Modeling the 'Unthought'

Chiara Palladino[1] ⓘ   Andreas Kuczera[2] ⓘ

Iian Neill[3]

[1] Furman University, Greenville, SC, USA
[2] University of Applied Sciences (THM), Gießen, Germany
[3] Academy of Sciences and Literature | Mainz, Germany

Abstract

This paper introduces Codex, a digital environment for modeling the practice of scholarly annotation of textual information. Codex is a text-as-graph solution that integrates a text-as-graph meta model via standoff property annotations. Codex has three distinctive characteristics. First, all standoff property annotations are connected to the user who created or edited them. Second, the affordances of standoff annotation in Codex's real-time standoff property text editor module (SPEEDy) make it convenient for users to overlay and overlap annotations, allowing for a multitude of interpretations to be captured without conflict. Third, the Statement meta model enables the user to formulate RDF-like assertions which can be applied to multiple entities. The use of trait statements constitutes a move from a 'class-based' ontology toward an 'aspect-oriented' ontology of discrete traits or qualities, which produces a classification that is both finer and more capable of functioning as a doxography that connects ontological claims to the historical sources from which they are derived. In this paper, we demonstrate a number of practical applications for Codex, focusing on how this solution contributes to the pragmatics of modeling the scholarly understanding of textual information. Using the Statement meta

model is crucial: it allows the modeling of recurring structural patterns of information by keeping their semantic differences intact, which in turn enables the researcher to keep the annotation model close to the research object and the associated research questions. Moreover, the flexibility of the graph extends to modeling hierarchies and combinations that may not have been thought of before, making it possible to 'model the unthought.' The final section concerns the usefulness of this model for the understanding of unstructured 'narrative' information and highlights its potential for the humanities and textual research in general.

1 Introduction: Modeling the 'Unthought'

Our title stems from the observation that there are certain types of information that cannot easily be incorporated into hierarchical structures. This information often originates from non-structured sources, such as textual narratives, wherein multiple meanings are possible: there can always be another interpretation that we have not anticipated – a yet unrealized possibility that we like to call the 'unthought.'

In the digital humanities, we model information to create formalized representations of non-formalized, non-structured data (McCarty, 2005). By these means, we are able to make explicit structures of knowledge that are usually left implicit in the sources being modeled. The formalization that occurs as a result must be able to support computational reasoning in order to ensure that machines can process and analyze the gathered information. In this sense, computational manipulability is an indispensable prerequisite for such models: they are useful for research precisely because they can be queried, inspected, and visualized in several ways that provide new insights into the material (Eide and Ore, 2019).

Modeling, however, is also a research method that allows scholars to investigate, interpret, and better understand the nature of a source through a systematic approach. The transformation of the information encoded in a textual/documentary medium into a logico-mathematical model implies a high level of abstraction, necessitating the use of a "deep structure" (Buzzetti, 2002, p. 76) that reifies superficial structures and their reciprocal connections. For this reason, modeling processes that involve textual and qualitative data generally rely on meta models, which represent abstract knowledge structures that support the scholarly process.

One meta model frequently used in the digital humanities is the graph

(Flanders and Jannidis, 2019).[1] Graphs represent knowledge in the form of relationships (also named edges) between entities (also named nodes), and are particularly appropriate for the representation of strongly networked information. At the core of this modeling approach is the idea that knowledge can be represented as a hypergraph, wherein classes of things are reciprocally connected by semantically meaningful relationships. The graph metaphor is also useful for representing the multi-dimensionality of textual data and the variety of its accompanying interpretations. Here the idea of multi-dimensional connectivity does not necessarily imply that we think of a knowledge domain visually as a graph, but rather offers us a way to map and model the knowledge that it expresses in a sophisticated way by emphasizing and empowering its relational nature.

While graphs offer a suitable meta model according to which we can map out our sources, graph databases provide a method for storing and querying the resulting relationships and object classes. Consequently, graph databases are rapidly gaining momentum among scholarly research communities and industry professionals, as they support the representation of multi-dimensional systems whose relationships are semantic in nature, and therefore exhibit significant potential for modeling qualitatively complex information (Robinson et al., 2013).

Finally, graphs ensure semantic flexibility: whereas the graph offers a generic abstract model for describing the data in the form of relational statements composed of nodes and edges, it does not provide any domain-specific vocabulary for describing classes of things and the nature of their relations. Vocabularies are usually provided by external schemas or ontologies, although they can also be built by the scholar who is modeling the source. In other words, they allow the modeling of new interpretations and of previously unthought structures.

