
Cognitive Modeling of Interpretations of Representations

Aaron Stockdill

Grecia Garcia Garcia

Peter C-H. Cheng

Department of Informatics, University of Sussex, Brighton, UK

A.A.STOCKDILL@SUSSEX.AC.UK

G.GARCIA-GARCIA@SUSSEX.AC.UK

P.C.H.CHENG@SUSSEX.AC.UK

Daniel Raggi

Mateja Jamnik

Department of Computer Science and Technology, University of Cambridge, Cambridge, UK

DANIEL.RAGGI@CL.CAM.AC.UK

MATEJA.JAMNIK@CL.CAM.AC.UK

Abstract

For research on cognition with external representations, we argue that more attention should be paid to the diversity of interpretations produced by different users of the same representations. We present a tool (RISE) for modeling memory structures that constitute interpretation which operationalizes a previously proposed theory of the interpretative structure of representations (RIST) and its accompanying notation (RISN). We demonstrate how the facility to build detailed models of interpretations may be used to examine the relative cognitive efficacy of representations for different users. This in turn highlights the differences between representations' cognitive properties, and provides an opportunity to recommend alternative, more effective representations for a specific user.

1. Introduction

Much human cognition involves interactions with external symbolic representations, such as: written natural language; visual information displays, including diagrams, graphs, charts and maps; and interfaces, like control panels and, of course, computer interfaces. Many advances have been made in the last 30 years in our understanding of reasoning, problem solving and learning with visual representations using empirical studies (e.g., Zhang, 1997; Carpenter & Shah, 1998; Egan & Schwartz, 1979; Cheng, 2004), computation cognitive models (e.g., Larkin & Simon, 1987; Larkin, 1989; Reed, 2019), and both together (e.g., Koedinger & Anderson, 1990; Peebles & Cheng, 2003). It is common in such studies to focus on the (imagined) performance of a prototypical user of the external representations. This presumes that users of a representation construct structures in memory that are essentially the same. Whilst this is a reasonable simplifying assumption for initial investigations of representation use, we contend that it does not adequately reflect the reality of representation use, in which different users may have substantially divergent mental representations of the same external representation. One obvious case is that between novices and expert. Experts will perceive large patterns that are meaningful relative to the content of the domain, whereas novices will see shallow configurations of graphical features (e.g., Chi et al., 1988; Egan & Schwartz, 1979). In contrast to novices, experts may mentally encode the perceived information as larger

chunks that are associated with strategic information (Chase & Simon, 1973) or as heterogenous *diagrammatic configuration schemas* (Koedinger & Anderson, 1990).

However, we contend that users' interpretations of representations are likely to differ in many situations, in which individual differences among users are less dramatic, which may include:

- Between two users with different levels of familiarity with a domain, although less dramatic than the difference between a novice and a full expert.
- Between two users with the same level of domain expertise but with different levels of experience with the type of representations being used.
- Within an individual user as they learn about a domain over hours or days.
- Within an individual user as they gain experience of a class of representations over many episodes of use.
- Within an individual user when they switch between different task goals on a given representation.
- Within an individual, moment to moment, from when they first attempt to decode a representation through to finding a solution to a given problem.

In other words, it appears that the conventional focus on a single prototypical interpretation of a representation in previous work masks a myriad of cognitive phenomena.

For example, in their seminal study, Larkin and Simon (1987) build cognitive models to explain the benefit that diagrammatic representations (sometimes) have over sentential representations. One of their production system models concerned problem solving with a diagram of a pulley system. The model centered upon components at the level of individual pulleys so that locational indexing of information could be exploited to link inferences, which was one part of their explanation of the superiority of diagrams. However, in an experiment to quantify the relative advantage of diagrammatic representations using the same problem domain, Cheng (2004) found that participants used other mental representations beyond the component-based representation of Larkin and Simon. On some problems, some of Cheng's participants used global symmetries spanning multiple pulleys to make strategic shortcut inferences. If those mental representations and inferences were incorporated into the original production system model, its performance could have been even more dramatic than for the sentential representation.

Further, in addition to the list of above situations that implicitly assume unchanging external representations, we may add circumstances in which the user is actively modifying the external representation (e.g., for instance drawing). These circumstances necessarily require the user to continually revise their interpretation of the external representation.

Consider a concrete example. Figure 1 is a multi-level pie chart that we will adopt as a running example through the paper. It was found on *Stack Overflow* (<https://stackoverflow.com/questions/48588312/labelled-multi-level-pie-chart>), which is a forum for programmers to share and discuss coding issues. Figure 1 shows the market share of different browsers over three years. It was initially posted in a thread seeking guidance on how to code the visualization in R. In addition to coding suggestions, responses also suggested using alternative

visualizations that would be more accurate and clearer, including a stacked bar plot and a line graph. Thinking about how Figure 1 is interpreted raises various questions: What are the meaningful components of the pie chart? What principles govern their organization in memory? How do the mental representations of the pie chart vary among viewers who have alternative interpretations? If we have a model of an interpretation, how can it be used to assess (and even quantify) the quality of Figure 1 compared to the alternative visualizations that were suggested? Figures 2 and 3 are cognitive models of alternative interpretations of the pie chart – concrete examples on how they differ (in terms of content and structure) are in Section 3.4.




Figure 1. A Multi-level pie chart.

All this suggests that studies of how representations are used should consider the interpretations made by individuals in specific circumstances, or at least the likely interpretation of people with similar experience, domain knowledge, and task goals. This paper reports on work that aims to support the cognitive modeling of interpretations of representations. Thus, our aims are twofold:

- A1. We examine requirements for cognitive models of users’ representation interpretations by summarizing the insights gained in our attempts, so far, to cognitively model interpretations. This includes the memory structures, rules for their composition, and emergent structures that are found across models for different types of external representations (*idioms*).
- A2. We introduce a graphical tool – *RISE* – that we are developing and allows analysts (such as cognitive science researchers and practitioners) to easily build these models. We will also note *RISE*’s analyst support features that identify representationally significant structures.

Our contribution is a general, practical modeling framework for cognitive scientists, representation designers and practitioners to model the memory structure of individuals’ interpretations of external graphical representations. Section 2 summarizes our cognitive theory and graphical notation for analyzing interpretations of representations, which we developed previously. The theory and notation will help to contextualize the explanation of the needs (A1/Section 3) and the modeling tool that operationalizes the theory and the notation (A2/Section 4). The paper ends with a discussion of the scope and limitations of our approach, and considers prospective lines of investigation and application.

2. Representational Interpretive Structure Theory (RIST) and Notation (RISN)

Cheng (2020a) and Cheng et al. (2022) describe the on-going development of *Representational Interpretive Structure Theory*, *RIST*, and its *notation*, *RISN*, for graphically describing interpretive structures of representations. Two lines of research drove the original development of *RIST*. The first was work on the design of novel representations to enhance problem solving and learning in conceptually demanding and informationally intensive domains (e.g., Cheng, 2002, 2011, 2020b).

