

This is a post-peer-review, pre-copyedit version of this article.
The final authenticated version is available online at:
https://doi.org/10.1007/978-3-031-22298-6_9.

Writing with (Digital) Scissors: Designing a Text Editing Tool for Assisted Storytelling using Crowd-Generated Content

Paulo Bala^{1,2}[0000-0003-1743-0261], Stuart James³[0000-0002-2649-2133], Alessio
Del Bue³[0000-0002-2262-4872], and Valentina Nisi^{1,2}[0000-0002-8051-3230]

¹ ITI-LARSyS, Portugal paulo.bala@iti.larsys.pt

² IST, Universidade de Lisboa, Portugal valentina.nisi@tecnico.ulisboa.pt

³ Visual Geometry and Modelling (VGM) Lab & Pattern Analysis and Computer
Vision (PAVIS), Istituto Italiano di Tecnologia, Italy
{stuart.james@iit,Alessio.DelBue}.it

Abstract. Digital Storytelling can exploit numerous technologies and sources of information to support the creation, refinement and enhancement of a narrative. Research on text editing tools has created novel interactions that support authors in different stages of the creative process, such as the inclusion of crowd-generated content for writing. While these interactions have the potential to change workflows, integration of these in a way that is useful and matches users' needs is unclear. In order to investigate the space of Assisted Storytelling, we designed and conducted a study to analyze how users write and edit a story about Cultural Heritage using an auxiliary source like Wikipedia. Through a diffractive analysis of stories, creative processes, and social and cultural contexts, we reflect and derive implications for design. These were applied to develop an AI-supported text editing tool using crowd-sourced content from Wikipedia and Wikidata.

Keywords: Digital Storytelling · Text Editing · Artificial Intelligence · Crowd-Sourced Content · Knowledge Graphs.

1 Introduction

As Artificial Intelligence (AI) technologies evolve and get entangled with our lives [14,18,32], decisions on how these technologies support us are made by those with the knowledge on how to develop them (e.g., engineers); most end users struggle to understand how AI supports them and have no influence on the design of such digital tools. Considering the field of Creative Support Tools (CST) [34,15,37,9], recent work have explored more human-centred approaches to the design of these tools. For example, Han et al. [22] leverage the experience of knowledge workers (such as researchers) in designing software that works across their different work activities (e.g., active reading, sensemaking, writing). Lately, HCI research has included AI in human-centred approach. Recent work [24] has shifted focus on AI to empower the creative process, leveraging the

“non human” qualities as a collaborative partner to humans. In this paper, our research question investigates how AI can support casual authors in storytelling activities.

Acknowledging the broad scope of CST [34,15,37,9] and possibilities of AI integration [24], we restrict the scope of exploration for this work. Firstly, we consider storytelling to be a skill developed over time, in which all users are proficient (but at different levels). Secondly, we consider writing as a foundational base for storytelling. Interfaces for text editing became a fruitful research topic as computers became mainstream and remain relevant in the present, as they impact millions of lives daily [21]. Finally, we constrain the context of the problem space (as done by other similar work [17,6]), to reflect on personal writing of memories and experiences of Cultural Heritage, enriched through searched or recommended content. Such restrictions create a setting where we can identify useful AI and retrieval features needed to support the task of storytelling and the creative process.

Our work is focused on exploring the design space of Assisted Storytelling according to a human-centred approach, and therefore the article is structured following research through design [44]. We set up a co-design workshop (described in section 3), where five participants (including two researchers) were asked to write a story about their Cultural Heritage and later editorialize it, using Wikipedia to add additional information and media to the story. Afterwards, participants presented the created artefact and discussed the story and activity within the group. Given the speculative nature of the activity and the critical making nature of the artefacts, we used a diffractive analysis as inspired by recent work in Human-Computer Interaction (HCI) [14,30,23,40,12]. Based on Barad’s diffraction as a metaphor for inquiry [2], the analysis (in section 4) focuses on differences; rather than striving for consensus, diffraction values the individual identity, the conflicts and the absences as a design material. Based on internal discussions, we highlight refinements to the methodology and implications for design. As such, this paper’s contribution to authoring interactive narratives relies on the insights for developing tools with Assisted Storytelling, especially tools using crowd-sourced content from Wikipedia and Wikidata. Additionally, we contribute with the design and implementation (in section 5) of an artefact system, a text editing tool with AI support, using the metaphor of digital scissors to create interactive elements of text that support writing and reading.

2 Related Work

Creative industries (such as design, film, entertainment, and others) have fostered the development of novel tools that stimulate creativity and support creative processes [34]. Palani et al. [34] have created a value framework for Creativity Support Tools (CST) based on creative practitioners’ experience, including values such as integration within the workflow, emotional connection and User Interface and Experience (UI/UX).

One such example in the CST landscape [15] are text editors tools, which have been a popular research topic since early HCI, and one of the cornerstones that made computing mainstream, as millions of users incorporated it into their work and personal lives [21]. In second wave HCI, digital tools were in particular studied for workplace use [4]; as such, these tools were designed to be efficient and generalizable to various contexts and activities [38]. Third wave HCI [4] refocuses on the user’s experience and desire, questioning how the user wants to use the tool. In addition, advancement in technology creates space for customization of interfaces [20]; therefore, tools become partners, not only being manipulable but responsive and supportive to the user’s workflow [20]. Through a media archaeology analysis of reading and writing interfaces, Emerson [13] exemplifies this with “readingwriting”, a practice of writing through the network, in which the algorithm is “constantly reading our writing and writing our reading” by influencing the user. While the text editing tools used in our daily lives may have not changed dramatically in form, recent research has re-centered interest on work with specific communities of users (to better suit their needs) and on integration of novel and complex interactions (to increase the creativity and expression of their writing). At times there is an intersection of both, with specific set of users (such as interactive narrative designers) requiring complex interactions (such as branching [19]).

On specific communities of users, Han et al. [22] queried patent workers and scientists on their practices, highlighting their need to manage multiple documents while searching, collecting, annotating, organizing, writing and reviewing. This led to the creation of *Passages*, where snippets of text (including relevant metadata such as origin or comments) can be detached from the original documents and fluidly move through the above activities. This work extended *Textlets* [21], in which text selections are treated as persistent interactive items; based on the Instrumental Interaction model [3] and inspired by interviews with legal professionals, textlets turn concepts (such as the selected text) into objects, that can be manipulated by instruments (commands) or meta-instruments (commands that act on instruments). This approach has also been extended to digital ink with *Style-Blocks+Ink* [39]. Structured note-taking (sketchnoting, self-tracking, or bullet journaling) is a (mostly hand-drawn) practice where pride in craftsmanship leads to a perceived increase in the value of the created artefacts. Romat et al. [39] support this practice digitally by treating digital ink as structured data in the form of interactive blocks, susceptible to change by instruments. Finally, Chen et al. [8] interviewed data professionals and identified a lack of connection between text and data, which can lead to writing errors. In their *CrossData* prototype, text-data connections are established automatically, and treated as objects. Using natural language as you write, these connections can be queried to retrieve data (a value for a participant), compute values (mean value) or identify mistakes (assessing the validity of a statement if the value has changed).

