

Authors' Values and Attitudes Towards AI-bridged Scalable Personalization of Creative Language Arts

Taewook Kim
taewook@u.northwestern.edu
Northwestern University
Evanston, IL, USA

Hyomin Han
hyomin@u.northwestern.edu
Northwestern University
Evanston, IL, USA

Eytan Adar
eadar@umich.edu
University of Michigan
Ann Arbor, MI, USA

Matthew Kay
mjkskay@northwestern.edu
Northwestern University
Evanston, IL, USA

John Joon Young Chung
jchung@midjourney.com
Midjourney
San Francisco, CA, USA

ABSTRACT

Generative AI has the potential to create a new form of interactive media: AI-bridged creative language arts (CLA), which bridge the author and audience by personalizing the author's vision to the audience's context and taste at scale. However, it is unclear what the authors' values and attitudes would be regarding AI-bridged CLA. To identify these values and attitudes, we conducted an interview study with 18 authors across eight genres (e.g., poetry, comics) by presenting speculative but realistic AI-bridged CLA scenarios. We identified three benefits derived from the dynamics between author, artifact, and audience: those that 1) authors get from the process, 2) audiences get from the artifact, and 3) authors get from the audience. We found how AI-bridged CLA would either promote or reduce these benefits, along with authors' concerns. We hope our investigation hints at how AI can provide intriguing experiences to CLA audiences while promoting authors' values.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**.

KEYWORDS

Authorial control, Creative language arts, Creative writing, Generative AI, Large language models, Scalable personalization

ACM Reference Format:

Taewook Kim, Hyomin Han, Eytan Adar, Matthew Kay, and John Joon Young Chung. 2024. Authors' Values and Attitudes Towards AI-bridged Scalable Personalization of Creative Language Arts. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3613904.3642529>

1 INTRODUCTION

Advancing large language model (LLM) technologies can generate text based on the user's specifications. With their generative

capabilities, LLMs have the potential to change *creative language arts* (CLA)¹. New LLM-based tools can provide authors with AI-generated suggestions based on the author's own writing or instructions [18, 58]. LLMs can also change the audience's experiences. For example, LLMs can personalize the style of existing writing, providing a different language arts experience from the original [26].

With an LLM's ability to generate texts with user-given constraints, one possibility to expand CLA experiences is to use AI models as an intermediary between authors and audiences. That is, instead of simply supporting the author's writing process with LLMs [36, 58], we explore the concept of AI-bridged CLA, where the author conveys their broad artistic vision to AI models and these models present personalized versions of the vision to the audience at scale, according to the audience's contexts and tastes. We ground the concept within the history of how technologies have transformed CLA. Personalization has been one strength of performing oral arts with a small audience, as the performer could adapt what they present to the audience's preferences and reactions [73]. However, with technologies that can distribute CLA to larger audiences (e.g., print media), these benefits may have been lost. That is, such technologies remove the role of intermediate performers and fix the content to one form. AI technologies, such as LLMs, might offer the benefits of both personalization and scalability, as they can take the author's vision and present it to the audience with personalizations they might enjoy.

With AI-bridged CLA, we strive to understand the perception of human authors about this new type of CLA media. We ask: How should AI-bridged CLA be used to facilitate what human authors value while not hindering their creative practice? What do human authors expect from AI-bridged CLA? We conducted a semi-structured online interview study with 18 authors from eight different CLA genres including poetry, novels, essays, screenplays, film scripts, pop-song lyrics, webcomics, and interactive fiction. Prior to the interview study, we designed 16 speculative scenarios (eight for lyric; eight for novel) of AI-bridged CLA as a slide deck (available in the supplementary material). We presented these scenarios to authors during the interviews to help them grasp the potential applications of AI-bridged CLA. We designed our study



This work is licensed under a Creative Commons Attribution International 4.0 License.

CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05
<https://doi.org/10.1145/3613904.3642529>

¹We use the term "creative language arts" to indicate creative, aesthetic linguistic artifacts that can focus on aspects such as narrative craft, character development, literary tropes, and traditions of poetry. Creative language arts can take various media forms, such as writing, stage performance, and audio or video recording

to reveal the potential benefits of AI-bridged CLA and the possible impacts on the dynamics between author and audience.

From the study, we found the dynamics between author, audience, and artifact and that the authors consider three types of benefits within these dynamics: 1) benefits authors get from the process (e.g., joy, therapeutic effect, attachment), 2) benefits audiences get from the artifact (e.g., empathy to the artifact, entertainment), and 3) benefits authors get from the audience (e.g., resonance from the audience, monetary reward). The authors expected AI-bridged CLA to impact these dynamics and benefits. For example, authors with high attachment to their own artifacts worried that AI might distort their intentions in the artifact. These authors often value the audience’s appreciation of the author’s unique characteristics in the artifact and worry that AI might convey distorted content to audiences. On the other hand, many other authors were okay with AI transforming aspects that are not core to their attachment, and acknowledged that AI-bridged CLA can add benefits, such as making the artifact more understandable and entertaining to the audience while increasing empathy and monetary reward authors get from the audience. Along with these findings and the authors’ concerns, we discuss potential approaches to support the creation (Section 6.1) and distribution (Section 6.2) of AI-bridged CLA and challenges & opportunities of AI-bridged CLA (Section 6.3). We hope our study attracts more researchers to investigate how to leverage generative AI for scalable personalization in other types of content, such as music and videos.

2 BACKGROUND OF CREATIVE LANGUAGE ARTS: TOWARDS AI-BRIDGED ONES

To introduce the concept of AI-bridged CLA, we start with the history of how CLA changed with different technology-powered media. In this section, we first briefly discuss those media from pre-computer eras. Then, we expand to computer-powered CLA and introduce AI-bridged CLA. We derived the history and characteristics of different media from a collection of previous work. We present the overall landscape of these media in Figure 1. While there can be many different aspects regarding these media and technologies, we mainly focus on three aspects that would be most relevant to the author-audience relationship: 1) whether they are personalizable to the individual audience (personalization), 2) the scale of the reachable audience (audience scale), and 3) if the author maintains the control over the content (authorial control). Note that our concept of an author would be either an individual (e.g., a sole novelist) or a collective team (e.g., a movie production team) that makes decisions on creating a CLA piece.

2.1 Creative Language Arts Before Computers

2.1.1 Early Oral Art. Before writing systems were invented (e.g., alphabets), performers conveyed narratives and poetry orally to others [73]. As oral means themselves (i.e., without any technologies) tend to be ephemeral, they would only reach limited audiences in the same time and space as the performer (**low audience scale**). Authorial control would be limited only with oral means, as the content can easily change from performer to performer (**low authorial control**). Moreover, the boundary between authors and performers would have been often unclear in the early oral arts. On

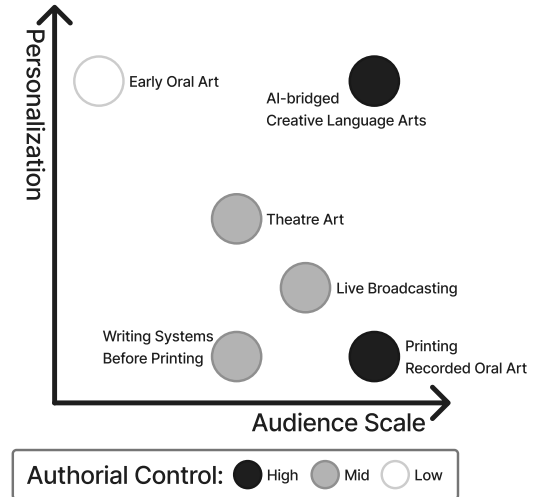


Figure 1: The landscape of creative language arts media powered by different technologies, along dimensions of 1) personalization, 2) audience scale, and 3) authorial control. The dimension values of AI-bridged creative language arts are based on recent advances in AI but assume idealistic technology within the direction of advances.

the other hand, oral forms tend to be empathetic, participatory, and personalizable [73], as performers directly face the audience and can adapt the content to what the audience would be interested in (**high personalization**).

2.1.2 Writing Systems Before Printing. With recording tools like brushes, pens, papers, and inks, writing systems, such as the alphabet, introduced visual ways of conveying CLA instead of using oral means. They removed the need for human performers to convey linguistic content. People can distribute written CLA as publications to different places and times. However, with early writing systems, the size of the reachable audience would still be limited (**mid-audience scale**). For example, with manual handwriting, creating or copying writing still requires significant effort. Writing systems could also have enabled authorial control over CLA, as the authors can record the original texts in permanent forms with credits to themselves. However, as people used the mixture of oral and written media for a long time [73], the introduction of writing systems would not immediately have introduced the concept of authorial control (**mid-authorial control**). Another characteristic of writing is that they are much less personalizable than oral means (**low personalization**). As writing is recorded, there would be no performers directly interacting with audiences. Due to that, the authors should have considered the reader as a fixed persona and the content could not be contextualized or adapted to different readers as oral art did.

2.1.3 Printing. Printing reinforced the characteristics of writing systems. As written content could be mechanically reproduced with little effort, printing greatly increased the audience scale of written CLA (**high audience scale**). With the massive distribution of printed writings, printing accelerated the transition from oral culture to writing culture. Printing also reinforced the notion of norm, credibility, correctness, and authorship, as authors can distribute a fixed content with minimal difference between copies [73] (**high authorial control**). However, as printing copies the same content, it minimizes the chance of personalizing the content to individual audiences (**low personalization**).

Note that variants of printed CLA allow a bit more extended personalization with high scalability. One type allows the reader to consume CLA non-linearly by listing multiple paths of consuming the content and allowing the audience to choose one of them. Choose-your-own-adventure style novels are examples of this type [20, 67]. With advanced distribution and publishing channels, printing different editions can be another variant. However, these still provide limited personalization compared to what oral arts could provide.

2.1.4 Theatre Art. Technology also changed CLA in oral forms. With improved and large stages, oral art forms, such as poetry recitals or plays, could be performed for larger and larger audiences. However, the audience still needs to be present at the performance, which is the limiting factor in the audience scale (**mid-audience scale**). Theater art became less personalizable than previous oral art. As the audience size grew, the performers could not react to each audience's context or taste. However, reacting to mass audiences is still possible, as the performer is present with the audience [71]. Moreover, while the final format of theater arts is in oral form, creating these often involves written scripts, which provide the high-level direction of the performance. With the fixed script and the credit of the playwright on it, the level of personalization to the performance would likely be less than when written scripts did not exist (**mid-personalization**). On the other hand, such written forms of oral CLA could have increased authorial control compared to early oral arts, as those scripts would limit transformations between performances (**mid-authorial control**). Note that authorial control of oral arts with massive audiences often tends to be distributed among multiple collaborators (e.g., directors, playwrights) and our concept of authorial control indicates the collective ones that are decided collaboratively during the planning (i.e., before the performance and personalization).

