



Introducing the Diagrammatic Semiotic Mode

Tuomo Hiippala¹ and John A. Bateman²

¹ University of Helsinki, Helsinki, Finland
tuomo.hiippala@helsinki.fi

² Bremen University, Bremen, Germany
bateman@uni-bremen.de

Abstract. As the use and diversity of diagrams across many disciplines grows, there is an increasing interest in the diagrams research community concerning how such diversity might be documented and explained. In this article, we argue that one way of achieving increased reliability, coverage, and utility for a general classification of diagrams is to draw on recently developed semiotic principles developed within the field of multimodality. To this end, we sketch out the internal details of what may tentatively be termed the *diagrammatic semiotic mode*. This provides a natural account of how diagrammatic representations organising lines of research integrate natural language, various forms of graphics, diagrammatic elements such as arrows, lines and other expressive resources into coherent organisations, while still respecting the crucial diagrammatic contributions of visual organisation. We illustrate the proposed approach using two recent diagram corpora and show how a multimodal approach supports the empirical analysis of diagrammatic representations, especially in identifying diagrammatic constituents and describing their interrelations in a manner that may be generalised across diagram types and be used to characterise distinct kinds of functionality.

Keywords: Diagrams · Semiotics · Multimodality · Corpora · Annotation

1 Introduction

Diagrams appear to be playing an ever greater range of roles in a similarly increasing range of application contexts. Several authors have consequently called for more efforts to characterise this diversity so as to establish more finely articulated accounts of just what kinds of diagrams there are and how they might serve different communicative and cognitive functions. Norman, for example, considers a two-dimensional characterisation in terms of ‘discretion’ and ‘assimilability’ to distinguish more clearly the role of diagrams among the traditionally drawn categories of descriptions, diagrams, and depiction [18]. Smessaert and Demey offer a typology of types of diagrams used in linguistics, focusing on more

content-oriented ‘linguistic parameters’ and specific semiotically-motivated diagrammatic parameters distinguishing iconic and symbolic representations [21]. Johansen *et al.* offer a typology of mathematical diagrams based on their use by mathematicians, identifying ‘resemblance’, ‘abstract’ and ‘Cartesian’ diagrams [15]. And Purchase and colleagues propose a multidimensional classification of infographics by analysing how users grouped a selection of 60 infographics [19].

Work of this kind raises important research questions, including questions of the particular cognitive capabilities demanded or supported by distinct diagram types [15, p. 107], the ways in which diagram usage has developed and expanded over time [15, p. 106], and how studies might be extended to address entire collections of diagrams, thereby focusing and organising lines of research on a broader scale than hitherto [19, p. 210]. Strengthening the empirical basis for diagrams research by drawing on broader sets of examples and providing more finely articulated characterisations of the properties of diagrams that go beyond existing categories, such as the fundamental distinctions offered by Peirce in terms of iconicity, indexicality and symbolicity and so on, are consequently now well established as aims. As Johansen *et al.* argue, Peirce’s functional definition, particular of iconicity, “is too broad and does not allow for making cognitively and practically meaningful distinctions in the category of diagrams” [15, p. 107]. Nevertheless, developing convincing classifications of diagrams – even in specific areas – has proved challenging.

Both the conceptually-based development of frameworks [8, 15, 18, 21, 24] and more bottom-up clustering based on human judgements [19] continue to face issues of exhaustivity, discriminability, and reliability. Several significant problems are noted by Johansen *et al.*:

“Although such a classification gave a more fine-grained resolution in the diagram classification . . . it turned out to be difficult to carry out in practice . . . Consequently, counting the number of diagrams of a specific type requires making judgements based on the visual appearance of diagrams to classify them correctly.” [15, p. 116]

Indeed, “we also encountered cases where we had to involve the textual or intellectual context of diagrams to classify them, and in other cases, we could only give educated guesses.” [15, p. 116]. Moreover, even in Purchase *et al.*’s ‘user-based’ classification, the Likert scores used to evaluate the infographics were found to be poor predictors of class with “their values bear[ing] little relation to the groupings created by the participants” [19, p. 216]; again, visual grouping of diagrams appeared to give stronger results.

Consequently, despite the generally positive reports and indications of considerable utility of working with diagram collections, it is less clear whether the classifications proposed to date can be scaled-up in a reliable fashion. Certain gaps appear to occur due to the context-dependent nature of any functional categories employed. In this article, therefore, we address these issues from a complementary perspective and argue that the resulting multiply-dimensioned classification promises a more robust approach to diagram classification,

capable not only of respecting visual appearance in a systematic and philosophically sound fashion, but also of providing a principled approach to context-dependent interpretations as well.

This approach draws on classification work in the field of multimodality research, an emerging discipline that studies how communication builds on appropriate combinations of multiple modes of expression, such as natural language, illustrations, drawings, photography, gestures, layout and many more [6]. One product of this work is a battery of theoretical concepts that strongly support *empirical* analysis of complex communicative situations and artefacts. We describe how this can now be applied directly to the analysis of diagrams within the context of the challenges set out above. For this, we define what we tentatively term the *diagrammatic semiotic mode*. This bridges discussions in the diagrams research and multimodality communities by introducing an explicitly multimodal, discourse-oriented perspective to diagrams research. We illustrate this in relation to two recently published multimodal diagram corpora.

