Hypertext as a Lens into Interactive Digital Narrative

David E. Millard[1][0000−0002−7512−2710]
and Charlie Hargood[2][0000−0002−5067−5941]

[1] University of Southampton, UK dem@soton.ac.uk
[2] Bournemouth University, UK chargood@bournemouth.ac.uk

Abstract. Interactive Narrative is blessed with a myriad of forms, this richness makes it hard to compare IDN systems or to develop general theories and tools as each example can seem like a special case. We take the approach of using hypertext as a method of inquiry to explore the similarities of different IDN forms. Using the Interactive Process Model to scope our analysis we systematically examine IDN from the perspective of hypertext structure. We show that hypertext can coherently explain the transition functions (the parts of the system that manages narrative state) across calligraphic, sculptural (storylets), adaptive, database driven, parser, and game narratives. In doing so we define a Hypertext Lens, made of layers of lexia state, story state, world model, and story engine. We also show how sculptural systems, parser fiction, and game narratives make use of interaction and presentation engines that complement and build upon these structures. Rather than trying to reconcile hypertext and IDN our approach instead presents hypertext as a useful thought pattern for approaching IDN that can bridge the gap between IDN forms and clarify their relationships to one another. Our analysis clearly shows a fluidity of form, encourages experimentation, and provides a mechanism through which theory can be applied widely.

Keywords: calligraphic hypertext, sculptural hypertext, adaptive hypertext, storylets, parser fiction, database narrative, narrative games

1 Introduction

Interactive Digital Narrative (IDN) has its foundations in both early text adventure games and the hypertext fiction of the 80s and 90s. Hypertext is therefore often seen as a precursor to modern IDN systems, or (as manifest in popular tools such as Twine) as a particular subset of IDN, based on textual nodes and navigational links. However, Hypertext research has continued over the last few decades, with many contributions around hypertext models, applications, usage in the wild, tools and standards. This knowledge could be a valuable resource to the IDN community if the relationship between Hypertext and modern IDN systems could be defined more clearly.

Rather than try to rationalise IDN as Hypertext, or vice-versa, in this paper we instead consider "Hypertext as Method", an approach that uses Hypertext

as a method of inquiry to understand different types of systems [5]. In this way Hypertext acts as a lens through which we can understand the common hypertextuality of different IDN systems, creating an analytical tool which can show how IDN forms are related to one another.

We also use Thue's Interactive Process Model [43] as a way to scope the lens and understand the boundaries of that hypertextuality. Rather than setting out hypertext as a restrictive view of IDN, we thus hope that this approach is a liberating one that could help us to understand how, through this common hypertextuality, theories developed for one form of IDN might translate to others.

2 Background

Hypertext has its roots in the work of its pioneers, Doug Engelbart who created NLS/Augment and considered hypertext as 'augmenting man's intellect' [14], and Ted Nelson and his conceptual system Xanadu, that imagined a global 'Permascroll' with digital technologies to seamlessly navigate and explore [34]. In both cases the central idea was to allow readers to navigate between text (or other media) by traversing navigational structures (such as links or trails). By the 1980s there were many digital implementations available such as IBIS, Intermedia, NoteCards, Hyperties, and ZOG [11], and by the end of the 1990s hypertext was an established research area, with distinct sub-communities working both on hypertext as a knowledge tool and hypertext as digital literature [45].

There were also attempts, working across both sub-communities, to formalise hypertext and create agreed models so that systems might interoperate [12, 13]. This led to an appreciation that there were different domains of hypertext, sets of models and behaviour focused on a particular task. *Navigational hypertext* based around nodes (containing media) connected via navigable links, *spatial hypertext* focused on spatial structures like lists or sets that could be dynamically arranged and identified by a spatial parser, and *taxonomic hypertexts* where conceptual hierarchies are arranged into alternative views that can be traversed [31].

In the last two decades Hypertext has gone feral, adopted and adapted by thousands of Internet communities [44], it has spawned studies in folksonomies and semantic graphs [27], been applied to the expanding web and social networks [4, 22], as well as the real world via both augmented and mixed reality [17, 42], and it continues to struggle with its own form and poetics [8, 35].

Interactive Digital Narratives share some of this early history, with key works published in early hypertext systems such as Notecards, and dedicated platforms such as Storyspace [7]. However, a mirror heritage in parser-based games, and a focus on high level narrative and content rather than low level associative structure, has given IDN research its own flavour, and ultimately a distinct community. Popular free platforms have been established such as Twine and Inform, encouraging a wider audience and experimentation that has driven craft knowledge [15]. In the last decade advances in game development frameworks has resulted in an explosion of independent narrative games, unconstrained by traditional forms, that are pushing IDN in a myriad of directions [16, 33, 38].

Not surprising then, that many in the IDN community consider hypertext to be either something from the history books or at best a tiny subset of what IDN can be. Millard attempts to reconcile the two, seeing 'literary hypertext as a subset of games, but narrative games as fundamentally hypertextual.' [32] But what does this *hypertextuality* actually mean?

