Chapter 9
Visualization Psychology: Foundations for an Interdisciplinary Research Program

Amy Rae Fox and James D. Hollan

Abstract What might a discipline of Visualization Psychology look like? If research on the psychological aspects of visualization were to coalesce, in the sense of a Lakatosian research program, what refutation-resistant theoretical commitments would magnetize its "hard core"? In this chapter, we argue that any interdisciplinary inquiry concerned with psychological aspects of visualization should situate its phenomena in the broader context of external representation, as a (triadic) semiotic activity achieved via information processing in a distributed cognitive system.

9.1 Introduction

Our goal in this chapter is *not* to provide a grand unified theory of visualization, nor to review all relevant work in the social and behavioral sciences. Rather, we offer a conceptual framework: a series of theoretical premises we argue should form the foundation of any interdisciplinary inquiry concerned with psychological aspects of visualization. We start by addressing the virtue of a hypothetical Visualization Psychology, arguing that the phenomenon of visualization is a fertile laboratory for exploring human cognition, that engineering and design-driven research can be improved via appropriate grounding in theories of perception and cognition, and that well-structured collaborations across disciplinary boundaries can foster a virtuous cycle beneficial to both traditions of research. In Sect. 9.3, we argue that such inquiry should situate visualization in the broader context of external representation (Sect. 9.3.1) as a (triadic) semiotic activity (Sect. 9.3.2) involving information processing (Sect. 9.3.3) in a distributed cognitive system (Sect. 9.3.4). In Sect. 9.4, we illustrate how this framework can be applied in both empirical and

A. R. Fox (✉) · J. D. Hollan
University of California San Diego, La Jolla, CA, USA

D. Albers Szafir et al. (eds.), *Visualization Psychology*,
https://doi.org/10.1007/978-3-031-34738-2_9

theoretical contexts, before concluding with a discussion of the role of Psychology in the history (and future) of Visualization.

9.2 Why Visualization Needs Psychology

The first workshop on Visualization Psychology was held in conjunction with the IEEE VIS conference in 2020, with the following Call for Papers (CFP) [71]:

> Before 2010, each VIS conference typically featured 0–2 papers on empirical studies. The VisWeek 2010 in Salt Lake City became a turning point, and since then more and more empirical study papers have been presented at VIS. Between 2016 and 2019, there were some 60 empirical study papers in VIS/TVCG tracks. Many young talents who are knowledgeable in both VIS and psychology emerged in the VIS community, while many colleagues in psychology are authoring and co-authoring such papers and attending VIS conferences. It is therefore timely to ask the two communities: is there a need for Visualization Psychology as a new interdisciplinary subject?
>
> There are many branches of applied psychology, such as clinical psychology, counseling psychology, educational psychology, forensic psychology, health psychology, industrial–organizational psychology, legal psychology, media psychology, music psychology, occupational psychology, sports psychology, and so on. Almost all of these are widely recognized academic subjects and have their own conferences and journals. Since interactive visualization and visual analytics encompass most human-centric processes in data science and real-world data intelligence workflows, many will argue for the necessity and feasibility for developing Visualization Psychology in a coherent and organized manner.
>
> This is the first workshop that will enable the experts in VIS and psychology to define the scope of this new subject of Visualization Psychology collectively and stimulate new research directions and activities in both fields. The goals of the workshop are:
>
> - To provide researchers in VIS with a significant platform to develop their theories and experiments in addition to acquiring knowledge from psychology
> - To broaden the scope of empirical research in VIS to involve more cognitive aspects in addition to considering visualization a vision or perception problem
> - To enable researchers in psychology to explore VIS as a rich playground and carry out research beyond the existing molds

(continued)

- To enable the development of the young talents in VIS and psychology through the development of a new interdisciplinary subject and the platforms for research communication, publications, and collaboration

This CFP solicits an intersection between two communities: "Psychology" and "VIS." In this context we can pragmatically identify the VIS community as scholars affiliated with publishing venues such as the VIS conference series[1] and the journal IEEE Transactions on Visualization and Computer Graphics. This is a community intellectually and institutionally grounded in the discipline of Computer Science, closely related to (for others, a subset of) Human–Computer Interaction (HCI). The term "visualization research" then is used to refer to work in the VIS community, historically centered in engineering and design perspectives: developing tools to solve problems.

Psychology is a much older, more expansive discipline of science, dating back to the mid-nineteenth century. Let us say for the sake of argument that our goal is to engage scholars of psychology already publishing in VIS venues and others whose work is sufficiently relevant to visualization phenomena. For this, we might constrain "the psychology community" to be scholars of the cognitive, perceptual, or educational branches of experimental psychology, as well as vision science, learning science and cognitive science, whose phenomena of interest include human interaction with visual-spatial representations of information.[2] Psychological research in this sense will include work published in venues outside VIS (such as journals and conferences of the Cognitive Science Society, Psychonomic Society, Association for Psychological Science, and International Society for the Learning Sciences, among others). For brevity, we use the term psychologist as a placeholder for members of this more diverse disciplinary milieu.

What might the goals of this new interdisciplinary community be? The following claims are made explicit in the VisPsych CFP and offer a first approximation of what a Visualization Psychology might hope to accomplish:

1. Visualization research should be informed by psychological theories.
2. Visualization research should emphasize cognitive as well as perceptual factors.
3. Visualization phenomena offer a rich playground for further developing psychological theory.

[1] Self-identified as the "premier forum for advances in visualization and visual analytics," VIS is sponsored by the IEEE (The Institute of Electrical and Electronics Engineers) Computer Society and Technical Committee (special interest group) on Visualization and Graphics (TCVG).

[2] One might also find research detailing interaction with graphics in other applied branches of Psychology—the use of multimedia graphics in the courtroom, for example—however the theories, models, and frameworks governing the basic science of such occurrences would likely come from cognitive, educational, or perceptual psychology.

The Role of Psychological Theory We suggest that the first point is true by virtue of epistemic relevance: the explanatory power and design impact of visualization research is improved when grounded in psychological theory, just as human interaction with a computer is better explained by theories of human psychology than formalisms governing the algorithms of the machine. For brevity, we use the term *psychological* as an umbrella for human aspects of interaction with visualizations; for example, how a reader perceives, forms a judgment from, or solves a problem with a visualization. This is in contrast to non-psychological questions, such as defining the algorithm for transforming a set of data into a particular representational form or how that computational system is engineered to afford input/output interaction. The latter questions may be required to enable the visualization phenomenon but neither necessitate nor explain human interaction with it. In this way research in visualization is like research in human–computer interaction. Psychological theories are needed to inform the design and evaluation of computational systems and to understand the dynamics of human interaction with them, but so too are contributions from the formal/mathematical science and engineering of computing. This is to say that VIS need not be subsumed *into* Psychology. Like HCI, visualization is a rich theoretical and empirical subject matter for interdisciplinary collaboration.

Situating Perception and Cognition The claim that empirical research in visualization should include cognitive in addition to perceptual theory is also trivially true, insofar as we are concerned with "cognitive" phenomena or behavior (i.e., beyond perceptual judgments). This is a question of levels of analysis and scope of phenomena. More often than not, empirical research in VIS (particularly investigations that center on the efficacy of some type of visualization or interactions with a visualization system) should be concerned with *cognitive* rather than *perceptual* phenomena. Accepting that the theoretical boundaries between perception and cognition are fuzzy, if we adopt an information-processing perspective from mainstream Cognitive Science, we can reasonably construe perception as some subset of cognition, concerned with stimulus-driven behavior, while the term cognition implies "higher order" processing, the influence of prior knowledge, or "what one does with" perceptual input. An empirical study with a task operationalized to measure constructs approximating perceptual processing is likely aimed at building and testing basic theory in perception, rather than evaluating the efficacy of a particular visualization.

