

Perceived Consistency and Coherence in Collaborative Story Worlds

Alan Tapscott

TESI DOCTORAL UPF / ANY 2017

Directors de la tesi
Dr. Josep BLAT
Dr. Carlos LEÓN

21 de juny del 2017



*The horizon melts into a limitless question mark, and
like the cartographers of old, we glimpse yawning
monstrosities and mind-forged utopias beyond the edges
of our paltry and provisional maps.*

— Erik Davis, *Techgnosis: Myth, Magic + Mysticism in
the Age of Information*

ABSTRACT

Current information technologies allow users to interact and collaborate to create fictional realities. In this thesis, we explore the role of perceived consistency and coherence in the collaborative creation of fictional *story worlds*. We conduct a series of studies focused on author contributions and interactions with story worlds. The studies include controlled experiments involving volunteer students using prototype non-linear timeline visualizations, semantic representations of stories and text generation systems. We also explore real world large-scale participative story world community sites, conducting a thorough analysis of one of the most popular ones. Our main findings imply that contributors and readers perceive one single continuity or version of facts, that they expect to be consistent and coherent. Timeline visualizations and semantic diagrams generated satisfactory collaborative user experiences for story worlds. The contribution sequence was a central factor, with initial content significantly influencing the perceived consistency and coherence. We successfully increased the reader perceived consistency by introducing a first approach to formal into the authoring prototypes. Our real story world analysis suggests that readers and authors read and contribute into specific *collaboration dimensions*. These content subsets are centered on specific elements or themes and reduce the need for general consistency and coherence. Based on our results, we propose a model for measuring contribution plot integration (hypothetically linked to consistency and coherence) in collaborative story world community sites.

Keywords: *Story World, Collaborative Writing, Consistency, Coherence, wiki.*

RESUM

Les tecnologies de la informació actuals permeten als usuaris interactuar i col·laborar per crear realitats fictícies. En aquesta tesi explorem el rol de la consistència i la coherència percebudes en la creació col·lectiva de móns narratius imaginaris. Conduïm una sèrie d'estudis enfocats en la interacció i les contribucions dels autors. Aquests estudis inclouen experiments controlats que involucren estudiants voluntaris que han fet servir prototips de visualitzacions de narrativa no lineal, representacions semàntiques d'històries i sistemes de generació de text. També explorem webs de comunitats d'internet participatives que contenen móns narratius de gran escala i fem un processament de text d'un d'aquests móns més recents i populars. Els nostres resultats principals impliquen, per una banda, que la majoria de participants i lectors perceben una sola continuïtat o interpretació dels fets, i per l'altra banda, que s'esperen que sigui consistent i coherent. Les visualitzacions de línia temporal i els diagrames semàntics han generat experiències d'usuari satisfactòries per a la interacció amb móns narratius. La seqüència de contribució d'autors ha resultat ser un factor central, el contingut inicial de la qual té una influència significativa en la percepció de la consistència i la coherència. Hem aconseguit incrementar la consistència que perceben els subjectes introduint una primera aproximació de model formal als prototips. Els nostres estudis en un context real, fora del laboratori, suggereixen l'existència de dimensions participatives de col·laboració no lineals. Aquestes dimensions, similars a arcs de personatge, estan centrades en elements o temes específics, i redueixen les contradiccions del món narratiu. Sobre la base d'aquests resultats, proposem un model que permet mesurar la integració argumental, hipotèticament vinculada a la consistència i la coherència, de móns narratius col·laboratius com els que estudiem.

Paraules clau: *Mon narratiu, Escripura col·laborativa, Consistència, Coherència, wiki.*

ACKNOWLEDGEMENTS

A **Joaquim Colás**, por aquella charla en el bar de Ciutadella que nos ha llevado a esto. Qué mala costumbre tenemos de venirnos arriba en el desayuno. Quiero aprovechar esta publicación para desmentir los rumores sobre nuestra relación o parentesco.

A **Josep Blat**, por haber aparcado la TARDIS un tiempo para supervisarnos con consejos que darían para un libro mejor y también por corregir mi enésimo “culteranismo” de bar.

A mi otro supervisor, el **León Carlos**, por explicarme innumerables veces de qué va esto de la inteligencia artificial, la ciencia y las cosas.

A **Ayman Moghnieh**, por perseguirme por la estación de Francia para que acabara el TFG y estar siempre preguntando por la tesis. Aquí la tienes, pero no te hagas daño levantándola (o leyéndola).

A **Pablo Gervás** y la **gente de NIL** por acogerme en la UCM con cariño. ¿Ya sabéis cómo puede el alquimista Sánchez servir plasma en vasos de vidrio? A **Javier Gómez** por su eterna hospitalidad y constante ayuda logística durante mi fase de vacío cuántico. A **Susana Bautista**, por el infinito refuerzo positivo y por hacerme un hueco en su departamento.

A **Xavier Ruiz Collantes**, por su curso intensivo de narratología que debería estar patentado. O en fascículos. ¿Quizás un audiolibro? Sí, eso, pero con su voz radiofónica. A **la gente de comunicación** por aceptarnos tal y como somos (informáticos).

A **Miguel Ángel Carralero** y **la banda del 305**. Puede que nos hayan robado el despacho, pero Benedict protege. Solía haber un agujero ahí, pero ya no está.

Al **personal del DTIC o ETIC** por su soporte indispensable. Mención especial para **Lydia García**, cuya buromancia, conocida más allá de Guadalajara, nos salvó demasiadas veces.

A **Valeria Righi** y **Yolanda Colás** por su aportación etnográfica y de ciencias sociales. Lo intenté.

A **mi familia**, por su apoyo. *I made it, you were right after all*. Este no es el libro que esperabais, pero no tengo otro. A **mis suegros** por el delicioso suplemento nutricional extremeño-galaico e hospitalidad general. También a **Mía y Luca** por los largos silencios.

Y por supuesto a la mejor compañera, **María Pérez**, por hacer más por mí que nadie con sus designios inescrutables y corregir esta página. Ya puedo pintar el pasillo.

Carme, Teresa, bon viatge. Al final he fet els deures.

Contents

ABSTRACT	iii
RESUM	iv
ACKNOWLEDGEMENTS	v
1 Introduction	1
1.1 Motivation and research area	1
1.2 Research questions and main results	6
2 Exploring the Collaborative Creation of Story Worlds	11
2.1 Writing Consistent Stories based on Structured Multi- Authored Narrative Spaces	14
3 Evaluating Visual Story Diagrams and Text Artefacts for Author Participation and Inspiration	39
3.1 Modifying Entity Relationship Models for Planning Sto- ries and its Impact on Potential Authors	42
3.2 Empirical Evidence of the Limits of Automatic Assess- ment of Fictional Ideation	58
4 Author Contribution and Dynamics in Large-scale, Long- term Conditions	77
4.1 We Built Our Own Worlds - Story Canonicity and In- direct Collaboration in a Shared Story World	80
4.2 Large-Scale Collaborative Story Worlds: Formalizing Content and Author Dynamics	104
5 Discussion	153
5.1 Consistency and coherence in a collaborative story world	153
5.2 Thematic contribution and navigation	156
5.3 Inter-author interaction	157
5.4 Author-story world interaction	159
5.5 Creativity, coherence and consistency	160
5.6 Innovative directions for collaboration in story worlds	160

7	Conclusions and future work	163
7.1	Conclusions	163
7.2	Future Work	165
	REFERENCES	167

Chapter 1

Introduction

1.1 Motivation and research area

The development of information technologies is changing our society in many, unprecedented ways. Online platforms and digital interactive media have been instrumental in the development of systems that bring together author and reader communities to compose and consume multi-authored stories through specialized platforms. In this context, the audience is not only interested in rich narratives, but also wants to participate in their development by adding and sharing their very own creations, compositions, and ideas. Nowadays, following current trends, people actively publish and share thousands of creative works, inspiring each other through the web. These collaboration dynamics keep evolving, with more authors expanding the content, structure, and knowledge value through innovative creation and composition processes, contributing with many new ideas that essentially redefine traditional creative processes. This context represents a fertile ground for research, with potential findings that might help to understand the processes involved in the collaborative construction of fictional realities, where the research of this thesis is framed.

Within this context, this thesis uses the term *story world* to describe the conceptual dimension of a story, containing the facts and plot elements (such as characters and locations) that constitute the emerging fictional reality established by the author(s). In the field of Artificial Intelligence (AI) and story planning, for instance, story world is used to refer to the information relative to the story itself, and not to the narrative plan used by a machine author [29, 35]. The term *storyworld* (with both words joined) has also been used for referring to the environment in which a video game takes place [3] or to

define games that embed the narrative in the environment [21]. This joined version of the term, is often found in studies of transmedia narrative to refer to a similar concept. [13], however, refused to use the term due to its ties to game design elements. A similar concept, *information space*, has been used in a different field: [31] proposed the following definition: "The set of concepts and relations amongst them held by an information system". This thesis uses the term story world understood as an information space composed by the set of elements (e.g., characters, locations, plots, motivations, rules) that constitute the world implicitly defined by a story, along with the relationships among them, in other words, a set of causally interlinked entities relevant to a specific story or set of stories. According to Tolkien and his conception of *secondary world* [49] (a term close to that of story world), all the elements (i.e., geography, characters, language and timeline) are interdependent and require internal consistency to suspend disbelief, becoming credible to the reader. In the story world a certain consistency and coherence is expected, even if facts never actually happened, or if they happen in a fictional reality. However, story worlds are usually constructs that emerge and develop with stories, but are rarely formulated explicitly. As a rare example, Tolkien's *Silmarillion* [48]-a collection of notes describing the fictional reality used in his novels-was published posthumously. In a context of collaborative writing, in which each author might have her/his own mental conception and potentially plan for the story, how does a story world come into play? The question is even more relevant when dealing with an online context, in which author communication and mediation is heavily conditioned by the site design. [1, 39] discussed extensively the *common information space* they introduced, defined as "...a central archive of organizational information with some level of "shared" agreement as to the meaning of this information (locally constructed), despite the marked differences concerning the origins and context of these information items." This concept is aligned with this thesis' conceptual framing of a shared story world, in the sense that it contains relevant data introduced by different authors and despite being potentially diverse, it must be agreed or coherent to some degree. [1] also discusses the dialectical nature of common information spaces and the challenge of putting information in common and interpreting it. Despite being related to other scenarios, most of the considerations for common information spaces are very likely to appear in a shared story world and must be addressed. [21] visionary work on *convergence culture* anticipated the deep implications of the internet users and their innovative activity, collaboratively curating popular media

in public sites. He highlighted the parallels with popular culture and discussed the potential repercussions. The current popular transmedia and cross-platform franchises seem to be following this vision. In the midst of an ongoing discussion regarding transmedia, [22] tried to provide a clear definition for the term: “Transmedia storytelling represents a process where integral elements of a fiction get dispersed systematically across multiple delivery channels for the purpose of creating a unified and coordinated entertainment experience. Ideally, each medium makes its own unique contribution to the unfolding of the story”. This unified and coordinated entertainment experience he mentions is another allusion to the shared story world we are framing in this thesis. The presence of story worlds in modern media discussions has been increasing in parallel with transmedia multimodal stories. Popular media extends through a wide array of platforms and formats, yet the essential story elements persist. These story elements, the story world, exist in the popular consciousness, detached from their original platform, in a similar way to ancient mythology or popular folk tales. Some of the most popular characters from relatively modern media have gradually permeated our culture. Each character belongs to its own story world, defined by their creator’s work, and can be examined and discussed partially detached from it, something that happens often in the eventual official or unofficial crossovers and referential works that get published as merchandise and extensively document these fictional characters and realities.

More recently, fandom has irrupted into the media scene. There are two distinct phenomena deeply rooted in fan culture that are worth mentioning. First, fans undertake a rigorous task of documentation with the object media of their fandom. They use wiki sites to document their favorite story worlds from popular movies, TV shows, comics or any other relevant media, following a crowdsourcing collaboration pattern to create online fiction encyclopedias. One of the most popular platforms, *Wikia*-a derivative wiki site-, is used to represent explicit descriptions of story worlds. As of May 2016, there are 30643 active Wikia sites¹, many of them used to document the elements from specific fictional story worlds. Using a platform designed and generally used as a digital encyclopedia, these communities extensively document fictional realities in these public platforms collaboratively. Another example is *A Wiki of Ice and Fire*², the website that holds most the information related to the *A Song of Ice and Fire*

¹<http://wikis.wikia.com/>

²<http://westeros.org/>

book series in its many articles. It represents another instance of a wiki-like website fed with content from multiple contributors, properly structured and published in a readable way. The site publishes encyclopedic articles that cover most of the books related to the canonical material created by the original author. The emergent story world from the original books is measured and explained by the many wiki curators. The Wikia example represents one of the most archetypical instances of an explicit story world. Some works have studied these fan-related phenomena, *Narractivity* [7], *Documentary Simulacra* [12] and wikis and participatory fandom [30]. The contents of these wiki sites describe the causal and temporal connection between actions to document story worlds, as in the concept of *narrative space* [31]. Collaboration on wiki platforms has also been the subject of many studies, as [27] perspectives on its value as a collaboration platform, [24] interpretation of the revision history as a collaboration network, [26] analysis of collaboration patterns and article quality, or [42] social dynamics. Overall, wikis have been often presented as a good example of collaboration, even to the point of being considered “good democracy” [27], and perhaps they represent a viable solution to the challenge of building collaborative story worlds.

The second case often found in fandom communities, is the exercise of expanding the narrative space of their favourite entertainment franchises. By introducing their own stories that are intertwined with a well-established narrative universe and its mythology, rich networks of original fan-driven fiction that coexist with the official material are created. Online information technologies provide momentum to these phenomena, exponentially augmenting their size, reach and impact. For instance, FanFiction³ or Quotev⁴, gather hundreds of thousands of users around hundreds of franchises, presenting fan-developed stories that extend the story world of popular media. They provide a good example of amateur and professional authors creating stories in a collaborative story world, contributing to it unofficially. Most of these stories, however, do not take into account the contributions of their fellow fan authors and represent a one-off extension of the original story world, limiting their potential to enrich the original material or even to span new fictional parallel realities or story worlds. There are a few exceptions in which these fan works develop a separate identity and characteristics, such as the *Fifty Shades of Gray* series, originally written as an amateur fan fiction from the *Twilight* series.

³<https://www.fanfiction.net/>

⁴<https://www.quotev.com/>

Without making any claims about its literary quality, we remark that *Fifty Shades of Gray* became a worldwide phenomenon, with more than 100 million books sold. Derivative works might be more than recreational creative activities, and might have a fundamental role in the future of media in the digital era. The fandom phenomenon has been extensively studied over the recent years. [14] discusses the main characteristics of fandom from the perspective of economics, highlighting how this phenomenon establishes an alternative culture that goes against the “official” one. This perspective reinforces that story worlds are used to collectively creating an alternative reality. [16] analyzed the *Lost* TV series and their fan activity on the net, discussing the encyclopedic and creative dynamics while focusing on the spoilers revealing crucial information to the audience during the chapters’ original run. The authors elaborate on the community’s struggle to establish some sort of order to the chaotic contributions from fans. The line between factual, hypothetical, speculative and original content was not a very clear one, especially when documenting fictional media that often resorted to mystery or intrigue tropes. Again, fan contributions are struggling to become more than a reflection of the original media. [40] presents a model that brackets the opposing potentialities of internet influence on offline society using large scale participatory fandom to centre her discourse. She discusses fan clubs, online producer-consumer affiliations and real-world legal controversies in the context of fans attempting to participate in the media they revere.

There are certain tools that support readers and writers who are contributing to a common story world. `articy:draft`⁵ and `Celtx`⁶ are both collaborative tools meant for structural creative story development and represent good examples of this trend. These kind of tools are often created by the same company responsible for the final narrative content, meaning that their design might be less generalist and more ad hoc, seeking to solve the specific problems of their usual narratives and media. It is worth noting that tools do not merely intend to support the construction of a story in the sequential, traditional way. They provide mechanisms that allow for free, divergent exploration of all the related information, supporting the non-linear growth of narrative spaces. `Storytron`⁷ is an interesting approach to developing a commercial tool that would allow users to design interactive stories. Although it is currently on-hold due to problems

⁵<https://www.nevigo.com/en/articydraft/>

⁶<https://www.celtx.com/>

⁷<http://www.storytron.com/>

regarding the learning curve (i.e., the complexity of building a whole interactive story with the tool), this approach is interesting in terms of decomposing the narrative space into a set of unitary elements, and defining the logic that relates them. Storyjacker⁸ is another interesting example closely related to the tradition of the Exquisite Corpse writing technique. This game proposes that its players first read a flash fiction (roughly between two or three hundred words) created by another writer with an explicit editorial challenge attached to it. Players rewrite the text answering the challenge and pass the result to the next player, introducing a new challenge of their own. While this approach is a game, the writing dynamics of its multi-author design are interesting and not very far away from what this thesis envisions.

1.2 Research questions and main results

Within the context described in the previous section, our focus is to explore and understand how *to support collaborative creation, composition, and consumption of multi-authored, story worlds*, when the latest information technologies and online communities are redefining these creative processes.

More precisely, we present next the research questions posed in this thesis.

- **Are the perceived consistency or coherence significant factors in the construction of collaborative story worlds?**

The first question is based on the hypothesis that it would be important to preserve consistency or coherence when several authors contribute to the same story world. The internal consistency or coherence has been highlighted in some of the works related to collaborative work, for instance in the case of a shared narrative space [1, 39] the Lost wiki site [30] or in The Million Penguins wikinovel [28]. If more than one author is involved, facts need to be communicated, coordinated and negotiated [2]. More specifically this thesis aims at assessing whether this concern about consistency or coherence is present in the people who collaborative create of a story world. The research tried to elucidate it through three small scale experiments, a large paper prototype without software support, and two additional ones involving a software prototype with an underlying model of narrative consistency.

⁸<http://www.storyjacker.net/>

- **How can we make the collaborative creation of story worlds less intrusive and more participative?**

The second question is based on the hypothesis that semantic diagrams and computational creativity can assist authors in building a collaborative story world. Specifically, in the research two different approaches to achieve this goal were implemented. In the first one authors were provided with a visual, intuitive mechanism based on semantic diagrams to plan stories. This was partially based on observations from previous work in which authors used diagrams to structure their story and link contributions successfully [43]. Semantic diagrams have been extensively used to document complex systems while allowing a certain degree of abstraction. The research chose this tool to avoid the bad reception of the model and its constraints by providing authors with a more flexible tool. The second approach delved into the field of AI and computational creativity to measure creativity on artefacts meant to boost author inspiration and quality. We hypothesized that using modern approaches to text generation with the right parameters might successfully inspire authors to be more productive.

- **Which are the content types and author dynamics of real collaborative story worlds?**

The third question is framed in complementing the small scale controlled experiments related to the first two questions, with the analysis of large scale collaborative narratives - live and uncontrolled, in the real world -. On one hand the work attempts to see whether the findings regarding consistency and coherence from the first question are replicated in a more general and real context involving many authors. On the other hand, the research tries to determine if semantic information is helpful for the authors and their contributions. A longitudinal study in a relatively controlled environment was performed. Second, we conduct an in-depth analysis of a popular internet site that allows users to collaboratively create a story world, using web parsing techniques. The research assumes that this generalizing process will allow to answer the question by generalizing the small-scale observations in real author communities and their story worlds.

Each of the chapters 2, 3 and 4 is focused in answering one of the previous questions.

Chapter 2 focuses on the specific problems that arise in a multi-author story context and specifically in the potential issues with consistency and coherence. We envisaged to analyse the participants' contributions within the framework of the shared story world concept previously discussed. The main goal was to isolate the exercise of creating a story set in a fictional reality collaboratively in a way that would allow us to observe significant factors for authors and their contributions; we tried to recreate the dynamics of large-scale multi-author franchises - such as using recurring characters and non-linear plots - in a controlled environment. We were specially interested in observing how authors managed potential contradictions. To the best of our knowledge, there is no literature that focuses on the author and reader perception of consistency and coherence.

In the paper contained in this chapter [43] an observation (including direct subject feedback) of an open process to create a shared story world and a first attempt to provide a tool specifically suited for collaborative authoring of consistent and coherent fiction are presented. Our approach relies in custom-built paper mock-up and software prototypes, purposely avoiding popular platforms meant for similar tasks. Through the observation of their actions and multiple questions, a better comprehension of their motivations was intended. Also, the small scope of the experiment allowed to interview and observe directly every participant. The results pointed towards a predilection for sequential visual metaphors to organize scenes and a predilection for a single continuity. In-situ visual observations and questionnaires suggested that consistency and coherence were one of the main motivations and concerns, implying that they are important factors in exercises of collaborative storytelling.

In the second section of the same paper, a formal model based on the results of previous experiments to enhance consistency was introduced. After its inclusion, authors valued the experience very negatively, but readers perceived the result as more consistent. The formal model was more popular amongst authors when used to provide consistency recommendations. There was a popular reception for storylines, explicit mechanisms to link author contributions and provide reading guidance. A certain territoriality and hesitation to modify other author's contribution for creative or artistic purposes was observed. This effect was not observed when the contributor was rearranging content or modifying it to correct what he or she she perceived as plot holes or causality discontinuities, mostly caused by the high amount of authors. The results were successful in establishing and isolating new and significant factors for the collaborative stories

in a shared fictional reality or story world. Mainly, that most authors considered that there was only one continuity for all contributions and that one of their main contributing motivations was to preserve the perceived consistency and coherence. There seems to be no publication devoted to studying such concerns in the context of collaborative fiction writing.

Chapter 3 focuses on understanding and supporting the author interaction observed in the initial research discussed in chapter 2. In this case, authors were provided with a tangible story world semantic diagram and computer-generated textual artefacts meant to boost quality and creativity.

Specifically, the first paper contained in the chapter [44] studies the usage of semantic diagrams to connect and enhance author contributions to a shared story. We propose a basic methodology for translating textual stories into an equivalent semantic diagram, and evaluate its efficiency for authors to understand and expand the plot compared to traditional text representation. We hypothesized that visual, tangible elements are favourable for the collaborative construction of a story. We used the basic terminology proposed by [9] to create a define a basic methodology meant to translate simple english sentences into its semantic model equivalent. The goal was to obtain an intuitive, informative representation of the story world that emerges from author interaction. This representation could be used later to communicate and extend stories, optimizing author interaction in the context of a collaborative story world. Results were generally favourable in terms of usability, comprehension and memorability, suggesting that semantic diagrams are a functional metaphor to represent and manipulate the contents of a small, sample story world.

In the second paper of the chapter [47], the evaluation of narrative artefacts to explore their perceived potential for stories and quality was explored. Particularly, trying to find a set of metrics that correlate to quality and narrative potential. The approach exploits a mixture of human evaluation and machine learning to obtain a reliable measure of these parameters generally considered highly subjective. Results suggested a strong correlation between narrative potential and quality. We did not obtain, however, the target metrics. The technological limitations found highlight the need for breakthroughs or distinct approaches to this problem. Still, we consider that a similar approach would be viable for generating textual artefacts perceived as consistent or coherent. This is based on the assumption that these magnitudes rely on repetition of patterns.

Chapter 4 presents an analysis of sites run by communities dedi-

cated to create fictional story worlds and stories that take place inside them. This chapter seeks to generalize findings from the previous research from the lab environment to the current online trends. Also, to formalize them in a model that helps in predicting and analysing any instance of collaborative narrative from the perspective of a shared story world.

The first paper in the chapter [45] describes a longitudinal study of students attempting to collaborate in a shared story world using an online tool specifically crafted to reinforce consistency and an in-depth analysis of existing large-scale multi-authored story worlds and their treatment of consistency and canonicity. Our approach is innovative in its focus: author dynamics and their impact on the perceived consistency and coherence. The resulting contributions were evaluated in a process aimed at obtaining qualitative and quantitative measures of content and author dynamics explained in detail. The results suggest that contributing authors are strongly influenced by material that was present before the experiment and favour one single canonical interpretation of events. Also, semantic links between contributions and the timeline visualization saw a positive reception. A last relevant observation was that in this context, users were not bothered by contradictions. Most users read little of the contributions introduced by other users (unlike in the previous studies), and used plot elements nobody else introduced in their stories. We suspect that some users might ignore contradictions because they are not knowledgeable on the contents of the story world.

In the second article of the chapter [46], popular, shared fictional story worlds from the internet are explored focusing on one of the newest, most popular ones-the SCP Foundation-to analyse its content and the author dynamics. The main findings suggest that the use of standardized formats and strong referential mechanisms, similarly to digital encyclopedias, encourage author interaction and collaboration. Other findings include the innovative collaboration dynamics represented by shared groups of interest and sub-canons, meant for authors to share common plot elements without intruding in each other's contributions. We also conclude that these sites follow a content and collaboration pattern we formalized, the Open Story World, which, amongst other aspects, has a balance of encyclopedic, narrative and navigational content, presents an exponential content distribution in terms of referencing and has a flexible canon model with dynamic coherence rules.

Chapter 2

Exploring the Collaborative Creation of Story Worlds

This chapter focuses on the specific circumstances of a multi-author story creation context, focusing in the role of consistency and coherence. The publication included in the chapter is the following:

Alan Tapscott, Joaquim Colás, Ayman Moghnieh, and Josep Blat. Writing Consistent Stories based on Structured Multi-Authored Narrative Spaces. In *Proceedings of the 4th Workshop on Computational Models of Narrative (CMN'13)*, volume 32, pages 277–292, 2013

The main goal is to study the creation of the *collaborative story world* we previously framed in a way that allows us to isolate significant factors of the process. We imitate the dynamics of large-scale multi-author franchises-using recurring characters and multiple, intertwined plots-in a controlled environment. By focusing on the perception of authors and readers, we provide an innovative view that to our best knowledge is rarely covered in the related literature. Despite some studies focusing on the big picture of computer-based collaborative writing [52, 23, 8] or on small-scale collaborative writing experiments [25, 37, 38], we found no study that focused in the story world and its consistency and coherence. We predicted that having multiple authors could lead to contradictions, and we were specially interested in observing how they were managed by authors with no specific instructions on the matter.

The study includes an observation of a participative process to create a shared story world (including subject questionnaires and

feedback) with prototype tools specifically suited for collaborative authoring of consistent and coherent fiction. Despite the extensive use of wikis and similar platforms for these kind of scenarios, we deliberately developed our own paper mock-up and desktop computer prototypes to avoid the authors' preconceived ideas influencing the collaborative writing process. We instructed a small group of volunteers to create story scenes using a closed set of paper elements (characters and objects) and to link them to an ongoing non-linear story, hanged, hanging them in a large glass wall. They could also draw lines between scenes of rearrange them as they saw fit. We made the paper elements large to improve our observations of subject attention and element manipulation. Their questionnaire answers provided a better picture of their motivations. The small scope of the experiment allowed us to interview and observe in person every participant. The results encouraged us to experiment with a computer prototype that attempted to recreate similar experimental conditions, mainly relying on direct mouse interaction with the visual representations of the story. The software prototype featured the same plot elements and allowed similar interactions. It also included an underlying formal model based on the results of previous experiments to enhance the consistent use of the fictional space. This model was based on a discretization of the fictional space, limiting the character movements between scenes. The computer prototype also included explicit mechanisms meant for authors to link their contributions around a common theme or character, called *storylines*, based on the classic narratological concepts of *actant* [34] or *narrative program* [17]. We then conducted two similar experiments with the computer prototype also involving small groups of volunteers.

Our observations point towards a predilection for sequential visual metaphors to organize the contributions and the perception of a single continuity. In-situ visual observations and questionnaires suggested that consistency and coherence were the main motivations for direct collaboration. We interpret this observation as an argument in favor of the hypothetic collective construction of a single story world. A single-consistent and coherent-continuity or interpretation of the fictional facts. Authors valued the inclusion of the consistency model and its constraints negatively, but readers perceived the results of its inclusion as more consistent. The model was more popular amongst authors when used to provide consistency recommendations instead of enforcing its rules. We saw a popular reception for the storylines amongst readers and authors. We also observed a certain author territoriality and hesitation to modify other author's contribution for

creative or artistic purposes. This effect was not observed when the contributor was rearranging content or modifying it to correct what he perceived as plot holes (conceptually close to consistency and coherence). The perceived consistency decayed over time as more authors participated, implying that the scalability of this contribution model might be troublesome with large author crowds. These results represent significant factors for collaborative story worlds, and they are explored more in depth in the following chapters. Also, [10] discusses similar experiments using the same prototypes, elaborating on some of the implications for HCI and shared narrative spaces.

Writing Consistent Stories based on Structured Multi-Authored Narrative Spaces

Alan Tapscott¹, Joaquim Colás¹, Ayman Moghnieh¹, and Josep Blat¹

¹*Grup de Tecnologies Interactives, Universitat Pompeu Fabra*

Abstract

Multi-authoring is currently a common practice in the field of contemporary storytelling, but producing consistent stories that share a common narrative space when multiple authors are involved is not a trivial task. Inconsistencies, which are not always well-received by readers are sometimes expensive to fix. In this work we attempt to improve the consistency of stories and narrative spaces by introducing a set of rules based on a formal model. The model takes into account the reader's concept of consistency in storytelling, and acts as a framework for building tools to construct stories grounded in a common narrative space with a reinforced sense of consistency. We define a model (the Setting) and deploy it through a tool (CrossTale); both based on previous research, and discuss some user evaluation, with an in-depth analysis of the results and their implications.

1 Introduction

The evolution of digital interactive media and information technologies has been instrumental in the development of systems that bring together authors and readers to compose and consume multi-authored stories through multiple media. In this context, the audience is not only interested in rich narratives, but also wants to participate in their development by adding and sharing their very own creations, compositions, and ideas. Nowadays people actively publish and share thousands of creative works (blogs, stories, songs...) on the web, often related to other original creations through relations that range from mere inspiration to direct referencing. Some of the works may be further developed by more authors, who expand their content, structure, and knowledge value through original creation and composition processes. On the other hand, there is an emerging interest

to support collaborative creation, composition, and consumption of multi-authored narratives that may grow in a shared information space for prosumers and professionals alike. We use a basic definition of information space: "The set of concepts and relations amongst them held by an information system" [13]. We believe narrative spaces are information spaces that ground all media based on the same characters, situations, plots or other casually interlinked entities, hence introducing a certain degree of consistency to the set as a whole. Narrative spaces are especially worth analyzing when dealing with collaborative storytelling since they establish many of the rules for the interaction among authors. The authors' awareness and interpretation of the narrative space will heavily condition their interaction with it. Fans often expand the narrative space of their favorite entertainment franchises by introducing their own stories deeply rooted in a well-established narrative universe and its mythology, creating rich networks of fanfiction (Fanfiction.net, referenced later, gathers hundreds of thousands of users around hundreds of franchises) that coexist with the official material. Also, Web and information technologies provide momentum to complex entertainment franchises created by dozens of authors to span across multiple media. In this context, there are certain tools that support readers and writers who are contributing to well-established narrative spaces. articy:draft [1] and Celtx [4] are both collaborative tools meant for creative story development and represent good examples of this emerging trend. Such tools may be created by the same company delivering the content, but this content is often the result of a collaborative effort undertaken by an author community. It is worth noting that tools do not merely intend to support the construction of a story in the sequential, traditional way. They provide mechanisms that allow for free, divergent exploration of all the related information, supporting the non-linear growth of narrative spaces.

There are some examples of narrative spaces worth mentioning. the FanFiction web site [7] is devoted to fan-developed stories within the narrative space defined by specific franchises, and provides a good example of amateur and professional authors creating stories in the same narrative space. Most of these stories, however, do not take into account the contributions of their fellow fan authors, only the original, canonical one. Another example is the website that holds most the information related to the *A Song of Ice and Fire* book series in its many articles [18]. It is fed with content from multiple contributors, properly structured and published in a readable way. The site also publishes articles that cover most of the books related to the canonical narrative space and a text-based roleplaying game

that allows players to introduce their own creations (e.g., characters, locations, and other elements around the original canonical narrative). Players can interact with each other while expanding the original setting. This site has the approval of the author of *A Song of Ice and Fire* who is known to be vocal against common fan fiction developed without consent. On the other hand, he created Wild Cards [11], a book series written by multiple authors under his editorial control. Chris Crawford's Storytron [6] is an interesting approach to developing a commercial tool for users that design interactive stories. Although it is currently on-hold due to problems regarding the learning curve (i.e. the complexity of building a whole interactive story with the tool), this approach is interesting in terms of decomposing the narrative space into a set of unitary elements, and defining the logic that relates them. Storyjacker [9] is another interesting example closely related to the tradition of the Exquisite Corpse writing game. This game proposes that its players first read a *flash fiction* (roughly between two or three hundred words) created by another writer with an explicit editorial challenge attached to it. Players rewrite the text answering the challenge and pass the result to the next player, introducing a new challenge of their own. While this approach is a game, the writing dynamics of its multi-author design are interesting and not very far away from what we propose in this paper.

In this paper we explore whether people are concerned with consistency when writing stories collaboratively, providing them with a tool meant for that purpose. First we discuss our focus on enhancing consistency, especially how it is perceived by authors and readers, followed by a brief state of the art of previous research on multi-authored narratives for similar scenarios. Next we describe some users' experiments we conducted. These experiments were designed to test mechanisms developed to increase narrative consistency. We then analyze and discuss the resulting experimental data. Finally we discuss these findings in relation to the approach proposed and introduce future research.

2 Supporting Narrative consistency

Complexity can easily scale with the developing size of narrative spaces, possibly increasing the difficulty of reading and authoring stories that reference each other. Each element in a narrative space, such as a character, location or event, is linked to other elements in the same narrative space through causal relations, providing a sense of continuity and consistency. Modifications introduced to the narrative space may cause contradictions in the logic of the network of elements and causal

links. This often leads to plot holes that may compromise the story's global consistency, potentially hindering the experience of authors and readers. Stories containing plot holes tend to have a bad reception amongst sophisticated readers [16].

If consistency is a key factor when dealing with multi-authored storytelling, some sort of mechanism designed to monitor and enhance its presence could result in a better experience for its readers. This work pursues a suitable method to assist multiple authors in developing narrative spaces with enhanced consistency. This might lead to stories which are more satisfactory to develop collaboratively and are also more enjoyable to read. When analyzing narrative spaces and their unfolding stories, we distinguish between two kinds of consistency measurements: - Firstly *structural as the recurring usage of recurring elements and patterns from the narrative space and their level of agreement*. This can be measured if the narrative space is mapped to a computational structure of some sort, validating the narrative space information against a formal model. - Secondly *reader-perceived as the level of consistency associated by readers to a specific story*. This can be obtained by asking readers to rate it after having read it. We think this distinction is necessary because of the subjective nature of some stories along with the existence of some literary techniques, such as the use of biased narrators that describe reality through perception and language. We use these two distinct conceptions of consistency to base our experiments. By analyzing the content of a narrative space and mapping it to a computable and evaluable structure, we can provide some recommendations or guidelines to increase the structural consistency of a narrative space. Starting with Propp and his structural approach to narrative [15], the field of semiotics is grounded on similar principles and has been an active discipline for decades. Its theoretical foundation, specifically the syntactic branch that deals with formal structures, has been a source of inspiration for our work. Deconstructing a narrative space into a computational structure based on a suitable model can be a challenging discretization process. Instead, the model we propose is based on observations regarding the author and reader perception and interpretation of consistency. Every author has a personal way to tell stories. This means that the perception of a story's consistency depends on the technique and structure of its discourse -not to mention the influence of genre-. Readers may find a story consistent or inconsistent regardless of the raw material from the narrative space used by the author. Also every reader's perception is heavily influenced by factors such as his/her cultural, academic and social background, which can be difficult to control and keep track

of. The user-perceived consistency of a specific story is measured simply by asking its readers to rate it. There are other more indirect methods, such as asking specific questions to check the comprehension of the story or to observe the reading process, trying to encode it into meaningful data. We have found these measurements difficult to operationalize and correlate to the reader-perceived consistency level. Our goal in this research is to determine whether monitoring and enhancing the structural consistency of a narrative space implies that stories based on it are perceived as more consistent by readers.