Atzenbeck and Nürnberg suggest the approach that we use in this paper. They argue that Hypertext is a method of inquiry; a way of thinking about systems that can provide new insights [5]. They identify three hypertext perspectives: first class structure, context dependent structure, and open ended structure, and point out that 'other fields may adopt features of some of these perspectives, but the primacy the hypertext literature places upon them is distinguishing.' They tend towards viewing hypertext as a knowledge tool, but in this paper we explore the other hypertext tradition, hypertext as digital literature.

We also draw on Thue's model of an interactive narrative process [43], this views IDN systems as three functions that control respectively observation (what the reader sees), action (what the reader does), and transitions (how those actions change narrative state). Our approach is to use the hypertext lens to analyse the transition functions of different IDN forms, arguing that the hypertext focus on structure helps to understand and contextualise those forms. By examining the observation and action functions we can also scope the lens, and see what important elements lie beyond the hypertextual analysis.

3 Hypertext as a Lens

In this section we systematically examine a series of IDN forms, looking at how a hypertext lens might help explain how they function and how they are similar or distinct to one another. We do this by exploring the way in which IDNs are *structured*, using hypertext terms and structural models.

3.1 Linking, Adaptation, and Guard Fields

When we use the term hypertext, most people will think of the navigational domain of hypertext, defined with nodes and links. Where these are explicitly defined this has also been called *Calligraphic Hypertext*, as the links are drawn deliberately between nodes to create navigational paths between them [6]. Links are emblematic of hypertext systems, so much so that Halasz referred to the 'Tyranny of the Link' to complain about the extent to which they exclude other types of structures (such as trails, or virtual documents) [19].

This *hyperstructure* is concerned with defining the possible states of the hypertext, and the ways in which a reader might move between them. It is therefore a way of describing the *transition function* of a system. In calligraphic hypertext it is tempting to think of the nodes of the hypertext as states, and the links as the transitions. But this confuses what is displayed on screen with the state of the story (the salient parts of the reader's understanding of the hypertext as modelled by the system). The story state only changes *after* the reader has

read the lexia[3] within a node. It is therefore the lexia that represent transitions between states of the story.

In calligraphic hypertext this distinction may seem moot, as there is a one-to-one mapping between the lexia and the possible states of the story (the lexia the reader has just read is sufficient to model their position in the story). However, *Adaptive Hypertext* breaks this strict connection. In adaptive hypertext additional rules dynamically adjust the lexia or the links or both [10]. The changes were classified by Brusilovsky into multiple methods (for example, removing content or links vs. changing presentation) [9]. These rules mean that in adaptive hypertext the current lexia is no longer sufficient to model the state of the story, and instead adaptive hypertext systems include some form of user model that is combined at runtime with the calligraphic hypertext model. The user model might be set up in advance (for example, selecting a learning preference in an educational hypertext system) and/or can be dynamically adjusted at runtime (for example, to track what the reader has already seen). The later means for example, that while calligraphic hypertexts can have circular structure, they cannot distinguish between repeat visits to a lexia (each visit returns to the same state), while an adaptive hypertext can do this (return visits can be different states).

StorySpace is an adaptive hypertext system, as it uses *guard fields* to control whether links are available or not [7]. Twine is also an example of an adaptive hypertext system where the author can define variables that are modified during reading [15]; both content and links within Twine can then be made dependent on these variables to achieve many of the effects set out by Brusilovsky.

Strict calligraphic hypertexts could therefore be defined as those where *lexia state* is sufficient to model the reader's position in a story, whereas adaptive hypertext systems - whilst still mostly calligraphic - require both lexia state and additional *story state* that can modify the presence and appearance of both content and structure[4].

3.2 Sculptural Hypertext and Storylets

Sculptural Hypertexts are those that use purely story state to manage the progression of the narrative. The lexia state is ignored. You can imagine that in a sculptural hypertext all nodes are potentially connected (which is why it is irrelevant which lexia you have just read), but at runtime the story state is used to sculpt away most of the connections [6]. For example, one node might assert that the reader has now 'met Alice' while another states that it is only available if 'met Alice' has been asserted.

[3] Although we use the term *lexia* this need not be textual content, and in fact our emphasis on lexia representing state changes means that it is their purpose within the overall narrative that is important rather than their form, making our use of the term very similar to Mateas' notion of 'dramatic beats' [30].

[4] The adaptive hypertext literature refers to a user model, as an amalgam of user preferences and current reading state, but in IDN non-diegetic user preferences (from outside of the user's interaction with the hypertext, e.g. age or expertise level) are less important, so we adopt the term story state to describe the same concept.