2   Codex

Codex is a web-based annotation environment that leverages the graph meta model to support the modeling of scholarly annotation of textual information. The foundational idea behind Codex is that information expressed in linguistic form is constituted by structural patterns of various kinds (spatial, narrative, etc.), and can be represented and investigated through the

---

[1]The idea of using graphs to represent humanities data is not new. It quickly gained a foothold in social network analysis and literary studies, in the wake of Franco Moretti's groudbreaking work (Moretti and Piazza, 2007; Moretti, 2011), but graphs have also seen use as a meta model to investigate and represent a wide variety of other knowledge structures, including neural networks, street maps, textual transmission, GIS and spatial modeling, and even as a new paradigm to model phenomena in ancient history (Malkin, 2011).

paradigm of multidimensional graphs.

2.1 Standoff Properties

Codex uses standoff properties to model textual annotations, based on concepts developed by Desmond Schmidt (Schmidt and Colomb, 2009). Standoff properties were chosen primarily for the simplicity they offer in recording freely overlapping annotations. They differ from markup languages like XML in that they are stored externally to the text and need not conform to a context-free grammar.[2] This enables the annotator to mix annotation types from different vocabularies, combining, for example, TEI codes with those of other domains, such as syntax, stylistics, custom semantics, etc. In addition, standoff is used to represent not only semantic encoding, but also style, presentation, and layout. A special tool called the Standoff Property Editor (SPEEDy) was devised for Codex to enable the user to mark up documents with standoff annotation and modify the text stream freely without invalidating the standoff property character positions. SPEEDy also interfaces with the Codex API to create entities in the Neo4j graph database corresponding to the semantic annotations added to the document.

2.2 Annotation Layers

As standoff properties are exogenous to the text stream and reference text ranges by character position indexes, one may think of properties which overlap the same text range as being layers of interpretation – provided, of course, that we regard annotation as an act of interpretation. These layers can be used to capture different perspectives on annotation which otherwise might be excluded from an inline markup process, either as a result of the definitions of the XML schema employed, or due to the technically complex nature of representing overlap in XML. For example, if two editors disagreed about the place signified by the deictic 'there' in a text, they could annotate their differing interpretations over the same text range without creating conflicting markups. Because the annotator's user identifier in Codex is recorded on each standoff property they create or modify, it becomes possible to separate and visualize interpretational layers by user.

2.3 Statements

A Statement in Codex is a data structure that enables the annotator to define a relationship between two or more entities within a given context. While an edge is sufficient for representing a relationship between two nodes, more complex relationships involving more than two entities require a hyperedge,

---

[2]https://github.com/argimenes/standoff-properties-editor

that is, an edge in a hypergraph model that connects at least two nodes. An example of an edge is a relationship like `subset of`, which is between two nodes only, whereas an example of a hyperedge in Codex is the `received` relationship which can involve a sender, a receiver, an object, a destination, etc. When this Statement hyperedge relates to an event or action, it is called an Event Statement – the assumption being that an event can be considered to be a kind of relationship. An event can be thought of as a relationship in its temporal context. For example, we might say that Andrea del Verrocchio taught Leonardo da Vinci and is thus connected to the latter by the `teacher of` relationship. However, this relationship is in itself a byproduct of a series of 'teaching actions,' as a result of which consensus at some point bestows the status `teacher of`. In Codex, the model that can be used to represent such 'teaching actions' is the Event Statement. The term 'statement' was chosen above a term like 'claim' as the latter can be taken to suggest (at the least) that the author intended a remark to be in some way 'factual'; whereas a 'statement' can be thought of as merely a formulation which is not inherently true or false. In this way, Codex contains a collection of hundreds of Event Statements drawn from remarks in primary source documents rather than 'claims' – which may be true or false – and thus leaves the evaluation of veracity to the annotator or indeed to the reader (or rather, the user) of the text corpus.