Text editing tools can also support the creative process by using crowd-generated content to spark or enhance writing. Using the semantic relationships

of structured data from Wikidata, Metilli et al. [31] made a semi-automatic tool to make narratives from a sequence of events. Tools can also empower users to create these semantic relationships. *Grannotate* [27] is a semi-automatic annotation tool for transcripts, that identifies entities and the relationship between them, creating knowledge graphs. These knowledge graphs can then be used to make prompts about the original text.

Advances in AI, such as language models like OpenAI’s GPT-3 [5], have given more agency for digital tools to act as co-authors. Crowd-generated content can also be aligned to language models to generate text that is not only fluent but also encodes factual knowledge that was not part of the original training data [29]. A common approach to this co-authorship with language models [6,11,43,17] is to structure the interaction as a dialogue between the author and the AI, where the AI responds to a text prompt by the user. New tools are emerging with even more complex interactions. For example, Chung et al.’s *TaleBrush* [10] is an ideation tool using line sketching of a character’s story arch and a pretrained GPT-based language model to create short stories. Zhang et al.’s *Storydrawer* [26] uses natural language processing of a child’s spoken narrative, extracts relevant entities and retrieves sketches from a dataset, creating a new drawing based on the child’s story.

While work on Human-AI collaboration can lead to novel interaction, researchers still struggle with understanding how to make it useful and desirable for users. Mina et al. [26] compiled a dataset of interactions between 63 writers with four instances of language models, varied writing tasks and varied prompts. Authors defend that analysis of large datasets such as this can help understand if language models are adding new ideas (or extending the user’s ideas) or help create better suggestions for interactive writing. Other existing work analyses these co-authoring relationships using specific writing content such as novels [6] or science writing [17].

The above works showcase the potential for text editing tools in Assisted Storytelling, especially considering reification (converting concepts into objects) from the Instrumental Interaction model [3]. Written text can be fragmented and acted upon through commands; crowd-generated content (from Wikipedia or Wikidata) can also be considered as objects (with metadata about its connections to other objects). Digital scissors are, therefore, instruments that can be used to construct new structures by acting on objects. How can users use digital scissors to write?

3 Co-Design Workshop on Assisted Storytelling

Addressing the Assisted Storytelling design space, we designed a co-design study to engage casual storytellers who have not been trained to work/design with AI or classical storytelling, such as play or screen writing, fiction or journalism. The workshop engaged participants in writing a story artifact related to Cultural Heritage (CH) and subsequently editorialise it with the support of Wikipedia. The prompt of Cultural Heritage worked well as a starting point for a storytelling

activity as it evokes personal memories and experiences for the participant. Furthermore, both tangible (e.g., monuments, buildings) and intangible elements (e.g., customs, traditions) are preserved and curated in digital archives such as Wikipedia or Wikidata [7], and therefore pliable for AI. The workshop was designed to explore how storytellers find information (crowd-generated structured information from Wikipedia or Wikidata) and incorporate that information into their storytelling practice.

3.1 Method

The process outlined for this workshop is indebted to multiple existing techniques and strategies such as cultural probes [16], traditional design studio critique [41] structured as focus groups and autoethnography [35]. Similarly, the “Magic Machines” workshop [1] uses multiple techniques to create speculative non-functioning physical artifacts about a prompted topic that are discussed with a group; the value of this method is not on the ingenuity of the artifact itself, but in the reflection of the prompted topic. In our method, while the activity of writing and editorializing will inform the design and development of an AI-driven authoring prototype, the story artifacts themselves are imbued with tensions related to the participant’s connection to Cultural Heritage (CH) and its representation in digital archives. The workshop was structured in the following stages:

1. **Recruitment & Prompt:** Some days before the session, recruited participants were asked to think of a memory or experience with a connection to CH; this could be expressed as a sketched outline of the story if they wished.
2. **Writing:** In the first 20 minutes, participants were asked to revise or compose their story draft without any external influence or support (e.g., searching online), in their digital/analogical medium of choice, highlighting the story elements that they would wish to clarify (e.g., checking the name of a location, or adding multimedia content).
3. **Editing:** For 40 minutes, participants were asked to edit their story, searching Wikipedia for elements to include. When no desired content could be found, they could describe the missing info via text or sketches.
4. **Description:** For around 5 minutes per participant, each participant presented their story to the group, projected on a public screen. Participants used this time to recount the story, reflecting on what they wrote, how they complemented it with Wikipedia, and on their experience of the workshop activity.
5. **Group Discussion:** For around 5 minutes per participant, researchers and other participants are free to ask questions to the presenting participant. Sample questions included:
 - “In the *Writing* stage, what were the main difficulties in thinking or writing?”
 - “Was there information (that could be added to the story) that you decided to not add? Why?”

- “Did you think if the information was reliable? Was there information that you decided to not add because of this?”
- “What was your strategy to adding content to your story? In searching for things to add to the story, did you start with one topic (like a town) and branch out? Did you search for specific content? Where you more interested in pictures? How much did you change the text you found?”
- “Some of you learned new things that you did not expect when browsing Wikipedia to support their story (e.g. crusades); how did this influence your experience? Was it enjoyable? Was it frustrating to diverge topic?”
- “If you started by looking at the Wikipedia information first, how would this affect the tone and structure of your story?”

3.2 Participants

Based on the goal and theme of the workshop, participants were recruited to ensure diversity of relationships to cultural heritage. The workshop includes a total of 5 people (recruited through convenience sampling), 2 females and 3 males, between 25 and 50 years of age. Some personal information (e.g., nationality) is disclosed below as it contextualizes the stories created. Four of the five participants are currently in a status of migration, living in a host country; this inclusion criteria was intentional as we wanted participants who would actively reflect about cultural heritage and identity, as migrants (living in a host country) do daily. All participants were asked to consent to the use of their story material for research purposes ⁴, including the sharing of the story material itself (see fig. 1). Participants were anonymized by substituting their names with letters and numbers (P1, P2, etc.) and any reference to identifiable personal data was erased or blurred in the reproduction of the stories.