2 A Multimodal Perspective on Diagrams

The framework of multimodality adopted here offers a common set of concepts and an explicit methodology for supporting empirical research regardless of the ‘modes’ and materials involved [6]. The result is capable of addressing all forms of multimodal representation, including diagrammatic representations of all kinds. The core theoretical concept within the framework is that of the *semiotic mode*, a graphical definition of which we show on the left-hand side of Fig. 1. Here we see three distinct ‘semiotic strata’ that the model claims are always needed for a fully developed semiotic mode to operate [3].

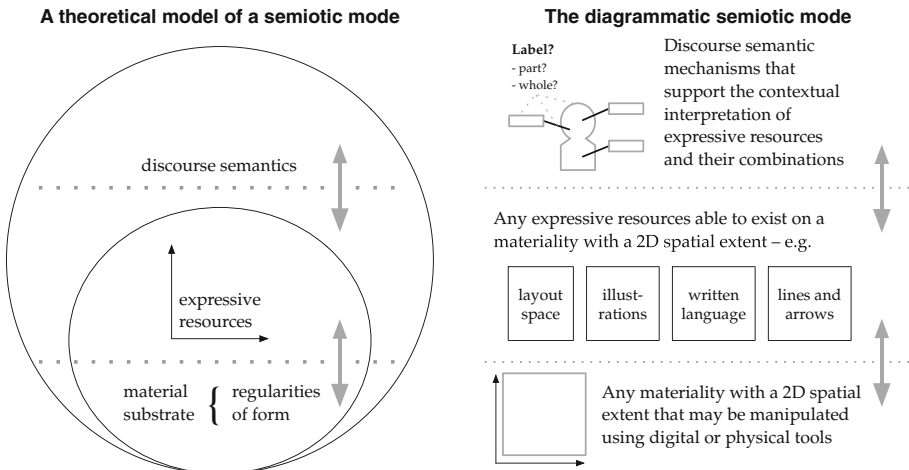


Fig. 1. A theoretical model of a semiotic mode and a sketch of the fundamentals for a diagrammatic semiotic mode [12, p. 408]

Starting from the lower portion of the inner circle, the model requires that all semiotic modes work with respect to a specified *materiality* which a community of users regularly ‘manipulates’ in order to leave traces for communicative purposes; second, these traces are organised (paradigmatically and syntagmatically) to form *expressive resources* that characterise the material distinctions that are specifically pertinent for the semiotic mode at issue; and finally, those expressive resources are mobilised in the service of communication by a corresponding *discourse semantics*, whose operation we show in a moment.

In general, no ordering is imposed on the flow of information across these three strata, although methodologically it can often be beneficial to begin with the more observable material traces. Different semiotic modes also provide differing degrees of constraint at the various levels: for example, whereas the semiotic mode of verbal language offers substantial form-driven constraints guiding discourse interpretation, pictorial semiotic modes often require more discourse constraints when selecting between perceptually plausible readings – Bateman, Wildfeuer and Hiippala [6, p. 33] discuss an example offered by Gombrich [9, p. 7] further from this perspective, showing how variability in interpretation is naturally supported. Finally, the model places no restrictions on the kinds of materiality that may be employed; for current purposes, however, we illustrate the approach by focusing on static two-dimensional diagrams. As Bateman [4] shows, however, the approach generalises equally to both dynamic and 3D cases.

Building on this scheme, we set out on the right-hand side of the figure an initial characterisation of the specific properties of the diagrammatic semiotic mode. The 2D materiality of this mode not only allows the creation of spatial organisations in the form of layout, but is also a prerequisite for realising many of the further expressive resources commonly mobilised in diagrams, such as written language and arrows, lines, glyphs and other diagrammatic elements, which also inherently require (at least) a 2D material substrate. An example of the corresponding expressive resources typical of the diagrammatic mode is offered by the “meaningful graphic forms” identified by Tversky *et al.* [25, p. 222], such as circles, blobs and lines. These can also be readily combined into larger syntagmatic organisations in diagrams such as route maps, as Tversky *et al.* illustrate [25, p. 223]; Engelhardt and Richards offer a similar set of ‘building blocks’ [8, p. 201]. However, theoretically, the diagrammatic semiotic mode can in fact draw on any expressive resource capable of being realised on a materiality with a 2D spatial extent, although in practice these choices are constrained by what the diagram attempts to communicate and the sociohistorical development of specific multimodal *genres* by particular communities of practice [2, 10]. Finally, it is the task of the third semiotic stratum of discourse semantics to make the use of expressive resources interpretable in context.

Embedding expressive resources into the discourse organisations captured by a discourse semantics is crucial to our treatment and a key extension beyond traditional semiotic accounts. Essentially, this enables the account to do full justice to the Peircean embedding of iconic forms within conventionalised usages [22]. It is this addition that explains formally how (and why) fundamental graphic

forms, such as those identified by Engelhardt and Richards, Tversky *et al.* and others, may receive different interpretations in different contexts of use – a problem noted for several of the classifications introduced above – while also allowing certain intrinsic properties of those forms (such as connectivity and directionality) to play central roles in finding interpretations as well. The combination of materiality, expressive forms and discourse interpretations then provides a robust foundation for considerations of diagrammatic reasoning quite generally.

3 Multimodal Diagram Corpora

We now illustrate the potential of a characterisation of diagrams drawing on our multimodal framework for dealing with collections of diagrams by considering two concrete, interrelated diagram corpora: AI2D [16] and AI2D-RST [11]. These corpora build on one another, as AI2D-RST covers a subset of AI2D. We describe the corpora and how they have been characterised and show how an increasing orientation to multimodality successively raises the accuracy and utility of the classification applied. We will argue that the characterisation provided supports a general methodology for building classifications for collections of diagrams.