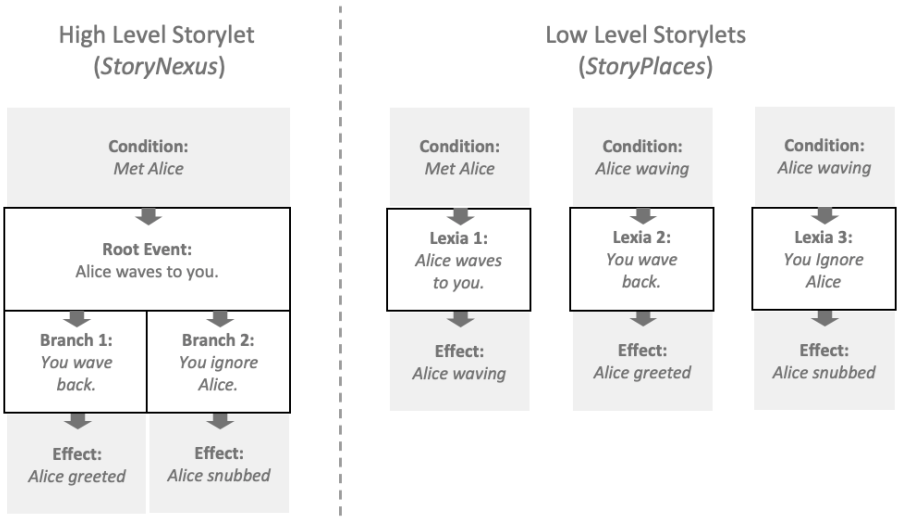

Fig. 1. One high level Storylet captures the behaviour of three low level Storylets.

The notion of Storylets emerged separately in the interactive fiction community [26]. Storylets are effectively the building blocks of a sculptural hypertext, they contain the lexia itself, as well as the conditions that govern whether it is available, and the behaviour that modifies the story state when the lexia is read. A collection of storylets *is* a sculptural hypertext (in the interactive fiction community these have also been referred to as quality-based stories [41]).

In the hypertext literature these building blocks have been conceptualised as a 'context-sourced link' (i.e a link where the source is a story state rather than a node) [46], meaning that storylets can be modelled consistently alongside calligraphic hypertexts. Storylets also exist at different levels of complexity. For example, in StoryPlaces [21] each storylet is a single lexia with conditions and effects, this low level approach is powerful but makes authoring branching structures difficult[5]; whereas in StoryNexus [1] the storylet is made up of a root event (a starting lexia) followed by branches with different effects. This difference is shown in Figure 1. StoryPlaces then is a strict sculptural hypertext - using only story state, whereas StoryNexus, although primarily sculptural, adds additional lexia state to make authoring branching choices and consequences simpler.

3.3 Story Engines and Database Narratives

Hypertext systems are more than a hypertext model, they also require a runtime system to apply those models and manage state. As described previously Thue's model of an interactive narrative process deconstructs this system into a

[5] Sculptural Hypertext is a good fit for locative literature systems like StoryPlaces, because location can be modelled as just another condition.

transition, observation, and action function [43]. We have already described the transition function in IDN as typically being fulfilled by a *story engine*, and we can say that in most IDN systems the observation function is fulfilled by some sort of *presentation engine*, and the action function by some sort of *interaction engine* (although in practice they may be implemented together).

In calligraphic hypertext the presentation engine displays the lexia and link anchors appropriate for the current state, the interaction engine registers the reader's clicks on link anchors, the story engine then uses the calligraphic model to select the next lexia, thus moving the story forward. Whereas in sculptural hypertext the presentation engine displays the lexia and the available storylets for the current state, the interaction engine registers the reader's selection of the next storylet, and the story engine uses the sculptural model to select the next lexia and alter the story state.

The hypertext lens is primarily concerned with the transition function, the story engine, and the models that drive it, but that does not mean that the other functions are not significant. For example, in sculptural systems we also see a lot of variety in the observation function, part of what Mateas calls the *content selection architecture* (Kreminski and Wardrip-Fruin provide a good overview of selection strategies [26]). We can see evidence of this in StoryPlaces, which exhaustively lists all available storylets given the current story state [21], whereas Card Shark shows a random selection of storylets, disabling (but still showing the details of) storylets whose conditions are not currently met [6]. These differences in the observation function have profound implications for how a narrative is experienced and designed (for example, this allows storylets to be added to Card Shark that can never be reached, but which influence the player when they appear in the selection as players do not know that they are unreachable).

In Sam Barlow's *Her Story* and *Telling Lies* the action function is not to click on links or select storylets, but to type search terms into a database of video. Gasque et al. call this approach 'database narratives' [16]. From the hypertext perspective database narratives are effectively a form of sculptural hypertext, but where the player rather than the system manages the transition function, in other words, database narratives require no story engine. Despite this, the underlying story model, used to design and plan the narrative, is still sculptural. But when lexia (or in this case video) is seen, it only changes the story state if the player themselves makes note of what new elements have been revealed. In the absence of a story engine Barlow uses the limitation of only returning the first 5 items (a modification of the observation function) in order to manage progression through the story [16].

3.4 Parsers, Narrative Games and World Models

Like sculptural hypertexts, parser fiction uses story state to manage its narrative, only here the interaction engine relies on interpreted typed commands rather than clickable links. These parser fictions often take the form of a traditional text adventure and express the story state in terms of a *world model*, a schema or super structure that provides a framework for a coherent (and

more complete) description of the story world. This allows authors to express themselves at a higher level than a set of variables - for example, in Inform we declare a set of rooms and their relationships, objects and their locations within rooms, and scenes in which story unfolds. There is then a grammar for interacting with this world model (moving between rooms, picking up and inspecting objects, etc.) The query interface of database narratives can also be considered as a simple grammar, but the superstructure of the world model makes a more complex grammar possible, as valid actions can be defined against whole classes of objects. Parser fiction also often uses a different presentation engine to traditional hypertext, as it doesn't show you explicit choices and instead expects your choices to be made diegetically by interacting with the world model [28].