This point is exemplified by the widespread misapplication of classic graphical perception studies by Bell Labs statisticians William Cleveland and Robert McGill (see [18–20]). When presented as stimulus a simplified statistical graph (e.g., a pie or divided bar chart, each with two segments marked with a dot), experimental subjects were asked to indicate "what percentage the smaller is of the larger": a *perceptual judgment*. The accuracy of subjects' responses (with respect to mathematical ground-truth) was evaluated and used to derive a ranking of relative accuracy for graphical encodings. From these results, one could conclude it is more effective to represent the quantitative difference between two values as a bar chart, rather than a pie chart. Unfortunately, this work has been generalized by some to the design heuristic, "bar charts are better than pie charts." If humans were perceptual computers with no prior knowledge, expertise, beliefs, or other individual differences, that *might* be the end of the story. However, a body of research in *graph comprehension* has demonstrated that if you use a different task in your study, for example, asking the graph reader to extract a specific value from the graph, to use it to make a decision, or perhaps a forecast, then the accuracy rankings do not necessarily hold (e.g., [63, 64]). This apparent contradiction arises from the insight that different task-demands require different "readings" of a graph: a (perceptual) judgment of relative size is different than extracting a data value which is different from detecting a trend, and so on. The more complex the behavioral task, the more "higher order" (i.e., resource-intensive, implicating prior knowledge) processing is required. While it would be appropriate to apply perceptual accuracy heuristics to design, for example, a simple graphic in a newspaper illustrating the quantitative month-over-month change in some economic report, it would be insufficient to rely solely on these heuristics to guide design of an interactive visual analytics system. Perceptual guidance for achieving simpler tasks is a useful starting point but does not encompass knowledge-driven interactions. Basic research on graph comprehension has clearly demonstrated that the effectiveness of graphical encodings arises not from the interaction of data and forms, but rather, data, forms, *individuals*, and *tasks*. This is not to say that perceptual processing is not relevant to complex cognitive activities or that there are no perceptual questions left to be answered. One of the most challenging, and in our view promising, areas ripe for theoretical development is along these fuzzy boundaries: exploring the factors that govern how stimulus-driven and knowledge-driven processes are integrated to produce behavior.

Visualization and the Virtuous Cycle The claim that visualization phenomena offer opportunities for advancing psychological research can be demonstrated from evidence. Grammars and frameworks (especially [9, 54, 74]) designed by Computer Scientists and implemented as libraries and interactive systems have made computer-based data visualizations accessible for researchers as tools for data analysis and presentation. For those whose research involves empirical study of human–information interaction, these are also tools for generating stimuli. The situations in which the stimuli might be used—for example, studying how a graph is used to make a decision, how a student leverages a chart and accompanying text to learn a concept, or how an analyst uses an interactive system to make a forecast—are all enabled by technologies borne of Computer Science-based visualization and computer graphics research. In turn, research on the human aspects of how and why and to what effect individuals interact with visualizations provides guidance for the appropriate design of visualization systems. Technology inspires new human activity, which offers the psychologist new subjects of inquiry. As is often the case with technology-driven endeavors, Psychology and Computer Science stand in relationship as a virtuous cycle: a positive feedback loop where progress in each stands to both improve in quality and volume the progress of the other. The relationship between Psychology and VIS is the relationship between Psychology and broader Human–Computer Interaction—appropriate considering that VIS grew from and is largely considered a part of HCI. There is, in fact, so much "psychology" *in* HCI and VIS and it is challenging to know where (and if) we should draw meaningful boundaries. We return to this issue in Sect. 9.5.

9.3 Elements of a Framework

> Lakatos's idea is to construct a methodology of science, and with it a demarcation criterion, whose precepts are more in accordance with scientific practice. (...) Instead of an individual falsifiable theory which ought to be rejected as soon as it is refuted, we have a sequence of falsifiable theories characterized by a shared hard core of central theses that are deemed irrefutable—or, at least, refutation-resistant—by methodological fiat. This sequence of theories constitutes a research program.—Musgrave and Pigden [44]

One way to conceptualize the structure of a Visualization Psychology is in terms of a *research program* in the tradition of post-positivist philosopher of science of Imre Lakatos [76]. Lakatos was skeptical of Kuhn's normative conception of science as progressing via successive stages where one research paradigm (i.e., a framework for approaching one's subject matter) is replaced by another. Lakatos characterized the practice of science as altogether messier, with multiple competing paradigms operating in parallel, in nonlinear cycles of progression (making theoretical and empirical progress) and degeneration (stagnating, and/or questioning core claims). For Lakatos, a research program was characterized not by a singular method, model, or theory, but rather a collection of basic (and by convention irrefutable) assumptions shared by its community. This *hard core* of theoretical commitments is surrounded by a protective *auxiliary belt* of hypotheses that constitute the work of science. Investigators rely on the shared language and lenses of the hard core to generate hypotheses in the auxiliary belt that might be shaped into theories or broken down and replaced. Progress is made so long as the auxiliary belt grows: theoretically, by extending the scope of theory to new empirical domains, or empirically, by finding corroborating evidence for theoretical claims (see [3, 76]).

What central theses—theoretical propositions resistant to refutation—might a Visualization Psychology have at its *hard core*? We propose a minimum of four central tenets. Individually, these ideas are not falsifiable theories, but rather perspectives and frameworks that have arisen from and give rise to empirically testable hypotheses.

1. **Visualization is external representation.** Visualization (as artifact) and visualization (as process) belong to the broader class of *external representation*.
2. **Meaning is constructed.** Interacting with a visualization is not a passive transmission of meaning (e.g., "extracted" from the artifact), but rather an active, interpretive semiotic process where knowledge is constructed.

3. **Information is processed.** Visualization is most effectively construed as the transmission of information across components of a system, via transformation between representational states.
4. **Cognition is distributed.** Intelligent action with a visualization is a function of a distributed cognitive system comprised of human actors and material artifacts situated in relation to their spatio-temporal environment.

We describe these perspectives in Sects. 9.3.1–9.3.4 and in Sect. 9.4 demonstrate how they can be applied in both empirical and theoretical research settings.

9.3.1 Visualization is External Representation

The language of representation is slippery and self-referencing. Shown a collection of marks on surfaces, you might label some as art, or pictures, others as diagrams, maps, or schematics, some charts, plots, or graphs, and others also as graphs, but you might use air quotes and call them "graph-theory graphs." Some you will identify as writing and others, like writing but not—some peculiar or particular system of notation. The linguistic labels you apply to each marking likely depend on your disciplinary background and are neither exhaustive nor mutually exclusive. Which of these are visualizations? (Fig. 9.1).

9.3.1.1 *On* Visualization

Let us start with definitions put forth in prominent VIS texts. In their venerated compilation of papers and essays, Card et al. [11] define Information Visualization as *"The use of computer-supported, interactive, visual representations of abstract data to amplify cognition"* (pg. 7). Stephen Few offers a functional definition, characterizing data visualization as *"an umbrella term to cover all types of visual representations that support the exploration, examination, and communication of data. Whatever the representation, as long as it's visual, and whatever it represents, as long as it's information, this constitutes data visualization"* [24, pg.12].

Such inclusive specifications may be effective for teaching but are less suitable guides for scientific inquiry. From this heuristic, we might conclude the words on this page constitute a visualization—but they would not be considered so by most visualization practitioners. Why? Because visualizations are somehow *graphic* in nature; from Ware, *"a graphical representation of data or concepts"* [72, pg.2]. Ware highlights how the term has transitioned in conventional meaning from "constructing a visual image in the mind" to *"an external artifact supporting*