3.1 The Allen Institute for Artificial Intelligence Diagrams Dataset

The Allen Institute for Artificial Intelligence Diagrams dataset (AI2D) was developed to support research on computational processing of diagrams [16]. AI2D contains a total of 4903 diagrams that represent topics in elementary school natural sciences, ranging from life and carbon cycles to human physiology and food webs, to name just a few of the 17 categories in the dataset. Because the diagram images were scraped from the web using school textbook chapter headings as search terms, the corpus covers a wide range of diagrams created by producers with various degrees of expertise with the diagrammatic semiotic mode, such as students, teachers and professional graphic designers. As the diagrams have been removed from their original context during scraping, little may be said about the medium they originated in. For this reason, it may be suggested that AI2D *approximates* how diagrams are used in learning materials realised using various media.

AI2D models four types of diagram elements: text, blobs (graphic elements), arrows and arrowheads. Although these elements cover the main expressive resources mobilised in these diagrams, no further distinctions are made between drawings, illustrations, photographs and other visual expressive resources [12]. Each diagram in the dataset is nevertheless provided with several layers of description. Instances of the four diagram element types were first segmented from the original diagram images by crowdsourced workers [16, p. 243]. The elements identified during this *layout segmentation* provide a foundation for a *Diagram Parse Graph* (DPG), which represents the diagram elements as nodes and semantic relations between elements as edges. Ten relations are used, drawing from the framework proposed by Engelhardt [7]. Crowdsourcing annotations

is a promising way of creating larger-scale collections of classified diagrams – it also, however, demands that the classifications constructed are sufficiently clear and well-defined to avoid the drawbacks observed by other approaches reported in the introduction above [11, pp. 683–684].

Figure 2 shows as an example the treatment given to a diagram originally scraped from the web, diagram #4210 in AI2D. Below the original shown at the top of the figure, we see the diagram’s crowdsourced layout segmentation and, at the bottom, its corresponding DPG. The original diagram represents a rock cycle, that is, transitions between different types of rock, using a combination of an illustration (a cross-section) whose parts are described using written language. These parts set up the stages of the rock cycle, which are then related to one another using arrows.

For the formation of the AI2D corpus, annotators were instructed to identify units and relationships. As the resulting layout segmentation image in the middle of the figure shows, text blocks and arrowheads were segmented using rectangular bounding boxes, whereas more complex shapes for arrows and various types of graphics were segmented using polygons. The layout segmentation illustrates a common problem with crowdsourced annotations: annotators tend to segment diagrams to quite uneven degrees of detail. Here the entire cross-section is assigned to a single blob (B0), although a more accurate description would be to segment separate parts of the cross-section, such as magma and various layers of rock. We will see shortly how such omissions readily compromise the accurate description of semantic relations in the DPG.

Referring again to the layout segmentation and DPG in the figure, we can see for example that the semantic relations carried by the edges in the DPG cover `ARROWHEADTAIL` between arrow A2 and arrowhead H2 in the upper part of the diagram, which together act as a connector in an `INTEROBJECTLINKAGE` relation between text blocks T1 (‘Magma flows to surface ...’) and T2 (‘Weathering and erosion’). As these relations illustrate, Engelhardt’s [7] relations cover local relations holding between diagram elements that are positioned close to one another or connected using arrows or lines. They neglect, however, the relations needed to describe the global organisation of the diagram, that is, relations between units that are made up of multiple elements [16, p. 239]. Crowdsourcing coherent graph-based descriptions of diagrams thus turns out to be a challenging task. AI2D DPGs often include isolated nodes and multiple connected components, as exemplified by the DPG in Fig. 2. Furthermore, the DPG does not feature a cyclic structure, although the diagram clearly describes a rock *cycle*. The AI2D annotation scheme provides the relation definitions necessary for describing this process in principle, such as `INTEROBJECTLINKAGE` and `INTRAOBJECTREGIONLABEL` [16, p. 239], but annotators following a more shallow visual grouping would not be led to this option from the diagram at hand. As we shall see, this is one of many problems that an explicit discourse semantic orientation can address.

The crowdsourced annotators were not explicitly instructed to decompose cross-sections or other visual expressive resources capable of demarcating meaningful regions, which results in an insufficiently detailed layout segmentation.