Two of the authors of the paper participated in the workshop. They are both foreign to the local culture; a British citizen living in Italy (P1) and an Italian long-term resident in Portugal (P5). The reason for their participation is two-fold. Firstly, diffractive analysis [40] often requires for researchers to be involved through first-person methods like autoethnography, as researchers must live the data to understand the diffraction and be able to design from it. Secondly, as a co-design workshop, researchers must be included to support participants in any knowledge gap and to direct the workshop (e.g., eliciting dialogue in the group discussion). The remaining participants are students attending a Portuguese university: a local Portuguese PhD student (P2), a visiting Dutch masters student (P3), and a Chinese PhD student (P4).

3.3 Analysis

The workshop involved three researchers as facilitators. Since two of these participated in the workshop, the third researcher observed and took notes of the activity. The *Description* and *Group Discussion* phases were audio recorded, and

⁴ <https://paulobala.github.io/ICIDS2022/>

later transcribed. Researcher notes, transcriptions, and the five created story artifacts were compiled and analyzed through a diffractive stance [14,30,23,40], highlighting differences and gaps among the participants' subjects, structure, methods of storytelling, use of Wikipedia as support and social and cultural contexts. Insights result from internal discussions between all authors.

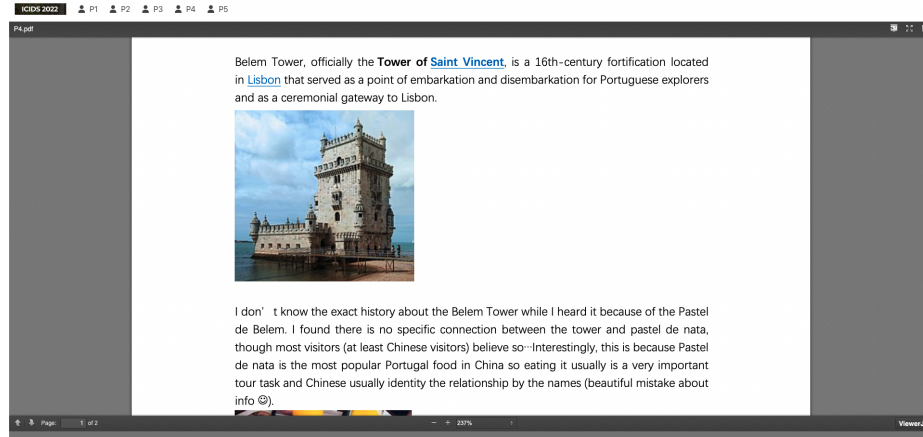


Fig. 1. Excerpt of P4's story. Stories can be consulted in the online supplementary material.

4 Diffractive Reflections on the Participants Storytelling Process

In this section, we present the researchers' reflections, based on the diffractive analysis [30,23,40,12] of the data (story artefacts, researcher notes and transcribed recordings) from the workshop.

4.1 Participant 1 (P1)

P1's story - P1, of British nationality, who has been living in Genoa (Italy, Genova in Italian) for five years, wrote about his experience and the changes the city has undergone. After mentioning the confusion between GenOva (in Italy) and GenEva (in Switzerland), P1 wrote about the Lantern, the iconic lighthouse monument of the city, which also stands as an image of Genoa. While consulting Wikipedia, to his surprise, P1 finds out that it was the second tallest lighthouse in Europe. Due to the topology of Genoa, whether by the sea or in the mountains, you can always see the Lantern. P1 mentioned you can always get lost in the contorted street of Genoa - but he could not find good pictures to exemplify this. Further details emerged from P1's story like the lack of parks

and the comparison between the architecture at the ground and higher levels. Reflecting on his time there, P1 mentions that it takes time to identify the city's hidden beauty and its historical significance.

P1's storytelling style & process In summary, this story tells the personal experience of a foreigner becoming accustomed to their new surroundings, including factual and anecdotal facts about the city's heritage. P1 used Wikipedia to find details and often learned new information about the city. P1 was satisfied with the search performance, mentioning that he found everything he was looking for, including most of the pictures, except the ones about the convoluted streets of Genoa. P1's story was anchored to tangible and officially recognized elements of heritage, which might have helped in finding content from Wikipedia (except for the more personal view of the streets of Genoa). The story was plainly formatted, as it met the participant's needs. There was no mention of specific formatting needs, or quality of visuals from external sources; this contrasts with other participants who disliked the linear formatting and were frustrated with the lack and quality of images. Unlike some participants, P1 did not compare his cultural identity to the host country; his story was firmly on the geolocalized heritage of the city.

4.2 Participant 2 (P2)

P2's story P2, born and living in Portugal, wrote a story about internal migration, recalling Alcáçovas, her family's hometown in Alentejo. Her mother and grandparents experienced a strong cultural dislocation when moving from a village to the capital city, and this is something that still permeates the younger descendants of that family. Through recounting a summer trip to Alcáçovas with her grandmother and brother, P2 writes a story of contrasts: the village culture versus the city culture and the past versus the present. On the latter, P2 described how urban renewal of transport lines affected her hometown, and complemented her story with information about old and new train stations. Throughout the story, P2 inserts information about her cultural heritage like Chocalhos (a Portuguese cattle bell that is now used in folk music, which she found out is recognized by UNESCO as an intangible cultural heritage) and the Treaty of Alcáçovas (a precursor to the Treaty of Tordesillas that split newly discovered lands in south America, between the Portuguese and the Spanish conquerors).

P2's storytelling style & process While structured as a free flow of memories, P2's story is linear in nature, using the topic of contrasts as a thread. When asked directly, P2 did not articulate any specific need in supporting the free-flowing storytelling style, but some reflections on the process emerged. P2 used Wikipedia to add facts and images but found it hard to find relevant content. In one case, she was disappointed that she couldn't find a Wikipedia entry for her hometown (regardless of the language of the entry). In another instance,

while searching for internal migration, she could find the places where people left and moved to, but could not quantify how many people it affected; she wished there was more of a human factor to the information present in Wikipedia. P2 started writing in a text editor tool, but when adding images and other facts from Wikipedia moved to a digital whiteboard as she wanted to emphasize the images. While narratively linear, the structure started to gain a waterfall aesthetic, as arrows were used to connect text and images, drawing relational meaning. P2 expressed a desire to add their images to further connect to the text.

