Journal of Universal Computer Science, vol. 25, no. 12 (2019), 1608-1624
submitted: 14/5/19, accepted: 31/10/19, appeared: 1/12/19 © *J.UCS*

Integrating 3D Objects in Collaborative Non-Linear Storytelling on the Web

Peter de Lange, Petru Nicolaescu, Jan Benscheid and Ralf Klamma
(Informatik 5, RWTH Aachen University
Ahornstr. 55, 52074 Aachen, Germany
{lange/nicolaescu/benscheid/klamma}@dbis.rwth-aachen.de)

Abstract: In both formal and informal learning scenarios, storytelling is a mean towards acquiring and sharing knowledge. On the Web, user generated content and digital representations of real-world artifacts contribute to the story's expressiveness. 3D objects are a good example of such representations, as they capture and distribute rich and complex information. Currently, a bridge between digital storytelling, social Web 2.0 features – such as tagging – and 3D objects is missing. We present an approach for the collaborative creation of non-linear stories in near real-time, centered around 3D objects. We use a metamodel, which acts as a basis for collaborative story creation. Stories are directly linked to 3D objects being attached to camera views and perspectives on a certain object. We instantiate story viewers from the collaboratively authored stories and use text and multimedia to annotate and browse the 3D objects. Our conducted evaluation proves the feasibility of the approach and promises good results in applying collaborative storytelling for 3D object browsing in order to scaffold learning.
Key Words: non-linear storytelling, 3D object annotation, collaboration
Category: H.3.5, L.1.0, L.1.1

1 Introduction

Storytelling is a well-approved method for conveying knowledge, a way to make diverse information cohere [Brown and Duguid, 2000]. A plot helps maintaining the listener's attention and facilitate learning by creating interconnections between the content [Bransford et al., 1999]. Furthermore, storytelling is an interactive process, during which the participants as part of a *Community of Practice* (CoP) [Wenger, 1998] collaboratively shape the story. Especially in many informal learning scenarios, enabling the collaborative creation of stories facilitates participation and knowledge negotiation [Lukosch et al., 2011]. Finally, through these collaborations, it is possible to identify the community's ignorance, which leads to a culture of evolution through collaboration, one of the fundamental principles of CoPs.

Many good examples for educational storytelling include the listener expecting a story and being told lessons along the way. However, in scenarios dedicated to knowledge transfer, there might be people wanting information, but being told a story. In this case, especially if the plot is dominant, but poorly authored, storytelling might be perceived as an unnecessary nuisance. A solution for this is delivering *non-linear stories*, as story consumers influence the direction of the narrative with their actions and creators have access to a richer mean of expressing their narration.

To further extent the means of expressing narrations, multimedia artifacts in digital non-linear storytelling can be used to improve conveying story content and offer a richer experience, beyond simple text [Schumann et al., 2013]. Thereby a lot of hidden, non-obvious information contained within multimedia artifacts can be enriched using annotations that highlight the most important, interesting facts. Recently, transforming real-world artifacts into a digital format has become easier and more cost effective, enabling access to a huge pool of potential users. Moreover, using modern Web technologies, it is now possible to display the resulting 3D objects directly in Web browsers. An example is given by 3D digital representations of real-world objects that are otherwise not reachable for the majority of learners, due to their perishable or expensive nature.

In this paper, we present a Web-based digital storytelling approach, which is used to create and share multimedia-based, non-linear stories around 3D objects in a collaborative manner. Therefore, we developed a well-defined metamodel according to which communities can collaboratively instantiate model-based stories in near real-time, directly in their Web browser. The focus of these stories is placed on learning, which is meant to be supported through the mixture between collaboration in the creation process, annotation of 3D object views and camera perspectives with various multimedia such as text, videos and images and novel, intuitive consumption of the resulting non-linear stories. We have evaluated and implemented our concept using a near real-time collaborative prototype, and our findings demonstrate that this blending of concepts is well perceived. With this, we want to answer the following two research questions:

R1: How can we support collaborative model-based storytelling processes around 3D objects on the Web?

R2: Does non-linear storytelling around 3D objects engage learners in knowledge creation and sharing?

2 Background and Related Work

Storytelling has a broad range of applications. In this contribution, we focus on the educational aspects, and how it can facilitate concrete knowledge transfer. We state three requirements a good educational story has to fulfill to be successful:

Attention Though entertainment is not the main purpose of educational stories, it is still a very important factor for the quality of a story. Good stories will create tension, which helps maintaining the consumer's attention. It's important to notice, that a bad story may lead to the exact opposite. In this case, the plot can be perceived as a distraction from the information that is to be conveyed.

Associations Drawing connections is very important for effective learning [Bransford et al., 1999]. Hearing a story will help interconnecting all the contained information with the plot and between themselves. There is a mnemonic (i.e. a technique for effective memorization) called the "Mnemonic link system "[1], in which one has to make up a story out of a list of arbitrary information. In a study by Bower and Duguid [Bower and Clark, 2013], this technique increased the memorization of random nouns by a factor of seven. Notably, those stories do not demand for any "narrative quality". It's even assumed (however disputed), that bizarreness may increase memorizing effects [McDaniel and Pressley, 2012].

