they can produce systematic biases when applied to flat, abstract representations.

Figure 10 Magritte, René. *The Promenades of Euclid*, 1955, Minneapolis Institute of Art.

Classic perceptual illusions provide clear evidence of this mismatch. Consider the Promenades of Euclid by René Magritte, shown in Figure 10. This painting uses the Ponzo illusion, in which converging lines on a flat surface appear to represent the ground plane receding into the distance. In this example, converging lines are used in two different ways, creating an engaging visual illusion.

Phenomena such as the Ponzo illusion and size – distance scaling reveal that the visual system interprets two-dimensional inputs through assumptions about three-dimensional structure. Artists have long exploited these mechanisms to create compelling depictions of depth and ambiguity on flat surfaces, highlighting both the power and the limits of perceptual inference.

Data visualizations operate in a similarly constrained visual space. They rely on simplified graphical marks (e.g., lines, bars, points, and areas) rendered on a two-dimensional plane to represent complex, often high-dimensional data. In this sense, visualizations can be understood as *impoverished stimuli*: they omit many of the cues available in natural scenes while still engaging perceptual mechanisms tuned for those environments. Thus, the same processes that support efficient perception in the real world may introduce systematic distortions in visualization contexts.

For example, perspective cues, intentionally or unintentionally embedded in a visualization, can induce depth interpretations that were not part of the intended encoding. Figure 11 shows an example of a population pyramid that might be perceived as a Ponzo illusion. Although the bars are meant to appear as a stack in a vertical plane, they give the impression of a road receding into the distance. It is not clear whether this effect would have a negative impact on viewers' comprehension of the data in the figure. For example, would viewers tend to overestimate the magnitude of the smaller bars because they interpret them as being farther away?

More broadly, visual features such as area, angle, and color are not perceived veridically; they are subject to known biases in estimation, comparison, and grouping.

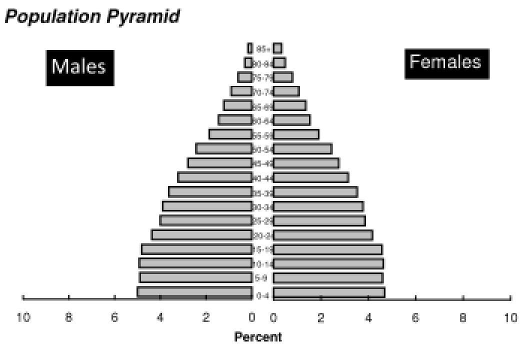

Figure 11 Example of a population pyramid visualization that could be perceived as a Ponzo illusion. Image obtained from the MASSVIS dataset [3].

Visualization research has repeatedly shown that judgments of position along a common scale are more accurate than judgments of area, angle, or thickness. What remains underdeveloped is an integrative account that links perceptual and cognitive mechanisms to the design and use of visualizations, particularly when perceptual mechanisms introduce visual effects that are not explicitly encoded in the data mapping.

The result is a critical gap: we lack an integrative account that explains how perceptual and cognitive mechanisms give rise to illusion-like effects in visualization, how these effects interact with higher-level reasoning, and when they meaningfully impact understanding. Without this connection, design guidelines risk remaining descriptive rather than explanatory, and opportunities to intentionally leverage perceptual biases remain underexplored.

7.1.4 Key Challenges & Open Questions

Advancing a research agenda on visualization illusions requires addressing several foundational and interrelated questions. These questions are not only empirical, but conceptual and methodological, and they demand coordination across visualization, psychology, and related disciplines.

1. **How prevalent are illusion-like effects in visualization?** Are such effects rare edge cases, or are they pervasive across commonly used chart types and design practices?
2. **Which designs, encodings, and contexts are most susceptible, and under what conditions?** How do chart type, visual encoding (e.g., position, area, color), layout, task, and framing interact to amplify or mitigate illusion effects?
3. **What are the measurable impacts of visualization illusions?** Do these effects influence not only accuracy, but also response time, cognitive load, confidence, trust, and downstream decision-making?
4. **For whom do these effects matter, and how do they vary across populations?** How do factors such as visualization literacy, domain expertise, cultural background, and task goals shape susceptibility or resilience to illusion-like effects?
5. **Can illusion-like effects be beneficial or even desirable?** Are there cases where perceptual or interpretive discrepancies support better outcomes – for example, by guiding attention, increasing engagement, or improving memorability?