3 Related Works

Meehan's TaleSpin is a system that generates stories via carefully crafted processes that operate at a fine level on story data [12]. It was one of the first attempts to model narratives as computational systems. Since it automatically generates stories, it holds a certain notion of computational causality and consistency. We also pursue a formal model with such notions, but Meehan's approach seems too constraining to support an open definition of a story. Brenda Laurel's doctoral dissertation described a complex framework for drama management [10] and might have inspired similar approaches that deal with structured narrative spaces. While it is meant for abstract depictions of large narrative spaces, it also provides a systematic representation for them. A key factor is its ability to introduce highly dynamic narrative structures. These structures support complex stories that hide the formal complexity from readers, something we wish to introduce in our approach. Thue [17] proposes an interesting approach that formally structures the story, favoring consistency monitoring and analysis. Player Modeling is a simple concept that attempts to personalize the story through several profiling techniques, enabling some of the user's personality traits to have certain impact on the resulting experience. Understanding the reader's perception of consistency is a concern we share. Some other approaches use a strictly formal definition to model stories. For instance, Cavazza proposed a character-based approach [3] that was adapted and improved by Pizzi to model a part of Madame Bovary [14]. This line of work is grounded on planning and the field of artificial intelligence. Interestingly enough, it deals with complex aspects of human nature such as emotions and feelings. The AI planning used in [14] is concerned with optimality, seeking to reach a target with economic operations and may not be adequate for our approach. This are intricate and complete attempts to discretize the narrative structure into a formal model, a goal we also

pursue. Next we discuss some existing formats and recent tools that allow modeling narrative entities independently from their story, that keep track of the flow of complex events, that impose constraints or rules to preserve consistency, that keep track of plot meta-data (such as character motivations, feelings or the literary theme and mood), and that are suited for collaborative development of a story. This discussion inspired the conception of our tool.

- Traditional scripts are often created by a single or a couple of authors. Large media franchises and episodic shows sometimes need to become heavily interrelated. Modern TV series seasons are also a good example, featuring interrelated scripts written by multiple authors. To some extent they represent one of the most popular instances of a multi-authored narrative with a strong need for consistency.
- There is a certain tradition of background books for rich fiction franchises, providing concept art, character profiles or even maps depicting fictional lands. These books, far from narrating a story in the traditional sense, describe a specific part of a fictional universe. We found these works interesting because they represent a set of characters, themes and plots in their original, protean form, not necessarily attached to the linear context of a traditional tale. They are often written by authors who were not creators of the original concepts, and represent an example of collaborative authoring.
- RPG books, such as the *ADD Monster Manual* [8] or *Vampire: The Requiem Coteries* [2] are interesting examples of narrative entities modeled independently. They provide a growing organic framework for authors to build their own adventures and share them with friends, adopting the role of a live storyteller in tabletop gaming sessions. The source material in these books can be used to enrich the session experience by introducing new characters, object or plot threads. While fairly similar to background books, RPG books provide guidelines that allow content to be used in the arbitrary context of a game with rules, which introduces a high degree of formality to the information.
- Certain tools such as Wikipedia or Wikia are effective means for storing and organizing data from a specific narrative space. Although they are commonly used to structure already-existing

background information, they represent some of the most popular tools that support collaborative writing. Their capability to deal with individual entities such as characters or locations is the trait we find more interesting. On the other hand, entities commonly depicted as linear, such as stories or plot threads, are not very intuitive to understand and follow using these tools. As shown in table 1, most of them possess some of the traits we introduced earlier, but no tool has got all of them as far as we know. Incorporating existing mechanisms that seem appropriate is part of our efforts to design a tool with all these traits.

Feature	Traditional Script	Background Book	P&P RPG Source Book	Wikipedia/Wikia
Atomic Narrative Elements	Not formally.	Yes, clearly differentiated.	Yes, clearly differentiated.	Yes, clearly differentiated.
History Log	Sequence is implicit on its description.	Yes, mostly inside the individual element descriptions.	Yes, mostly inside the individual element descriptions.	Yes, mostly inside the individual element descriptions.
Consistence Constraints	No	No	Yes, enforced by the game's rules.	No
Plot Meta-Data	No	Yes, mostly inside the individual element descriptions.	Yes, abstract data and plot hooks are provided.	Yes, only attached to individual element descriptions.
Suited for Collaborative Development	No	No	No	Yes

Fig. 1: Multi-authoring narrative supports comparison

4 Experiments

We carried out three experiments to understand better narrative spaces and the stories based on them in terms of user perception. For each of them we introduce its purpose, the experimental tool specifically designed for it, the experimental design and the most significant results.

4.1 Experiment I - Understanding the Sharing of Narrative Spaces

Our first experiment aimed at understanding how users perceive a narrative space and its associated stories while contributing and navigating through it. We intended to understand their mental model and to measure it. Some arbitrary conventions were introduced, such as an initial set of scenes already connected or a limited set of characters and



Fig. 2: A Story Wall

objects. This was done to encourage participation, providing a certain sense of narrative immersion and to reduce the effort required from subjects in order to participate. A fairy tale was chosen, including its most typical elements (e.g., a king, a princess, a castle, and a dragon among others) along with atypical ones (e.g., a robot, aliens, and a starship among others).

4.2 Experiment I: A Collaborative Story Wall

A large glass wall was used as a space to develop and visualize a collaborative narrative (figure 2) composed of scenes and transitions. Our purpose was to provide a canvas for authors to freely interact with the story. The scenes were sheets of paper with a collage of images (obtained by mixing characters and props picked from a set) and text written to describe the scene more explicitly. Scenes could be added anywhere on the wall and connected by transition arrows drawn on the glass, as a directional indicator, providing a sequential order by connecting them. We provided an initial story as a starting point for users who, in succession, could modify what was on the wall: change or delete scenes, alter the structure (erasing and drawing transition arrows, and moving scenes to new positions), and place their own scenes in any point of the unfolding story. We introduced 7 initial scenes narrating the beginning of the kidnapping and rescue of the Princess.

4.3 Experiment I in Detail

16 subjects were invited to participate in the experiment one after another sequentially. There was no special consideration in the demographics involved. A non-imposed average elapsed time of 12 minutes was measured.

Subjects were asked to read the existing narrative which was the

result of the accumulative modifications made by previous subjects. They were also interviewed after they finished reading the existing narrative on how they had chosen to read the story (order, objects and concepts they had followed, etc.), along with their opinion on some specific matters such as the literary value and consistency perceived.

Next they were offered the possibility to contribute to the narrative, and allowed to modify or delete previous scenes, to alter the structure of the story structuring (erasing and drawing transition arrows, and moving scenes to new positions), and to place their own scenes at any desired point. Finally all subjects answered a series of questions designed to learn more on how they interacted with the story, such as the nature of their contributions (according to them) along with their driving motivation or purpose. We also asked some open questions on some subjects such as if it was a fun experience or if they would enjoy doing the same with their friends through a social network.

4.4 Experiment I: Results

The story resulting from the experiment contained 29 scenes connected through two main branches that converged towards their end. Each participant added either one or two scenes to the growing narrative. No subject eliminated scenes from previous participants, but modifications on existing scenes were common: half of the participants inserted their scenes between existing ones and/or altered the direction of arrows; over one third created convergence between two or more isolated branches (for example two characters gathering at one point, or one event affecting the story of another author). A few subjects claimed to focus exclusively on solving inconsistencies during the authoring phase of the experiment. Maintaining consistency in the evolving narrative was stated as the principle reason for 8 out of the 15 contributors. The notion of conflicting scenes was stated 4 times as something disliked in the interviews. According to subjects all of the changes made to previously existing elements were for the sake of consistency. Other contributions were centered mainly on extending existing plot arcs instead of creating new ones. Consistency seemed to be key in user motivation and overall experience. The subject-perceived level of narrative consistency (figure 3) tends to be on the middle-high portion of the scale but decays slowly. As the initial story is different for each user, the results cannot be easily compared but subsequent experiments allow for comparison.

According to the interviews, the literary value of the narrative was of little concern to the subjects. Interestingly, individual scenes and

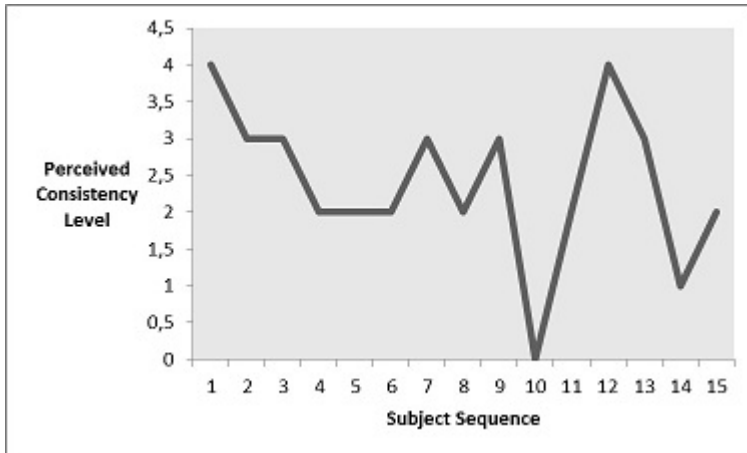


Fig. 3: Story on a Wall reader-perceived consistency level

small narrative branches had greater entertainment value than the overall narrative. Since the sequence of events can only be guessed through the spatial layout of the scene and the arrows network, some conflicting notions appeared on what was happening before, after, or simultaneously to a given scene when dealing with parallel stories. This suggested that scenes could be arranged in some sort of linear organizational structure to provide an improved sense of sequence and causality. Our close observation of how scenes related with each other and how participants authored existing characters, revealed that each character was considered the same entity throughout the whole narrative, almost always labeled with the same name. The experiment also showed that the authors faced a complexity which scaled if they tried to maintain the structural consistency of the story. The more scenes it contained, the harder it was to introduce new material without contradicting or violating existing established facts. On the other hand, the decreasing reader-perceived consistency of stories containing a large amount of scenes indicated that the reading process became more difficult as well. Some people were motivated by the unfolding implicit collaboration, and nobody stated openly to be bothered by it. In fact, contributing to the narrative was not mandatory but all of the subjects added scenes, and they actively searched for an interesting entry point and modified the whole context, changing and rearranging scenes connected to their contributions, instead of just attaching them to the end of a story thread. More than half of the subjects expressed their interest in repeating the process later and many of them returned after their contribution to see how the

narrative was evolving. A good number of people who just happened to pass by stopped to read the whole story, many of whom asked to participate in subsequent iterations of the experiment.

4.5 Experiment II - Measuring the Impact of Consistency Constraints

The purpose of the second experiment was to measure the impact of an underlying formal model to user contributions and their overall interaction with a multi-authored non-linear narrative. This formal model was designed to provide structural consistency to the narrative space, hopefully reinforcing key factors that enhance the production of stories that are perceived as more consistent. We introduced some constraints into the interaction to prevent subjects from creating scenes that somehow violated the rules proposed by the underlying model. We used a platform we developed, [5], to be used on a laptop individually in an isolated lab, which meant changing to a much more private environment.

4.6 Experiment II: A Setting that provides an Underlying Consistency Model

Our proposed model, the Setting, tries to provide an underlying formal model that resembles the author's mental construction of a narrative space. We used data from the previous experiment to map their understanding of the story into an assessable and measurable model. The Setting serves to monitor and enhance the consistency of the stories unfolding within it. It provides a common ground for authors to interact by building stories in the same narrative space. Its informal definition is the following:

- The Setting contains *timeframes*, *locations*, *scenes*, *character storylines* and *plot storylines*.
- Timeframes have a sequential order.
- Every location is connected to other locations. The distance from location A to B is the minimum number of locations needed to go from location A to B.
- Every scene takes place in a location and timeframe, containing one or more characters.

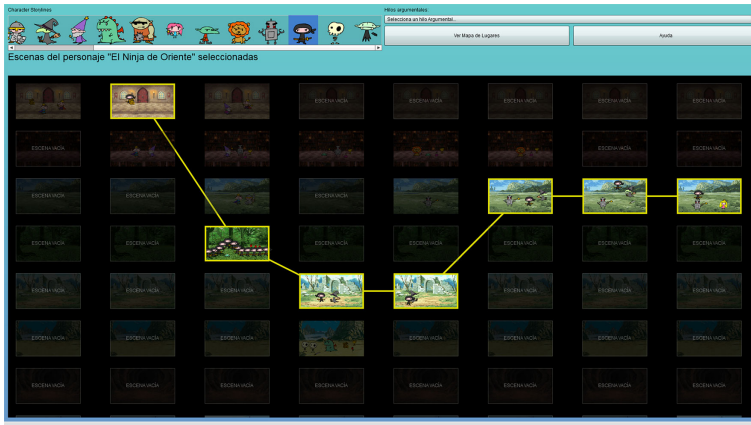


Fig. 4: CrossTale interface

- Scenes can belong to plot storylines or character storylines.
- Storylines contain one or many scenes.
- Character storylines contain all the scenes that contain a specific character.
- Plot storylines contain all the scenes tagged with a specific plot.
- Characters may only appear once per timeframe in a scene.
- Characters may only appear once in the same scene.
- Characters may only move once between directly connected locations during consecutive timeframes.

These rules were designed to provide a certain sense of consistency, which can be measured, monitored and enhanced, on the basis of the results of the Story on a Wall experiment, attempting to predict and enforce the factors actively pursued by users through their contribution. Our goal was not to evaluate this definition as a generalist model capable of describing any narrative; instead we wished to measure the impact of using a formal model in a multi-authoring scenario in terms of the consistency of the resulting stories.

4.7 Experiment II: Introducing the platform

CrossTale (see figure 4) is a software prototype featuring a visualization that follows a distribution similar to that of Story Wall, adding the

rules imposed by the Setting into the authoring process. In fact, its main context is a dashboard with two axes, one for time and one for place. Users can scroll at will to navigate the dashboard. By selecting existing scenes they can view their images and read the descriptive texts. Specific characters and storylines can be selected, enabling users to read all the scenes involving that character or storyline in a sequential order. The grid also highlights scenes belonging to the selected entity and connects them with an arrow line to reflect their sequential order. There is also a secondary context that enables users to create scenes, providing a set of components (characters, objects and plot storyline tags) along with a visual representation of the location where the scene takes place and a text box to introduce the description. These scenes are added to one of the Setting timeframes and locations and are treated as an integral part of the narrative space. Violations of the Setting were not allowed in this experiment, and the user got a message requesting him/her to resolve the conflict before saving the scene.

4.8 Experiment II in Detail

20 subjects of similar characteristics as those in the first experiment took part. Two groups of 10 were created randomly. The control group used the tool to read and contribute to the existing narrative, and the experimental group had consistency constraints based on the Setting. The order of contribution was sequential, as each user found the story in the situation left by the previous one. No time limits were provided and the average time of the users was 20 minutes. A CrossTale prototype was deployed with an initial set of scenes describing the start of a fairy tale. The 8 initial scenes introduced were almost identical to the ones used in the previous experiment, introducing a Princess, her kidnapping by a witch and the Prince trying to rescue her. Each subject was asked to read the story which was composed by the initial scenes provided plus the contributions made by previous subjects. No specific method was imposed. The completeness of the reading process was optional. Reading character storylines could be a strategy amongst others. A brief interview was conducted to understand how users read and understood the whole scene set regarding the storylines. Then they were invited to add one or more scenes to the existing ones. After they were done, a second interview was conducted to understand what kind of additions and modifications they had made, their motivations, the intended influence on the previous state of the story, and any other relevant details of the interaction between the subjects and the




Fig. 5: CrossTale reader-perceived consistency level with and without consistency constraints

story. The whole experiment was recorded for further coding and observations. Subjects were aware of the collaborative nature of the tool, but did not have physical contact with the rest of the subjects before, during or after the experiment.

4.9 Experiment II: Results

Results were analyzed independently for each group. It appears that subjects were not very concerned with reading the whole narrative before interacting with it. Users only read a fraction of the existing content. No user read the whole story. The most common interaction recorded during the reading phase involved the user selecting one or two storylines and reading its content before moving on to the contribution phase. The perceived consistency (figure 5) was rated high in both groups, with a slight tendency to decay towards the end in the group without constraints. The difference did not seem very significant. Both groups ended up with a story composed of 28 scenes and 10 storylines. The average scene contribution was 2 scenes per user. Most users placed their scenes inside one and only one storyline. No user modified scenes created by other authors. The rating of the user experience was positive (average 4.4 out of 5) as well as of the application design (average 4 out of 5). We asked subjects if they would use CrossTale regularly with an average 3.6 out of 5 and if they would like to have a similar tool to create and share narratives in the context of a social network, with an average 3.7 out of 5.

The focus of the experiment was to observe if the introduction of consistency constraints derived from the Setting caused any interesting effects. The most remarkable observation was that the perceived consistency seemed to decay more quickly over each contribution for the group without constraints, although the resulting data isn't very significant. This could mean that enforcing certain notions of time and space through the scenes tends to produce more consistent results, supporting our initial hypothesis. A larger subject group in future experiments could validate or refuse this claim. Adding the constraints seems to have an annoying effect on the experience of users who felt limited all the time (as seen during the video codification, where they complained almost every time a constraint blocking message popped up). This might be caused by the way messages themselves are displayed in CrossTale. It could be an interesting line for future research. Joining the data from both groups also revealed some interesting facts. The use of a computer program to conduct the experiment might have affected the user experience, limiting the user's freedom when compared to the previous experiment. The story in this experiment was read on a screen and embedded inside a software program instead of being on a glass wall. Subjects were less inclined to interact with the existing scenes; no user modified scenes created by other authors. Subjects spent less time interacting with the narrative (the decreased time could either be an indicator of a less pronounced learning curve, a good interaction design or a decrease in the motivation of subjects). Also, according to the interviews, they were less concerned by narrative inconsistencies. As previously mentioned, the story was now stored in a computer program. We believe this might have caused users to be less aware of the story as a whole and therefore less concerned with its global consistency. In fact, the reader-perceived consistency of the narrative was larger for both groups of users compared to the previous experiment. This might also be related to the fact that users never read the whole story. Users aren't concerned with the consistency of scenes they haven't read. We chose to follow a cumulative contribution sequence as in the first experiment on both groups. This was done to gain some insight on the evolution and scalability of the story while comparing the results with the previous experience. We are aware that this decision prevents us from comparing subjects' individual performance in terms of consistency.

4.10 Experiment III - Measuring the Usage of Storylines

The third experiment explored the use of storylines further. Namely, we were interested in measuring certain aspects such as the number of storylines read by subjects, the degree of comprehension after reading, the performance when creating new storylines and their consistency. Moreover, we wanted to cross measures between reading and contributing phases and find any significant correlations. Also, we reduced the Setting's restrictions, switching from usability constraints to recommendations.

4.11 Experiment III in Detail

This experiment was fairly similar to the previous one. The main difference was that user contributions were not cumulative, every subject found the same initial set of scenes and there was only one group. Every subject started their contribution with the initial 12 scenes we provided. The story was the same fairy tale. The initial scenes introduced 3 main storylines that explained the events through the prince, princess and the witch's own viewpoints. CrossTale was used with the same rules derived from the Setting, the derived consistency constraints from the Setting were always active; its application was not enforced, only warning messages existed. There were minor usability refinements to CrossTale. We provided users with the ability to zoom in and out (using the mouse wheel) when viewing the scene grid. We also allowed users to scroll through the scene grid by dragging the mouse anywhere, not only the scrollbars. These additions were introduced to provide more visibility and accessibility to the existing scenes inside CrossTale. 16 subjects of similar characteristics as those involved in the previous experiments took part. An average time of 10 minutes of involvement with the system was measured. The experiment began with each subject reading the story. CrossTale provided several mechanisms to do it: reading individual scenes, following specific storylines according to plot threads or characters. Users were free to read only a part if they wished. The interactions with the reading interface were registered, and a brief interview was conducted afterwards to analyze their reading experience. The next phase was the contribution. Every subject was asked to add more scenes to the same existing story if they wanted to. Their interaction was registered and a brief questionnaire was administered. This questionnaire was used to rate the users general impression of the story when contributing to it. Subjects were asked to rate the warning messages, the story in

terms of consistency and amusement through Likert scales, and also to propose one or more titles. In both phases the proceedings were run by a collaborator not directly involved in the research, who coded the interactions as well. Unlike the previous experiments, modifications to the scene set were not cumulative between subjects, so the consistency measurement was done through a 4 person jury evaluation of each subjects' contribution.

4.12 Experiment III: Results

Regarding the reading phase, most subjects partially read the existing scenes through the usage of storylines. Readers selected an average of 7.77 storylines to read. 83% of them were read from start to finish. 43.59% of the initial character storylines were read and 1.38% of the initial plot storylines were read. The average contribution per subject was 2.6 scenes. The number of scenes read seems to be correlated with the number of plot storylines used. There is a medium-high correlation between the number of titles for the story proposed by subjects and the number of characters mentioned in those titles. Also there's another medium-high correlation between the number of plot storylines referenced in the proposed titles and the amount of plot tags used later during the authoring phase. There's a positive correlation between the number scenes created, the number of storylines read and diversity of characters used in the created scenes. A few message warnings informing of violations of the Setting rules were displayed (Warnings appeared in 24% of the composed scenes). Of these warnings, only 17% made the authors change the story. According to the jury, the resulting inconsistency level measured was an average of 1 inconsistency per subject overall contribution, or 0.46 inconsistencies per scene. The number of inconsistencies remained stable during each user's session. In those cases, having the same author for all the contributions also ensured a more accessible and scalable development. The small size of the initial narrative, along with the improvements and refinements to the CrossTale user experience were also instrumental for this to happen. This also could explain certain measurements, such as the average reduced time for each subject's interaction with the story. Subjects seemed generally more inclined to use character storylines to read the provided story. There's a tendency towards a character-driven exploration of the story. Nearly no subject read scenes without using storylines. We believe they proved to be a good mechanism to explore non-linear narratives such as the one we created in this experiment. Some users made extensive usage of the tool to create a large amount

of scenes, which allowed us to briefly analyze the scalability of the system in terms of consistency.

While these measurements might make it difficult to correlate the structural consistency of the narrative space with the consistency perceived by reader, the jury evaluation and the qualitative analysis of the stories suggest some major critical inconsistencies were avoided thanks to the Setting's warnings. Since we lack more evidence to sustain such a claim, we are already pursuing new experiments to provide more data in this direction. It is worth noting all elements tagged as incoherent by the Setting's rules were not considered very incoherent by the jury evaluating the consistency of subject's resulting narrative.

5 Discussion and Conclusions

In this section relevant issues that emerged from the three experiments are discussed together with considering other interesting points for the near future research.

5.1 The Role of Consistency

Consistency appeared as a relevant factor during collaborative narrative composition, and it influences on the way stories are read and written in multi-authored scenarios. Let us recall that in the first experiment, authors introduced quite a few modifications to the overall story when it was necessary to maintain the consistency of the plot arc they were developing or to correct a discontinuity in the overall narrative consistency. Consistency provides stories with a sense of causality and makes them more accessible for new authors and enjoyable for readers.

We believe there is a certain cultural common knowledge of what is consistent and what represents a plot hole, defined by Ryan [16] as an inadvertent inconsistency in the logical and motivational texture of a story. In our model, a plot hole is a discontinuity in the cause-effect logic of the story discourse. Further experiments are needed to validate this hypothesis of the relevance of causal links.

However, in the second and third experiments authors were not as clearly concerned by consistency as in the first experiment. We believe this is due to the experimental settings, as the use of a more focused and constraining software prototype meant incoherences were less visible to the users. The introduction of an underlying formal model with its own rules, and of reading mechanisms, which were absent on the first experiment, probably led to the reduced interest in

providing consistency. CrossTale ensured consistency preservation in an effective way, and reduced the users' concerns.

However, consistency is not the only issue worth tracking when building stories collaboratively: the lack of visibility of scenes or the constraining effect of the model on creativity were not our focus in the experiments and should be further studied.

The distinction between the two types of consistency has been an effective way to formulate our research. The Setting provided an objective measure of consistency based on our model, and its impact in the perceived consistency level could be assessed.

5.2 Monitoring and Enhancing Consistency through the Setting

The Setting aimed at dealing with the user's concerns about consistency observed during the first experiment. These concerns seemed to mean that time and space limitations had to be enforced, and therefore, the Setting only deals with these aspects of stories. It established a framework for developing narrative collaboratively, with a simplistic definition of what is consistent and what is not. Forcing users to follow the Setting rules during the scene composition process was not a very popular design decision among authors, but the stories built under these conditions apparently provide better reading experiences. Therefore we illustrate an interesting trade-off; constraining scene composition under a Setting-like model may lead to more consistent results while hampering the authoring process. No specific observations were made on creativity aspects, but we feel that the Setting could easily decrease the creativity of the stories it supports. This should be properly tested in subsequent experiments.

The Setting in the second experiment proved to be a double-bladed sword: authors were aware of some of the things they needed to take into account that might have ignored so far, but they also felt less able to express their creativity due to the constraining nature of the consistency rules. The implementation of the Setting in the third experiment was more restrictive and generally ignored; authors were always aware of violations to the Setting rules, but they could react in different ways. Some deliberately ignored the warnings, while others (a minority) prioritized such violations and removed them. Ultimately, we believe there is no formal model valid and complete for all possible narratives. Our future attempts to provide support and guidance in building consistent multi-authored stories will probably involve the authors in the construction of their formal model. What might be

consistent in one narrative space, such as involving magic characters, might be inconsistent in others, and there is no one better suited to establish these discriminations than the individuals who are creating the stories. Future experiments could even introduce inconsistency generators, based on approaches that generate events and situations, possibly reducing the user-perceived consistency but maybe providing some inspiration to the authors.

It is important to remark that the results coming from experiments where the modifications to the narrative persist and those where every subject deals with the same exact set of scenes are not directly comparable.

Group dynamics are another relevant aspect of multi-authoring scenarios. The Setting essentially stated the game rules, which each author had to follow to enter into the game of story creation. On the other hand, each author introduced modifications to the narrative space that needed to be respected by subsequent authors, meaning that the learning time needed by the following author increased. A possible improvement could be to provide better communication amongst authors to support their coordination. This could improve cooperation during narrative composition and introduce specializations such as committing specific authors to preserve consistency by stating the fundamental consistency rules and reorganizing structured content.

5.3 Very Human and Causal Storylines

Human-generated stories within a narrative space, as those observed in the first experiment, are not random. Most contributions followed existing plots, commonly associated with a character or some abstract concept, such as a motivation or a specific theme. The introduction of explicit, storylines in the second and third experiments was meant to reinforce the sense of causality and continuity, trying to predict the authors' behavior to ultimately enhance the user experience. After analyzing their use during the experiments, it is safe to say that they meant a difference to the results. The reader has to follow the clear cause-to-effect relationship made explicit. The story exists in a specific region of the narrative space. Users embraced this storyline mechanism to explore and understand the narrative space, and in most cases avoided the free scene selection in favor of the sequential reading order provided. They also used this mechanism to link new scenes into existing storylines or even to start new storylines from scratch to propose new ways to read the content of their creations. This might have been one of the key reasons for the increase in the

reader-perceived consistency measured in the experiments that used CrossTale.

We believe the use of storylines as tools to communicate stories is fundamental in the exchange between a storyteller and its audience. From the Setting computational point of view, storylines are not necessary for the narrative space to exist. However, without them, the information has, arguably, a less narrative quality. Even if storylines did not formally exist in the Setting information architecture, any story introduced by human beings would probably have cause-to-effect relationships.

Another interesting finding that we will probably introduce in future attempts to map a story to a formal model is that readers prefer storylines based on characters to those based on plots, as they chose the former almost always. Apparently, in the context of a non-linear story, users find more natural to follow specific characters instead of plots. One possible explanation is that in most of our stories (and in many stories found on contemporary media) a character only appears in one plot with a main role. While s/he could appear (seldom) in additional storylines, the character would then have a minor role. Some of the most popular Semiotic models [15] are built around characters and their roles, rarely depicting meaningful entities that display human-like behavior. We will explore this approach in the future.

5.4 Conclusions and Other Future Work

Narratives are highly subjective, as any product of an artistic discipline. There is an implicit notion of causality in any story. Scenarios involving cooperation between authors might suffer from discontinuities in their causal relationships, which produce less satisfactory stories for their readers. We believe consistency plays a fundamental role and we presented experimental data that supports our belief. Our approach introducing a formal model that imposes consistency constraints derived from the narrative space was tested; showing it was capable to monitor and increase the structural consistency of the multi-authored narrative space as intended. This apparently translated into stories with an enhanced reader-perceived consistency. However, the negative reaction from authors when facing constraints imposed by the model requires further exploration. We believe some media (such as TV, films, comics amongst others) have the difficulties of collaboration amongst multiple authors discussed throughout the paper, and we plan to extend to them the methods introduced.

There are also some possible paths for future work that deal with

some secondary factors observed. Regarding creativity, subjects from all experiments seem to perceive scenes created by authors with a background in communication or arts as generally more creative but not necessarily more consistent. The relation between creativity and consistency is not clear at all in our observations. A more specific experimental design, possibly involving subjects with specific backgrounds and narrative expertise, could shed more light into the matter and maybe provide some details on the hypothetical correlation between creativity and consistency.

On the other hand little attention was paid to the interaction and aesthetic design of CrossTale. This is an interesting line of research that deals mainly with usability and user experience, potentially improving the CrossTale results.

References

- [1] Articy draft. www.nevigo.com/en/articydraft/overview. Accessed: May, 2013.
- [2] Kraig Blackwelder, Jacob Klunder, Matthew McFarland, and Will Hindmarch. *Vampire: The Requiem, Coteries*. White Wolf Pub Print, 2004.
- [3] Marc Cavazza, Fred Charles, and Steven J. Mead. Character-based interactive storytelling. In *IEEE Intelligent Systems*, volume 17, pages 17–24. University of Teesside, UK, 2002.
- [4] Celtx. www.celtx.com. Accessed: May, 2013.
- [5] Joaquim Colás, Alan Tapscott, Ayman Moghnieh, and Josep Blat. Crosstale: Shared narratives as a new interactive medium. *MMEDIA*, 2012.
- [6] Chris Crawford. Storytron. www.storytron.com. Accessed: May, 2013.
- [7] Fan fiction. www.fanfiction.net. Accessed: May, 2013.
- [8] Gary Gygax. *Advanced Dungeons and Dragons: Monster Manual*. Wizards of the Coast Print, 2012.
- [9] David Jackson. Storyjacker. www.storyjacker.net. Accessed: May, 2013.
- [10] Brenda Laurel. *Towards the Design of a Computer-Based Interactive Fantasy System*. PhD thesis, Ohio State University, Columbus, Ohio, 1986.
- [11] George R.R. Martin. *Wild Cards*. Gollancz Print, London, 2012.
- [12] James Meehan. TALE-SPIN, an interactive program that writes stories. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*. University of California, Irvine, 1977.
- [13] Gregory Newby. Metric multidimensional information space. *NIST Special Publications SP*, pages 521–536, 1997.
- [14] David Pizzi, Fred Charles, Jean-Luc Lugrin, and Marc Cavazza. Interactive storytelling with literary feelings. In *The second International Conference on Affective Computing and Intelligent Interaction (ACII2007)*, University of Teesside, UK, 2007.
- [15] Vladimir Propp. *The Morphology of the Folktale*. University of Texas Press, Austin, 1968.

-
- [16] Marie-Laure Ryan. Cheap plot trick, plot holes, and narrative design. *Narrative*, 17(1):56–75, 2009.
- [17] David Thue, Vadim Bulitko, Marcia Spetch, and Eric Wasylshen. Interactive storytelling: A player modelling approach. In *Proceedings of the 3rd Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE07)*, University of Alberta, Alberta, 2007.
- [18] Westeros. www.westeros.org. Accessed: May, 2013.

Chapter 3

Evaluating Visual Story Diagrams and Text Artefacts for Author Participation and Inspiration

This chapter focuses on supporting the collaborative construction of story worlds we observed in our initial research with innovative solutions. In order to improve author communication and coordination, we compare the performance of a traditional written story against its visual semantic representation. We also attempt to find new ways to stimulate author creativity by assessing a text generation system and the artefacts it produces in terms of narrative potential and quality. The two sections of the chapter are composed by the two following publications:

Alan Tapscott, Joaquim Colás, Ayman Moghnieh, Josep Blat, and Universitat Pompeu Fabra. Modifying Entity Relationship Models for Collaborative Fiction Planning and its Impact on Potential Authors. In *Proceedings of the 5th Workshop on Computational Models of Narrative (CMN'14)*, volume 41, pages 209–221. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, jan 2014

Alan Tapscott, Javier Gómez, Carlos León, J Smailovi, M Znidaršič, and Pablo Gervás. Empirical Evidence of the Limits of Automatic Assessment of Fictional Ideation. In *Proceedings of the 5th International Workshop on Computational Creativity, Concept Invention, and Gen-*

eral Intelligence, pages 58—71, Bozen-Bolzano, Italy, 2016

In the first section we experiment with the usage of semantic diagrams for coordinating and communicating stories between authors. Building up from the previous chapter, we hypothesize that visual, interactive elements are favorable for the collaborative planning of a story. Some studies back this hypothetical belief, stating the benefits for tangible interfaces in general [20], specifically for children [37] or as the base for conceptual frameworks [50]. Similarly to how [19] used the base E-R diagram to model English sentences, we propose to use a modified nomenclature (based on the original one proposed by Chen [9]) to structure and communicate a story. We also introduce a methodology to convert English sentences to their semantic diagram visual representation. Using a small pool of subjects, we compare their comprehension and memorability of a written Wikipedia movie synopsis, its semantic diagram equivalent, and both simultaneously. We also compare their performance when using the provided medium to expand the story. Results are generally favorable in terms of usability, suggesting that semantic diagrams are a functional metaphor for authors to read and extend a small, sample story. Also, the introduced methodology proved to be intuitive and popular amongst subjects. The combination of both mediums, that is, textual story plans and semantic story plans, in our limited experimental conditions, increased subject comprehension. The mapping we propose between a tool meant for conceptual communication (with a highly structured formulation typically used in software engineering) and a textual description of a story (such as the plot of a movie) could be helpful and perhaps replicable for similar scenarios.

In the second section of the chapter, we delve into the evaluation of computer-generated textual artefacts, seeking a set of metrics that correlate the perceived quality and narrative potential. Our approach exploits a mixture of human evaluation and machine learning to obtain a reliable measure of these parameters generally considered highly subjective (narrative potential and quality). This is a divisive matter [5, 6], but the general perception is that creativity is a cultural construct and therefore, there are many limitations to measure it. Regardless, some studies have attempted similar feats with mixed results [33, 36, 51]. Our own results suggest a strong correlation between narrative potential and quality, yet they fail to reveal the correlating metrics. Our limited success highlights the need for breakthroughs that improve existing knowledge bases, text recognition systems, emotional models and formal narratological models.