Participation A plot does not only help by keeping the listener interested. A study by Paul [Annie Murphy Paul, 2012] showed, that while listening to a narrative, the brain does not only try to picture what is told, but actually reenacts the story. This may encourage thinking about own solutions for the problems depicted in the story, and thereby assist actual understanding, instead of pure memorizing. Non-linear and therefore interactive narration may further intensify this effect, due to the actual involvement of the consumer in the course of events [Prensky, 2003].

Stories can, with increasing influence on the course of events by the consumer, broadly be categorized into *linear-, non-linear and adaptive* ones. Linear stories follow a fixed stringent plot, on which the consumer has no influence on. Examples for linear stories are ubiquitous (e.g. books, movies). Their dramatic structure has been analyzed for centuries, in order to understand what makes a story compelling. The most prominent example is *Freytag's Triangle* [Freytag, 1872], dividing a plot into *exposition* (where the setting is introduced, and the problem arises), leading via *rising action* to the *climax* (the peak of tension, accompanied by the turning point). From there on, *falling action* finally results in *resolution* or denouncement (Fig. 1 (a)). The idea of dividing a story into three parts dates back to Aristotle [Kaplan, 1993]: "A whole is what has a beginning and middle and end." Christopher Booker gets more concrete, by reducing most existing storylines to "The Seven Basic Plots" [Booker, 2004]. These cover a huge part of known stories and can be a useful tool for authors who need an orientational template for their plot. A very famous linear story template is "The Hero's Journey" [Campbell, 2008]. In 17 stages, it describes the typical story of an inconspicuous person going on an adventure, coping with

[1] `https://en.wikipedia.org/wiki/Mnemonic_link_system`

a difficult task, and returning as a hero. The world in which the story takes place is divided into the *known*, where the hero comes from and returns to, and the *unknown*, where the adventure takes place (Fig. 1 (b)). The pattern is highly recognizable within many popular movies, including prominent examples like *Star Wars* and *The Matrix* [Campbell, 2008]. Non-linear stories also follow a

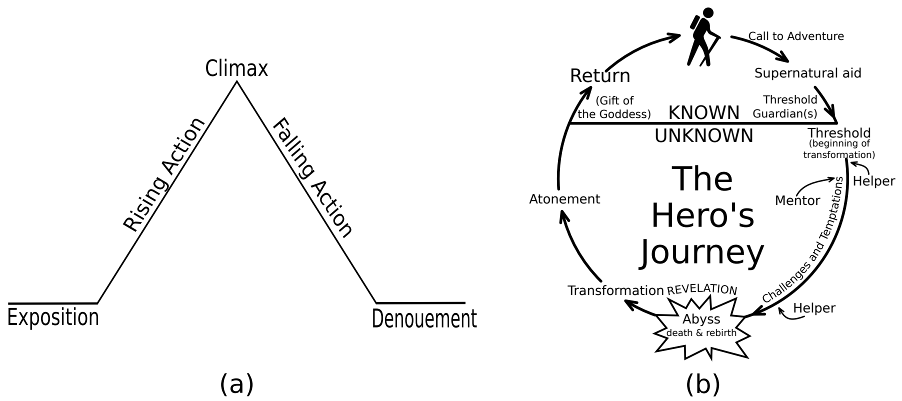

Figure 1: (a) Freytag's triangle; (b) Simplified version of "The Hero's Journey".

plot, but it can take different paths, depending on the consumer's feedback. The definition is broad, and reaches from books with multiple endings to complex video games. With increasing complexity of the possible storylines, the challenge arises, how to maintain a coherent and exciting plot along all of them. The obvious approach of embellishing all possible paths along the story tree may lead to the most individual, while conclusive experiences, but may also quickly become infeasible, as the number of endings can grow exponentially in the worst case (Fig. 2 (a)). A method which is especially popular e.g. in game design, is the *String of Pearls* [Schell, 2014], where the story is interrupted by periods of player freedom of arbitrary complexity, while eventually returning to the main storyline (Fig. 2 (b)). While the overall story technically is still linear, the player's individual experience is non-linear. The fact, that a story is non-linear, influences the type of stories that are told. When a person tells a personal story (e.g. from their life), it will most likely be linear, because that's how it occurred. Non-linear deviations in the story flow might be induced by a listener asking questions, but those will rarely have a dramatic influence on further progress of the story. A higher potential for strong non-linearity (in terms of deviations in the story path) lies in fictional stories, where the consumer can slip into the

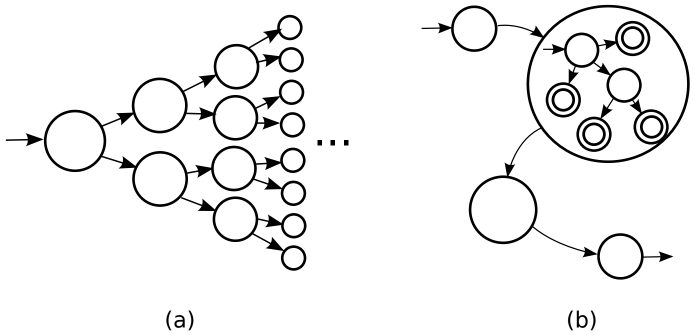

Figure 2: (a) Non-linear story with exponentially growing complexity; (b) "String of Pearls" story.

role of the protagonist. Finally, adaptive storytelling is the idea of making the system telling the story robust to a highly dynamic environment with unpredictable events (e.g. in MMORPGs[2]) in a way, that the author has to consider those as little as possible [Paul et al., 2011]. If for example a player needs an important item from a dungeon, but another one already took it, there has to be another way of obtaining it, as for example by trading or looking somewhere else.