2.1.5 Live Broadcasting. Additional technological advances, such as radio, provided the oral arts access to an even larger audience scale through live broadcasts. Through live radio or television shows, live broadcast performers did not have to be in the same space as the audience (**high audience scale**). However, with limited face-to-face interactions with their audience, personalization became even more challenging [41] (**mid-personalization**). For live broadcasts, the script would still give high-level direction, but the existence of performers who can make changes would limit the authorial control unless the performers themselves are the authors (**mid-authorial control**). Again, due to the collaborative nature of oral arts media that faces massive audiences, authorial control

would likely be a collective one, and we consider it as collective decisions made before the performance and personalization.

2.1.6 Recorded Oral Art. Audio and video recording provided yet another shift within the oral arts. These served the role similar to writing systems and printings—the performance could be recorded and replicated an infinite number of times in the same form. It freed oral performance from time and space constraints, making the medium reachable to a larger audience (**high audience scale**). However, similar to writing systems and printings, the possibility of personalization is greatly reduced as the recording forces a fixed form [71] (**low personalization**). Moreover, as the author can make decisions on the final form of the artifact presented to the audience, the authors could have high authorial controls compared to live settings where the performer can bring in changes that misalign with the author's intention (**high authorial control**).

2.1.7 Summary of Creative Language Arts Before Computers. Oral traditions have seen significant evolution due to various technological advances that allowed the form to be distributed to larger audiences, beyond the constraints of time and space. Moreover, authorial control was strengthened, as technologies like printing allowed language arts to maintain their forms as the original author intended. However, as a trade-off, the personalization of early oral art gradually disappeared or weakened. As these media forms are recorded, copied, or broadcasted, performers who could previously recognize the audience's tastes and reactions and provide personalization either lost their roles (e.g., reading does not require performers) or were restricted due to the scale of the audience (e.g., the historical equivalent of streamers who cannot react to all individual audiences). Yet, personalization has the potential for interesting effects in language arts, such as adapting to the audience's reactions or fitting the content to the audience's taste. However, with the technologies mentioned in this section (before computers), personalizable CLA remains in forms that can reach only a limited set of audiences (e.g., oral arts performed for a small audience).

2.2 Creative Language Arts With Computers: Towards AI-Bridged Ones

We discuss how computers shifted CLA practices and how recent AI technologies like LLMs would make AI-bridged CLA more feasible. We first provide background on how computer technologies including AI shaped our CLA experiences in fixed written media. Then, we give a description of how computers, including AI, facilitated another type of CLA media, which can personalize their content to the audience. We extend this discussion to the concept of AI-bridged CLA: where AI models can accurately recognize the author's artistic vision with flexible personalization on how such vision is presented to the audience.

2.2.1 Using Computers for Fixed Written Media. For producing CLA in fixed, written formats, there have been many computer-based tools. They span from word processors such as Microsoft Word or Google Docs to more specialized tools, including grammar corrector [55], machine-learning-based thesaurus [34], crowd-powered writing assistants [5, 42, 47, 48, 70], or tools for specific types of writings, such as stories [11, 30], poems [35], lyrics [99],

help requests [43], journalism [64], mental support [82], and even affectionate messages [49].

Researchers and practitioners also investigated how we can leverage generative AI technologies to provide support in producing CLA in fixed forms. While generation of CLA has been explored with various technical backbones, including template-based approaches [19, 39], symbolic planning [56, 65, 81, 87, 98], case-based reasoning [38, 83, 85, 91, 95], or character simulation [15, 61], LLM technologies powered by transformer architecture [96] have brought in a large leap in flexibility and accuracy of text generation, as these models could be “prompted” to serve arbitrary natural language tasks [12, 74, 94]. With these LLM capabilities, many writing tools have been introduced, and one type of tool is those that suggest text phrases to the user’s writing, which is often called human-AI co-writing [2, 36, 53, 66, 103]. Researchers studied how these LLM suggestions can change people’s writing and found that generated texts could spark new ideas [14, 18, 36], lower grammatical errors [58], and increase vocabulary diversity [58], while introducing cognitive challenge of integrating generated texts into the user’s writing [14, 90]. Researchers also studied how specific designs of LLM, such as the number of suggestions [13] or allowing users to input instructional prompts or not [25] can impact writing. Moreover, the researchers investigated the sense of ownership of writings with AI suggestions [28]. However, AI suggestions were not the only approach to support human writing. For example, generating a summarization of the user’s writing to provide an external viewpoint can be one approach [24]. Researchers also studied socio-technical aspects of how writers interact with LLM-powered writing tools [37] and how the writer’s own value would collide with the use of AI [7]. Moreover, researchers introduced approaches to evaluate LLMs from the perspectives of interacting with human writers [59]. While this large body of literature helps us understand how LLMs can support writing fixed artifacts, they did not study the use of LLMs for personalizable media.

2.2.2 Using Computers for Personalizable Creative Language Arts. With computer technologies, the personalization of linguistic content has become more feasible. Personalization can take two forms [1]: 1) recommending different content to users [16, 32, 45], or 2) transforming the content to the user [51]. We scope our discussion of AI-bridged CLA to transforming the content to users, as LLMs would extend transformative characteristics of such media.

One early type of computer-powered personalized CLA is adaptive hypertext [51], which links the user to different content based on their contexts or direct input. Interactive narrative [69] is one specific type, that provides the reader with a set of pre-authored options during the narrative progression so that the audience can explore a set of multiple story plots. These are analogous to choose-your-own-adventure novels and have been popularized as forms of digital games, visual novels, or interactive drama. Interactive poetry, which engages user direct manipulation into the presentation of poetry, is another [4].

By linking the user to different content based on contexts and inputs, adaptive hypertext introduces the potential of crafting personalizable CLA. At the same time, as hypertext authors create different paths of content by themselves, authorial control would be largely maintained. Moreover, as digitized adaptive hypertexts

can be distributed through the internet, these contents can reach a large audience when the audience wants to. However, personalization of adaptive hypertext would still have limited flexibility. With adaptive hypertext, personalizing contents need to be pre-authored, and hence, the number of possible personalizations would be constrained to a finite set.

AI can open new opportunities for AI-bridged personalizable CLA, which conveys the author’s intentions to the audience at scale, with adequate transformations that align with the audience’s contexts and reactions [86]. AI-bridged personalizable CLA extends the flexibility in personalizing the content, beyond the user selecting one of the options/links in hypertext. One thread of research explored approaches to recognize the user context, instead of getting their direct input. Often in digital games, researchers and practitioners explored approaches to dynamically change content elements based on various aspects [77] such as user preferences [93, 100–102], and affective states [46]. However, these approaches either required a lot of technical knowledge in authoring the content or had limited flexibility in personalization. For example, with planning-based approaches, AI could select which story path to follow, but each story plot element should have been manually authored.

With extended intelligence, LLMs would allow more flexible experiences in authoring and personalization. These models can receive flexible input from audiences and authors, such as prompts [12] or preferences [76]. Moreover, with generative capabilities, LLMs do not restrict personalized outputs to a pre-authored set of contents. However, one concern with this extended flexibility is losing authorial control with the generation, similar to how those controls were absent with frequent transformations in early oral arts. However, with some approaches to align LLM generations to the authors’ intentions, such as prompting [12] or tuning with human feedback [76], the authors can increase the probability of aligning generated content with the authors’ intentions. Moreover, from the authors’ perspectives, generative LLMs can lower the required efforts in personalizing content, as they do not have to pre-author all possible content options. The scale of the audience would be high, as they can be distributed digitally at the moment when the audience wants. That is, ideally, with advanced AI technologies, AI-bridged CLA could achieve all of the high personalization, audience scale, and authorial control (Figure 1).

Researchers and practitioners explored leveraging extended AI capabilities in AI-bridged personalizable CLA. For example, AI dungeon was designed to allow users to enjoy AI-powered interactive fantasy stories with free-form text inputs [97]. Some allowed the design of artificial characters and world elements that can autonomously interact with the user [52, 78]. While previous work shows the feasibility of personalizable CLA extending the audience experiences, less is known about what the author’s expectations would be regarding these emerging types of media. Thus, we investigate what authors expect, desire, and imagine with AI-bridged CLA using speculative scenarios.

3 SCENARIO DESIGN

We interviewed CLA authors to investigate their expectations toward AI-bridged CLA (Section 4). However, they may have limited prior experience in AI-bridged CLA and related technologies. To

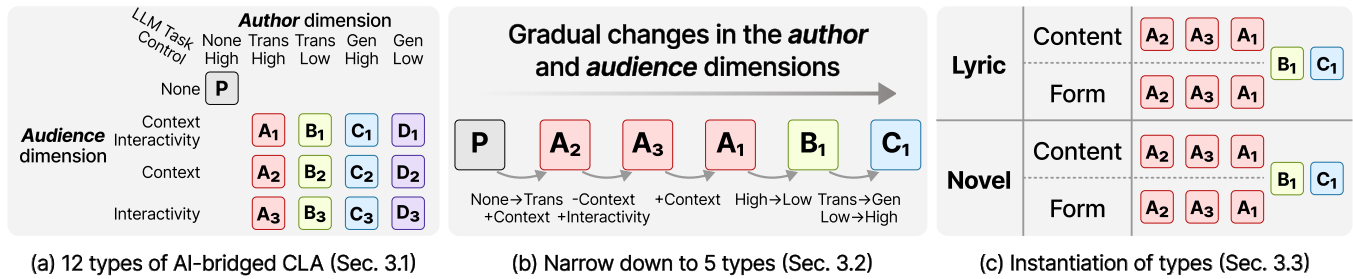


Figure 2: Overview of our scenario design process. Trans and Gen indicate transfer and generation, respectively. We first identified 12 types of AI-bridged CLA along author dimensions (LLM task, control) and audience dimensions (context, interactivity). We narrowed these into five types, gradually adding changes to the author’s current practice (*P*). Then, we instantiated five types in genres of lyrics and novels while focusing AI transformations on content, form, or both.

help them imagine tangible examples of AI-bridged CLA, we developed several representative scenarios to use as a probe [44, 72] when interviewing them. We characterized the speculative space with dimensions around two stakeholders: *author* and *audience*.

3.1 Dimensions: Author and Audience

3.1.1 Author dimensions. We first focused on two types of LLM-centered authoring tasks: *transfer* and *generation*. *Transfer* tasks are those in which LLMs change parts of the author’s original artifacts (e.g., stylistic transfer [57]). In contrast, *generation* tasks are those in which the LLMs create parts of artifacts that did not exist yet (e.g., story generation [17] or poetry generation [75, 84]). In interacting with LLMs, authors can set different levels of authorial control: *high* and *low*. With *high* control of a transfer task, for example, the author’s story would ensure that the central events of the story, its background, and the main characters remain unchanged. With *low* control, the transfer could change these elements. Generation with *high* control would ensure that the produced content strictly followed the author’s specifications. Based on the combination of *LLM tasks* and *levels of authorial control*, we can derive four combinations: Transfer/High (*A*), Transfer/Low (*B*), Generation/High (*C*), and Generation/Low (*D*).