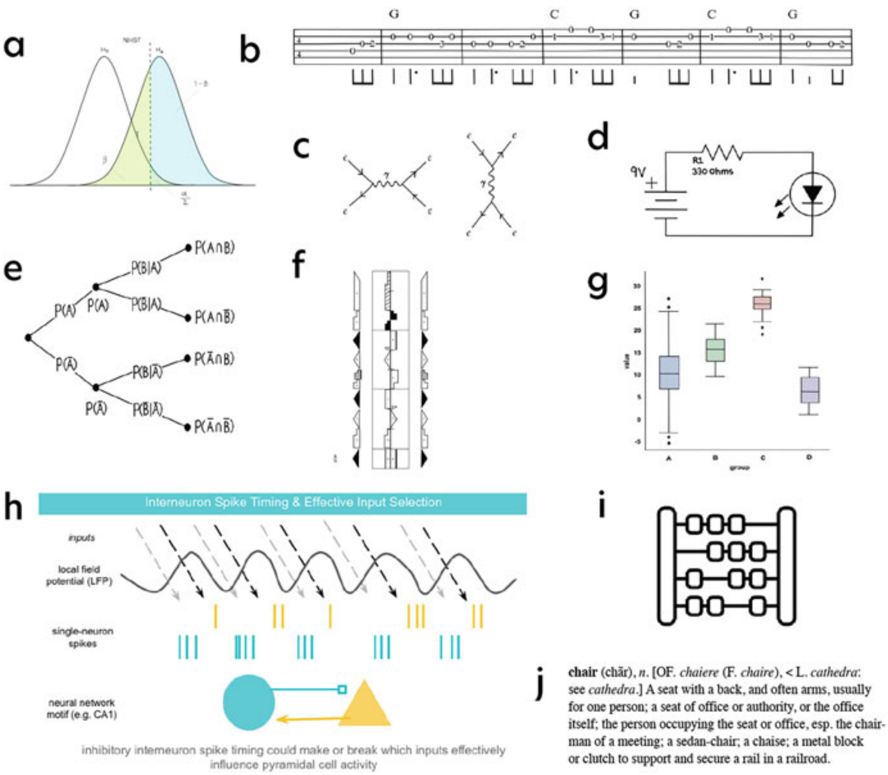

Fig. 9.1 A group of visual-spatial external representations. (**a**) a conceptual diagram indicating key concepts in null hypothesis significance testing; (**b**) portion of the song 'You Are My Sunshine' in guitar tabs notation; (**c**) Feynman diagram for an interaction between an electron and anti-electron with exchange of a photon; (**d**) schematic of a circuit depicting a 9V battery in configuration with a single resistor and LED; (**e**) tree diagram used in solving Bayesian reasoning problems; (**f**) Laban notation representing a ballet exercise; (**g**) boxplot depicting mean, interquartile range and outliers for 4 groups; (**h**) a figure from a neuroscience presentation that combines multiple representations of related phenomena to orient readers to both the research method and analysis of results; (**i**) an icon of an abacus-note that the object the icon represents would also be considered an external representation of number; (**j**) image of the words in a dictionary definition of the word chair (inspired by the conceptual art piece 'One and Three Chairs' by Joseph Kosuth)

decision making." Bertin (in English translation [5]) also refers to graphic representation, distinguishing this from relational imagery (e.g., art, photography) and mathematics (e.g., symbolic notation). In a more recent text, Munzner pragmatically characterizes the purpose of *computer-based visualization systems* as providing *"visual representations of data sets designed to help people carry out tasks more effectively"* [43, pg.1].

If we start with the assertion that a visualization (noun) is a type of representation, then to arrive at a useful definition we should characterize the properties of what representations can be admitted to this type. An externally available representation is accessed via some sensory modality; that visualizations are subject to *vision* is the only shared property of the aforementioned definitions. Similarly, we lack clarity as to whether visualizations need to be *graphic* (or indeed, what being graphic entails), if they need to be interactive, generated by a computer, whether they can refer to any information, or only "data," and perhaps only that which is deemed "abstract." Alternatively, we can try to infer a shared conceptualization of such terms based on how a community organizes itself. VIS[3] explores the sometimes fluid distinction within these properties in the organization of annual conference tracks, paper types, and sessions. Whether the referent of a visualization is abstract data or has some physical/geometric invariant is the (historical) distinction between the Information Visualization (InfoVIS) and Scientific Visualization (SciVIS) conferences. If the purpose of an artifact is to support an interactive, analytical process, then it would likely be called a visualization and fall into the Visual Analytics (VAST) conference. If the referent is more "conceptual" than data-driven, research is more likely to be published outside of VIS, such as in the (multidisciplinary) *International Conference on the Theory and Application of Diagrams*, and if the use of the representation is primarily for learning, then the research is likely evaluated in either disciplinary education (e.g., Chemistry Education, Math Education) or Learning Science. The number of paper types (and submissions) at VIS implicating computer systems, prototypes, and algorithms suggests a strong preference toward the computer as a presentation medium or "physical substrate." Though there is exciting growth in the topic of data *physicalization* and exploration of alternative sensory modalities for representing data, this area has yet to emerge as a large enough topic to warrant its own conference session in the past five years. Research on data *sonification* or *tactilization* are more likely to be found in the broader ACM SIG-CHI or topical journal like ACM Transactions on Applied Perception.

Definitions, as terminology, serve as tools for communicating and conceptualizing one's subject matter [10]. We draw on these definitions of Information and Data Visualization, not in critique of their notable contributions, but rather to call attention to a puzzling inconsistency in the foundation of the field. Our objects of inquiry are altogether over-specified and under-defined. Which of the artifacts in Fig. 9.1 are visualizations? We argue that to the visualization psychologist, it should not really matter. They are *all* instances of the larger class: external representations. Just as psycholinguists are concerned with the psychological and neurobiological factors that enable humans to acquire, use, comprehend, and produce *language* (not English, or "languages using the roman alphabet," or "languages written from left-to-right"), visualization psychologists should be concerned with the factors that enable humans to make use of *external representations* (not just the "graphic,"

[3] Referring to the annual IEEE combined conferences on Information Visualization (InfoVIS), Scientific Visualization (SciVIS), and Visual Analytics Science and Technology (VAST).

"data-driven," or "computer generated" variety). In this sense, designers and engineers of visualization systems have the luxury of specialization. But insofar as we believe that the interaction with visualization relies on general purpose cognitive mechanisms, psychologists do not. To understand how these artifacts function—to study how they are used by humans to construct meaning in support of complex cognitive activities—we must climb up the ladder of abstraction.

9.3.1.2 *On* External Representation

> The power of the unaided mind is highly overrated. Without external aids, memory, thought, and reasoning are all constrained. But human intelligence is highly flexible and adaptive, superb at inventing procedures and objects that overcome its own limits. The real powers come from devising external aids that enhance cognitive abilities. (...) It is things that make us smart.—Norman [47, pg. 43]

The term *external representation* came to prominence in the late 1970–80s, as the new discipline of Cognitive Science emerged from information-processing psychology with a common focus on the existence and nature of mental representation (see [8, 38, 46, 49]). But when the researchers focused solely on the mental, they needed unnecessarily complex mechanisms to explain behavior. Although AI and the mental imagery debate would ensure that mental representation remained a focus of mainstream Cognitive Science, the need to distinguish *internal* from *external* meant the birth of a new research area.

The complexity of external representation, however, was not immediately appreciated. In his treatise on cognitive representation, Palmer argued that mental representations were *"exceedingly complex and difficult to study,"* so one might start with the examination of *"noncognitive"*[4] representations, as they are *"simple, and easy to study"*[5] [48, pg. 262]. Subsequent elaboration of representational systems demonstrated there is much to explore with respect to the nature and function of such "noncognitive" structures (see [37, 55]).

Like research on visualization, however, empirical work on external representation was lacking in the explicit definition of terms. A study on problem solving with a diagram might refer to the diagram as an external representation and rely on the reader to draw the same antonymic implication as Palmer: an external representation is a representation that is *not* internal. The sensory modality, encoding media, presentation substrate, and communicate purpose are left under-specified, allowing the term to serve as a category for *things that can be perceived, that refer to other things.* Such things might be presented via any medium, in any encoding structure,

[4] Palmer reserves the qualifier *cognitive* for internal representations, designating the external as "noncognitive." Following a distributed cognitive perspective, we would characterize *both* as cognitive representations and prefer the term "mental" to describe those representations not perceivable to others.

[5] More "accessible" is perhaps the more generous characterization.

via any sensory modality, referring to anything (real or imagined), for any purpose. Zhang and Norman explicitly described external representations as *"knowledge and structure in the world, as physical symbols (e.g., written symbols, beads of abacuses, etc.) or as external rules, constraints, or relations embedded in physical configurations (e.g., spatial relations of written digits, visual and spatial layouts of diagrams, physical constraints in abacuses, etc.)*" [77, pg.3].

We refine this definition:

> *An **external representation** (noun) is the form of information, purposefully encoded as structures in material artifacts that serve a semiotic function as part of an interpretive process.*

Information is encoded externally via forms and structures that can be described along a continuum of implicit to explicit, depending on how much effort, or inference, is required in their use (see [34, 35]). We remove reference to knowledge in the world, preferring the constructivist premise that knowledge does not exist *in* the environment but is actively constructed by the individual via interaction *with* their environment. The kinds of constraints and structures described by Zhang and Norman are constituent parts of representations and of how they work. Most importantly, we clarify the scope of external representations as being constructed by some actor, for some purpose, thus grounding external representation in the context of communication—though broadly construed. (Many of the external representations we construct are meant for communication not with others, but our future selves.) Here, we admit visualization as a subset of external representation: an active construction of meaning via the exchange of information between actor and artifact. What is crucial is that we orient ourselves equally toward the artifact and the interactive process: representation as noun and representation as verb.