Additionally, regardless of the perceived narrative potential or quality, the deployed system was capable of producing relatively human-like yet economic textual artefacts. Perhaps, we could implement a similar consistency or coherence generator using recurring plot elements and their relationships from an ongoing collaborative story world. With the results of the previous sections in mind, the authors' structuring of the story world or the semantic diagrams could be used to inform the system of specific story world elements and their idiosyncrasies. The correlation between narrative potential and quality is another argument in favor of participative methodologies for collaborative writing, something we explore more in depth in the following chapter.

Modifying Entity Relationship Models for planning stories and its impact on potential authors

Alan Tapscott¹, Joaquim Colás¹, Ayman Moghnieh¹, and Josep Blat¹

¹*Grup de Tecnologies Interactives, Universitat Pompeu Fabra*

Abstract

Semantic models are often used for the communication of conceptual information. We propose a modification of the Entity Relationship model syntax, traditionally used for software engineering, to store and share plot data. Its flexibility demonstrated by decades of use for modelling a wide variety of software designs, suggests that it could be used by authors to represent a wide array of stories while keeping a certain degree of computational potential. The E-R model syntax is changed to suit better story plans, switching the emphasis on generic types to instanced story entities, retaining the original relationships and attributes. We conducted an experiment to evaluate the comparative performance of authors in understanding and contributing into a pre-existing story using traditional text and our modified E-R diagrams. The result analysis revealed that the E-R model is as effective as a written text plan in terms of reading comprehension, memorability and ease of contribution. Participants also achieve better comprehension with the combination of the two formats, always within the frame of our experiment. We discuss potential applications of these findings.

1 Introduction and related work

There have been many attempts to provide computational models for narrative and storytelling, pioneered by Propp's morphology of the folk tale [8]. Narrative models adequate for planning stories should deal with several aspects. First, different kinds or genres of narrative need different types of rules. Second, stories should be innovative and original. Computational models for stories often obstruct the creative development of the authors' contributions. In this paper we introduce a narrative model flexible enough to support a wide variety of stories while laying a strong foundation for all sorts

of contributions supporting their internal coherence. It is based on providing a story visual representation, inspired by the Entity-Relationship (E-R) model, widely used in software engineering. It draws part of its inspiration on Lehnert [7] high level analyses of stories in terms of plot units as arcs in a graph that encodes the plot of the story.

Our model exploits the analogy between the requirements of an information system, and the plan for a story. To achieve this, we modify a well-established model for software design by Chen [1] and its advances [5], the Entity-Relationship (E-R), which provides a semantic representation of the requirements. The key modification introduced is to focus on instances instead of classes. The model should support authors' communication for collaborative writing and as a first step we are interested in testing its impact on potential authors.

In what follows, first we explain more in detail our approach within supporting collaborative story planning and also explain the modifications to the E-R model to make it suitable for story representation. We then elaborate on the experimental setting and analyse the results of the experiment. A discussion of the results and other related work, followed by future work, ends the paper.

2 Strategy to support authors in story planning

The aim of this study is *to support authors in developing and maintaining their own semantic story representation when planning their stories collaboratively*. They should be able to include as much relevant information as they need. We make them responsible of finding the optimal structure for the planning of their own story using our modified E-R syntax and methodology. Which characters, feelings, or locations should be formalized is better left to the authors. Flexibility is a very desirable quality in our proposed model syntax. Author should determine what aspects or traits are important for the other authors to be aware or keep track of, adding them into the model to support the story structure and development. This approach might increase the internal coherence of the story, enhancing author collaboration dynamics.

3 The suitability of the E-R model and our modifications to enable it for story planning

The E-R model [1], introduced in the seventies, is still widely used by engineers to design data structures holding real-world input. It is flexible enough for any kind of quantitative data set, regardless of its anatomy, as it addresses any potential scenario. Thanks to this flexibility, authors could use

and adapt the E-R model to fit their needs. The model successfully represents semantic information into a computable structure, ready to be implemented, and could formally support the relevant information of an unfolding narrative plan.

More precisely system architects gather a requirements list for an information system and translate it into an E-R model through a process called data modelling; the output fits each requirement within a globally coherent system formally formulated. Our approach intends to exploit the analogy between system requirements and author story plans, both expressed in plain English sentences. Authors could develop both the written story plan (in plain sentences) and the E-R diagram simultaneously, each medium mirroring the mother one.

The E-R formulation we propose, uses Chen's original syntax changing its meaning, as software engineering and story planning have distinct purposes.

Entities represent the *agents* of the story. Any item with any degree of abstraction could fall into this category. In information systems entities usually denote classes or types, such as animal races or vehicle models. Stories deal with *specific* characters and thus we switch the focus from data *classes* to data *instances*. If two entities were to share a name, authors should add enough information, by extending the entity name and/or providing differentiating attributes, for instance, to avoid confusion and disambiguate them.

Relationships represent links between entities, for instance, informing of a fact, such as a contract of marriage between two characters. Since most story entities are instances, relationship cardinality is removed. A friendship relationship between two characters implies two persons, without additional cardinal data.

Attributes provide additional information regarding an entity or relationship. The common E-R formulation uses labels and values, but stories generally provide more values than labels. Entities in stories have few attributes in common. Thus, we avoid labels and store attributes as values. For instance, instead of having a *personality* attribute with *kind* as its value, a character might have the attribute *kind*. This modification sacrifices standardization in favor of flexibility.

The following example illustrates the differences between the information system and story planning modelling. *Employees have a Name and ID number. Every Employee has a Payroll assigned. Payrolls have a Gross income value and a Tax deduction value.* This might be modelled by an E-R diagram included in Fig 1.

In a story plan, a more typical statement would be: *Mike is an unhappy employee with a poor payroll*, included in Fig 2.



Fig. 1: Example of an information system E-R modelling



Fig. 2: Example of a story planning modified E-R modelling

Chen proposed a set of rules to translate system requirements formulated as English sentences into E-R diagrams [2], which can be used to translate explicit sentences from a story plan into its E-R model. Specifically the first four rules are simple and easy to use in the context of narratives. They convert *common nouns* into entity types, *transitive verbs* into relationship types, *adjectives* into entity attributes and *adverbs* into relationship attributes. The tenth rule proposed by Chen (meant to convert clause sentences into a group of interconnected sub-entities) can help in organizing nested plot data. We propose following a three-step strategy:

1. Formulate the story plan in plain explicit sentences; the narrative plan will be made of “story requirements”.
2. Translate the sentences into an E-R model using Chen’s rules [2]
3. Merge the E-R models and disambiguate any conflicts.

The merging process involves combining the new information with the already modelled one, and disambiguating any potential contradiction. It involves understanding the new entities and establishing their relationships to existing ones. It is a process that can be almost impossible to automatize or assist due to its subjective nature. For instance sometimes an entity must be transformed into other one, sometimes entities are duplicated or even merged. An author with a good conception of a story plan can perform such task.

4 Experimental setting

We tested whether modified E-R modelling supports the collaborative planning of stories by comparing against traditional written text. We had 3 groups with distinct experimental conditions: 1) Received a written story plan without the E-R model; 2) Received the same story plan along with its 'equivalent' E-R model; 3) Received only the E-R model. We asked each subject to read and extend the story plan, measuring several aspects of their contributions. Measurements include comprehension and memorability of the provided story plan, as well as the time spent in each of the phases of the experiment.

The experiment consisted of: 1) E-R training for those receiving an E-R model; 2) Reading of the received mediums, the text, the E-R model or both; 3) Comprehension test; 4) Memory test; 5) Contribution phase. Each subject could take as long as he or she wanted to complete the tasks.

4.1 Preparation: Story seed and E-R model

We chose the first part of the *Stagecoach* movie synopsis (taken from Wikipedia [4]) as our experimental source material. A volunteer Computer Science graduate created the E-R model from this synopsis, on a desk using pen and paper. It him approximately one hour to complete the task.

Subjects and groups 35 subjects were asked to read and contribute to the story. The experiment took place individually in an isolated desk. Their ages ranged from 20 to 65 years and they had from none to substantial experience in E-R modelling and story writing. They were divided into three groups:

- (experimental group 1) 10 subjects received both, the movie synopsis and the E-R model on a paper sheet.
- (experimental group 2) 13 subjects received only the E-R model on a paper sheet.
- (control group) 12 subjects received only the movie synopsis on a paper sheet.

Each group had approximately the same proportion of subjects with previous knowledge of E-R modelling (33% for only text, 38% for E-R and 40% for text and E-R).

E-R training

Subjects from experimental groups received a brief training on E-R modelling, including a basic explanation of the entity, attribute and relationship

concepts, along with some examples of their use, similar to the descriptions from the previous section of this paper.

Reading

Every subject received the corresponding printed material (story, E-R diagram, or both) and was given the briefing: *This is an incomplete story plan. Please read it.*

Comprehension test

After each subject finished with the reading, he or she filled a short questionnaire meant to determine the comprehension of the story plan. A pool of questions that required reading the story, understanding its content and being able to relate their concepts to answer them correctly (e.g., “What share in common characters X and Y?”) was created. This pool was evaluated by four judges who rated each question on a scale from 1 to 4 on its usefulness to determine the comprehension of the story. The test determined the top five questions on both higher consensus among judges, and higher overall rating. They were all open questions, the same judges rated the answers, and we used free-marginal Kappa coefficients to determine the agreement among judges.

Memory test

Without the written material, subjects were asked to answer a questionnaire by marking each sentence as either true or false, to determine how much was remembered of the story. The sentences were extracted from the original text and modified slightly for the false ones (e.g., “Character X is Y” to “Character X is not Y”).

Contribution phase

After returning the original reading materials to each subject along with a pencil and blank paper sheets, he or she was given the following briefing: *We would like you to contribute to this story plan? Any modification is allowed.* They were free to contribute as much as they wished, in any part of the original text or E-R diagram. No additional guidelines were provided.

4.2 Results

E-R training

Average time duration for this phase was of 2 minutes 36 seconds with a standard deviation of 28 seconds. No significant difference was found across the two groups (E-R and text + E-R groups).

Time

We found a significant difference between the time used by the three groups during the reading phase (ANOVA $p=0.0016$ $F(2,27)=8.2061$). A post-hoc analysis using t-tests in pairs revealed no significant difference between the groups with only text and Only E-R model, but each of these

two groups had significant difference with the group using both materials, which took more time to its subjects. We did not find statistically significant differences between the three groups in the time used in the phases of comprehension (ANOVA $p=0.1614$ $F(2,26)=1.9575$), memory (ANOVA $p=0.1954$ $F(2,32)=1.7186$), and contribution (ANOVA $p=0.9346$ $F(2,18)=0.0677$). We eliminated some subjects (between 2 and 3 per group) because they spent a drastically increased amount of time (more than ten times the average) to complete the contribution phase.

Comprehension and Memory

The ANOVA test reveals significant difference between the three group means with regards to comprehension ($p=0.0132$ $F(2,30)=5.0075$). The post-hoc t-tests revealed no significant difference between the two groups with either text or E-R, but a very significant difference between each and the group with both. The group that used text+E-R achieved significantly better results than the others. The ANOVA test on memory results did not uncover significantly different results between the three groups ($p=0.9341$ $F(2,30)=0.0682$).

Contributions

The contributions of the subjects (except the 6% that chose to make no contributions) were measured along a variety of parameters. There is a very weak positive correlation (coefficients ranging from 0.36 and 0.5) between reading time and 1) *word count* for text contributions, 2) *new entities* for E-R contributions, 3) *new but related entities* for E-R contributions and 4) *total new characters introduced* for all the contributions. All subjects who received only text contributed using text. From the individuals who received only the E-R model, 23% contributed using text and 69% used an E-R diagram. The individuals in the group that received both E-R and text contributed across both mediums equally (30% used text, 30% used E-R and 40% used both formats). There is a strong positive correlation between *text contribution word count* and *E-R contribution attributes and relations introduced* for subjects who contributed both text and E-R diagrams (coefficients around 0.8), broadly indicating that they contributed in similar proportions in both ways not privileging one of them. We did not find correlations between *comprehension* or *memory* and any of the *contribution measurements* (length, number of elements, etc.). The ANOVA test did not show significant differences between the contributions of the three groups in terms of text contribution *word count* (ANOVA $p=0.8372$ $F(2,19)=0.1793$), E-R contribution *entities / relationships used amounts* (t-test p ranging from 0.33 to 0.77 for the different measures) and overall *amount of old characters used* (ANOVA $p=0.2015$ $F(2,32)=1.6848$) and *new characters used* (ANOVA $p=0.8286$ $F(2,32)=0.1890$) in either contribution (when comparable).

Other aspects of the experiment

From our in-situ participant observations:

- 18% of the subjects within the group text+E-R model corrected or extended the original E-R model to better reflect the written story plan. They were all computer scientists.
- 27% of the group with only written story plan complained about grammar and confusing sentence construction (a couple of them even made corrections).
- 26% of the group with only the E-R model complained about the lack of reading order; 21% complained about the confusing relationship syntax, specifically the lack of roles for entities.
- 25% of the subjects complained about the amount of English names found in the story and/or E-R diagram.
- 5% of the subjects found out that the story was from *Stagecoach*.

5 Discussion

Our experiment intended to perform a first assessment of the effect of introducing E-R modelling to support authors in their writing through story planning. *Contributions* is the most directly related parameter; *Comprehension* is also important to ensure that authors understand what others did; *Memory* is a more secondary but important issue towards the support for writing. *Time* used is related to efficiency, which is rather secondary for this study.

There was no significant difference between the group with only text and the group with only E-R regarding reading time, comprehension and memory. Thus, it seems that the E-R model could replace the text, without an apparent impact; let us remark that this statement concerns only our limited experimental conditions.

On the other hand, the group which had both E-R diagrams and text took a significantly longer time in reading than the other two groups, but the comprehension was significantly enhanced as well. This seems to indicate benefits when introducing both text and E-R diagrams; better comprehension probably leads to enhanced coherence for collaborative authoring. The observations that some E-R experts corrected the diagrams, or some people fixed the text might be an indicator of motivation for quality / consistency; but it might be due to other factors, such as linguistic rigor.

The results did not indicate differences between experimental groups in terms of contribution amount, which suggests that the collaboration was not

stimulated when introducing E-R models (alone or with text); however it was not reduced either. Participants with only the E-R model contributed the least time; this suggests that extending the story in a modified E-R model requires less effort. On the other hand, the strong positive correlation between text length and amount of entities and relationships found in the group with both mediums suggests that they did it in parallel, textual content extensions mirroring E-R diagram extensions. We did not analyse qualitatively the contributions, an important factor in collaborative creative writing.

A substantial part (26%) of the group who only had the E-R diagram complained about the lack of a reading order. In information systems, E-R models represent a snapshot of a data set structure. Transformations are a fundamental part of stories and that is somewhat opposed to the nature of standard E-R models. Using multiple E-R models, perhaps one per chapter, episode or page, could provide an answer to this problem. Another alternative would be using a syntax to map the transformations and story progression into one single E-R model, such as the one proposed by Klopprogge [6]. Also within this group, 21% complained about the confusing syntax of relationships; the roles in relationships were removed (e.g., which character hates and which one is hated in a hate relationship), and this makes a story harder to understand. We did not anticipate this, and perhaps using Chen's original relationship role labelling or Cortman et al approach [3], introducing directed networks of relations, could make the E-R model more understandable for story planning.

A possible explanation to the increased comprehension of the text+E-R group might lie in the existence of the induced paths Corman et al. found in their work [9]. Such paths might mean that the E-R model suggests some implicit information to readers, as no significant difference was found among the time answering the tests by the groups, nor a correlation between individual's reading time and comprehension score, meaning that participants that read for a longer time did not necessarily understand the story better. The pre-existing knowledge on E-R models might have had some impact on the results, and the subject's background should be better controlled on subsequent experiments.

The plan used in the experiment was relatively small. The scalability of the E-R model for stories has yet to be tested. The computational characteristics of the model to support story coherence were not explored. A narrative coherence qualitative evaluation of the resulting contributions would be necessary. New experiments dealing with scalability and coherence are intended.

6 Future work

As indicated earlier, support for the temporal dimension seems key because of the transformational nature of stories, and adding roles to the relationships is also necessary. Our modified E-R model could be extended to include recent E-R improvements such as the ones proposed by Hartmann et al [5]. Semantic structures could be introduced to streamline story planning and assist the authors in their task, without compromising the model flexibility.

The author's perception of the story itself could be instrumental in creating useful tools that support collaboration, especially when it deals with creativity. Allowing authors to also tweak the syntax to support their own understanding of the story might provide them with more adequate information systems. This could be mixed with solutions from other fields, such as language processing, enabling less expert profiles to inform computational systems. Their contributions could represent computational models that describe the morphology of different genres (beyond the limited scope of our western). Potential pre-made skeletons could include genre tropes and recurrent writing techniques, formulated in a syntax ready to be integrated into the author's story representation.

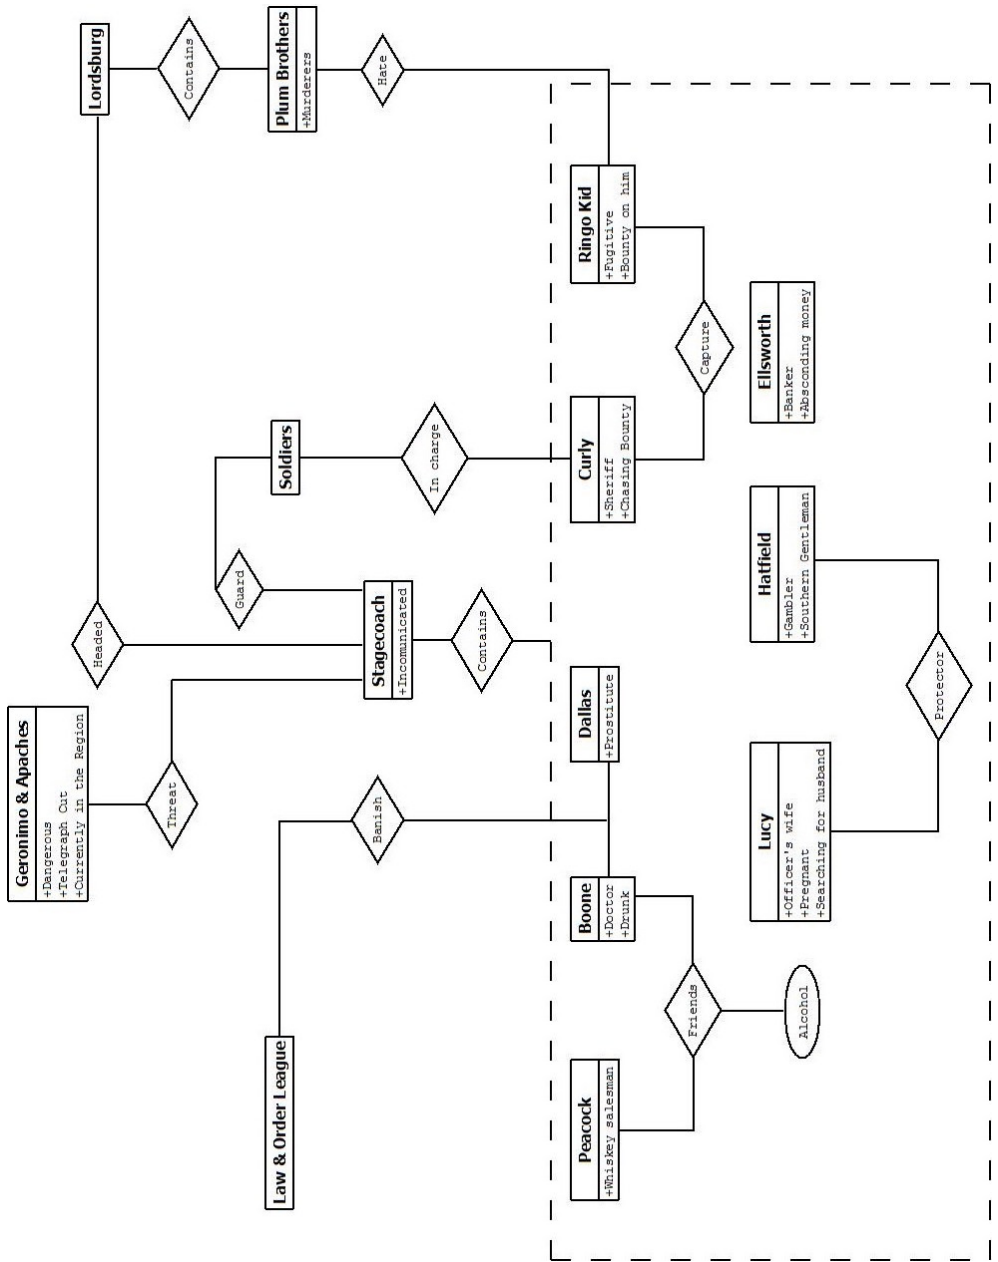
Acknowledgements We sincerely thank Xavier Ruiz Collantes, Miguel Angel Carralero, Jose Emilio Lavilla and Daniel Giralt for their help, as well as all the people who volunteered for for the experiment.

References

- [1] Peter Pin-Shan Chen. The entity-relationship model toward a unified view of data. *ACM Transactions on Database Systems (TODS)*, 1(1):9–36, 1976.
- [2] Peter Pin-Shan Chen. English sentence structure and entity-relationship diagrams. *Information Sciences*, 29(2):127–149, 1983.
- [3] Steven R. Corman, B. Hunter Ball, Kimberly M. Talboom, and Gene A. Brewer. Assessing Two-Mode Semantic Network Story Representations Using a False Memory Paradigm. In Mark A. Finlayson, Bernhard Fisseni, Benedikt Löwe, and Jan Christoph Meister, editors, *2013 Workshop on Computational Models of Narrative*, volume 32 of *OpenAccess Series in Informatics (OASICs)*, pages 53–61, Dagstuhl, Germany, 2013. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [4] John Ford. Stagecoach, 1939.
- [5] Sven Hartmann and Sebastian Link. English sentence structures and eer modeling. In *Proceedings of the fourth Asia-Pacific conference on Conceptual modelling-Volume 67*, pages 27–35. Australian Computer Society, Inc., 2007.
- [6] Manfred R. Klopprogge. Term: An approach to include time dimension in the entity-relationship model. In *Proceedings of the Second International Conference on the Entity-Relationship Approach to Information Modeling and Analysis*, ER '81, pages 473–508, Amsterdam, The Netherlands, The Netherlands, 1983. North-Holland Publishing Co.
- [7] Wendy G Lehnert. Plot units and narrative summarization. *Cognitive Science*, 5(4):293–331, 1981.
- [8] Vladimir Propp. *The Morphology of the Folktale*. University of Texas Press, Austin, 1968.

A Appendix: Story seed text, E-R diagram and result tables

In 1880, a motley group of strangers boards the east-bound stagecoach from Tonto, Arizona Territory to Lordsburg, New Mexico Territory. These travelers are unremarkable and ordinary at first glance. Among them are Dallas, a prostitute who is being driven out of town by the members of the "Law and Order League"; an alcoholic doctor, Doc Boone; pregnant Lucy Mallory, who is traveling to see her cavalry officer husband; and whiskey salesman Samuel Peacock. When the stage driver, Buck, looks for his normal shotgun guard, Marshal Curly Wilcox tells him that the guard has gone searching for fugitive the Ringo Kid. Buck tells Marshal Wilcox that Luke Plummer is in Lordsburg. Knowing that Kid has vowed to avenge the deaths of his father and brother at Plummer's hands, the marshal decides to ride along as guard. As they set out, U.S. cavalry Lieutenant Blanchar informs the group that Geronimo and his Apaches are on the warpath and his small troop will provide an escort until they reach Dry Fork. As they depart, the stagecoach is flagged down to pick up two more passenger, gambler and Southern gentleman Hatfield as well as banker Henry Gatewood, who is absconding with \$50,000 embezzled from his bank. Along the way, they come across the Ringo Kid, whose horse became lame and left him afoot. Even though they are friends, Curly has no choice but to take Ringo into custody.



Tab. 1: Average time spent in the different phases

	Reading phase	Comprehension test	Memory test	Contribution phase
Text group	124 seconds $\sigma=52.8$	261 seconds $\sigma=80.5$	70 seconds $\sigma=35.4$	288 seconds $\sigma=133.4$
E-R group	142 seconds $\sigma=72.4$	233 seconds $\sigma=71.6$	92 seconds $\sigma=31.2$	272 seconds $\sigma=111.7$
Text + E-R group	236 seconds $\sigma=69.6$	315 seconds $\sigma=107.5$	93 seconds $\sigma=34.2$	293 seconds $\sigma=92.5$

Tab. 2: Time used in the comprehension phase

	Text vs. E-R	Text vs. Text+E-R	ER vs. Text+E-R
p=0.5170 (no significance) t=-0.6625	p=0.0011 (very strong significance) t=-3.9894	p=0.0101 (strong significance) t=-2.8756	

Tab. 3: Memory and comprehension test averages

	Comprehension test (0 to 3 points)	Memory test (0 to 9 points)
Text group	2.325 points $\sigma=0.329$	7.417 points $\sigma=1.443$
E-R group	2.254 points $\sigma=0.438$	7.384 points $\sigma=1.387$
Text + E-R group	2.680 points $\sigma=0,167$	7.5 points $\sigma=1.354$

Tab. 4: Comprehension test judge agreement

Average item-total correlation between ratings:	Percent of overall agreement Po	Fixed-marginal kappa	Free-marginal kappa
0.8900	0.6647	0.3602	0.553
Overall Re-sult			
Moderate agreement			

Tab. 5: Comprehension test t-tests

	Text vs. E-R	Text vs. Text+E-R	ER vs. Text+E-R
	p=0.6124 (no significance) t=0.5150	p=0.0046 (strong significance) t=-3.259	p=0.0088 (strong significance) t=-3.0763

Tab. 6: Contribution per group

	Text word count	Text sentence count	sen- tences	E-R new entities	E-R old characters used	Total new characters introduced
Text	74.6	4.6				
E-R	56	2.6	0.8	1.9		
Both	80.7	4.3	0.5	0.9		
	E-R old entities related	E-R new entities related	Total old characters used	Total new characters introduced		
Text			5.2	0.3		
E-R	4.1	1.1	3.4	0.5		
Both	5.6	1.6	5.5	0.5		

Empirical Evidence of the Limits of Automatic Assessment of Fictional Ideation

A. Tapscott¹, J. Gómez¹, C. León¹, J. Smailović², M. Žnidaršič²,
and P. Gervás¹

¹Facultad de Informática, Universidad Complutense de Madrid

²Department of Knowledge Technologies, Jožef Stefan Institute

Abstract

Automatic evaluation of fictional ideation systems and their output is a topic relevant to Computational Creativity. Models and techniques have been proposed for this task, but their applicability is limited to the field of fictional ideation. In this paper we describe an evaluation procedure for fictional ideation, which compares human validation of the ideas with a number of automatically generated metrics obtained from them. We report on the observed limits of this procedure. The results suggest that, besides technical limitations, providing a stable evaluation method is fundamentally incomplete unless the full creative phenomenon is modelled, including aspects that are beyond current technical capabilities.

1 Introduction

Evaluation of creative processes and artefacts is key to computational creativity. Explicitly reflecting on the relative value and novelty is crucial if machines are to produce content that would be *deemed creative* [6]. As such, addressing evaluation is fundamental for computational creativity that can successfully fulfill human needs.

This crucial aspect contrasts with the relative scarcity of systems explicitly generating rich evaluation of their own generated material or inner processes. Some systems arguably control the quality of their artifacts by carrying out a process that ensures a minimum relative quality, but an explicit evaluation arguably represents a qualitative

advantage, both theoretical (as studied by computational creativity frameworks [29]) and practical ([4]).

Although the semantics of creativity are elusive and usually problematic, the vision that quality and novelty influence the perception of the creativity of an artifact (at least from the point of view of observation) is commonly accepted. Still, quality and novelty vary depending on the domain and context. Theoretical discussion on this exists and it is seminal in the field [1, 2], while other works attempt to offer either formal or procedural techniques for evaluating creativity [18, 25, 30]. These efforts address the evaluation of creativity in generic terms, and they are of limited applicability for the evaluation of the quality of specific artifacts generated automatically. It might be the case that the assumption that there is a global definition of creativity applicable to every creative domain is not possible, but we still need more empirical evidence supporting whether this is so.

Moreover, even when working within a domain in which there is an agreed definition of characteristics assumed to play a role in creativity (let us say *quality*), addressing explicit automatic evaluation can be a costly task, even more costly than creating the generative system that is being evaluated. It is not uncommon that being able to generate appropriate artefacts is doable, while yielding an explicit, measurable evaluation is not (for instance, in images generated by evolutionary computing [15]).

This paper reports on an empirical study in which the output of an automatic ideation system is assessed by computational means. When compared to human evaluation, the conceptual and practical limits of the approach were evidenced. This led to an in-depth analysis of the challenges, which is provided in Section 5.

2 Previous Work

While all scientific exploration requires thorough evaluation of the steps taken, doing so in creativity represents a challenge. How to assess creativity itself is a commonly discussed aspect of the whole phenomena of creative generation. While most authors agree on the correlation between a number of features and the perception of creativity, there is no consensus either on what these features are or how they really correlate. Moreover, adding computers to the problem makes it even more difficult to know whether a system has been successful or not. There is still a debate on what parts should be evaluated, the influence of the programmer on the output, the very definition of creative behavior, the decision of whether to focus on the

process or the artifacts (or both), and many others.

The few examples present in the literature describing actual evaluation of automatic creative systems usually focus on less ambitious, more measurable aspects. This makes these systems less useful from a general perspective, but they nonetheless provide insight on the current capabilities of computer systems to assess their own production.

There is, however, a number of proposals that try to provide guidelines to evaluate creative systems. For instance, Ritchie [24, 25] addresses the issue of evaluating when a program can be considered creative by outlining a set of empirical criteria to measure the creativity of the program in terms of its output. He makes it very clear that he is restricting his analysis to the questions of what factors are to be observed, and how these might relate to creativity, specifically stating that he does not intend to build a model of creativity. Ritchie's criteria are defined in terms of two observable properties of the results produced by the program: *novelty* (to what extent is the produced item dissimilar to existing examples of that genre) and *quality* (to what extent is the produced item a high-quality example of that genre). To measure these aspects, two rating schemes are introduced, which rate the typicality of a given item (item is typical) and its quality (item is good). Another important issue that affects the assessment of creativity in creative programs is the concept of *inspiring set*, the set of (usually highly valued) artifacts that the programmer is guided by when designing a creative program. Ritchie's criteria are phrased in terms of: what proportion of the results rates well according to each rating scheme, ratios between various subsets of the result (defined in terms of their ratings), and whether the elements in these sets were already present or not in the inspiring set. Ritchie's criteria have been used in subsequent evaluations of creative systems output [7, 21, 8].

Pease et al. [19] discuss relevant factors to evaluating systems in terms of creativity. The proposed framework mainly takes into account input provided, output produced and process employed. Each of these categories are detailed in depth, detailing their required measures. Before detailing the measurement methods, Pease et al. provide assumptions regarding creativity, also admitting their 'somewhat arbitrary' nature. The evaluation tests proposed deal with two main aspects: how close does the test predict human evaluation of creativity and how possible and practical it is to apply the test to a system. Overall, this work suggests that the very definition of creativity is subjective and that evaluating systems in a general way is problematic.

Colton et al. [5] propose an extension of Ritchie's criteria [24] that

attempts to determine the impact of the input data on the creative artifact produced by a system. This more agnostic approach attempts to obtain an objective measure by comparing the output of the system to the inspirational material used as input. This investigation attempts to discriminate systems that overfit or shuffle input data (fine-tuning) instead of producing genuine novel artifacts. Among other conclusions, the authors state that comparing creative systems might not be viable, suggesting their criteria to be used as guidelines for program construction rather than post-hoc evaluation.

The creative tripod framework, proposed by Colton [3], is built around the premise that a creative system must demonstrate skill, imagination and appreciation. These qualities are not required to be possessed by the system, but rather to be perceived as possessed by the system. This is an important remark by Colton to avoid debates around the definition of creativity. The framework also includes the programmer, the system and the consumer, however Colton is only interested in the program's behavior.

Pease and Colton [18] propose an alternative to the Turing Test to assess computational systems' creativity, the FACE (Frame, Aesthetic, Concept, Expression of concept) and IDEA (Iterative Development Execution Appreciation) model. The model includes creative acts and audiences, with relevant measures such as popularity, appeal, provocation, opinion, subversion and shock. Putting the focus on the reaction produced by the creative artifact, this model attempts to avoid the shortcomings of the Turing Test by going further than merely assessing the capacity of a creative system to imitate human behavior. By including the audience into the model, this approach acknowledges the highly subjective nature of creativity evaluation.

SPECS [9], introduced by Jordanous as "a standardised and systematic methodology for evaluating computational creativity", represents a substantial effort to provide a standard for evaluating the creativity of a system in the field of computational creativity and address the multi-faceted and subjective nature of creativity. Its flexible nature allows SPECS to adapt to the demands of the researchers' field, applying the required demands and standards. The methodology informs researchers of their system's strength and weaknesses, providing useful feedback for achieving creative results.

2.1 Evaluation of Automatically Generated Narrative

Automatic generation of narratives has been a long-standing goal of Artificial Intelligence since its very beginning. There are a number of

systems described in the literature, but the evaluation of these systems – be it its output, its creative process or whatever other aspect – is seldom found. This is most likely due to the fact that the average quality or variety of the generated stories is not really comparable to those written by most humans, not necessarily professional writers.

The Mexica system [23] includes procedures for the dynamic assessment of the novelty of a story in progress with respect to previously known stories. Novelty is considered in terms of how the stories differ in terms of the actions they include and their frequency of appearance.

In Pérez et al [22] three different characteristics are considered as relevant for measuring story novelty: sequence of actions, structure of the story, and use of characters and actions.

Peinado & Gervás [20] carried out an empirical study of how generated stories were perceived by a set of human volunteer evaluators. Human judges blindly compared one of the generated basic stories to two alternatives: one rendered directly from a stored fabula of the knowledge base and another randomly generated. Values were collected for: *linguistic quality* (how well is the text written), *coherence* (how well is the sequence of events linked), *interest* (how interesting is the topic of the story for the reader) and *originality* (how different is the story from others).

León & Gervás [11] propose a model, intended as a tool to drive automatic story generation, of how quality is evaluated in stories. This paper proposes a computational model for story evaluation in which an evaluation function receives stories and outputs a value as the rating for that story. The value for this function is computed from values assigned to: accumulation of contributions from individual events depending on the meaning of the event – aspects such as whether the reader wants to continue reading the story, or how much danger or love the reader perceives in the story –, appearance of patterns or relationships between the events of a story – aspects such as causality, humour or relative chronology – and inference – which captures the ability to interpret stories by adding material to explain what they are told even if it is not explicitly present in the story. The evaluation function has been implemented as a rule based system.

Ware, Young et. al. [27] propose a formal model for narrative conflict with seven dimensions from various narratological sources meant to aid in distinguishing one conflict from another: participant, subject, duration, balance, directness, intensity and resolution. Their experimental results [28] suggest the model predicts these seven dimensions of narrative conflict similarly to human criteria. Their good results predicting human-perceived narrative conflict suggest a similar

approach may be viable for measures related to creativity.