3.1.2 Audience dimensions. We speculated that AI-bridged CLA could vary in two main aspects for audiences: *context* and *interactivity*. Context refers to the audience’s traits and situations, such as their background, preferences, personalities, device information, location, and time restrictions. Video streaming recommendations (e.g., Netflix) are an example of using context to personalize content (without transforming the content itself). Interactivity reflects the amount of possible audience interaction with AI during the scalable personalization of CLA. AI Dungeon [97] is an example of incorporating interactivity, as the audience can intervene to input text while the AI unfolds the story. We excluded scenarios that used neither audience context nor interactivity, as personalization cannot happen if we do not have at least one (i.e., there would be no input for personalization). From these, we can surface three combinations: Context/Interactivity (S_1 : Y/Y, S_2 : Y/N, and S_3 : N/Y) ($S \in \{A, B, C, D\}$). In summary, we obtained 12 combinations of scenarios (4 author dimensions \times 3 audience dimensions) (see Table 1).

Table 1: 12 possible scenario types for AI-bridged scalable personalization of creative language arts.

	Author		Audience	
	LLM task	Control	Context	Interactivity
A_1	Transfer	High	Y	Y
A_2	Transfer	High	Y	N
A_3	Transfer	High	N	Y
B_1	Transfer	Low	Y	Y
B_2	Transfer	Low	Y	N
B_3	Transfer	Low	N	Y
C_1	Generation	High	Y	Y
C_2	Generation	High	Y	N
C_3	Generation	High	N	Y
D_1	Generation	Low	Y	Y
D_2	Generation	Low	Y	N
D_3	Generation	Low	N	Y

3.2 Narrowing Down Combinations

As 12 scenarios were infeasible to cover in a single interview, we restricted ourselves to five for the sessions. To select these, we first considered how authors’ current writing practice (*P*) maps to the author and audience dimensions described in Section 3.1. In the author dimension, *P* is *without LLM* tasks but with *high authorial control*, as authors have complete control over the output. In the audience dimension, *P* does *not* reflect the *audience context*, such as audiences’ background, and has *no interactivity*. Thus, authors’ current writing practice (*P*) is [n/a, high, N, N] (Figure 2 (a)).

Having *P* as a baseline, we chose scenarios in a sequence with gradual changes in the author and audience dimensions. We picked A_2 : [transfer, high, Y, N] as the first scenario because it is distinct from *P* in one aspect (i.e., context), except for the presence of an LLM task. Then we chose A_3 : [transfer, high, N, Y] to examine distinctions in the audience dimensions. Next, we selected A_1 : [transfer, high, Y, Y], in which both audience dimensions are reflected. We selected B_1 : [transfer, low, Y, Y] as it shares all other aspects with A_1 except the authorial control. Lastly, we picked C_1 : [generation, high, Y, Y], to show generation scenarios. In summary, we picked five types of scenarios ($A_2, A_3, A_1, B_1,$ and C_1 —see Figure 2 (b)).

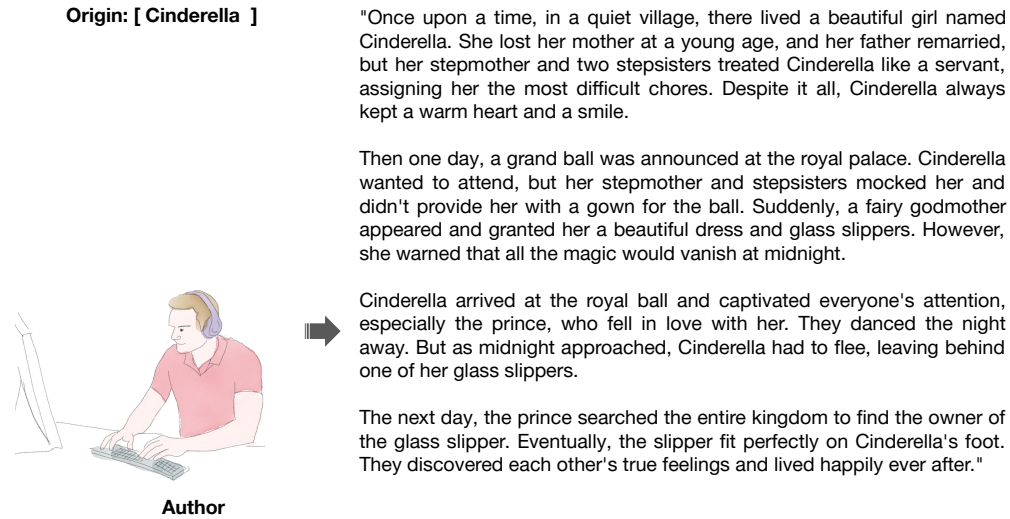


Figure 3: We used *Cinderella* story to create AI-bridged CLA examples in novels. This is a translation from Korean to English.

3.3 Instantiation of Types

3.3.1 Two genres: lyrics and novels. For these five types, we brainstormed feasible and realistic scenarios for lyrics and novels, respectively, to show each version to relevant authors. As we wanted to ensure that all authors would know the original artifact, we selected *Twinkle Twinkle Little Star* (lyrics) and *Cinderella's story* (novel) as a stand-in for an author's original content (see Figure 3).

3.3.2 Derivatives: content and form. We drew derivative scenarios based on two aspects — *content* and *form* — across all five combinations to ensure a broad coverage. In the literary critique, *content* and *form* are widely considered the core aspects of artifacts [27, 88]. Content (*what* it tells) includes properties such as theme, character, setting, plot, thought, etc., while form (*how* it tells) considers syntax, narrative, versification, diction, imagery, etc. [88]. We developed scenarios using these properties. However, it is not always clear whether certain properties affect solely content or form. For example, the change of narrative perspectives would affect the content. Hence, we differentiated the example scenarios based on the extent to which either *content* or *form* was relatively more modified.

We finally instantiated five types for two genres (*i.e.*, lyric and novel) with two aspects (*i.e.*, content and form). However, we found that B_1 : [transfer, low, Y, Y] (transfer with low control) and C_1 : [generation, high, Y, Y] (generation without the original shape) make the distinction between content and form aspects more ambiguous. Therefore, for these two types (B_1 and C_1), we instantiated scenarios without differentiation of the content and form aspects. As a result, we projected 16 speculative scenarios of AI-bridged CLA (eight for lyric; eight for novel; see Figure 2 (c)). We used GPT-4 [74] in designing these 16 scenarios. We manually edited some of the results from GPT-4 to improve the quality as a research probe. The entire scenario deck is available as supplementary material (an example in Figure 4).

4 METHODOLOGY

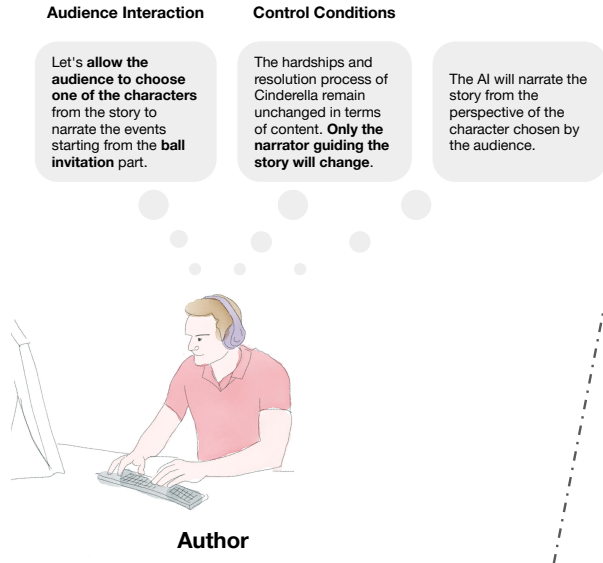
Our study aims to understand authors' views and expectations towards AI-bridged CLA. We interviewed CLA authors, presenting five speculative scenarios to guide our discussion (see supplementary material). Specifically, we looked into the authors' (a) writing practices (*how* they write), (b) a ranking of the derived benefits of work (*why* they write), (c) reactions to AI-bridged CLA, and how (a), (b), and (c) interact. To capture a broad perspective, we included authors from various CLA genres, ranging from traditional literature like poetry and novels to modern forms like (web)comics and interactive fiction, and considered authors at different career stages.

4.1 Recruitment and Participants

We recruited 18 authors (8 males and 10 females) working with eight different genres, including poetry, novels, essays, screenplays, film scripts, pop song lyrics, webcomics, and interactive fiction. To find these, we reached out to the CLA author community in South Korea, leveraging the first author's ties with them. We applied purposeful sampling to capture maximum variation [80]. First, we selectively distributed our brochures to representative members across heterogeneous genres. We then asked them to recommend other authors who might be interested in our study, following a snowball sampling approach. While interviewing participants, we simultaneously recruited more authors to constantly evaluate our theory with new samples (see Section 4.3 for details).

We tabulated details of our participants in Table 2. All participants had officially debuted as CLA writers and are currently active. Their experience ranged from 1 year to 41 years. We counted the participants' experience based on the year their first work was published. We filtered out those who have less than a year of experience in their active genre. All participants confirmed having a substantial history of publications within their respective genres. Some of them are active in more than one genre. For example, P5

Type 2: Switch Narrative Perspectives (Interaction)



"Once upon a time, in a quiet village, there lived a beautiful girl named Cinderella. She lost her mother at a young age, and although her father remarried, her stepmother and two stepsisters treated Cinderella like a maid, making her do the hardest chores in the house. Despite this, Cinderella never lost her warm heart and smile.

[From the Step-mother's Perspective]

One day, an invitation to a royal ball arrived. My two daughters were thrilled, and Cinderella also wanted to go. However, **I told her she couldn't because she had no dress. In truth, I wished for her to enjoy the ball and be happy, but in my own way, I was trying to protect her.**

Yet, **Cinderella found a way and attended the ball.** There, she caught everyone's attention, even dancing with the prince. But before midnight, she returned home, leaving behind one glass slipper. The next day, the prince traveled throughout the village with that glass slipper. **I was worried deep inside.** But when I realized the slipper belonged to Cinderella, I genuinely wished for her happiness. **The two of them went on to spend a lifetime of joy together. As a woman who tried to love and protect her family in her own way, I cheered for their happiness."**

The reader chose the character: **Stepmother** → It was reinterpreted from the stepmother's perspective.

Figure 4: A_3 -type ([transfer, high, N, Y]) derivative of *Cinderella*. In this example, we explain to authors that audiences are allowed to interactively choose and switch narrative perspectives. Note that this is a translated version from Korean to English.