Despite a dearth of precise terminology in the proceeding decades, researchers[6] took up the challenge of discovering how humans *think with things*, studying how various forms of external representations (ERs) influence thinking for various ends. An early focus was the fashion in which graphic/diagrammatic ERs influence thinking in contrast to natural language, such as in problem solving [37, 77], learning [2], design [21], and scientific discovery [14]. Distinctions were drawn between encoding structures: the sentential/propositional (symbols), and graphic/diagrammatic (images), where the latter class was taken up by its own interdisciplinary community

[6] Particularly in Cognitive Science, Learning Science/Educational Psychology, and disciplinary education like Math, Chemistry, and Physics.

in the early 2000s.[7] Educational Psychologists and Learning Scientists turned their attention to multimodal and multimedia representations [41, 56]. A particularly impactful contribution was made by Michael Scaife and Yvonne Rogers in [55], wherein they proposed a "new agenda" for research on graphical representations and in considerable detail and sophistication demonstrate how such research promises to improve the design of future technologies while simultaneously advancing theories of cognition. By the late 2000s, sufficient interest across allied disciplines warranted a special issue of the journal TopiCS in Cognitive Science, dedicated to visual-spatial representations, with milestone contributions on visual analytics [25], graph comprehension [62], and diagrams [13], as well as reviews of how visual-spatial representations serve as tools for thinking [70] and corresponding implications for design [28]. These are indicative of the work we believe should be at the theoretical core of visualization psychology research.

Thus, we have moved from the study of computer-generated interactive data graphics to any externalization of thought. What we are left with, it seems, is a Goldilocks problem. The idiomatic conception of visualization is too narrow and an exhaustive conception of external representation too broad. Fortunately, there are dimensions along which this metaphorical problem space can be surveyed. We might think of these dimensions as ranges along which we can attune our attention, progressively expanding or narrowing our scope of inquiry depending on the state of theoretical and technological advancement.

On Encoding Medium Though we have noted the lack of precision in defining the scope of visualizations, there has been no lack of effort in cataloging [27] and taxonomizing them, from general descriptive frameworks [6, 15, 50, 65, 69] to those concerned with specific domains of data [1, 4, 7]. Two particularly useful (and under-appreciated) are those of Engelhardt [23] who offers an atomic, generative framework deserving of its characterization as a language of graphics and Massironi [40] who offers both a taxonomy and an evolutionary timeline. While most taxonomies deal with some intersection of graphical structure and data type (e.g., geographic maps, relational networks), the more common distinction in the cognitive and learning science literature is the continuum from descriptive to depictive, roughly analogous with symbolic to analog, or propositional to graphic. These terms refer to a semiotic modality (also medium), which indicates the degree of convention in the relation between a representation and thing to which it refers. While the poles of a depictive–descriptive continuum can be easily identified, there lays betwixt a murky medium. At what point of abstraction does an icon become a symbol? When it is no longer identifiable as its referent without convention? In whose judgment? We are more accurate in describing our scope of inquiry as multimedia than "primarily graphic." We propose that while origins of VIS as a field lie in the distinction of graphics from text, fundamental questions about

[7] The International Conference on the Theory and Application of Diagrams is a biennial gathering held since 2000, attended by a cross-section of Philosophers, Psychologists, Mathematicians, and Computer Scientists.

framing, persuasion, and even comprehension rely on understanding the function of text *alongside* graphics. It is rarely the case that external representations of the visual graphic variety are not accompanied by some form of linguistic proposition or sentential notation. Indeed, a visualization without a title and labels may be worth no words at all.

On Sensory Modality External representations can be constructed for any sensory modality, though by far the most attention has been paid to the visual. Deservedly so, as visuals are the most pervasive information artifacts, and the sensory modality about which we have the most understanding. Though we are surely far from exhausting the wellspring of questions to be asked about visual representations, we suggest that we accept within our scope multi-sensory representation. From a theoretical stance, this requires broader inclusion of expertise across perceptual psychology, though the applications are consequential. In an increasingly visualization-driven world, equality and accessibility demand informationally equivalent tools for those without visual perception. Notably, we can trace this view back to the inception of visualization in HCI:

> It should be noted that while we are emphasizing visualization, the general case is for *perceptualization.* It is just as possible to design systems for information *sonification* or *tactilization* of data as for multiple perceptualizations. Indeed, there are advantages in doing so. But vision, the sense with by far the largest bandwidth, is the obvious place to start, and it would take us too far afield to cover all the senses here.—Card et al. [11, pg.7]

On Representational Purpose or *Communicative Context* VIS texts describe the purpose of visualization as being to "amplify" cognition [11, 24, 72] though research in Cognitive Science suggests the story is more nuanced (see [36, 47]). External representations *enable* cognition and can change the nature of the task we are performing. This is not to say that one cannot think without external representations, but rather, there are certain kinds of thinking that are not possible without the right representations to think them.

The most generic purpose is to simply record: to offload from internal memory to external cognition. In terms of communication, to inform—for example, the boxplot in a manuscript, where one aims to inform the reader of some aspects of the underlying information—in a clear and simple manner.[8] But one might design that artifact differently if one intends for you to explore the data, undertake an analysis, or make a decision, a plan, or a forecast. An author might change their strategy if they want to strongly persuade you or, alternatively, want you to use the representation to learn. There are entire systems of diagrams designed for solving particular kinds of problems, and the design of representations to support conceptual change is the focus of specific subdisciplines in STEM education. We use the term *communicative context* to refer to the "cognitive activities" the designer of a representation intends the user to perform. The structure of these activities

[8] Note that clarity and simplicity *do not* imply truth. The designer of a representation has a voice that is echoed in every design decision, from what information to include to how to encode it.

has not been taxonomized, though a compelling framework for their hierarchical, emergent structure is detailed by Sedig and Parsons [60]. The relevant insight is that certain parameters of a representation, such as the computational efficiency, or relative explicitness of certain aspects of data, need to be tuned in accordance with the task the reader is expected to perform. Bertin (in English translation [5, pg. 183]) writes *“A graphic is never an end in itself; it is a moment in the processes of decision making.”* To this, we add *“… or reasoning, or learning, or problem solving, or sensemaking, or analyzing, or planning, or forecasting… ”* The graphic in the moment is thus deeply intertwined with the individual, their situation, and task contexts.

9.3.2 Meaning Is Constructed

> All meaningful phenomena (including words and images) are signs. To interpret something is to treat it as a sign. All experience is mediated by signs, and communication depends on them.—Chandler [12, pg. 23]

If external representations are things purposefully constructed to refer to other things, then understanding their referential function falls squarely within the realm of *semiotics*. Semiotics is the study of signs, where a sign is construed as “something which stands for something else”—*aliquid stat pro aliquo* [12]. Note this is a larger class of phenomena than external representations which we have (pragmatically) constrained as being purposefully constructed. Signs, conversely, can be naturally occurring: a trail of footprints in the snow or mud puddles following a heavy rain. The crux of the semiotic puzzle is that to *be* a sign, is to be *interpreted*. Phenomena become signs when meaning is assigned to them. You may have the intuition that to implicate semiotics is to open a Pandora’s box where terms like *represent* and *signify* become so complex they risk losing any consistent meaning—and you would be right.[9] Our task is to introduce the elementary constructs of a particular semiotic approach that can be productively applied to understanding the function of external representations in distributed cognitive systems.

Imagine you encounter a line graph in a newspaper. Your job as a reader is to develop an understanding (interpretant) of what the graph (the representamen) indicates about some state of the world (the referent). The terms referent, representamen, and interpretant are drawn from American philosopher Charles Sanders Pierce, and his general account of the relations that govern representation, reference, and the meaning of signs [30]. Peirce’s basic claim is that a “sign” consists of three

[9] “No treatment of semiotics can claim to be comprehensive because, in the broadest sense (as a general theory of signs), it embraces the whole field of signification, including “life, the universe, and everything,” regardless of whether the signs are goal-directed (or interpreted as being so)” [12, pg. xvi].