While parser based adventure games adopt new interaction and presentation engines, narrative games can take this even further. Whilst still built on a story engine (and often a world model) their interactions can include rich ludonarrative mechanics [2], with elements of the story being delivered through a variety of channels, such as cutscenes, scripted events/character barks, log entries, interactive dialog, and environmental storytelling (reminiscent of transmedia [23]).

The element of environmental storytelling means that, like database narratives, narrative games and parsers do not necessarily track and manage all of the state through their story engines, and that there is the possibility that the transition function is shared between player and machine. To some extent then, all three forms are epistemic, not in Ryan's motivational sense (in that players are "driven by the need to know") [40] but rather they are at least partially driven by *what* the player knows.

As we have considered these different forms of IDN - Calligraphic, Adaptive, Sculptural, Database, Parsers, and Games - we have moved through and then beyond the hypertext lens. Particularly in these last two forms, where there is increased importance of the observation and action functions as ways of delivering story and managing the experience. Thus a design theory of narrative games cannot be solely based on hypertext, but these hypertextual models can be used to understand the transition function - the story engine and underlying structural models for such experiences. Hypertext models can therefore be used as analytical tools to understand how narrative games work, and hypertext systems can be used as authoring tools for much of their narrative structure.

4 Discussion

Through our hypertextual description of IDN forms we have developed a simple layer model of the Hypertext Lens (shown in Figure 2).

Central to this is a finite state model where each lexia that the reader encounters represents a transition between states of the narrative. Note that lexia represent transitions, not states (see Section 3.1). However, it is possible that narratives are managed solely via *lexia state*, in which case there is a one-to-one correlation between the lexia state and the state of the narrative (e.g. the state might be: 'You have just read Page 1'). Calligraphic Hypertext exists at this

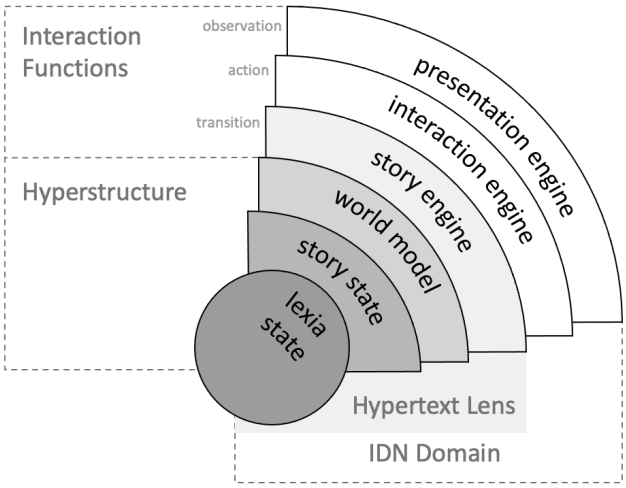

Fig. 2. The Hypertext Lens as Applied to the Layers of IDN

level, reading a lexia in the hypertext transitions you to a matching state, and links in that lexia represent the next possible lexia and state transitions. Authors are effectively directly writing the narrative state graph. Because of this one-to-one match Calligraphic Hypertext is easy to understand and is therefore one of the most accessible forms of IDN, although if the narrative is very open (meaning that the reader has many choices) then the density of the state graph can become high, and the hypertext becomes difficult to author and visualise.

A way to solve this problem is to model *story state* explicitly, and to use this and not lexia state to model the narrative. Figure 3 shows a simple story sequence modelled as a calligraphic hypertext (using only lexia state) and as a sculptural hypertext (using only story state). Lexia still act as transitions between story states (e.g. reading 'You take the key.' leads to the state: 'has key') but now, assuming that not all of the information in the lexia will affect the onward story[6], different lexia can transition you to the *same* story states, which simplifies the state graph (in Figure 3, Original Story, states 2 and 3 in the calligraphic model have been consolidated into one state in the sculptural model). This is the basis of Sculptural Hypertext, which exclusively uses story state to manage the reader's progress.

Sculptural Hypertext has a further advantage. Authors writing calligraphic hypertext have no choice but to use the whole of the state information when creating transitions, because the states of the story are collapsed down into a single piece of information (the current lexia). For example, if I want to add a lexia where the reader can 'examine the key' to the calligraphic system (Figure 3, Extension) then I have to create transitions from all states where the reader has