Table 2: Background of participants. † Int.Fiction denotes interactive fiction. *We specified the experience with LLM; the former refers to any experience with LLM, while the latter denotes the use of LLM in their creative writing process.

#	Genre	Age	Training	Experience	Exp. w/ LLM*
P1	Poetry	57	Informal	21 years	Y / N
P2	Poetry	59	Informal	30 years	N / N
P3	Poetry	57	Informal	31 years	Y / N
P4	Poetry	48	Formal	15 years	Y / Y
P5	Novel	67	Informal	41 years	N / N
P6	Novel	45	Formal	21 years	N / N
P7	Novel	31	Formal	10 years	Y / N
P8	Novel	39	Formal	1 year	Y / N
P9	Essay	61	Informal	15 years	N / N
P10	Essay	28	Formal	3 years	Y / N
P11	Essay	58	Informal	10 years	Y / N
P12	Screenplay	54	Formal	23 years	Y / Y
P13	Film scripts	43	Formal	12 years	Y / N
P14	Song lyrics	44	Informal	15 years	Y / N
P15	Comics	29	Informal	3 years	Y / Y
P16	Comics	35	Informal	1 year	Y / Y
P17	Comics	42	Informal	17 years	Y / Y
P18	Int.Fiction†	33	Informal	3 years	Y / Y

made her first debut as a novelist, then expanded to poetry. We requested participants to specify their primary genre. During the interview, we instructed them to refer to this primary genre when answering questions.

We asked participants about their background. First, we asked them what training they received to become a CLA author (see the fourth column in Table 2). We indicate those who earned relevant higher degrees specifically in creative writing (e.g., B.A., M.A., M.F.A., or Ph.D.) as formal training. We did not label participants (e.g., P2, P3, P5) having higher degrees in literature as formal training because they claimed that their degree programs were academic (e.g., analysis of CLA), with nothing to do with creative writing practices. The most common types of informal training are self-study and group study. Furthermore, we asked participants if they had previous experience with LLMs. We specified whether their experiences with LLM were part of the artifact creation process. Even if someone had never employed LLM for their literary creation process, those who used it for inspiration or information search were classified as having used LLM in their creative process.

4.2 Interview Protocol

Using slides of the five speculative scenarios, we conducted semi-structured individual interviews in Korean via Zoom with CLA authors from eight different genres. Although we did not require video, all but one (P10) had cameras on for the interview. With the consent of the participants, we began recording the interview for data analysis using theoretical coding (see Section 4.3). The interviews took approximately an hour on average (M: 68.3 minutes, SD: 11.9 minutes). We paid each participant a virtual gift card in the amount of 70,000 KRW which is the equivalent amount of \$55 (USD) as compensation for their time.

At the beginning of the interview, we gave the participants a brief overview of the interview process. Before showing them our example scenarios designed in Section 3, we first asked questions

about (a) their writing practices (*how* they write) and (b) their prioritized value of work (*why* they write) for about 25 minutes. Then we introduced the five scenarios slides by screen sharing to see (c) their reactions to AI-bridged scenarios of scalable personalization. In some cases, participants were completely unfamiliar with LLMs. For example, P9 had never heard about LLMs and no understanding of their capabilities. Therefore, we explained each scenario one by one to ensure that they understood the concept of AI-bridged CLA. We also asked participants to imagine scenarios where their own artifacts are applied. We informed the participants that they could interrupt us at any time if they had any questions since the ideas introduced could be unfamiliar to them. For (c), we specifically asked questions like – *What do you like/dislike about this scenario? What do you wish to have (if possible) in this scenario? What kind of elements do you want to control to achieve your wish? In what contexts do you expect to have this scenario? What are the potential cases of misuse? Do you have other concerns?* This study protocol was reviewed and approved by the Institutional Review Board (IRB).

4.3 Analysis

We transcribed interviews using automatic transcribing software as soon as we finished each interview. Two authors who are bilingual in Korean and English translated them into English. We analyzed interview transcripts using theoretical coding that includes open coding, axial coding, and selective coding [68]. For the open coding process, two authors independently read the transcript, found quotes relevant to the research focus, and placed them on cards with labels (*i.e.*, low-level codes) in English. Then we collaboratively reviewed these low-level codes to resolve conflicts. We conducted this open coding once each interview was finished. We started the axial coding process when we finished about six interviews of three different genres (poetry, novels, and essays). In the axial coding process, we investigated relationships between low-level codes. Similar quotes were merged into high-level categories, and different ones were compared. These high-level categories resulted in: i) the benefits authors receive from the *process*, ii) benefits audiences get from the *artifact*, and iii) benefits authors get from the *audience*. Through analysis, we identified the dynamics of these benefits from the authors' perspective (see Figure 5). We repeated this analysis process concurrently while interviews were ongoing, to continue recruitment until diversity saturation was reached. We iteratively tested and refined our intermediate theories on the new data we collected. We regularly met to discuss, compare, and reshape codes and resolve disagreements through discussions. Finally, we derived two diagrams about (A) the dynamics of author-audience without AI, and (B) how AI-bridged CLA would influence these dynamics, as illustrated in Figure 6.

5 FINDINGS: AUTHORS' VALUES AND ATTITUDES

Figure 6 summarizes the author-audience dynamics identified by our interview analysis. We found three types of derived values: i) benefits authors receive from the *process*, ii) benefits audiences receive from the *artifact*, and iii) benefits authors receive from the *audience*. In addition, the authors showed a spectrum of reactions to AI-bridged CLA. These reactions are tied to values that the authors

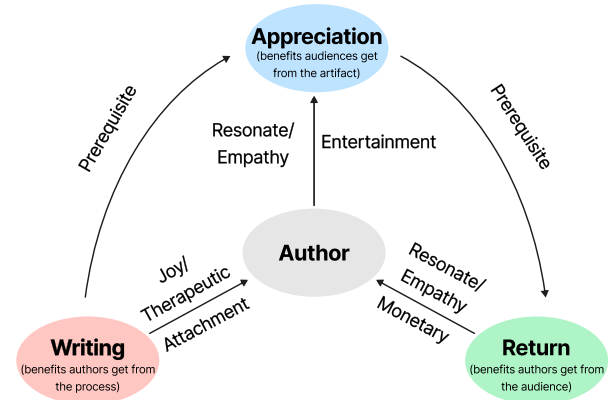


Figure 5: Benefit dynamics from the authors' perspective.

prioritize within the author-audience dynamics. We explain both the details of the author-audience dynamics without AI (Section 5.1) and with (Section 5.2). Lastly, we introduce the author's concerns on the AI-bridged CLA in Section 5.3.

5.1 Author-Audience Dynamics without AI

Our participants look for diverse values within their profession. What they prioritize varies as well. We categorized these benefits based on where they come from (see the colored legends in Figure 5). Note that the authors tended to prioritize some—but often more than one—potential benefits. They all acknowledged and emphasized that they aim to achieve multiple benefits through CLA. We highlight those instances where authors notably prioritized certain derived benefits over others.

5.1.1 Benefits The Author Gets from The Process. For some authors, the benefit derived from the artifact creation process is highly important. These authors often find **therapeutic** (*e.g.*, P8, P9) and **joyful** (*e.g.*, P4, P9) experiences during the process of creating artifacts. Although they acknowledged other positive benefits, these authors highlighted the process-derived benefits. For example,

P8: “Writing has a therapeutic value for me, which seems to be the most important. Through writing, I reflect on my life and introspect. Such self-reflection helps me heal myself, understand myself better, and in turn, understand others as well.”

Some authors emphasized the imbuing or realization of their own intrinsic artistic values in their artifacts, feeling **attachment** to those (*e.g.*, P1, P12, P17). They were particularly driven to increase these aspects (*e.g.*, uniqueness, completeness, infused intentions) within artifacts. For example,

P1: Literary value and uniqueness are the ones I most care about. Hence, if someone changes even a single letter, that work is no longer mine. I consider it a blemish on my originality.

5.1.2 Benefits The Audience Gets from The Artifact. There are authors who emphasized the benefit that audiences get from the artifact as a motivating factor. The authors infuse values into the artifact, often hoping that those will reach the audience (Figure 6 (A)).

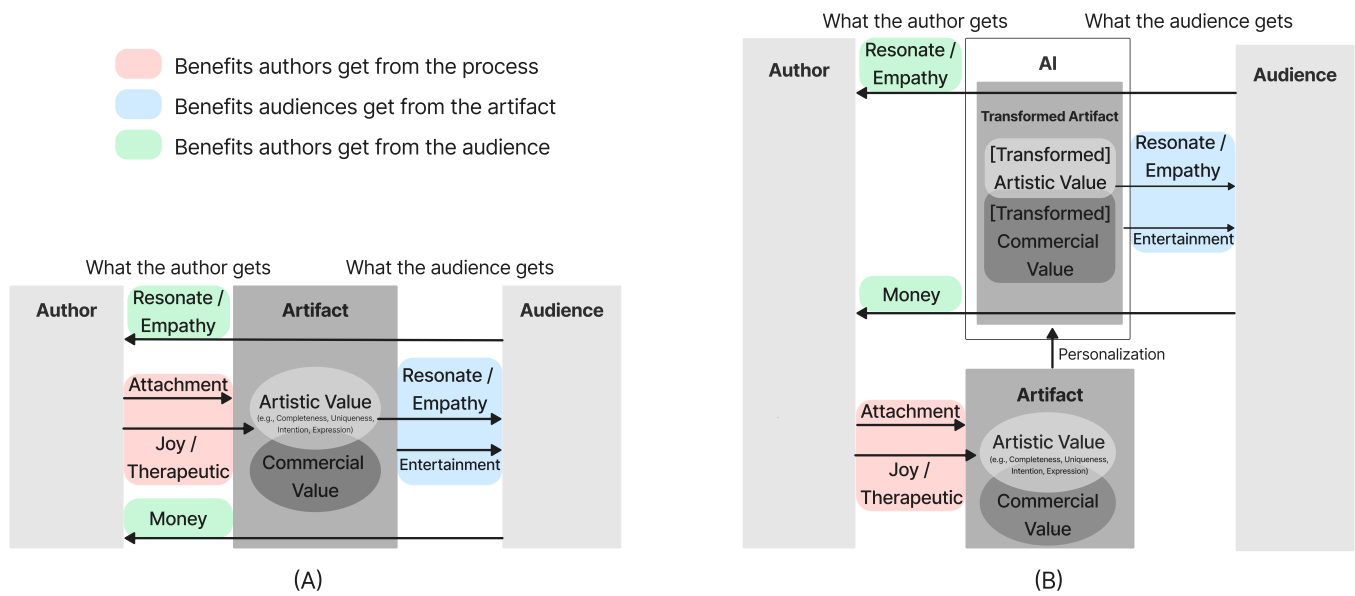


Figure 6: (A) Author-audience dynamics without AI. (B) Author-audience dynamics with AI-bridged CLA. Authors can find values from the process of creating the artifact (red). Authors infuse certain values into the artifacts in the hope that audiences would find these values from artifacts (blue). Lastly, authors might receive values from the audience (green).