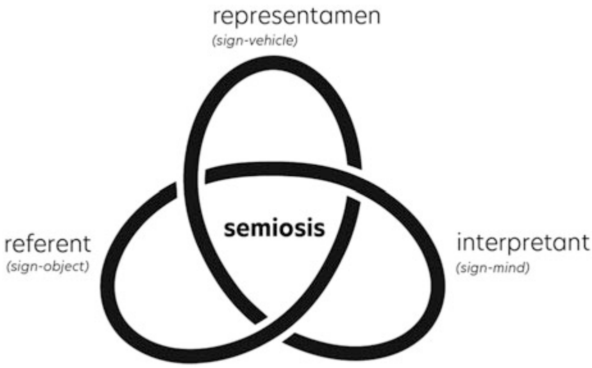

Fig. 9.2 The three components of a Peircean sign (referent, representamen, and interpretant) are irreducibly triadic

parts: (1) an object (referent) that is the thing being signified, (2) an element that signifies (representamen): that which does the referring, and (3) the interpretant: understanding that is made of the referent-representamen relation. Importantly, the entire triadic relation is referred to as a sign or representation and the dynamics of the relation semiosis or signification. Though Peirce's own terminology changed over the development of his ideas, to avoid confusion, we choose here three terms not commonly employed outside of semiotics: *referent* (also, sign-object, or signified), *representamen* (also, sign-vehicle, signifier), and *interpretant* (also, sign-mind, understanding) (see Fig. 9.2). The labels we colloquially apply to the material substances that comprise external representations—representation, sign—are in semiotic terms explicitly *not* equated with the material component of the sign. That is to say, the "representation" is not *the* representation, but only a *part* of it. The sign-relations are *irreducibly* triadic, and while we might for sake of analysis wish to isolate the relation between sign-object and sign-vehicle (for example, how a designer chooses to encode some information) or sign-vehicle and sign-mind (for example, how a reader interprets the encoding), their function is only constituted as a property of all three. This is perhaps more intuitive in psychological terms: constructing meaning is a combination of top-down (knowledge-driven) and bottom-up (stimulus-driven) interpretative processing. To examine how a reader interprets an encoding, we must consider their interaction with the encoding, and prior knowledge of the information being encoded.

Peirce's triadic semiotic is significant to the psychology of visualization in two ways. First, it makes explicit the constructive nature of meaning. Peirce's interpretant brings into the signifying function someone or something that does the interpreting: an intelligent process that constructs the translations between signifying elements of the representamen, in order to arrive at some approximation of the referent. In this way, the relation between the "thing" and the "representation" is not a direct and determined mapping, but entirely subjective, based on the interpretation of the observer. Second, Peirce's semiosis is dynamic, relying not on the entirety of that which acts as the representamen, but only on the elements relevant in signifying. Later accounts elaborate on subdivisions in the referent and interpretant that pertain to stages of processing in an unfolding chain of meaning [30]. This aspect has a distinctly cognitive appeal, as it suggests a distribution of meaning-making between the observer and environment; one that occurs via a process in time, not contained solely within artifacts or minds. In the context of cognition, together these features of Peirce's approach are consistent with what we know about the influence of prior knowledge and individual differences in the determination of meaning.

9.3.3 Information Is Processed

> "There is no information without information vehicles. Information vehicles are the carriers of information, the physical material in which the information-for-the-interpreter is encoded."—Nauta [45]

In an age of grounded, embodied, and extended cognition, it is rather fashionable to discount information-processing psychology as outdated. However, there is a difference between studying psychological phenomena *as* the processing of information and studying *the phenomenon* of information processing. The classical conception of information-processing regards the mind as a computational system manipulating symbols to enact representational states. The information-processing psychologist might seek to explain all psychological phenomena through this lens—behavior resulting from the propagation of representations, disregarding the influence of the body, modal systems, or environment. Contemporary theories that situate cognition beyond the mental are extraordinarily applicable to human interaction with external representations. But so too are some constructs from information processing. In a Visualization Psychology, we are directly concerned with how humans interact with information via representations. To the extent that we rely on the notion of information, we cannot escape the notion of its processing. Importantly, we are not proposing that to adopt an information-processing view of visualization requires commitment to a computational theory of mind, nor any strictly sentential/propositional symbol manipulation in the brain. One problem with information-processing models of cognition was that they paid "scant regard" to the external world of artifacts and information (see [53]). By exploring phenomena that require processing of multimedia (i.e., text and graphic) information, we expect that the Visualization Psychologists can improve on these theories by directly addressing the interface between external and internal information, especially in the construction of meaning.

9.3.4 Cognition Is Distributed

> It does not seem possible to account for the cognitive accomplishments of our species by reference to what is inside our heads alone. One must also consider the cognitive roles of the social and material world. But how shall we understand the relationships of the social and the material to cognitive processes that take place inside individual human actors? This is the problem that distributed cognition attempts to solve.—Hutchins [32, pg. 2071]

As behavioral scientists, we are concerned not only with the design and efficacy of external representations but also with their mechanisms: how and why they function (or not). These functions are enacted between the artifact(s) and person(s), embodied and situated in their environments and complex social structures. This complexity demands a distributed perspective of cognition, one that extends functions of the mind beyond the individual's skin and skull (see [16, 17]) and distributes them through time and space via material artifacts and members of society (see [31, 32]). Unlike traditional theories, *distributed cognition* extends the reach of what is considered "cognitive" beyond the individual to encompass interactions between people and with resources and materials in the environment.

The applicability of a distributed cognitive perspective to research in visualization [39] and human–computer interaction more broadly [29] has been successfully argued, and corresponding methods of cognitive ethnography are now widely accepted in VIS and HCI publications. Through cognitive ethnographic techniques (e.g., interviewing, participant observation, in-situ recording), a researcher can determine *what things mean* to the participants in an activity and to document *the means by which* these meanings are created. In this way, cognitive ethnography yields data for exploring cognitive mechanisms, while also feeding distributed cognitive theory by adding to the corpus of observed phenomena the theory should explain.

A distributed perspective on cognition is particularly relevant to the psychology of visualization because it not only provides an overarching framework for investigating representations and representational processes but actively encourages integration of ethnographic and experimental approaches as well. While the study of cognition *in the wild* can answer many kinds of questions about the nature of human cognition in real workplaces, the richness of real-world settings places limits on the power of observational methods. This is where well-motivated experiments are necessary. Having observed phenomena in natural settings, the researchers can set about designing more constrained experiments to systematically explore specific aspects of observed situated behaviors. Importantly, distributed cognition does *not* require that every aspect of a cognitive system be examined in every interaction: levels of analysis still apply. But a distributed cognitive perspective does require that the most highly operationalized inquiries of basic processes are contextualized as only parts of a more complex system of factors that taken together, explain behavior.

In every area of science and technology, the choices made about units of analysis have crucial consequences. Boundaries are often a matter of tradition in a field. Sometimes the traditionally assumed boundaries are exactly right for investigating a specific issue. For other phenomena, however, the boundaries either span too much or, more frequently, too little. The failure to reevaluate the unit of analysis as sciences advance and technology changes can fundamentally restrict development. A common critique of distributed cognition in psychological traditions is the *necessity* of extending the unit of analysis to the environment. From Wilson, for example, *"The fact that causal control is distributed across the situation is not sufficient justification for the claim that we must study a distributed system. Science is not ultimately about explaining the causality of any particular event. Instead, it is about understanding fundamental principles of organization and function"* [75, pg. 630]. We obviously disagree with this claim and argue that insofar as the function of the mind is to control real-time action in dynamic environments, any sufficient understanding of its organization requires theoretical and methodological approaches that directly address the environment as an active participant in cognition. Fortunately, today the lens of distributed cognition is part of an emerging zeitgeist that appreciates the central importance of closing the divide between computationally focused disciplines and disciplines concerned with understanding people and sociotechnical systems.