Allowing the creation of non-linear stories provides a larger degree of freedom for both the authors and consumers. But this freedom comes at the cost of an increased complexity in creating the story. *Movement Oriented Design* (MOD) [Sharda, 2007] is an approach assisting non-linear storytellers in wrapping up subject matter into compelling stories. A story is seen as a collection of *Story Units*, where each unit has three components called *Begin*, *Middle* and *End* (BME). The begin of a story unit is meant to "hook" the consumer. The middle carries the main information, and the end either terminates the story or links it to the next story unit. This reduces the nature of a story to the arc of suspense in a concluded story unit. By enforcing this *BME*-scheme, authors are nudged into regarding a basic narrative structure, while having no upper limit on dramatic intensity. This provides a convenient workflow of scaffolding the story in an incremental fashion.

An example of a non-linear implementation of MOD is given in *Media Integrated Story-Telling* (MIST) [Spaniol et al., 2006]. Here, the authors solve the problem with the ambiguous affiliation of multiple BME components by not allowing the recursive split of a BME component directly. Instead, if e.g. a begin should be further split into a begin, middle and end, those have to be combined

[2] Massively Multiplayer Online Role-Playing Game

under a new story unit, which is then associated with the initial begin. The problem of unambiguous story paths is solved by forcing the author to define all transitions manually. The connections have to be drawn in a way, that each possible path through the story would result in a correct linear MOD story. *YouTell* [Cao et al., 2008] is a storytelling system that transfers the features of MIST to the Web and supplements them with a management interface, providing search functionalities and social features like story rating and expert-finding. Liu et al. [Liu et al., 2010] compare the effects of non-linear storytelling and linear storytelling by applying it as a Web 2.0 application of an animated picture book to third graders. They measure a steeper learning curve and higher engagement for the non-linear approach.

General model-driven development had a first peak already during the early 1980s, with the emergence of Computer-Aided Software Engineering (CASE) [Teichroew and Hershey, 1977] tools, but since then model-driven approaches focus mostly on software design and development [Pastor et al., 2008], rather than on story creation. With regard to model-driven methodologies for storytelling, Madni et al. developed a model-driven interactive storytelling approach for analyzing interactions and dependencies within complex systems [Madni et al., 2014]. Recently, also commercial software companies are engaging more and more in creating tools that facilitate storytelling. *Adobe Slate* is a tool for assembling digital photos into linear photo stories for the Web. *Twine* is an open-source tool for creating and sharing interactive, non-linear stories in a text adventure like fashion. Their storyboard graphs allow for flexible plot design. Especially social media platforms also increase their focus on storytelling. *Snapchat* enabled sharing of linear autobiographic stories early on, a feature which lately has been integrated into *Facebook* and *WhatsApp* as well. With *Custom Stories*, Snapchat is currently trying to establish a new focus on collaborative storytelling.

Using dynamic 3D objects for teaching has benefits in terms of vividness compared to 2D objects, as for example static photos of such an object. It can also be argued that similar to mnemonics, like the *Method of Loci* [O'Keefe and Nadel, 1979], learning with 3D objects can be beneficial by engaging spatial memory. Using 3D scanning techniques, rare, fragile or immobile objects can be archived and made accessible for everyone. This third dimension adds more interaction possibilities with the objects and bears the possibility to engage the learner with the object in a way that 2D representations can not provide.

Anatomy 2.0 [Nicolaescu et al., 2015] is a *Personal Learning Environment* (PLE) for the collaborative exploration and annotation of 3D objects for the Web. Users can tag objects in near real-time and gather in online course rooms in which the presentation of the object is guided by a teacher. Especially in the domain of cultural heritage, virtual museums are a large research area (e.g. [Kiourt et al., 2016]). Among others, the *Virtual Museum* by VR3D is a freely walkable virtual museum on the Web, which focuses on ancient Vietnamese relics. It provides explanations and stories around the models, synced with an automatic camera guidance system.

3 Use Case

For describing our approach, we consider two scenarios, story consumption and story creation, to showcase the requirements for different scenarios and user roles our approach supports, and exemplify its contributions. It should be noted, that these roles community members take are not exclusive and can overlap. A storyteller can also for example have the role of a 3D artist (digitizing real-world objects as 3D digital representations) and even be a learner in another area of the CoP's domain. This is in-line with CoP theory, where members can take multiple roles and learn from each other through collaboration, in our case both building stories together and consuming other's stories.