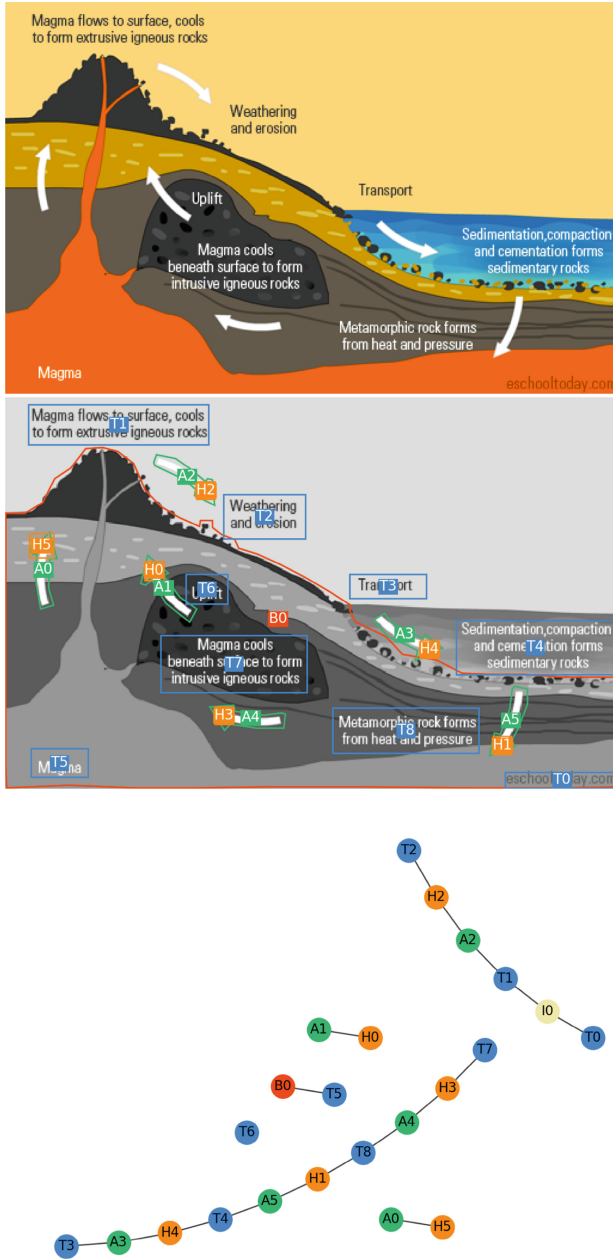


Fig. 2. Original diagram image (top), layout segmentation (middle) and Diagram Parse Graph (bottom) for diagram #4210 in the AI2D corpus. In the layout segmentation, the original image has been converted into grayscale to highlight the crowd-sourced layout segmentation. Each layout segment is coloured according to diagram element type (blue: text; red: blob; arrow: green; arrowhead: orange) and assigned a unique identifier. These colours and identifiers are carried over to the Diagram Parse Graph. (Color figure online)

The blob B0, which covers the entire cross-section shown in the diagram, is as a consequence not segmented into its component parts – i.e., the stages of the rock cycle with labels such as ‘Magma’ (T5) and ‘Metamorphic rock forms from heat and pressure’ (T8) even though each of these picks out a particular region of the cross-section through visual *containment* [7, p. 47] as necessary for defining the stages of the cycle. The cross-section (B0) instead constitutes a single unit and so an otherwise applicable relation such as `INTRAOBJECTREGIONLABEL` cannot be used to pick out corresponding regions simply because those regions are not present in the inventory of identified elements. As such, the description is not sufficiently detailed to represent a cyclic structure.

These challenges relating to decomposing diagrammatic representations relate to the well-known problem of identifying ‘units’ mentioned above and discussed in multimodality theory for many visually-based semiotic modes. In general an annotator, be that an expert analyst or a crowdsourced non-expert worker, will not know on purely visual grounds whether it is necessary, or beneficial, to segment areas presented in a diagram. As we shall see, this is precisely where we need to engage a corresponding notion of discourse semantics for the semiotic mode at issue. The discourse semantics simultaneously supports *decomposing* larger units into component parts and *resolving* their potential interrelations, always with the goal of maximising *discourse coherence* [5, p. 377]. In the next section, we show how this approach can be used for a more effective design of a multimodal corpus of diagrams.

3.2 AI2D-RST – A Multimodally-Motivated Annotation Schema

The second corpus considered here, AI2D-RST, covers 1000 diagrams taken from the AI2D corpus and is annotated using a new schema by experts specifically trained in the use of that schema [11]. The primary goal here was precisely to compare the original corpus, with its style of classification, to a corpus adopting a classification more explicitly anchored into the requirements raised by the diagrammatic semiotic mode. The development of AI2D-RST was motivated by the observation that the AI2D annotation schema introduced above conflates descriptions of different types of multimodal structure, such as implicit semantic relations and explicit connections signalled using arrows and lines, into a single DPG [13]. These can now be separated multimodally so as to better understand how such structures contribute to diagrammatic representations.

To achieve this, AI2D-RST represents each diagram using three distinct graphs corresponding to three distinct, but mutually complementary, layers of multimodally motivated annotations: grouping, connectivity and discourse structure. Figure 3 shows examples of all three graphs for the diagram from Fig. 2. To begin, the grouping layer (top right) organises diagram elements that are likely to be perceived as belonging together into visual perceptual groups loosely based on Gestalt properties [26]. The resulting organisation is represented using a hierarchical tree graph, with grouping nodes with the prefix ‘G’ added to the graph as parents to nodes grouped together during annotation. Such grouping nodes

can be picked up in subsequent annotation layers to refer to a group of diagram elements and thereby serve as a foundation for the description of both the connectivity and discourse structure layers.