9.4 *On Doing* Visualization Psychology

We propose the following definition for *Visualization Psychology*:

> ***Visualization Psychology** is a scientific research program at the intersection of computing, behavioral and social sciences. It is characterized by the application of theories of perception, cognition and behavior to predict and/or explain the nature of human interaction with visualization systems, and by the use of visualization phenomena to inform theories of perception, cognition, and behavior.*

This definition emphasizes that (1) VisPsych should be a *scientific* endeavor: though it may involve close collaboration with designers and engineers, the intellectual goal of the research is generating knowledge, and (2) the flow of insight in VisPsych should be bidirectional: benefitting from and contributing to work in engineering or design-oriented aspects of visualization. Research in Visualization Psychology can contribute to the design and evaluation of visualization systems, while the design and engineering of visualization systems can provide sites of inquiry for both basic and applied psychological research. It has elements of both basic and applied science, employing basic theories to explain specific (visualization) phenomena, the outcomes of which may serve as data for (re)constructing basic theory. In this sense, much Visualization Psychology might be most accurately characterized as *use-inspired basic research*, à la Stokes [67] Pasteur's quadrant conception of a two-dimensional relation between basic and applied research. Finally, (3) we have purposefully characterized VisPsych as a *program* of research, rather than a singular discipline or distinct community.

According to Lakatos, the measure of a research program is its ability to make both theoretical and empirical progress. Theoretical progress is made by building on foundational tenets to develop theory and apply it to new empirical domains. Empirical progress is made by evaluating theory. Across the VIS and visualization-adjacent literature in psychology/learning science/cognitive science we find diverse examples of such progress.

Extending Theory to New Domains An example of theoretical progress can be found in the research agenda of Cognitive Psychologist Priti Shah, whose early research applied psychological theory from reading comprehension to the emerging topic of graph comprehension. Shah and colleagues built upon Pinker's [52] information-processing theory of graph comprehension by re-construing perception of a graph as *reading* and invoking constructs from the prevailing Construction-Integration [33] theory of text and discourse comprehension (see [61]). From this application came new testable hypotheses about the role of prior knowledge and individual differences in comprehension and the temporal dynamics of information processing in graph comprehension more broadly. Although this work was not explicitly situated in the context of distributed cognition or triadic semiotics, it is consistent with both lenses insofar as it situates graph comprehension as *discourse* between the designer and the reader of a visualization, differentially influenced by factors acting on the display, the individual, and task(s).

Evaluating Theory in New Domains A powerful example of the virtuous cycle between basic psychological and applied-visualization oriented research can be found in the recent movement to (re)connect research in visualization with vision science. While some of the earliest empirical work in visualization was concerned with visual perception (e.g., [18, 66]), modern interdisciplinary research in vision science offers both new methods (see [22]) and theoretical constructs. *Ensemble perception*, for example, refers to how the visual system extracts summary statistical information from groups of similar objects, ostensibly as a way of dealing with spatial and temporal processing constraints [73]. Szafir and colleagues have applied the theoretical arguments for ensemble perception to the domain of data visualization and argued for how it may serve as mechanism for some of the most common perceptual tasks we perform when interpreting visualizations, including identifying outliers, detecting trends, and estimating means [68]. While the development of ensemble perception as a construct is not necessarily grounded in distributed or semiotic perspectives, its application to visualization *is*: first, by providing an account for how differing interpretations can arise from the same visual stimuli (an explicit acknowledgement of the triadic nature of semiotic discourse), and second, by positioning the research as a *contribution* rather than *determinant* of design heuristics (an implicit acknowledgement that ensemble perception is a valuable piece rather than sole factor in the puzzle of visualization behavior).

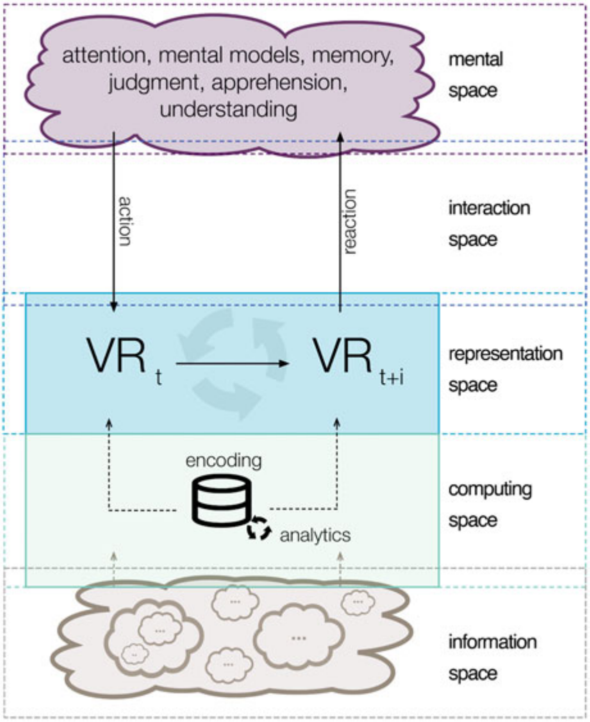

Fig. 9.3 The human–information interaction epistemic cycle (adapted from original draft with permission of author, Paul Parsons). We cannot overemphasize the importance of conceptualizing these spaces as metaphorical and not simultaneously reifying the layers as physical systems with linear exchanges of information. In practice, information processing emerges dynamically, simultaneously across the material components that constitute the system. This diagram can be construed as a snapshot of this dynamic processing, linearly unfolded in time from left to right

This contextualization is crucially important in ensuring basic theory is applied appropriately in design-driven research.

Model Building to Support Innovation An exemplar for progress that supports research in both basic and applied dimensions is the EDIFICE framework[10] developed by Sedig and Parsons. As a conceptual model, it provides a structure for thinking about the processing of information (such as goal-directed interaction with a visualization) distributed through the components of a cognitive system. In Fig. 9.3, we find five (metaphorical) spaces that together form a *human–information interaction epistemic cycle* (see [57–59]).

The *information space* consists of the set of information with which users might interact and the *computing space* its storage and manipulation (i.e., machine computation). In the *representation space*, encoded information is made available for perception. (The "space" of representation is an abstraction, but is reified in computers as "the interface.") The *interaction space* affords exchange of information via action and perception: where the interpreter performs actions and receives reactions. In the *mental space* exist the mind and mental operations that contribute to

[10] Epistemology and Design of human–InFormation Interaction in Cognitive activitiEs.

but importantly do not entirely constitute the construction of knowledge. The model is clearly grounded in the perspectives of information processing and distributed cognition. Though it was conceived in the context of interaction with complex visualization tools, its abstractions can be fruitfully applied to the wider space of multimodal and multimedia external representations. Most importantly, it makes explicit that the design of a visualization tool is a communicative act between designer and user.

The EDIFICE framework offers a productive nomenclature for designating which components of a distributed cognitive system we might be addressing in the context of a particular research project, allowing us to more accurately characterize limitations and desired integrations for future work. For example, a new visualization system that uses machine learning to recommend graph encodings would primarily involve the design of algorithms in the *computing space* and resultant productions in the *representation space.* A user-study of such a tool would involve measuring the outcome of operations in the *mental space* when an individual interacts with the application (via the *interaction space)*. Most importantly, the framework serves as tool for thinking about how the processing of information is distributed across a system of human-visualization interaction: a problem of substantial importance to designers and researchers alike. The authors have applied the framework to describe the relative distribution of information processing across machine and human actors [51], to characterize the construct of interactivity [60], and as the backbone for a pattern-language to aid conceptualization of novel visualization designs [57, 58].

9.5 The History and Future of Visualization Psychology

In *From Tool to Partner: The Evolution of Human–Computer Interaction* [26], Jonathan Grudin provides a comprehensive history of HCI. But this is not a commentary on the growth of a discipline, rather he illustrates how HCI (as a topic of study) emerged as a practice across communities in computer science, human factors, information systems, and information science. This is a telling editorial choice, revealing how entrenched institutional structures in academic disciplines interact with the moderately more pliable boundaries of professional societies to endow structural support to emerging subjects of inquiry that necessitate cross-disciplinary contribution. The cover illustration for the volume (penned by Susie Batford) can be read as deeply metaphorical. Over an undulating sea rise distinct mountain peaks, bearing the labels of various computing-related fields, including MIS (management information systems), HF (Human Factors), CS (Computer Science), and LIS (Library and Information Science). Running down and over and across the peaks, ostensibly nourishing rich research ecosystems, are bright blue rivers fed by an enormous raining cloud—labeled *psychology*.