Story Consumption. When consuming a story, the user takes the role of a learner. The learner enters the *Story Consumer View* via her Web browser, chooses a 3D object and selects an attached story. After the story is loaded, she starts her journey through it. The stories are split into units, each unit corresponding to a story page and to a specific view on the 3D object. After each story page, there is either a single transition to the next page, or a decision between multiple possible steps, from which learners select one, thereby experiencing different story paths. While clicking through the story pages, the camera automatically moves around the 3D object, focusing on different parts of it according to the story content. Besides these automatic camera movements, the learner is free to also move the camera manually to have a closer look at certain aspects of interest. The connection between the current point in the story and the 3D object is marked with visual tags appearing on the object. By clicking them, learners can get detailed information on specific features using text, images, links to external Web pages and videos. When the learner reaches one of the possibly multiple endings of the story, she has the possibility to revisit certain parts of the story by using a go-back functionality to experience an alternative path at a fork encountered in previous story pages.

Story Creation. Here, our user is an expert on a certain topic and she wants to communicate and share her knowledge with her community. Thus, she takes the role of a storyteller. Since the story should be told around a certain object she has to have access to a digitized 3D version of it. This could come from other members of the group, then taking the role of a media producer or 3D artist. In the story consumer view, she chooses to create a new story, which directs her to the *Story Editor View*. First, she connects the 3D model with the newly created story. Then, she starts to scaffold her story by first modeling the coarse structure according to MOD, without adding media content yet. She notices some lack of knowledge about a certain aspect of the object and asks another expert for help. The expert joins the story editor view and after a brief introduction (either via the chat provided by the editor itself or third party communication tools) he starts building his sub-story. The users start adding their multimedia content to the story graph and extending it, all in near real-time on the Web. They place tags on certain places on the 3D object to highlight those, and connect the tags with appropriate story sections. They also define camera settings to highlight relevant aspects of the object. Once the story is in a presentation-ready state, the users publish it, so it can be consumed by other members of the learning community. If later on, at any time, a community member discovers a part of the story she can contribute to, it is possible for her to switch back to editing the story and add either additional content or even sub-stories connected to the overall plot to enhance it with new information.

4 Metamodel for Non-Linear MOD Around 3D Objects

We created a metamodel (cf. Fig. 3) that can specify non-linear stories for 3D objects and connections between these 3D objects and media, tags and 3D properties such as position, camera angle and perspective. Due to its suitability for Web 2.0 collaborative storytelling, we choose to reflect the non-linearity using the interpretation of MOD from the MIST approach presented in Sec. 2.

The central entity of our metamodel is the *Story Unit*. A *Story Unit* has to have exactly one *Begin*, *Middle* and *End*, which derive from a *Movement*. All *Movements* have a title and the option for the author to leave a note. The recursive splitting of a *Movement* into further *Begin*, *Middle* and *End* is achieved by associating another *Story Unit* with the *Movement*, which is then split up. The *Media Object*s contain the actual media content presented to the consumer. It is an abstract entity with a title, a caption and again a note, from which the concrete entities *Image*, *Video* and *Text* derive. They constitute the media nodes, that the author will utilize. For dictating the story's flow, there is the *follows* relation between *Media Object*s. Each *Media Object* can have an arbitrary number of successors. These *follows* relations can have a tag and an option to only use the tag-clicking way to trigger the transition. Each *Media Object* can be associated with an arbitrary number of *3D Tag*s. A *3D Tag* has a title, a description, a position, a color and a reference to a remote media file as attribute media. Each *Media Object* can also have at most one *3D View*, whose only attribute is the camera setting. From a conceptual point of view, it is also necessary to be able to associate the *3D Tags* and *3D Views* with the *3D Object* they belong to.

As an experimental feature on top of the MIST layout, we introduce requirements (called *requires* in the metamodel). Their purpose is to enable the story

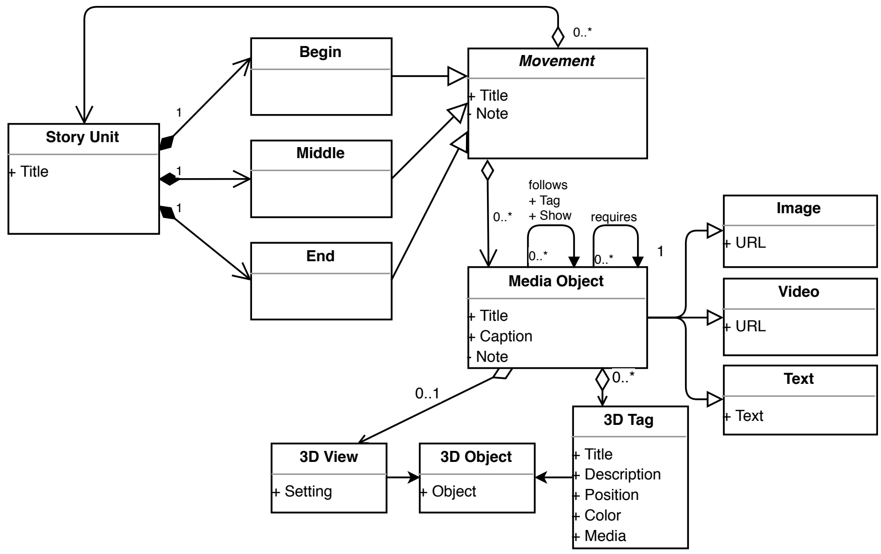

Figure 3: Metamodel for non-linear MOD around 3D objects

author to save some redundancy in the story graphs. If two nodes A and B are connected as $A \xleftarrow{Requires} B$, then during story consumption, B will only be accessible, if the viewer has already visited node A. This reduces the complexity of story setups, where two paths lead to the same or very similar states, which however result in a different progress of the story. Note, that this does not have an impact on the validity of a story, as each story graph using requirement arrows, can be easily converted into one without such, while the story does not change from the consumer's perspective. An example of this is given in Fig. 4.