3 Evaluating Automatic Ideation

Original ideation is central to any creative process. Coming up with innovative ideas that potentially trigger the creation of new material is fundamental to human creativity. It is not uncommon to focus creative processes on the identification of a single, valuable idea that unlocks new paths leading to finished artifacts. Although human creative teams usually rely on pure ideation to foster creativity, there have only been a few small, ad-hoc studies of how to automate ideation until recent times. Section 3.1 describes an effort to provide a system able to produce novel ideas.

3.1 The What-If Machine

Llano et al. have recently proposed an automatic ideation system [13, 14, 12]. This computational system is designed to produce relatively valuable and novel ideas autonomously. This system, the *What-If Machine*¹, includes a module for analysing the ideas and generating narrative metrics, and a module for computing a predictive machine learning model. This model is trained against collected human evaluations of what-ifs, and is intended to learn a robust function from narrative metrics to perceived overall quality. Two main hypotheses guide the design of the What-if Machine and the presented research:

1. There is a strong correlation between the perceived *overall quality* and the perceived *narrative potential*, in the sense that if the audience perceives high narrative potential, it will also perceive a high overall quality. The overall quality is defined in terms of the analyzed response from humans (i.e. no specific model beyond what humans say about quality is assumed), and the narrative potential is assumed to be directly proportional to the amount and quality of the stories a certain what-if can trigger or inspire.
2. There is a set of computable metrics whose values correlate (directly or indirectly) with the overall quality and the narrative potential.

¹ The What-if Machine: <http://www.whim-project.eu/>.

The What-If Machine is, to the best of our knowledge, the only attempt to implement a computer system able to produce novel what-if ideas. The What-If Machine is a distributed computer system in which several modules collaborate in order to output rendered what-ifs. Five modules compose the system:

1. The **ideation module** produces, using a knowledge base, what-if ideas formalized as *mini-narratives*.
2. The mini-narratives are fed into the **narrative-based metric generation**, which generates values for a set of metrics which hypothetically have a correlation with human perception of quality. These metrics are based on narrative properties of the what-ifs.
3. The mini-narratives, now enriched with its corresponding metrics, are sent to a **crowd-sourcing evaluation module**, which applies machine learning to create and refine models for predicting overall quality against human ratings.
4. The **world view** creation, providing knowledge for what-if generation, story creation and metric computation.
5. The finished, filtered what-ifs are finally passed to a **rendering module**, which creates artifacts from the final what-ifs (stories, texts or images, for instance).

A subset of the What-If Machine (modules 1, 2 and 3) was used to generate the material for the study, which is described in detail in Section 4.

4 Study

A pilot study was performed to determine the feasibility of predicting the perceived *quality* and *narrative potential* in the artifacts created by a computable creative system. Both magnitudes have been introduced in the previous section, and in order to avoid influencing our subjects, no definition for them is provided in the questionnaires (as seen in Fig. 1). This naive approach is a result of our focus on the model and its capability to predict human assessment instead of introducing our own views or definitions. The study was conducted to obtain the human rating of perceived *quality* and *narrative potential*.

Using both measures, a machine learning process will search for correlations between some metrics (detailed in the next section) and

the perceived *quality* and perceived *narrative potential*. This should allow us to determine what measures are relevant to predict human-perceived *quality* and *narrative potential* to produce what-ifs that present both qualities to human observers.

4.1 Metrics

Since we have no certainty about what metrics extracted from each what-if's mini-narrative may impact over the perceived *quality* and *narrative potential*, we focused on generating the maximum amount of computable features. The impact of these features on the perceived *quality* and *narrative potential* may be obtained with machine learning techniques (we refer to these features as *metrics*). This approach is similar to the one used by Nowak for image classification [17] that generates a high number of arbitrary features from each image.

A mini-narrative is a structure that contains a set of **narrative points** linked to schemas like *setting* or *resolution*. Each **narrative point** is a set of **narrative statements** that provide information about characters or events through predicates (e.g., *dog is old* or *dog learns to play a piano*). **Narrative statements** may be related to one another (*caused by* or *inferred by* another statement).

The next list includes the set of implemented features along with their description:

- **Length**: mini-narrative **narrative points** amount.
- **SettingQuality**: Amount of schemas divided by 3.
- **ExplicitFact**: the amount of **narrative statements** in the mini-narrative.
- **RatioCharacters**: the character/statement ratio.
- **Originality**: hits returned by the full text of the mini-narrative in the *Bing* search engine.
- **OriginalityAccurate**: hits returned by the **exact** full text of the mini-narrative in the *Bing* search engine.
- **Divergence**: average hits returned by the mini-narrative statements in the *Bing* search engine.
- **DivergenceMinimum**: minimum hits returned by the mini-narrative statements in the *Bing* search engine.

- **Evolution:** amount of `learnTo` predicates found in the mini-narrative.
- **Handicap:** amount of negated `capableOf` predicates found in the mini-narrative.
- **InterestingLife:** amount of negated `doesFor` predicates found in the mini-narrative.
- **TotalStoriesGenerated:** amount of stories generated by the story generator from the current mini-narrative.
- **StoryCharacters:** average number of characters in the generated stories.
- **Names:** StanfordNLP [16] queries for the what-if's names.
- **NamesRatio:** *Names/ExplicitFact* ratio.
- **Valence:** Sum per statement, each statement codified as +1 if a fact is positive, -1 if negative and 0 otherwise).
- **ValenceAverage:** *Valence/ExplicitFact* ratio.
- **JointWordsProbability:** joint probability average for each set of words using *ngrams*. For this metric we use the Project Oxford² services.
- **JointWordsProbabilityMinimum:** the minimum joint probability for the set of words using *ngrams* from Project Oxford.
- **RealityDistortionRatio:** events in the mini-narrative that negate a fact from the *knowledge base* are considered a *reality distortion*. This metric provides the *reality distortion* amount/*ExplicitFact* ratio.
- **FictionalAdditionsRatio:** any event in the mini-narrative that is missing from the *knowledge base* is considered a *fictional addition*. This metric provides the *fictional addition* amount/*ExplicitFact* ratio.
- **FictionalRatio:** *reality distortion* amount plus *fictional addition* amount/*ExplicitFact*.

² <https://www.projectoxford.ai/>

- **ResolutionTriggerRatio:** *resolution events solve conflicts* from the mini-narrative. Provides the *resolution event amount/ExplicitFact* ratio.
- **MainCharacterEventsRatio:** *protagonist statements* are statements in which this actor plays any role. This metric provides the *protagonist statement amount/ExplicitFact* ratio.

4.2 Methodology

A set of 890 what-ifs were generated by the What-If Machine. All of their source mini-narratives were processed by the metric generation system. A total of 15 different questionnaires were created, each including 10 what-ifs rendered as text from the original set of 890. 150 what-ifs were included in the evaluation set. 101 volunteers received a link that randomly redirects to one of the 15 possible questionnaires through email. Given the simplicity of the questions, Google Forms was our platform of choice. The platform was robust and stable and all of the answers were successfully stored in a Google Sheet document automatically. There was no active supervision for each subject given the remote nature and limitations of the Google Forms platform.

4.3 Questionnaire

The questionnaire informed subjects about their participation in a study related to computer-generated content (Figure 1). Some demographic information was queried (age, gender and English level) and then they were asked to evaluate the overall quality (on a 0-5 Likert scale) of each what-ifs plus its narrative potential (yes/no binary answer). A text box accepting any comment was also provided in order to gather additional qualitative information.

You are about to evaluate some of the preliminary results of the “WHIM: The What-If Machine” research project from the European Union. The overall objective of the What-If Machine is to automatically generate fictional ideas with cultural value. You will be presented a number of what-if style ideas and we kindly ask you to rate them according to the following features:

- Overall quality: from 0 (no quality) to 5 (superb quality).
- Narrative potential (yes/no).
- Any observation you can provide.

Completing the questionnaire should not take more than 10 minutes. We really appreciate your contribution to the project.

Fig. 1: Information presented to the user in the evaluation questionnaire.

4.4 Results

101 subjects participated in the study. Statistical analysis of the results revealed no significant differences between evaluators in terms of English level, age or gender. For instance, the quality (Q) for gender yielded $\mu(Q)_{male} = 2.66$, $\sigma(Q)_{male} = 0.75$; $\mu(Q)_{female} = 2.69$, $\sigma(Q)_{female} = 0.89$. The corresponding results for English and age are comparable.

Questionnaires provided 1,007 *Quality* and 1,004 *Narrative Potential* rankings for the 150 *What-Ifs* used. *What-Ifs* were ranked between 1 and 27 times. For the *Narrative Potential* (P) measurements, we mapped “Yes” to +1, “Not sure” to 0, and “No” to -1. Overall measures resulted in $\mu(Q) = 2,4$ and $\sigma(Q) = 1,3$ for *Quality* and $\mu(P) = -0,05$ and $\sigma(P) = 0,89$ for *Narrative Potential*. Individual *What-Ifs* aggregated ranking values were used for calculating:

- Pairwise correlations between perceived *Quality* and perceived *Narrative Potential*, perceived *Quality* or perceived *Narrative Potential* and the metrics, and between individual metrics.
- Global measure of attribute importance for these metrics in predictive modeling of the average perceived *Quality* or perceived *Narrative Potential*.

Pairwise correlations Metrics that provided the same values for all *What-Ifs* in the dataset were discarded. Correlation coefficients were calculated with the Pearson Product-Moment. There is a strong positive correlation between *Quality* and *Narrative Potential* averages (0.83) and medians (0.758). As seen in table 1, both measures correlate positively with some metrics, such as *MainCharacterEventsRatio* and *RatioCharacters* and correlate negatively with others, such as *ExplicitFact* and *Length*.

Importance for Predictive Modeling In order to determine the importance of each metric in predicting perceived *Quality* and *Narrative Potential* we used the Relief measure [10, 26], which is a method commonly used for feature selection in machine learning. This measure does not assume independence among the metrics, but takes their possible interdependence into account. The more the Relief scores are positive, the more a metric contributes to prediction of a target value (in our case, the value of average *Quality* or the average *Potential*). The ones that scored close to zero or negative are irrelevant and those with negative values have even a negative impact.

According to the results in Table 2 it seems that most of the metrics have no use in predictive models of average Quality. For the average Narrative Potential, however, most of the metrics seem to be slightly informative. According to Relief ranks for the metrics results, usefulness of the metrics for average Quality is to some extent inversely proportional to their usefulness for the average Narrative Potential. The absolute values of the Relief scores depend on the characteristics of data and the parameters of the assessment, which makes it difficult to use absolute thresholds for judgements on the relevance of features. However, a strong correlation among the Quality and Narrative Potential values and a mismatch of the Relief scores of metrics for these two targets provide an indication that also the contributions of the positively scored metrics are likely to be too low to be considered relevant.

5 Relative Limits of Evaluating Quality

The results previously presented evidence that there is a strong correlation between narrative potential and perceived overall quality of a what-if, which indicates that focusing on narrative plausibility as one of the main factors of quality can lead to better results. Moreover, some of the metrics are weakly correlated to narrative potential. However, these results are still inconclusive, and there is a number of aspects worth mentioning for their influence on the results.

Automatically generating stories and computing useful values for metrics is heavily dependent on the available knowledge. The outcome of the system is constrained by the use of ConceptNet. The amount of relations that can be safely used in ConceptNet is small and the richness and depth of the chains of properties is limited regarding to its use as a source for narrative processing. This makes it necessary to address knowledge management from a different perspective. The WHIM project currently includes a whole module for providing robust knowledge to the rest of the modules, and the impact of the application of this subsystem on the creation and evaluation of what-if ideas will be reported once the results are ready.

The generation process (for the what-ifs, the stories and the metrics) strongly influences the overall outcome. Many design decisions have been taken in order to provide a working, implemented prototype able to generate actual what-ifs, and these decisions set the kind of what-ifs generated, the complexity of the stories and many other aspects. The provided results are then the outcome of a specific implementation which does not claim any generality. However, the ap-

proach itself (namely the generation-metric computation-evaluation process) is presented as a generally applicable method for producing novel what-if ideas.

The used metrics for labeling narrative properties do not cover all computable features. There is a large number of aspects that can be extracted from a what-if, and the narrative-based feature extraction module of the What-If Machine does not currently provide coverage for all of them. This is considered to be not strictly relevant with regard to the methodology and scope of the study. To test the second hypothesis (the existence of a correlation between a certain set of metrics and the overall quality and plausibility), the metrics must be improved. For that purpose, the presented study gives valuable insight on which direction to go next.

The weak correlation between our metrics and the quality perceived by humans suggested that considering more sophisticated metrics was necessary. Some of them were considered:

1. **Humanization:** An approximation of how much human-like the main character is, assuming that fictional scenarios use characters that, while behaving like humans, can be non-human.
2. **Empathy:** How much empathy will a reader feel about the characters.
3. **Tragedy:** The amount of tragedy in the story.
4. **Reality:** How real and current the context is. An approximation of fictionality in terms of context.
5. **TimeSpan:** The time span the story covers. It could be minutes, days or years.

Modelling and implementing these metrics proved to be beyond technical capabilities because it required complex, rich knowledge bases (1, 4), reliable text understanding systems (5), sophisticated emotional models (2) or formal versions of narratological models (3). All of these resources are currently not available.

6 Conclusions

The current paper has presented a pilot study trying to gain insight on two hypotheses, namely that (1) human evaluation on overall quality of what-if ideas correlates to the perception of narrative potential and

that (2) there is a set of computable metrics that also correlate to this perception. The study has evidenced that there is a strong correlation between quality and narrative potential for humans (1), but failed to prove such a strong correlation between the current metrics and the human ratings. These results have been analysed and discussed in terms of the limited potential of the current implementation of both the fictional ideation procedure and the method employed to evaluate it. Actual implementations lack the required complexity to approximate evaluations with a relatively acceptable level of accuracy, mainly due to the limited technical capabilities of current computational solutions.

Tab. 1: The correlation coefficient between average/median *Quality (Q)* or *Narrative Potential (P)* labels and the metrics. The values are sorted by correlation coefficient values of the average *Quality*.

	Avg <i>Q</i>	Mdn <i>Q</i>	Avg <i>P</i>	Mdn <i>P</i>
MainCharEventsRatio	0.371	0.346	0.379	0.329
RatioCharacters	0.354	0.296	0.368	0.307
ResolutionTriggerRatio	0.342	0.303	0.305	0.261
TotalStoriesGenerated	0.312	0.250	0.321	0.264
JointWordsProbMin	0.308	0.289	0.367	0.314
...
ValenceAverage	-0.219	-0.188	-0.296	-0.249
ValenceSum	-0.258	-0.234	-0.323	-0.276
StoryCharacters	-0.283	-0.269	-0.327	-0.285
ExplicitFact	-0.379	-0.336	-0.406	-0.345
Length	-0.379	-0.336	-0.406	-0.345

Tab. 2: Relief measure results for average *Quality (Relief Avg Q)* and average *Narrative Potential (Relief Avg P)*. Rows sorted by Relief Avg *Q*. The best three results are in bold and the worst three are in italics.

Metric	Relief Avg <i>Q</i>	Relief Avg <i>P</i>
Handicap	0.027	<i>-0.009</i>
MainCharacterEventsRatio	0.007	0.004
NamesRatio	0.001	0.006
DivergenceMinimum	0.000	0.000
JointWordsProbabilityMinimum	0.000	<i>0.000</i>
Divergence	0.000	<i>0.000</i>
Originality	-0.006	0.013
...
FictionalAdditionsRatio	-0.075	0.028
InterestingLife	-0.116	0.045
TotalStoriesGenerated	-0.116	0.045
OriginalityAccurate	-0.126	0.024
FictionalRatio	-0.142	0.039
RatioCharacters	-0.142	0.039
SettingQuality	<i>-0.147</i>	0.024
Names	<i>-0.147</i>	0.024
ValenceSum	<i>-0.174</i>	0.033

References

- [1] Boden, M.: Computational Models of Creativity. *Handbook of Creativity* pp. 351–373 (1999)
- [2] Boden, M.: *Creative Mind: Myths and Mechanisms*. Routledge, New York, NY, 10001 (2003)
- [3] Colton, S.: Creativity Versus the Perception of Creativity in Computational Systems. *Proceedings of the AAAI Spring Symposium on Creative Systems (Colton 2002)*, 14–20 (2008)
- [4] Colton, S.: The painting fool: Stories from building an automated painter. *Computers and Creativity* 9783642317, 3–38 (2012)
- [5] Colton, S., Pease, A., Ritchie, G.: The effect of input knowledge on creativity. *Technical Reports of the Navy Center for* (2001)
- [6] Colton, S., Wiggins, G.: Computational creativity: The final frontier? *ECAI* (2012)
- [7] Gervás, P.: Linguistic creativity at different levels of decision in sentence production. In: *Proceedings of the AISB 02 Symposium on AI and Creativity in Arts and Science*, 3rd-5th April 2002, Imperial College. pp. 79–88 (2002)
- [8] Haenen, J., Rauchas, S.: Investigating artificial creativity by generating melodies, using connectionist knowledge representation. In: *The Third Joint Workshop on Computational Creativity* (2006)
- [9] Jordanous, A.: A Standardised Procedure for Evaluating Creative Systems: Computational Creativity Evaluation Based on What it is to be Creative. *Cognitive Computation* 4(3), 246–279 (2012)
- [10] Kira, K., Rendell, L.: A practical approach to feature selection. In: *Proceedings of the ninth international workshop on Machine learning*. pp. 249–256 (1992)
- [11] León, C., Gervás, P.: The Role of Evaluation-Driven rejection in the Successful Exploration of a Conceptual Space of Stories. *Minds and Machines* 20(4), 615–634 (2010)
- [12] Llano, M.T., Colton, S., Hepworth, R., Gow, J.: Automated Fictional Ideation via Knowledge Base Manipulation. *Cognitive Computation* pp. 1–22 (2016)

- [13] Llano, M.T., Cook, M., Guckelsberger, C.: Towards the automatic generation of fictional ideas for games. *Experimental AI in Games Workshop* (2014)
- [14] Llano, M.T., Hepworth, R.: Automating fictional ideation using ConceptNet. *Proceedings of the AISB14 symposium on computational creativity* (2014)
- [15] Machado, P., Martins, T., Amaro, H., Abreu, P.: Beyond interactive evolution: Expressing intentions through fitness functions. *Leonardo* (2015)
- [16] Manning, C.D., Surdeanu, M., Bauer, J., Finkel, J.R., Bethard, S., McClosky, D.: The stanford corenlp natural language processing toolkit. In: *ACL (System Demonstrations)*. pp. 55–60 (2014)
- [17] Nowak, E., Jurie, F., Triggs, B.: *Sampling Strategies for Bag-of-Features Image Classification*. pp. 490–503. Springer Berlin Heidelberg (2006)
- [18] Pease, A., Colton, S.: On impact and evaluation in computational creativity: A discussion of the Turing test and an alternative proposal. *AISB 2011: Computing and Philosophy* pp. 15–22 (2011)
- [19] Pease, A., Winterstein, D., Colton, S.: Evaluating machine creativity. In: *Workshop on Creative Systems, 4th* (2001)
- [20] Peinado, F., Gervás, P.: Evaluation of Automatic Generation of Basic Stories. *New Generation Computing, Computational Paradigms and Computational Intelligence*. Special issue: *Computational Creativity* 24(3), 289–302 (2006)
- [21] Pereira, F.C., Hervás, R., Gervás, P., Cardoso, A.: A Multiagent Text Generator with Simple Rhetorical Abilities. In: *Proc. of the AAAI-06 Workshop on Computational Aesthetics: AI Approaches to Beauty and Happiness, July 2006*. AAAI Press (2006)
- [22] Pérez, R.y., Ortiz, O., Luna, W., Negrete, S.: A system for evaluating novelty in computer generated narratives. *Creativity* (2011)
- [23] Pérez y Pérez, R.: *MEXICA: A Computer Model of Creativity in Writing*. Ph.D. thesis, The University of Sussex (1999)
- [24] Ritchie, G.: Assessing creativity. In: *Proceedings of the AISB Symposium on AI and Creativity in Arts and Science*. pp. 3–11. York, UK

- [25] Ritchie, G.: Some Empirical Criteria for Attributing Creativity to a Computer Program. *Minds & Machines* 17, 67–99 (2007)
- [26] Robnik-Šikonja, M., Kononenko, I.: An adaptation of relief for attribute estimation in regression. In: *Machine Learning: Proceedings of the Fourteenth International Conference (ICML 1997)*. pp. 296–304 (1997)
- [27] Ware, S.G., Young, R.M.: Validating a Plan-Based Model of Narrative Conflict. In: *Proceedings of the International Conference on the Foundations of Digital Games*. pp. 220–227. ACM Press, New York, New York, USA (2012)
- [28] Ware, S.G., Young, R.M., Harrison, B., Roberts, D.L.: Four Quantitative Metrics Describing Narrative Conflict.pdf. pp. 18–29. Springer Berlin Heidelberg (2012)
- [29] Wiggins, G.: A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19(7) (2006)
- [30] Wiggins, G.: Searching for Computational Creativity. *New Generation Computing, Computational Paradigms and Computational Intelligence*. Special Issue: Computational Creativity 24(3), 209–222 (2006)

Chapter 4

Author Contribution and Dynamics in Large-scale, Long-term Conditions

This chapter seeks to generalize findings from our previous research, generally conducted in controlled lab environments, in scenarios that reflect better the real world. The two sections of the chapter are composed by the two following publications:

Alan Tapscott, Joaquim Colás, Valeria Righi, Carlos León, and Josep Blat. We Built Our Own Worlds - Story Canonicity and Indirect Collaboration in a Shared Story World. In *Proceedings of COLLA 2016 the Sixth International Conference on Advanced Collaborative Networks, Systems and Applications*, pages 30–38, Barcelona, Spain, 2016. ThinkMind Alan Tapscott, Joaquim Colás, Valeria Righi, Carlos

León, and Josep Blat. Large-scale collaborative story worlds: Formalizing content and author dynamics. "Unpublished manuscript", 2017

The first section describes a longitudinal study a larger subject pool (22 groups made of 3 or 4 volunteer students) collaborating in a shared story world using a prototype tool over three weeks. Most of the similar studies we found [32, 18, 41, 15] are mostly focused on pedagogy and the effect of these creative exercises on the students themselves. We observe how participating authors contribute to a common fictional reality, establishing a collaborative story world. The tool, called *Chronoverse*, allows users to create original scenes in a public timeline and link them with each other using simple tags (sim-

ilar to those often found in internet pages). The results were used for two additional evaluation phases. The first evaluation involved a three person jury and rated each group's contribution in terms of plot, tone, consistency and coherence. The second evaluation involved 40 volunteers that compared the each group's contributions with each other.

The overall results suggest that contributing authors adhere strongly to the scenes that were present before the experiment. These scenes were included to establish an initial background, and were based in the most common contributions from previous courses. Subjects also favor one single canonic interpretation of events, and avoid contradictions (generally by introducing their own, complementary, plot elements). Also, semantic links between contributions -implemented as public, inclusive tags- and the timeline visualization used as a navigational mechanism were well received and extensively used. Readers rated the contributions considered less consistent and coherent by our jury poorly. Our conclusions are similar to the ones from previous research [43]. In the context of a shared story world, the contribution order is essential in establishing the overall perceived consistency and coherence. Authors prefer to avoid conflicting with established plot by contributing in distinct, divergent, directions yet readers prefer the most intertwined stories.

In the second section of the chapter, we explore popular, shared fictional story worlds from the internet. First we research the history of creative writing and fandom communities. Then, we use the previous framing to introduce the *hybrid collaborative narrative and story world sites*. These sites and their communities have a twofold goal: to document a story world and to write tales that take place in it. After a brief overview of the most popular hybrid sites, we conduct a more through analysis of popular and recent hybrid site-the SCP Foundation- to learn more of its author dynamics and contributions. Finally, we attempt to crystallize all of the findings into a model, the Open Story World.

Amongst other observations, we found that these hybrid sites seem to evolve mirroring the evolution of the internet, involving more the authors over time, including editorial and content structuring tasks. We also saw a strong tendency to establish a central element type for the encyclopedic content (e.g., solar systems, anomalies and fictional manuscripts). The theme of these sites is often mysterious, introducing tropes from urban legends or fantasy that reinforce reader engagement while reducing the need for consistency and coherence. The most relevant observation of the SCP Foundation is the introduction

of a flexible canon model, implemented via simple tags. Most contributions are linked through tags that involve some sort of semantic or plot connection. These *collaboration dimensions*, with a heightened internal consistency and coherence (mainly groups of interest and sub-canons) might overlap or simply ignore each other. Authors can acknowledge or ignore other existing contributions as they please while remaining consistent (by using referencing recurrent encyclopedic content) and coherent (by agreeing to the facts established by the encyclopedic content they reference). This is a similar observation to how our volunteers linked their contributions to the initial background content through tags while ignoring each other's contributions [45]. Also, after analysing the tags of the SCP Foundation we established a content categorization for collaborations that includes encyclopedic, narrative, index and community content. The connections of these content subsets suggest that indeed encyclopedic content is the most referenced and navigational, or index content is used as a nexus for accumulating references. We also discuss how the characters, often considered central in the literature and apparently absent from the SCP encyclopedic segment, are replaced by other kinds of encyclopedic content (supernatural anomalies performing with agency) and public inclusive factions (groups of interest, including characters). Finally, the OSW we introduce is a first attempt to provide a formal definition of similar hybrid sites. It categorizes the content according to its function, supports multiple, overlapping author collaboration dimensions and measures the plot integrity (hypothetically linked to consistency and coherence) of contributions, authors and collaboration dimensions.

We Built Our Own Worlds - Story Canonicity and Indirect Collaboration in a Shared Story World

*Alan Tapscott¹, Joaquim Colás², Valeria Righ², Carlos León¹ and
Josep Blat²*

¹*Facultad de Informática, Universidad Complutense de Madrid*

²*Grup de Tecnologies Interactives, Universitat Pompeu Fabra*

Abstract

In this work, we conducted a longitudinal study to understand better how authors collaborate when building shared story worlds. To accomplish this goal we deployed Chronoverse, a tool specifically designed for this purpose that provides authors with a common story timeline visualization and character faction tags. The study had three distinct phases. In the first phase, undergraduate students (the authors we study) used Chronoverse to develop their stories in a common context during a lab that lasted four weeks. In the second phase, a jury rated the authors' contributions in terms of coherence and consistency, attempting to measure the integrity of the shared story world. In the third phase, a larger crowd of readers rated the stories according to their preferences and shared their opinion of what stories belong to a common overarching story. The results suggested that the initial story, introduced by the researchers, was given priority and considered more canonical or "official" by the authors, jury and readers. Author groups did not reference each other's contributions directly, but achieved consistent and coherent results (according to jury measures) indirectly by adhering to the initial story in terms of plot and tone. The usage of tags in the design of Chronoverse was positively received by authors and enhanced the plot and tone consistency perceived by readers.

1 Introduction

Collaborative writing is a challenging task. However, the potential benefits of collaboration are motivating enough to encourage studies that seek to understand and improve the co-authoring of all sorts of texts. Writing fiction collaboratively using a shared story world represents a specific kind of challenge. Authors must share their conceptions

of the story world's contents to write consistently, maintaining the impression that all media belongs to the same continuity. For the most part, story worlds do not exist explicitly; they emerge from the facts established by the author or authors in the story. Large fiction franchises, often involving multiple authors, solve this through the usage of internal documents that keep track of the plot developments, documenting the past, present and even future of the fiction, along with relevant character biographies and other fictional encyclopedic content. For instance, TV series often rely on a confidential bible maintained and used by all screenwriters as reference to collaborate consistently. This seeks to avoid alienating the audience with inconsistencies and incoherences. While this kind of solution has been successfully used in scenarios with over a dozen authors, there is a clear challenge in a hypothetical large-scale online, crowdsourced scenario.

In creative writing, authorial style and artistic vision are important, as well as the capability to build interesting fictional worlds, especially for fantastic genres. Facts contained in the narrated story deviate from our reality up to some degree, presenting landscapes and characters that only exist in our imagination. For the current study, we consider the story world the set of elements (e.g., characters, locations, plots, motivations, rules) that constitute the world implicitly defined by a story along with the relationships between each other. According to Tolkien and his conception of secondary world [1] (a definition we believe to be close to our story world), all the elements (i.e., geography, characters, language and timeline) are interdependent and require internal consistency to suspend disbelief, becoming credible to the reader. Schmidt and Bannon have an extensive publication track that introduces the Common Information Space [2][3], defined as "...a central archive of organizational information with some level of 'shared' agreement as to the meaning of this information (locally constructed), despite the marked differences concerning the origins and context of these information items." We believe that this definition is conceptually aligned with our conception of a shared story world in the sense that it contains relevant data introduced by distinct authors and despite being potentially different, it must agree or be coherent to some degree. Bannon and Bødker discuss the dialectical nature of common information spaces and the challenge of putting information in common and interpreting it [4]. Despite presenting distinct scenarios, most of the considerations for common information spaces are very likely to appear in a shared story world and must be addressed. Mainly, Bannon and Bødker insist on the need for a common information space to be accessible and malleable while providing reliable

information. Most of the relevant literature focuses either on the technical implementation of concurrent writing systems or on the narratological study of story worlds. Despite the emerging relevance of fiction transmedia story worlds, both in commercial and amateur contexts, to our best knowledge there is no formal study focused on collaborative authoring and story worlds published. This work aims to understand better the dynamics of collaborating authors when contributing to a shared story world, using computer-assisted collaboration to monitor and measure relevant aspects, such as cross-referencing, coherence and plot/tone consistency.

The study is structured in the following way. In Section 2, we provide a brief discussion of works that deal with relevant problems. Next, in Section 3, we describe our three months' longitudinal study where lab students contributed into a shared story world using Chronoverse, a prototype tool we designed and deployed to support the co-authoring of shared story worlds. Also in Section 3 we describe the jury and reader evaluation of the results from the longitudinal study. We present the results of the evaluation in terms of consistency, coherence and co-existence of an overarching story. Finally, in Section 4, we discuss the implications of the results and conclude with the insight learned from the whole study, along with potential directions for hypothetical subsequent research.

2 Related Work

There have been many works that have studied the usage of computer-based platforms to support collaboration. ShrEdit by Olson et al. [5] was a shared collaborative text editor meant to aid designers in brainstorming ideas. According to the authors, groups that used the editor produced less ideas ranked as more creative. This might imply that in a similar context, the usage of a digital platform for collaboration might prioritize quality over quantity. Posner and Baecker [6] present a taxonomy based on interviews that describes joint writing in the following terms: roles played in the collaboration, activities performed in the writing process, document control methods used, and writing strategies employed. This was later expanded by Lowry et al. [7] with collaborative writing activities, document control modes, roles, tools and work modes, a categorization that could be helpful in designing an adapted user experience. Google Docs, the popular collaborative word processor, has been studied in some works from the educational perspective [8], the longitudinal perspective of a large, diverse organization [9] and a tool to write educational papers [10] amongst others.

Overall the design of Google Docs seems very appropriate for collaborative work in a computer platform.

Some other published works deal with similar scenarios such as Robinson's exploration of collaboration in authoring multimedia stories through specific devices [11]. Krowne and Bazaz discuss authority and territoriality in a study of collaborative editing systems [12], this is especially relevant in collaborative systems that deal with creativity and authoring. In the context of a shared story world and multiple collaborating authors, territoriality could be a very important factor to take into account. Likarish and Winet attempted to reproduce the surrealist *Exquisite Corpse* writing game on Twitter [13], attempting to "... understand the practical pitfalls of synchronous community-based authorship and to recommend methods of avoiding them." Besides reporting successful participation, Likarish concludes by acknowledging the need for "...providing structure via a wiki or suit of tools to enable authors to track details as well as the importance of community self-policing..." We share the belief for tools that introduce structure into the creative process. Thomas and Mason bravely attempted to write a novel in an open process [14] using a wiki platform. The wiki proved to be a competent and useful tool for structuring the narrative thanks to its familiar and accessible nature. According to the authors inter-author collaboration dynamics were challenging, citing content deletions or major restructuring performed by a single author as the source of conflicts. Among the other relevant remarks from the authors, there seems to be a dichotomy between contribution order and creativity, highlighting the importance of contribution sequence and timing when building this kind of systems.

Relevant works that seek to explore collaboration have resorted to the usage of timelines to provide adequate user experiences. Thiry et al. use a timeline in Project Greenwich [15], a tool meant for people to author their own personal digital timelines. In their work, Thiry et al. study the usage of the timeline as a vehicle that helps multiple authors in connecting past and present contributions. The capability to collaborate over time seems to be especially desirable in the context of an evolving shared story world. Some other interesting usages of a timeline include a programming interface by Cardoso et al. [16], medical records overviews by Reddy and Dourish [17], an adaptive timeline interface to personal history data by Ajanki et al. [18] and as an aid for history learning by Pyshkin and Bogdanov [19].

3 Study

The following section describes the longitudinal study, including its motivation, the prototype used and the methodology and results for each of its phases.

3.1 Context and motivation

Since the academic course of 2012-2013, Computer Science and Audiovisual Engineering undergraduates have attended the Audiovisual Language and Interactive Storytelling subject at UPF, participating in the lab and creating their own visual novels. Participants are distributed in groups of three to four students. Over two months, each of the groups creates a visual novel, including its design, script and implementation. The teachers were surprised by the convergence of the plots and tones of the stories, even though they never encouraged collaboration. Despite the lack of hard evidence to support the claim, informal observations suggested some students were collaborating indirectly, using common elements in their stories. Social and cultural trends had an impact in the creative process, as well as the university environment. For instance, every year featured multiple post-apocalyptic stories (a popular recurring theme in current popular fiction) and roughly half of the plots happened in the student's university (UPF). This observation served as motivation for the deployment of Chronoverse, a digital collaborative tool meant to reinforce the collaborative dynamics and the consistency of their collaborations.

3.2 Chronoverse

Our prototype, Chronoverse, was an online tool meant to help authors in collaborating to produce media based in a shared story world.



Fig. 1: Chronoverse interface

Chronoverse allowed authorized users to introduce their own original scenes. As seen in Fig. 1, every scene contained a title, a short

description, a date, an optional picture and optional tags describing the involved factions or groups of interest. On the bottom of the screen users could see a timeline with all the existing scenes. The usage of a timeline metaphor for the chronology of all scenes was meant to reinforce, for every contribution, the sense of belonging to the same continuity. Since all scenes had a date and were rendered inside the same timeline, we expected to provide an enhanced sense of coexistence in the same story world for all scenes. Also, as seen in the previous section, timelines are useful to connect past and present resources [15]. We made authors split their stories into scenes to promote the intertwining of plotlines. By providing a common canvas for contributions that took place inside the shared story world, with its content visible to all users, we expected to promote inter-author story awareness.

Chronoverse's timeline stacked scenes vertically, in up to three lanes, allowing scenes to be placed close in the horizontal axis without much overlapping. Also the splitting of the story into scenes was favorable for the timeline, allowing users to perceive the story progression over time. The provided initial scenes in Chronoverse (described more in depth later) included factions. We chose to use factions as the main actors of the story instead of characters to avoid the problem of author territoriality mentioned in the previous section [12], preventing conflicts and promoting indirect collaboration. Contributors could create their own original characters and integrate them into the story world more easily via faction membership without necessarily conflicting with other authors and their characters. Other potential organizations (e.g., nations, clubs, nobiliary houses) are perfectly valid to achieve the same results. Users could filter the timeline to see scenes that only involve a specific faction, providing a navigational mechanism closely linked to the story world. This mechanism helps in establishing thematic links between scenes created by distinct authors who are not necessarily collaborating directly. We also introduced a pop-up that displays the relationship between the distinct initial factions in a graph.