One can imagine a similar scene for a history of Visualization. Research involving the creation, systematization and situated use of (primarily, though not entirely, graphic) visual-spatial representations of information is taking place across the sciences and humanities. Such research is enabled by *both* computing technology and theories of human behavior. By virtue of its name, the VIS community claims epistemic authority over visualization and serves as a pragmatic "home base" for technological innovation. But basic psychological theory rains upon disciplinary peaks like chemistry, physics and mathematics, education, communication studies, and even philosophy. Scholars in these disciplines are not merely *using* visualizations as tools in their work, but doing work that centers (representation) design, development, and evaluation, as well.

Through their Call for Papers, the organizers of the 2020 VisPsych workshop articulated a vision for a new subject, one that would catalyze an interdisciplinary community in pursuit of new research directions of benefit to both VIS and psychology. We begin this chapter by detailing the grounds on which we agree with this premise: that visualization (as a phenomenon) is a fertile laboratory for exploring human cognition, that engineering and design-driven research in visualization can be improved via appropriate grounding in psychological theory, and that well-structured collaborations across disciplinary boundaries can foster a virtuous cycle of mutual benefit. Where we diverge from this vision, is in characterizing the subject as *new*. Rather, we see the intersection of visualization and psychology as tracing back to the origins of human–computer interaction. Furthermore, relevant study of external representations permeates beyond the present institutional boundaries of VIS. We believe that the psychology of visualization is so fundamental to our progress that a call for a new interdisciplinary community should both catalyze a dedicated research program *and* re-center and expand the boundaries of visualization as a field.

Writing from the hallowed halls of Xerox PARC in the late 1990s, Stuart Card, Jock Mackinlay, and University of Maryland colleague Ben Shneiderman compiled what was to become the first *de facto* textbook for a burgeoning field—*Readings in Information Visualization: Using Vision to Think* (1999). Compiled a decade after the NSF-sponsored report that spawned the formal discipline [42] this now-venerated collection of papers and essays documented the state of VIS research at the close of its "foundational period," its table of contents betraying its continued entanglement with human–computer interaction, human factors, and computer graphics communities. As a technology, visualization opened new frontiers for presenting data in multiple dimensions with real-time interactions that the newly affordable PC platforms could render. Visualization was a tool for exploring the new information structures digital computers afforded, for supporting user interaction within the document-application paradigm of the time, and for conceptualizing and building the very graphical user interfaces we take for granted today. And there at the very beginning of visualization, there was *the psychology of external representation*. Card, Mackinlay and Shneiderman saw fit to begin their introductory chapter with a narrative of cognition outside the mind, describing how visual external representations like Arabic numerals, slide rules, and navigational charts could be used to support computation distributed through the environment.

But VIS was and would remain first and foremost a constituent of Computer Science. Like HCI and Human Factors, the early contributions of Psychology would be primarily psychophysics and empirical measures of "usability." While these areas are not to be dismissed, in the interceding decades, scientists have come to embrace perspectives that ground cognition in a body, situated in an environment, distributed through an ecosystem. There is a milieu in which these perspectives intersect and inform research as disparate as how expert mathematicians invent notations for new concepts, how animations of 3D models help or inhibit learning in chemistry, and how multiple modalities can be leveraged to engage diverse audiences in museums. These questions too are about humans interacting with representations of information; they are *like* but not *quite* VIS material. We believe that as a field, visualization should re-center itself in this space, taking a step back from Computer Science and toward social and behavioral sciences more broadly, "zooming out" from the interactive, abstract, computer-based caveats of (traditional) visualization to the first principles that apply across these phenomena. If we shift our focus from *visualization as a method of computing* to *external representation as a tool for thinking*, we find a framework for giving structure to the factors that exert causal influence on the phenomena we study; concepts that considered in isolation appear idiosyncratic may in fact be part of a more predictable, coherent whole.

Acknowledgments We offer thanks to Pamela Riviere for offering a diagram of neural data in Fig. 9.1, to Arvind Satyanarayan, Paul Parsons, David Kirsh, and Oisin Parkinsoncoombs for productive discussions on these topics, and to anonymous reviewers for productive feedback.

References

1. W. Aigner, S. Miksch, W. Muller, H. Schumann, and C. Tominski. Visualizing time-oriented data. a systematic view. *Computers & Graphics*, 31(3):401–409, June 2007.
2. S. Ainsworth. DeFT: A conceptual framework for considering learning with multiple representations. *Learning and Instruction*, 16(3):183–198, June 2006.
3. W. Bechtel. *Philosophy of Science: An Overview for Cognitive Science*. Psychology Press, 1988.
4. F. Beck, M. Burch, S. Diehl, and D. Weiskopf. A Taxonomy and Survey of Dynamic Graph Visualization. *Computer Graphics Forum*, 00(00):1–27, 2016.
5. J. Bertin. *Semiology of Graphics: Diagrams, Networks, Maps*. University of Wisconsin Press, Madison, WI, 1983.
6. A. Blackwell and Y. Engelhardt. A Meta-Taxonomy for Diagram Research. In *Diagrammatic Representation and Inference*, pages 47–64, 2002.
7. T. Blascheck, K. Kurzhals, M. Raschke, M. Burch, D. Weiskopf, and T. Ertl. Visualization of Eye Tracking Data: A Taxonomy and Survey. *Computer Graphics Forum*, 36(8):260–284, 2017.
8. M. Boden. *Mind As Machine: A History of Cognitive Science Two-Volume Set*. Oxford University Press, 2006.
9. M. Bostock, V. Ogievetsky, and J. Heer. D^3 Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), Dec. 2011.

10. M. Cabré Castellví. *Terminology: Theory, Methods, and Applications*, volume 1. John Benjamins Publishing, 1999.
11. S. Card, J. Mackinlay, and B. Shneiderman. *Readings in Information Visualization. Using Vision to Think.* Morgan Kaufmann Publishers, San Francisco, CA, 1999.
12. D. Chandler. *Semiotics: the basics*. The basics. Routledge, New York, NY, third edition, 2017.
13. P. C. Cheng. Probably good diagrams for learning: Representational epistemic recodification of probability theory. *Topics in Cognitive Science*, 3(3):475–498, 2011.
14. P. C. Cheng and H. A. Simon. Scientific discovery and creative reasoning with diagrams. In S. M. Smith, T. T. Ward, and R. A. Finke, editors, *The Creative Cognition Approach*, pages 205–228. The MIT Press, 1995.
15. E. Chi. A taxonomy of visualization techniques using the data state reference model. In *IEEE Symposium on Information Visualization 2000. INFOVIS 2000. Proceedings*, pages 69–75, 2000.
16. A. Clark. *Being There: Putting Brain, Body and World Together Again*. MIT Press, 1997.
17. A. Clark and D. Chalmers. The Extended Mind. *Analysis*, 58(1):7–19, 1998.
18. W. S. Cleveland and R. McGill. Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods. *Journal of the American Statistical Association*, 79(387):531–554, 1984.
19. W. S. Cleveland and R. McGill. An Experiment In Graphical Perception. *International Journal of Man-Machine Studies*, 25(5):491–500, Nov. 1986.
20. W. S. Cleveland and R. McGill. Graphical Perception: The Visual Decoding of Quantitative Information on Graphical Displays of Data. *Journal of the Royal Statistical Society Series A-Statistics in Society*, 150:192–229, 1987.
21. E. de Vries and C. Masclet. A framework for the study of external representations in collaborative design settings. *International Journal of Human-Computer Studies*, 71(1):46–58, Jan. 2013.
22. M. A. Elliott, C. Nothelfer, C. Xiong, and D. A. Szafir. A Design Space of Vision Science Methods for Visualization Research. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1117–1127, Feb. 2021.
23. Y. Engelhardt. *The Language of Graphics*. PhD thesis, University of Amsterdam, 2002.
24. S. Few. *Now You See It: Simple Visualization Techniques for Quantitative Analysis*. Analytics Press, 2009.
25. B. Fisher, T. M. Green, and R. Arias-Hernández. Visual Analytics as a Translational Cognitive Science. *Topics in Cognitive Science*, 3(3):609–625, 2011.
26. J. Grudin. *From Tool to Partner: The Evolution of Human-Computer Interaction*, volume 10 of *Synthesis Lectures on Human-Centered Informatics*. Morgan & Claypool Publishers, 2017.
27. R. L. Harris. *Information Graphics: A Comprehensive Illustrated Reference*. Oxford University Press, 1999.
28. M. Hegarty. The Cognitive Science of Visual-Spatial Displays: Implications for Design. *Topics in Cognitive Science*, 3(3):446–474, 2011.
29. J. Hollan, E. Hutchins, and D. Kirsh. Distributed cognition: Toward a new foundation for human-computer interaction research. *Transactions on Computer-Human Interaction (TOCHI)*, 7(2):174–196, June 2000.
30. J. Hoopes. *Peirce on Signs: Writings on Semiotic by Charles Sanders Peirce*. The University of North Carolina Press, Chapel Hill, 1991.
31. E. Hutchins. Cognition in the Wild. *MIT Press*, pages 1–5, 1995. ISBN: 0262581469.
32. E. Hutchins. Cognition, Distributed. In *International Encyclopedia of the Social & Behavioral Sciences*, pages 2068–2072. Elsevier, 2001.
33. W. Kintsch. The role of knowledge in discourse comprehension: a construction integration model. *Psychological Review*, pages 163–182, 1988.
34. D. Kirsh. When is Information Explicitly Represented? In *The Vancouver Studies in Cognitive Science*, pages 340–365. Oxford University Press, 1990.
35. D. Kirsh. Implicit and Explicit Representation. *Encyclopedia of Cognitive Science*, pages 478–481, 2006.