5 Realization

We implemented our concept using state-of-the-art frontend and backend technologies, resulting in a modular and flexible microservice-based architecture. Fig. 5 provides an overview of this architecture. We achieve separation of functionalities by using Web widgets. These are standalone interface components with a clear-cut functionality, developed in JavaScript using the Google Polymer library[3], an implementation of the Web components W3C standard. The widgets are hosted inside the "Responsive Open Learning Environment" (ROLE) Web

[3] https://www.polymer-project.org

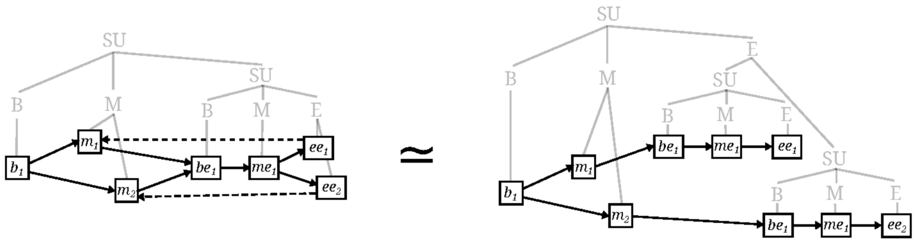

Figure 4: Equivalence of stories with (left side) and without requirement relations (right side). Requirement relations are depicted as dotted arrows.

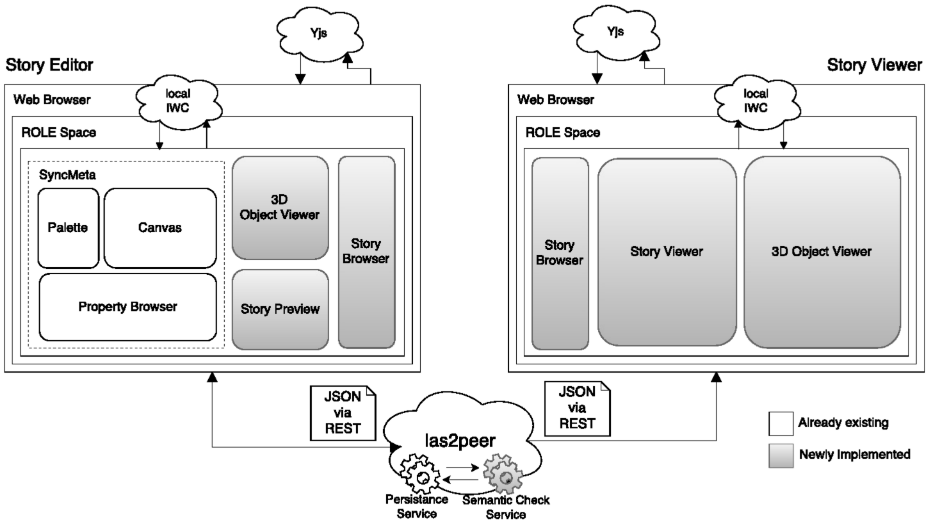

Figure 5: Architecture diagram. Note, that the widgets to be created in the story editor are congruent to the ones in the story viewer.

platform, which offers an SDK to realize inter-widget communication (IWC) [Renzel et al., 2015]. For story modeling, we make use of a Web-based metamodeling framework called SyncMeta [Nicolaescu et al., 2018], in which we integrated the metamodel described in Sec. 4. It also enables the collaborative near real-time capabilities of the editor by integrating yjs, a p2p shared editing framework[4]. We use the X3DOM[5] JavaScript library to display an object,

[4] https://github.com/y-js/yjs
[5] https://www.x3dom.org

adjust camera positions and place tags. This framework allows for integrating and manipulating X3D[6] scenes directly in the browser. In the backend, RESTful microservices are developed on a distributed microservice infrastructure for informal learning communities [de Lange et al., 2018], called las2peer. They ensure the persistence of stories and perform semantic model checks upon changes. While the persistence functionality was already present as an extension of the SyncMeta modeling framework, we added a separate semantic check service that communicates with the persistence service. When a semantic error occurs in a story, meaning that the story in its current state is not compatible with the metamodel, a detailed description is passed to the frontend to inform the author. This can for example happen if a certain requirement relation cannot be fulfilled or if (a part of) the story graph does not completely implement the BME schema.