4.3 Participant 3 (P3)

P3's story P3, of Dutch nationality and who has only been living in Portugal for three months, chose to write his story comparing the gastronomic cultures of both countries, accentuating their differences. For example, while people in the Netherlands have more economic power, they rarely eat out and view “food as nothing more than sustenance”; Portuguese, on the other hand, frequently and lengthily spend their time in restaurants, enjoying food with friends and family. The rest of the story identifies several traditional foods in Portugal, and questions if the Mediterranean diet is still maintained in Portugal. P3 complemented this line of reasoning with information from Wikipedia, and linked it back to the type of foods common in Northern European countries.

P3's storytelling style & process Similar to P2, P3's story uses contrasts or comparison as a guiding motif to his structure. P3 mentioned that he usually writes at the same time as he searches, using parallel windows. When questioned about the process of writing for this activity, P3 answered “I'm always interested in the background, of how things are the way they are. This is hard to find in Wikipedia. All the information is fairly general, lots of fact ridden, but not written like a history book, where they explain how this affects that: how there is this event and then there is that event, this kind of explanation, there is no link - so I couldn't really figure out why, why this is so different?”. This highlights his need for causal relationships between information. P3 was also frustrated about reaching dead-ends in his process of searching. In particular, the Wikipedia page for the Mediterranean diet mentioned it being a paradox (eating more red meat but by less prone to cardiovascular diseases), but he could not investigate further, leaving his questions unanswered.

4.4 Participant 4 (P4)

P4's story P4, of Chinese nationality, has lived in Portugal for six months and wrote a story exploring a cultural connection between Portugal and China. The Pastel de Nata (custard tart), also called Pastel de Belém, is famous worldwide, especially in countries with historical ties to Portugal such as China. In writing, P4 tries to understand how exactly Belém, a mandatory stop for Chinese

tourists, is connected to the pastry. Switching between Chinese and English versions of Wikipedia, he is unable to understand why, but complements the story with information about the colonial nature of the Age of Discoveries and the architectural styles of monuments such as Tower of Belém and Monument of the Discoveries.

P4's storytelling style & process P4's story is structured as a mystery story with a question that guides the story and the search process in Wikipedia, even though he cannot answer the question by the end. P4 laments the lack of transparency about the sources and the information he is retrieving (yet this is the nature of Wikipedia information as a crowd-generated source). The trustworthiness, but mostly the provenance of the information, is something only this participant has brought attention to. P4 also expressed a desire to use his images since they would better express his story and he was concerned with copyright issues (as he was unaware that Wikipedia images are public domain). Finally, P4 also expressed cognitive overload from the activity as he had to deal with several languages at a time; while he wrote in English, he searched the English and Chinese Wikipedia entries that had Portuguese names. Moreover, by going back to the Chinese Wikipedia, he feels he is defeating the intent of his story, which is to double-check and reference the information he gathered from the Chinese sources with information from other sources, finding an answer to his question.

4.5 Participant 5 (P5)

P5's story P5, of Italian nationality, who has lived in Portugal for ten years, previously sketched her story about the sea culture of Italy and Portugal. The story revolves around some self-reflections about migration, and the comparison of the two cultures, highlighting differences and similarities. The size of the two bodies of water, the Mediterranean Sea and the Atlantic Ocean, impacted how the different cultures explored the world. As such, P5 searched for and complemented the story with artefacts such as nautical maps from different cultures.

P5's storytelling style & process Since P5 had a drafted story, this activity mostly consisted of editing: adjusting words, looking for synonyms, cutting and adding some information and images. While looking up facts on Wikipedia, P5 finds very interesting new threads of information that work as a rabbit hole - where she gets curiously lost - the Portuguese discoveries and the nautical maps representing the limits of their knowledge. P5 tries and fails to complement the story with very specific images (convivial situations in Portugal and in Italy, like family gatherings at Christmas), so P5 uses maps to highlight the differences between the Mediterranean sea and the Atlantic Ocean. After spending some time looking for geographies and maps of the world, P5 manifested interest in adding a branch or an aside to the story to write about the evolution of

cartography. Overall, P5 found pleasure in getting lost in the information search, following branching of curiosities and details.

4.6 Findings

After analysing and discussing the data among the researchers, we highlight findings along two main categories: i) Reflections to the methodology used in the workshop, and ii) Insights for the design of the AI assisted authoring tool.

Reflections on Methodology

Timing Issues - While the allotted *Writing & Editing* phases were correctly timed, the *Presentation & Group Discussion* phases were rushed for participants. Further iterations should account for more time for discussion of the artefacts and the writing process, as these were useful in determining the needs and desires of participants. Furthermore, its important to consider that taking individual time to reflect on each participant, makes this method hard to scale to larger groups.

Prompt - After receiving the prompt, participants asked questions about what to write and in which style. A clear brief or prompt for the story is helpful to get people started right away, but allocating some time to discuss the prompt in the group is also needed, as we can not anticipate how clear the brief would be for all participants.

Group Discussion - When the participants presented, they tended to talk about the story rather than the process. After, during the *group discussion*, the group was not very articulated. The facilitators were needed to step in and ask several questions to clarify and expand on their process, but the timing did not allow for probing deeper, although the answers seemed to converge. Incorporating the facilitators as participants helped as they had first hand experience of the activity and its difficulties, eliciting conversation within the group.

Alternative Tools - Participant wrote the stories and edited it on their own laptops. While logistically easier to use the computer, a paper version of the study, as suggested by participants, could add some tangible and fun dimensions to it. Operating the editorializing as a paper collage with scissors and glue might invite more sharing and group reflections.

Limitations - A limitation of our work is the small number of participants and the inclusion of researchers as participants due to the exploratory nature in terms of both data and methodology. Future studies should capture a more heterogeneous set of participants' needs in regard to the use of "digital scissors". Furthermore, while the inclusion of researchers as participants is useful for contrast in our initial design exploration in "living" with the data, future iterations should focus on the inclusion of participants with different levels of writing experience and cultural connections.

Insights for Design

Access to Information - Using Wikipedia as an auxiliary source was frustrating for participants due to the lack of information and images. Participants would have liked access to more precise information, more variety and better quality of images than what they actually got. Crowd-sourced content, while large in scope, does not guarantee complete coverage or the highest quality. For digital tools leveraging on services as Wikipedia or AI, it's important to manage user's expectations. This follows existing insights on using design strategies to deal with failures and breakdowns [18], creating a space for users to adopt a more understanding stance towards technology and its imperfections. A possible avenue to explore is to promote contributions to the crowd-sourced content. Participants were keen on adding their own images, and were often knowledgeable about the information they were searching for, so AI could support a symbiotic relationship between the authoring tools and the information services.