7.1.5 Methods, Measures, and Infrastructure

Addressing these challenges will require a combination of methods, including controlled experiments measuring accuracy, response time, and cognitive load; qualitative studies examining interpretation and explanation; and longitudinal or ecologically valid studies in educational and applied settings. There is also an opportunity to curate and analyze large-scale corpora, and to automatically detect or predict illusion-prone configurations.

Systematic Audits of Visualization Practice. To understand prevalence, the field needs large-scale audits of real-world visualizations across domains such as journalism, science, policy, and education. These audits should identify recurring patterns where illusion-like effects are likely to arise, linking specific design choices (e.g., encodings, layouts, annotations) to known perceptual and cognitive mechanisms. Such efforts would parallel corpus-based approaches in other areas of visualization and could help establish benchmarks for how frequently and in what forms these effects occur in practice.

Controlled Experimental Paradigms. There is a need for carefully designed experiments that isolate illusion-like effects while preserving ecological validity. This includes adapting classic paradigms from vision science (e.g., comparative judgment tasks, threshold estimation) to visualization contexts, as well as designing new tasks that capture interpretation, reasoning, and decision-making. Importantly, these studies should move beyond binary correctness to measure response time, confidence, cognitive load, and attention patterns.

Integrative Theoretical Models. A key gap is the lack of models that connect perceptual mechanisms to visualization outcomes. Developing such models will require synthesizing insights from vision science (e.g., cue integration, perceptual scaling), cognitive psychology (e.g., heuristics, biases, mental models), and visualization theory (e.g., encoding effectiveness, task taxonomies). The goal is not only to describe when illusions occur, but to explain why they arise and to predict their impact under different conditions.

Individual Differences and Cross-Context Studies. Future work must account for variability across users and contexts. This includes examining how visualization literacy, expertise, cultural background, and prior knowledge influence susceptibility to illusion-like effects. Cross-cultural and in-the-wild studies are particularly important for assessing generalizability, as many existing findings are based on narrow participant populations and controlled settings.

Design Interventions and Trade-off Exploration. Research should explore design strategies that mitigate harmful effects while potentially leveraging beneficial ones. This includes testing alternative encodings, interaction techniques, and annotation strategies, as well as explicitly studying trade-offs between accuracy, efficiency, engagement, and memorability. Such work can inform more nuanced and context-sensitive design guidance.

Tooling and Computational Support. There is an opportunity to develop tools that help designers detect, diagnose, and reason about potential illusion-like effects. This could include automated linting systems, simulation tools that model perceptual biases, or design environments that surface risks and trade-offs in real time. These tools would bridge research and practice, making theoretical insights actionable for a broader community.

Shared Datasets, Benchmarks, and Taxonomies. Progress in this area will benefit from shared resources, including curated datasets of visualization examples, standardized experimental tasks, and evolving taxonomies of illusion types. While this paper does not propose a fixed taxonomy, iterative efforts toward shared vocabularies and benchmarks will be essential for cumulative knowledge building and reproducibility.

7.1.6 Opportunities for Collaboration

This subfield naturally invites collaboration between Visualization, Vision Science, Cognitive Psychology, Learning Science, Artists, Designers, Educators, and Visualization Practitioners. Joint efforts could focus on shared experimental paradigms, pedagogical materials, or open repositories of illustrative examples and effects. Another opportunity for collaboration would be exploring richer, more naturalistic representations of data that would have stronger alignment with the strengths of human visual perception.