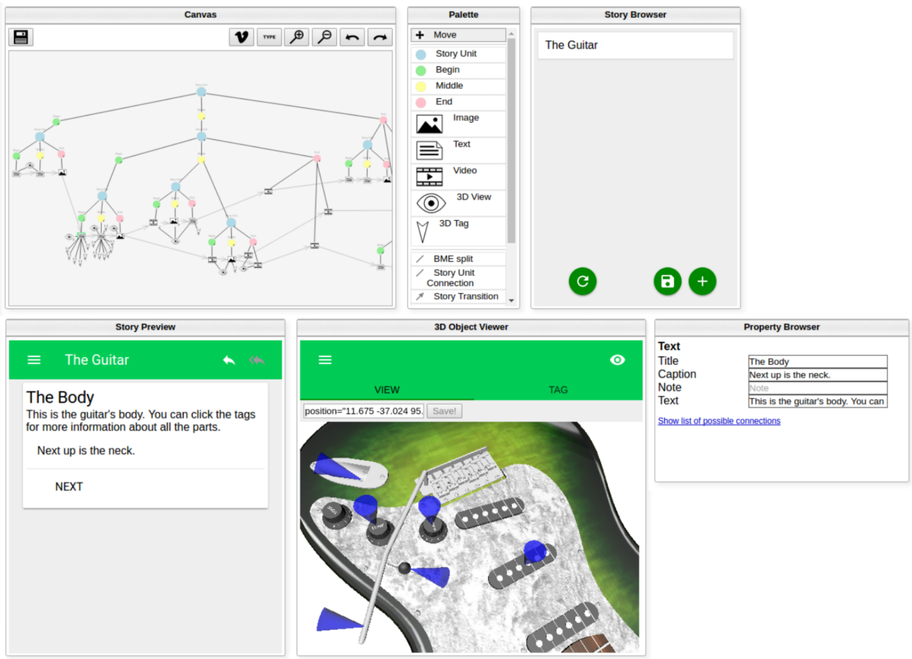

Figure 6: Screenshot of the story editor view

Fig. 6 shows a screenshot of the story editor view. Stories can be created collaboratively in near real-time using the *Palette* and *Canvas* widgets, which

[6] `https://www.web3d.org/standards`

reflect an instantiation of the metamodel presented in Sec. 4. The *3D Object Viewer* widget contains the 3D object and its modification features. By clicking on a certain tag, a window containing text, images and/or videos opens. A "story transition" tag triggers a switch of the story page. When switching the story page, the camera automatically moves to the 3D view associated with it (if available) and the corresponding tags are shown. Users can also move the camera manually around the 3D object. The *Story Preview* widget displays the individual pages of the story and provides the navigation. Its content is updated in near real-time when changes in the model occur, in order to provide a live preview of the current state of the story. The *Story Browser* enables save and load functionality of stories. The story consumer view (not depicted in the screenshot) contains the *Story Browser* widget and additionally two widgets very similar to the *Story Preview* and the *3D Object Viewer* widget. It only lacks the editing options in the latter two, which are only available in the story editor view, that can be accessed at any time via the story browser widget. The system is available as a collection of open source projects on GitHub[7,8].

6 Evaluation

Methodology. We evaluated our approach in two ways. We recruited our first participant group from university students and staff of different fields. Our goal was to evaluate the compatibility of storytelling and 3D objects and how the collaborative aspect of our solution is adopted and rated by the end-users. The user evaluation was conducted in group sessions of one hour with two participants each and a total number of 12 participants, resulting in six sessions. After a short introduction to our prototype, the participants started to work on predefined tasks. The first task was to consume a demo story. In order to evaluate later how intuitively the story viewer could be used, we did not provide a detailed UI-explanation on the task sheet. One of the pages of the story simply contained a "To do" message. When both participants of a session walked through the story to one of the endings, we asked them to go back to this page. They then had to switch to the story editor view, in order to collaboratively replace the mentioned page with meaningful content. The task sheet contained different sub-tasks for each participant, ensuring that they collaboratively created a sub-story, integrated into the overall story. This also included tagging certain aspects of the 3D model. At the end of the session, the participants were asked to fill out a questionnaire.

We also evaluated our approach with a domain expert in the field of teaching medicine in order to assess the applicability of our approach for learning scenarios. Here, we did not prepare a fixed set of tasks, but rather provided a series of

[7] `https://github.com/rwth-acis/3DStorytellling`
[8] `https://github.com/rwth-acis/las2peer-3DStorytelling`

tutorial videos[9] on how to use the editor and the concept of MOD. This way, the domain expert was provided with enough information to acquire experience with the tool autonomously during a longer usage period. After two weeks, he was asked to author a custom story from scratch. This allowed a least-biased view on the system, giving us as much insight into its perception and room for improvement as possible. We prepared semi-structured interviews, with a focus on open questions, which contained specific questions around the concept of story structuring and inquired about the functionalities of the editor.