3.3 Study structure and settings

As in the previous years, the 2014-2015 Audiovisual Language and Interactive Narrative subject lab took place. We decided to deploy Chronoverse and measure how did it impact on the results. The following phases were planned:

1. *Authoring phase*: Conduct the lab with the author students,

having every group introduce the story plan for their visual novel into Chronoverse over four weeks. Conduct a short questionnaire to rank the usage of Chronoverse, according to its users.

2. *Jury evaluation phase*: A small jury ranks the contributions, rating coherence and plot/tone consistency. Comparisons include each contribution against each other and each contribution against the initial story.
3. *Reader evaluation phase*: A larger crowd reads all the stories and ranks their preferred ones. In order to validate the scores from the jury evaluation, they are also guess what stories belong to a common overarching story.

3.4 Authoring phase

We provided access to Chronoverse to all student groups and asked them to introduce the outline of their stories in it. The goal was to have every group plan their visual novel's plot before writing the final, extended script, encouraging collaboration and convergence. Since the platform was hosted in a public web server, every group could see other groups' work and potentially reference each other directly or indirectly. The resulting story plans, described by the scenes introduced into Chronoverse, would then be adapted into scripts for the visual novels. Ideally, this would allow authors to develop their story in a common frame and influence each other before being tied to the cumbersome work of developing a whole script.

Chronoverse had six initial story scenes (ISS) created by the teachers. The scenes told a single post-apocalyptic story in the university where several factions struggled for power. The first scene explained very vaguely the downfall of civilization for an unknown reason. Each of the following five weeks contained one scene narrating how each of the three fictional student factions attempted to rule the university until a tense and unstable peace is reached in the end. The results section details how the initial story and factions influenced author contributions. The story contained some of the most frequent elements from previous years' contributions, such as a post-apocalyptic world, the students' university, politics and mysterious factions. This story world was very open and purposefully vague, providing flexibility to allow the groups to fully develop their own ideas.

Overall, 22 groups composed by between three and four undergraduates participated. Each group received a textual description of a story world or setting in a brief text document that described the main

characters, factions, events and also included some mysteries and “plot hooks” that could be freely developed or used at the authors’ discretion. Groups were asked to create at least three scenes in the timeline that represented their visual novel’s plot. For every weekly lab and during three weeks, the teacher told them to add at least one scene. This was done to provide groups with enough time to read existing contributions. Reading or even referencing existing material was always optional. No amount of contribution or collaboration was enforced. Three months after the lab sessions ended and each group had produced their Chronoverse scenes along with the final visual novel, we asked the author students to take a short questionnaire meant to get some feedback on their user experience. Specifically, we were interested in learning about their opinion on the usage of a default story world and the Chronoverse platform.

3.5 Authoring phase results




Fig. 2: Author group scene and tag amount




Fig. 3: Total word count per group

Table I, Fig. 2 and Fig. 3 show an overview of the results. Groups contributed an average of 2,8 scenes, 2 tags and 283,7 total words. Half of the groups did not use any tag at all and there were a couple of significant outliers (one groups used 9 tags and another group introduced 9 tags). Fig. 4 shows when author groups introduced scenes into Chronoverse. All contributions were introduced between the 7th




Fig. 4: Total scene contributions by date

and the 22nd of June. The main contribution peaks are around the lab deadlines (8th and 14th of June).

3.6 Author questionnaires results

Due to the optional nature of the questionnaire only 17 of the 84 participating authors submitted answers. Results, rated from 1 to 5, are summarized in Tables II and III. Subjects found easy to read and write ($\bar{x}=3,87$ with $\sigma=0,83$ and $\bar{x}=3,59$ with $\sigma=1,18$ respectively). The provided initial story world (scenes and deliverable text document) scored as moderately useful to write $\bar{x}=3,65$ with $\sigma=0,70$, followed by images $\bar{x}=3,53$ with $\sigma=0,68$, and finally dates $\bar{x}=2,76$ with $\sigma=1,18$. Explicit dates are the less useful and less enjoyed part of the initial story. Some of the open answers provided corroborate this. Amongst all the results from the questionnaire we found a few worth mentioning. 88% of the subjects read some scenes and 12% none. Most of the subjects explicitly were not bothered by contradictions and state they did not influence their writing at all. 94% of the users claimed to have used the initial scenes provided in Chronoverse as inspiration for their contribution.

3.7 Jury evaluation phase

Once the lab was finished and we gathered all the contributions, we were interested in evaluating the results in terms of consistency and coherence. These measures were meant to determine the integrity of the set of stories as a story world. In order to rate the stories, we created a jury made of three members. Two were teachers who actively designed and supervised the labs and the third was a Ph.D. student who was unfamiliar with the subject or the lab. We were especially interested in determining the relationship between each group's contributions and the initial scenes. Each groups' Chronoverse scenes (previously described) were joined to create an author group scene

set (AGSS). The initial scenes created by us were joined into an initial scene set (ISS). The jury ranked every AGSS after the following measures:

- AGSS unitary coherence (is the scene set contributed by the author group coherent?)
- AGSS unitary image coherence (are the images used coherent?)
- AGSS tonal consistency with the ISS (is the author group scene set tonally consistent with the initial scene set from Chronoverse?)
- AGSS plot consistency with the ISS (is the author group scene set consistent with the initial scene set from Chronoverse in terms of plot?)
- AGSS tonal consistency with other AGSSs (is the author group scene set tonally consistent with other authors' scene sets?)
- AGSS plot consistency with other AGSSs (is the author group scene set consistent with other authors' scene sets in terms of plot?)

Coherence of a text referred to the internal logic of its discourse and image. Consistency referred to the text's similarity, in terms of narrative plot and tone, to another text. To ensure the jury evaluation criteria was unified, the measures were discussed informally and we conducted a pre-evaluation with some random Chronoverse contributions. We found no significant differences on the pre-evaluations. Despite the fact that a member of the jury who was not involved in the lab experiment, the criteria for their ranking apparently was uniform. For the main jury evaluation, instead of comparing every AGSS to the rest of the set, each AGSS was compared to the ISS and 6 random AGSSs more. This cut was necessary to reduce the cost of the jury evaluation.

3.8 Jury evaluation phase results

Results, with the factors evaluated in a 1 to 4 scale, are summarized in Table IV. AGSS unitary plot coherence was somewhat high (average 3,02) while AGSS images were considered coherent with the story created by authors (average 3,45). AGSSs were moderately consistent with the ISS in terms of plot and tone (average 2,55 and 2,89).

Inter-AGSS tonal consistency was also moderate (2,44/4) while inter-AGSS plot consistency was low (1,86/4). Next, we run a Pearson correlation analysis on the measures. Specifically, we wanted to find potential AGSS inter-relationships and between the 22 AGSSs and the ISS. The jury's measures of consistency should provide some insight on the collaboration dynamics of the participating groups. The results can be seen in Table V. There is not a significant correlation between AGSS-ISS plot consistency and unitary AGSS coherence, suggesting that adhering to the initial story world elements did not lead to either more or less coherent stories. AGSS unitary coherence was usually high, independently of their consistency with the ISS. There is a significant correlation ($r(N) = 0.63, \rho(N) = 0.002, r^2(N) = 0.4$) between AGSS-ISS plot consistency and inter-AGSS average plot consistency. The same happens with AGSS-ISS tone consistency and inter-AGSS average tone consistency ($r(N) = 0.62, \rho(N) = 0.002, r^2(N) = 0.4$). So groups trying to remain consistent with the initial story world in terms of plot and tone scored also high tonal and plot consistency with other author group contributions. It seems that there is also a strong correlation ($r(N) = 0.66, \rho(N) = 0.0008, r^2(N) = 0.4$) between AGSS-ISS plot consistency and AGSS-ISS tone consistency. Therefore, according to the jury, tonal and plot consistency with the ISS seem to imply each other to some degree. An interesting observation by our jury suggested that AGSS contributions consistent with other AGSSs referenced factions, places and events provided by the ISS, but never elements created by other author groups. AGSSs were ranked as consistent with each other due to their usage of initial ISS elements, not new AGSS elements introduced by other author contributions.

3.9 Reader evaluation phase

We conducted an evaluation to determine each scene perceived co-existence to the same overarching story and overall preference, according to external readers. We displayed the contributions in pairs and asked whether or not each pair of displayed texts belonged to the same story. We expected this to reflect the integrity of a hypothetical story world described by the crossover of both stories, hypothetically implying the perceived consistency and coherence. Text pairings marked as belonging to the same story would argue in favor of the contributing authors collaborating in a shared story world. 40 volunteers participated in an online questionnaire, with ages ranging from 18 to 63 and a gender distribution of 43% females and 57% males. The volunteers were not author students from the first phase or juries from the sec-

ond phase. The questionnaire was a simple website we developed using basic HTML and a Django backend. The questionnaire presented 10 pairs of texts, displaying only two at a time. To reduce the cost of the evaluation, pairings included i) the three AGSS that the jury ranked as more consistent with the initial story, ii) the three AGSS ranked as the less consistent with the initial story, and iii) two AGSS set in the middle range. The initial scene set (ISS) was also added to the questionnaire. Each volunteer was asked to rank the stories from 1 to 5 on a Likert scale where 1 meant “I do not like the story” and 5 was “I like the story a lot”. We then asked them if paired texts belonged to the same story (only yes or no, closed reply). The simple web questionnaire can be seen in Fig. 5.

Fuga de cerebros

Desde un aula de programación, un profesor se desmaya. Al año de más tiene un diagnóstico oncológico, de saber muy bien que fallecerá. Cuando empieza la clase se le cae el papeo en la cara. Se pone a reír y luego piensa por la actividad durante todo el día. Una vez montado todos los papeos de la clase se pone a programar. El profesor y el curso quedan con el alma agridulce.

Desde luego los profesores no saben muy bien que les pasará, y se olgan primero a la segunda Práctica y posteriormente a Plus Catalán. A medida que se recupera la memoria se va dando cuenta de que la mayoría de cosas que le han pasado y lo que le va a pasar, y se siente responsable por el profesor y por la clase. Pero cuando empieza en su primer y último día de clase del curso se da cuenta de que son días increíbles.

No se gana nada No gana nada

El Rumor

El año pasado de la segunda práctica, aquella donde colapsaron los sistemas informáticos, aunque otros sistemas no hubieran estado, nuestro protagonista Cati es un estudiante de la UPEL que decide no irse entre su profesor Práctico que aplica la empresa de Cati, pero que accidentalmente acaba su práctica de manera.

Una vez que el profesor Práctico, Cati decide irse al profesor Práctico, y se va a ir a la empresa de Cati.

Cati empieza la carrera de la UPEL en un momento de la práctica. Luego le da el profesor Práctico, pero cuando empieza a trabajar en la empresa de Cati, se da cuenta de que el profesor Práctico, Cati se esfuerza al máximo en su año universitario, al máximo.

No se gana nada No gana nada

Ambos textos pertenecen a la misma historia Los textos pertenecen a historias diferentes Muéstrame otro par

Fig. 5: Online reader questionnaire

3.10 Reader evaluation phase results

Table VII shows how what stories were marked by readers as belonging to a common story. AGSSs marked by readers as belonging to the same story than the ISS tended to be also marked as belonging to a common story among each other. This result seems to be in line with our jury evaluation, where AGSSs that were consistent with the ISS were also consistent with each other and once again suggests indirect collaboration.

3.11 Global results comparison

After conducting the three phases of our study, we can compare each of the AGSSs metrics to each other. This should allow us to validate the distinct evaluations by finding correlations. Table VII summarizes the whole study and contains all the relevant comparisons with the resulting Pearson coefficient. Pearson correlations revealed a moderate correlation between the AGSSs previously ranked by the jury as highly consistent (both, in terms of plot and tone) with the ISS and the tendency to be marked by readers as belonging to the same story

than the ISS ($r(N) = 0.52, \rho(N) = 0.01, r^2(N) = 0.27$ for both cases). This suggests the jury's consistency measures and the reader similarity measures are relatively aligned. There is also a moderate correlation between reader score of the AGGs and AGG-ISS plot consistency and AGG-ISS tone consistency (respectively $r(N)_{rv-pc} = 0.61, \rho(N)_{rv-pc} = 0.002, r^2(N)_{rv-pc} = 0.37$ and $r(N)_{rv-tc} = 0.47, \rho(N)_{rv-tc} = 0.02, r^2(N)_{rv-tc} = 0.22$). Apparently readers liked more stories consistent with the initial story world in terms of plot and tone. There is a moderate correlation between reader score of the AGGs and tonal consistency with other AGGs ($r(N) = 0.53, \rho(N) = 0.01, r^2(N) = 0.28$). Being tonally consistent with other author contributions was liked by readers (according to the provided scores). Unlike the previous case, this is not extensible to plot consistency. There is a moderate-high correlation between the usage of tags and reader score ($r(N) = 0.54, \rho(N) = 0.001, r^2(N) = 0.29$) and a high correlation between the reader score and the tendency to be marked by readers as belonging to the same story than the ISS ($r(N) = 0.82, \rho(N) = 3e-6, r^2(N) = 0.67$). Also, AGSSs that used more tags revealed high correlations with plot consistency and tone consistency with other AGSSs (respectively $r(N)_{atpc-apc} = 0.78, \rho(N)_{atpc-apc} = 3e-5, r^2(N)_{atpc-apc} = 0.6$ and $r(N)_{attc-atc} = 0.65, \rho(N)_{attc-atc} = 0.001, r^2(N)_{attc-atc} = 0.42$) and the ISS (respectively $r(N)_{atpc-ipc} = 0.92, \rho(N)_{atpc-ipc} = 0, r^2(N)_{atpc-ipc} = 0.84$ and $r(N)_{attc-itc} = 0.77, \rho(N)_{attc-itc} = 2e-5, r^2(N)_{attc-itc} = 0.59$). According to the jury, people contributing closer to other scenes in terms of plot and tone felt more inclined to label their usage of existing character factions with tags.

4 Discussion

Our goal was to understand better the dynamics of collaborating authors when contributing to a shared story world. Overall, despite giving freedom to our authors, we believe there have been two main outcomes. Some author groups (the majority) have followed and extended the initial story, while a few others have ignored completely the initial stories and their author colleagues. This means that there is collaboration happening, but it is not the collaboration we were expecting. Author groups that have followed the initial story collaborate indirectly with other author groups that have done the same, creating consistent stories with each other in terms of plot and tone. According to our readers, consistent contributions to Chronoverse are enjoyable and belong to a common shared story. Results have been extensively documented in the previous sections and annex, however there are a few key aspects we believe worth discussing in depth.

4.1 One story canon

Our main finding implies readers (and perhaps writers) have one single canonical or “official” version of facts. Most authors contributed consistently with the initial story and most readers marked stories consistent to the initial content as belonging to the same overarching story. Overall, it seems the initial scenes were considered more “official” or canonical than those introduced by the contributing author groups. Contributors seem to embrace the notion of an initial explicit story world and integrate its plot hooks and elements into their own creations. Our authors did not read many of the other authors’ contributions and they never explicitly integrated them into their own contributions. The tone, however, was considered consistent between all authors’ contributions. Perhaps the initial content, being written by the teachers, was given a special consideration. Maybe timing is the key factor, contributing authors might not feel confident or comfortable referencing and intertwining their content with the unfinished and ongoing contributions of other authors. The sequence in which the contributions occur might be key to determine what parts of the shared story world are perceived as more central or canonical.

This might explain the problems faced by Thomas and Mason [14] in their Wikinovel, with authors colliding as they attempt to impose a main plot structure. When compared to our own research it seems this phenomenon is similar to previous results [20] in which authors converged, establishing a main continuity or central interpretation of the story. The joint conception of a story world might require some information hierarchy, providing new contributors some solid narrative background or baseline from which to start. Information canonicity should be central in any further study of this nature, that is, finding mechanism for authors not only to contribute to a common story world, but also to ensure those ideas are well integrated and accessible. This should make possible that subsequent contributions by other authors are perceived as consistent, something positive for a shared story world according to our own results and by other works on collaborative information spaces [3].

4.2 Tags and inclusive character groups

The usage of character group or faction tags seems to increase story world consistency. Stories that included tags seem to be more consistent with all the rest of stories in terms of plot and tone. Also stories that used tags have been ranked better by readers. Overall, tags seem

to be a good explicit mechanism to structure content in a shared story world. This somewhat implies a benefit in using inclusive groups or categories with an active role in the story contained in a story world (such as the ones found in Chronoverse). Our goal was to avoid author territoriality [12] or the lack of structure in a collaborative environment described by Likarish [13]. This kind of collaborative meta-data, often found in collaborative platforms such as wikis, might be beneficial for integrating original characters from multiple authors while enhancing the global consistency.

4.3 The usage of Chronoverse

The author questionnaire informal results seem to point towards a positive user experience. The usage of a timeline seems to be favorable for the construction of a shared story world, similarly to some related works [21]. Regarding other author contributions, authors did not read many other group's contributions before writing and did not care about contradictions. Coping with other authors' stories is apparently not a priority. We believe the authors were fairly motivated students, but this is not necessarily the average user in a shared story world-building scenario. We suspect that more proficient users, such as professional writers, might be more inclined to contribute with consistent and coherent contents. On the other hand, students might be more open to novel scenarios, such as collaborating in an online platform to build a story world together. Still, we believe replicating the experiment with a professional audience would be paramount to generalize the results to a more common scenario, such as a team of scriptwriters writing a TV show season or a story anthology.

5 Conclusions and Future Work

Our main conclusion from this study is that most participants across all evaluation phases perceived one main continuity for the story world, and rated consistency and coherence around it. Also, this main continuity is mainly established by the pre-existing material initially introduced into the story world. Contributions are more likely to reference and connect to pre-existing content than to ongoing contributions by other authors. In a shared story world, the sequence of the contributions seems to be critical in establishing information canonicity. Older contributions are more referenced than newer less established ones. Our results and observations also point towards the aptness of timeline visualizations, scene tags and inclusive character factions

for collaborative story world building scenarios. There are two main directions for this research to continue, generalizing its findings and extending its applications. On one hand, we are already trying to replicate these observations in a large-scale real scenario, a successful online community that builds and maintains a rich story world. On the other hand, these findings could be used to build a cognitive model or architecture of a collaborative story world.

Acknowledgment

Supported by WHIM (611560) and PROSECCO (600653), funded by the European Commission, Framework Program 7, the ICT theme, and the Future Emerging Technologies FET program. This research is partly funded by the Spanish Ministry of Economy and Competitiveness under RESET (TIN2014-53199-C3-3-R) and the Maria de Maeztu Units of Excellence Programme (MDM-2015- 0502). This work would not have been possible without the invaluable collaboration of the students from the Audiovisual Language and Interactive Storytelling lab, the volunteer reader crowd and the Ph.D. student jury collaborator.

References

- [1] J. R. R. Tolkien, “Beowulf: The Monsters and the Critics,” *Interpretations of Beowulf: A Critical Anthology*, p. 31, 1991.
- [2] L. J. Bannon and K. Schmidt, “CSCW: Four characters in search of a context,” *Studies in Computer-Supported Cooperative Work*, pp. 3–16, 1991.
- [3] K. Schmidt and L. J. Bannon, “Taking CSCW seriously - Supporting articulation work,” *Computer Supported Cooperative Work*, vol. 1, no. 1-2, pp. 7–40, 1992.
- [4] L. J. Bannon and S. Bødker, “Constructing Common Information Spaces,” *Proceedings of the Fifth European Conference on Computer Supported Cooperative Work*, pp. 81–96, 1997.
- [5] J. S. Olson, G. M. Olson, M. Storøsten, and M. Carter, “Groupwork close up: a comparison of the group design process with and without a simple group editor,” *ACM Transactions on Information Systems*, vol. 11, no. 4, pp. 321–348, 1993.
- [6] I. R. Posner and R. M. Baecker, “How People Write Together,” *Proceedings of the Twenty-fifth Annual Hawaii International Conference on System Sciences*, vol. 4, pp. 127–138, 1992.
- [7] P. B. Lowry, A. Curtis, and M. R. Lowry, “Building a taxonomy and nomenclature of collaborative writing to improve interdisciplinary research and practice,” *Journal of Business Communication*, vol. 41, no. 1, pp. 66–99, 2004.
- [8] I. Blau and A. Caspi, “What Type of Collaboration Helps? Psychological Ownership, Perceived Learning and Outcome Quality of Collaboration Using Google Docs,” *Quality*, pp. 48–55, 2009.
- [9] Y. Sun, D. Lambert, M. Uchida, and N. Remy, “Collaboration in the cloud at Google,” *Proceedings of the 2014 ACM conference on Web science - WebSci '14*, pp. 239–240, 2014.
- [10] R. Watson, “Extending Google Docs to Collaborate on Research Papers,” *The University of Southern Queensland, Australia*, pp. 1–11, 2006.
- [11] S. Robinson, M. Jones, E. Vartiainen, and G. Marsden, “PicoTales: collabora-

- tive authoring of animated stories using handheld projectors,” *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work - CSCW '12*, p. 671, 2012.
- [12] A. Krowne and A. Bazaz, “Authority Models for Collaborative Authoring,” *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*, vol. 00, pp. 18–24, 2004.
- [13] P. Likarish and J. Winet, “Exquisite Corpse 2.0,” *Proceedings of the Designing Interactive Systems Conference*, pp. 564–567, 2012.
- [14] B. Mason and S. Thomas, “A Million Penguins Research Report,” Institute of Creative Technologies, De Montfort University, Leicester, United Kingdom, pp. 1–36, 2008.
- [15] E. Thiry, S. Lindley, R. Banks, and T. Regan, “Authoring personal histories: Exploring the timeline as a framework for meaning making,” *Proceedings of the 2013 SIGCHI conference on Human Factors in computing systems (CHI 2013)*, pp. 1619–1628, 2013.
- [16] B. Cardoso and T. Romão, “The timeline as a programming interface,” *Conference on Human Factors in Computing Systems*, pp. 1423–1428, 2014.
- [17] M. Reddy and P. Dourish, “A finger on the pulse: temporal rhythms and information seeking in medical work,” *Proceedings of the ACM Conference on Computer Supported Collaborative Work CSCW*, pp. 344–353, 2002.
- [18] A. Ajanki, M. Koskela, J. Laaksonen, and S. Kaski, “Adaptive timeline interface to personal history data,” *Proceedings of the 15th ACM on International conference on multimodal interaction*, pp. 229–236, 2013.
- [19] E. Pyshkin and N. Bogdanov, “Learning history with timelines: Use cases, requirements and design,” *Computer Science and Information Systems (FedCSIS), 2014 Federated Conference on*, pp. 273–280, 2014.
- [20] A. Tapscott, J. Colás, A. Moghnieh, and J. Blat, “Writing Consistent Stories based on Structured Multi-Authored Narrative Spaces,” *Proceedings of the 4th Workshop on Computational Models of Narrative*, vol. 32, pp. 277–292, 2013.
- [21] M. Bohøj, N. G. Borchorst, N. O. Bouvin, S. Bødker, and

P.-O. Zander, “Timeline collaboration,” *Proceedings of the 28th international conference on Human factors in computing systems CHI 10*, pp. 523–532, 2010.

Appendix - Data Tables

Tab. 1: AUTHORING PHASE DESCRIPTIVE STATISTICS

	Scenes	Tags Used	Total Word Count
Sample Size:	22	22	22
Mean:	2.81	2.09	283.77
Median	3	0	268
Minimum:	1	0	87
Maximum:	7	9	581

Tab. 2: AUTHOR QUESTIONNAIRE RANK QUESTION RESULTS

	Initial story was helpful in writing	Dates helped write	Images helped writing	Liked using dates	Liked using images	Reading from Chrono-verse was easy	Contributing to Chrono-verse was easy
AVG	3,65	2,47	3,53	2,76	3,71	3,87	3,59
STDEV	0,7	1,01	0,99	1,03	1,14	0,24	0,06

Tab. 3: DETAILED AUTHOR QUESTIONNAIRE PERCENTAGE QUESTION RESULTS

	1 to 5 existing scenes	5 to 10 existing scenes	No existing scene
I read...	53%	35%	12%
My reading order was...	randomly 33%	chronologically 27%	guided by title 33%
	before writing	after writing	during the whole experience 13%
I read existing scenes...	33%	53%	typical 38%
Initial story was...	fun 25%	inspiring 31%	didn't bothered I liked them
	were a minor annoyance 24%	I didn't care 35%	me at all 29%
Contradictions in the story...	made writing difficult 6%	didn't influence me 88%	inspired me 6%
Those contradictions...	based on existing story world 94%	from scratch	
Wrote our story...		6%	

Tab. 4: JURY EVALUATION PHASE RESULTS

Plot coherence	Image coherence	Plot consistency with other AGSSs	Tone consistency with other AGSSs	Plot consistency with ISS	Tone consistency with ISS
AVG 3,02	3,45	1,86	2,44	2,55	2,89
STDEV 0,87	0,69	0,47	0,56	0,95	0,69

Tab. 5: JURY EVALUATION PHASE RESULTS PEARSON CORRELATION COEFFICIENTS

	Plot coherence	Image coherence	Plot consistency with other AGSSs	Tone consistency with other AGSSs	Plot consistency with ISS	Tone consistency with ISS
Plot coherence	1	0,23	-0,01	-0,01	0,09	0,16
Image coherence	0,23	1	-0,16	0,16	-0,3	-0,18
Plot consistency with other AGSSs	-0,01	-0,16	1	0,23	0,63	0,33
Tone consistency with other AGSSs	-0,01	0,16	0,23	1	0,28	0,62
Plot consistency with ISS	0,09	-0,3	0,63	0,28	1	0,66
Tone consistency with ISS	0,16	-0,18	0,33	0,62	0,66	1

Tab. 6: READER EVALUATION PHASE AGSSs DESCRIPTIVE STATISTICS

	Reader Score	Same story than ISS?
Sample Size:	8	8
Mean:	3.10625	0.498125
Median:	3.08	0.52
Minimum:	2.75	0.08
Maximum:	3.4	875

Tab. 7: AGSS/ISS AVERAGE SIMILARITY

Same story?	AGSS1	AGSS2	AGSS3	AGSS4	AGSS5	AGSS6	AGSS7	AGSS8	ISS
AGSS1	1	0,4	1	0	0,22	0	0,4	0,42	0,71
AGSS2	0,4	1	0,3	0,29	0	0,14	0	0,33	0,08
AGSS3	1	0,3	1	0,36	0	0	0,67	0,4	0,88
AGSS4	0	0,29	0,36	1	0,38	0,13	0,6	0,2	0,33
AGSS5	0,22	0	0	0,38	1	0,33	0	0	0,14
AGSS6	0	0,14	0	0,13	0,33	1	0	0,14	0,15
AGSS7	0,4	0	0,67	0,6	0	0	1	1	0,83
AGSS8	0,42	0,33	0,4	0,2	0	0,14	1	1	0,88
ISS	0,71	0,08	0,88	0,33	0,14	0,15	0,83	0,88	1

Tab. 8: GLOBAL PEARSON CORRELATION COEFFICIENTS PER AUTHOR GROUP

	Reader evaluation phase			Jury evaluation phase					Authoring phase		
	Reader score	Same story than ISS	Plot coherence	Image coherence	Plot coherence	Tone consistency with other AGSSs	Plot coherence	Tone consistency with other AGSSs	Average word count	Total word count	Total tags
Reader Score	1										
Same story than ISS	0,27	1									
Plot coherence	-0,16	0,19	1								
Image Coherence	0,35	0,19	-0,08	1							
Plot consistency with other AGSSs	0,37	0,52	0,62	-0,01	1						
Tone consistency with other AGSSs	0,53	0,52	0,04	0,16	0,74	1					
Plot consistency with ISS	0,61	0,47	-0,05	-0,24	0,66	0,82	1				
Tone consistency with ISS	0,47	0,36	-0,03	-0,04	0,64	-0,33	0,06	1			
Average word count	0,21	-0,45	-0,52	-0,04	-0,33	-0,33	0,06	-0,07	1		
Total word count	0,2	0,77	-0,38	-0,14	-0,28	-0,4	0,07	0,73	0,98	1	
Total tags	0,54	0,62	0,14	-0,15	0,78	0,65	0,92	0,77	0,19	0,22	1

Large-Scale Collaborative Story Worlds: Formalizing Content and Author Dynamics

Alan Tapscott¹, Joaquim Colás¹, Carlos León², and Josep Blat¹

¹*Grup de Tecnologies Interactives, Universitat Pompeu Fabra*

²*Facultad de Informática, Universidad Complutense de Madrid*

Abstract

Shared collaborative narratives are an emerging media, in parallel with the so-called Web 2.0 and Web 3.0 and along with other forms of enriched user-generated content. Participative creative writing and the documentation of fictional story worlds represent two instances of this trend. Hybrid story world sites mix both goals into a single platform. In this article, we do an overview of some of the most successful ones, and study one of them in depth, with the specific goal of understanding their content structure and the author dynamics. Indeed, our analysis reveals the ways these sites allow large author communities to create their own stories and document the story world they contribute to in an intertwined loop of synergies between both activities. Reflecting the web general evolution, in these sites the author dynamics have changed over time, opening their editorial control progressively, delegating the plot overall control to general users and providing semantic networks to link contributions. We also show how hybrid sites resort to certain stylistic resources, such as unreliable narrators and mysterious backgrounds, which are very different from those introduced in encyclopedias. These resources imitate popular folklore and urban legends, and seek to avoid potential inconsistencies between authors while resulting into a more engaging read. We analyse more in detail a recent and successful site—the SCP Foundation¹-. By examining the semantic tags of the pages and link distribution we conclude that most content is either an encyclopedic article, a tale, a link hub or a community discussion. We also discuss how characters roles are displaced from traditional characters to supernatural anomalies and “group of interest”, inclusive character factions. The SCP Foundation implements the legendary status of hybrid story worlds through a multi-level canon model that allows authors to be coherent with each other to the degree they wish to. To conclude our work we present a more formalised model (which we call Open Story World) of this type of content structure and dynamics found in

¹ www.scp-wiki.net/

hybrid story world sites. The model introduces metrics inspired by popular page relevance algorithms, meant to determine content and author plot integration and inter-author affinity through simple wiki page magnitudes. We conclude this study with a discussion of the implications of the model for author dynamics and their contributions.

1 Introduction: Collaborative Creative Writing

Shared collaborative narratives are an emerging media, in parallel with the so-called web 2.0 and web 3.0 and along with other forms of enriched user-generated content. In this article, we aim to get a better understanding of how large author communities collaborate to write stories and document original story worlds in online sites created for that specific purpose. We use the term story world to describe the conceptual dimension of a story, containing the facts and plot elements (such as characters and locations) that constitute the emerging fictional reality established by the author(s). Yet, a story world can be created as an original entity, in a process that mixes creative writing and encyclopedic documentation. Moreover, if more than one author is using the same plot elements, the creation of the story world is a collaborative exercise. The term collaborative writing is often used to describe a project in which more than one person is involved in the joint creation of a written document. Coordinating collaborative writing generally benefits from a strategy, a methodology. From pairs of authors working closely to large crowds that communicate scarcely, there are many possible roles and configuration of participants. Amongst the participants, there is at least one who writes or formulates the actual words while other contributors might participate in the writing, conceptualization, edition and revision. The result of such collaboration might be superior or not, yet few authors feel the process is something natural or intuitive [47]. On the other hand, each genre has its own culture and style, sometimes even involving specific multi-author methodologies. Academic articles, for instance, are often signed by multiple authors, but narrative and similar forms of storytelling are rarely attributed to more than one author, despite going through extensive editorial and processes. Individual novels or tales signed by multiple authors are rare, often experimental occurrences: the creative role in narrative works is very often attributed to a single author. Fictional narrative relies on invented facts and stylistic resources, generally used more coherently through a single authorial voice. Indeed, anthologies are one of

the few common collaborative' narrative formats, compiling the work of multiple authors into a single publication. However, each chapter typically contains a tale by a single author, still maintaining the boundaries between authors and their creations. Tales contained in an anthology might or might not converge in plot or tone and very rarely occur in a common fictional reality with shared plot elements. Not all the instances of collaboration occur in a single document. The backbone of human culture is a collection of popular tales and folklore, that indirectly influences our perception of reality and also our every creation. In this sense, collaborative writing can be understood as a more general activity that occurs implicitly and is hardly distinguishable from cultural dynamics. Ancient mythology tales feature intertwined plot elements, often crossing characters and themes in a similar way than contemporary media and large fiction franchises. For instance, Greek mythology had a recurring cast of heroes and gods who behaved in a more or less consistent way, establishing a pantheon with recurring themes, characters and plots. Tales from classic Greek mythology were often anonymous but it is assumed that they were created in different moments of history by multiple authors. Roman mythology picked up the same plot elements and further developed the mythos, extending the mythos even further. Later examples include the Arthurian cycle legends, written by multiple authors from the 12th to the 15th century and the Cthulhu mythos, by H.P. Lovecraft and his circle of correspondents during the early 20th century. Despite the existence of collaborative writing in our culture, most studies on this subject are relatively recent and scarce. [10] uncover some of the earlier works, such as an early essay covering teams of husband and wife writing children books. A more recent and general work [22] compares ancient mythology to modern fan fiction. In a influential work on small group research [18], the author acknowledges the complexities inherent to a successful integration of results in the context of a collaborative study effort. He also insists on the lack of research concerning this problem, although scientific publications have been the subject of the majority of the studies on collaborative writing. Typically more concerned with the product than the process, these studies deal with specific aspects of the collaboration such as effects of author name ordering [56, 36], a comparison of different types of textual measurements commonly used in attribution studies [16] and analyses of the causes and effects of collaborative research and publication [39, 38, 32]. The subject of collaborative writing has also been explored in the learning field, with relevant publications such as the ones that compose the cooperative learning body of research

[20, 43, 37]. More recently, there have been many studies that also explore collaborative writing from the perspective of learning, particularly studies involving students producing a jointly written text [47], student perception of wiki use for collaboration and writing [30, 54], or the usage of wikis as collaborative writing tools for students [17]. The wiki is used in most of these studies, however some works resort to other platforms, such as blogs [50] or Google Docs [48, 23].