We highlight several concrete opportunities for collaboration:

- **Visualization+Psychology Shared Experimental Paradigms.** Visualization researchers and psychologists can co-develop experimental frameworks that bridge controlled perceptual studies with ecologically valid visualization tasks. This includes adapting classic illusion paradigms to data representations, as well as designing new studies that capture interpretation, reasoning, and decision-making in realistic settings.
- **Pedagogical Integration.** Educators across visualization, psychology, and design can collaborate to develop teaching materials that move beyond isolated examples toward a more principled understanding of perception and interpretation. This includes interactive demonstrations, curricula, and assessment tools that help learners recognize when and why illusion-like effects occur.
- **Design and Artistic Exploration.** Artists and designers have long explored perceptual ambiguity, depth, and visual deception. Collaborations with these communities can surface new classes of effects, inspire novel visualization techniques, and challenge assumptions about clarity, accuracy, and aesthetics. Such partnerships can also help reframe illusions not only as problems, but as expressive and exploratory design elements.
- **Tooling and Practice-Oriented Research.** Engagement with visualization practitioners creates opportunities to translate theory into practice. This includes developing design guidelines, linting tools, and interactive systems that help identify and reason about illusion-like effects during the design process. Practitioners, in turn, can provide insight into real-world constraints, trade-offs, and recurring failure modes.

7.1.7 Roadmap & Next Steps

We envision a progression from defining the phenomenon, to measuring it, to ultimately shaping how visualization is designed and taught. However, realizing this research agenda will require coordinated, staged efforts that build from shared foundations to empirical understanding and, ultimately, integration into practice.

Next 1 year: Establish Foundations

- **Consolidate terminology and scope.** Develop a shared vocabulary and working definitions to enable consistent communication across disciplines.
- **Curate example sets and corpora.** Compile annotated collections of visualization examples that exhibit illusion-like effects, including those commonly used in pedagogy and real-world practice.
- **Seed the community.** Organize workshops, panels, or special interest groups to bring together researchers from visualization, psychology, and design.
- **Pilot studies and replications.** Conduct small-scale experiments to validate key effects and adapt classic perceptual paradigms to visualization contexts.

Next 3 years: Build Evidence and Theory

- **Systematic empirical studies.** Investigate illusion-like effects across chart types, encodings, tasks, and diverse user populations, including cross-cultural and in-the-wild settings.
- **Develop integrative models.** Formulate and test theoretical accounts that link perceptual and cognitive mechanisms to visualization outcomes.
- **Establish benchmarks and datasets.** Create shared experimental tasks, stimulus sets, and evaluation protocols to support reproducibility and comparison across studies.
- **Map design spaces and trade-offs.** Identify when illusion-like effects are detrimental, negligible, or beneficial, and how they interact with factors such as task and context.

Next 5 years: Integrate and Translate

- **Embed in education and training.** Incorporate findings into visualization curricula, literacy assessments, and practitioner-facing resources.
- **Inform design guidance.** Move beyond heuristic "best practices" toward evidence-based, context-sensitive recommendations.
- **Develop computational tools.** Build systems for detecting, simulating, and reasoning about illusion-prone designs, including integration into design workflows.
- **Expand to broader contexts.** Extend insights to emerging domains such as immersive analytics, AI-assisted visualization, and interactive systems.

References

1 Terek S Amer and Sury Ravindran. The effect of visual illusions on the graphical display of information. *Journal of information systems*, 24(1):23–42, 2010.
2 TS Amer. Bias due to visual illusion in the graphical presentation of accounting information. *Journal of Information Systems*, 19(1):1–18, 2005.
3 Michelle A Borkin, Azalea A Vo, Zoya Bylinskii, Phillip Isola, Shashank Sunkavalli, Aude Oliva, and Hanspeter Pfister. What makes a visualization memorable? *IEEE transactions on visualization and computer graphics*, 19(12):2306–2315, 2013.
4 Steven L Franconeri, Lace M Padilla, Priti Shah, Jeffrey M Zacks, and Jessica Hullman. The science of visual data communication: What works. *Psychological Science in the public interest*, 22(3):110–161, 2021.
5 Christophe Hurter, Claus-Christian Carbon, Mauro Martino, and Bernice E. Rogowitz. Art, Visual Illusions, and Data Visualization (Dagstuhl Seminar 24301). *Dagstuhl Reports*, 14(7):81–114, 2025.
6 Brian Rogers. *Perception: A very short introduction.* OXFORD University press, 2017.