User Evaluation Results. One main focus regarding usability was the intuitiveness of the story consumer view. This was rated with a good score of 4.33 ($\sigma = 0.49$) on average, with comments being very positive ("Very simple and intuitive"). As a critique point, the participants commented that it was sometimes not obvious, when one has to decide between multiple options for the next story page. During the session we made the interesting observation, that many users tend to only click the rightmost button in the first run. Afterwards, we tried to already deal with this by simply highlighting the buttons to better separate them from the background and each other. However, we cannot rule out that this could have influenced the perception of our participants during the evaluation. For the story editor view, a user stated that she understood and valued the usefulness of the tool after about 20 minutes into working with it. The live preview aspects of the widgets in the story editor view were rated 4 ($\sigma = 1.13$). Although rather high, the feature requires slight implementation improvements due to its prototypical nature during the evaluation ("Widgets sometimes need to be reloaded", "Just sometimes I had to refresh the canvas myself"). One of our goals was to evaluate the acceptance of non-linear stories for educational purposes. This was a complicated task, as the participants only had limited time to get familiar with the concept. The participants seemed to prefer the idea of non-linear stories both over linear ones, structured information without a story, and completely unstructured information. An average agreement-score of 3.58 ($\sigma = 1.13$) for "no structure" was the lowest of those options. Also very contentious was the question, whether the plot was helpful for keeping the participant interested. It was answered with an average score of 3.4, but with $\sigma = 1.43$, it sadly allows no profound conclusion. Again, factors that might have influenced this result were the limited time of the sessions and of course the perception of the individual demo story we used. With an average rating of 4.5, the participants said that they appreciated the functionality for having a second look at certain information and, with a very high score of 4.75 ($\sigma = 0.45$), for choosing in alternative story path. This highlights one fundamental challenge with non-linearity in digital storytelling. If one's goal is to learn, one will feel like missing something, as soon as one has to decide between various

[9] `https://www.youtube.com/watch?v=enKijrMpYe0`

different paths. One participant clearly stated “Because of the not linear story I have not seen every part of the information and thus have not learned everything that I could have”. Another point within the conceptual questions was to find the ideal use case for our system. We asked the participants how much they would appreciate non-linear storytelling for learning purposes in different scenarios. The results showed an increase in acceptance with decreasing formality, with a lowest average score of 2.92 for university education, and a highest average score of 4.42 for museum-like websites. A question regarding both future work and the potential of near real-time collaborative story editing, was about the different editing workflows. While the collaboration in general was perceived very positive, the participants seemed to appreciate workflows in which they interfere with each other the least. Their most appreciated workflow was working separately on own parts of the story, closely followed by having different roles (like “Story Scaffolder” and “Media Producer”, as we also provided in the given tasks). Working unorganized had the least approval. Organization seems to be a very important aspect of the collaborative editing, with one participant for example suggesting the use of a voice chat. As a matter of fact, one observation we made was that the teams finished the tasks faster and more accurate, the more they communicated with each other.

Domain Expert Evaluation Results. The results of our second evaluation indicate that overall, the reviewer felt comfortable in using the editor after viewing the tutorial videos. While very positive about non-linearity of stories, he was doubtful about MOD. It turned out, that he intended to use our system for creating content in the direction of flexible online courses, rather then “explicit” storytelling. Although he appreciated the idea of a BME structure, he wanted each of those components to potentially be longer then just one media object. Especially in the case of a middle object, he wanted to use multiple consecutive ones, to do more extensive explanations. Another suggestion was, to create “excursions” inside the story. These would be optional sub-stories, that return to their entry point when finished. While this concept is surly interesting to invest, it is hardly compatible to MOD. Aside from that, he appreciated the idea of also building stories around multiple objects and separately displayable parts of the same object. He also liked the idea of having the possibility to hide certain parts of the story graph in order to reduce its complexity while editing.

Discussion and Implications for Future Work. It can become challenging to tell a long story using only a single 3D object. We consider the possibility for including multiple 3D objects into a single story to be a relevant direction for future work. However, this has to be reflected in the metamodel and the complications introduced into the story modeling process need to be computed. Approaches that introduce different modeling views could be a solution for this. Non-linear storytelling for learning purposes appears to be complicated without providing a way to revisit previous decisions. Without MOD, this problem would be manageable by allowing story branches to return to their origin by design. It is then up to the author to structure the story in a way, that no important information can be missed, and up to the consumer to trust the author. With MOD this is not the case, as backward transitions are a clear violation of the paradigm. The implementation of some way to return to previous story sections seems to be inevitable for a good user experience and for effective learning. Another challenge for future work is the possibility for optional excursions mentioned in the evaluation. We also recognized the need for the possibility to provide multiple media/pages per movement. Finally, we are currently working on enhancing story consumption with augmented reality, where the 3D objects can be explored in even more engaging ways, while consuming the story attached to them.

7 Conclusions

In this paper, we presented an approach for near real-time collaborative, non-linear storytelling around 3D objects on the Web, based on a well-defined meta-model. The collaborative aspect of our approach enables new ways of how stories can be created and enhanced. Future work in this domain may highlight the benefits of this kind of collaboration and provide additional valuable insights in the dynamics of the collaborative aspects. The connection of storytelling with 3D objects is novel and our evaluation provides first insights into the usefulness of this approach. We strongly believe that both storytelling as well as digitizing real-life artifacts are two great ways for improving learning processes, and the combination of both provides interesting and promising results, which build the basis for future work in this domain. With our contribution, we provide a first step in this direction and give first impressions of the power and expressiveness this new way of interactive non-linear storytelling provides. We are confident that future research in this domain and applications in the form of interactive stories available via the Web, will provide great new ways of learning with collaborative, non-linear storytelling.

References

[Annie Murphy Paul, 2012] Annie Murphy Paul (2012). Your Brain on Fiction.

[Booker, 2004] Booker, C. (2004). *The Seven Basic Plots: Why We Tell Stories.* A&C Black.