The authors were split on the nature of this goal. Some are keen on creating artifacts that audiences can deeply **empathize and resonate** with (e.g., P2, P7, P10, P14), while others are more interested in creating artifacts that can provide pleasure and **entertainment** (e.g., P6, P13, P15, P16, P18). Both are similar in that they desire their artifacts to hypothetically elicit certain reactions (e.g., entertainment, empathy) from audiences. In other words, when certain properties are infused with the authors' intention, authors expect that their artifacts could evoke such reactions from audiences.

P10: "I strive to ensure that my essays don't end as self-confessional stories, but rather as writings that can resonate with the reader's heart. I want to write pieces that allow readers to have their own realizations."

P16: "I value entertainment and coherence. I eliminate content and scenes that seem uncomfortable to audiences. For example, in the case of the protagonist, even if they are in danger, I make sure it gets resolved in a short amount of time."

When making an impact on audiences through artifacts, authors wanted to correctly and effectively convey their intended vision within the artifacts to their audience. The authors described various strategies to achieve this. For example, some authors added universal objects (P4) or relatable elements (P7, P18) to provide familiarity. However, in the case that the author wants the audience to empathize, simply making every element familiar to the audience might not work as empathy presupposes a difference between the author and the audience. Hence, some authors mentioned that it was important to present familiar things in an unfamiliar way (i.e., *defamiliarization* [22]), but to the level that ensures the audience's understanding (P6, P10, P15). Others mentioned making

catchphrases (P14) and receiving peer feedback (P13, P16) as ways to help audiences understand the authors' intentions effectively.

5.1.3 Benefits The Author Gets from The Audience. Some authors placed great importance on the benefits derived from the audience. Among these authors, some aim for occasional but tangible evidence (e.g., social movement, a fan letter) of understanding and empathy from their audience (e.g., P1, P2, P8, P10, P11), while others seek consistent monetary returns (i.e., money) (e.g., P15, P18). Authors could receive such returns from the audience, when their artifacts attract a broader audience, offer genuine entertainment, or resonate with the audience. That is to say, when values within the artifact are correctly and effectively delivered, authors can expect appreciation both in sentiment and financially, from their audience. For example,

P1: "I occasionally receive emails from random fans, and I appreciate it. They said they've found themselves through my poetry. My work had been healing for them. I will keep it up to make such work."

P18: "I produce entirely commercial works. To be honest, creating work that can earn money is my top priority. My number one goal is to create content that can go viral and appeal to a broad audience."

These authors make various efforts to achieve these benefits from the audience. First, the same effort to facilitate the audience to empathize/entertain with the author's work would also apply here, as the audience will more likely return empathy and monetary reward if they appreciate the artifact. In addition, some placed emphasis on creating their own unique artifact (e.g., P1, P2) because they believe that such benefits from the audience can only be attained from work in the authors' unique tones. Others even

use social media platforms such as personal blogs and Instagram to communicate with audiences (e.g., P10, P11). They may also use these platforms to learn what is popular so as to create their own viral artifacts (e.g., P18).

5.2 Author-Audience Dynamics with AI

Our participants showed a spectrum of reactions to AI-bridged CLA. Their responses vary depending on benefits they prioritize within the author-audience dynamics. In Figure 6 (B), we illustrate how the author-audience dynamics change if AI is embedded. The most significant change is that the benefit that the audience gets from the *artifact* and the benefit that the authors get from the *audience* are mediated by the AI-transformed artifact—not the authors' completely original work. We emphasize this distinction in our analysis below. Note that the author-created “artifact” can both be the author's own piece to be transferred or the specification that is handed to AI for generation. Additionally, authors may individually have ambivalent or inconsistent attitudes toward AI-bridged CLA.

5.2.1 Prioritizing Author Benefits from The Process. As AI-bridged CLA transforms the author's artifact, they presented varied responses depending on how much of an attachment they had to the created artifact. With high attachment, the authors increased their consideration of whether the transformed outputs of AI-bridged CLA preserved the authors' original intentions. However, there were differences among authors in terms of the intended aspects that were important. For example, some authors counted content and contextual properties (e.g., backgrounds and personalities of characters in novels, materials in poetry) as a part of their intent (e.g., P15). Others (e.g., P1, P3, P17) considered not only the content but also the style and format of the writing (e.g., poetic lineation, narrative perspectives). Therefore, some did not like changes in content and contextual properties, but were unconcerned about changes in format. Those who preferred not to change both content and format considered that every factor of the artifact represents the authors' intention. For example,

P1: “The fact that the content of my poem can be consumed in any altered format is displeasing. If it has been changed in any way, it is not my work.”

P3: “I don't like these transformations. In CLA, messages/topics/themes should be aligned with the format. There's always the right format for each content.”

Interestingly, while these authors showed a strong attachment to their artifacts, some mentioned that any changes through AI-bridged CLA could be acceptable, as long as their copyright and ownership were clearly ensured and the public would be fully aware of it (e.g., P1, P17).

In other cases, there were authors with very little attachment to their artifacts, who were fine with any changes in both content and format (e.g., P2, P4). They even mentioned that they do not want to claim ownership of their artifacts. For example,

P2: “I simply transcribe poetry that I receive from a spiritual entity. I have no interest in gaining fame from it.”

P4: “It would be nice if audiences could read my poems more freely and extensively. I don't want to claim

my copyright and ownership for my work because I already found my happiness while writing my poems. I am happy with the act of writing itself.”

However, both P2 and P4 worried about intentional misuse of their work. For example, P2 expressed concerns that AI could be misused to enforce standardization of individuals, like fascism, rather than facilitate diversity and personalization for individuals.

5.2.2 Prioritizing Audience Benefits from The Artifact. Authors imagined several potential benefits and harms to the audience from artifacts of AI-bridged CLA. First, many authors noted that the audience could be more entertained through interactive features and personalized content of AI-bridged CLA (e.g., P3, P8, P12, P17).

P12: “Audience co-creating the story with AI, it can be very interesting and fun [for audiences].”

P17: “If I were a reader, I'd probably engage actively and be quite ambitious about it. From that point on, it might feel like a game to me.”

P17 also mentioned that if the audience's interaction with AI-bridged CLA could have a significant impact on the story, it would be more fun and entertaining. Otherwise, audiences might not derive value from the interactive features.

Second, authors expected that personalized stories of AI-bridged CLA might facilitate audiences' appreciation and empathy with the content of the artifact (e.g., P2, P10). They found that audiences could enjoy more immersive experiences from the artifact of AI-bridged CLA. For example,

P2: “The original purpose of poetry is to personalize it as audiences read. How to personalize it should be determined by individuals, not by the poet. If AI can help audiences achieve this, it would be wonderful.”

P10: “Switching perspectives and adapting my work pieces from each individual's own perspectives might enrich the appreciation of the essay much more effectively. It can enhance their level of immersion.”

P8 introduced another benefit of personalized stories. It can increase the accessibility of the audience's understanding. She gave an example of a story about people smoking on airplanes, which was still allowed over 15 years ago. However, a young reader, having never lived in such an era, did not think it was a reflection of past society. The reader thought it was the author's imagination. The author of the story had never even considered that some readers might not understand or relate to such a true story.

P8: “Young readers could struggle with understanding old pieces due to unfamiliar elements. In such cases, if AI-bridged CLA could personalize such unfamiliar elements into more relevant contemporary elements for the readers, that would be great. I believe it's essential for the sustainability of literature.”

Lastly, other authors expected that audiences could be exposed to more diverse perspectives from artifacts of AI-bridged CLA (e.g., P2, P5, P8, P10). They also hypothesized that individual audiences might be able to share their personalized content with each other and have extensive discussions about the original artifact.

P8: “Transformation that shows another character’s perspective is interesting as it could facilitate people to discuss what is not shown directly in the piece.”

The authors also mentioned the potential harm to the audience. Specifically, they were concerned that some audiences could be exposed to inappropriate content (e.g., age inappropriate) by AI-bridged CLA (e.g., P5, P16). To prevent such incidents, they suggested that AI-bridged CLA should be able to suppress and facilitate different aspects of the content based on the audience profile. We compiled and discussed their concerns about uncertain benefits for audiences in Section 5.3.3.

5.2.3 Prioritizing Author Benefits from The Audience. In this case, the authors expected a higher likelihood of *return* (Figure 5) from the audience. Some predicted that authors could have increased financial benefits from AI-bridged CLA. These tools could potentially improve the commercial value of artifacts (e.g., P9, P15, P18) and officialize audiences’ secondary creation (e.g., fan fiction) on an AI-bridged CLA platform (e.g., P16, P18).

P16: “Currently, derivative works created by readers tend to be shared on unofficial channels. It would be wonderful if, instead, anyone could officially produce derivatives via AI-bridged CLA platforms, and a part of the profits could be shared to the original authors.”

P18: “It feels like AI can actively assist with creating fan fiction. Even now, when you look at those who write fanfiction, they take creative liberties. I’d rather see more opportunities arising for the author to take character royalties through in-app purchases.”

Others highlighted that authors would have more opportunities to receive empathy and reactions from audiences, as a consequence of increased exposure (e.g., P1, P8, P15) and enhanced impression to audiences (e.g., P10) by AI-bridged CLA.

However, some participants were concerned that their artifact may not resonate in the way they would intend if AI-bridged CLA is applied. For example, both P1 and P12 said that the uniqueness augmented in their artifact will be diminished by AI-modification. Then their audiences, who appreciated the author’s voice, would no longer enjoy their work.

P1: My readers read my works because they want to see the uniqueness of the author. If they can no longer see my voice in the poems, they won’t want such poems. They’d have no reason to modify my poems through interactions. They probably wouldn’t want to read my poems anymore.

P1 and P12 view artifacts as a medium to express their unique individuality. Their primary goal is not mass appeal; instead, they aim for a niche audience that values and appreciates the authors’ distinct voices and styles. They look for feedback and reactions from these exclusive audiences. Therefore, for these authors, it would seem to be a disruptive machine that could even blemish their unique voice as expressed through their artifacts. In contrast, P10 has a different perspective. For her, artifacts are more than just her self-reflection. She is eager to influence her audiences profoundly, using her essays as tools to encourage changes in their lives. Therefore, for P10, AI-bridged CLA could be an assistant to

help achieve her goal by disseminating her work personalized to individuals at scale. The main difference between them would be that the former authors have a strong attachment to the artistic values they imbue in the artifact, and hope those to be accurately conveyed to the audience, while the latter prioritizes seeing the impact of the artifact on audiences.