8 Artificial Intelligence, Psychology, and Visualization

8.1 Assessing and Enhancing AI's Capability in Visualization

David Ebert (University of Arizona – Tucson, US, ebertd@arizona.edu)
Steffen Frey (University of Groningen, NL, s.d.frey@rug.nl)
Wolfgang Jentner (University of Arizona – Tucson, US, wjentner@arizona.edu)
Stephen Kobourov (TU München – Heilbronn, DE, stephen.kobourov@tum.de)
Harald Schupp (Universität Konstanz, DE, harald.schupp@uni-konstanz.de)
Laura Pelchmann (Universität Köln, DE, laura.pelchmann@uni-koeln.de)
Tatiana von Landesberger (Universität Köln, DE, tlandesb@uni-koeln.de)

The integration of large language models (LLMs), visual language models (VLMs), and Agentic AI into everyday workflows underscores the importance of visualization in data analysis and communication. However, the suitability of visualizations for AI remains largely unexplored. This work aims to assess the similarities and differences between AI models and human responses to visualizations, examining how various models interpret visual inputs and where systematic errors or biases emerge.

Key challenges and open questions, such as replicating human subject studies with LLMs, identifying visual deficiencies in LLMs, and making AI aware of visual metaphors, abstraction, and analogy, are addressed. Potential mitigation strategies for various visualization types are discussed. The report concludes with a roadmap for future experiments on visual channels, optical illusions, Gestalt principles, and bias experiments, emphasizing the need for further research to enhance the accuracy and applicability of AI in visualization tasks.

8.1.1 Introduction & Motivation

Visualization is increasingly integral to the functionality of large language models (LLMs), visual language models (VLMs), and Agentic AI, influencing various professional domains such as data science, health, engineering, business intelligence, education, and software development. The reliance on AI to generate, interpret, and explain visualizations for tasks like data analysis and decision support underscores the need to assess the suitability of visualizations for artificial intelligence [4]. Although psychology has extensively studied human perception and processing of visual information, the extent to which these empirical regularities apply to machine perception remains underexplored.

The primary research goal is to evaluate the similarities and differences between AI models and human responses to visualizations. This involves examining how different models interpret visual input, their alignment with human perception, cognition, and reasoning, and identifying systematic errors or biases. Quantitative evaluations on synthetic benchmark datasets [1, 7] and real-world visualizations are proposed to measure model accuracy across various visual artifacts, data types, and tasks. Our work will also explore the potential for AI-based systems to detect and mitigate visual artifacts, focusing on high-quality model detection accuracy and independence from data and task specifics.