Trust and Traceability - While only one participant highlighted concerns about transparency and trust towards the sources of information, trust is a common issue for many storytellers, from journalists, to historical, fiction and biography writers alike. Trust is also relevant regarding crowd-sourced content. Kuznetsov et al. [25] experimented with visual trust indicators for Wikipedia; while these can increase trust about the content, they can also have an opposite effect as it exposes the vetting process behind the content. When using user-generated content to complement narratives, its background information is important; it's equally important that the use of auxiliary sources is transparent to future readers. Therefore, AI supported digital tools should keep track of the information provenance as well as if and how it has been manipulated. This insight resonates with previous work on provenance [22], highlighting the need to consider text (or fragments of it) as objects, capable of having metadata to track origin and manipulation.

Embracing Connections among Cultures - Writing about CH often involves connecting, referencing, comparing and deriving meaning from more than your own heritage. It may also involve thinking about cultural identity at a supranational or subnational level (e.g., P2 searching for specific details about the family's hometown). In our study, several participants make use of comparisons to structure their story; this entails supporting comparison and connections across cultures and languages. Searching for a topic from one viewpoint, does not acknowledge the existence of "information borders". Ochigame and Ye [33] mapped Google search results based on multiple geolocations, languages and user profiles; these results showed that search algorithms use cultural assumptions, establishing cultural "filter bubbles". Crowd-sourced content is also susceptible to this, since it's likely that the users from a certain culture are the contributors and consumers. AI Assisted digital tools could foster reflection on cultural connections by exposing users to different language entries; our participants did this naturally by switching Wikipedia entries.

Story Structure and Genre - Structuring information in linear, non linear, hierarchical and rhizomatic structures is part of the authors' craft and choices on how to better support their story. The genre of a story can affect its structure, as well as the process of writing. P5's story is structured as a mystery solving quest answering a main question and therefore, exploring different (successful or unsuccessful) avenues when searching. P2's story, adopts a free-flowing reminiscence style about biographical content. P5's story follows the participant discovering new information from auxiliary sources and expressing desire to have a branching narrative. In these cases, the structure and genre of the story is deeply shaped by the process of writing. Assisted storytelling tools could support diverse genres and structures by focusing on the stages of writing. For example, Elicit ⁵ is an AI research assistant that uses GPT-3 language model and a database of 175 million papers to answer user's research question with key takeaways from abstracts; the workflow for this tool is based on the building blocks of research (e.g., search, summarization, classification).

Story Context - As participants write their stories on cultural heritage, the content is often grounded on some existing physical locations, including images such as monuments or maps to give the reader a sense of context. While previous insights have been focused on supporting the process of writing, one can also think about supporting the process of reading. Assisted Storytelling can leverage existing external services to contextualize a story, offering additional content for the reader. For example, maps can be used to represent entities mentioned in the story and to connect to other cultural heritage sites in those area.

From the diffractive analysis, discussions often delve into "what is not there". On stories about CH, the social and cultural context is needed for readers to better understand the stories. Assisted Storytelling with crowd-sourced content can create these connections for the writer or reader. Furthermore, it can create spaces to critically reflect on why information is not present or not included in a story. For example, some stories identified monuments about the Portuguese discoveries in a positive tone, while the negative effects of colonialism are not presented.

5 Design & Implementation

Starting from basic text editing needs (e.g. formatting, layout, images) and the insights gathered in the workshop, we designed and implemented a minimum viable prototype. While this prototype might not support every insight found at the moment, it can be used for further studies with participants.

The prototype was built using Vue ⁶, a javascript framework for user interfaces, and a TipTap ⁷, a rich text WYSIWYG editor and wrapper for Prose Mirror ⁸. TipTap treats content as customizable nodes, separating data from

⁵ <https://elicit.org/>

⁶ <https://vuejs.org/>

⁷ <https://tiptap.dev/>

⁸ <https://prosemirror.net/>

presentation; therefore, nodes can be acted upon whether its visually (e.g., formatting such as underlining or more complex user interface elements like pop-ups indicating provenance) or internally (e.g., transforming into another node or keeping track of changes to the data of a node). With TipTap, we can define a structure to the document, supporting saving and loading without losing internal information about the nodes. Making text as nodes allows to treat it with digital scissors, cutting parts and assembling it in a new structure.

Similar to how P3 uses parallel windows for writing, our prototype (see fig. 2 to 5) uses resizable parallel panels matching the different stages of the writing and editing process. A left panel corresponds to the main text editor, where the user can write and transform nodes. In this panel, we support base editing functions like formatting commands (e.g., bold, italic) and structural commands (e.g., headings, lists, dividers, quotes), as well as history commands (undo/redo). A right panel has multiple views that assist in storytelling:

- **Preview** - Corresponds to how readers would view the story. Nodes can have different behaviour depending on whether they are in a writing or reading panel. This allows for future exploration in how information is shown to a reader. For example, hovering a node can display the original text (see fig. 4) or show CH sites connected to a physical location.
- **Manual Search** - With multiple sources of information and various modalities (images, text), we need ways of effectively accessing this information within the prototype. In its most straightforward format, this can be a simple keyword-based retrieval (as was used in the co-design workshop). As this is the most ubiquitous approach, it makes sense to include it as the first step to finding relevant content. The co-design workshop focused on Wikipedia alone as a resource, and as participants noted, this provides only one level of granularity. While a general search across the internet can help to resolve this, the integrity of the information is uncontrolled. We, therefore, opt for a two-level approach provided by Wikidata, where Wikidata can satisfy the connectivity of information and the ability to see relevant and similar items. Found content is treated as a node and can be dragged to the main editor window for further editing, keeping track of provenance and changes to content. We expect this mode to be used more in a *writing* phase, where users actively look for content to kickstart the creative process.
- **Automatic Search** - We assist in editorializing stories in two ways:
 - *Keyword spotting*: Machine learning approaches for language, Natural Language Processing (NLP), has progressed significantly in recent years, especially for the task of Named Entity Recognition (NER) or keyword spotting [28]. Therefore, it is now possible to use State-of-the-Art approaches for other tasks, such as storytelling. For example, using a NER can access both Wikipedia and Wikidata [36], running continuously over the full text as the user writes or edits the story. Therefore, we can present related content that the user can explore and incorporate into their story by dragging it to the main editor.
 - *Context & Connectivity*: Based on the connections that are added to the story by NER or manual search while writing, it can create a virtual

subset of Wikidata. As this subset of relevant nodes increases, a contextual search is possible as highly interconnected nodes can have increased relevance. Such an approach can be similar to a simple clustering on Knowledge Graphs, where the more common connections are, the more likely they are relevant. As Wikidata also provides multiple languages, this can be a way to access alternative perspectives (if the user knows the language).