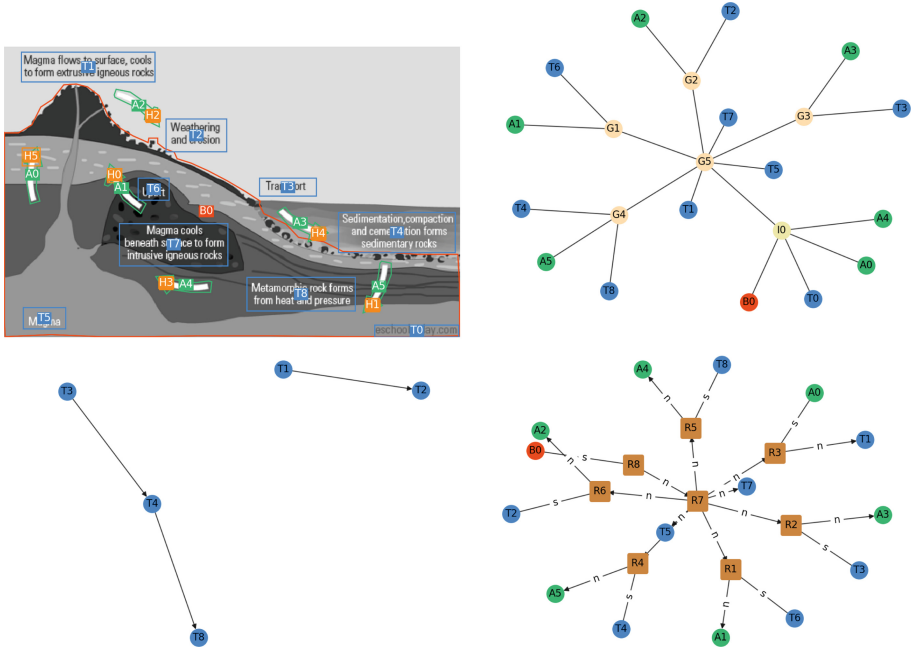


Fig. 3. The original crowd-sourced layout segmentation from AI2D (top left) and AI2D-RST grouping (top right), connectivity (bottom left; with two subgraphs) and discourse structure (bottom right) graphs for diagram #4210. Note that unlike AI2D, AI2D-RST does not model arrowheads as individual units, which is why they are absent from the graphs. This information can be retrieved from the original AI2D annotation if needed.

The connectivity layer (bottom left) is represented using a cyclic graph, in which edges represent visually explicit connections signalled using arrows and lines in the diagram. As the connectivity graph in Fig. 3 shows, it is important that these cover explicit connections only since this reveals the diagram to leave several gaps in its characterisation of the rock cycle, namely between the stages represented using text blocks T7 (‘Magma cools beneath surface ...’) and T1 (‘Magma flows to surface ...’), and between T2 (‘Weathering and erosion’) and T3 (‘Transport’). It is consequently left to the viewer to fill in such connections during discourse interpretation. Not including such connections in the description of connectivity allows us to capture discrepancies between explicit visual signals, such as arrows and lines, and implicit meanings that are only derivable from the discourse structure.

In AI2D-RST, such implicit discourse relations are handled by the third layer, that of discourse structure, which uses Rhetorical Structure Theory (RST) [17, 23] to describe semantic relations between diagram elements. RST was originally developed as a theory of text organisation and coherence in the 1980s [17] and has frequently been applied subsequently to the description of discourse semantics in multimodality research as well [2]. Originally, RST attempted to describe why well-formed texts appear coherent, or why individual parts of a text appear to contribute towards a common communicative goal [23], and so this is a relatively natural perspective to take on diagrams and other forms of multimodal communication. RST defines a set of ‘rhetorical relations’ that are intended to capture the communicative intentions of the designer, as judged by an analyst. AI2D-RST applies these relations to diagrams from the AI2D dataset to provide an alternative annotation schema offering a more multimodally informed description of the intended functions of diagrammatic representations [11].

The relations defined by RST are added to the discourse structure graph of diagrams in the corpus as nodes prefixed with the letter ‘R’ as shown in the graph bottom right in Fig. 3; the edges of the graph describe which role an element takes in the discourse relation, namely nucleus (‘n’) or satellite (‘s’). This notion of *nuclearity* is a key criterion in definitions of semantic relations in RST. Following the original RST definitions, AI2D-RST represents the discourse structure layer using a strict tree graph: if a diagram element is picked up as a part of multiple rhetorical relations, a duplicate node is added to the graph to preserve the formal requirement of tree structure.

In Fig. 3, the specific rhetorical relations in the bottom right graph include IDENTIFICATION (R1–R6), CYCLIC SEQUENCE (R7) and BACKGROUND (R8). Since AI2D-RST still builds on the inventory of diagram elements provided by the original layout segmentation in AI2D, this requires some compromises in the RST analysis. Here the original annotator of the diagram had concluded that most text instances serve to identify what the arrows stand for, namely stages of the rock cycle. The image showing the cross-section (B0), in turn, is placed in a BACKGROUND relation to the CYCLIC SEQUENCE relation. The definition of a BACKGROUND relation [17] states that the satellite (B0) increases the ability to understand the nucleus (R7), which is the top-level relation assigned to the diagram’s representation of the entire cycle.

Although this offers a first incremental step for including discourse information in a diagram corpus, building directly on the original AI2D corpus and its segmentation is also a severe limitation. In fact, this offers only a rather crude description of the discourse structure of the diagram in Fig. 3 because the cross-section B0 is actually providing far more information. This information is crucial for understanding what the diagram is attempting to communicate *but we cannot know that such a decomposition is necessary without considering the rhetorical discourse organisation of the diagram as a whole*. The particular decomposition of diagrams must often be pursued in a top-down direction therefore, emphasising the discourse structure from the outset [5]. Without methodologically prioritising the analysis of discourse structure, it is difficult to know which aspects of

the diagrammatic mode are being drawn on and which elements should actually be included in the description of discourse structure.

This is one of the basic problems underlying several of the limitations discussed for previous diagram classifications above. A visually-accessible cross-section such as the one shown in Fig. 2 is, in fact, very likely to use *illustration* or other expressive resources capable of representing and demarcating meaningful regions in 2D layout space [20]. This possibility makes the question of whether the capability is actually being drawn on pertinent and, if the capability is used, raises further the issue of the extent to which the illustration must be decomposed so as to achieve the inventory of elements needed for making appropriate inferences about the discourse structure. Analytical problems arising from the original layout segmentation are consequently still being propagated from AI2D to AI2D-RST.