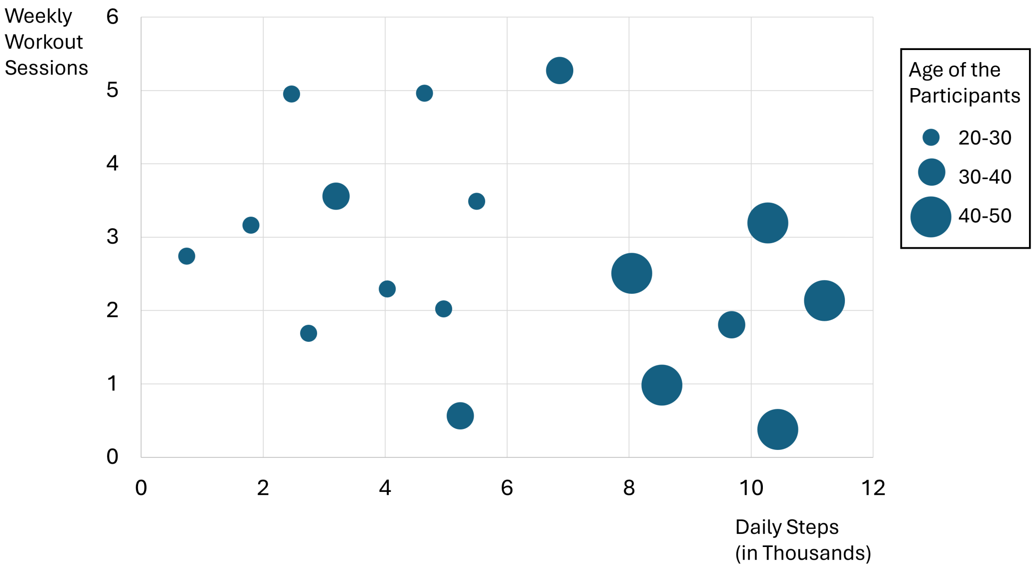

Figure 12 Shows a scatterplot demonstrating the Ebbinghaus illusion [6] in a real-life data visualization. In this example, identically sized mid-sized points appear larger or smaller depending on the size of the surrounding points, illustrating how contextual elements influence perceived object size.

8.1.2 Scope & Definitions

The scope of this research encompasses various types of architectures and models, with the aim of understanding their capabilities and limitations in handling visual artifacts. The investigation extends to the development of AI-based detectors and mitigation strategies for these artifacts. Key strategies include adding labels and annotations, changing visual representations and orientation, and reassembling visual elements to improve human readability. The success of these strategies is measured by the reduction of detected artifacts in adapted representations.

8.1.3 Relevant Theories & Empirical Foundations

Understanding the baseline for human behavior and visual channels is crucial [2, 8]. Psychological theories provide the foundational knowledge necessary for perceptual grounding in visualization theories. Our work emphasizes the importance of systematic manipulation and evaluation of visual illusions [3, 9, 10, 5] to create more realistic information visualization scenarios. Geometric illusions that lead to biases in area, angle, and length are particularly relevant, as they can and indeed do occur real-life examples of bubble charts, treemaps, node-link diagrams, and multi-series line charts (see Figure 12).

8.1.4 Key Challenges & Open Questions

Several open questions guide this research.

- When replicating human-subject studies with LLMs, what are the similarities and differences?
- In what visual tasks do LLMs outperform humans, and what are the visual deficiencies of the LLMs?
- Can AI-based methods be made aware of visual metaphors, abstraction, and analogy?