5.3 Concerns in AI Quality, Ethics, Audiences, and Profession

Despite many authors seeing potential benefits in AI-bridged CLA, our participants also raised several concerns. We categorized these by entities within the author-audience dynamics; AI itself (Section 5.3.1), CLA artifacts (Section 5.3.2), audiences (Section 5.3.3), and authors (Section 5.3.4), respectively.

5.3.1 Unreliable Quality of AI Performance. Several participants raised concerns regarding the inconsistent quality of AI in handling artifacts. They questioned AI’s capability to accurately identify and enhance the distinct characteristics of artifacts. The authors emphasized that the artistic merit and unique attributes of their artifacts are crucial for their value. However, they worried that the unreliable performance of AI might risk altering or even potentially misrepresenting the essence of their work.

P12: “I have doubts regarding the ability of AI to produce high-quality stories. I’m concerned about whether AI could preserve an author’s unique style and generate detailed aspects within the story.”

Some authors corroborated this concern with an example. P7 said that LLM performance in direct *telling* is comparable to humans, while it shows extremely bad quality in *showing*, where P7 defined *showing* as indirectly implying contextual information through descriptions. For example, ‘*he was shocked*’ is an example of *telling*, while ‘*he dropped his mug*’ is an example of *showing*. They explained:

P7: “Audiences who have been trained in ‘showing’ through long periods of writing or reading can often detect the shortcomings in ChatGPT’s writing, feeling its lack and thirst, and they can easily identify that the text was written by ChatGPT.”

Without higher-quality LLMs, they doubt that AI-bridged CLA could be successful.

5.3.2 Misuse of AI-bridged CLA. Some authors raised a concern that AI-bridged CLA might accelerate the mass production of mundane artifacts. They noted that many authors overly reproduce similar patterns popular to the general audience. The authors were concerned that such practice might become easier and faster with AI—resulting in the audience seeing many pieces that are quite similar to each other (e.g., P3, P5, P15).

P3: “I’m afraid that AI would ‘reproduce’ already similar content that is somewhat not very high-quality. It might bring detrimental effects on the overall quality of the art. Ultimately, it would be desirable if such democratization of creative arts could happen with the assurance of the overall quality.”

The authors also worried about the potential misuse of AI-bridged CLA for distributing a uniform ideology to society. P2 viewed this as similar to how fascism leveraged mass media to spread its messages.

P2: “What I am afraid of is the standardization by AI. If one uses AI to break individuals’ uniqueness, and standardize us, then it becomes like ‘Big Brother.’”

Many authors often take great pride and responsibility in their work (e.g., P1, P14). In that sense, the deliberate misuse of AI-bridged CLA potentially entailing the mass production of low-quality work and dissemination of unintended messages could seriously tarnish the author’s pride and reputation.

5.3.3 Uncertain Benefits for Audiences. Some authors were skeptical that there would be tangible benefits of AI-bridged CLA for audiences. They wondered why audiences would want to engage in the interactive features (e.g., P1, P7, P17). P1 worried that AI-bridged CLA would hinder audiences from empathizing with the authors’ intent correctly. Personalization would also negate the benefit of defamiliarization, which enables audiences to expand their experiences and knowledge (P1, P10). P17 claimed that some audiences would prefer a passive and relaxing experience over actively engaging in interactive features when they enjoy CLA.

P17: “Audiences aged 30-40 read my comics to chill out. However, if AI continuously requires audiences to make decisions, it might be too exhausting. But teenagers, with their boundless energy, might enjoy it. I think preferences might vary depending on age.”

P7 corroborated this concern, mentioning the decrease of interactive features on a video streaming platform (Netflix²) possibly due to low audience usage.

5.3.4 AI’s Invasion of the Profession. Authors had various concerns about their job security due to AI. While they acknowledged the inevitable impact of AI on their career, common reactions were about damage to pride (e.g., P1, P8) and loss of opportunities for human authors (e.g., P5, P7, P10).

P10: “If AI reinterprets this world from various angles, our job opportunities would decrease.”

Another thread was about copyright and ownership of artifacts (e.g., P12, P14). All participants wanted to know how their copyright would be handled if this type of new media appeared in the future. Some questioned if they have to take the copyright and full responsibility of AI-bridged CLA content as well. It is at least clear that authors want to distinguish the boundaries of their ownership to protect themselves from AI-related issues.

P12: “Copyright issue also would be relevant to the issue of responsibility.”

P14: “Responsibility is important as the society would react to created content, and if the generated content is inappropriate to the society, it can be problematic.”

Some authors stated that authors should quickly accept the inevitable reality of AI invasion, and then proactively consider what they can do with AI. They shared a few ways that AI might have a positive impact, such as leveraging AI in their creative process

(P2, P7, P12, P17) and intelligent matching between authors and audiences based on shared interests (P4).

6 DISCUSSION

Authors value the connection that they establish with audiences through their artifacts. Authors receive benefits from the creation process (red in Figure 6) and imbue values into the artifact, expecting audiences to also benefit (blue in Figure 6). In return, authors hope to receive back benefits from the audience (green in Figure 6). As AI intervenes within this dynamic, authors seek to maximize these benefits while not distorting the values they imbue into their artifacts. We discuss approaches to maximize the benefits of AI-bridged CLA by supporting the creation and distribution, challenges and opportunities of AI-bridged CLA, expanding AI-bridged scalable personalization to other media, and limitations.

6.1 Supporting Creation of AI-bridged CLA through Authorial Controls

Authors often have imbued values and intentions (e.g., defamiliarization (P1, P10)) into their artifacts that they want to preserve in variations created by AI (Section 5.2.1). Hence, AI-bridged CLA creation systems should enable authors to have some controls during the authors’ creation process. First, AI-bridged CLA creation systems should enable authors to have flexible control for **constraining boundaries of variations**. For example, an author can categorically oppose any personalization that would create inappropriate versions (e.g., P5, P16 in Section 5.2.2), such as turning a fairytale for kids into a pornographic narrative. In such cases, the system must be able to enforce strict restrictions to prevent these unwanted transformations (e.g., using methods to prevent toxic degeneration of an LLM [33]).

However, some authors might struggle to envision what kind of variations would be possible. To facilitate the authors’ speculation, AI-bridged CLA creation systems could **generate estimated variations**. Furthermore, such systems could provide valuable information on how different groups of audiences could react to these variations, through which authors could preview *the benefits from the audience* (e.g., AI-powered reaction simulation [79]). By engaging in several iterations of this reflection with the AI-generated estimations and their potential impact, authors might effectively preserve or even enhance the connection within author-audience dynamics with AI-bridged CLA.

6.2 Supporting Distribution of AI-bridged CLA through Socio-technical Approaches

Authors want AI-bridged CLA to strengthen *the benefits authors get from the audience* (Section 5.2.3). They want to monitor and control unforeseen variations and leverage the reactions of audiences even after distribution. To achieve these needs, socio-technological approaches should be carefully considered in designing distribution platforms of AI-bridged CLA.

For example, distribution platforms could **record a stream of variations**. This could provide authors with a transparent view of how their artifacts are personalized and monetized at scale. Sharing records of variations on a blockchain could allow real-time monitoring and ensure transparency throughout the distribution

²<https://help.netflix.com/en/node/62526>

process [3, 31]. Furthermore, given the records of variations, the platforms may need to **prevent unauthorized deviations** beyond the original boundaries set by authors (e.g., deliberate misuse of AI-bridged CLA (Section 5.3.2)) even after the distribution. Watermarks for LLMs [50] could be a technical option for this. For example, a platform could embed watermarks to personalized variations upon distribution, to restrict any further modification after distribution. In addition, the distribution platforms could provide community-based **interaction between authors and audiences**. For example, commentary interaction on video streaming platforms could allow authors to monitor the reactions of audiences. It could enable authors to learn what kind of variation audiences enjoy and discuss within the community. Authors can then incorporate what they have learned from these reactions into their future creations [40].

However, any technological approach (i.e., blockchain, watermarks) would likely be only temporarily effective and become obsolete as adversarial technologies also evolve. Thus, it is crucial to foster a more foundational discussion on cultivating social environments, e.g., through education and policy, to develop a healthy culture around such technology.

6.3 Challenges and Opportunities of AI-bridged CLA

Authors (e.g., P5, P7) are concerned about AI taking over their profession (Section 5.3.4), which has also been clearly shown with the strike by the Writers Guild of America [23] against using AI to replace the writers' positions. This recalls the challenges traditional painters faced with the emergence of photography. As photography influenced a transformative shift in artistic expression [9] (e.g., impressionism [92], hyper-realism [10]), AI could similarly push the boundary of creative language arts. That is, authors might try to shun what AI models easily generate and create a new trend in creative language arts. However, this speculation assumes that AI would generate some types of writing more easily than others—potentially those more frequently found in the training data. As AI becomes more sophisticated, it could become 'creative' in the sense of generating content beyond its training data distribution. In such cases, the traditional role of professional authors could diminish.

We believe, however, that humans will play a crucial role in the art world even with such advanced AI technologies. This is because arts are to be consumed by humans. With perpetual human roles as consumers and highly advanced AI technologies, prosumers—those who both create and consume content [54]—might be prevalent. This shift could democratize content creation. However, this also presents challenges, such as the potential to reinforce filter bubbles [8], where prosumers may predominantly create content that aligns with their own preferences and biases. Facilitating interactions between prosumers (e.g., sharing personalization [6]) could be a solution to overcome such challenges. Moving towards prosumer media, AI-bridged CLA would be an intermediate format, as it can allow audiences to actively express ideas while still requiring the role of human authors. Prosumer media would, ultimately, merge the roles of authors and audiences in the dynamics of AI-bridged CLA, leaving two entities: prosumers and AI.

6.4 AI-bridged Scalable Personalization in Other Creative Arts

As generative AI has shown potential in diverse media formats such as audio [21], image [89], and video [29], we can further explore the opportunities for AI-bridged scalable personalization in other types of creative media, like music and video arts. What if AI provides scalable personalization for music? For example, AI could automatically transform the mood of music to provide personalized experiences for listeners depending on context [63]. Also, AI could synthesize the voices of listeners' favorite singers in playing any music to provide more personalized experiences for them [60]. Such transformations could be made at scale. However, similar to our study, we would need to first understand creators' perspectives on these technologies: how would music composers consider AI-bridged music, what kind of authorial controls do they want to have, how would their music-creating process be changed, and how would AI-bridged music influence the music composers' perceptions of their ownership [62]? Creators might have domain-specific concerns and desires for authorial control, which are understudied at the moment. Therefore, as we integrate AI into creative arts for scalable personalization, it is crucial to correctly understand the subtle differences in each medium.