Collaborative writing has also been used for the authors' own amusement. Some recreational examples of multi-authoring stories include the surrealist tradition of the exquisite corpse in which authors take turns to extend an existing story. [1] provides a thorough analysis of the exquisite corpse as a technique and strategy. Despite its experimental nature, [24] provide a unique academic and artistic view on the exquisite corpse. Another example of a similar technique is the round-robin story, in which every author takes turns to write a chapter of the same story. This technique is not as popular as the exquisite corpse and is often used as a tool to teach how to read instead as a writing technique [11]. In these examples, the original author delegates the creative authority to other participating authors, and the driving motivation is not to produce valuable stories, but to amuse the authors. In this sense, multi-authoring can sometimes be meant for authors instead of readers. This does not imply that experimental, participative methodologies for collaborative writing cannot produce valuable stories. The internet and information technologies seem to follow a similar path, relying more on user participation and content creation. Collaborative writing-for education, creative writing and as a leisure activity-has developed over the last years as a large-scale activity. Most studies involve relatively small groups compared to the large-scale real communities engaging in this kind of activities, often involving hundreds or thousands of users reading and writing. Generally, these large-scale studies of collaborative writing are focused on Wikipedia, ignoring other collaborative exercises that deal with fictional knowledge such as derivative online encyclopedias (e.g., Wikia² or Wikidot³) or creative writing communities (e.g., FanFiction⁴ or Protagonize⁵). Amongst the works that study these platforms, most explore fandom and fan culture [12, 40] or their literacy and language learning potential [4, 25]. Wikias have been featured in some works that explore their large-scale nature, such as studies of impact

² <http://http://www.wikia.com/>

³ <http://www.wikidot.com/>

⁴ <https://www.fanfiction.net/>

⁵ <http://blog.protagonize.com/>

of membership overlap on the survival of online communities [55] or recurring subgraphs in co-author networks [3]. We found no literature focused on the collaborative construction of fictional realities or story worlds. The focus of this work are essentially unprecedented exercises of large-scale collaborative writing and the implicit story world building.

As we mentioned previously, there are many potential methodologies to write collaboratively and document types to produce. The latest developments in information technologies have made collaboration more accessible and viable for large amounts of remote users. This work seeks to study how authors collaborate in writing fiction together in the context of revolutionary information technologies, and more specifically, to understand how they coordinate and build fictional worlds together. Our main concerns are to determine collaboration dynamics, and which types of contributions are generated related to that collaboration. First, we do a brief overview of sites with either an active large-scale participating community or a singular approach to this scenario. We cover two distinct kinds of collaborative sites: those which promote creative writing, and those which document story worlds developed in popular media, followed by an analysis of the hybrid sites, which appeared later and mix both goals; indeed they are collaborative sites that promote creative writing and the documentation of the emerging original story world in parallel. Our analysis is focused on content types (its function and relationships) and author dynamics (collaboration nature and traits). Some of these sites have been less successful than others, and we analyze in detail the SCP Foundation which is hybrid and one of most successful ones, using more exhaustive web crawling to get more precise qualitative and quantitative results. From a more abstract view of these, we propose a model that helps in describing and analyzing original collaborative fictional story worlds, the Open Story World. Our proposed model includes formulas that inform of content and author plot integration as well as inter-author collaboration affinity.

2 Collaborative Stories and Story Worlds in the Web

This section presents an overview of different types of online sites for collaborative writing, starting with participative creative writing sites and collaborative fictional world encyclopedias. Then we focus more in detail on the hybrid collaborative narrative and story world sites, discussing the main ones, in terms of content types and author dynamics.

2.1 Participative Creative Writing Sites

The Internet and the rise of information technologies have brought about public and accessible information platforms, including some that encourage users to participate by contributing with fictional writings of their own. The cost of exchanging information and collaborating has decreased dramatically, allowing large crowds to contribute and share their creations. In this context, there has been a wide array of emerging opportunities for users to become authors of all sorts of fictional stories and share them with each other. Indeed, there have been many sites that promote collaborative fictional writing, implementing a wide array of rules and methodologies. We discuss some of the popular or singular ones along with the content they host(ed), reader's capabilities and author dynamics.

WikiStory⁶ was one of the pioneering initiatives that allowed authors to collaborate online. It provided users with a wiki in which they could write their own stories. Contributions were public and any user could participate in the attached discussion section. Any author could opt to participate in the multi-authoring segment of the site, allowing any other user to edit or comment the story. The site's own guidelines suggested edits could be either story extensions or editorial revisions, but had no sort of editorial control besides the basic profanity restrictions to allow readers of all ages. Optionally, authors could protect their stories to avoid other author's modifications. Readers could also comment stories as long as they registered into the site. StoryMash⁷ is another creative writing community meant for authors and readers. Authors can create new fiction stories or introduce a new chapter into an existing one. The focus on chapters is perhaps one of the most unique traits of this site. Chapter additions can go anywhere in the narrative sequence (preludes and interludes are allowed and encouraged). The web shares part of the advertisement revenues with the authors and holds writing contests with monetary prizes, enhancing the allure for amateurs and professionals. The site's community has a reputation-perhaps due to some professional authors charging readers-of being quite critic with the authors, however there are no restrictions for uploading contributions. Readers can rate stories and provide in-depth feedback, as well as purchasing the non amateur stories. Fabulate⁸ was another similar site and community with a different premise. Authors submitted individual pages to ex-

⁶ wikistory.com, offline, reached via archive.org/web/

⁷ <https://winningwriters.com/resources/storymash>

⁸ www.fabulate.co.uk, offline, reached via archive.org/web/

tend an ongoing book. The site's guidelines were loose: pages had to "contain up to 500 words and follow the previous page and make sense in the context of the book". Readers could also become reviewers by applying for the role. Reviewing readers could rate individual pages and post reviews for recent contributions. The staff then decided if a page was approved or rejected according to the reviews and scores. Addventures seem to be a recurring format for internet collaborative fiction writing. According to Wikipedia, an addventure is: "... a type of online interactive fiction that combines aspects of round-robin stories and Choose Your Own Adventure-style tales. Like a round-robin story, an addventure is a form of collaborative fiction in which many authors contribute to a story, each writing discrete segments. However, like a gamebook, the resulting narrative is non-linear, allowing authors to branch out in different directions after each segment of the story. The result is a continually growing work of hypertext fiction." One Million Monkey's Typing⁹ was a site dedicated to collaborative addventures. Authors could submit their story snippets for users to read, either by starting a new story or branching an existing one by allowing readers to alter the direction of the plot to follow their contribution, in a what-if? scenario of sorts. Readers could explore any story, interacting with the plot whenever a question was asked by choosing the direction of the events. Readers

Protagonize¹⁰ was another successful collaborative writing community that used the addventure format. The community originated in a site meant for addventures called Choose your own Schizophrenia, which evolved into Protagonize, a platform that supports all kinds of creative writing. One of the most salient features was the capability of the original author to establish the author guidance section. The author guidance section contained notes on the plot outline, setting, characters, narrative mode, intended length, and inspiration amongst other relevant aspects. This did not enforce subsequent authors to comply with it, but it did empower the original author with explicit informal authority over the content. Neither the original author, collaborating authors or readers could exercise editorial control of any sorts over the contributions. As usual, users could comment stories, rate them or even recommend them to other readers.

Overall, most sites present a decentralized informal editorial process, in which readers rate their favorite stories and explicitly comment them. The only users with real control over what is published are the

⁹ 1000000monkeys.com, offline, reached via archive.org/web/

¹⁰ www.protagonize.com, offline, reached via archive.org/web/

staff, and generally only intervene to prevent profanity and unpleasant behavior. Authors and readers are free to contribute and criticize with little constraints. Another common trait observed was that most contributions were individual, rarely involving collaborations with other authors. Crossovers with shared plot themes or elements were rare, i.e., despite users were collaborating via reader rating and feedback, a collaborative story world is generally not established.

2.2 Fandom and Fictional Story Worlds

The irruption of the fandom phenomenon has brought new perspectives on creative writing and the documentation and creation of fictional story worlds. Fans use the internet and digital encyclopedias—generally derivative wiki sites—to document their favorite media. As of May 2016, there are 30643 active Wikia sites run by dedicated communities¹¹. The result is an explicit formulation of an existing story world, most often curated by users unrelated to its original authors. Some works have studied these fan-related phenomena, Narrativity [5], Documentary Simulacra [7] and wikis and participatory fandom [33]. The contents of these wiki sites describe the causal and temporal connection between actions to document story worlds, as in the concept of narrative space [34]. Collaboration on wiki platforms has also been the subject of many studies, as [29] perspectives on its value as a collaboration platform, [26] interpretation of the revision history as a collaboration network, [28] analysis of collaboration patterns and article quality, or [49] social dynamics. [15] analyzed the Lost TV series and their fan activity on the net, discussing the encyclopedic and creative dynamics while focusing on the spoilers revealing crucial information to the audience during the chapters' original run. The authors elaborate on the community's struggle to establish some sort of order to the chaotic contributions from fans. Overall, wikis have been often presented as a good example of collaboration, even to the point of being considered “good democracy” [29]. The collaborative documentation of a story world by a large-scale community has been made possible for the first time thanks to the tool itself—the wiki site-, and therefore users might not conceive performing it with a different one. The original purpose of a wiki site was not to document fictional knowledge from multiple authors—with its potential gaps in consistency and coherence-, i.e., its design might have room for improvement and other hypothetical tools might be suitable for the task. There are alternatives to wiki sites for large-scale collaborative documentation of

¹¹ <http://wikis.wikia.com/wiki>

story worlds. Traditional written encyclopedic documents have been used for this purpose as internal documents—such as a franchise bible—and merchandise—such as a sourcebook—. [articy:draft](https://www.nevigo.com/en/articydraft/)¹² and [Celtx](https://www.celtx.com/)¹³ are both commercial collaborative tools meant for structural creative video game story world development. Their design is partially focused on the specifics of the video game medium. These kind of tools are often created by the same company responsible for the final narrative content, meaning that their design might be less generalist and more ad hoc, seeking to a specific medium and a specific set of problems. Other similar commercial tools include editors designed to design non-linear branching narratives and interactions such as [Twine](https://twinery.org/)¹⁴ or [ChatMapper](http://www.chatmapper.com/)¹⁵ or more general purpose semantic diagrams such as [yED](https://www.yworks.com/products/yed)¹⁶ and meant for digital information maps, such as [VUE](http://vue.tufts.edu/)¹⁷. The general purpose of these tools allows designers to craft interactive stories, but they are not meant to document or extend story worlds and mostly avoid the collaborative dimension we are exploring in this paper. A skilled writer with the capability to abstract data might be able to document a story world with them, but they are neither accessible nor scalable for this specific use case. Fans, however, do more than document existing story worlds. Using participative creative writing sites, similar to the ones we have seen in the previous section, they also extend popular media with their own contributions. By intertwining their own stories with a well-established story world and its mythology, networks of original fan fiction coexist with the official material are created. Online information technologies provide momentum to these phenomena, exponentially augmenting their size, reach and impact. For instance, the [FanFiction](https://www.fanfiction.net/) or [Quotev](https://www.quotev.com/)¹⁸, gather hundreds of thousands of users around hundreds of franchises, presenting fan-developed stories that extend the story world of popular media. Most of these stories, however, do not take into account the contributions of their fellow fan authors and represent a one-off extension of the original story world, limiting their potential to enrich the original material or even to span new fictional parallel realities or story worlds. The fanfiction phenomenon has been extensively studied over the recent years. [52] provides a brief overview of the nature

¹² <https://www.nevigo.com/en/articydraft/>

¹³ <https://www.celtx.com/>

¹⁴ <https://twinery.org/>

¹⁵ <http://www.chatmapper.com/>

¹⁶ <https://www.yworks.com/products/yed>

¹⁷ <http://vue.tufts.edu/>

¹⁸ <https://www.quotev.com/>

of online fanfiction communities, studying their general literacy practices within forums, chatting, role-playing and the discussions strands of the community. [13] discusses the main characteristics of fandom from the perspective of economics, highlighting how this phenomenon establishes an alternative culture that goes against the “official” one. This perspective reinforces that story worlds are used to collectively creating an alternative reality. The line between factual, hypothetical, speculative and original content was not a very clear one, especially when documenting fictional media that often resorted to mystery or intrigue tropes. Again, fan contributions are struggling to become more than a reflection of the original media. [44] presents a model that brackets the opposing potentialities of internet influence on offline society using large scale participatory fandom to center her discourse. She discusses fan clubs, online producer-consumer affiliations and real-world legal controversies in the context of fans attempting to participate in the media they revere.

2.3 Hybrid Collaborative Narratives and Story Worlds Sites

2.3.1 Hybrid sites and similar initiatives

The communities discussed in sections 2.1 and 2.2, where creative writing and documentation of fictional story worlds were separated, inspired the creation of hybrid sites merging both at some point. This idea had some precedents, such as multiplayer games in which players build the world themselves (e.g., Multi User Dungeons-MUDs-or Second Life communities). It also seems somewhat similar to the “novel scriptural world” suggested by [7] in its analysis of the Harry Potter Wikia site. The crossroad they describe between fictional regime and documental regime seems to apply in these original collaborative story worlds. In creative writing, the overwhelming success of wiki-like platforms for collaborative tasks encouraged communities to develop their story worlds and fiction jointly in a common wiki site. The result is a collaborative site with a repository of tales that take place in a shared story world and also an encyclopedic collection of articles that describe the very same story world defined by the tales. The hyperlinking capabilities of the web allow inter-page links that readers and authors use to navigate in an integrated experience. From the author’s perspective, this represents an opportunity to contribute in fundamentally different ways. S/he can write a new story and upload it to the site, or alternatively, an encyclopedic entry, describing an element that fits the theme of the site but not elaborating much into

its role or history. This also means that authors can write stories that feature the existing encyclopedic elements or also write encyclopedic entries for the elements introduced by existing stories. In creative writing, these use cases are innovative and made possible thanks to the hybrid nature of these sites. The wiki also provides the default mechanisms for author and reader interaction, including discussion, versioning, reviewing and rating. The hybrid sites we discuss next include large author communities and substantial content both narrative and encyclopedic. There are several reasons for an extended discussion of these sites: they combine both writing and documentation, and they have been successful in involving large communities and generating extensive content. One of the sites (The SCP Foundation) seems especially successful in terms of the number of authors/readers and the content created and/or documented. Thus, we undertake its detailed analysis in the next section. We discuss first several of them, which have been successful as well, but to a lesser degree, to provide a background for this analysis and avoid it to be too unilateral.

2.3.2 Overview of the sites

Tables 1 and 2 reflect a descriptive overview with some measures of participation and content of the different sites (as of May 2016).

The Holder's Series¹⁹ contains a collaborative collection of supernatural Holders. Every Holder is an entity (physical or metaphorical) and its entry describes a methodology to acquire it. The steps often involve creepy or unsettling actions written with an urban legend vibe. The site also has a section for creative fiction that involves the Holders mentioned in the collection. Galaxiki²⁰ is presented to readers as a fictional map of the galaxy, allowing to read and author the description of specific solar systems, planets or stars. The site also has a section for stories of any sort that take place in the galaxy. Its business model encourages users to purchase existing solar systems with real money, whose content can be only modified by the buyer. The Orion's Arm Universe Project²¹ describes itself as a hard science fiction collective world building effort. It contains the *Encyclopaedia Galactica*, a large collection of articles that describe with detail the fictional universe. Users may submit new science fiction articles that either rely on real science publications or are grounded on a plausible chain of events in the future. The site promotes many forms of

¹⁹ <http://theholders.org/>

²⁰ <http://www.galaxiki.org/>

²¹ <http://www.orionsarm.com/>

art, such as renders, music or writings, and often hosts contests and publishes magazines and anthologies to disseminate the art produced. Some of these publications are not free. The SCP Foundation²² is a fictional organization to secure, contain and protect humanity from all sorts of supernatural anomalies (the SCPs themselves). Each SCP entry, created by a participating user, follows a structured template and uses scientific terms to describe the supernatural object or occurrence. Pictures and complimentary reports are commonly included in each SCP. There is also an extensive collection of fiction linked to the documented SCPs and the Foundation. The Wanderers' Library²³ is a spin-off site originated from the SCP Foundation based on the same fictional material, which holds all sorts of fantastic manuscripts instead of documenting supernatural occurrences. As its sister site, it also contains stories that involve the documented manuscripts. These two sites exist in the same shared story world, with some elements (such as certain groups of interest) featuring in both. The sites summarized were created between 2000 or after 2010, and we could not find similar sites that started before or later. While their golden age might have passed, these sites that are still active and receive new contributions. Only the oldest sites (The Holder's Series and The Orion's Arm Universe) use their own engines, while the most recent ones implement an existing collaborative engine. Their functionality is very similar to regular Wikis, allowing users to create their own pages.

2.3.3 Content analysis

There are two main trends regarding plot and tone of the contents, sci-fi and urban legends/paranormal phenomena. Both are multi-author friendly. Sci-fi and space exploration are very favorable for authors who wish to carve their own region and story without conflicting with others. Urban legends borrow from oral tradition and ancient storytelling dealing with fictional facts that could occur in our daily lives. As stated by Tolkien [53], fantasy that does not diverge much from reality is friendlier to most readers. Furthermore, urban legends rarely cross or reference each other, granting more freedom to authors to contribute without taking the other existing contributions into consideration to prevent contradictions or incoherences. In general, these sites seek to portray themselves as extensions of our own reality, although through very different means. The usage of these purposely

²² <http://www.scp-wiki.net/>

²³ <http://wanderers-library.wikidot.com/>

vague background might reduce the need for all stories of a common story world to be coherent with each other. A shared story world with stories that tie to each other perfectly might be at odds with having many authors with strong creative visions free to contribute in the way they see fit. One of the most distinguishing traits of legends is their implicit unreliability. Even if the narrator does not make it explicit, the audience acknowledges that they might or might not have happened [8]. This legendary trait is used to avoid the problem of coherence in large-scale collaborative story worlds. While it is far from a new perspective, it is innovative when applied to a digital encyclopedia. It is debatable if it merely represents a way to avoid the problem (incoherences are still present and might hamper the reader's experience) but nevertheless it is a functional and successful solution. Also, the usage of in-universe strategic omissions and unreliable narrators reinforce this strategy, suppressing the need for precise information and rigor with stylistic resources that even increase the engagement of the text. Despite being hosted in wikis (essentially web 3.0 software [14]), these sites go beyond regular electronic collaborative encyclopedias usually used to document knowledge or fictional knowledge, providing a framework for the organic growth of fictional cultures or mythologies in a process that is reminiscent of our own cultural development.

Semantic tags only make an appearance in the SCP Foundation and the Wanderer's Library. These mechanism, often seen in modern sites, allow content to link each other following specific subjects or themes. Besides allowing to clearly distinguish the encyclopedic from the narrative content, these non-exclusive tags allow users to navigate through specific author content, similar SCPs, or even content related to specific groups of interested or canons. Authors can not create new tags, however, if they follow the guidelines, they might introduce new groups of interest or canons, effectively altering the topology in ways the other sites cannot replicate. We go more in depth to explain this aspect of the Foundation later. Overall, we see a trend here to imitate the semantic web 3.0 by adding other metadata to the user contributions. Our analysis of encyclopedic content revealed that the majority of the sites have a section that focuses on one type of encyclopedic element . While there might be other element types featured, most of the content is of the same kind. Only the Orion's Arm does not feature a standardized encyclopedic content, meaning that the site encyclopedia is of general purpose. The Holders' Series encyclopedia is about Holders, Galaxiki focused on galaxies, the SCP Foundation is meant to collect SCP reports and the Wanderer's Library is a manuscript

library. This central element is based on a standard template with some required fields to be filled as the author sees fit and is useful to provide some order and writing guidelines while enhancing consistency and coherence. This implies that the encyclopedic segment of the sites might not be an actual encyclopedia (with entries describing relevant objects, characters locations), despite using an encyclopedic platform and format. The serialized nature seems to favor author participation, providing a basic guideline for new encyclopedic entries. The Orion's Arm encyclopedic section is the only that follows a classical encyclopedia formula. Regarding narrative content, there seems to be more encyclopedic content than tales (with the exception of Galaxiki, although our analysis could not cover all solar systems, so this might not be conclusive). Overall, the encyclopedic original purpose of a Wiki site seems to drive authors to contribute with non-narrative information. There is, however, a strong link between tales and encyclopedic content. Our informal observations suggest most tales feature some element from the encyclopedic content, but many encyclopedic pages are not featured in any tale. Most references are explicit hyperlinks, however implicit references without a hyperlink are not unheard of (for instance, SCPs from the SCP Foundation often refer to each other by codename rather than hyperlink). Only the SCP Foundation and the Wanderer's Library offer authors the capability to connect content through custom hub pages. Generally, hub pages provide a list of links to content with a common author or canon while providing additional context (in-universe information and/or from the real authors' viewpoint). Ultimately, custom hubs allow authors to go beyond mere contributions to provide new ways for readers to navigate the site and the story world. These semantic interpretations of the content are in line with modern creations closer to the web 3.0 paradigm.

2.3.4 Editorial control, feedback mechanisms, interaction

In terms of editorial control, all of the sites have some sort of guidelines on the kind of content expected from contributors, which are used to determine which contributions are accepted. Galaxiki is by far the less strict when it comes to content, allowing anyone to post new content and only enforcing minimal standards (a profanity filter not related with the science fiction theme). Only the Holder's Series has a non-participative editorial process in which the site staff has the last word on whether a new contribution should be erased, while the rest of the sites have a public process in which all users can par-

ticipate. The Orion's Arm Universe requires new submissions to be posted in the message board to allow other users to post their reviews, although the staff finally decides. The site's writing guidelines remark the need for any submission to be backed by the laws of science, citing rigorous publications if needed. This is a deliberate attempt to keep the Orion's Arm Universe in the genre of hard science fiction instead of fantasy or space opera. The SCP Foundation and the Wanderer's Library require new contributions to be publicly posted on the board, but the ultimate decision is (democratically) determined by users' scores. There is a trend to move content authority and responsibility from the site staff to the participating users. Only the Orion's Arm Universe lacks feedback mechanisms allowing readers to rate and comment specific pages, but drafts must appear on the message board for other users to criticize. The more modern sites allow users to say their opinion, assuming that this usually means more user activity and involvement and perhaps more user attachment to the contributions. Indeed, Galaxiki system was rewritten in 2010 to include feedback. The message boards to support interaction are always present. More evolved systems are those of Galaxiki, that introduce a public blog that allows users to upload their own news and chats provided by the SCP Foundation and the Wanderer's Library, which support as well running recurrent contests to promote new submissions (the Orion's Arm Universe, too). More community interaction options are thus available in more recent sites. Formal support for collaborative authoring of individual pages is only supported by the more modern sites, the SCP Foundation and the Wanderer's Library. These sites include a tag for these contributions, clearly highlighting their different nature. They also provide access to the whole page history through a version control system to keep track of changes so that multi-authoring is more stable and secure.

2.3.5 Activity

Table 3 presents quantitative data on the sites as of May 2017, based on semrush.com and similarweb.com

The activity data we have gathered suggest the SCP Foundation has the most user activity in terms of visits, average visit duration, average pages visited per user and links received from other sites. The previous analysis suggested that the SCP Foundation and the Wanderer's Library were the most open and participative sites, yet the latest seems to be underperforming compared to her sister site. Also, the Orion's Arm Universe displays a high activity despite not

implementing many of the mechanics found in modern participative communities.

2.3.6 Summary and discussion

While we have no data from our own sources to corroborate that the presented activity measurements are reliable, we suspect the snowball effect found in many online communities and other factors unrelated to the interaction design and author dynamics might be very influential in the sites overall success. Perhaps the SCP Foundation relation to the extremely popular internet board 4chan (the backlink measurements could reinforce this hypothesis), as well as the potential splintering of the community when the Wanderer's Library was created could explain some of these results. The analysed hybrid sites were created over the last decade, perhaps highlighting a trend from that specific time period. Unlike some of the creative writing sites we saw in the previous overview, they still exist and remain active. Our analysis of these hybrid sites has showed how they have evolved over time to progressively allow users to actively participate in the editorial process and its discussion (following the web 2.0 trend in storytelling [2]) and to add structure, semantic navigation and other meta-data (following the web 3.0 to introduce semantic web content, specifically "contextually relevant and easily interpretable content" [45]). We anticipate a tendency for these sites to imitate or integrate with popular social content platforms such as Tumblr or reddit in the immediate future. It remains to be seen if other popular, participative activities or platforms could be integrated to introduce other innovations beneficial to the creation of a collaborative story world. For instance using social network friend circles to explicitly associate author circles or blending the concept of news aggregation or trending topic with narratives arcs.

Overall, These sites represent examples of how author communities can create original story worlds collaboratively while being successful, overcoming all possible difficulties. Their hybrid nature might be reflecting our own culture deeply grounded in documented knowledge and myths. A microculture of sorts that reflects our own. We have also seen how in terms of author dynamics, generally speaking, openness and participation in terms of editorial process, contribution and data structuring, seems to be the predominant trend. Perhaps with a closer look into one of these sites we might learn more about the content contributions in terms of function and relation to each other.

In the next section we conduct a more thorough analysis of one of them (the SCP Foundation).

3 The SCP Foundation

As detailed in Table 3, from the collaborative websites analyzed, the SCP Foundation website is the most visited. Additionally these visits also take longer than the ones to the other sites, on average, and, as described in the previous section, it holds more content, both encyclopedic and tales. For these reasons, the SCP Foundation site has been chosen for a specific qualitative and quantitative analysis.

3.1 Overview of the SCP Foundation site

The SCP Foundation is a wiki site with a large community dedicated to creating and curating reports and stories that take place inside a common fictional story world. Authors contribute, discuss, peer review, rank and edit pages, in a public and participatory process. Site editors try to keep a minimum order in the message board and the profanity level low while enforcing the basic rules. Given their full permissions to edit the content, they can erase or improve the content, or ask the original user to apply some modification to her contribution. The SCP Foundation site's own description provides a detailed history of its inception. The site started after some users from the paranormal board of the 4chan community²⁴ began posting reports of fictional paranormal anomalies in a standardized way, including pseudo-scientific descriptions along with creepy and eerie (sometimes manipulated) real pictures. These reports were named Special Containment Procedures (SCPs). The community eventually moved to its own site. We lack clear evidence to establish if they ever considered a non-wiki platform. The site added a new background section that developed the idea of a government-backed organization devoted to the scientific study of paranormal objects and entities, the SCP Foundation. SCP stands for *Secure, Contain and Protect*, the mission statement of a fictional shady organization dedicated to the location, study and containment of paranormal anomalies. Over time and thanks to ongoing contributions, the content of the SCP Foundation slowly grew involving loosely related set of fictional mythos that included multiple relevant organizations and recurring anomalies. The community eventually began to push the boundaries to not

²⁴ <http://boards.4chan.org/x/>

only create new SCP reports, but also to write tales involving the anomalies. Instead of setting up another site for the tales, they were included in the same wiki site, using the hyperlinking capabilities to knit the content together following references and allusions.

3.2 Content analysis

The SCP Foundation site is a network of information with distinct purposes inside a common fictional story world. Its reference material is self-contained, meaning that instead of documenting external knowledge or media, the source material is the encyclopedia itself. In this section we explain the main content segments along with their purpose.

3.2.1 Content function

The SCP Foundation writer's guide establishes this general classification for the site's content:

- SCP Series
- Tales
- Canons
- Groups of interest

The *SCP Series* reports document the anomalies that represent the focus of the SCP Foundation story world. SCP reports are always labeled with a serialized common denominator (SCP-xxxx) and follow a standard template. The SCP report (or just SCP) itself must be written using objective language, mimicking a scientific publication or research notes. The standard template includes the object classification (based on how dangerous it is), the necessary containment procedures and a general description. Additional media is often included, generally in the form of a picture of the anomaly but sometimes including audio or even video. The description is intentionally vague, providing intriguing aspects of the object without revealing their fundamental nature, origin or logic. Such omissions include avoiding the subject, citing poor experimental conditions or "black tape" found covering critical information. Authors also often include secondary articles such as interviews with people who have interacted with the anomaly, complementary reports on incidents involving the anomaly or detailed notes on fictional experiments.

SCP-173

Item #: SCP-173

Object Class: Euclid

Special Containment Procedures: Item SCP-173 is to be kept in a locked container at all times. When personnel must enter SCP-173's container, no fewer than 3 may enter at any time and the door is to be relocked behind them. At all times, two persons must maintain direct eye contact with SCP-173 until all personnel have vacated and relocked the container.

Description: Moved to Site-19 1993. Origin is as of yet unknown. It is constructed from concrete and rebar with traces of Krylon brand spray paint. SCP-173 is animate and extremely hostile. The object cannot move while within a direct line of sight. Line of sight must not be broken at any time with SCP-173. Personnel assigned to enter container are instructed to alert one another before blinking. Object is reported to attack by snapping the neck at the base of the skull, or by strangulation. In the event of an attack, personnel are to observe Class 4 hazardous object containment procedures.

Personnel report sounds of scraping stone originating from within the container when no one is present inside. This is considered normal, and any change in this behaviour should be reported to the acting HMCL supervisor on duty.

The reddish brown substance on the floor is a combination of feces and blood. Origin of these materials is unknown. The enclosure must be cleaned on a bi-weekly basis.

SCP-1460

Item #: SCP-1460


Object Class: Euclid

Special Containment Procedures: SCP-1460 is kept in a standard humanoid containment cell at Site █. Contact or experimentation with SCP-1460 outside of standard medical care may only be performed with prior permission from at least two (2) senior research personnel.


SCP-1460 is to be administered a psycholethal/analgesic regimen daily. Until further notice, use of medication outside of that required for basic medical care has been suspended.

Elements of Mobile Task Force Gamma-5 ("Red Herring") have been deployed to the areas affected by SCP-1460 and have standing authorization to administer amnesia as necessary to maintain secrecy and minimize disruption of the civilian population.

Description: SCP-1460 is a 32-year old male Asian-American identified as [REDACTED], a former resident of the city of █, █, USA. SCP-1460 has been in a coma for the past eight (8) years, of which the last seven (7) have been in Foundation care, as a result of a traffic accident on █, █, █ in which a public transit bus struck and critically injured SCP-1460 while he was walking to his place of employment. SCP-1460 was transferred to a regional hospital due to a lack of appropriate trauma facilities at the local hospital, and came



SCP-173 in containment



File photo of SCP-1460 shortly after being admitted to [REDACTED] Regional Medical Center.

SCP-1018


Item #: SCP-1018

Object Class: Safe

Special Containment Procedures: SCP-1018 is currently contained in Storage Unit █, located at Site █.

All components of SCP-1018 are to be kept in 1m x 1m x 3m containers and kept under guard by two (2) members of security at all times. No liquids are to be introduced to SCP-1018 components without the express permission from one (1) member of Level 3 personnel.

Description: SCP-1018 is the overarching designation given to a group of three statues hereafter referred to as SCP-1018-1, SCP-1018-2 and SCP-1018-3. All components of SCP-1018 depict an emaciated or elderly human male and are composed primarily of concrete and gravel. A red luminescence is visible on several points on the surface of SCP-1018 components, mostly present on the head and throat. SCP-1018-1's anomalous effect is activated by applying pressure to the back of its neck while SCP-1018-2 and SCP-1018-3's are activated by contact with liquid.



SCP-1018-1

Fig. 1: Sample SCP Series reports.

Tales are stories that involve the SCP Series in some capacity. These pages contain fiction that uses classic narrative tropes more akin to the kind of content found in a novel or a short story anthology. Some contributions are more eclectic, including songs and poems, but most belong to the narrative genre. The Foundation tales feature and reference SCPs from the SCP Series or other Foundation Tales, enhancing the reader's immersion in the global story world and encouraging its exploration of other content. Most tales are contributions that reference (directly or indirectly) SCP reports, although a subset of them can be considered extensions of the narrative content. These extensions, instead of being a free narration with some connection to the Foundation anomalies, complement specific SCP reports with experiments, explorations, incidents, interviews or general supplements.

Canons are content subsets of the SCP story world created by author groups to reinforce the consistency and coherence of their content. The SCP Foundation allows contradictions and incoherences, mainly through canons. The site's guidelines explicitly states this:

“The idea that there is no canon is a bit silly at times. It’s not that we don’t have any. It’s that we have a multitude which touch, cross, and dip into each other. It’s up to you, as the reader, to decide what you believe and what you embrace as the heart of the universe. That doesn’t mean, though, that authors lack intent or design, and collaboration is the heart of innovation.”

Contributions introduced into a SCP canon present a reinforced sense of consistency dealing with recurring plot elements (mostly SCP anomalies) and an enhanced coherence. Internal canon contradictions are possible but rare—they defeat the purpose of canons-. While most SCPs and tales rarely reference each other (directly or indirectly), canons often feature closely related narratives, with overarching plots involving common anomalies, characters and locations. Canons might include a wide array of contribution types, including SCP reports, tales and other miscellaneous ones, for instance a hub page (or more) meant to explain their premise, content along with reading aids such as timelines or relevant characters lists. Also, canons often feature contributions from multiple authors.

Groups of Interest or *GoIs*, are another content subset featuring fictional character groups or factions relevant to the SCP Foundation itself. They are similar to Canons in the sense that they often feature multiple authors and include multiple types of content like SCP reports, tales and hub pages with reinforced consistency and coherence. The main difference with canons is that instead of being focused around any plot element, GoIs tie content necessarily to a character group. Canons and GoIs work in distinct dimensions, meaning that Canons might feature one or more GoIs and vice versa.

3.2.2 Content tags

All SCP Foundation pages can be tagged. Tags are single words that contextualize a page and are set by the users when they create or modify the content.

At the time of this analysis, the SCP Foundation site has 2915 pages tagged as SCP reports, 1972 tagged as tales, 327 tagged as belonging to a specific canon and 898 as belonging to a specific GoI. Some canons (7 out of the existing 23) do not tag their pages, and therefore our count is not completely reliable regarding the amount of pages in canon subsets. Despite the fact that tags are not exclusive by definition, there is no overlap between SCP reports and tales, but

Fig. 2: SCP Foundation content tag distribution

there might be between the rest of categories. We can establish the basic distribution seen in Fig 2.

Table 4 enumerates the most common tags for pages that are not tagged as SCP reports or Tales.

Supplements and experiments are pages meant to complement the regular SCP series, therefore they fulfill a similar role. Pages tagged as **author**, **hub** and **goi-format** are meant to serve as structural hubs or index of sorts, connecting pages and presenting reading guidance to users. For instance, **goi2014** is a tag for one of the recurring writing contests run by the community, according to our data the most popular one. The workbench is meant for drafts to be rated and evaluated by the community.

Data found in Table 5 suggests that some contributions serve a distinct function in the context of a collaborative fictional story world. Some of the content (tagged as SCP series, experiments or supplements) documents the story world, functioning as *encyclopedic content*. Next, some of the content (tagged as tales) tells a story functioning as *narrative content*. Also, some of the content (tagged as **hub**, **author** or **goi-format**) provides structure to the site, and order to readers, functioning as a *hub or index content*. Part of the site, the *community comment* and discussion pages, provide community interaction but were not parsed and counted. This is because votes and user comments are attached to pages that are already serving another purpose, making the parsing and comparison troublesome. From all the previous observations, we summarize the following classification:

- Encyclopedic content: Content meant to document the story world with fictional information. (includes either tag SCP, supplement and experiment).
- Narrative content: Stories that take place inside the fictional story world (includes tale tag).
- index hub content: Navigation guidance for the story world content (includes hub tag).
- Community: Content meant for the discussion and criticism of the story world.
- Other content: Content that does not fit in any of the previous categories.

Fig 3 presents a graph for the SCP Foundation content according to its function (community feedback and discussion content is not included).




Fig. 3: SCP Foundation content by function.

Table 5 shows the tag distribution amongst the overall and main designated content sets by function. The encyclopedic and narrative content represents the majority of the site contributions.

Content links

We have used web parsing and the previous content by function designation to analyse the SCP Foundation's content link distribution. First we analysed the site contributions, following our previous functional categorization, to determine what content had the most incoming or outgoing links. The analysis was constrained to source and target pages from inside the SCP Foundation. Encyclopedic content receives the most links by a wide margin and index content has the most outgoing links, also by a wide margin. We also counted mentions to specific SCP series as additional outgoing links. This explains the high count detected in the index pages, often dedicated to explaining the role of the mentioned anomalies.