- **Visual Search** - Wikidata commonly offers one image per entity, therefore, leaning itself to visual search. In contrast, to perform an image search on Wikipedia, each document would need to be parsed, the images extracted and encoded into a descriptor [42], which would significantly increase the number of images resulting in a less accurate search. Therefore, one solution is to build a representation for each entity in Wikidata; users can then query by the image when they want a similar image. As this process can be expensive, this can also incorporate filters provided by NER and the relevant virtual subset of nodes.

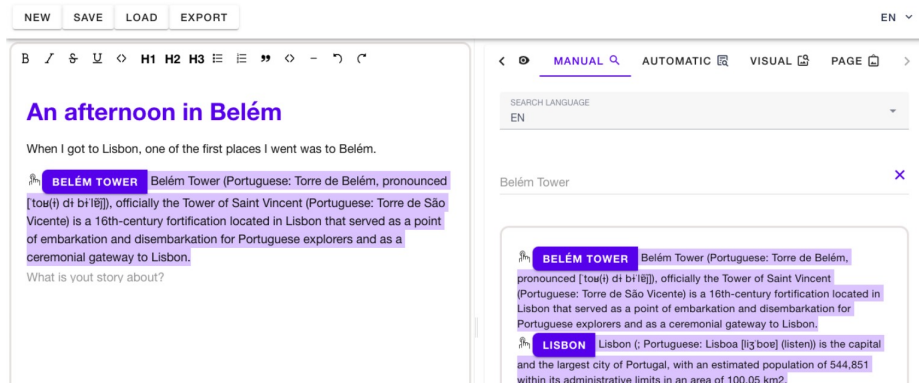


Fig. 2. Prototype tool with Assisted Storytelling – a user manually searches for a topic, and drags the node (content) to the main editor on the left.

6 Conclusion & Future Work

While advances in AI can lead to new ways to empower users, recent work suggests that we should recenter this as how do users want to be empowered by AI? In this paper, we explore how AI can be used for Assisted Storytelling by taking a human-centred approach. Restricting this scope to Cultural Heritage allows us to (1) engage participants in a personal (and sometimes controversial) topic and (2) engage AI in a topic with abundant data, but that is often too complex or nuanced to be understood. Based on a co-design workshop where

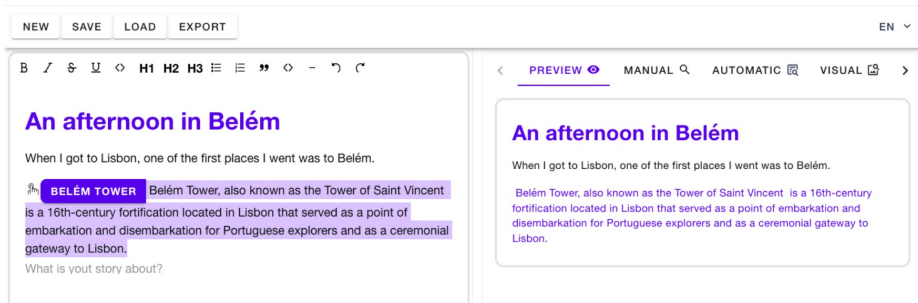


Fig. 3. Prototype tool with Assisted Storytelling – a user rewrites the found node and previews how readers would see it.

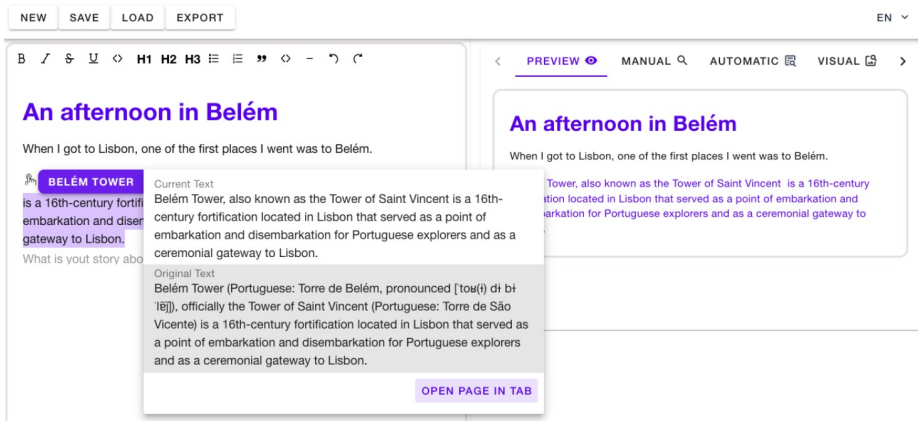


Fig. 4. Prototype tool with Assisted Storytelling – a user can track origin and manipulation by selecting nodes.

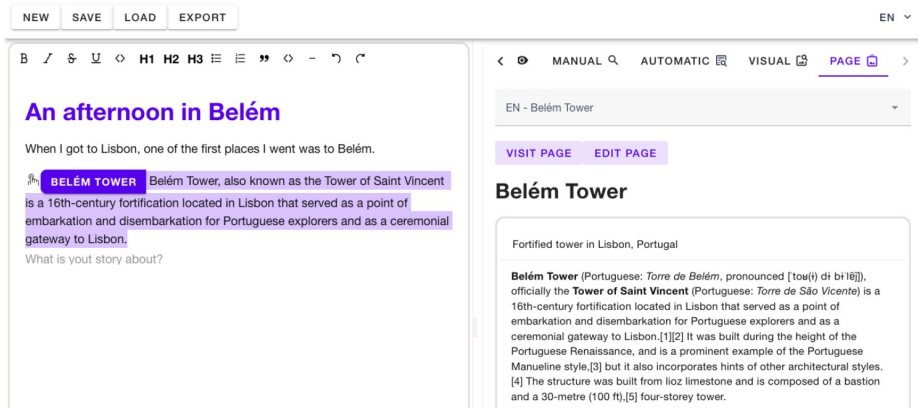


Fig. 5. Prototype tool with Assisted Storytelling – a user can consult more information on a node and compare it in different languages.

participants wrote and reused crowd-generated content, we reflect on how AI can be designed to assist writing/editing and the problems and opportunities of using crowd-generated content. We designed and implemented a digital tool leveraging the concept of digital scissors to assist in different stages of the creative process.

Future work will involve a formal evaluation with users to better understand and refine how AI can empower users in writing, as pointed out by previous work [26,17,6]. Future work will also involve further refinements to the prototype, taking advantage of crowd-generated content to identify connections (or lack of connections) in what users write and taking advantage of nodes to create more complex interactive narratives (that include casual authors in determining how they should work). Aligning AI and crowd-generated content allows recasting the creative process as “readingwriting” [13], where a user can influence and be influenced in a collaborative partnership with digital tools.