6.5 Limitations

Our study has several limitations. First, we only interviewed authors. While incorporating the insights from the audiences would be informative for a richer understanding as our participants noted (Section 5.3.3), we primarily focused on exploring the authors' perception of AI-bridged CLA. Therefore, future studies would be necessary. Second, our interview findings are based on hypothetical scenarios. Our speculative scenarios are just a few representative ones, but cannot entirely cover the wide spectrum of real cases. We may not know what the real applications would look like or how the authors would consider them. Thus, it might be necessary to investigate authors' and audiences' reactions to tangible applications of AI-bridged CLA in future studies. Finally, authors' reactions to AI-bridged CLA may have discrepancies depending on cultural or temporal contexts. Some authors wanted AI-bridged CLA to correctly preserve authors' original intent (e.g., P15), while others argued that audiences should freely consume the content in their own ways (e.g., P2). We consider that the degree of such variations can depend on cultural or temporal contexts. For instance, if an author and their audiences are in a society that encourages audiences to have a free and independent interpretation of the artifact rather than figuring out the author's intention and thoughts, the spectrum of reactions might differ.

7 CONCLUSION

Rapidly advancing generative AI technologies introduce the potential for AI-bridged CLA. These new types of interactive media can bridge the author and the audience by personalizing the author's vision to the audience's context. In this work, we investigate the author's perspective on AI-bridged CLA, the benefits they consider or anticipate, and their attitudes more broadly. For this purpose, using hypothetical scenarios of AI-bridged CLA, we conducted an

interview study with 18 authors from eight genres. We found dynamics between the author, the audience, and the artifact. We also reflect on the benefits authors and audiences are likely to receive. We further identify how AI tools could add, reduce, or moderate the benefits that authors find important. Along with the authors' concerns, we discuss design implications for future AI-bridged CLA.

ACKNOWLEDGMENTS

Special thanks to Mijung Kim for her support throughout this work. We thank Duri Long and the anonymous reviewers for their valuable feedback. Also, our appreciation goes to the participants for their time and sincere responses. This work was supported by the Northwestern Advanced Cognitive Science Fellowship.

REFERENCES

- Eytan Adar, Carolyn Gearig, Ayshwarya Balasubramanian, and Jessica Hullman. 2017. PersaLog: Personalization of News Article Content. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 3188–3200. <https://doi.org/10.1145/3025453.3025631>
- Kenneth C. Arnold, Krysta Chauncey, and Krzysztof Z. Gajos. 2020. Predictive Text Encourages Predictable Writing. In *Proceedings of the 25th International Conference on Intelligent User Interfaces* (Cagliari, Italy) (IUI '20). Association for Computing Machinery, New York, NY, USA, 128–138. <https://doi.org/10.1145/3377325.3377523>
- Rita Azzi, Rima Kilany Chamoun, and Maria Sokhn. 2019. The power of a blockchain-based supply chain. *Computers & Industrial Engineering* 135 (2019), 582–592. <https://doi.org/j.cie.2019.06.042>
- Bad Quarto. 2023. Taper. <https://taper.badquarto.to/>
- Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. 2010. Soylent: A Word Processor with a Crowd Inside. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology* (New York, New York, USA) (UIST '10). Association for Computing Machinery, New York, NY, USA, 313–322. <https://doi.org/10.1145/1866029.1866078>
- Md Momen Bhuiyan, Carlos Augusto Bautista Isaza, Tanushree Mitra, and Sang Won Lee. 2022. OtherTube: Facilitating Content Discovery and Reflection by Exchanging YouTube Recommendations with Strangers. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 204, 17 pages. <https://doi.org/10.1145/3491102.3502028>
- Olof C. Biermann, Ning F. Ma, and Dongwook Yoon. 2022. From Tool to Companion: Storywriters Want AI Writers to Respect Their Personal Values and Writing Strategies. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference* (Virtual Event, Australia) (DIS '22). Association for Computing Machinery, New York, NY, USA, 1209–1227. <https://doi.org/10.1145/3532106.3533506>
- Engin Bozdag. 2013. Bias in algorithmic filtering and personalization. *Ethics and Information Technology* 15, 3 (01 Sep 2013), 209–227. <https://doi.org/10.1007/s10676-013-9321-6>
- Marta Braun. 1977. Painting and photography. *Afterimage* 5, 4 (10 1977), 14–15. <https://doi.org/10.1525/aft.1977.5.4.14>
- Horst Bredekamp. 2006. *Hyperrealism - One Step Beyond*. Tate Museum, Publishers, London, UK.
- Write Brothers. 1994. Dramatica@The Next Chapter in Story Development.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models Are Few-Shot Learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada) (NIPS'20). Curran Associates Inc., Red Hook, NY, USA, Article 159, 25 pages.
- Daniel Buschek, Martin Zörn, and Malin Eiband. 2021. The Impact of Multiple Parallel Phrase Suggestions on Email Input and Composition Behaviour of Native and Non-Native English Writers. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 732, 13 pages. <https://doi.org/10.1145/3411764.3445372>
- Alex Calderwood, Vivian Qiu, Katy Ilonka Gero, and Lydia B Chilton. 2020. How Novelists Use Generative Language Models: An Exploratory User Study. In *HAI-GEN+ user2agent@IUI*.
- Marc Cavazza, Fred Charles, and Steven J. Mead. 2001. Characters in Search of an Author: AI-Based Virtual Storytelling. In *Proceedings of the International Conference on Virtual Storytelling: Using Virtual Reality Technologies for Storytelling* (ICVS '01). Springer-Verlag, Berlin, Heidelberg, 145–154.
- P.R. Chesnais, M.J. Mucklo, and J.A. Sheena. 1995. The Fishwrap personalized news system. In *Proceedings of the Second International Workshop on Community Networking 'Integrated Multimedia Services to the Home'*. 275–282. <https://doi.org/10.1109/CN.1995.509583>
- John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. TaleBrush: Sketching Stories with Generative Pretrained Language Models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 209, 19 pages. <https://doi.org/10.1145/3491102.3501819>
- Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A. Smith. 2018. Creative Writing with a Machine in the Loop: Case Studies on Slogans and Stories. In *23rd International Conference on Intelligent User Interfaces* (Tokyo, Japan) (IUI '18). Association for Computing Machinery, New York, NY, USA, 329–340. <https://doi.org/10.1145/3172944.3172983>
- Simon Colton, Jacob Goodwin, and Tony Veale. 2012. Full-FACE Poetry Generation. In *Proceedings of the Third International Conference on Computational Creativity, ICC3 2012, Dublin, Ireland, May 30 - June 1, 2012*, Mary Lou Maher, Kristian J. Hammond, Alison Pease, Rafael Pérez y Pérez, Dan Ventura, and Geraint A. Wiggins (Eds.). computationalcreativity.net, 95–102.
- W.W. Cook and P. Collins. 2016. *Plotto: The Master Book of All Plots*. Tin House Books.
- Jade Copet, Felix Kreuk, Itai Gat, Tal Remez, David Kant, Gabriel Synnaeve, Yossi Adi, and Alexandre Défossez. 2023. Simple and Controllable Music Generation. *arXiv preprint arXiv:2306.05284* (2023).
- Lawrence Crawford. 1984. Viktor Shklovskij: Différance in Defamiliarization. *Comparative Literature* 36, 3 (1984), 209–219.
- Ashley Cullins and Katie Kilkenny. 2023. As Writers Strike, AI Could Covertly Cross the Picket Line. <https://web.archive.org/web/20230504040017/https://www.hollywoodreporter.com/business/business-news/writers-strike-ai-chatgpt-1235478681/>
- Hai Dang, Karim Benharrak, Florian Lehmann, and Daniel Buschek. 2022. Beyond Text Generation: Supporting Writers with Continuous Automatic Text Summaries. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 98, 13 pages. <https://doi.org/10.1145/3526113.3545672>
- Hai Dang, Sven Goller, Florian Lehmann, and Daniel Buschek. 2023. Choice Over Control: How Users Write with Large Language Models Using Diegetic and Non-Diegetic Prompting. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 408, 17 pages. <https://doi.org/10.1145/3544548.3580969>
- Debarati Das, David Ma, and Dongyeop Kang. 2023. Balancing the Effect of Training Dataset Distribution of Multiple Styles for Multi-Style Text Transfer. In *Findings of the Association for Computational Linguistics: ACL 2023*. Association for Computational Linguistics, Toronto, Canada, 3932–3943. <https://aclanthology.org/2023.findings-acl.243>
- Amy J. Devitt. 1993. Generalizing about Genre: New Conceptions of an Old Concept. *College Composition and Communication* 44, 4 (1993), 573–586.
- Fiona Draxler, Anna Werner, Florian Lehmann, Matthias Hoppe, Albrecht Schmidt, Daniel Buschek, and Robin Welsch. 2023. The AI Ghostwriter Effect: Users Do Not Perceive Ownership of AI-Generated Text But Self-Declare as Authors. *arXiv preprint arXiv:2303.03283* (2023).
- Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. 2023. Structure and Content-Guided Video Synthesis with Diffusion Models. *arXiv:2302.03011* [cs.CV]
- Fictional Devices. 2021. Plottr. <https://plottr.com/>
- Julie Frizzo-Barker, Peter A. Chow-White, Philippa R. Adams, Jennifer Mentanko, Dung Ha, and Sandy Green. 2020. Blockchain as a disruptive technology for business: A systematic review. *International Journal of Information Management* 51 (2020), 102029. <https://doi.org/10.1016/j.ijinfomgt.2019.10.014>
- Evgeniy Gabrilovich, Susan Dumais, and Eric Horvitz. 2004. Newsjunkie: Providing Personalized Newsfeeds via Analysis of Information Novelty. In *Proceedings of the 13th International Conference on World Wide Web* (New York, NY, USA) (WWW '04). Association for Computing Machinery, New York, NY, USA, 482–490. <https://doi.org/10.1145/988672.988738>
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. RealToxicityPrompts: Evaluating Neural Toxic Degeneration in Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 3356–3369. <https://doi.org/10.18653/v1/2020.findings-emnlp.301>