3   An Example of Modeling the 'Unthought'

After transcribing Robert W. Carden's edition of Michelangelo's letters (Carden, 1913) in Codex, Iian Neill applied various annotation layers such as 'persons,' 'places, 'concepts,' 'traits,' 'sentiment,' parts of speech, and morphosyntactics to the texts in question. Figure 1 shows one of the letters with highlighted sentiment and syntax analysis. All of these different annotation layers are combined in the graph model. In Figure 2, we see the entity tab of Codex, which displays all annotated entities in the letter, including pronouns. These pronouns are annotated in a semi-automatic way, with Google syntax analysis nominating a pronoun as a candidate for annotation and the user then assigning an entity reference. Once the pronoun has been annotated, the reference can be copied across to other pronoun co-references in the document, thus combining manual and automatic annotation methods. In this manner, all references to an entity in a given document are captured, whether by name or by pronoun.

Two major advantages of this method are that one can find every reference to an entity in the graph even if that entity is not directly named, and that the accuracy of entity statistics in a text corpus is significantly increased. To take

a random example from the Codex edition of the Michelangelo Letters, in the letter from Michelangelo to his brother Buonarroto dated 5 September 1512, Michelangelo's name appears only once, in the letter's signature. But when all of the pronouns had been annotated, it became clear that there were 14 references to Michelangelo in total, which in turn constitutes 34% of all the references to various people found in the letter.

Bearing these capabilities of overlapping syntax and semantics in mind, we asked ourselves if we could find a sentence in the corpus where Michelangelo talks about himself in the second person. Figure 3 shows the interface for this query, and Figure 4 presents the query results. Two such sentences were found, both occurring in a letter in which Michelangelo remonstrates with his brothers, Buonarroto and Giovansimone:

> [I]f thou hadst loved me in very truth thou wouldst have written to me after this manner: 'Michelagniolo, spend the three thousand crowns on *yourself* in Rome, for *you* have given us so much that we already have sufficient for our needs. We have more care for *your* person than for *your* property.'

Thanks to the coexistence of syntax and semantic layers in one system, it becomes possible to solve this exemplary research problem in the matter of a few seconds. However, when we originally added semantic and syntax annotations to the corpus, we did not set out to answer questions like this – the ability to query for overlapping annotations led us to this discovery only after the fact. As it turned out, we had inadvertently 'modeled the unthought.'

4   Ancient Spatial Narratives in Codex

Spatial narratives represent a perfect example of implicit knowledge structures that we can model with Codex. The description of landscapes and way-finding processes in textual narratives is, by definition, much less structured than the cartographic representation of a region or itinerary through an established grid of coordinates. Premodern spatial narratives, in particular, are based on strongly non-cartographic cognitive processes, and challenge our prevalently Cartesian perspectives in many ways (Palladino, 2016; Barker et al., 2016).

Spatial narratives, however, can be considered as knowledge systems with their own concepts and paradigms, which are encoded through specific linguistic and semantic tools (Thiering and Geus, 2014). We can define these tools as regular expressive patterns describing different types of functional information related to space and navigation. In this sense, 'modeling the unthought' means to make explicit thought processes that are largely implicit

in premodern navigational systems – whether it be the semantic typologies of these patterns, the different kinds of spatial relations they express, or the formal-linguistic constructs used to convey this information.

To clarify this idea, let us consider a passage from the *Stadiasmus of the Great Sea*, the earliest portolan known to be in existence (Arnaud, 2017):

114. Ἀπὸ Λέπτεως εἰς Θερμὰς στάδιοι ξ΄· κώμη ἐστί· τὸν δ᾽ αὐτὸν τρόπον καὶ ὧδε βράχη ἐστὶ δυσκατάγωγα. 115. Ἀπὸ Θερμῶν πλεύσας σταδίους μ΄ ὄψει ἀκρωτήριον ἐπ᾽ αὐτῷ ἔχον δύο νησία ἐσκολοπισμένα· ὕψορμός ἐστιν. 116. Ἀπὸ τοῦ ἀκρωτηρίου ὄψει Ἀδραμύτην τὴν πόλιν ἀπὸ σταδίων μ΄· ἀλίμενος. 117. Ἀπὸ Ἀδραμύτου ἐπὶ τὴν Ἀσπίδα στάδιοι φ΄· ἀκρωτήριόν ἐστιν ὑψηλὸν καὶ περιφανές, οἷον ἀσπίς. Ἐπ᾽ αὐτὴν πλέε τὴν ἄρκτον [ὡς] παραφαίνειν ἐξ εὐωνύμων.