To advance that work, a systematic method to capture users' understanding of the meaning of representations is needed. The second drive is a project – *rep2rep* – that is designing a mechanized

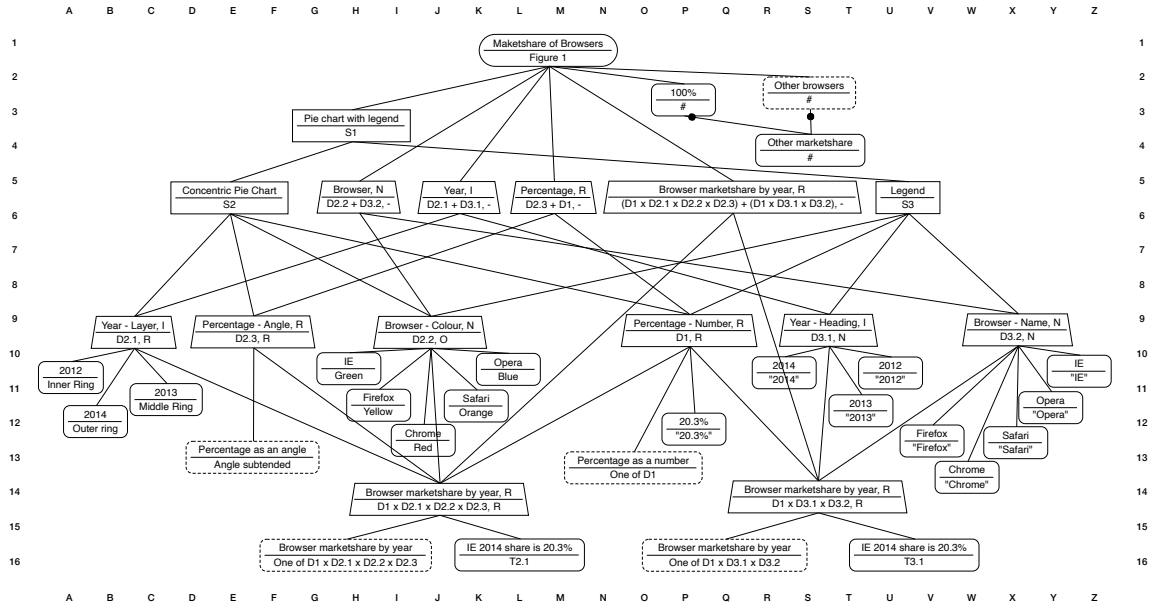


Figure 2. A RISN model for an interpretation of Figure 1. Surrounding letters and numbers are a coordinate grid for our convenience in describing aspects of the model.

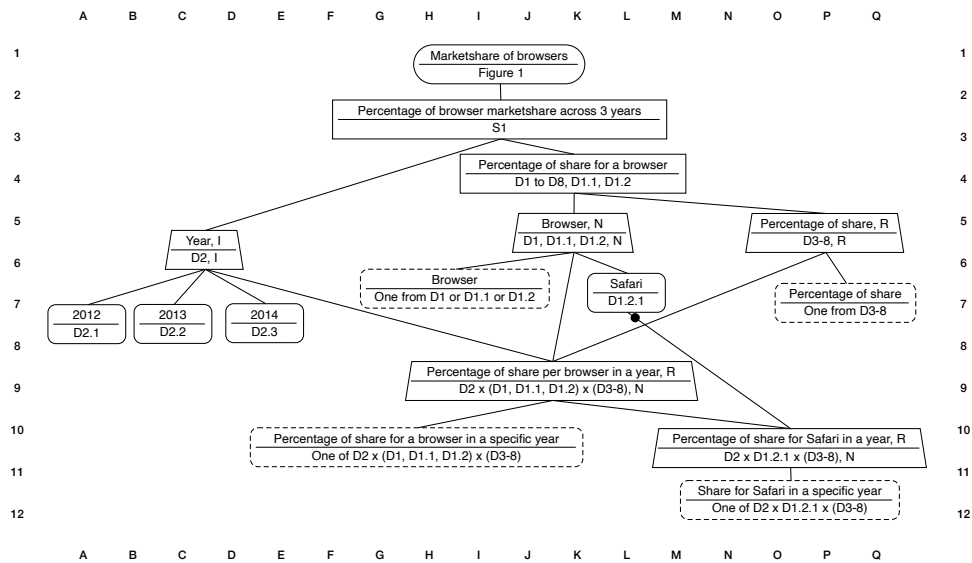


Figure 3. A RISN model showing an alternative interpretation of Figure 1.

approach for the automatic selection of representations for AI tools (Jamnik & Cheng, 2021; Raggi et al., 2020). The overall goal of the rep2rep project is to evaluate the potential efficacy of alternative representations in terms of their informational suitability and their cognitive adequacy with respect to a given problem and individuals. To evaluate cognitive adequacy, we analyze how representations fare in terms of a framework of cognitive properties of representations (Cheng et al. 2021), but this requires some means to model the mental structures that users build as they interpret representations: that is the purpose of RISP.

This section gives an overview of RISP and RISP. The key question that RISP tackles is: How are interpretations of external graphical representations encoded in memory? RISP takes inspiration from previous works that have examined what memory structures are needed to encode heterogeneous information spanning (i) domain facts and relations and (ii) perceived images from an external representation. Work on chess expertise introduced the idea of *perceptual chunking* (Chase & Simon, 1973) and then *templates* (Gobet & Simon, 1996) as schema-like structures to explain thinking in chess without resort to propositional or mental imagery forms of inference. Koedinger and Anderson (1990) proposed *diagrammatic configuration schemas* to explain the elevated performance of expert geometry problem solvers. Diagrammatic configuration schemas are heterogeneous structures that combine an image of a component of the external representations with declarative information about what inferences can be made with that component and the circumstance under which those inferences are permitted. Cheng (1999, 2002) used the idea to explain some of the benefits of *Law Encoding Diagrams* for problem solving and learning in other domains.

RISP makes three core claims about the nature of mental representations that result from interpretations of external representations (ERs). (1) The basic components of the mental representations are schemas that simultaneously encode conceptual information about the domain and graphic objects from the ER. (2) Four types of these basic schemas are proposed that encode information at different levels of conceptual and graphical granularity. (3) An interpretation of an ER is comprised of a hierarchical network of these four types of schemas. Alternative interpretations of an ER have different structures of schemas.