Fig. 4: SCP Foundation page incoming links

In order to determine how each segment is connected to the other ones we then break down the outgoing links. According to our results, links from the index to the encyclopedic content are the most common




Fig. 5: SCP Foundation page outgoing links




Fig. 6: SCP Foundation incoming and outgoing link averages

ones, followed by links from index to narrative and index content respectively.

3.2.3 Content canon

Story canons are an important part of any large fiction franchise, transmedia or multimodal narrative. They determine the level of reliability or how official are the different story contributions to a common story world. The main authors contributions' are generally considered more canonic, and subsequent contributions to the story world will be more coherent with them. Fan contributions are generally not official, and therefore, the main authors of the story world will ignore them. The SCP Foundation introduces a canon model that works on multiple levels and is both loose and strict simultaneously. All contributions belong to the top general level of the SCP Foundation background canon, acknowledging its existence and basic tenets. However, contributions might be interconnected with other contributions by the through the common author, group of interest or subcanon tags. This model might not be very intuitive for readers but allows authors to participate with a flexible measure of consistency and coherence with the rest of the story world. The purpose of having multiple levels of canonicity is to allow a flexible unobtrusive collaboration model for authors. We call this phenomenon *collaboration dimensions*.

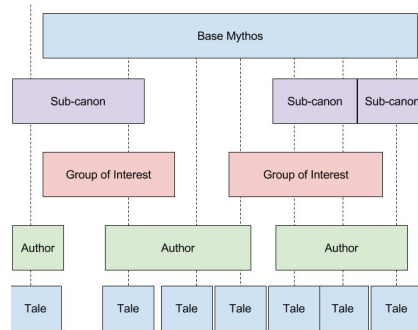


Fig. 7: SCP Foundation canon model

- *SCP Background*: a shady organization meant to secure, contain and protect us from paranormal anomalies (any contribution).
- *Sub-canon*: story world subset involving a specific theme or concept (content with the *canon* tag).
- *Group of interest*: fictional organizations with an specific interest in the SCP Foundation or its documented anomalies (content with the *goi* tag).
- *Author*: contributions from a common author (content by a common author).

All collaboration dimensions have a certain overlap, meaning a tale might belong simultaneously to the SCP background canon level, and an author's level, a specific interest group's level and an explicit sub-canon level. An author might contribute with a single contribution following the SCP background and ignoring or even contradict every group of interest and canon. Any contribution will belong to the SCP background and an author's level, yet its relationship with specific sub-canons and GoIs is flexible.

3.3 Activity

Our analysis revealed that the SCP Foundation has a relatively high amount of active contributors. Specifically 22 (Apr-May 2016), 31 (Mar-Apr 2016), 29 (Feb-Mar 2016), and 14 (Gen-Feb 2016). We consider a contributor active if she has contributed more than 4 times over the last month. For the top 1000 Wikia sites ranked by article amount, the average amount of active contributors (following our same




Fig. 8: Created and contributed pages per SCP author

criteria) is 5,7²⁵. Logged user activity data (seen in Table 7 and Table 6) also revealed that most of the content is created by a minority. Authors create an average of 2,17 pages with a standard deviation of 10,47 and a variance of 109,76. Regarding contributions authors contribute to an average of 12,83 pages with a standard deviation of 89,62 and a variance of 8032,72. The most prolific author created 283 pages, and the most active editor contributed to 1972 pages.

Our measures of the SCP Foundation author activity also revealed some interesting correlations that highlight the importance of social activity, such a mid-high correlation between page rating and discussion amount (0.7), a low-mid correlation between page discussion post amount and contributor amount (0.4) and a medium correlation between revision amount and contributor amount (0.6). Additional mid-high correlations include that pages that link more SCP reports have more ongoing revisions (0.5 correlation coefficient) and that pages that mention more SCP reports in their text without hyperlinking them also receive more revisions (0.7 correlation coefficient).

3.4 Summary and discussion

The quantitative and qualitative measures of the site content obtained through crawl analysis lead to a better picture of how the large-scale collaboration and hybrid nature of the site influence its content. While some of the analysis shows traits similar to other sites, there are differential (novel) aspects specific of the collaborative creative writing (and reading): we discuss the specific function that collaboration di-

²⁵ <http://wikis.wikia.com/wiki>

mensions (story canons and groups of interest) have with respect to consistency and coherence; and how hub pages guide the readership. These aspects probably explain the popularity of the site, which conceals its structural complexity behind attractive and simple mechanisms to appeal new and veteran users alike. We discuss as well the apparent contradiction of the small role played by characters.

The main implication we derive from the results is that the site's content is divided among exclusive encyclopedic, narrative, hub and community segments. The results agree with our function analysis; hub pages have the most outgoing links (they provide content guidance) and encyclopedic content has the most incoming links (it is used as reference material by other contributions).

The site emerged organically from the interaction of users on a paranormal message board. Apparently the urban legend theme (typically introducing an unsettling story set in a familiar context) might have been appropriate for the transition into an integrated story world due to the simple and familiar context and premise. The fictional SCP anomalies supposedly happen in our contemporary reality, not a distant, original invented world. Also, the usage of serialized standard templates (the SCP reports) seems to have promoted a certain consistency in terms of creative direction and reader expectations. In every SCP Series report, the authors present an original, intriguing and engaging read instead of simply presenting knowledge from an existing source to raise the interest of potential readers. This kind of content blurs the line between story and encyclopedia as it tries to document interesting and mysterious fictional knowledge in a structured, technical way.

The results helped us in understanding better the site's overall topology and link distribution, and reinforcing our content categorization according to its function. Since index pages have the most outgoing links and SCP Series reports have the most incoming links.

3.4.1 Multi-level canonicity

The term canon was used beyond its biblical origins (sacred texts accepted as genuine) to distinguish Sir Arthur Conan Doyle's original Sherlock Holmes stories from pastiches created by other authors. According to [9], a canon is "... a selection of texts that represent the supposed essence or highest quality examples of their forms or phenomenon". Typically, in modern transmedia narratives, a history is considered to be canonic if the original author acknowledges its existence in his creations-often because it was also written by him-. This

is often seen for instance in tv shows. Every season, its plot direction firmly controlled by the main writer and showrunner, drives the narrative forward, by often, there are other releases that expand the story in different directions. These secondary narratives, often written by other authors for other media, might be canonic (for instance a prequel to the main story) and will impact the next episode or season from the tv show. Alternatively, the production might be non-canonic (a sidestory) and will be ignored by future releases. Typically, these tales have vague connections to the main narrative, allowing the main author to ignore them without explicitly labeling as canonic or non-canonic. In a scenario where many authors use elements from a shared story world, with no clear author hierarchy, canonicity determines the level of accordance between all the interweaving stories. A strict canon model encourages all authors to produce coherent and consistent stories with each other, while a loose or nonexistent canon model overlooks plot incoherencies by allowing authors to contradict each other. According to Butler et al “the true power of wikis lies in the fact that they are a platform that provides affordances which allow for a wide variety of rich, multifaceted organizational structures” [6].

The SCP Foundation canon model makes collaboration an explicit multi-dimensional activity that makes possible for multiple authors to build a story world together, being consistent and coherent only up to the degree they prefer to without alienating much the readers. This model allows great flexibility when contributing to the SCP Foundation, it requires very little knowledge for newcomers but can establish sub-story worlds with complex inter-relationships of stories and narratives elements. The more layers the story belongs to, the more consistency constraints the author will find. Another strength of this model is that there’s room for incoherences and contradictions. It also helps in avoiding the problem of author territoriality and makes the overall reading and participation to the SCP Foundation more intuitive. From the reader’s perspective, this model provides a mean to navigate the SCP Foundation content by specific themes, groups and continuities. The tags found at the bottom of the page inform of what canon layers the page belongs to. The SCP Foundation’s canon model seems to be an instance of a complex multifaceted structure that takes advantage of the wiki software. The two main collaboration dimensions, Sub-canons and groups of interest, reduce the time investment required to participate in an ongoing story world such as the SCP Foundation, providing freedom for creative authors while remaining consistent to some degree.

In essence, the SCP Foundation delegates to the authors the responsibility or the power to enforce or ignore plot continuity through their interaction of the multiple canon model layers. While we might not go as far as to herald a new web paradigm, we believe this is certainly revolutionary for collaborative creative works. Participatory fandom often deals with the need for a canon model due to the implicit need to distinguish original story elements from the ones they create. For instance, the Lostpedia (the Lost TV series encyclopedia) at one point considered introducing a model based on the catholic one, featuring canonical, noncanonical, deuterocanonical, apocrypha and ex cathedr contributions [33]. The Lostpedia ultimately established a model that featured canon (original authors of the show), fanon (facts established by fans based on their inventions), theory (facts product of pure speculation) and parody. Canons are not necessarily expected by all users, the need for canons in literature has been questioned [42]. Canons have also been defined as “slippery things” [41] or as “a moment of the life of a story within a community” [46] Ultimately, the Foundation’s canon can not implement a model based on a religious one (it lacks a deity or prophet with authority to establish facts) or based on fandom (it lacks original authors and fans).

Collaboration dimensions and the SCP canon model seem to reflect the conclusions from [33], in the author’s own words “I want to conclude by highlighting the potential of the wiki architecture to overcome and blur boundaries and hierarchies between fiction and truth, canon and fanon.” In our opinion, the multi-layer canon model along blurs lines between realities and hierarchies, proposing an innovative perspective on fictional story world building that introduces middle ground for consistency and coherence, maintaining the fundamental illusion of one single reality in the story world to its users while allowing many partial ones.

3.4.2 Where are the characters?

As shown in Table 5 the SCP Foundation site has an extensive list of non-hierarchical tags for pages, however there is no tag for designating characters. Characters are a fundamental part of any story. Despite the presence of characters in tales, they are not fully featured with their own tag in the fictional encyclopedic content of the SCP Foundation. We found surprising that a site with hundreds of tales featuring characters does not use its encyclopedic capabilities to explicitly document such characters. This is especially rare when

compared to existing fictional wiki sites such as wiki of Ice and Fire²⁶ (a fan-created encyclopedia from George RR Martin's book series) where out of the more than 7000 articles half of them are character profiles or the One Piece Wiki²⁷ (documenting the Japanese cartoons and their Story World) where out of the 4450 articles more than 1100 are character profiles. Many anomalies reported on the SCP Series have human traits. *humanoid*, *sentient*, *alive* and *sapient* are some of the most common tags used in the SCP reports. The fantastic nature of the reported anomalies blur the line between object and character, meaning SCPs can in fact play a character role in tales or SCP reports. Despite this being explicitly discouraged in the site's writing guidelines, the extensive use of human properties in the SCP description might point towards the natural tendency of humans to focus narratives around human characters. Tales that feature SCPs with humanoid qualities often see those SCPs playing character roles with agency and an impact on the story. This is a bit misleading, mainly because the tags would be assumed to refer to humanoid characters by a casual reader, and instead they refer to alien, strange anomalies or entities. Given that the site emerged from SCP reports contributed into a paranormal message board, it is reasonable that most of the tags revolve around them. We would like to analyse more original shared story worlds in depth to see if this focus on one kind of element is a recurring tendency, or if it only represents an anomaly product of our reduced sampling.

The group of interest collaboration dimension also has a large role in the SCP Foundation. Its tags (e.g. *goi2014*, *dr-wondertainment*, *marshall-carter-and-dark*) are used very frequently in the tale segment. Groups of interest can also perform with agency and fulfill the role of characters, driving the story forwards while interacting with each other. Another interesting trait of these character groups, is that their control is not relegated to any author or group of authors. The actions of the group are constituted by the actions of its members, and generally, the creators of each character determine their actions. We believe the structural focus on groups of interest instead of characters is an intentional attempt to make a more approachable and inclusive story world for potential new contributors. The switch from individual protagonist to protagonist character groups should also help in switching the protagonism from one author to many. Despite the absence of character encyclopedic content pages, we believe

²⁶ <http://westeros.org/>

²⁷ <http://onepiece.wikia.com>

this has reinforced the collaborative authoring dynamics of the site. Author territoriality is a problem in a collaborative writing scenario as seen in an analysis of wiki collaboration for a fictional story world [19]. The same work discusses how characters are amongst the pages that gather more author collaboration and discussion. By replacing characters with inclusive organizations, the story world becomes more accessible and less territorial. This might help authors in creating stories in the story world that feel significant by making their original characters relevant, avoiding the need to use existing characters other authors might feel more attached to or protective about. The general implication for story worlds is that participative plot element groups are good to promote inter-author collaboration, bypassing their creative differences and democratizing the story control.

3.4.3 Author activity

Author activity data suggests a similar pattern to the 90% rule or 90-1-9 principle that states only 1% of the visitors to a website create content, 9% sporadically participate in the related discussions and 90% just observe. This is also related to the participation inequality coined by Nielsen [35]. Most content is created and maintained by a minority of contributors, who already are a minority amongst the visitors. This phenomenon is in line with [27] that explicitly claims that “Wikipedia articles are often maintained by a dominant few”. Encyclopedic content (the SCP report pages) represents the majority of the site. Index hub pages have a higher amount of contributors and revisions per page, they are less discussed, tagged or rated than the encyclopedic content pages. Index hub pages have a very high editing activity mainly due to the need to update them every time new content is introduced. Despite being a site meant to provide engaging fiction authors are more active in the fictional encyclopedic part of the site than the rest of it. This might be related to the user’s preconception of what a wiki site is used for. Another way to interpret this phenomenon is that creating SCPs is a similar exercise to other collage or pastiche creative activities that can be performed in social fan art sites such as DeviantArt²⁸. The main similarities are that there’s a part of reference material (the original intellectual property or inspiring media) a part of social interaction (usually in the shape of a social network) and some creative process involved. It is a common practice in fan sites for authors to introduce their own ideas (such as a new character description or drawing) that other

²⁸ <http://www.deviantart.com/>

users might use or link in their own contributions. These popular fan collaborations have some similarities to the writing paradigm found in the SCP Foundation; an author's original ideas might directly or indirectly influence other author's subsequent creations. The main difference is that there's no external intellectual properties explicitly referenced in the SCP Foundation; authors and readers are their own fans. Once again we suspect this creative process is something new that hints towards new authoring dynamics and has emerged thanks to the latest information technologies.

According to our measurements (Fig 4, Fig 5 and Fig 6 showing page connectivity), the SCP reports found on the fictional encyclopedic segment are consistently cited by tale pages and other SCPs. Traditionally wikis have always been used for citation and referencing, but given the creative nature of the SCPs, referencing might be considered as an innovative instance of co-authoring. An author often cites her own SCP's when creating new tale or SCP pages, however that it is not always the case, some of the references are made by distinct authors. SCP Report pages are used as reference material, receiving more links than pages from other segments. Our data also reveals that pages that link more SCP reports have more ongoing revisions (0.5 correlation coefficient). Pages that mention more SCP reports in their text without hyperlinking them also receive more revisions (0.7 correlation coefficient). This seems to follow a similar pattern than the one described by [21], that reinterpreted Wikipedia as a social network, claiming that social dynamics (including the page connectivity) are more important than direct collaborative aspects (such as the amount of contributors or page ranking, two magnitudes that we have failed to correlate to any relevant measures). This phenomenon also somewhat implies a collaboration trend in which authors link material rather than co-author it, collaborating indirectly. The tendency to cite a small set of old, well-established encyclopedic elements is reminiscent of how some disciplines structure knowledge (for instance biblical scholars favoring antiquity or author over other concerns such as internal consistency or coherence). Other works have highlighted how most participation seems to be focused on a few authors in wikis [26] and online communities [35]. Despite being an original fictional world, authors prefer referencing old, well-established encyclopedic content than newer or more narrative content. This is somewhat reminiscent of the principle of seniority, in which the oldest surviving knowledge is granted the most reliability. We found interesting that despite the modern nature of wiki platforms, with their hyperlinking capabilities, social affordances and semantic data, the Foundation

presents such tendency.

Due to the distinct nature of the message board, the threads were not included in our quantitative analysis. The Message board or forum is generally used for discussions relevant to the site, its content and policies. The discussion of specific content and its rating is done in the comments sections of each contribution. This includes the democratic editorial process of each contribution. Any subsequent attempt to perform an in-depth analysis of a site should include this data to elucidate the editorial process and dynamics.

4 A model proposal: The Open Story World

In this section we try to crystallize all the previous insight derived from observing hybrid story world sites. We propose a model meant with similar content types and author dynamics, the Open Story World. The OSW seeks to formalize large-scale collaborative sites meant for building story worlds while describing how these explicit collaborative story worlds are structured. It aims to provide a quantified measure of general *plot integration* to its content and participants, based of the recurrent use of specific elements and user activity. We base our measure of plot integration in the hypothesis that each contribution's author activity and connectivity with other contributions implies influence or integration into the overall plot consistency and coherence of the story world. In other words, *contributions that are well connected to other contributions and are actively modified, commented and positively rated by users will be perceived as more consistent and coherent*. It also introduces inclusive *collaboration dimensions*, a content subset meant for indirect collaboration, supporting flexible consistency and coherence in large-scale collaborative story worlds. Last, we briefly discuss the resulting hierarchical canon content distribution, the *canon pyramid*.

4.1 Contributions by function

Contributions to the OSW should fall into one of the following categories:

- *Encyclopedic contributions*
- *Narrative contributions*
- *Index contributions*

- *Community contributions*

These content segments follow the same guidelines than the ones we used in the SCP Foundation content analysis. Authors contribute with content that falls into one of the categories. These categories establish a synergistic loop that describes the collaborative creative process. Creating new content for the fictional story world, narrating tales, structuring content or interacting with fellow authors. New fictional encyclopedic content produces tales, and new tales produce new encyclopedic content. Index hubs provide structure to both collections and keep the site accessible to readers. Community contributions helps authors in staying close to the story world's continuity, reinforcing consistency and coherence.

The encyclopedic content should generally document the same sort of relevant arbitrary object (such as locations, weapons, reports or manuscripts), easy to use and reference. The standardized nature of the template provides an intuitive entry point for new authors. Also, having some sort of template guidelines helps in focusing the discussion of whether the content is appropriate or not for the OSW, facilitating agreement between authors. Despite being primarily informative, encyclopedic content is also an implicit form of fiction itself, since it describes fictional items. It might also resort to tropes such as unreliable narrators, intentional omissions and episodic arc structures to engage the reader without abandoning the illusion of rigor. Information gaps and omissions in the encyclopedic content provides opportunities for tales and more encyclopedic content. Centering the encyclopedic content on characters is not forbidden, but is not recommended. Instead, the objects of the documentation could be capable of acting with a certain agency, with the capability to drive the narrative forward.

Content from the tales category includes any form story or narrative that exists in the same fictional reality that the rest of the OSW. Generally, it will involve contributions from the encyclopedic content segment. The featured plot elements serve as linchpins for collaborations to emerge, and also as entry points for readers, not familiar with the story or the author but interested in the elements shared with the OSW.

Index hubs serve as a mechanism that links tales and fictional encyclopedic content amongst each other, providing entry points and contextual information to potential readers and assisting with the non-linear structure of the site. They help authors in organizing the content of the OSW and provide a congruent mechanic for readers

to explore it. Typical examples of index hubs include chronologies of tales, author pages, tag pages or thematic pages. The authoring of an index hub is a unique opportunity to frame a subset of the OSW in a specific way, providing a specific configuration to readers to influence their exploration for a specific purpose. The nonlinear nature of the OSW contributions provides a great deal of flexibility to the content of index hubs.

Community content exists with for the purpose of user interaction. The editorial process requires some sort of mechanic for communication. Also, to promote author participation, community content also introduces participation opportunities such as thematic cycles and contests. Community content also must resort to the best information technologies to blend into the content itself, providing appropriate context to keep the discussion focused.

4.2 Contribution and author plot integration

With many authors contributing in diverging directions, the main question still stands: What contributions are the most integrated into the story world plot, influencing its consistency and coherence rules? Based on our previous observations, we propose a discreet measure seen in eq.1.

$$\begin{aligned}
 \text{Individual OSW contribution plot integration} = & \\
 & \text{outgoing links} \\
 & + \text{incoming links} \\
 & + \text{discussion amount} \times \text{discussion participant amount} \\
 & + \text{number of edits} \times \text{number of editors} \\
 & + \text{reader score}
 \end{aligned} \tag{1}$$

4.3 Collaboration dimensions

Collaboration dimensions are public, participative subsegments of the OSW with a common plot element that links all of its contributions together with an enhanced consistency and coherence. OSW contributions might not belong to a single story or plotline, they coexist in a shared fictional reality that is weaved by the juxtaposition of all of them. In our previous observations, this coexistence in a collaboration dimension was implemented through a common semantic tag. This explicit tags not only inform of the link between contributions, but also provide navigation mechanisms between them, often through

an intermediate hub page. The dimension can represent many things, including in-universe factors, such as a common organization, location or timeframe or other community motivations such as author circles with common interests. An implicit advantage of using public collaboration dimensions is that they promote author collaboration, direct and indirect.

Following the same measurements we used before, we might determine the plot integration of a collaboration dimension by summing the integration of its contributions, as show in eq. 2.

$$\text{OSW collaboration dimension plot integration} = \sum \text{dimension contribution integration} \quad (2)$$

A discrete measure of author collaboration affinity can also be obtained through a similar formula that takes into account direct and indirect collaboration (eq. 3).

$$\begin{aligned} \text{OSW}_{\text{author}_1} \xrightarrow{\text{collaboration affinity}} \text{author}_2 &= \sum \text{co-edited contributions} \\ &+ \sum \text{co-commented contributions} \\ &+ \sum \text{collaboration dimension coincidences} \end{aligned} \quad (3)$$

4.4 Canon pyramid

Contributions from the OSW can be categorized into a hierarchy according their plot integration. If the content is to follow our previous analysis, the result should be exponentially distributed. The resulting measurements, the *Canon Pyramid*, reflect what content is more important to be consistent or coherent with, explicitly formulating the overall story world canon.

In the high level, contributions receive the most references, comments and edits and constitute the core backbone of the OSW's background. These contributions state the main rules to achieve coherent and consistent contributions with the overall background. Also, readers might base their expectations around these contributions. On the other hand, low level contributions receive the less scores, edits, comments and references. While still being technically part of the OSW's background, they have little influence on authors and their contributions or in readers and their general perception. This does not mean

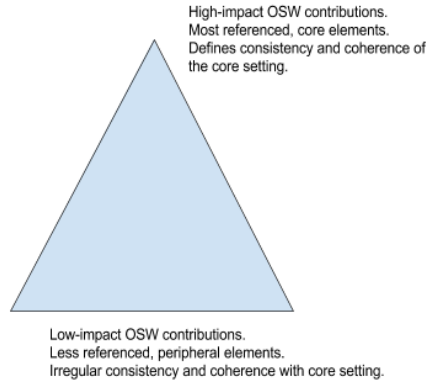


Fig. 9: OSW Canon Pyramid

their consistency or coherence with the high plot integrity OSW contributions is necessarily low, but they are more likely to deviate from the overall fictional reality established by the site.

4.5 Contribution function and antiquity

Our provided measures omit a critical part of our previous analysis. According to their function, contributions are connected more or less often. Also, the oldest content is often the most referenced one [51]. The measure of plot integrity should take this into account by weighting more these part of the equation. The specific weight values should be based in a quantitative analysis. We omitted including them due due to the limited scope of our data. The weights should be determined after including more similar sites into our analysis.

5 Discussion

The internet has been used consistently by online communities to engage in innovative creative activities. We have seen how some communities have mixed well established site types to create a new type meant to built original fictional story worlds and write stories that place in that same story world. The result is a hybrid site with a unique scope and goal. In these sites, the amount of participants with an active role in the original fictional story world (who either write, edit, comment or vote) is unprecedented in the history of creative collaborative writing. Despite the inherent complexities associated to building a story world with so many authors, with coexisting authorial visions and conflicting narratives, the usage of a participative

wiki-like site seems to make this possible. Making sense of such large collection can be very challenging, just like in the case of any “big data” scenario [31]. This is applicable for authors and readers, and even for we the researchers. The formal structuring of a collaborative site, moreso a wiki site, is a great opportunity for the study of user information interaction, with many implications for multiple fields such as social activity, knowledge and culture. Our contribution has been focused on fictional story worlds and their consistency and coherence, yet there are many potential directions for other studies and their potential generalizations at large for our society.

We observed a tendency in encyclopedic contributions of the SCP Foundation to omit critical information, contradicting the purpose of a traditional encyclopedia trading informational completeness for engaging reads. In terms of consistency and coherence, this ad hoc strategy might have solved one of the main challenges for establishing a collaborative story world, yet as previously stated, this strategy might not work for a hard science fiction section like the Orion’s Arm Universe. Our preliminary observations of the Orion’s Arm Universe suggested that real science was used to establish the main rules, while contributions had to take into account previous entries, exponentially increasing the contribution cost for new, uninformed authors. While this approach might work, we lack the empirical data to check if it is realistically scalable. Overall, the idea that an information space requires vagueness and contradictions frontally conflicts with the general belief that optimal information systems are complete and logically stable. This characteristic trait of the SCP Foundation might be an obstacle in formalizing its contents’ coherence and consistency in a potential model. The OSW avoids this problematic by simply basing its measurements in content referencing, content connectivity and author activity. Therefore, the OSW provides a measure of plot integration based on magnitudes we hypothesize to be conceptually related to consistency and coherence.

The classification and organization introduced by the OSW model is a first step towards in formalizing collaborative story worlds and their consistency and coherence. Our study was limited by the nature of the medium (a wiki site) and our resources, therefore lacking reader data that could help us in ensuring the integration measure reflects not only author or contributor criteria, but also reader. These measurable magnitudes allow us to apply formulas (based on the previous categorization of contents) to determine each contribution and each author’s overall plot integration in the story world contained in the OSW. The resulting distribution for contribution plot integration

was exponential, pointing towards a very reduced and influential core contribution set. This result is reminiscent most of our current on-line information spaces and also other traditional and modern story worlds, such as a specific series's Wikia or the ancient mythology from a specific culture.

6 Conclusions and Future Work

The evolution of participative creative writing sites and story world collaborative documentation has produced an innovative hybrid site. The SCP Foundation structure reflects the essential nature of a hybrid encyclopedic and narrative site while extending this stipulation with additional index and community content segments. The SCP Foundation introduces participative collaboration dimensions (groups of interest and sub-canons), meant for a guided exploration of its contents and flexible consistency and coherence to promote author participation. The resulting canon model is implemented with page tags. Characters, typically central in a story world, have delegated its protagonism to the central entity of the SCP Foundation (SCP anomalies) and participative collaboration dimensions (groups of interest). We introduced a model that structures collaborative story worlds, introducing content functions and dynamic collaboration dimensions along with a flexible canon model. The model also provides a measure of consistency and coherence through its formulas for contribution, author and collaboration dimension plot integration. A formula for inter-author collaboration affinity is also provided.

The SCP Foundation has been analyzed in depth. The obvious direction to generalize our claims is to perform a similarly rich analysis of other hybrid sites, such as *The Holder's Series*, *Galaxiki*, *The Wanderers' Library* and *the Orion's Arm Universe*. Our proposed model for an *Open Story World* could greatly benefit from the results. The sites studied in this work are relatively modern, but we suspect collaboration has taken new forms over the last years. Newer communities such as *thumblr*, *reddit* or similar social media might represent the latest incarnation of collaborative story worlds. Subsequent efforts in this direction should take into serious consideration newer online trends and compare the content and dynamics with the ones we have analyzed. The study is lacking important data; detailed visitor metrics. Although our analysis was centered on the authoring dynamics, the visitor amounts and frequency could enrich our perspectives and improve many of our findings. Obtaining such data might necessarily involve the community of the observed site, but could justify the

effort with improved results. We still suspect there are group dynamics to be revealed, specifically the existence of active author circles and hierarchies. Despite our readings failing to support our theory, A clustering process using inter-author affinity might help in isolating these hypothetical collectives and their activity.

References

- [1] Elza. Adamowicz. *Surrealist collage in text and image : dissecting the exquisite corpse*. Cambridge University Press, 1998.
- [2] Bryan Alexander and Alan Levine. Web 2.0 Storytelling: Emergence of a New Genre. *Educause Review*, 43(6):40–56, 2008.
- [3] Thomas Arnold, Johannes Daxenberger, Iryna Gurevych, and Karsten Weihe. Is Interaction More Important than Individual Performance?: A Study of Motifs in Wikia. In *WWW (Companion Volume)*, pages 1609–1617. International World Wide Web Conferences Steering Committee, 2017.
- [4] Rebecca W. Black. Fanfiction Writing and the Construction of Space. *E-Learning and Digital Media*, 4(4):384–397, dec 2007.
- [5] Paul Booth. Narrativity and the narrative database: Media-based wikis as interactive fan fiction. *Narrative Inquiry*, 19(2):372–392, 2009.
- [6] Elizabeth Butler, Brian and Jacqueline Pike. Don't look now, but we've created a bureaucracy: the nature and roles of policies and rules in wikipedia. *Proceedings of the SIGCHI conference on*, pages 1101–1110, 2008.
- [7] Caroline Courbières and Sabine Roux. Wikia: Between Documentary Simulacra and Documented Fictions. volume 2, 2015.
- [8] N. DiFonzo and P. Bordia. Rumor, Gossip and Urban Legends. *Diogenes*, 54(1):19–35, feb 2007.
- [9] Mark Duffett. *Understanding fandom: an introduction to the study of media fan culture*. 2013.
- [10] Lisa S. Ede and Andrea A. Lunsford. *Singular Texts/plural Authors: Perspectives on Collaborative Writing*. Southern Illinois University Press, 1990.
- [11] J. Lloyd Eldredge, D. Ray Reutzell, and Paul Hollingsworth. Comparing the effectiveness of two oral reading practices: Round-robin reading and the shared book experience. *Journal of Literacy Research*, 28(2):201–225, jun 1996.
- [12] Despoina N. Feleki. Wikia Fandom Craze: Connecting, Participating, Creating, and Re-negotiating Boundaries. *Gamma: Journal of Theory and Criticism*, 23(0):45–59, 2016.

- [13] John Fiske. The Cultural Economy of Fandom. In *The Adoring Audience: Fan Culture and Popular Media*, pages 30–49. Routledge, 1992.
- [14] Christian Fuchs, Wolfgang Hofkirchner, Matthias Schafranek, Celina Raffl, Marisol Sandoval, and Robert Bichler. Theoretical Foundations of the Web: Cognition, Communication, and Co-Operation. Towards an Understanding of Web 1.0, 2.0, 3.0. *Future Internet*, 2(1):41–59, feb 2010.
- [15] Jonathan Gray and Jason Mittell. Speculation on Spoilers: Lost Fandom, Narrative Consumption and Rethinking Textuality. *Particip@tions*, 4(1):1–28, 2007.
- [16] J. Grieve. Quantitative Authorship Attribution: An Evaluation of Techniques. *Literary and Linguistic Computing*, 22(3):251–270, may 2007.
- [17] Said Hadjerrouit. Wiki as a collaborative writing tool in teacher education: Evaluation and suggestions for effective use. *Computers in Human Behavior*, 32:301–312, 2014.
- [18] A. Paul Hare. *Creativity in small groups*. Sage Publications, 1982.
- [19] Rik Hunter. Erasing ” Property Lines” A Collaborative Notion of Authorship and Textual Ownership on a Fan Wiki. *Computers and Composition*, 28(1):40–56, 2011.
- [20] David W Johnson and Roger T Johnson. Research shows the benefits of adult cooperation. *Educational Leadership*, 45(3):27–30, 1987.
- [21] Gerald C. Kane. It’s a Network, Not an Encyclopedia: A Social Network Perspective on Wikipedia Collaboration. *Academy of Management Proceedings*, 1(1):1–6, aug 2009.
- [22] Tony Keen. Are fan fiction and mythology really the same? *Transformative Works and Cultures*, 21(21), 2016.
- [23] Greg Kessler and Dawn Bikowski. Developing collaborative autonomous learning abilities in computer mediated language learning: attention to meaning among students in wiki space. *Computer Assisted Language Learning*, 23(1):41–58, feb 2010.
- [24] Kanta Kochhar-lindgren, Davis Schneidermann, Tom Denlinger, and Nebraska Press. *The Exquisite Corpse . Chance and Collaboration*

- in Surrealism 's Parlor Game. *Papers of Surrealism*, (9):1–7, 2011.
- [25] Annette Lamb and Johnson Larry. an information skills workout: wikis and collaborative writing. *Teacher Librarian*, 34(5):57, 2007.
- [26] David Laniado and Riccardo Tasso. Co-authorship 2.0: patterns of collaboration in Wikipedia. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia - HT '11*, pages 201–210, New York, New York, USA, 2011. ACM Press.
- [27] Jung Lee and DongBack Seo. Crowdsourcing not all sourced by the crowd: An observation on the behavior of Wikipedia participants. *Technovation*, 55-56:14–21, sep 2016.
- [28] Jun Liu and Sudha Ram. Who Does What: Collaboration Patterns in the Wikipedia and Their Impact on Article Quality. *ACM Trans. Manage. Inf. Syst.*, 2(2):11:1–11:23, 2011.
- [29] Sarah Ann Long. Exploring the wiki world: the new face of collaboration. *New Library World*, 107(3/4):157–159, mar 2006.
- [30] Andreas Lund. Wikis: a collective approach to language production. *ReCALL*, 20(01):35–54, jan 2008.
- [31] Lambert M. Surhone, Mariam T. Tennoe, and Susan F. Henssonow. *Big Data*. 2010.
- [32] Mildred B. Mitchell. Trends Toward Multiple Authorship in Scientific Publications. *The Journal of Psychology*, 52(1):125–131, jul 1961.
- [33] Jason Mittell. Sites of Participation: Wiki Fandom and the Case of Lostpedia. *Transformative Works and Cultures*, 3(3):1–10, 2009.
- [34] Gregory B Newby. *Metric Multidimensional Information Space*. 1996.
- [35] Jakob Nielsen. Participation inequality: Encouraging more users to contribute. *Jakob Nielsen's alertbox*, pages 1–5, 2006.
- [36] AE Nudelman and CE Landers. The Failure of 100 Divided by 3 to Equal 33-1/3. *American Sociologist*, 7(9):9, 1972.
- [37] Angela M. O'Donnell, Donald F. Dansereau, Thomas Rocklin, Judith G. Lambiotte, Velma I. Hythecker, and Celia O. Larson. Cooperative Writing: Direct Effects and Transfer. *Written Communication*, 2(3):307–315, jul 1985.