[Bower and Clark, 2013] Bower, G. H. and Clark, M. C. (2013). Narrative Stories as Mediators for Serial Learning. *Psychonomic Science*, 14(4):181–182.

[Bransford et al., 1999] Bransford, J. D., Brown, A. L., and Cocking, R. R. (1999). *How People Learn: Brain, Mind, Experience, and School.* National Academy Press.

[Brown and Duguid, 2000] Brown, J. S. and Duguid, P. (2000). *The Social Life of Information.* Harvard Business Review Press.

[Campbell, 2008] Campbell, J. (2008). *The Hero with a Thousand Faces*, volume 17. New World Library.

[Cao et al., 2008] Cao, Y., Klamma, R., and Martini, A. (2008). Collaborative Storytelling in the Web 2.0. In *First International Workshop on Story-Telling and Educational Games*.

[de Lange et al., 2018] de Lange, P., Göschlberger, B., Farrell, T., and Klamma, R. (2018). A Microservice Infrastructure for Distributed Communities of Practice. In *EC-TEL'18*, pages 172–186. Springer.

[Freytag, 1872] Freytag, G. (1872). *Die Technik des Dramas*. Hirzel.

[Kaplan, 1993] Kaplan, T. J. (1993). Reading Policy Narratives: Beginnings, Middles, and Ends. *The argumentative turn in policy analysis and planning*, pages 167–185.

[Kiourt et al., 2016] Kiourt, C., Koutsoudis, A., and Pavlidis, G. (2016). DynaMus: A Fully Dynamic 3D Virtual Museum Framework. *Journal of Cultural Heritage*, 22:984–991.

[Liu et al., 2010] Liu, C.-C., Liu, K.-P., Chen, G.-D., and Liu, B.-J. (2010). Children's Collaborative Storytelling with Linear and Nonlinear Approaches. *Procedia-Social and Behavioral Sciences*, 2(2):4787–4792.

[Lukosch et al., 2011] Lukosch, S., Klebl, M., and Buttler, T. (2011). Utilizing Verbally Told Stories for Informal Knowledge Management. *Group Decision and Negotiation*, 20(5):615–642.

[Madni et al., 2014] Madni, A. M., Spraragen, M., and Madni, C. C. (2014). Exploring and assessing complex systems' behavior through model-driven storytelling. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1008–1013. IEEE.

[McDaniel and Pressley, 2012] McDaniel, M. and Pressley, M. (2012). *Imagery and Related Mnemonic Processes: Theories, Individual Differences, and Applications*. Springer Science & Business Media.

[Nicolaescu et al., 2018] Nicolaescu, P., Rosenstengel, M., Derntl, M., Klamma, R., and Jarke, M. (2018). Near Real-Time Collaborative Modeling for View-Based Web Information Systems Engineering. *Information Systems*, 74:23–39.

[Nicolaescu et al., 2015] Nicolaescu, P., Toubekis, G., and Klamma, R. (2015). A Microservice Approach for Near Real-Time Collaborative 3D Objects Annotation on the Web. In *ICWL'15*, pages 187–196. Springer.

[O'Keefe and Nadel, 1979] O'Keefe, J. and Nadel, L. (1979). The Hippocampus as a Cognitive Map. *Behavioural and Brain Science*, 2(4):487–494.

[Pastor et al., 2008] Pastor, O., España, S., Panach, J. I., and Aquino, N. (2008). Model-Driven Development. *Informatik-Spektrum*, 31(5):394–407.

[Paul et al., 2011] Paul, R., Charles, D., McNeill, M., and McSherry, D. (2011). Adaptive Storytelling and Story Repair in a Dynamic Environment. In *Interactive Storytelling*, pages 128–139. Springer.

[Prensky, 2003] Prensky, M. (2003). Digital game-based learning. *Comput. Entertain.*, 1(1):21–21.

[Renzel et al., 2015] Renzel, D., Klamma, R., Kravčík, M., and Nussbaumer, A. (2015). Tracing Self-Regulated Learning in Responsive Open Learning Environments. In *ICWL'15*, pages 155–164. Springer.

[Schell, 2014] Schell, J. (2014). *The Art of Game Design: A book of Lenses*. CRC Press.

[Schumann et al., 2013] Schumann, J., Buttler, T., and Lukosch, S. (2013). An Approach for Asynchronous Awareness Support in Collaborative Non-Linear Storytelling. *Computer Supported Cooperative Work*, 22(2-3):271–308.

[Sharda, 2007] Sharda, N. (2007). Applying Movement Oriented Design to Create Educational Stories. *International Journal of Learning*, 13(12):177–184.

[Spaniol et al., 2006] Spaniol, M., Klamma, R., Sharda, N., and Jarke, M. (2006). Web-Based Learning with Non-linear Multimedia Stories. In *ICWL'06*, pages 249–263. Springer.

[Teichroew and Hershey, 1977] Teichroew, D. and Hershey, E. (1977). PSL/PSA: A Computer-Aided Technique for Structured Documentation and Analysis of Information Processing Systems. *IEEE Transactions on Software Engineering*, 3(1):41–48.

[Wenger, 1998] Wenger, E. (1998). *Communities of Practice: Learning, Meaning, and Identity.* Learning in doing. Cambridge University Press, Cambridge, UK.