[6] In our example we are not interested in modelling the state of the guard

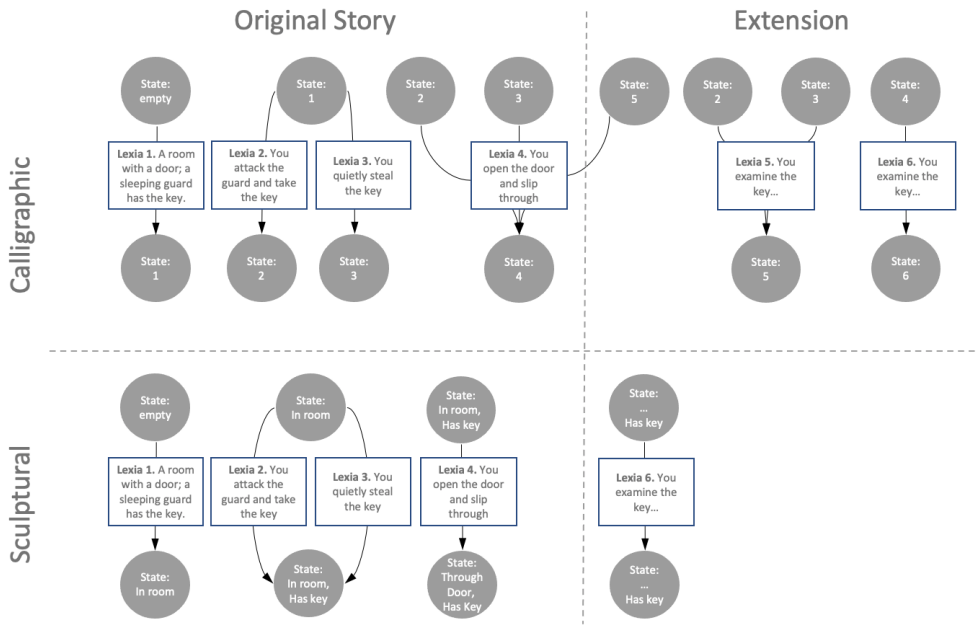

Fig. 3. Calligraphic Hypertext relies on Lexia State and authors directly write the state model, whereas Sculptural Hypertext uses Story State and users write at a higher level

that key (states 2, 3, and 4), in addition I will need two versions of the lexia, one if they examine it in the room (state 5), and another for after the door (state 6) in order to present the correct set of onward transitions. However, in a sculptural system I can define the transition based on partial story state. I can do this by creating a single storylet that transitions from any state where 'has key' is set, this will then be available from all states in the underlying state model where this is true. Sculptural hypertext authors are not therefore directly authoring the state graph, but working at a higher level of abstraction (effectively they are creating parallel story state models that each track a different aspect of the story and can be checked and altered independently, e.g. a state model for the key, another for the room). This is what makes sculptural hypertext so powerful, but also more complex to conceptualise and design.

Adaptive Hypertext (which is what most node/link systems actually represent) is a way to balance this power and complexity, here we use lexia state to model the bulk of the hypertext, but story state to simplify that graph when needed (e.g. using guard fields [7]). Parser-based fiction also uses story state abstractions. Often these are text adventures that also provide a comprehensive schema to help manage story state in terms of the story world, we call this a *world model*. For example, the world model provided by Inform includes rooms and their connections, objects and inventory, and scenes and their progression. All of this could be replicated in a sculptural hypertext with many individual

state variables, but authoring using the world model is an easy way to define complex state quickly and in a manner that is familiar to writers [36].

All of these systems use a *story engine* [3] to keep track of state and manage the available transitions (fulfilling the transition function). An exception to this are database-driven stories, these operate using story state but do not provide the machinery to track state. Instead the story state is left in the head of the reader, typically in the form of keywords that they might wish to pursue further (the transition function is undertaken by the reader, not the machine). This weakens the connection between lexia and state changes, as the reader must notice and remember that state for the transition to occur. Lexia in database stories thus represent the *possibility* of state transitions. In addition, there are state changes that can occur entirely outside of the system (for example, if the reader forgets a crucial piece of information, or learns something from an external source). Therefore in database driven stories, although the author might design an ideal story state model (that becomes embedded in the lexia), the actual states and transitions experienced by the reader could be different.

Hypertext works well to examine these inner layers of the model. However, there are clearly limits to its ability to describe IDNs. For a start (and following Adam's terminology [3]) there are a wide variety of interaction engines (fulfilling the action function) and presentation engines (fulfilling the observation function), and it is possible to play with these in order to get specific effects. Calligraphic hypertexts convey subtle information about choices via the placement of link anchors. Sculptural hypertexts typically list available lexia and use explicit names or prompts to provide clues to the reader. Parsers often don't present story options at all, but provide a set of actions and a grammar to interact with the world model by typing commands. Database narratives provide a search interface where the users type keywords directly.

Variations in interaction and presentation engines are at the very edges of the Hypertext space, what Bernstein calls 'Strange Hypertexts' [6], and at their extremes result in the sophisticated action functions we see in games - where the interaction engine constitutes a complex set of mechanics [7]. Similarly hypertext says very little about the presentation engine that is used to translate lexia and player options to the screen. In this the Hypertext Lens outlined above matches the classic Dexter Model of hypertext [18], which separates systems into different layers, and places hypertext in a central 'Storage Layer'. The Lens we have set out (lexia state, story state, world model, and story engine) correspond to this 'Storage Layer'. Whereas interaction engines are part of the 'Runtime Layer' and presentation engines part of the 'Within Component Layer', both of which Dexter considers to be outside of the core hypertext design.

The Hypertext Lens does not therefore cover the entirety of IDN systems. However, it helps to explain what Millard called a 'core hypertextuality' in narrative games [32], evidenced in attempts to characterise Bandersnatch as a Gauntlet [37] or to map the structure of The Walking Dead [24]. We go beyond this

[7] Even in calligraphic hypertext, the observation and interaction engines convey meaning, e.g. Mason and Bernstein's work on the poetics of contemporary link usage [29]

claim and have begun to unpack this hypertextuality and to relate it to different IDN forms. Thue's model of an interactive narrative process helps to explain that the Hypertext Lens describes the transition function of IDN, whereas games design better covers the observation and action functions. This means that when those functions are relatively simple the hypertext lens is a more complete description than when they are complex. While the explanatory power of the lens is focused on the interactive narrative elements of games, it does not follow that it is limited to games where narrative is seen as a minor element, just that its scope is limited within those games. In these cases the lens provides clarity precisely because it allows you to focus on one part of a more complex medium.