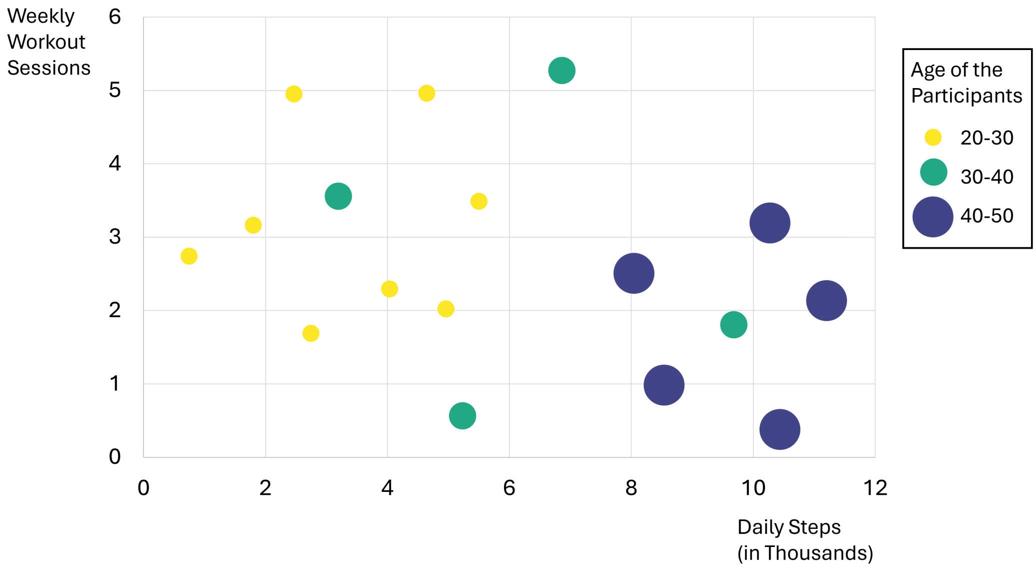

Figure 13 Shows a scatterplot mitigating the Ebbinghaus illusion through color differentiation. In addition to varying point sizes, points are colored to reduce the perceptual bias, making the actual sizes more discernible and highlighting the importance of using multiple visual cues in data representation.

Our work explores whether LLMs can meaningfully serve as human subjects in visualization experiments and if perceptual cognition-aware LLMs can be developed to detect and repair visual deficiencies, ultimately enhancing human-readable visualizations (see Figure 13).

8.1.5 Methods, Measures, and Infrastructure

The study adopts a sequential structure, beginning with identification of limitations in both humans and AI. This lays the groundwork for fine-tuning AI models to enhance their accuracy and developing mitigation strategies for specific visual representations. Potential mitigation strategies include changing layouts, changing visual appearance, and adding labels and annotations for various visualization types.

8.1.6 Roadmap & Next Steps

The immediate next steps involve running experiments on visual channels and adapting experiments from existing literature to repeat with LLMs. Our work references key studies on perceptual orderability of visual channels and the evaluation of LLMs' visualization literacy, underscoring the importance of continuous research and collaboration in this evolving field.

References

1 Nan Chen, Yuge Zhang, Jiahang Xu, Kan Ren, and Yuqing Yang. Viseval: A benchmark for data visualization in the era of large language models. *IEEE Transactions on Visualization and Computer Graphics*, 31(1):1301 – 1311, January 2025.

2 David H. S. Chung, Daniel Archambault, Rita Borgo, Darren J. Edwards, Robert S. Laramee, and Min Chen. How ordered is it? on the perceptual orderability of visual channels. *Computer Graphics Forum*, 35(3):131–140, 2016.

3 Alex Gomez-Villa, Adrián Martín, Javier Vazquez-Corral, Marcelo Bertalmío, and Jesús Malo. On the synthesis of visual illusions using deep generative models. *Journal of Vision*, 22(8):2–2, 07 2022.
4 Jiayi Hong, Christian Seto, Arlen Fan, and Ross Maciejewski. Do llms have visualization literacy? an evaluation on modified visualizations to test generalization in data interpretation. *IEEE Transactions on Visualization and Computer Graphics*, 31(10):7004–7018, 2025.
5 Christophe Hurter, Claus-Christian Carbon, Mauro Martino, and Bernice E. Rogowitz. Art, Visual Illusions, and Data Visualization (Dagstuhl Seminar 24301). *Dagstuhl Reports*, 14(7):81–114, 2025.
6 Ted Jaeger and Kyle Klahs. The ebbinghaus illusion: New contextual effects and theoretical considerations. *Perceptual and Motor Skills*, 120:177–182, 02 2015.
7 Yuyu Luo, Nan Tang, Guoliang Li, Chengliang Chai, Wenbo Li, and Xuedi Qin. Synthesizing natural language to visualization (nl2vis) benchmarks from nl2sql benchmarks. In *Proceedings of the 2021 International Conference on Management of Data*, SIGMOD '21, page 1235 – 1247, New York, NY, USA, 2021. Association for Computing Machinery.
8 Caitlyn M. McColeman, Fumeng Yang, Timothy F. Brady, and Steven Franconeri. Rethinking the ranks of visual channels. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):707 – 717, January 2022.
9 Jianyi Yang, Junyi Ye, Ankan Dash, and Guiling Wang. Illusions in humans and ai: How visual perception aligns and diverges, 2025.
10 Yichi Zhang, Jiayi Pan, Yuchen Zhou, Rui Pan, and Joyce Chai. Grounding visual illusions in language: Do vision-language models perceive illusions like humans? In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5718–5728, Singapore, December 2023. Association for Computational Linguistics.