RISP’s four types of schemas are (in order of increasing granularity): *Representation*, *R-Scheme*, *R-Dimension*, and *R-Symbol*. Their icons in RISP are shown in Figure 4. They all associate a *domain concept* with a *graphic* in order to capture an interpretive mapping of a graphical object in the ER with a mental concept about the domain. In our running example, the overall Representation schema is presented at the very top of the model in Figure 2, and it links the concept of ‘Marketshare of Browsers’ with an external representation that is in ‘Figure 1.’ R-Schemes are used when different types of information need to be coordinated together; for example, the ‘Legend’ R-Scheme (Figure 2, U5) coordinates (is the parent of) four concepts that are of different types: ‘Percentage – Number’, ‘Year – Heading’, ‘Browser – Name’ and ‘Browser – Colour’. The alphanumeric annotations in the bottom half of some schema icons (e.g., “S1”, “D2.2+D2.3”) refer to labels




Figure 4. Schemas of RISP: Representation, R-Scheme, R-Dimension, and R-Symbol (from left to right).

the analyst has added to their own copy of Figure 1.¹ That is, each concept is (potentially²) associated with a part of the specific representation being considered, tying the interpretation to the representation. In contrast to the R-Scheme, the R-Dimension schema represents information that is of the same type. Figure 2 shows several R-Dimensions such as ‘Browser’, ‘Year’, ‘Percentage’, and ‘Browser marketshare by year’ (Figure 2, H5–Q5). Each concept and graphic in an R-Dimension has a quantity scale (one of Stevens’ (1946) quantity scales: Nominal, Ordinal, Interval, and Ratio; the first letter replaces Q in Figure 4). The quantity scales of the concept and the graphic are independent: just because the concept has one quantity scale does not mean the graphic must have the same quantity scale – we shall discuss this further in Section 3.1. R-Symbols come in two flavors: a ‘single’ R-Symbol (referred to simply as an R-Symbol), and a ‘class’ R-Symbol. Class R-Symbols stand in for many closely related R-Symbols that the analyst considers to be distinct, but only in an uninteresting way. For example, a single R-Symbol is ‘2012’ represented in the Inner Ring (Figure 2, A11); while an example of a class R-Symbol is the ‘Percentage as an angle’, which is shown in RISN with a dashed border (E13) because there are many percentages represented by an angle, differing only in value.

Schemas may be connected in two ways: they are part of a hierarchy (e.g., an R-Symbol is one element of an R-Dimension); or one concept is anchored beneath another (e.g., a region is anchored by bounding curves). Hierarchical connections are largely unconstrained; anchoring connections can only occur below R-Symbol schemas. While hierarchy is the ‘typical’ connection, an anchoring connection is very specific: a schema is anchored below an R-Symbol when the concept of the anchored schema is induced by the concept within the R-Symbol. For example, in Figure 2, the analyst considers the concepts of ‘100%’ and ‘Other browsers’ (P3, S3) in spite of these not being explicit in the representation; these two concepts help to deduce (induce) that there is ‘Other marketshare’ (S4) for other browsers. Complete, valid RIST models are directed acyclic graphs.

Examples of full models are in Figure 2 and Figure 3. In Figure 2, the analyst has focused on the perceived complexity of the ER’s dimensions, positioning four R-Dimensions (H5–Q5) below the root Representation schema (M1) before breaking them down by their appearance. Also, below the root, they position the overall R-Scheme for the ‘Pie chart with legend’ (H3), which breaks down into two sub-R-Schemes: the ‘Concentric Pie Chart’ (E5) and the ‘Legend’ (U5). Each sub-R-scheme breaks down into its constituent R-Dimensions, each of which contain R-Symbols. Then, the analyst combines the dimensions (B9–X9) to read off ‘Browser marketshare by year’ in two different ways, one for the pie chart (J14) and one for the legend (S14). The analyst only uses anchoring to encode the inference that there must be some unrepresented market share (S4) because the percentages do not add to 100%, and that there must be other browsers. In contrast, in Figure 3, the analyst had a more abstract interpretation of the representation and coded the duplication of information using the ‘graphic’ part of the schemas. The analyst highlighted the key concepts of ‘Year’, ‘Browser’ and ‘Percentage of share’ (Figure 3, C5–P5) – the latter two coordinated by a sub-R-Scheme (K4), and all coordinated by an overarching R-Scheme (J3). Then, these key

¹ In this model the analyst uses *S* to denote R-Schemes, and “S1” is the first of them. *D* denotes R-Dimensions, with numbers indicating the relations among them; the label “D2.2+D2.3” refers to a graphic composed of two sub-R-Dimensions, D2.2 and D2.3, whose parent is the S2 R-Scheme.

² An analyst might assign a schema the graphical reference “#” – that is, there is *no* graphical object associated with this concept. See Section 3.2 for a more detailed discussion.

concepts are combined to read off the share of the browsers per year (K9); this schema has a sub-R-Dimension that only selects data for Safari (P10) and it is anchored by the ‘Safari’ R-Symbol (L7).

Many accounts of the structure of representations have been proposed that do not explicitly distinguish internal mental and external graphical components of representations (e.g., Kosslyn 1989; Zhang, 1996; Card et al., 1999). One exception is Pinker’s (1990) theory that considers the memory representations needed for graph comprehension in the form of schemas that encode information about spatial relations and visual properties (e.g., *part, height*) of components (e.g., *line, bar*). RIST, in contrast, posits four types of schemas that perform functions at different levels with generic hierarchical links among them. The aim with RISN models is to abstract away from all the low-level minutia of an ER whilst still capturing the essential structure that encodes the meanings which the reader extracts during the process of interpretation.

3. Analysis of Cognitive Properties of Representations in RISN Models

To illustrate the value of RIST for analyzing representation interpretations, and to address aim A1, this section considers how RISN models can be used to examine cognitively important features of representations that affect their relative efficacy for users. In turn, this reveals new questions for research about cognitive properties of representations. We will consider: (1) the encoding of quantities; (2) the mapping of domain entities to graphical objects; and (3) generic interpretive structures of representations (*idioms*). The section finishes with a note about the evaluation of RIST and RISN.

3.1 Quantity Scales

Classically, Stevens (1946) identified four quantity scales: Nominal, Ordinal, Interval and Ratio. The character of inferences with each type of scale is distinct and, for example, they are used to differentiate classes of statistical tests. In our pie chart example, three of the scales are present: nominal – browser types; interval – years; and ratio – the percentage share. The importance of quantity scales is recognized in accounts of representations (e.g., Zhang, 1996; Card et al., 1999).

Domain variable concepts and graphic objects can be assigned quantities. When both quantities align, reasoning may progress smoothly, with the ER supporting mental inferences well, because the types of inferences that a person may desire to make are reflected by operations with the ER. However, when they are different, reasoning may be obstructed, because permissible operations on the ER do not have a valid corresponding inference, or because no operation on the ER directly matches the desired inference. For example, in Figure 2, the analyst identified that the ring representation of Year does not match the quantity scale of the concept (B9, and Figure 4): the concept of Year has an Interval scale, but the diagram represents it as ring regions, the areas of which have a Ratio scale.




Figure 4. An R-Dimension from Figure 2 (B9) in which the quantity scale for the concept does not match the graphical quantity scale.

The definition of the R-Dimension schema in RIST includes slots for the quantity scale of both the domain concept and its corresponding graphic object, and RISN icons explicitly notate the quantity scales (the *Qs* in the R-dimension icon in Figure 4 are replaced by one of *N, O, I* or *R*).