- [38] R A Y Over and Susan Smallman. Maintenance of Individual Visibility in Publication of Collaborative Research by Psychologists. *American Psychologist*, 28(February):161–166, 1973.
- [39] Ray Over. Collaborative research and publication in psychology. *American Psychologist*, 37(9):996–1001, 1982.
- [40] Olivia Riley, Susan Riley, and Meredith Gill. *Archive of Our Own and the Gift Culture of Fanfiction*. PhD thesis, 2015.
- [41] AR Rotruck and W Brooker. Using the Force: Creativity, Community, and Star Wars Fans. 2005.
- [42] Peter Seixas. Who Needs a Canon? In *Beyond the Canon*, pages 19–30. Palgrave Macmillan UK, London, 2007.
- [43] Shlomo Sharan and Yael Sharan. *Small-group teaching*. Educational Technology Publications, 1976.
- [44] Elana Shefrin. Lord of the Rings , Star Wars , and participatory fandom: mapping new congruencies between the internet and media entertainment culture. *Critical Studies in Media Communication*, 21(3):261–281, sep 2004.
- [45] Gabriele Siegert, Kati Förster, Sylvia M Chan-Olmsted, and Mart Ots. Media Branding 3.0: From Media Brands to Branded Entertainment and Information. In *Handbook of Media Branding*, pages 11–32. 2015.
- [46] Gerrie Snyman. Intertextuality, story and the pretense of permanence of canon. *205 OLD TESTAMENT ESSAYS*, 8:205–222, 1995.
- [47] Neomy Storch. Collaborative writing: Product, process, and students’ reflections. *Journal of Second Language Writing*, 14(3):153–173, sep 2005.
- [48] Carola Strobl. Affordances of Web 2.0 Technologies for Collaborative Advanced Writing in a Foreign Language. *CALICO Journal*, 31(1):1–18, 2014.
- [49] Bongwon Suh, Ed H. Chi, Bryan A. Pendleton, and Aniket Kittur. Us vs. Them: Understanding Social Dynamics in Wikipedia with Revert Graph Visualizations. In *2007 IEEE Symposium on Visual Analytics Science and Technology*, pages 163–170. IEEE, oct 2007.

- [50] Yu-Chih Sun and Yu-jung Chang. Blogging to Learn: Becoming EFL Academic Writers Through Collaborative Dialogues. *Language Learning & Technology*, 16(1):43–61, 2012.
- [51] Alan Tapscott, Joaquim Colás, Valeria Righi, Carlos León, and Josep Blat. We Built Our Own Worlds - Story Canonicity and Indirect Collaboration in a Shared Story World. In *Proceedings of COLLA 2016 the Sixth International Conference on Advanced Collaborative Networks, Systems and Applications*, pages 30–38, Barcelona, Spain, 2016. ThinkMind.
- [52] Angela Thomas. Fan Fiction Online: Engagement, Critical Response and Affective Play through Writing. *Australian Journal of Language and Literacy*, 29(3):226–239, 2006.
- [53] J. R. R. Tolkien. Beowulf: The Monsters and the Critics. *Interpretations of Beowulf: A Critical Anthology*, page 31, 1936.
- [54] Yildiz Turgut. EFL learners' experience of online writing by PbWiki. In *EdMedia: World Conference on Educational Media and Technology*, volume 2009, pages 3838–3847, Honolulu, HI, USA, 2009. Association for the Advancement of Computing in Education.
- [55] Haiyi Zhu, Robert E. Kraut, and Aniket Kittur. The impact of membership overlap on the survival of online communities. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, pages 281–290, New York, New York, USA, 2014. ACM Press.
- [56] Harriet A. Zuckerman. Patterns of Name Ordering Among Authors of Scientific Papers: A Study of Social Symbolism and Its Ambiguity. *American Journal of Sociology*, 74(3):276–291, nov 1968.

Tab. 1: Hybrid site descriptive statistics I

	Creation year	Theme	Wiki engine	Editorial control	Reader feedback	Community interaction
The Holder's Series	2007	Urban legends, cursed artifacts	Custom built	Automatic acceptance, site staff approval	Rating system	Message board
Galaxiki	2007	Sci-fi universe	CornelIOS	Automatic acceptance, profanity check	Comment and rating systems	Message board, blog
The Orion's Arm Universe	2000	Sci-fi universe	Custom built	Public election, science and logic focus	None	Message board, contests
The SCP Foundation	2008	Reports on paranormal anomalies	WikiDot	Public election, minimum rating	Comment and rating systems	Message board, contests, chat
The Wanderer's Library	2010	Collection of natural documents	WikiDot	Public election, minimum rating	Comment and rating systems	Message board, contests, chat

Tab. 2: Hybrid site descriptive statistics II

	Version control system	Multi-user contributions	Encyclopedic content	Specialized encyclopedic content	Narrative content	Dynamic hubs	Semantic tags
The Holder's Series Galaxiki	No	Not supported	660 holders	Yes	~200 stories	No	No
The Orion's Arm Universe	No	Not supported	47 solar systems ~1000 assorted articles	Yes	~77 stories	No	No
The SCP Foundation	Yes	Supported	2900 anomalies reported	Yes	Commercial books and magazines 1927 tales	Yes	Yes
The Wanderer's Library	Yes	Supported	307 collected documents	Yes	24 archive tales	Yes	Yes

Tab. 3: Hybrid site activity

	The Holder's Series	Galaxiki	The Orion's Arm verse	The SCP Foundation	The Wan- derer's Library
Monthly visits	29300	18200	38500	2800000	18900
Avg visit duration	6:46	0:06	6:10	11:25	0:28
Avg pages per visit	7.17	1.16	5.65	11.34	1.5
Backlinks	801	550	1000	1100000	3200

Tab. 4: Misc page tag incidence

Tag Name	Page Amount	Overall Percentage of Others
supplement	266	25%
author	229	21,60%
hub	87	8,20%
experiment	68	6,40%
goi-format	62	5,80%
goi2014	51	4,80%
workbench	38	3,10%

Tab. 7: SCP page author amounts and averages

	Overall Pages	Encyclopedia	Narrative	Index
Page Amount	5975	3182	1972	318
Participating Author Amount	2670	2100	963	766
Avg Created per Author	2,17	1,16	0,71	0,14
Avg Contributed per Author	12,83	7,26	3,43	1,12

Chapter 5

Discussion

5.1 Consistency and coherence in a collaborative story world

In this thesis, we often use the terms consistency and coherence in the context of a story world. Consistency is used to refer to the usage over time of the same elements, be it plot elements (e.g., characters, locations and objects) or more stylistic resources (e.g., plot, tone, theme or tropes). Coherence is used to refer to the agreement of facts in the fictional story world, whether the succession of events is logical in its context. Indeed, most of our initial observations and inquiries pointed towards an interest in the preservation of consistency, but the role of coherence was not so clear. While very few of the observed subjects questioned whether the story “made sense” the contributions generally adhered to the plot and tone not only consistently, but also coherently. Although both measures are highly subjective, consistency is apparently easier to preserve by simply ensuring the same elements are used over time. Without going much into the nature of these concepts, we suspect most of our subjects purposely avoided to question or discuss whether the unfolding story was logical, therefore, they rarely mentioned coherence at all. However, in hindsight, we believe their actions were also trying to reinforce the internal logic of the story.

One of the first and more significant findings of this thesis, was that in the multiple experiments we conducted to build a story set in a collaborative story world, participants perceived one single or “canonical” version of facts. Despite their power to extend the story, they very rarely contributed contradicting existing facts. This has several implications for the collaborative construction of fiction. Whatever

initial story we introduce will have a deep impact on the results. We very rarely observed any subject rewriting existing content, and in all the sites analysed that supported collaboration, a special consideration was awarded to the original author. If authors tend to write coherently and consistently, and the initial content cannot be modified, it will have a substantial impact on all subsequent contributions. The sequence of contribution, therefore, is essential in establishing a collaborative story world. This might explain why our analysis of a popular story world building site suggested that a few, older content was more referenced and ranked with a higher score. In a certain sense it is also reminiscent of the problems mentioned by [28] in their Wikinovel, involving authors colliding as they attempt to impose a main plot structure. Maybe, if some author laid the foundation for the tale, instead of relying on the free interaction of the community, the result would have been more bearable as a novel. However, this would also compromise its unique nature. When compared to our own research it seems that this phenomenon is similar to previous results [43, 45] in which authors converged, establishing a main continuity or central interpretation of the story. Another implication of this assumption of one single fictional reality, is that as more authors contribute to the story world, they might need to take into account all the previous contributions. This also implies that the scalability of a collaborative story worlds is troublesome. Our analysis of the SCP Foundation suggested the existence of *collaboration dimensions*, content subsets with enhanced internal consistency and coherence. These subsets allow authors to contribute only partially acknowledging the rest of the contributions. They represent a viable strategy to reduce the scalability issue. However, they can produce an isolating effect, associated to author territoriality [23] or “walled gardens” [28]. It is a trade off of sorts, in which highly consistent and coherent stories must take into account more contributions, increasing the cost for its author (forcing him to read more).

During the early stages of our study, participants seemed to be concerned with the preservation of consistency in an evolving collaborative story. In the first experiment, according to their questionnaire answers and our in-situ observations, some participants focused their contributions on correcting discontinuities in the overall narrative consistency, preserving what appeared to be causality chains. Particularly, we were surprised to observe that the only modifications done to other participants’ original works, were made with the purpose of reorganizing the story sequence and to be consistent with facts introduced by other contributions. However, in the following ex-

periments, participants were not concerned by consistency as earlier. This switch in participant behavior was apparently triggered by the modifications we did to the experimental procedure. We changed the platform from a highly visible, public participative glass wall, to software prototypes. Despite trying to replicate the original experience as much as we could, the transition was still a substantial modification to the process of collaboratively creating a story world. After reviewing the results, the coded interaction and each subject's questionnaire and interview answers we had a clearer picture of how the context impacts collaborative construction of stories set in a common story world. When participants saw each other (in the first experiment) they assumed the process was evolving and felt more responsible in developing it in a consistent way. In the following experiments, set in a computer prototype with no visible presence of the collaborators, participants were less aware of the collaborative nature of the procedures, and cared less about being consistent. Also, our prototype designs in this stage of the research were focused on providing a good reader experience, therefore providing a zoomed perspective, narrow compared to a wall any bystander can approach freely. This stance towards consistency was more less the trend in our following experiments involving computer prototypes and isolated users. For this reason, towards the end of this thesis, we introduced partially synchronous groups with awareness of each other. Despite the results pointing towards certain visibility, we failed to observe the same high interest on maintaining the story's consistency in any of our subsequent works. Not even the introduction of an underlying model of consistency, the usage of semantic diagrams or our analysis of hybrid sites for collaborative story worlds replicated the observed concern for consistency. Our hypothesis at this point, was that the limited scope of the initial story and the scenes disposition on the wall were interpreted as an implicit limitation on the type of possible contributions. The only data backing this hypothesis was that the first experiment results were mostly introduced into two main storylines, the same two initial ones.

The inclusion of a formal model to enhance consistency produced two main effects. First, users who were collaborating to the story were bothered by its presence. The negative reception was mitigated when we switched from hard constraints to usability recommendations. On the other hand, the produced stories were perceived as more consistent. Ultimately, we introduced this model to make the authoring task less focused on consistency and produced a worse user experience. Some of our results suggest consistent and coherent stories

are received better, but it is clear that the prototype's user experience was annoying to potential authors. In hindsight, this is hardly surprising, especially after our exploration of computational creativity and the technical limitations it entailed. A knowledge-intensive approach-similar to the we tested-could be implemented for consistency, but judging from our previous results, it might be challenging. A potential approach could involve relatively objective measures of consistency (characters moving through a measured space). Regarding our first attempt at a formal model of consistency, there was no clear agreement on its validity for the unfolding narrative. Involving the authors in the definition of a hypothetical model could help in producing a tool that agrees with a specific demographic. Our analysis of popular sites meant for the creation of story worlds, suggests that allowing users to freely interact and discuss can provide a certain measure of agreement, and perhaps, of consistency and internal coherence.

5.2 Thematic contribution and navigation

In our first experiment, the branching story had narrative tropes, typically used in fairy tales and space opera. Interestingly, the plot developed following two parallel branches, one with plot, tone and tropes closer to a fairy tale and another one with plot, tone and tropes closer to a space opera. After close observation, we determined that each branch not only had its own protagonist, but was fundamentally an independent story with common links to the other one. From a structural perspective, both stories followed a classical hero journey, but their tone and plot elements were different. In other, more recent experiments, we observed how even within a shared story world, there is a tendency to establish independent story arcs with its own recurring elements. For instance, in *Chronoverse*, certain scene tags were associated to specific subsets of the story world. The twilight choir tag was used for scenes that involved dark and surreal themes, while Pompeu runners was used to portrait heroes, hope and themes of rising against an oppressive rule. While this was somewhat implied by the original setting provided to the experiment participants, the contributions included prominently these factions and their stories and characters were structured around them, apparently in a consistent and coherent way. This was similarly observed in the SCP Foundation site and their usage of tags, particularly for the collaboration dimensions: groups of interest and canons.

Our observations regarding how the story was structured around characters traditionally associated to the protagonist role—the prince, the princess, and to a lesser extent the ninja—, motivated us to explore narrative and character arcs in the context of collaborative story worlds. We introduced explicit storylines into our subsequent prototypes to analyse how they were used. Both, general purpose storylines and character-specific storylines were popular amongst readers and writers. We were surprised as we failed to replicate this observation for characters in real large-scale collaborative story worlds. Despite the existence of reading guides, or storylines of sorts, these were not structured around characters. Apparently, these stories are structured similarly to the classic character arcs, but instead of being structured around a character, they are centered on assorted entities that can act with agency. What is more, these central entities (e.g., planets, locations, supernatural objects or anomalies or manuscripts) are used by multiple authors partially avoiding any territorial issues. Also, inclusive collaboration dimensions (e.g., character factions or solar systems) are effective in integrating other secondary creations by assorted authors (e.g., characters, planets). The resulting knowledge subsets, the collaboration dimensions, can be then used by readers for a more accessible and rewarding guided thematic navigation and contribution of the story world .

Our hypothesis (backed by our experimental data and analysis of a popular site) is that collaborative story worlds require consistency and coherence, but as they scale, authors need to structure knowledge around multiple thematic dimensions that neither overlap nor contradict each other. It is this phenomenon that apparently emerges naturally in this kind of context. Our proposed model, the Open Story World, attempts to crystallize these assumptions into a model that explains this phenomenon and provides a measure of plot integration for the authors, contributions and contribution dimensions. To some extent, this hypothetical information structure is not very distinct than popular social networks that deal with trends, such as Twitter and its use of hashtags to structure the large-scale nonlinear conversations. It also avoids the the lack of structure in a collaborative environment described by some authors [25, 4].

5.3 Inter-author interaction

Another of the most recurrent observations was the reluctance of authors to modify contributions from other authors. The main exception

to this phenomenon were authors who played more a “curator” or “editor” of sorts, reorganizing the story or introducing minor revisions of an existing contribution to preserve the global consistency or coherence. Authors feel uneasy when modifying contributions originally created by others, especially if there is no objective reason, such as following the site’s guidelines or correcting a typo. The role of editor might be something more natural in our culture than a direct collaborative writing. Despite the author’s influence over each other is more or less direct in a collaborative story world (and the degree of influence is highly subjective), we found very few instances of authors creating the same story or plot element collaboratively, splitting the writing with each other. This includes both our experiments and the analysis of the SCP Foundation. On the other hand, our analysis of popular story world sites showed how users participate in debates regarding the validity of contributions and overall enjoy acting of peer reviewers or just commenting submissions. Large-scale democratic reviews, made possible thanks to information technologies, might be one of the most innovative and functional tools for the collaborative creation of story worlds. In this sense, the inclusion of an standardized template for the encyclopedic content in hybrid narrative sites, provides the rules of interaction for participative editorial processes.

Generally speaking, most authors participating in our experiments were not very interested in reading all the existing contributions and even less to integrate their contributions directly with anyone. This, might be partially due to the limited visual scope of a software prototype compared to the glass wall. Even with the assistance of content filtering, dynamic zooms and storylines to guide their reading, our attempts to create a story collaboratively in a computer setting were unsuccessful in encouraging the exploration of fellow author’s contributions or the direct integration and crossing of stories. The collaboration, however, was happening indirectly, as they shared plot elements and integrated their contributions to shared collaboration dimensions. Even if this process is not completely aware to the contributors, the result is a relatively consistent and coherent story world with its associated stories that follow common themes (established by the collaboration dimensions).

Demography might be influencing our experimental results. On one hand, experienced fiction writers might be more proficient at writing stories and building story worlds, potentially being more motivated to contribute more due to the professional background. On the other hand, students might be more open to novel scenarios, such as collaborating in an online platform to build a story world together.

5.4 Author-story world interaction

As previously discussed, contributing to a collaborative world can be tricky and writing with other authors requires some coordination. Our experiments with semantic diagrams, using them as planning tools, complementing written text were promising. They performed efficiently in terms of time and contribution, and are good vehicles for communicating a story plan. They increase the comprehension of readers when paired with text and were considered a flexible and intuitive collaboration mechanism. Despite its shortcomings (such as confusing relationship cardinality or reliance on some sort of zooming and filtering mechanism) they might help in enabling authors to collaborate more directly when building a collaborative story world. Particularly, in large-scale scenarios, they could help by providing visual support to the complex, multi-dimensional information, such as in the case of the SCP Foundation. In such scenario, explicit links or tags are popular, and might be good candidates to become integrated into some sort of semantic visualization. Chen's original relationship role labelling or Corman et al approach [11], introducing directed networks of relations, could make the E-R model more understandable for story planning. The main limitation of the E-R model for story modeling, is its inability to represent the progression of the story, with the transformations of its entities. This means that it might be a good tool to represent a snapshot of a story world, but not to represent its progressions. Our tests also neglected the requirements of a large-scale story world. If used in a context similar to our latest experiments and real world analysis, the diagram might require additional features such as zoom, filtering and clustering.

Most of our prototypes resorted to a timeline visualizations of some sort to render information. The sequential metaphor (backed with tags to inform of collaboration dimensions) has performed well across all our experiments. It is also occasionally used in the SCP Foundation site to organize data. The syntax of the E-R model could be easily paired with a formal representation of the story world and allow authors to interact directly with the rules that govern the consistency and coherence principles. Incidentally, the only case in which the subjects complained about the lack of reading order was when there was no timeline, that is, the experiments with E-R diagrams. Perhaps, the combination of both, could be a good way to visually represent the story world information.

5.5 Creativity, coherence and consistency

This thesis has been focused on author interaction and the role of story consistency and coherence. Creativity is a fundamental part of any artistic exercise, including the collaborative construction of story worlds. We suspect that this focus on establishing consistent and coherent stories might neglect the need for stories to be interesting, distinctive or memorable amongst many other additional factors related to entertainment. There is an implicit trade off between being consistent and original. Stories that are consistent are resorting to the same recurrent elements, and therefore, are repetitive up to some degree. So, there is an inevitable question; how creative can we be if we want to stay in a common story world? The answer is beyond the scope of this thesis, but there are a few leads worth mentioning. Our Chronoverse results suggest that there was a moderate correlation between a contribution score, its consistency with the initial story and overall usage of tags. Our interpretation is that the jury enjoyed stories that were well integrated into the story world. Following our previous reasoning, this implies that they were enjoying more scenes that were fundamentally more similar. This argument questions the need for creativity in the context of a collaborative story world. Perhaps users understand the distinct nature of the contributions and reduce their expectations of surprise. Without delving much more into this discussion, we want to remark that certainly there is much to be explored regarding novelty, originality and reader enjoyment. This is partially reminiscent of how some contemporary movies that take place in a shared story world, are generally bashed by critics (who might be looking for a more innovative or distinct product) and perform outstandingly in the box office (provably thanks to loyal fans that enjoy the overall continuity).

5.6 Innovative directions for collaboration in story worlds

Our foray into the field of computational creativity and the generation of narrative artefacts was used to explore the computational approach to a creative practice, with its many potential implications for digital platforms meant for authors. The main motivation was to formalize a perceived magnitude-creativity-that is generally considered subjective, conceptually close to the focus of this thesis-consistency and creativity-using a familiar methodology-crowdsourcing questionnaires-

. We concluded that there is a strong correlation between narrative potential and perceived overall quality of a creative text artefact. This observation is similar to the evaluation of our longitudinal experiment, in which the jury consistency ratings were positively correlated to reader quality ratings. Breaking down the perceived quality into more specific parameters should help elucidating its implied meaning in the context of creative writing and how it relates to other subjective measures we attempt to formalize. We were unable to isolate a set of metrics that help us in predicting and assisting creativity, however, the results were stable for the artefacts that used the most familiar language and structure. The main implication is that there is a collective or converging conception of creativity, an observation we have been able to replicate in our experiments, in which subjects were in agreement when assessing consistency or coherence. The secondary implication is that assessing consistency, a magnitude based on the recurrent use of elements (therefore making them more familiar to the reader), might be more viable and stable. Understanding the reason for this inability to isolate critical factors involved in the perceived creativity is paramount if we are to formalize other relevant measures to a collaborative story world. Automatically generating stories and computing useful values for metrics is heavily dependent on the available knowledge networks. This makes it necessary to address knowledge management from a different perspective.

Chapter 7

Conclusions and future work

7.1 Conclusions

This thesis attempted to determine the role of consistency and coherence in the context of a collaborative story world. The main result is that collaborative story worlds are expected to be coherent and consistent by readers. This is related to another finding, the general perception that story worlds have one single reality, continuity and interpretation of facts. This generalized perception is used as a cornerstone around which consistency and coherence are evaluated. Our attempts to introduce a formal model for consistency based on reader perception (focused on character location over time in a measured space) in the creation of a collaborative story worlds were poorly received by authors but helped in creating stories perceived as more consistent. Despite the subjective nature of perceived consistency and coherence, we were partially successful in predicting and enhancing them in limited scope of our prototypes.

We learned more on how large amount of participating contributors achieve consistency and coherence in collaborative story worlds. By compartmenting knowledge in specific subspaces or collaboration dimensions centered around specific topics or themes of the story world, authors can intertwine their contributions without being too intrusive with each other. The nonlinear nature of the story world allows them to link their contributions following relevant themes which might be from inside the story world's fiction (e.g., factions, topics, timeframes and locations) or the contributing community (e.g., author circles, contests or seasonal celebrations). This multi-dimensional

continuity allows to establish a canon that is flexible and solid simultaneously, containing contributions that might be loosely connected, deeply intertwined or that just ignore each other. This flexibility is achieved thanks to the usage of the collaboration dimensions, in which authors may acknowledge or ignore other existing contributions as they please while remaining consistent (by explicitly referencing specific contributions or collaboration dimensions) and coherent (by integrating the facts from the referenced content into their discourse). Authors can acknowledge or ignore other existing contributions as they please while remaining consistent (by using referencing recurrent encyclopedic content) and coherent (by agreeing to the facts established by the encyclopedic content they reference). Also, by introducing specific narrative tropes into the story world contributions, such as unreliable and subjective narrators and intended omissions to rise the reader's curiosity, the story world can cope better with conflicting information. Additionally, the participative nature of modern sites, allows for rich discussions and democratic editorial processes to take place. We also observed how most content in a large-scale story world either belongs to the encyclopedic, narrative, indexing or community segments. These content groups are either meant to respectively document knowledge, narrate a tale, organize information under collaboration dimensions or promote the community interaction dynamics.

Navigation in collaborative story worlds seems to closely resemble traditional story arcs. In our small lab experiments we saw a predilection for character arcs, however, in real large-scale hybrid narrative sites, the focus of the arcs used for navigation seemed to switch to the main encyclopedic content, which incidentally contains few characters. Instead, factions or other inanimate, less human elements were used both, to structure the bulk of the encyclopedic knowledge and the site navigation. This seems to be meant to encourage participation and avoid author territoriality (a phenomenon closely related to the authorship of characters). When contributing to a collaborative story world, authors rarely modify contributions originally created by others. Authors also prefer to reference and integrate their contributions with the oldest ones. This implies that the initial story and the sequence of collaboration are essential in establishing the base rules that will determine the general perceived consistency and coherence. Finally, we suspect that author presence, either physical or digital, seems to reinforce overall interest in preserving the story consistency and coherence.

7.2 Future Work

Future attempts to assist authors in building consistent and coherent multi-authored stories could also involve the authors in the construction of a formal model that reflects their story worlds. The highly subjective nature of the task qualifies them for more optimal results. Another possible improvement could be to provide better communication amongst authors to support their coordination and communication. This could improve cooperation during narrative composition and introduce specializations such as consistency, plot, tone or theme curators. Also, enhancing the digital presence and visibility of authors through proper visual metaphors could result in an enhanced community sense.

Most of our experiments' volunteers were college students. Our real story world site studies, while being a considerable step forward, featured online communities, which are very complex to control from the demographic point of view. We believe replicating the experiments with a demographic made of professional writers would help in generalizing the results. A potential candidate subject pool could be a team of authors writing a TV show season or a story anthology.

Regarding the semantic diagram use, we did not exploit the computational characteristics of the model to support either coherence or consistency of a story. Also, we did not experiment with long term collaboration and its scalability. A potential venue for future work could be to attempt to replicate the findings with these modifications. Also, semantic diagrams could support complex, multi-dimensional story worlds, such as the SCP Foundation. In this kind of context, explicit links and tags are popular, and might be good candidates to become integrated into some sort of semantic visualization meant for visual representation of entities and relationships. Particularly, a similar visual representation to the E-R model could be linked with a formal representation of the story world, allowing authors to interact directly with the rules that govern the consistency and coherence principles. Our study of computational creativity also opened some potential roads for future work. We could implement a similar methodology to assess or generate similar metrics that have some sort of correlation with perceived consistency or coherence in a collaborative story world. Also, we could try to assess whether there is a correlation between narrative potential, perceived overall quality and perceived consistency or coherence. This process, however, should be the product of another study with the aim of obtaining the right metrics and performing an adequate evaluation. The used metrics

for labeling narrative properties do not cover all computable features. There is a large number of aspects that can be extracted from an text artefact, and our narrative-based feature extraction module does not currently provide coverage for all of them. This could be a good opportunity for coming up with new, relevant metrics to story worlds and consistency or coherence.

Bibliography

- [1] Liam Bannon and Susanne Bødker. Constructing Common Information Spaces. In *Proceedings of the Fifth European Conference on Computer Supported Cooperative Work*, pages 81–96, Dordrecht, 1997. Springer Netherlands.
- [2] Paul Benjamin Lowry Aaron Curtis Michelle René Lowry and Paul Benjamin Lowry. Building a Taxonomy and Nomenclature of Collaborative Writing To Improve Interdisciplinary Research and Practice. *Journal of Business Communication*, 41(1):66–99, 2004.
- [3] Jim Bizzocchi. Games and Narrative: An Analytical Framework. *Loading*, 1(1):1–10, 2007.
- [4] Rebecca W. Black. Fanfiction Writing and the Construction of Space. *E-Learning and Digital Media*, 4(4):384–397, dec 2007.
- [5] Margaret A Boden. Computational Models of Creativity. In *Handbook of Creativity*, pages 351–373. 1999.
- [6] Margaret A Boden. *The Creative Mind: Myths and Mechanisms, Second Edition*. Routledge, 2004.
- [7] Paul Booth. Narrativity and the narrative database: Media-based wikis as interactive fan fiction. *Narrative Inquiry*, 19(2):372–392, 2009.
- [8] Susan Bryant, Andrea Forte, and Amy Bruckman. Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. *Group '05*, pages 1–10, 2005.
- [9] Peter Pin-Shan Chen and Peter Pin-Shan. The entity-relationship model—toward a unified view of data. *ACM Transactions on Database Systems*, 1(1):9–36, mar 1976.
- [10] Joaquim Colás, Alan Tapscott, Ayman Moghnieh, and Josep Blat. Shared narratives as a new interactive medium: Crosstale as a pro-

- tototype for collaborative storytelling. *Volume 6, Number 1 & 2, 2013*, 6(1 & 2), 2013.
- [11] Steven R. Corman, B. Hunter Ball, Kimberly M. Talboom, and Gene A. Brewer. Assessing two-mode semantic network story representations using a false memory paradigm. In *2013 Workshop on Computational Models of Narrative, CMN 2013*, pages 52–60, Hamburg, Germany, 2013. Schloss Dagstuhl-Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing.
- [12] Caroline Courbières and Sabine Roux. Wikia: Between Documentary Simulacra and Documented Fictions. volume 2, 2015.
- [13] Christy Dena. Transmedia Practice: Theorising the Practice of Expressing a Fictional World across Distinct Media and Environments. *Media and Communications*, PhD:373, 2009.
- [14] John Fiske. The Cultural Economy of Fandom. In *The Adoring Audience: Fan Culture and Popular Media*, pages 30–49. Routledge, 1992.
- [15] Yiannis Gabriel and N a D Connell. Co-creating stories: Collaborative experiments in storytelling. *Management Learning*, 41:507, 2010.
- [16] Jonathan Gray and Jason Mittell. Speculation on Spoilers: Lost Fandom, Narrative Consumption and Rethinking Textuality. *Particip@tions*, 4(1):1–28, 2007.
- [17] Algirdas Julius Greimas. *Sémantique structurale: recherche et méthode*. Larousse, Paris, 1966.
- [18] Sibylle Gruber. Re: Ways we contribute: Students, instructors, and pedagogies in the computer-mediated writing classroom. *Computers and Composition*, 12(1):61–78, jan 1995.
- [19] Sven Hartmann and Sebastian Link. English sentence structures and EER modeling. In *Conferences in Research and Practice in Information Technology Series*, volume 67, pages 27–35, 2007.
- [20] Hiroshi Ishii and Brygg Ullmer. Tangible bits. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '97*, pages 234–241, New York, New York, USA, 1997. ACM Press.
- [21] Henry Jenkins. *Convergence culture : where old and new*

- media collide*. New York University Press, 2006.
- [22] Henry Jenkins. *Trans-media 202: Further Reflections*, 2011. Author's official website - <http://henryjenkins.org>.
- [23] Aaron Krowne and Anil Bazaz. Authority Models for Collaborative Authoring. In *Information and Communication Technologies: From Theory to Applications*, volume 00, pages 1–7. IEEE, 2004.
- [24] David Laniado and Riccardo Tasso. Patterns of collaboration in Wikipedia. *22nd ACM Conference on Hypertext and Hypermedia HT 2011*, 2(2):201–210, 2011.
- [25] Peter Likarish and Jon Winet. Exquisite Corpse 2.0. In *Proceedings of the Designing Interactive Systems Conference on - DIS '12*, page 564, New York, New York, USA, jun 2012. ACM Press.
- [26] Jun Liu and Sudha Ram. Who Does What: Collaboration Patterns in the Wikipedia and Their Impact on Article Quality. *ACM Trans. Manage. Inf. Syst.*, 2(2):11:1–11:23, 2011.
- [27] Sarah Ann Long. Exploring the wiki world: the new face of collaboration. *New Library World*, 107(3/4):157–159, mar 2006.
- [28] Bruce Mason and Sue Thomas. A Million Penguins Research Report, 2008.
- [29] Jr Meehan. TALE-SPIN, An Interactive Program that Writes Stories. In *Proceedings of the 5th international joint conference on Artificial intelligence*, pages 91–98. Morgan Kaufmann Publishers, 1977.
- [30] Jason Mittell. Sites of Participation: Wiki Fandom and the Case of Lostpedia. *Transformative Works and Cultures*, 3(3):1–10, 2009.
- [31] Gregory B Newby. *Metric Multidimensional Information Space*. 1996.
- [32] Susanna Nordmark, Lisbeth Frolunde, and Marcelo Milrad. Promoting the creation of digital stories by teacher students: A pilot project on curriculum development in a swedish teacher education program. In *Society for Information Technology & Teacher Education International Conference*, pages 1153–1160. Association for the Advancement of Computing in Education (AACE), 2010.

- [33] Alison Pease and Simon Colton. On impact and evaluation in computational creativity: a discussion of the turing test and an alternative proposal. *Proceedings of the AISB symposium on AI and Philosophy*, page 39, 2011.
- [34] Vladimir Yakovlevich Propp. *Morphology of the folktale*,. University of Texas Press, Leningrad, 1928.
- [35] Mark O Riedl and R Michael Young. Open-World Planning for Story Generation. *Proceedings of the International Joint Conference on Artificial Intelligence*, 2005.
- [36] Graeme Ritchie. Some Empirical Criteria for Attributing Creativity to a Computer Program. *Minds and Machines*, 17(1):67–99, jul 2007.
- [37] Kimiko Ryokai and Justine Cassell. StoryMat. In *CHI '99 extended abstracts on Human factors in computing systems - CHI '99*, page 272, New York, New York, USA, may 1999. ACM Press.
- [38] Johann W Sarmiento and Gerry Stahl. Group Creativity in Interaction: Collaborative Referencing, Remembering, and Bridging Group Creativity in Interaction Sarmiento and Stahl. *journal of Human-Computer Interaction Intl.*, 24(5):1–13, 2008.
- [39] Kjeld Schmidt and Liam Bannon. Taking CSCW seriously - Supporting articulation work. *Computer Supported Cooperative Work*, 1(1-2):7–40, 1992.
- [40] Elana Shefrin. Lord of the Rings , Star Wars , and participatory fandom: mapping new congruencies between the internet and media entertainment culture. *Critical Studies in Media Communication*, 21(3):261–281, sep 2004.
- [41] Neomy Storch. Collaborative writing: Product, process, and students’ reflections. *Journal of Second Language Writing*, 14(3):153–173, sep 2005.
- [42] Bongwon Suh, Ed H. Chi, Bryan A. Pendleton, and Aniket Kittur. Us vs. Them: Understanding Social Dynamics in Wikipedia with Revert Graph Visualizations. In *2007 IEEE Symposium on Visual Analytics Science and Technology*, pages 163–170. IEEE, oct 2007.
- [43] Alan Tapscott, Joaquim Colás, Ayman Moghnieh, and Josep Blat. Writing Consistent Stories based on Structured Multi-Authored

- Narrative Spaces. In *Proceedings of the 4th Workshop on Computational Models of Narrative (CMN'13)*, volume 32, pages 277–292, 2013.
- [44] Alan Tapscott, Joaquim Colás, Ayman Moghnieh, Josep Blat, and Universitat Pompeu Fabra. Modifying Entity Relationship Models for Collaborative Fiction Planning and its Impact on Potential Authors. In *Proceedings of the 5th Workshop on Computational Models of Narrative (CMN'14)*, volume 41, pages 209–221. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, jan 2014.
- [45] Alan Tapscott, Joaquim Colás, Valeria Righi, Carlos León, and Josep Blat. We Built Our Own Worlds - Story Canonicity and Indirect Collaboration in a Shared Story World. In *Proceedings of COLLA 2016 the Sixth International Conference on Advanced Collaborative Networks, Systems and Applications*, pages 30–38, Barcelona, Spain, 2016. ThinkMind.
- [46] Alan Tapscott, Joaquim Colás, Valeria Righi, Carlos León, and Josep Blat. Large-scale collaborative story worlds: Formalizing content and author dynamics. "Unpublished manuscript", 2017.
- [47] Alan Tapscott, Javier Gómez, Carlos León, J Smailovi, M Znidaršič, and Pablo Gervás. Empirical Evidence of the Limits of Automatic Assessment of Fictional Ideation. In *Proceedings of the 5th International Workshop on Computational Creativity, Concept Invention, and General Intelligence*, pages 58—71, Bozen-Bolzano, Italy, 2016.
- [48] J. R. R. Tolkien. *The Silmarillion*. 1917.
- [49] J. R. R. Tolkien. Beowulf: The Monsters and the Critics. *Interpretations of Beowulf: A Critical Anthology*, page 31, 1936.
- [50] B. Ullmer and H. Ishii. Emerging frameworks for tangible user interfaces. *IBM Systems Journal*, 39(3.4):915–931, 2000.
- [51] Geraint A. Wiggins. Searching for computational creativity. *New Generation Computing*, 24(3):209–222, sep 2006.
- [52] Dennis M. Wilkinson and Bernardo A. Huberman. Assessing the Value of Cooperation in Wikipedia. feb 2007.