36. D. Kirsh. Thinking with external representations. *AI and Society*, 25(4):441–454, 2010.
37. J. Larkin and H. Simon. Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 99:65–99, 1987.
38. P. H. Lindsay and D. Norman. *Human information processing: an introduction to psychology*. Academic Press, New York, 1972.
39. Z. Liu, N. Nersessian, and J. Stasko. Distributed cognition as a theoretical framework for information visualization. *IEEE transactions on visualization and computer graphics*, 14(6):1173–80, 2007.
40. M. Massironi. *The Psychology of Graphic Images*. Psychology Press, 2001.
41. R. Mayer. Cognitive Theory of Multimedia Learning. In *The Cambridge Handbook of Multimedia Learning*, pages 31–48. Cambridge University Press, 2005.
42. B. H. McCormick, T. A. DeFanti, and M. D. Brown. Visualization in Scientific Computing. *Computer Graphics*, 6(21):81, Nov. 1987.
43. T. Munzner. *Visualization Design and Analysis*. A.K. Peters visualization series. CRC Press, Taylor & Francis Group, 2014.
44. A. Musgrave and C. Pigden. Imre Lakatos. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, 2016.
45. D. Nauta. *The Meaning of Information*. Mouton, The Haugue, 1972.
46. U. Neisser. *Cognitive Psychology*. The Century psychology series. Appleton-Century-Crofts, New York, 1967.
47. D. A. Norman. *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1993.
48. S. E. Palmer. Fundamental Aspects of Cognitive Representation. In E. Rosch and B. B. Lloyd, editors, *Cognition and Categorization*, pages 259–302. Erlbaum, Hillsdale, NJ, 1978.
49. S. E. Palmer and R. Kimchi. The Information Processing Approach to Cognition. In *Approaches to Cognition: Contrasts and Controversies*, page 38. Lawrence Erlbaum Associates, 1986.
50. P. Parsons and K. Sedig. Common Visualizations: Their Cognitive Utility. In *Handbook of Human Centric Visualization*, pages 671–691. Springer, 2014.
51. P. Parsons and K. Sedig. Distribution of Information Processing While Performing Complex Cognitive Activities with Visualization Tools. In *Handbook of Human Centric Visualization*, pages 671–691. Springer, 2014.
52. S. Pinker. Theory of Graph Comprehension. In R. Freedle, editor, *Artificial Intelligence and the Future of Testing*, pages 73–126. Erlbaum, Hillsdale, NJ, 1990.
53. Y. Rogers. When the external entered HCI: Designing effective representations. In T. Erickson and D. W. McDonald, editors, *HCI Remixed: Reflections on Works That Have Influenced the HCI Community*, pages 275–280. MIT Press, Cambridge, MA, Jan. 2008.
54. A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-Lite: A Grammar of Interactive Graphics. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):341–350, 2017.
55. M. Scaife and Y. Rogers. External cognition: how do graphical representations work? *International Journal of Human-Computer Studies*, pages 185–213, 1996.
56. W. Schnotz. An Integrated Model of Text and Picture Comprehension. In *The Cambridge Handbook of Multimedia Learning*, pages 49–69. Cambridge University Press, 2005.
57. K. Sedig and P. Parsons. Interaction Design for Complex Cognitive Activities with Visual Representations: A Pattern-Based Approach. *AIS Transactions on Human-Computer Interaction*, 5(2):84–133, 2013.
58. K. Sedig and P. Parsons. *Design of Visualizations for Human-Information Interaction: A Pattern-Based Framework*. Synthesis Lectures on Visualization. Morgan & Claypool Publishers, 2016.
59. K. Sedig, P. Parsons, and A. Babanski. Towards a Characterization of Interactivity in Visual Analytics. *Journal of Multimedia Processing and Technologies, Special Issue on Theory and Application of Visual Analytics*, 3(1):12–28, 2012.

60. K. Sedig, P. Parsons, M. Dittmer, and R. Haworth. Human-Centered Interactivity of Visualization Tools: Micro and Macro-level Considerations. In *Handbook of Human Centric Visualization*, pages 671–691. Springer, 2014.
61. P. Shah. Graph Comprehension: The Role of Format, Content and Individual Differences. In M. Anderson, B. Meyer, and P. Olivier, editors, *Diagrammatic Representation and Reasoning*, pages 173–185. Springer London, London, 2002.
62. P. Shah and E. G. Freedman. Bar and line graph comprehension: An interaction of top-down and bottom-up processes. *Topics in Cognitive Science*, 3(3):560–578, 2011.
63. P. Shah, E. G. Freedman, and I. Vekiri. The Comprehension of Quantitative Information in Graphical Displays. In P. Shah, editor, *The Cambridge Handbook of Visuospatial Thinking*, pages 426–476. Cambridge University Press, New York, NY, 2005.
64. P. Shah and J. Hoeffner. Review of graph comprehension research: Implications for instruction. *Educational Psychology Review*, 14(1):47–69, Mar. 2002.
65. B. Shneiderman. The Eyes Have It: A Task By Data Type Taxonomy for Information Visualization. *Proceedings 1996 IEEE Symposium on Visual Languages*, pages 336–343, 1996.
66. D. Simkin and R. Hastie. An Information-Processing Analysis of Graph Perception. *Journal of the American Statistical Association*, 82(398):454–465, 1987.
67. D. E. Stokes. *Pasteur's Quadrant: Basic Science and Technological Innovation*. Brookings Institution Press, 1997.
68. D. A. Szafir, S. Haroz, M. Gleicher, and S. Franconeri. Four types of ensemble coding in data visualizations. *Journal of Vision*, 16(5):11–11, Mar. 2016.
69. M. Tory and T. Moller. Rethinking Visualization: A High-Level Taxonomy. In *IEEE Symposium on Information Visualization*, pages 151–158, Oct. 2004.
70. B. Tversky. Visualizing Thought. *Topics in Cognitive Science*, 3(3):499–535, 2011.
71. VisPsych Committee. Visualization Psychology — 2020 IEEE VIS Workshop, 2020.
72. C. Ware. *Information Visualization: Perception for Design*. Elsevier Inc., San Francisco, CA, 2004.
73. D. Whitney and A. Y. Leib. Ensemble Perception. *Annual Review of Psychology*, page 25, 2017.
74. L. Wilkinson. *The Grammar of Graphics*. Statistics and Computing. Springer-Verlag, New York, 1999.
75. M. Wilson. Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4):625–36, Dec. 2002.
76. J. Worrall and G. Currie, editors. *The Methodology of Scientific Research Programmes: Philosophical Papers*, volume 1. Cambridge University Press, Cambridge, 1978.
77. J. Zhang and D. Norman. Representations in Distributed Cognitive Tasks. *Cognitive science*, 18(1):87–122, Jan. 1994.