The construction of a RISN model for a representation thus forces the analyst to consider the compatibility of the quantity scales for every R-Dimension in the model. Given complete models for different representations we can potentially judge their relative merits by considering the number of types of mismatches for each of the quantity scales. However, this possibility reveals gaps in our current knowledge about the cognitive costs of making inferences with the different quantity scales and incompatible quantity scales. For example, how much harder is it to reason about ordinal relations over nominal relations, or ordinal compared to ratio, and so forth? In the rep2rep project we have conducted experiments to assess the cost of inferences with the quantity scales across a variety of reasoning tasks. Aggregating over all the tasks, it was discovered that cost increases in an approximately linear fashion across the quantity scales from nominal to ratio. (The experiments will be presented in a future publication, currently in preparation.)

3.2 Concept Mapping

The nature of the mapping between mental concepts and their associated graphic objects in the ER is another important, well-acknowledged cognitive property of representations (Barwise & Etchemendy, 1995; Gurr, 1998; Moody, 2009). We call this property *concept mapping* (Cheng et al., 2021). Following Moody’s (2009) terminology, the five types of concept mapping are: *isomorphism*, *redundancy*, *excess*, *overload*, and *deficit*. Isomorphism happens when a concept in the domain has a distinct graphical symbol in the representation. A representation has redundancy if there exist at least two distinct graphical symbols that encode the same concept. Excess happens when some graphical symbol is not associated with any concept in the domain, acting as ‘chart junk’. A concept mapping is overloaded if two distinct concepts map to the same graphical symbol in the representation; that is, a graphic object is overloaded if it stands for two different concepts. A representation is deficient when there is a concept for which there is no graphical symbol; thus, there are concepts in the domain which do not occur in the representation, forcing the reader to mentally keep track of the missing concepts, rather than offloading them to the representation. Each of these concept mappings is identifiable within a RISN model.

Let’s consider two examples, again using Figure 1. RIST models deficit by having schemas with a concept, but no graphic. In Figure 2, the analyst identified the concepts ‘Other browsers’ and ‘100%’, and marked the lack of graphic with ‘#’ in both cases (Figure 2, P3, S3; reproduced in

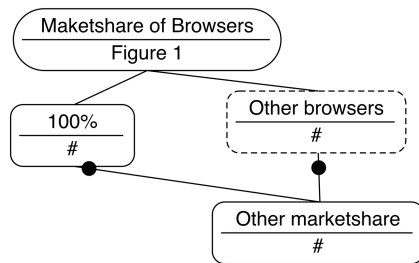


Figure 5. Schemas from Figure 2 that do not have a graphic reference, making the representation deficient.

Figure 5). That is, the analyst noted that the five listed browsers are not exhaustive and that the given percentages do not add to 100%. Thus, the representation is deficient with respect to these concepts. Redundancy can be identified in RISN models when (a) a concept is reused in different schemas and assigned a different graphic, or (b) a schema has more than one graphic annotation. For example, in Figure 2 (reproduced in Figure 6), we see the concept of ‘IE 2014 share is 20.3%’



Figure 6. Schemas from Figure 2 that share the same concept, making the representation redundant.

being reused in different schemas: it is used for the legend of the pie chart (V16) and again, separately, for the annotation of the concentric pie chart (L16). In the case of Figure 3 (and Figure 7), the concept of ‘Browser’ (H7) has more than one graphic annotation, which correspond to the legend and the concentric pie chart.




Figure 7. A schema from Figure 3 (H7) that has several graphic references, making the representation redundant.

As the examples show, RIST models encode the types of concept mappings that an interpretation of an ER invokes. This raises the possibility of rating the relative efficacy of alternative representations by analyzing the numbers of each type of concept mapping and their location within the overall schema hierarchy of their models. Superior representations will in general be more isomorphic (e.g., Gurr, 1998). However, like the quantity scales property, there are gaps in our knowledge about the concept mapping cognitive property of representations. For example, what is the relative cognitive cost of different types of mappings? Might it be worthwhile to tradeoff some redundancy in order to eliminate some deficits? How does the severity of different forms of mapping vary across the levels of schemas from R-symbols to R-schemes? We are starting to address such question empirically in the rep2rep project.

3.3 Generic representational structures – *idioms*

A pertinent consequence of being able to build models of representations in RISN – which highlights the value of RIST – is our discovery of common patterns of schemas across divergent representations; we call these *idioms* (Cheng et al., 2022). To date, the idioms we have identified fit into three classes: *collections* of elements, how *dimensions* are composed and decomposed, and how *coordinate systems* can be read. A list of currently discovered idioms was given in Cheng et al. (2022), and these play an important role in the *insights* features of the RISN Editor (Section 4).

Some of these idioms are present in Figures 2 and 3. For example, the collection idiom *pick* is found when an R-Dimension acts as a parent of an R-Symbol and (typically) a class R-Symbol. It means that there may be many concepts, but the analyst is highlighting a concept as being of particular interest. Figure 2 shows a pick idiom with the schemas in P9, P12 and N13 (reproduced in Figure 8(a)); Figure 3 shows this idiom in K5, H6, and L7 (reproduced in Figure 8(b)). The dimension idiom *product* combines two or more R-Dimensions so they can be read as a composite. Examples of this idiom can be found in Figure 2, with the R-Dimension schemas C9–P9 and its product R-Dimension J14 (reproduced in Figure 9(a)); in Figure 3, the schemas are C6, K5, P5 and K9 (reproduced in Figure 9(b)). Both models show examples of coordinate system idioms: these are

structures that coordinate information. An example of an *implicit coordinate system* is shown when an R-Scheme is the parent of a group of R-Dimensions. In Figure 2, this can be seen in E5 and C9–P9 (reproduced in Figure 10).

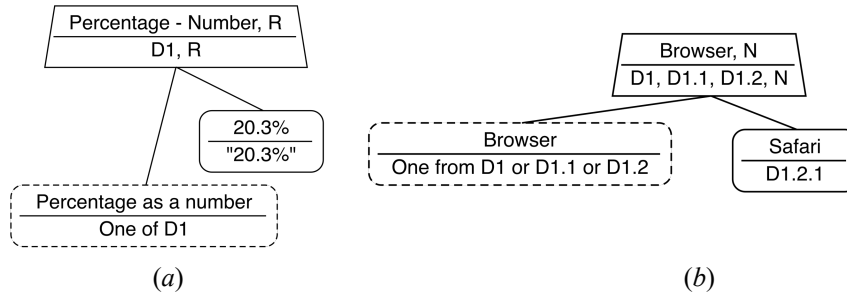


Figure 8. Examples of the pick idiom from Figure 2 (in (a)) and Figure 3 (in (b)).