3.3 Next Step: Adding Discourse-Driven Decomposition to AI2D-RST

To solve the analytical problems described above, we propose an alternative, discourse-driven layout segmentation that overcomes the limitations discussed above by incorporating the distinctions provided by our definition of a semiotic mode (see Fig. 1). Figure 4 shows a decomposition motivated by discourse structure for diagram #4210, which picks out relevant parts of the cross-section. In contrast to the crowdsourced segmentation in Fig. 2, the cross-section has been decomposed with the goal of maximising the coherence of discourse structure, which involves making available all the elements needed for such a representation of the diagram and its communicative intentions using the AI2D-RST annotation schema.

This is shown in Fig. 4, which applies the AI2D-RST annotation schema to the diagram elements identified through discourse-driven decomposition. When provided with a sufficient inventory of diagram elements, the grouping graph more accurately reflects key structural properties of the diagram. The grouping graph (top right) contains two subgraphs, whose root nodes G10 and I0 correspond to the cross-section and cycle, respectively. Keeping in mind that the grouping graph seeks to capture visual groupings, this already provides a strong cue for two visually distinct configurations, which the AI2D-RST annotation schema refers to as *macro-groups*. These constitute established configurations of the diagrammatic mode that may be flexibly combined in diagrams [11, p. 681]. To summarise, the grouping graph then already pulls these macro-groups apart and provides a foundation for their further analysis. We will shortly show how these macro-groups are integrated in the discourse structure graph.

The connectivity graph (bottom left) reveals that the diagram makes perhaps surprisingly limited use of arrows and lines as an expressive resource despite the intention that the diagram represents a cycle. This is one of the typical complicating factors contributing to the problems for annotation mentioned in the introduction above. The diagram does use arrows to set up connections between some individual elements and their groups, but the connectivity graph

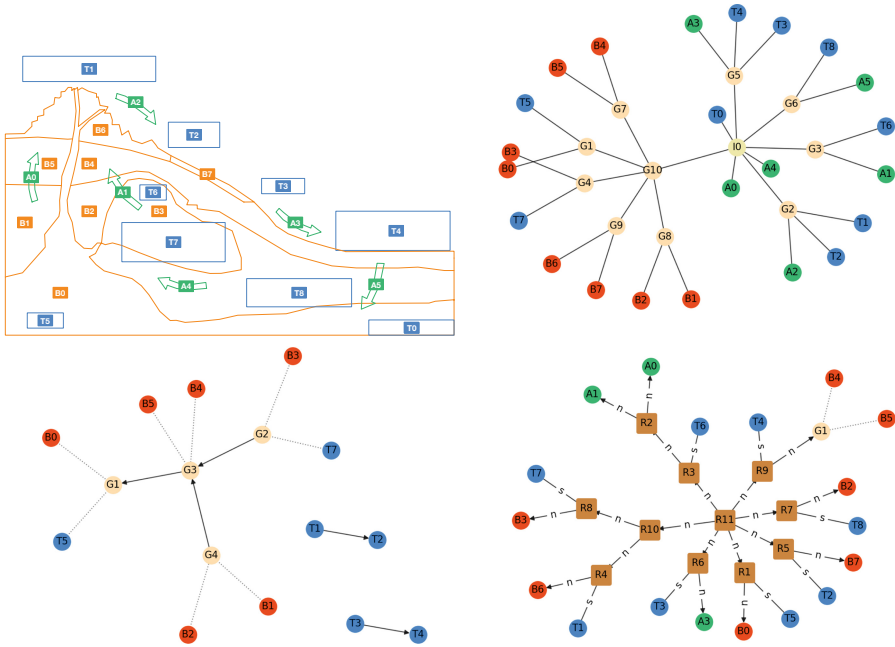


Fig. 4. A discourse-driven decomposition of diagram layout (top left) with grouping (top right), connectivity (bottom left) and discourse structure (bottom right) graphs for diagram #4210.

does not exhibit a cyclic structure. Some arrows, such as A2, have clear sources (T1; ‘Magma flows to surface ...’) and targets (T2; ‘Weathering and erosion’), whereas other arrows, such as A4, do not. This encourages two alternative frames of interpretation for arrows [1]: some clearly signal transitions between stages (A2, A3), whereas others indicate the overall direction of the cycle (A4, A0).

The disconnections in the connectivity graph raise a crucial question: how does an interpretation involving a cyclic structure emerge if it is not clearly signalled using arrows? The answer to this question lies in the discourse structure of the graph as a whole, which here relies largely on *written* language as an expressive resource. This allows the diagram to describe stages of the rock cycle explicitly using clausal structures, e.g. “Metamorphic rock forms from heat and pressure”, but does not express the relationships diagrammatically using arrows. The verbal descriptions are instead placed in relation with specific regions of the cross-section, as shown in the discourse structure graph (bottom right).

The discourse structure graph illustrates how the cross-section and the cycle, which form separate subgraphs in the grouping graph, are tightly integrated in the discourse structure graph, capturing their *joint* contribution towards a shared communicative goal and moving beyond the information in visual grouping alone. The specific rhetorical relations in Fig. 4 and criteria for their application, based loosely on Bateman [2, pp. 149–162], are given in abbreviated form in

Table 1. Beginning from the top of the table, several IDENTIFICATION relations are used to name regions (R1) and arrows (R6, R3). In relation R3, identification is extended to both arrows A0 and A1, which are joined together using the JOINT relation R2. ELABORATION relations R4–R5 and R7–R9, which assign descriptions to specific regions of the cross-section. This explains most of the phenomena depicted in the diagram.