114. From Leptis to Thermae sail for 60 stades; it is a town, and here in the same way the shoals make the sailing difficult. 115. From Thermae sailing 40 stades you will see the promontory against which are two islands staked out with pilings; there is an anchorage. 116. From the promontory you will see the city of Adramyte, distant 40 stades: it does not have a harbor. 117. From Adramyte to Aspis 500 stades; it is a high and conspicuous promontory, shaped like a shield. From there sail north [so that] it appears on the left.

Some of the key concepts of this passage are relatively easy to isolate. First, the core of any spatial description is an open class of spatial entities, which include place names (Aspis, Thermae), as well as anaphoric, deictic, and pronominal structures ('here,' 'which'), or common nouns defining locations ('town,' 'promontory'). The text is structured through a set of relations connecting these entities: these are mainly distances (including expressions of extent) and indications of orientation. Both of these categories appear to be deceptively simple to classify, when in fact distances can be expressed in a wide range of formats, with enormous variations that depend on several factors, including the mode of transportation, the season, the type of route, the context, etc. (Palladino, 2016). Expressions of orientation, on the other hand, seem to follow two main reference frameworks: they can function through absolute reference points, such as the compass rose or the movement of the sun ('sail north'), or through relative reference points ('against which,' 'on the left'). Their expressive structures, however, may vary widely, and represent one of the most difficult aspects of spatial descriptions to model in a controlled framework. Equally important are the semantic categories that are projected onto spatial entities, which can be expressed as place types

('the city of Adramyte'), various kinds of qualities ('a high and conspicuous promontory,' 'there is an anchorage') or even abstract concepts ('shaped like a shield'). Finally, a scholar may be interested in including some real-world information through the association of entities to concrete locations.

Spatial narrative is a multi-dimensional system, where multiple layers of information interact with each other to create meaning, which is precisely what makes the graph structure so well-suited for modeling this kind of semantically complex network. In the following sections, we will show how it is possible to annotate and model this information using the graph structure provided by Codex.

4.1 Modeling Spatial Relations with the Statement Meta Model

In Codex, we use the Statement meta model to annotate spatial information gathered from texts. A Statement usually takes the form of an event, represented as a verb phrase with prepositional agents (Neill and Kuczera, 2019). The idea is that any information in a text can be modeled as a semantically classified set of relations between different kinds of entities: the entities are represented as nodes fulfilling a specific function in the relation (e.g. as subject or object), and the Statement can be associated to a semantic label expressing its function (e.g. a distance).

Typically, spatial relations connect two features of the landscape (Place A and Place B) through different types of functional linking, such as distances or directions. Distances can be expressed as connection between two places: the place from which the distance is calculated and the destination. However, the model is almost never as simple as (From Place A to Place B) + (Distance/Direction). Additional semantic information is contributed by the value of the distance (when present) and the unit expressing it. Moreover, grammatical structures, like verbal or adverbial constructs, can contribute important information, e.g. the transportation mode ('sailing,' 'walking'), the type of position ('to lie,' 'to be adjacent'), and so on. In other words, spatial relations in texts have more in common with hyperedge structures that connect different functional nodes than with simple subject-predicate-object triples (see above).

Figure 5 shows how a relation of the type 'distance' can be modeled in the graph structure of Codex. The sentence, as it appears in the text, is modeled as a :Statement node, or Event, and the participating entities are linked to it as agent nodes with different labels, depending on their function in a given context (e.g. :Place). Relationships between the nodes are semantically labeled in the graph structure according to the role of each node: for example, in a distance relation, the place from which the distance is calculated acts as a Sub-

ject, and the destination acts as an Object. Additional concept nodes indicate other information, such as the distance quantification and measurement unit ('40 stades') and the mode of transport ('sailing'). The entire statement is then classified through a `:Relation-type` node, which is labeled according to an external or user-defined vocabulary. Figure 6 shows how these relationships appear in the graph database, where they can be queried by the scholar for further investigation.

Directions work in a similar way, with the additional issue that they tend to follow different reference systems, since they can use a variety of reference points as landmarks for navigation. The choice of such features depends on a variety of (potentially subjective and at least partly unpredictable) cultural and environmental factors. We indicate these relationship types as direction_absolute and `direction_relative`, with further specifications added if and when available. Figure 7 shows examples of how directional relations appear in the Codex graph database.