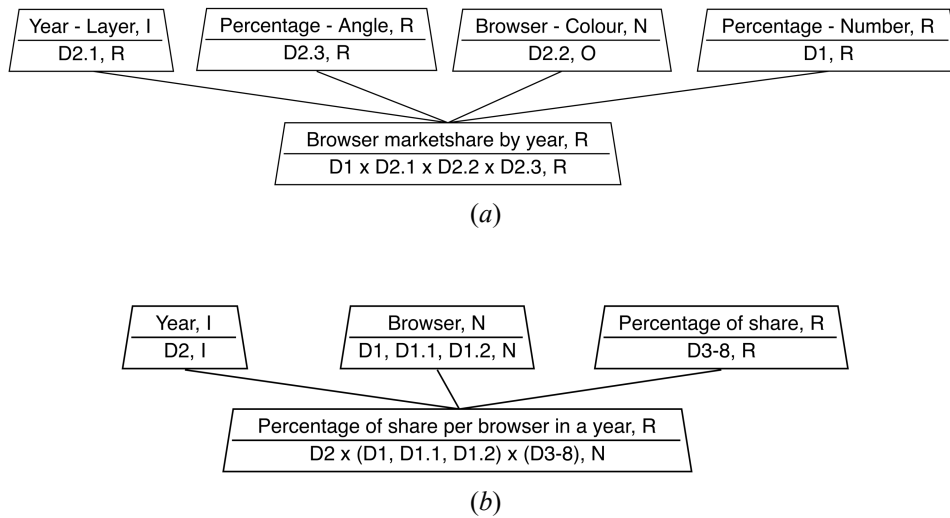


Figure 9. Examples of the product idiom from Figure 2 (in (a)) and Figure 3 (in (b)).




Figure 10. An example of the implicit coordinate system idiom from Figure 2.

In summary, it appears that there is an intermediate layer of meaningful cognitive structures that operates between the level of elementary concepts and full representations. This contrasts with previous accounts of representations that analyze representations as networks of elementary components (e.g., Pinker, 1990) and also contrasts with accounts that posit heterogeneous schemas (e.g., Koedinger & Anderson, 1990). In the latter case, the primary functional layer is one of schemas, whereas in RIST, subnetworks comprised of several schemas may do substantial cognitive work. The psychological reality of idioms needs to be fully established in experiments that examine whether users of representations perform cognitive operations at this level in the network of schemas. For example, will the analysis of concurrent verbal protocols show that users' attention focuses on clusters of schemas that are idioms? Will verbal descriptions of representations written to emphasize idioms be more readily comprehended than those that highlight other types of relations?

3.4 Evaluation of RIST and RISN

Figures 2 and 3 demonstrate that RIST and RISN can capture alternative readings of Figure 1 by two analysts. These models differ in the concepts they encode and in their structures. For example, Figure 2 (and Figure 5) shows that the analyst noted the misrepresentation of the market share and so included the concepts of 'other browsers' and '100%' (P3, S2, S4). Figure 3's analyst missed this point. Also, the model in Figure 2 makes a high-level split between the Concentric Pie Chart (E5) and the Legend (U5), reflecting the distinct representational functions that this analyst attributed to the two parts of Figure 1. In contrast, Figure 3 omits detailed mappings of elementary graphic objects to concepts and adopts a more abstract perspective focusing more on the overall form of the interpretation without concern about redundant graphic objects. Further, there are differences in how the analysts conceptualize the coordinate systems; in Figure 3 the system includes Browser and Percentage R-dimensions as an explicit sub-system alongside Year (K5 & P5 versus C6), whereas the same R-dimensions are given equal conceptual status in one system in Figure 2 (H5-M5).

Quantity scales, concept mapping and idioms are three aspects of the nature of representations that are brought to the fore as we modeled interpretations of representations using RIST. In each case, our attempts to model representations has either re-confirmed the importance of certain cognitive properties of representations or lead to the discovery of a new aspect of representations. In doing so, new questions about the nature of representations have been raised, which argues for the value of RIST and the utility RISN, if not for their validity.

Although formal experimental evaluations of the acceptability of the theory are yet to be conducted, we have developed and run small workshops with potential users of the approach. The workshops lasted from 2 to 3.5 hours. In this time, participants were given an overview of RIST, RISN, practiced annotating representations (see Section 4.2) and familiarized themselves with the RISN Editor, RISE. From the feedback and the quality of the models that the participants produced, it was clear that they understood the underpinning theoretical ideas, the constraints of the notation and how to drive the editor.

4. RISE – Graphical Tool for Constructing RISN Models

So far, we have described the theory and notation (Section 2), and illustrated how RISN models can be used to judge representations in terms of properties that may determine their cognitive efficacy, which in turn raises new questions about the nature of representations for this area of study (Section 3). In this section, we address aim A2 by considering practical matters: (a) the RISN modeling tool – RISE; and (b) our recommendations about how to approach modeling interpretations of representations. A key aspect of RISE is the incorporation of feedback into the modeling process to support the construction of valid models and the identification of potential idioms.

4.1 The RIS-Editor

Although a generic vector graphic drawing package plus a word processor were suitable for our initial model building (Cheng, 2020a), we have developed a dedicated tool, RISE, for building RISN models. The motivation for developing RISE is twofold. First, we wish to make the process as straightforward as possible: RISE integrates the construction of a graphical network and the filling of the schema slots. Second, RIST is complex, with four types of schemas and rules for assembly of valid networks. Thus, RISE checks for the completeness of schemas and validity of subnetworks. Further, RISE includes methods to intelligently identify idioms during model construction, which can prompt modelers to consider whether they might use these common structures. In other words, RISE may be considered an operationalization of RIST in software. The editor is a React web app. Let us explore features of the editor that support RISN model building.

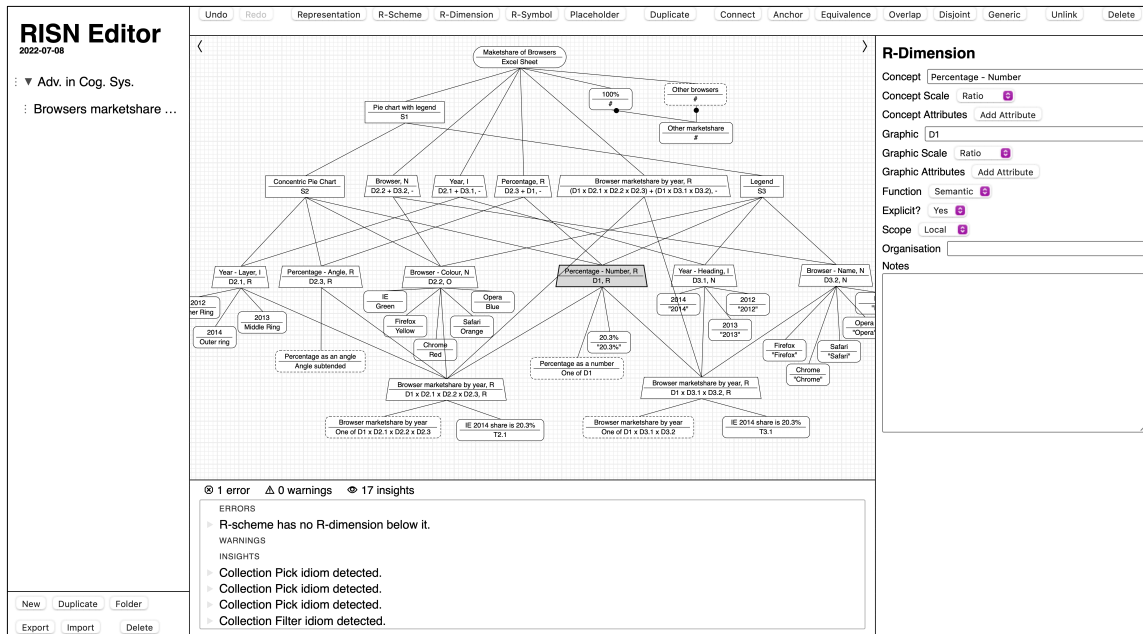


Figure 11. Screenshot of RISE being used to edit the model in Figure 2.