Acknowledgements. This research was supported by MEMEX (MEMories and EXperiences for inclusive digital storytelling) project funded by the European Union’s Horizon 2020 research and innovation program under grant agreement No 870743, and by LARSyS (Project UIDB/50009/2020).

References

1. Andersen, K., Wakkary, R.: The magic machine workshops: making personal design knowledge. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. pp. 1–13 (2019)
2. Barad, K.M.: Meeting the universe halfway: quantum physics and the entanglement of matter and meaning. Duke University Press, Durham (2007)
3. Beaudouin-Lafon, M.: Instrumental interaction: an interaction model for designing post-WIMP user interfaces. In: Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '00. pp. 446–453. ACM Press, The Hague, The Netherlands (2000). <https://doi.org/10.1145/332040.332473>
4. Bødker, S.: When second wave HCI meets third wave challenges. In: Proceedings of the 4th Nordic conference on Human-computer interaction changing roles - NordiCHI '06. pp. 1–8. ACM Press, Oslo, Norway (2006). <https://doi.org/10.1145/1182475.1182476>, <http://portal.acm.org/citation.cfm?doid=1182475.1182476>
5. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners (2020)
6. Calderwood, A., Qiu, V., Gero, K., Chilton, L.B.: How novelists use generative language models: An exploratory user study. In: HAI-GEN+user2agentIUI (2020)
7. Cameron, F.: The future of digital data, heritage and curation in a more-than human world. Routledge, Abingdon, Oxon ; New York, NY (2021)

8. Chen, Z., Xia, H.: CrossData: Leveraging Text-Data Connections for Authoring Data Documents. In: CHI Conference on Human Factors in Computing Systems. pp. 1–15. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491102.3517485>
9. Chung, J.J.Y., He, S., Adar, E.: The Intersection of Users, Roles, Interactions, and Technologies in Creativity Support Tools. In: Designing Interactive Systems Conference 2021. pp. 1817–1833. ACM, Virtual Event USA (Jun 2021). <https://doi.org/10.1145/3461778.3462050>, <https://dl.acm.org/doi/10.1145/3461778.3462050>
10. Chung, J.J.Y., Kim, W., Yoo, K.M., Lee, H., Adar, E., Chang, M.: TaleBrush: Visual Sketching of Story Generation with Pretrained Language Models. In: CHI Conference on Human Factors in Computing Systems Extended Abstracts. pp. 1–4. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491101.3519873>
11. Clark, E., Ross, A.S., Tan, C., Ji, Y., Smith, N.A.: Creative Writing with a Machine in the Loop: Case Studies on Slogans and Stories. In: 23rd International Conference on Intelligent User Interfaces. pp. 329–340. ACM, Tokyo Japan (Mar 2018). <https://doi.org/10.1145/3172944.3172983>
12. Devendorf, L., Andersen, K., Kelliher, A.: Making Design Memoirs: Understanding and Honoring Difficult Experiences. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. pp. 1–12. ACM, Honolulu HI USA (Apr 2020). <https://doi.org/10.1145/3313831.3376345>, <https://dl.acm.org/doi/10.1145/3313831.3376345>
13. Emerson, L.: Reading writing interfaces: from the digital to the bookbound. No. 44 in Electronic Mediations, University of Minnesota Press, Minneapolis (2014)
14. Frauenberger, C.: Entanglement HCI The Next Wave? ACM Transactions on Computer-Human Interaction **27**(1), 1–27 (Jan 2020). <https://doi.org/10.1145/3364998>, <https://dl.acm.org/doi/10.1145/3364998>
15. Frich, J., MacDonald Vermeulen, L., Remy, C., Biskjaer, M.M., Dalsgaard, P.: Mapping the Landscape of Creativity Support Tools in HCI. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. pp. 1–18. ACM, Glasgow Scotland Uk (May 2019). <https://doi.org/10.1145/3290605.3300619>
16. Gaver, B., Dunne, T., Pacenti, E.: Design: Cultural probes. Interactions **6**(1), 21–29 (Jan 1999). <https://doi.org/10.1145/291224.291235>, <https://dl.acm.org/doi/10.1145/291224.291235>
17. Gero, K.I., Liu, V., Chilton, L.: Sparks: Inspiration for Science Writing using Language Models. In: Designing Interactive Systems Conference. pp. 1002–1019. ACM, Virtual Event Australia (Jun 2022). <https://doi.org/10.1145/3532106.3533533>, <https://dl.acm.org/doi/10.1145/3532106.3533533>
18. Giaccardi, E., Speed, C., Redström, J., Ben Allouch, S., Shklovski, I., Smith, R.C.: AI and the Conditions of Design: Towards A New Set of Design Ideals. Bilbao, Spain (Jun 2022). <https://doi.org/10.21606/drs.2022.1078>, <https://dl.designresearchsociety.org/drs-conference-papers/drs2022/editorials/27>
19. Green, D., Hargood, C., Charles, F.: Contemporary Issues in Interactive Storytelling Authoring Systems. In: Rouse, R., Koenitz, H., Haahr, M. (eds.) Interactive Storytelling, vol. 11318, pp. 501–513. Springer International Publishing, Cham (2018). https://doi.org/10.1007/978-3-030-04028-4_59, http://link.springer.com/10.1007/978-3-030-04028-4_59, series Title: Lecture Notes in Computer Science
20. Grudin, J.: From tool to partner: The evolution of human-computer interaction. Synthesis Lectures on Human-Centered Interaction **10**(1), i–183 (2017)