4.2 Use of the Statement for Conceptual Relations

The projection of cultural or conceptual categories on spatial entities can take the form of simple place-type declarations ('it is a town'), indications of geographical characteristics ('it does not have a harbor'), or culturally influenced conceptual categorizations ('shaped like a shield'). Although the relational nature of these statements is not empirical like in navigational relations, they can still be modeled as relational statements between nodes expressing places and concepts.

Figure 8 shows all the conceptual statements referring to the place Aspis in our passage, classified according to different semantic types. The structure of the statement is similar to distances and directions: the subject node is the place to which the concept is associated (Aspis), while the node expressing the `:Concept` is modeled as the statement name. Meanwhile, a `:Type` node is associated to the entire statement, functioning as a semantic label that indicates the nature of the relationship between place node and concept node.

4.3 External Properties

Finally, external properties are used to add further information concerning specific locations in order to improve data mining and support map-based visualizations of the network. External properties in Codex follow the usual key-value format of graph databases (the property latitude, for example, has the value '36.837800,' as can be seen in Figure 10).

5 Conclusion

In the previous sections of this paper, we have shown how Codex and the underlying graph model can be used to annotate and investigate a diverse range of data, including grammatical, expressive, and spatial information. In a machine-actionable model, this information typically constitutes different layers of data. In the original non-digital source, however, every part of every layer overlaps and interacts: grammatical information appears in expressions of sentiment; external coordinates frame places in spatial relations; pronouns overlap with semantic information; and so on. Thanks to the graph model, it is possible to represent this overlapping structure and its interactions. In other words, Codex offers a number of tools for representing this multi-dimensionality in a practical way. Standoff properties allow individual users to alter and improve their annotations over time, and to capture any associated changes. The non-destructive nature of standoff annotation, which does not change the text stream, greatly facilitates collaboration amongst scholars. Standoff annotation also enables the free mixing of annotation schemata, which are normally constrained by the hierarchical grammars of XML documents. One can combine (or overlay) annotation types from TEI-XML, corpus linguistics, syntax analysis, and so on. What is more, standoff is also used to anchor Event Statements to their textual origins, thereby bringing the graph model and the text stream into a kind of interoperability which allows the user to pass from text to graph model and vice versa. As we have demonstrated in this paper, the Event Statement is a flexible, extensible, yet structured model for capturing the manifold morpho-syntactical and semantic relationships encountered in complex texts. We believe that the combination of these tools gives researchers increased flexibility in modeling their scholarly process, and therefore a way to bring the 'Unthought' in their research to the surface. Scholars are often confronted with a situation where they cannot store every piece of information that they come across in their research process. Codex, with its combination of standoff properties and the multidimensional power of the graph, makes it possible to freely annotate texts and combine various annotation hierarchies and modeling perspectives. In this way, the full range of knowledge gathered during the research process can be preserved and made accessible to other researchers and research communities.

Figures

From Rome, August 21ˢᵗ, 1563.

To Lionardo di Buonarroto Simoni,/

in Florence.

LIONARDO, — I see from thy letter that thou has lent thine ear to certain envious and rascally persons who, finding they can neither rob nor deceive me, have written thee a lot of lies. They are a gang of greedy robbers, and thou art a fool to listen to what they tell thee about me, as though I were a baby. Drive them from thy sight, like the envious scandal-mongers and evil livers they are. As to the unsatisfactory manner in which I am being looked after and the other things of which thou writest; as to the former, I tell thee I could not

☑ sentiments

☐ Adp   ☐ Punct   ☐ Adj   ☐ Pro
☐ Pron   ☐ Verb   ☐ Conj   ☐ X
☑ Noun   ☐ Adv   ☐ Det   ☐ Num