4.1.1 Editing RISN Models

The fundamental feature of RISE is manipulating schemas and connections. Most of the screen is devoted to the model canvas, where the RISN model will appear and can be interacted with (see Figure 11). Along the top of the editor is the toolbar, which contains buttons for adding schemas and connections of each variety, as well as deleting and duplicating them.

On the model canvas, typical actions can be performed with a combination of mouse actions and keyboard shortcuts. The model can be navigated by clicking and dragging on the canvas to pan and use the scroll wheel to zoom in and out. Schemas can be selected individually, or as multiple selections. Double-clicking on either the concept (text above the dividing line) or the graphic (text below the line) of a schema allows the analyst to change the value.³ Hotkeys can be used when the cursor is over the canvas. For example, if the analyst presses the ‘D’ key⁴ on their keyboard while their cursor is over the model canvas, an R-Dimension schema is inserted below their cursor. Further, this new R-Dimension is automatically connected hierarchically to any selected schemas.

While we have not performed any usability evaluations, it is obvious that this tool is superior to the previous model-building method using the drawing capabilities of Microsoft PowerPoint. The combination of mouse interaction and keyboard shortcuts allows the analyst to produce a significant portion of their RISN model using just the model canvas. For the remaining schema slots required by RIST, but not visualized in RISN models, we provide the *inspector*.

4.1.2 The Schema Inspector

RIST specifies many *slots* for analysts to fill in so that they comprehensively capture the interpretation of the representation. To support the analyst, we provide the *inspector* panel, which lists each slot for the selected schema. For example, in Figure 11 we see the analyst has selected an R-Dimension schema, so the inspector panel (on the right) is showing the appropriate slots. The slot content is evaluated ‘live’ – for example, if the analyst edits the ‘Concept’ slot in the inspector panel, the corresponding schema on the model canvas is immediately updated to show the new concept.

Each slot consists of a name and a value. The slots have been detailed in previous work by Cheng (2020a). If the value is free-form, we provide a text entry field; if the value is restricted, we provide a dropdown box of all possible values. The name of each slot can be hovered over to reveal a brief description of the slot. This reduces the memory load on the analyst: they do not need to remember every slot, and if they forget the purpose of a slot then the information is available.

The inspector can display the slots for a single schema, or for many schemas if many are selected. The inspector also provides a ‘notes’ textbox for each schema and connection, and for the entire RISN model if nothing is selected. Analysts may use the textboxes to document a model – for instance, to record explanations of specific modeling decisions as needed.

³ Currently the analyst cannot edit R-Dimensions’ quantity scales this way; they must be set via the inspector.

⁴ The hotkeys are displayed when hovering over the toolbar buttons.

4.1.3 Supporting Analysts

The features we have outlined so far make it easy to build RISN models, but inevitably analysts can make mistakes: RISN models are non-trivial, with many rules and slots to keep track of. Thus, we support analysts with feedback inside the editor: RISE continually monitors the model that the analyst is building, and provides feedback along the bottom of the screen in three categories: *errors*, *warnings*, and *insights* (see Figure 11). Errors are things that are *definitely* wrong, warnings are things that are *possibly* wrong, and insights are things that we believe are meaningful to consider. Feedback computations run on every meaningful change.⁵ We now briefly consider each kind of feedback.

The editor can quickly identify *errors* in models. Common errors include forgetting to fill in a slot, introducing a cycle in the model, having more than one root, using an anchor connection below a schema that is not an R-Symbol, or using a hierarchy connection from an R-Symbol to a non-R-Symbol schema. To see which errors have been detected, the analyst can click on the ‘errors’ count at the bottom of the screen; the error can be selected such that the affected schemas are highlighted in the model canvas. Many errors come with recommendations that will correct the error. When the analyst successfully corrects the error, the error message is automatically removed.

Sometimes the editor detects cases in RISN models that are not ‘wrong’, per se, but are surprising and extremely unlikely. We call these *warnings*, and are displayed by clicking the ‘warnings’ count on the bottom of the screen. Common warnings include leaving the default Concept or Graphic value unchanged, or introducing a ‘transitive hierarchical connection’ – that is, a direct hierarchical connection from a grandparent schema to its grandchild. These cases can be valid: you might truly have a concept ‘#Sym#’ (the default R-Symbol concept value), or the hierarchical connection might necessarily connect a grandparent to its grandchild. However, we consider them sufficiently unlikely that we highlight them for the analyst to reconsider. If the analyst is certain that this *is* what they want in their model, they can mark the warning as ‘ignored’; it will not be included in the warning count (although it will be counted as an ignored warning) and will appear greyed out in the warnings list. As with errors, if the cause of the warning is eliminated from the RISN model, the warning is automatically removed.

The final aspect of feedback is *insights*. Currently, insights show the analyst which idioms the editor has detected within their model. As with errors and warnings, analysts can view insights by clicking the ‘insights’ button at the bottom of the screen. Each has a description explaining what it means, and selecting an insight then highlights the relevant schemas in the model. Idioms are a powerful aspect of RIST, providing an intermediate level of abstraction between the atoms of schemas and the entire model of a representation; surfacing them supports the analyst in understanding the interpretation they are modeling. We are considering adding further insights: R-Dimensions where the quantity scale differs between concepts and graphics, cases where the concept or graphic occur in more than one schema, and potentially measures of the model’s complexity.

⁵ A ‘meaningful change’ is an action like adding or deleting schemas or connections, or modifying their slots. Something that would *not* be a meaningful change would be repositioning a schema’s icon, for example.

4.2 Building models

RISE provides the tool in which to build models; now, we turn to the process by which an analyst builds a model. We emphasize that, because RISE is continually monitoring the correctness of the analyst’s model, they are free to focus on the nature of the interpretation, rather than ensuring they are following all the rules of RIST. Analysis is a cyclic process that consists of four key states: *identify*, *describe*, *annotate*, and *build*. Previous guidance was outlined by Cheng (2020a); here we refine this guidance and examine how it relates to an analysis of Figure 1.

Identifying the ER is more than just picking an ER. It is important to consider its context of use. Who is the reader likely to be? What is their goal in working with this ER? These questions inform the interpretation we will be modeling, and so they impact our description. When writing the description, analysts must be careful to ensure they are capturing a single interpretation. This is particularly important when updating the interpretation during the annotation and building phases of the model, but is equally important early on, too. Analysts must fix the level of abstraction: are we modeling this ER, or are we modeling this class of ERs? In Figure 2, our model is of this specific chart, not the more general class of concentric pie charts. The specific content of the description is often the important symbols, the coordinate systems in place, and the dimensions being discussed.