Table 1. Rhetorical relations in the discourse structure graph in Fig. 4

Identifier(s)	Relation	Nucleus	Satellite
R1, R3, R6	IDENTIFICATION	Identified	Identifier
R2	JOINT	No constraints	No constraints
R4–5, R7–9	ELABORATION	Basic information	Additional information
R10	DISJUNCTION	Two or more alternatives	–
R11	CYCLIC SEQUENCE	Repeated steps	–

All of these descriptions contribute towards an interpretation involving a cycle, which requires not only world knowledge, but is also supported using cohesive ties between lexical elements, such as the nouns ‘magma’ and ‘rock’ and the verb ‘to form’. The cycle itself is represented by the CYCLIC SEQUENCE relation R11, which joins together the individual descriptions that form its steps. The cycle also includes two possible alternatives, that is, whether magma cools below or above ground to form rocks, which is also explicitly captured by the DISJUNCTION relation R10 visible in the figure.

This analysis illustrates several of the methodological benefits of adopting a discourse-driven approach to unpacking the structure of diagrammatic representations. We can now move in a principled fashion beyond visual grouping and the individual sources of information in any diagram analysed to produce classifications more sensitive to the likely functions of the diagram as a whole.

4 Discussion

We now briefly discuss some of the principal implications of our analysis for diagrams research more generally. The analysis has shown how a multimodal perspective can yield valuable insights into diagrammatic representations by drawing on the broader basis provided by an appropriately differentiating view of the diagrammatic semiotic mode. Instead of building pre-defined inventories of diagrammatic elements, for example, which are rapidly exhausted when faced with data that do not fall neatly into the categories defined, one can focus more on mapping the expressive resources available to the diagrammatic semiotic mode and describing the kinds of discourse structures they participate in.

This can be approached both empirically and with respect to existing proposals for the graphical elements and properties of diagrams. A recent example of

such a proposal is that of Engelhardt and Richards, who seek to define “universal building blocks of all types of diagrams and information graphics” [8, p. 201]. However, this still excludes “context-related aspects” of diagram use [8, p. 203], which, as we have seen above, can be problematic when characterising larger collections of diagrams. A multimodal perspective is inherently geared towards addressing all of the aforementioned aspects of diagrammatic representations and naturally spans from form to contextually-motivated use. Furthermore, such frameworks can be applied reliably to diagrams, as exemplified by substantial inter-annotator agreement achieved for the AI2D-RST corpus [11, pp. 674–679].

Multimodality research can also contribute towards a deeper understanding of *signification* in diagrams, as this is precisely the work that expressive resources perform as part of the diagrammatic mode. As our analysis shows, diagrams that represent cycles do not necessarily need to draw on arrows for this purpose: the diagrammatic mode provides alternatives, such as written language, whose structural features (here: cohesive ties) may be used to cue a discourse semantic interpretation involving cyclicity. This allows a fine-grained decomposition of the proposed building blocks of diagrammatic representations [8, 14]. Conversely, multimodality research is likely to benefit from the concepts developed in diagrams research for producing systematic descriptions of expressive resources. This will, however, require a significant effort in triangulating what has been done previously in multimodality and diagrams research, and aligning their theoretical concepts as necessary. Previous approaches to diagram classification as described above are the logical place to start such investigations.

Finally, our findings also carry implications for the computational modelling of diagrams. In particular, problems with the AI2D annotation [16] echo the need remarked on for mathematical diagrams above by Johansen and colleagues for domain expertise in describing the diagrammatic mode in order to achieve a description that respects its specific features. When applied to diagrams, computer vision tasks such as instance-level semantic segmentation and visual question answering must acknowledge particular characteristics of the diagrammatic mode. They should not be based simply on assumptions concerning how such tasks are defined for processing pictorial representations, since pictures constitute a quite different family of semiotic modes and exhibit rather different properties. Particularly important here is the issue of the appropriate level of semantic segmentation, that is, to what extent the mode in question needs to be decomposed into its components. Developing appropriate descriptions of the diagrammatic mode for computational modelling is therefore a task that needs to involve research communities working on both diagrams and multimodality.

5 Conclusion

We have introduced a multimodal perspective on diagrammatic representations, and presented a description of the diagrammatic semiotic mode, exemplifying the proposed approach using two recent multimodal diagram corpora. Multimodal analysis involves decomposing diagrammatic representations into their

component parts, and we have argued for supporting decompositions driven by discourse structure – that is, what the diagrammatic representations attempt to communicate and how their organisations explicitly guide readers to candidate interpretations. Capturing segmentations of this kind explicitly in appropriately designed corpora ensures that the necessary diagrammatic elements are available for further analysis. We suggest that given the widespread use of diagrams and their variation in different domains, an extensive programme of corpus-driven research of the kind we have proposed is now essential for developing an empirically-motivated account of the diagrammatic semiotic mode.