Figure 1: Segment of text with various annotation layers
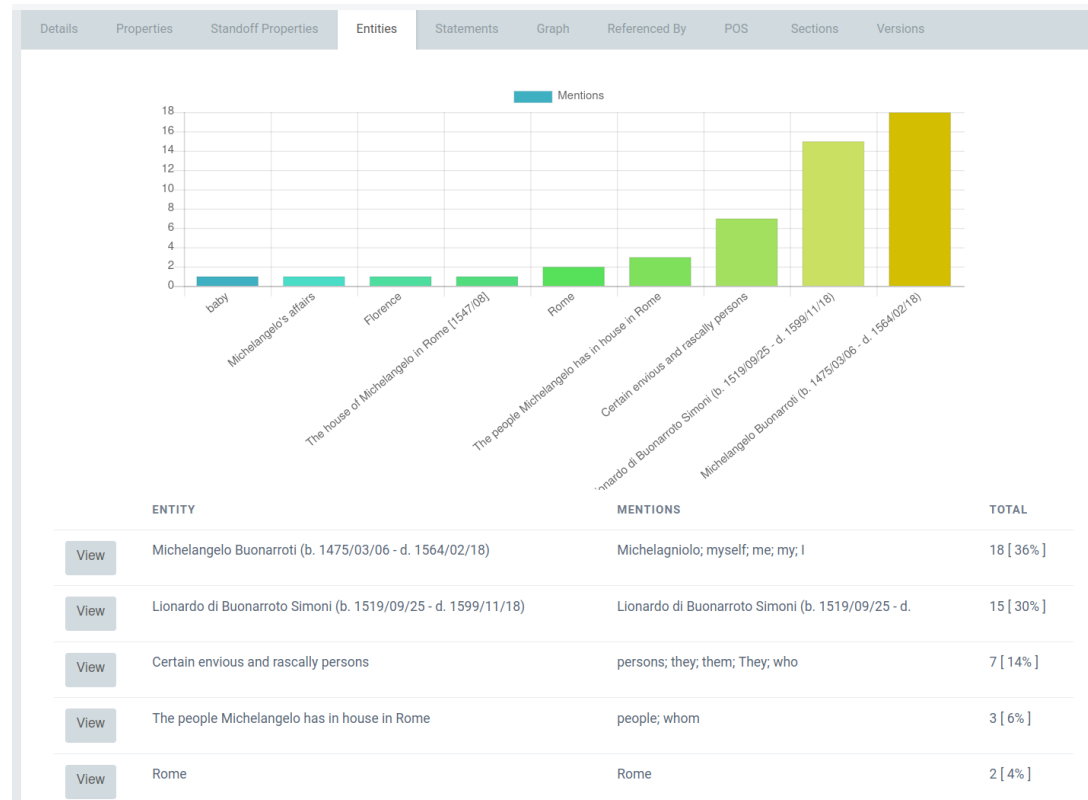
Figure 2: List of entity statistics from a single document

Figure 3: Interface for querying the second person of the entity
Michelangelo

Part of Speech

Person: Second

Page | Rows | Sort by | Sort direction
1 | 10 | By Name | ○ ○

Clear | Search

TEXT | **SENTENCES**

View | 1546/02/06: To Lionardo di Buonarroto Simoni

Characters 845-1077

View | Edit

If thou hadst loved me in very truth thou wouldst have written to me after this manner: "Michelagniolo, spend the three thousand crowns on yourself in Rome, for you have given us so much that we already have sufficient for our needs.

Characters 1079-1136

View | Edit

We have more care for your person than for your property."

Figure 4: Sentences in which Michelangelo talks to himself in the second person
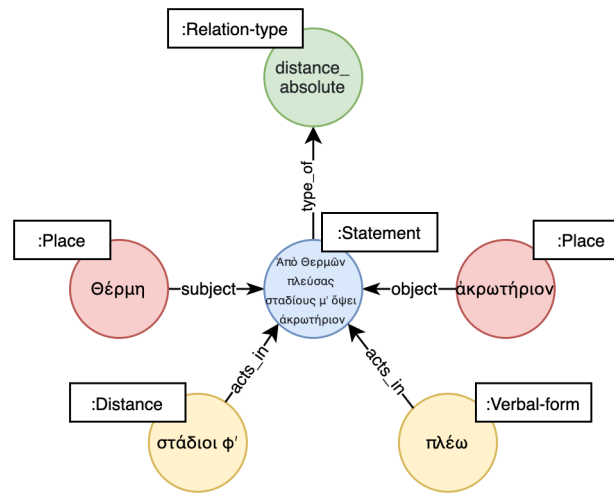
Figure 5: An example of a distance relation from the text cited on p. 4 (Ἀπὸ Θερμῶν πλεύσας σταδίους μʹ ὄψει ἀκρωτήριον), expressed in the graph meta model. The sentence translates as "From Thermae sailing 40 stades you will see the promontory."

| | Details | Properties | Standoff Properties | Entities | **Statements** | Graph | Referenced By | POS | Sections | Versions |

Name

Claim Type

Place

Section

Type

Page

Rows

Sort by

Sort direction

1

5

By Date Added

○ Ascending ● Descending

Clear **Search**

	ROLE	TYPE	NAME	ENTITIES	TIME	ADDED
View		distance_absolute	Ἀπὸ Θερμῶν πλεύσας σταδίους μ΄ ὄψει ἀκρωτήριον	(According To) Stadiasmus of the Great Sea (Of) σταδίους μ΄ (With) πλέω (Object) ἀκρωτήριον (Subject) Θέρμη		2021/1/5

Figure 6: Example as a statement in the graph database of Codex. Note that even the source itself can be modeled as an additional node.