After writing an initial *description* of the representation, the analyst *annotates* the ER. They add labels to the given representation such that they can unambiguously reference specific graphic objects or structures in their model. What they annotate is informed by their description, however it is common that the annotating process reveals implicit concepts which should have been described. The analyst is free to update their description, but must take care to not alter the overall interpretation. In Figure 2, we see the analyst often uses literal strings for their annotations, e.g., “2013”, including the quotation marks (at T12); more abstract entities get annotations such as ‘S3’ for the R-Scheme of the legend (at U5), or ‘D3.1’ for the years R-Dimension in the legend (at S9). The analyst uses an ad-hoc system for other annotations, such as ‘x’ or ‘+’ for combining dimensions (e.g., at Q5) – which can be documented in the notes for each schema.

Finally, the analyst begins the *model-building* process inside RISE. Using their description and annotations, they create schemas for the concepts they have identified, and insert connections between the schemas. RISE ensures that all the slots are filled in for each schema, and that the connections between the schemas are valid. To summarize, RISE can surface errors, warnings, and insights; the analyst is free to focus on the interpretation they are modeling, not the syntax of RIST, before returning to fix any issues that might have come up.

5. Discussion

In the introduction we made the case that the study of cognition with ERs should attempt to model the structure of individual’s interpretations of representations in relation to the specific circumstance of their use. This contrasts with much prior work that tend to implicitly assume the occurrence of a single prototypical interpretation of an ER. We contend that readers’ interpretations and the concomitant memory structures may vary for many reasons, ranging from large scale individual differences (e.g., novice–expert differences) down to the local active goal during ongoing problem solving. To address the myriad of interesting cognitive phenomena arising from differences in interpretations we proposed an approach to modeling interpretations of representations that includes

a theory (RIST), modeling notation (RISN), and tool (RISE). For selected cognitive properties (Cheng et al., 2021) of representations, we outlined how RISN models might be used to assess the relative efficacy of different representations, which in turn revealed gaps in our understanding of cognitive factors that impact the difficulty of using representations (Section 3).

The intended application of RIST spans interpretations in which graphical components of ERs have a substantive role in encoding domain relevant information by iconic, spatial, geometric, and topological means. In addition to graphs, charts, tables, diagrams and the like, we include control panels and computer graphical user interfaces among such ERs. For pragmatic reasons, we do not anticipate using RIST to build models of ERs that are essentially propositional (e.g., written natural language) or situations in which the ER encodes little domain content (e.g., mental representations), because such applications would rather trivially set aside much of the functionality that is built into the schemas of RIST. In other words, no fundamental theoretical claim is being made that cognition with graphical ERs is inherently different to propositional thinking, with or without an ER. Rather, we assert that models of such interpretations in RISN would yield few insights that have not already been obtained with more conventional accounts of mental representations.

This section reflects upon four issues: (a) the nature of interpretation modeling as a cognitive modeling activity; (b) the nature of RIST’s schemas in contrast to prior schema accounts; (c) the challenges of building valid models of interpretations; and (d) the purposes to which interpretation modeling may be put.

Clearly, the construction of RISN models is distinct from typical forms of cognitive modeling, such as production systems (e.g., Peebles & Cheng, 2003; Reed, 2019), task analysis, or even verbal models and process flow charts (e.g., Carpenter & Shah, 1998). RISN modeling also contrasts with the rarer models that use only visual representations for inferences, such as Furnas’s (1992) BITPIC and Kunda, McGreggor & Goel’s (2013) affine model of solutions to Raven’s progressive matrices intelligence test. All these focus on the processes used by an agent to transform its current state of knowledge into new states, in the context of task goals. In contrast, the target of RIST models is to understand the structure of the mental representation that a user of an ER builds when using an ER. RIST models are static snapshots of a user’s state knowledge at a given moment, so have parallels to descriptions in the form of traditional propositional or semantic networks, or schema theory. However, RIST differs in that it makes specific claims about the nature of its schemas and how they are connected: all four types of schema coordinate information about domain concepts and objects in the ER. Further, RISE, the editor for RISN models, operationalizes the claims of RIST in the rules that govern the construction of networks. In the manner that cognitive architectures encode theoretical claims about cognitive processes, that in turn constrain the possible models that can be built, RISN and RISE embody theoretical claims of RIST about the nature of interpretations of representations. Of course, one can imagine modeling inferential processes with a succession of RISN models to show the gradual transformation of the network of schemas. However, that is work for the future.

The second issue concerns the nature of the schemas proposed by RIST. A theoretical claim is that four types of schemas are necessary and sufficient to encode the memory structures that are interpretations of external graphical representations. Other accounts of cognition with visual representations have also invoked the idea that specialized schemas are employed by the mind to encode non-propositional information, including Chase and Simon’s (1973) *perceptual chunks*, Gobet &

Simon’s (1996) *templates*, and Koedinger & Anderson (1990) *diagrammatic configurations schemas*. RIST schemas differ from those prior accounts in three ways. First, as RIST aims to be a general theory of interpretation of representations, rather than an account for a particular domain, it theorizes four different types of schemas (Representation, R-scheme, R-dimension and R-symbol) that each play different functions in interpretations, rather than proposing one class of schemas. Second, and of particular importance, is the idea that a core function of the schemas is to associate conceptual (propositional) information about a topic with mental images (visuospatial) information perceived from the external representation. The ability of RISN to express varying types and levels of connection between concepts and the graphical components of an ER gives RISN the flexibility to model subtle variations between different interpretations. The third distinction to the previous schema accounts is that the four schemas constitute a general language for describing interpretations, which may have the potential to model diverse forms of representations and representations that are rich and intricate.

The third issue we consider concerns the validity of RISN models. There are two aspects to validity. One is theoretical – their consistency with the claims of RIST; the other is empirical – how realistically do the RISN models capture users’ mental models? One of the motivations for the development of RISE was precisely to embody the theoretical claims in a tool that constrains the form of models that can be built. The empirical aspect is inherently challenging, because interpretations are mental representations and, as noted in the introduction, interpretations can easily vary if we change the goals of our task. Thus, we are investigating converging methods from multiple angles to address the challenge. The first angle is our recommended method for building models (Section 4.2), which aims to promote precision in the specification of interpretations. Analysts should provide as inputs to the modeling process: (1) a verbal task-oriented description of the representation that matches the given level of experience of the target user; (2) an annotated version of the ER that shows the meaningful graphical objects of the representation at multiple levels of granularity and abstraction. Nevertheless, the quality of a model will depend upon the accuracy of the analyst’s knowledge of what the target users understand. This, of course, introduces an unknown amount of subjectivity into the modeling process. Thus, future research will need to develop methods that will directly elicit users’ interpretations; for example, by having users annotate their own representations whilst giving concurrent verbalizations for later analysis.