The Hypertext Lens also reveals the fluidity of IDN forms. This might inspire us to consider IDNs that mix forms together. For example, we might imagine database stories with some sculptural elements, sculptural stories with a more robust world model, or calligraphic stories that include some free text query and therefore push some story state outside of the story engine. We have also seen how hypertext can help explain complex IDNs that already mix up these forms. This is true even in the case of commercial narrative games, for example Supergiant's narrative roguelike *Hades* could be seen as a sculptural hypertext using large grain storylets with calligraphic internal structure, coupled with a presentation engine that selects storylets based on a priority selection strategy.

Understanding the hypertextual commonalities of these forms also enables them to benefit from theories and craft knowledge developed elsewhere. Database narratives may seem like unknown design spaces, but much of the craft knowledge from storylets could be applied to them (e.g. by utilising sculptural patterns in their design [20]). Understanding whether the narrative model of a game is more calligraphic or sculptural enables us to apply the right narrative design principles, and to pick an appropriate IDN tool in which to write the script.

Hypertext is a structure-centric view of IDN (after all, Atzenbeck argued that a focus on structure was the defining perspective of hypertext [5]), and while it does not fully capture all of the elements it does allow us to think about the structures separately from the other parts of the experience. Thus it has real value as a *thought pattern*, a way of approaching IDN that sits alongside similar analytical models such as SPP [25] or the Double-Hermeneutic Circle [39] and can be used both descriptively for analysis and prescriptively for design.

5 Conclusion

In this paper we have adopted Atzenbeck's approach of using Hypertext as a Method of Inquiry in order to analyse IDN forms through a hypertext lens. We have also used Thue's Interactive Process Model as a means to scope that lens, and have shown that it helps to deconstruct the transition function of IDN systems (how the system manages narrative state). In doing so we have revealed how different IDN forms relate to one another. Calligraphic hypertext purely uses lexia states and transitions (resulting in a one to one mapping between lexia and narrative state), sculptural hypertext (a set of storylets) purely uses

story states and transitions, and adaptive hypertexts mix the two, using lexia state to manage most of the narrative, with story state employed to simplify overly complex sections. In all three cases a story engine manages the transition function. The lens also shows that database IDNs are sculptural in nature, but instead of a story engine the transition function occurs within the mind of the reader; and it reveals parser fiction as structurally sculptural but defined against a world model (which makes authoring complex state easier). Throughout, we have given small or abstract examples of the lens in action, and in future work we hope to use the lens to explore specific IDN works in more detail.

The hypertext lens is a thought pattern that puts structure first, it thus moves the focus away from the presentation and interaction engines that fulfil the observation and action functions (how narrative and choices are presented, and the ways in which the reader/player interacts with them). In calligraphic and adaptive hypertext these are relatively simple, and the lens provides an almost complete picture, but in sculptural systems a more complex presentation engine makes decisions about which storylets to reveal next, and in parser fiction a more sophisticated interaction engine manages a complex grammar made possible by the world model. In narrative games there is a high level of experimentation with presentation and interaction engines - with multiple channels reminiscent of transmedia, and complex mechanics that yield only to game design theory. Nevertheless, the hypertext lens continues to usefully highlight the structures within the transition function and the activity of the game's story engine.

This coherent hypertextual view of IDN suggests a fluidity of forms, with the potential for theory to be applied consistently across them. When new ideas emerge (such as database narratives) if they can be explained in terms of existing hypertext models they immediately benefit from design theory that already exists. Hybrid approaches that seem at first esoteric and strange (for example, the located and contextual nodes of StoryPlaces) can be explained in terms of what has gone before. It also implies a core set of narrative design skills that could be taught, and would be relevant across many different forms.

Finally it is worth acknowledging that the hypertext lens is a technical deconstruction of IDN, it does not tell us how to write dramatic plots, how to manage agency effectively, or how to build believable characters, but it does explain the structural narrative architecture of IDN forms and reveals the fluidity between them. This comprehensive explanation of a part of IDN is valuable. If we understand the hypertextual similarities between IDNs then when we see best practice in one form we can translate it to another, even though the lens doesn't tell you what that best practice itself might be. We hope that by using hypertext to set out a coherent view of IDN forms, and by describing a lens that can be used in future to analyse new works and approaches, that we will empower the next generation of writers, designers, and tool developers to explore these other questions, and in doing so, develop answers that can be applied widely.

Acknowledgements The authors would like to thank David Thue and Joey Jones for providing feedback on an earlier draft of this paper.