9 Concluding Remarks

The contributions in this report outline the contours of the emerging field of **Visualization Psychology**. This field is defined by its integrative perspective, linking perceptual and cognitive theory with the design and use of visualizations, while also using visualization as a lens for studying cognition in realistic, task-driven settings. Each chapter highlights complementary facets of this space, from foundational theory-building and conceptual alignment, to learning and friction, decision-making, storytelling, perceptual effects, and human – AI interaction.

A key insight across these efforts is that visualization effectiveness must account for a broader range of cognitive and behavioral outcomes, including confidence, trust, engagement, and action. At the same time, variability across individuals, cultures, and contexts challenges assumptions of universality and calls for more inclusive and ecologically valid approaches. These insights point to the need for new theories, methods, and evaluation frameworks that reflect the full complexity of human-centered data interaction.

The establishment of Visualization Psychology as a coherent field represents both a challenge and an opportunity. It requires sustained interdisciplinary collaboration, shared infrastructure, and a commitment to integrating knowledge across traditionally separate domains. However, it also offers the potential to fundamentally reshape how we design, study, and deploy visualizations – enabling systems that are not only more effective, but also more aligned with the ways people perceive, think, and act in an increasingly data-driven world.

Participants

- Alfie Abdul-Rahman
King's College London, GB
- Katerina Batziakoudi
INRIA Saclay – Orsay, FR
- Peter Baudains
King's College London, GB
- Leslie Blaha
AFOSR – Arlington, US
- Rita Borgo
King's College London, GB
- Anne-Flore Cabouat
INRIA Saclay – Île-de-France, FR
- Min Chen
University of Oxford, GB
- Evanthia Dimara
Utrecht University, NL
- David S. Ebert
University of Arizona – Tucson, US
- Mennatallah El-Assady
ETH Zürich, CH
- Yuri Engelhardt
University of Twente – Enschede, NL
- Sara Irina Fabrikant
Universität Zürich, CH
- Brian D. Fisher
Simon Fraser University – Surrey, CA
- Amy Rae Fox
University of California – San Diego, US
- Steffen Frey
University of Groningen, NL
- Lane T Harrison
Worcester Polytechnic Institute, US
- Petra Isenberg
INRIA Saclay – Orsay, FR
- Wolfgang Jentner
University of Arizona – Tucson, US
- Stephen Kobourov
TU München – Heilbronn, DE
- Kuno Kurzhals
Universität Stuttgart, DE
- David H. Laidlaw
Brown University – Providence, US
- Nadine Malcolm
Quantellia – Denver, US
- Mauro Martino
MIT-IBM Watson AI Lab – Cambridge, US
- Laura Matzen
Sandia National Labs – Albuquerque, US
- Eva Mayr
Donau-Universität Krems, AT
- Andrew Meso
King's College – London, GB
- Carolina Nobre
University of Toronto, CA
- Alvitta Ottley
Washington University – St. Louis, US
- Laura Pelchmann
Universität Köln, DE
- Margit Pohl
TU Wien, AT
- Lorien Pratt
Quantellia – Denver, US
- Brian Rogers
University of Oxford, GB
- Jan Rummel
Universität Heidelberg, DE
- Filip Sadlo
Universität Heidelberg, DE
- Antonia Schlieder
Universität Heidelberg, DE
- Karen Schloss
University of Wisconsin – Madison, US
- Harald Schupp
Universität Konstanz, DE
- Barbara Tversky
New York, US
- Tatiana von Landesberger
Universität Köln, DE
- Florian Windhager
Donau-Universität Krems, AT
- B. L. William Wong
Auckland University of Technology, NZ
- Michael Wybrow
Monash University – Clayton, AU