Finally, let us consider the potential use of RISN for modeling. Obviously, we assert that research on cognition with ERs should be more explicit about the diversity of interpretations that can exist of any representations, because users and task contexts impact how users read meaning into external representations. In the rep2rep project, mentioned above, we plan to explore whether teachers will be able to produce RISN models, which may enable them to consider the suitability of chosen representations more explicitly for topics in a curriculum. Further, in that project, we are investigating how to use RISN models to quantify the cognitive adequacy of representations in terms of their cognitive properties (Cheng et al., 2021). In the Introduction we listed situations in which users’ interpretations of representations are likely to differ. Obviously, to model alternative readings of an ER by different users, alternative RISN models would be built. In contrast, for interpretations of a single user that evolve over different times scales, one might build sequences of models to capture the interpretations at successive stages. For the process of reading over a short timescale the models might reflect the gradual addition of content into the slots of the schemas and

the instantiation of schemas in the network. For learning over an intermediate timescale, subnetworks from a detailed model might be replaced by single schemas in order to reflect the chunking of components. For models of the development of expertise, alternative models with dramatic restructuring of the overall form for their networks might be built to reflect radical conceptual changes.

We also envisage a potential role for RISN models in the design of representations and information visualizations. Designers might build RISN models in order to systematically question whether particular designs of representations are effective and suitable for the intended audience. If the developer who posted Figure 1 on Stack Overflow had built a RISN model of it, they might have realized its problems themselves and have chosen a better alternative even before asking for coding advice.

Acknowledgements

This work was supported by the EPSRC grants EP/R030642/1, EP/T019603/1, EP/T019034/1 and EP/R030650/1.

References

- Barwise, J., & Etchemendy, J. (1995). Heterogenous logic. In J. Glasgow, N. H. Narayanan, & B. Chandrasekaran (Eds.), *Diagrammatic Reasoning: Cognitive and Computational Perspectives* (pp. 211-234). Menlo Park, CA: AAAI Press.
- Card, S., MacKinlay, J., & Shneiderman, B. (1999). Information Visualization. In S. Card, J. MacKinlay, & B. Shneiderman (Eds.), *Readings in Information Visualization* (pp. 1-34). Mahwah, N.J.: Lawrence Erlbaum Associates.
- Carpenter, P. A., & Shah, P. (1998). A model of the perceptual and conceptual processes in graph comprehension. *Journal of Experimental Psychology: Applied*, 4(2), 75-100.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1), 55-81.
- Cheng, P. C.-H. (1999). Networks of law encoding diagrams for understanding science. *European Journal of Psychology of Education*, 14(2), 167-184.
- Cheng, P. C.-H. (2002). Electrifying diagrams for learning: principles for effective representational systems. *Cognitive Science*, 26(6), 685-736.
- Cheng, P. C.-H. (2004). Why diagrams are (sometimes) six times easier than words: benefits beyond locational indexing. In A. Blackwell, K. Marriot, & A. Shimojima (Eds.), *Diagrams 2004* (pp. 242-254). Berlin: Springer-Verlag.
- Cheng, P. C.-H. (2011). Probably good diagrams for learning: Representational epistemic re-codification of probability theory *Topics in Cognitive Science* 3(3), 475-498.
- Cheng, P. C.-H. (2020a). A Sketch of a Theory and Modelling Notation for Elucidating the Structure of Representations. In: Pietarinen, AV., Chapman, P., Bosveld-de Smet, L., Giardino, V., Corter, J., Linker, S. (Eds.), *Diagrams 2020*. LNCS, vol. 12169. Springer, Cham.
- Cheng, P. C.-H. (2020b). Truth Diagrams versus extant notations for propositional logic. *Journal of Language, Logic and Information*, 29, 121-161.

- Cheng, P. C.-H., Garcia Garcia, G., Raggi, D., Stockdill, A., & Jamnik, M. (2021). Cognitive Properties of Representations: A Framework. In A. Basu, G. Stapleton, S. Linker, C. Legg, E. Manalo, & P. Viana (Eds.), *Diagrams 2021*. (pp. 415-430). Cham: Springer.
- Cheng, P. C.-H., Stockdill, A., Garcia Garcia, G., Raggi, D., & Jamnik, M. (2022). Representational Interpretive Structure: Theory and Notation. Accepted for publication, *Diagrams 2022*.
- Chi, M. T. H., Glaser, R., & Farr, M. J. (Eds.). (1988). *The Nature of Expertise*. Hillsdale, N.J.: L. Erlbaum Associates.
- Egan, D. E., & Schwartz, B. J. (1979). Chunking in recall of symbolic drawings. *Memory and Cognition*, 7(2), 149-158.
- Furnas, G. W. (1992). Reasoning with diagrams only. In H. Narayanan (Ed.), *AAAI Technical Report on Reasoning with Diagrammatic Representations (Report No. SS-92-02)* (pp. 118-123). Menlo Park: CA: AAAI.
- Gobet, F., & Simon, H. A. (1996). Templates in chess memory: a mechanism for recalling several boards. *Cognitive Psychology*, 31, 1-40.
- Gurr, C. A. (1998). On the isomorphism, or lack of it, of representations In K. Marriott & B. Meyer (Eds.), *Visual Language Theory* (pp. 293-306). New York, NY: Springer-Verlag.
- Jamnik, M., & Cheng, P. C.-H. (2021). Endowing machines with the expert human ability to select representations: why and how. In S. Muggleton & N. Chater (Eds.), *Human-Like Machine Intelligence*. Oxford: Oxford University Press.
- Koedinger, K. R., & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science*, 14, 511-550.
- Kosslyn, S. M. (1989). Understanding charts and graphs. *Appl. Cogn. Psychol.*, 3, 185-226.
- Larkin, J. H. & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65-100.
- Larkin, J. H. (1989). Display-based problem solving. In D. Klahr & K. Kotovsky (Eds.), *Complex Information Processing: The Impact of Herbert A. Simon* (pp. 319-341). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Moody, D. L. (2009). The “physics” of notations: toward a scientific basis for constructing visual notations in software engineering. *Software Engineering, IEEE Transactions on*, 35(6), 756-779.
- Peebles, D. J., & Cheng, P. C.-H. (2003). Modelling the effect of task and graphical representations on response latencies in a graph-reading task. *Human factors*, 45(1), 28-45.
- Pinker, S. (1990). A theory of graph comprehension. In R. Freedle (Ed.), *Artificial Intelligence and the Future of Testing* (pp. 73-126). Hillsdale, NJ: Lawrence Erlbaum.
- Raggi, D., Stapleton, G., Stockdill, A., Jamnik, M., Garcia Garcia, G., & Cheng, P. C.-H. (2020). How to (re)represent it? In S. Pan (Ed.), *32nd International Conference on Tools with Artificial Intelligence: IEEE*.
- Reed, S.K. (2019). Modeling visuospatial reasoning. *Spat. Cogn. Comput.*, 19(1), 1-45.
- Stevens, S. S. (1946). On the theory of scales of measurement. *Science*, 103(2684), 677-680.
- Ullmann, J. R. (1976). An algorithm for subgraph isomorphism. *Journal of the ACM*, 23(1), 31-42.
- Zhang, J. (1996). A representational analysis of relational information displays. *International Journal of Human-Computer Studies*, 45(1), 59-74.

Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science*, 21(2), 179-217.