	ID	ROLE	SUBJECT ENTITY	TAG	OTHER ENTITIES	NAME
Select	b66ee2a4-ea18-4a3d-beed-a1d515fa8d9a		(Subject) ἀκρωτήριον	direction_left	(According To) Stadiasmus of the Great Sea (With) παραφαίνω (Object) εὐώνυμος	παραφαίνειν ἐξ εὐωνύμων
Select	431ad10e-1c85-416d-ac59-eb2ece4b3d20		(Subject) Ἀσπίς	direction_north	(According To) Stadiasmus of the Great Sea (With) πλέω (Object) ἄρκτος	Ἐπ’ αὐτὴν πλέε τὴν ἄρκτον
Select	4c409fd8-32cf-4d2c-a4b8-42489f7f1629		(Subject) ἀκρωτήριον	direction_relative	(According To) Stadiasmus of the Great Sea (With) ἐπί (Object) νησία	ἐπ’ αὐτῷ ἔχον δύο νησία ἐσκολοπισμένα

Figure 7: Some examples of directions as they appear in the graph database, including sentences like: "the promontory appears on the left," "From there (Aspis) sail north," and "against which (the promontory) are islands."

	ROLE	TYPE	NAME	AGENTS
View	Trait	quality_metaphor	ἀσπίς	(According To) Stadiasmus of the Great Sea (Subject) Ἀσπίς
View	Trait	quality_visibility	περιφανές	(According To) Stadiasmus of the Great Sea (Subject) Ἀσπίς
View	Trait	quality_elevation	ὑψηλὸν	(According To) Stadiasmus of the Great Sea (Subject) Ἀσπίς
View	Trait	place-type_promontory	ἀκρωτήριόν	(According To) Stadiasmus of the Great Sea (Subject) Ἀσπίς

Figure 8: The concepts associated to Aspis in our example. Aspis is associated with the image of a shield, the qualities of high visibility and elevation, and finally with the place type of a promontory.
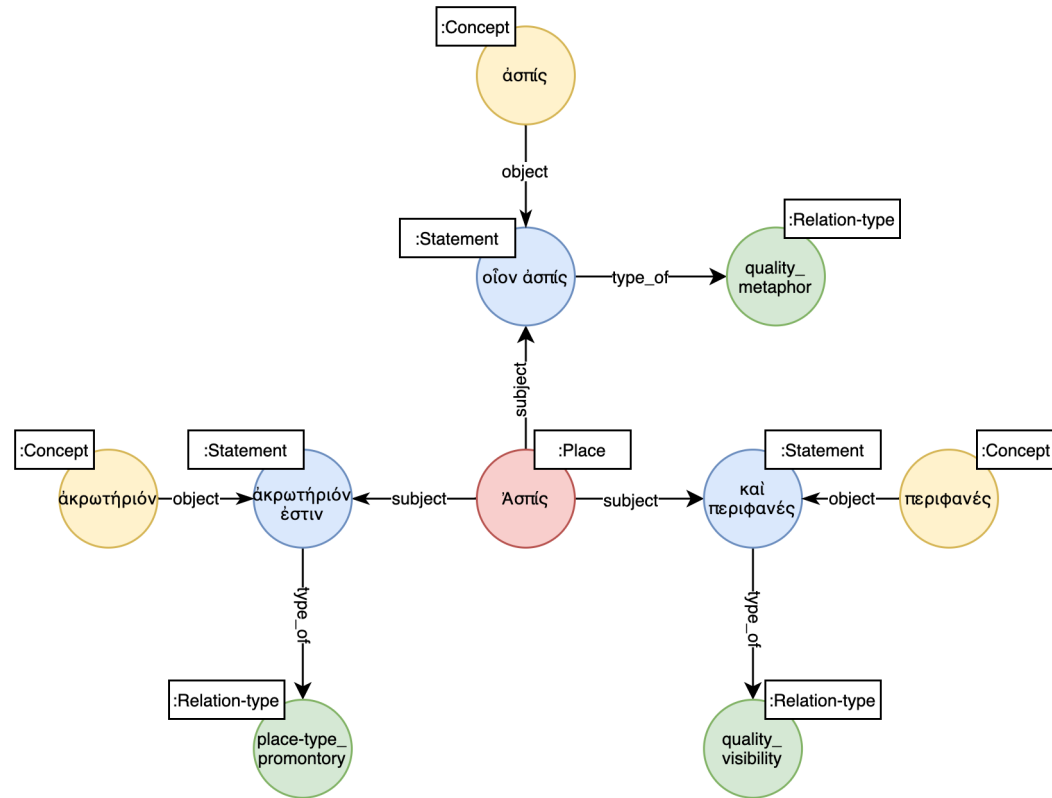
Figure 9: The same data represented as a graph meta model