References

1. Storynexus: Reference guide, version 1.04. Tech. rep., Failbetter Games (2012), http://wiki.failbettergames.com/start
2. Aarseth, E.: A narrative theory of games. In: Proceedings of the International Conference on the Foundations of Digital Games. p. 129–133. FDG '12, Association for Computing Machinery, New York, NY, USA (2012). https://doi.org/10.1145/2282338.2282365
3. Adams, E.: Fundamentals of Game Design. New Riders Publishing, USA, 3rd edn. (2014)
4. Anderson, M., Carr, L., Millard, D.E.: There and here: patterns of content transclusion in wikipedia. In: Proceedings of the 28th ACM Conference on Hypertext and Social Media. pp. 115–124 (2017)
5. Atzenbeck, C., Nürnberg, P.J.: Hypertext as method. In: Proceedings of the 30th ACM Conference on Hypertext and Social Media. p. 29–38. HT '19, Association for Computing Machinery, New York, NY, USA (2019). https://doi.org/10.1145/3342220.3343669
6. Bernstein, M.: Card shark and thespis: Exotic tools for hypertext narrative. In: Proceedings of the 12th ACM Conference on Hypertext and Hypermedia. p. 41–50. HYPERTEXT '01, Association for Computing Machinery, New York, NY, USA (2001). https://doi.org/10.1145/504216.504233
7. Bernstein, M.: Storyspace 1. In: Proceedings of the Thirteenth ACM Conference on Hypertext and Hypermedia. p. 172–181. HYPERTEXT '02, Association for Computing Machinery, New York, NY, USA (2002). https://doi.org/10.1145/513338.513383
8. Bernstein, M.: On hypertext narrative. In: Proceedings of the 20th ACM Conference on Hypertext and Hypermedia. pp. 5–14 (2009)
9. Brusilovsky, P.: Adaptive hypermedia. User modeling and user-adapted interaction **11**(1), 87–110 (2001)
10. Brusilovsky, P., Kobsa, A., Vassileva, J.: Adaptive hypertext and hypermedia. Springer (1998)
11. Conklin, J.: Hypertext: An introduction and survey. Computer **20**, 17–41 (10 1987). https://doi.org/10.1109/MC.1987.1663693
12. Davis, H.C., Millard, D.E., Reich, S., Bouvin, N., Grønbæk, K., Nürnberg, P.J., Sloth, L., Wiil, U.K., Anderson, K.: Interoperability between hypermedia systems: The standardisation work of the ohswg. In: Proceedings of the Tenth ACM Conference on Hypertext and Hypermedia: Returning to Our Diverse Roots: Returning to Our Diverse Roots. p. 201–202. HYPERTEXT '99, Association for Computing Machinery, New York, NY, USA (1999). https://doi.org/10.1145/294469.294904
13. Davis, H., Reich, S., Millard, D.: A proposal for a common navigational hypertext protocol. Technical report, Dept. of Electronics and Computer Science, University of Southampton, 1997 (1997)
14. Engelbart, D.C.: Augmenting human intellect: A conceptual framework. Menlo Park, CA (1962)
15. Ford, M.: Writing interactive fiction with Twine. Que Publishing (2016)
16. Gasque, T., Tang, K., Rittenhouse, B., Murray, J.: Gated Story Structure and Dramatic Agency in Sam Barlow's Telling Lies, pp. 314–326 (10 2020). https://doi.org/10.1007/978-3-030-62516-0_28
17. Grønbæk, K., Kristensen, J.F., Ørbæk, P., Eriksen, M.A.: "physical hypermedia" organising collections of mixed physical and digital material. In: Proceedings of the fourteenth ACM conference on Hypertext and hypermedia. pp. 10–19 (2003)

18. Halasz, F., Schwartz, M., Grønbæk, K., Trigg, R.H.: The dexter hypertext reference model. Commun. ACM **37**(2), 30–39 (Feb 1994). https://doi.org/10.1145/175235.175237
19. Halasz, F.G.: "seven issues": Revisited, hypertext'91 closing plenary. In: Proc. Third ACM Conference on Hypertext (Hypertext'91) (1991)
20. Hargood, C., Hunt, V., Weal, M., Millard, D.: Patterns of sculptural hypertext in location based narratives. In: HT'16: Proceedings of the 27th ACM Conference on Hypertext and Social Media. pp. 61–70. ACM (July 2016), https://eprints.soton.ac.uk/390748/
21. Hargood, C., Weal, M., Millard, D.: The storyplaces platform: Building a web-based locative hypertext system. In: HT '18 Proceedings of the 29th ACM Conference on Hypertext and Social Media. pp. 128–135. ACM (2018), https://eprints.soton.ac.uk/421122/
22. Helic, D., Strohmaier, M., Granitzer, M., Scherer, R.: Models of human navigation in information networks based on decentralized search. In: Proceedings of the 24th ACM conference on hypertext and social media. pp. 89–98 (2013)
23. Javanshir, R., Millard, D., Carroll, E.: Structural patterns for transmedia storytelling. PLoS ONE **15**(1) (January 2020), https://eprints.soton.ac.uk/438130/
24. Killham, E.: Here's a chart of every choice in the walking dead: Season 1 (image) (Mar 2013), https://venturebeat.com/2013/03/31/the-walking-dead-season-one-plot-graph/
25. Koenitz, H.: Towards a theoretical framework for interactive digital narrative. vol. 6432, pp. 176–185 (11 2010). https://doi.org/10.1007/978-3-642-16638-9_22
26. Kreminski, M., Wardrip-Fruin, N.: Sketching a map of the storylets design space. In: Rouse, R., Koenitz, H., Haahr, M. (eds.) Interactive Storytelling. pp. 160–164. Springer International Publishing, Cham (2018)
27. Marlow, C., Naaman, M., Boyd, D., Davis, M.: Ht06, tagging paper, taxonomy, flickr, academic article, to read. In: Proceedings of the seventeenth conference on Hypertext and hypermedia. pp. 31–40 (2006)
28. Mason, S.: On games and links: Extending the vocabulary of agency and immersion in interactive narratives. In: Koenitz, H., Sezen, T.I., Ferri, G., Haahr, M., Sezen, D., Catak, G. (eds.) Interactive Storytelling. pp. 25–34. Springer International Publishing, Cham (2013)
29. Mason, S., Bernstein, M.: On links: Exercises in style. In: Proceedings of the 30th ACM Conference on Hypertext and Social Media. p. 103–110. HT '19, Association for Computing Machinery, New York, NY, USA (2019). https://doi.org/10.1145/3342220.3343665, https://doi.org/10.1145/3342220.3343665
30. Mateas, M., Stern, A.: Structuring content in the façade interactive drama architecture. In: Proceedings of the First AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment. p. 93–98. AIIDE'05, AAAI Press (2005)
31. Millard, D.E., Moreau, L., Davis, H.C., Reich, S.: Fohm: A fundamental open hypertext model for investigating interoperability between hypertext domains. In: Proceedings of the Eleventh ACM on Hypertext and Hypermedia. p. 93–102. HYPERTEXT '00, Association for Computing Machinery, New York, NY, USA (2000). https://doi.org/10.1145/336296.336334
32. Millard, D.: Games/hypertext. In: ACM Conference on Hypertext and Social Media (13/07/20 - 15/07/20). pp. 123–126 (July 2020), https://eprints.soton.ac.uk/442652/