21. Han, H.L., Renom, M.A., Mackay, W.E., Beaudouin-Lafon, M.: Textlets: Supporting Constraints and Consistency in Text Documents. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. pp. 1–13. ACM, Honolulu HI USA (Apr 2020). <https://doi.org/10.1145/3313831.3376804>
22. Han, H.L., Yu, J., Bournet, R., Ciorascu, A., Mackay, W.E., Beaudouin-Lafon, M.: Passages: Interacting with Text Across Documents. In: CHI Conference on Human Factors in Computing Systems. pp. 1–17. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491102.3502052>
23. Homewood, S., Hedemyr, M., Fagerberg Ranten, M., Kozel, S.: Tracing Conceptions of the Body in HCI: From User to More-Than-Human. In: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. pp. 1–12. ACM, Yokohama Japan (May 2021). <https://doi.org/10.1145/3411764.3445656>, <https://dl.acm.org/doi/10.1145/3411764.3445656>
24. Hwang, A.H.C.: Too Late to be Creative? AI-Empowered Tools in Creative Processes. In: CHI Conference on Human Factors in Computing Systems Extended Abstracts. pp. 1–9. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491101.3503549>
25. Kuznetsov, A., Novotny, M., Klein, J., Saez-Trumper, D., Kittur, A.: Templates and Trust-o-meters: Towards a widely deployable indicator of trust in Wikipedia. In: CHI Conference on Human Factors in Computing Systems. pp. 1–17. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491102.3517523>
26. Lee, M., Liang, P., Yang, Q.: CoAuthor: Designing a Human-AI Collaborative Writing Dataset for Exploring Language Model Capabilities. In: CHI Conference on Human Factors in Computing Systems. pp. 1–19. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491102.3502030>
27. Lee, Y., Chung, J.J.Y., Kim, T.S., Song, J.Y., Kim, J.: Promptiverse: Scalable Generation of Scaffolding Prompts Through Human-AI Hybrid Knowledge Graph Annotation. In: CHI Conference on Human Factors in Computing Systems. pp. 1–18. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491102.3502087>
28. Li, J., Sun, A., Han, J., Li, C.: A survey on deep learning for named entity recognition. *IEEE Transactions on Knowledge and Data Engineering* **34**(1), 50–70 (2022). <https://doi.org/10.1109/TKDE.2020.2981314>
29. Logan, R., Liu, N.F., Peters, M.E., Gardner, M., Singh, S.: Barack’s Wife Hillary: Using Knowledge Graphs for Fact-Aware Language Modeling. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 5962–5971. Association for Computational Linguistics, Florence, Italy (2019). <https://doi.org/10.18653/v1/P19-1598>, <https://www.aclweb.org/anthology/P19-1598>
30. Lupton, D., Watson, a.A.: Towards more-than-human digital data studies: developing research-creation methods. *Qualitative Research* p. 146879412093923 (Jul 2020). <https://doi.org/10.1177/1468794120939235>, <http://journals.sagepub.com/doi/10.1177/1468794120939235>
31. Metilli, D., Bartalesi, V., Meghini, C.: A Wikidata-based tool for building and visualising narratives. *International Journal on Digital Libraries* **20**(4), 417–432 (Dec 2019). <https://doi.org/10.1007/s00799-019-00266-3>, <http://link.springer.com/10.1007/s00799-019-00266-3>
32. Nicenboim, I., Giaccardi, E., Redström, J.: From explanations to shared understandings of AI. Bilbao, Spain (Jun 2022). <https://doi.org/10.21606/drs.2022.773>, <https://dl.designresearchsociety.org/drs-conference-papers/drs2022/researchpapers/293>

33. Ochigame, R., Ye, K.: Search Atlas: Visualizing Divergent Search Results Across Geopolitical Borders. In: Designing Interactive Systems Conference 2021. pp. 1970–1983. ACM, Virtual Event USA (Jun 2021). <https://doi.org/10.1145/3461778.3462032>, <https://dl.acm.org/doi/10.1145/3461778.3462032>
34. Palani, S., Ledo, D., Fitzmaurice, G., Anderson, F.: "I don't want to feel like I'm working in a 1960s factory": The Practitioner Perspective on Creativity Support Tool Adoption. In: CHI Conference on Human Factors in Computing Systems. pp. 1–18. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491102.3501933>
35. Rapp, A.: Autoethnography in Human-Computer Interaction: Theory and Practice. In: Filimowicz, M., Tzankova, V. (eds.) *New Directions in Third Wave Human-Computer Interaction: Volume 2 - Methodologies*, pp. 25–42. Springer International Publishing, Cham (2018). https://doi.org/10.1007/978-3-319-73374-6_3, http://link.springer.com/10.1007/978-3-319-73374-6_3, series Title: *Human-Computer Interaction Series*
36. Ravi, M.P.K., Singh, K., Mulang, I.O., Shekarpour, S., Hoffart, J., Lehmann, J.: Cholan: A modular approach for neural entity linking on wikipedia and wikidata. In: *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. pp. 504–514 (2021)
37. Remy, C., MacDonald Vermeulen, L., Frich, J., Biskjaer, M.M., Dalsgaard, P.: Evaluating Creativity Support Tools in HCI Research. In: *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. pp. 457–476. ACM, Eindhoven Netherlands (Jul 2020). <https://doi.org/10.1145/3357236.3395474>
38. Rogers, Y.: HCI Theory: Classical, Modern, and Contemporary. *Synthesis Lectures on Human-Centered Informatics* 5(2), 1–129 (May 2012). <https://doi.org/10.2200/S00418ED1V01Y201205HCI014>, <http://www.morganclaypool.com/doi/abs/10.2200/S00418ED1V01Y201205HCI014>
39. Romat, H., Marquardt, N., Hinckley, K., Henry Riche, N.: Style Blink: Exploring Digital Inking of Structured Information via Handcrafted Styling as a First-Class Object. In: CHI Conference on Human Factors in Computing Systems. pp. 1–14. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491102.3501988>
40. Sanches, P., Howell, N., Tsaknaki, V., Jenkins, T., Helms, K.: Diffraction-in-action: Designerly Explorations of Agential Realism Through Lived Data. In: CHI Conference on Human Factors in Computing Systems. pp. 1–18. ACM, New Orleans LA USA (Apr 2022). <https://doi.org/10.1145/3491102.3502029>, <https://dl.acm.org/doi/10.1145/3491102.3502029>
41. Schön, D.A.: *The reflective practitioner: how professionals think in action*. Ashgate, Farnham (2013)
42. Yang, M., He, D., Fan, M., Shi, B., Xue, X., Li, F., Ding, E., Huang, J.: Dolg: Single-stage image retrieval with deep orthogonal fusion of local and global features. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 11772–11781 (October 2021)
43. Yuan, A., Coenen, A., Reif, E., Ippolito, D.: Wordcraft: Story Writing With Large Language Models. In: *27th International Conference on Intelligent User Interfaces*. pp. 841–852. ACM, Helsinki Finland (Mar 2022). <https://doi.org/10.1145/3490099.3511105>, <https://dl.acm.org/doi/10.1145/3490099.3511105>

44. Zimmerman, J., Forlizzi, J.: Research Through Design in HCI. In: Olson, J.S., Kellogg, W.A. (eds.) *Ways of Knowing in HCI*, pp. 167–189. Springer New York, New York, NY (2014). https://doi.org/10.1007/978-1-4939-0378-8_8, http://link.springer.com/10.1007/978-1-4939-0378-8_8