Figure 10: External properties associated with Aspis in the Codex graph database

References

Arnaud, P. M. (2017). Un illustre inconnu: le Stadiasme de la Grande Mer. *Académie des Inscriptions & Belles-Lettres, Comptes Rendus*, 2:701–727.

Barker, E., Isaksen, L., and Ogden, J. (2016). Telling Stories With Maps: Digital Experiments With Herodotean Geography. In Barker, E., Bouzarovski, S., Pelling, C., and Isaksen, L., editors, *New Worlds from Old Texts: Revisiting Ancient Space and Place*, pages 181–224. Oxford University Press, DOI: `10.1093/acprof:oso/9780199664139.003.0009`.

Buzzetti, D. (2002). Digital Representation and the Text Model. *New Literary History*, 33(1):61–88, `https://www.jstor.org/stable/20057710`.

Carden, R. W. (1913). *Michelangelo Buonarroti: Michelangelo: A Record of His Life as Told in His Own Letters and Papers*. Constable & Co., London.

Eide, Ø. and Ore, C.-E. S. (2019). Ontologies and Data Modeling. In Flanders, J. and Iannidis, F., editors, *The Shape of Data in the Digital Humanities: Modeling Texts and Text-Based Resources*, pages 178–198. Routledge, Taylor & Francis Group, London/New York.

Flanders, J. and Jannidis, F., editors (2019). *The Shape of Data in the Digital Humanities: Modeling Texts and Text-Based Resources*. Digital Research in the Arts and Humanities. Routledge, Taylor & Francis Group, London/New York.

Malkin, I. (2011). *A Small Greek World: Networks in the Ancient Mediterranean*. Oxford University Press.

McCarty, W. (2005). *Humanities Computing*. Palgrave Macmillan UK, London.

Moretti, F. (2011). Network Theory, Plot Analysis. *New Left Review*, 68:80–102.

Moretti, F. and Piazza, A. (2007). *Graphs, Maps, Trees: Abstract Models for Literary History*. Verso, London/New York.

Neill, I. and Kuczera, A. (2019). Codex: An Atlas of Relations. In Kuczera, A., Wübbena, T., and Kollatz, T., editors, *Die Modellierung des Zweifels – Schlüsselideen und -konzepte zur graphbasierten Modellierung von Unsicherheiten*, volume 4 Special Issue of *Zeitschrift für digitale Geisteswissenschaften*. DOI: `10.17175/sb004_008`.

Palladino, C. (2016). New Approaches to Ancient Spatial Models: Digital Humanities and Classical Geography. *Bulletin of the Institute of Classical Studies*, 59(2):56–70.

Robinson, I., Webber, J., and Eifrem, E. (2013). *Graph Databases - New Opportunities for Connected Data*. O'Reilly, Beijing/Cambridge/Farnham/Köln/Sebastopol/Tokyo.

Schmidt, D. and Colomb, R. (2009). A Data Structure for Representing Multi-Version Texts Online. *International Journal of Human-Computer Studies*, 67(6):497–514, DOI: 10.1016/j.ijhcs.2009.02.001.

Thiering, M. and Geus, K. (2014). *Features of Common Sense Geography: Implicit Knowledge Structures in Ancient Geographical Texts*, volume 16 of *Antike Kultur und Geschichte*. LIT, Berlin/Münster/Wien/Zürich/London.