33. Moss, A.: The link out. In: Interactive Storytelling: 11th International Conference on Interactive Digital Storytelling, ICIDS 2018, Dublin, Ireland, December 5–8, 2018, Proceedings. vol. 11318, p. 206. Springer (2018)
34. Nelson, T.H., et al.: Literary machines: The report on, and of, project xanadu, concerning word processing, electronic publishing, hypertext, thinkertoys, tomorrow's intellectual revolution, and certain other topics including knowledge, education and freedom (1981)
35. Pisarski, M.: New plots for hypertext? towards poetics of a hypertext node. In: Proceedings of the 22nd ACM Conference on Hypertext and Hypermedia. pp. 313–318 (2011)
36. Reed, A.A.: Telling stories with maps and rules: using the interactive fiction language "inform 7" in a creative writing workshop. Creative writing in the digital age: theory, practice, and pedagogy pp. 141–152 (2015)
37. Rezk, A., Haahr, M.: The case for invisibility: Understanding and improving agency in black mirror's bandersnatch and other interactive digital narrative works (11 2020). https://doi.org/10.1007/978-3-030-62516-0_16
38. Roth, C., Koenitz, H.: Bandersnatch, yea or nay? reception and user experience of an interactive digital narrative video. In: Proceedings of the 2019 ACM International Conference on Interactive Experiences for TV and Online Video. pp. 247–254 (2019)
39. Roth, C., van Nuenen, T., Koenitz, H.: Ludonarrative Hermeneutics: A Way Out and the Narrative Paradox, pp. 93–106 (11 2018). https://doi.org/10.1007/978-3-030-04028-4_7
40. Ryan, M.L.: Interactive narrative, plot types, and interpersonal relations. In: Spierling, U., Szilas, N. (eds.) Interactive Storytelling. pp. 6–13. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
41. Short, E.: Beyond branching: Quality-based, salience-based, and waypoint narrative structure (Apr 2016), https://emshort.blog/2016/04/12/beyond-branching-quality-based-and-salience-based-narrative-structures/
42. Sinclair, P., Martinez, K., Millard, D.E., Weal, M.J.: Links in the palm of your hand: tangible hypermedia using augmented reality. In: Proceedings of the thirteenth ACM conference on Hypertext and hypermedia. pp. 127–136 (2002)
43. Thue, D.: What might an action do? toward a grounded view of actions in interactive storytelling. In: Bosser, A.G., Millard, D.E., Hargood, C. (eds.) Interactive Storytelling. pp. 212–220. Springer International Publishing, Cham (2020)
44. Walker, J.: Feral hypertext: when hypertext literature escapes control. In: Proceedings of the sixteenth ACM conference on Hypertext and hypermedia. pp. 46–53 (2005)
45. Wardrip-Fruin, N.: What hypertext is. In: Proceedings of the Fifteenth ACM Conference on Hypertext and Hypermedia. p. 126–127. HYPERTEXT '04, Association for Computing Machinery, New York, NY, USA (2004). https://doi.org/10.1145/1012807.1012844
46. Weal, M.J., Millard, D.E., Michaelides, D.T., Roure, D.C.D.: Building narrative structures using context based linking. In: In Hypertext '01. Proceddings of the Twelfth ACM conference on Hypertext, Aarhus, Denmark. (01/08/01). pp. 37–38 (August 2001), https://eprints.soton.ac.uk/256136/