References

1. Alikhani, M., Stone, M.: Arrows are the verbs of diagrams. In: Proceedings of the 27th International Conference on Computational Linguistics, Santa Fe, New Mexico, USA, pp. 3552–3563 (2018)
2. Bateman, J.A.: *Multimodality and Genre: A Foundation for the Systematic Analysis of Multimodal Documents*. Palgrave Macmillan, London (2008)
3. Bateman, J.A.: The decomposability of semiotic modes. In: O'Halloran, K.L., Smith, B.A. (eds.) *Multimodal Studies: Multiple Approaches and Domains*. Routledge Studies in Multimodality, pp. 17–38. Routledge, London (2011)
4. Bateman, J.A.: Dimensions of materiality: towards an external language of description for empirical multimodality research. In: Pflaeging, J., Wildfeuer, J., Bateman, J.A. (eds.) *Empirical Multimodality Research: Methods, Evaluations, Implications*, pp. 35–64. De Gruyter, Berlin and Boston (2021)
5. Bateman, J.A., Wildfeuer, J.: Defining units of analysis for the systematic analysis of comics: a discourse-based approach. *Stud. Comics* **5**(2), 373–403 (2014)
6. Bateman, J.A., Wildfeuer, J., Hiippala, T.: *Multimodality: Foundations. Research and Analysis - A Problem-Oriented Introduction*. De Gruyter Mouton, Berlin (2017)
7. Engelhardt, Y.: *The language of graphics: a framework for the analysis of syntax and meaning in maps, charts and diagrams*. Ph.D. thesis, Institute for Logic, Language and Computation, University of Amsterdam (2002)
8. Engelhardt, Y., Richards, C.: A framework for analyzing and designing diagrams and graphics. In: Chapman, P., Stapleton, G., Moktefi, A., Perez-Kriz, S., Bellucci, F. (eds.) *Diagrams 2018*. LNCS (LNAI), vol. 10871, pp. 201–209. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-91376-6_20
9. Gombrich, E.: *Art and Illusion: A Study in the Psychology of Pictorial Representation*. Pantheon Books, New York (1960)
10. Hiippala, T.: *The Structure of Multimodal Documents: An Empirical Approach*. Routledge, New York and London (2015)
11. Hiippala, T., et al.: AI2D-RST: a multimodal corpus of 1000 primary school science diagrams. *Lang. Resour. Eval.* **55**(3), 661–688 (2020). <https://doi.org/10.1007/s10579-020-09517-1>
12. Hiippala, T., Bateman, J.A.: Semiotically-grounded distant view of diagrams: insights from two multimodal corpora. *Digit. Scholarsh. Human.* **37**(2), 405–425 (2021). <https://doi.org/10.1093/llc/fqab063>
13. Hiippala, T., Orekhova, S.: Enhancing the AI2 diagrams dataset using rhetorical structure theory. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018), pp. 1925–1931. European Language Resources Association (ELRA), Paris (2018)

14. Hullman, J., Bach, B.: Picturing science: design patterns in graphical abstracts. In: Chapman, P., Stapleton, G., Moktefi, A., Perez-Kriz, S., Bellucci, F. (eds.) *Diagrams 2018*. LNCS (LNAI), vol. 10871, pp. 183–200. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-91376-6_19
15. Johansen, M.W., Misfeldt, M., Pallavicini, J.L.: A typology of mathematical diagrams. In: Chapman, P., Stapleton, G., Moktefi, A., Perez-Kriz, S., Bellucci, F. (eds.) *Diagrams 2018*. LNCS (LNAI), vol. 10871, pp. 105–119. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-91376-6_13
16. Kembhavi, A., Salvato, M., Kolve, E., Seo, M., Hajishirzi, H., Farhadi, A.: A diagram is worth a dozen images. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9908, pp. 235–251. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_15
17. Mann, W.C., Thompson, S.A.: Rhetorical structure theory: toward a functional theory of text organization. *Text* **8**(3), 243–281 (1988)
18. Norman, J.: Differentiating diagrams: a new approach. In: Anderson, M., Cheng, P., Haarslev, V. (eds.) *Diagrams 2000*. LNCS (LNAI), vol. 1889, pp. 105–116. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-44590-0_13
19. Purchase, H.C., et al.: A classification of infographics. In: Chapman, P., Stapleton, G., Moktefi, A., Perez-Kriz, S., Bellucci, F. (eds.) *Diagrams 2018*. LNCS (LNAI), vol. 10871, pp. 210–218. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-91376-6_21
20. Richards, C.: Technical and scientific illustration: picturing the invisible. In: Black, A., Luna, P., Lund, O., Walker, S. (eds.) *Information Design: Research and Practice*, pp. 85–106. Routledge, London (2017)
21. Smessaert, H., Demey, L.: Towards a typology of diagrams in linguistics. In: Chapman, P., Stapleton, G., Moktefi, A., Perez-Kriz, S., Bellucci, F. (eds.) *Diagrams 2018*. LNCS (LNAI), vol. 10871, pp. 236–244. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-91376-6_24
22. Stjernfelt, F.: Diagrammatology. An investigation on the borderlines phenomenology, ontology, and semiotics. *Int. J. Semiot. Law* **21**, 297–301 (2007). Springer, Dordrecht
23. Taboada, M., Mann, W.C.: Rhetorical structure theory: looking back and moving ahead. *Discourse Stud.* **8**(3), 423–459 (2006)
24. Tversky, B.: Diagrams: cognitive foundations for design. In: Black, A., Luna, P., Lund, O., Walker, S. (eds.) *Information Design: Research and Practice*, pp. 349–360. Routledge, London (2017)
25. Tversky, B., Zacks, J., Lee, P., Heiser, J.: Lines, blobs, crosses and arrows: diagrammatic communication with schematic figures. In: Anderson, M., Cheng, P., Haarslev, V. (eds.) *Diagrams 2000*. LNCS (LNAI), vol. 1889, pp. 221–230. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-44590-0_21
26. Ware, C.: *Information Visualization: Perception for Design*, 3rd edn. Elsevier, Amsterdam (2012)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