- [34] Katy Ilonka Gero and Lydia B. Chilton. 2019. How a Stylistic, Machine-Generated Thesaurus Impacts a Writer's Process. In *Proceedings of the 2019 Conference on Creativity and Cognition* (San Diego, CA, USA) (C&C '19). Association for Computing Machinery, New York, NY, USA, 597–603. <https://doi.org/10.1145/3325480.3326573>
- [35] Katy Ilonka Gero and Lydia B. Chilton. 2019. Metaphoria: An Algorithmic Companion for Metaphor Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300526>
- [36] Katy Ilonka Gero, Vivian Liu, and Lydia Chilton. 2022. Sparks: Inspiration for Science Writing Using Language Models. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference* (Virtual Event, Australia) (DIS '22). Association for Computing Machinery, New York, NY, USA, 1002–1019. <https://doi.org/10.1145/3532106.3533533>
- [37] Katy Ilonka Gero, Tao Long, and Lydia B Chilton. 2023. Social Dynamics of AI Support in Creative Writing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 245, 15 pages. <https://doi.org/10.1145/3544548.3580782>
- [38] Pablo Gervás, Belén Díaz-Agudo, Federico Peinado, and Raquel Hervás. 2005. Story Plot Generation based on CBR. In *Applications and Innovations in Intelligent Systems XII*, Ann Macintosh, Richard Ellis, and Tony Allen (Eds.). Springer London, London, 33–46.
- [39] Hugo Gonçalves Oliveira, Tiago Mendes, Ana Boavida, Ai Nakamura, and Margaret Ackerman. 2019. Co-PoeTryMe: Interactive poetry generation. *Cognitive Systems Research* 54 (2019), 199–216. <https://doi.org/10.1016/j.cogsys.2018.11.012>
- [40] Qingyu Guo, Chao Zhang, Hanfang Lyu, Zhenhui Peng, and Xiaojuan Ma. 2023. What Makes Creators Engage with Online Critiques? Understanding the Role of Artifacts' Creation Stage, Characteristics of Community Comments, and Their Interactions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 556, 17 pages. <https://doi.org/10.1145/3544548.3581054>
- [41] William A. Hamilton, Oliver Garretson, and Andruid Kerne. 2014. Streaming on Twitch: Fostering Participatory Communities of Play within Live Mixed Media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). Association for Computing Machinery, New York, NY, USA, 1315–1324. <https://doi.org/10.1145/2556288.2557048>
- [42] Chieh-Yang Huang, Shih-Hong Huang, and Ting-Hao Kenneth Huang. 2020. Heteroglossia: In-Situ Story Ideation with the Crowd. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376715>
- [43] Julie S. Hui, Darren Gergle, and Elizabeth M. Gerber. 2018. IntroAssist: A Tool to Support Writing Introductory Help Requests. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173596>
- [44] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B. Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, Nicolas Roussel, and Björn Eiderbäck. 2003. Technology Probes: Inspiring Design for and with Families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA) (CHI '03). Association for Computing Machinery, New York, NY, USA, 17–24. <https://doi.org/10.1145/642611.642616>
- [45] F.O. Isinkaye, Y.O. Folajimi, and B.A. Ojokoh. 2015. Recommendation systems: Principles, methods and evaluation. *Egyptian Informatics Journal* 16, 3 (2015), 261–273. <https://doi.org/10.1016/j.eij.2015.06.005>
- [46] Kostas Karpouzis, Georgios Yannakakis, Ana Paiva, Jeppe Herlev Nielsen, Asimina Vasalou, and Arnab Jhala. 2013. User Modelling and Adaptive, Natural Interaction for Conflict Resolution. In *2013 Humaine Association Conference on Affective Computing and Intelligent Interaction*. 719–721. <https://doi.org/10.1109/ACII.2013.131>
- [47] Joy Kim, Justin Cheng, and Michael S. Bernstein. 2014. Ensemble: Exploring Complementary Strengths of Leaders and Crowds in Creative Collaboration. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Baltimore, Maryland, USA) (CSCW '14). Association for Computing Machinery, New York, NY, USA, 745–755. <https://doi.org/10.1145/2531602.2531638>
- [48] Joy Kim, Sarah Sterman, Allegra Argent Beal Cohen, and Michael S. Bernstein. 2017. Mechanical Novel: Crowdsourcing Complex Work through Reflection and Revision. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) (CSCW '17). Association for Computing Machinery, New York, NY, USA, 233–245. <https://doi.org/10.1145/2998181.2998196>
- [49] Taewook Kim, Jung Soo Lee, Zhenhui Peng, and Xiaojuan Ma. 2019. Love in Lyrics: An Exploration of Supporting Textual Manifestation of Affection in Social Messaging. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 79 (nov 2019), 27 pages. <https://doi.org/10.1145/3359181>
- [50] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023. A Watermark for Large Language Models. In *Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 17061–17084.
- [51] Evgeny Knutov, Paul De Bra, and Mykola Pechenizkiy. 2009. AH 12 years later: a comprehensive survey of adaptive hypermedia methods and techniques. *New Review of Hypermedia and Multimedia* 15, 1 (2009), 5–38. <https://doi.org/10.1080/13614560902801608>
- [52] Max Kreminski, Melanie Dickinson, Michael Mateas, and Noah Wardrip-Fruin. 2020. Why Are We Like This?: The AI Architecture of a Co-Creative Storytelling Game. In *Proceedings of the 15th International Conference on the Foundations of Digital Games* (Bugibba, Malta) (FDG '20). Association for Computing Machinery, New York, NY, USA, Article 13, 4 pages. <https://doi.org/10.1145/3402942.3402953>
- [53] Max Kreminski, Melanie Dickinson, Noah Wardrip-Fruin, and Michael Mateas. 2022. Loose Ends: A Mixed-Initiative Creative Interface for Playful Storytelling. In *Proceedings of the Eighteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (Pomona, CA, USA) (AIIDE'22). AAAI Press, Article 15, 9 pages. <https://doi.org/10.1609/aaide.v18i1.21955>
- [54] Bodo Lang, Rebecca Dolan, Joya Kemper, and Gavin Northey. 2021. Prosumers in times of crisis: definition, archetypes and implications. *Journal of Service Management* 32, 2 (01 Jan 2021), 176–189. <https://doi.org/10.1108/JOSM-05-2020-0155>
- [55] Claudia Leacock, Michael Gamon, Joel Alejandro Mejia, and Martin Chodorow. 2022. *Automated grammatical error detection for language learners*. Springer Nature.
- [56] Michael Lebowitz. 1983. Creating a Story-Telling Universe. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence - Volume 1* (Karlsruhe, West Germany) (IJCAI'83). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 63–65.
- [57] Dongkyu Lee, Zhiliang Tian, Lanqing Xue, and Nevin L. Zhang. 2021. Enhancing Content Preservation in Text Style Transfer Using Reverse Attention and Conditional Layer Normalization. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Online, 93–102. <https://doi.org/10.18653/v1/2021.acl-long.8>
- [58] Mina Lee, Percy Liang, and Qian Yang. 2022. CoAuthor: Designing a Human-AI Collaborative Writing Dataset for Exploring Language Model Capabilities. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 388, 19 pages. <https://doi.org/10.1145/3491102.3502030>
- [59] Mina Lee, Megha Srivastava, Amelia Hardy, John Thickstun, Esin Durmus, Ashwin Paranjape, Ines Gerard-Ursin, Xiang Lisa Li, Faisal Ladhak, Frieda Rong, Rose E. Wang, Minae Kwon, Joon Sung Park, Hancheng Cao, Tony Lee, Rishi Bommasani, Michael Bernstein, and Percy Liang. 2023. Evaluating Human-Language Model Interaction. arXiv:2212.09746 [cs.CL]
- [60] Jinglin Liu, Chengxi Li, Yi Ren, Feiyang Chen, and Zhou Zhao. 2022. DiffSinger: Singing Voice Synthesis via Shallow Diffusion Mechanism. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 10 (Jun. 2022), 11020–11028. <https://doi.org/10.1609/aaai.v36i10.21350>
- [61] Sandy Louchart and Ruth Aylett. 2007. Building Synthetic Actors for Interactive Dramas. In *AAAI Fall Symposium: Intelligent Narrative Technologies*. 63–70.
- [62] Ryan Louie, Jesse Engel, and Cheng-Zhi Anna Huang. 2022. Expressive Communication: Evaluating Developments in Generative Models and Steering Interfaces for Music Creation. In *27th International Conference on Intelligent User Interfaces* (Helsinki, Finland) (IUI '22). Association for Computing Machinery, New York, NY, USA, 405–417. <https://doi.org/10.1145/3490099.3511159>
- [63] Chien-Yu Lu, Min-Xin Xue, Chia-Che Chang, Che-Rung Lee, and Li Su. 2019. Play as You Like: Timbre-Enhanced Multi-Modal Music Style Transfer. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (Jul. 2019), 1061–1068. <https://doi.org/10.1609/aaai.v33i01.33011061>
- [64] Neil Maiden, Konstantinos Zachos, Amanda Brown, George Brock, Lars Nyre, Aleksander Nygård Tonheim, Dimitris Apsotlou, and Jeremy Evans. 2018. Making the News: Digital Creativity Support for Journalists. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3173574.3174049>
- [65] James R. Meehan. 1977. TALE-SPIN, an Interactive Program That Writes Stories. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence - Volume 1* (Cambridge, USA) (IJCAI'77). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 91–98.
- [66] Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-Writing Screenplays and Theatre Scripts with Language Models: Evaluation

- by Industry Professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 355, 34 pages. <https://doi.org/10.1145/3544548.3581225>
- [67] R.A. Montgomery, S. Sundaravej, and K. Thongmoon. 2005. *Journey Under the Sea*. Choseco.
- [68] Michael J. Muller and Sandra Kogan. 2012. Grounded Theory Method in Human-Computer Interaction and Computer-Supported Cooperative Work. In *The Human Computer Interaction Handbook* (3 ed.), Julie A. Jacko (Ed.). CRC Press, Boca Raton, FL, Chapter 44, 1003–1024. <https://doi.org/10.1201/b11963-51>
- [69] Janet Horowitz Murray. 1997. *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*. The Free Press, USA.
- [70] Michael Nebeling, Alexandra To, Anhong Guo, Adrian A. de Freitas, Jaime Teevan, Steven P. Dow, and Jeffrey P. Bigham. 2016. WearWrite: Crowd-Assisted Writing from Smartwatches. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 3834–3846. <https://doi.org/10.1145/2858036.2858169>
- [71] Charles Neuringer and Ronald A Willis. 1987. The psychodynamics of theatrical spectatorship. *Journal of Dramatic Theory and Criticism* (1987), 95–110.
- [72] Renee Noortman, Britta F. Schulte, Paul Marshall, Saskia Bakker, and Anna L. Cox. 2019. HawkEye - Deploying a Design Fiction Probe. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3290605.3300652>
- [73] W.J. Ong. 2002. *Orality and Literacy: The Technologizing of the Word*. Routledge.
- [74] OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- [75] Aitor Ormazabal, Mikel Artetxe, Manex Agirrezabal, Aitor Soroa, and Eneko Agirre. 2022. PoELM: A Meter- and Rhyme-Controllable Language Model for Unsupervised Poetry Generation. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 3655–3670. <https://doi.org/10.18653/v1/2022.findings-emnlp.268>
- [76] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 27730–27744.
- [77] Panagiotis D. Paraschos and Dimitrios E. Koulouriotis. 2023. Game Difficulty Adaptation and Experience Personalization: A Literature Review. *International Journal of Human-Computer Interaction* 39, 1 (2023), 1–22. <https://doi.org/10.1080/10447318.2021.2020008>
- [78] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 2, 22 pages. <https://doi.org/10.1145/3586183.3606763>
- [79] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 74, 18 pages. <https://doi.org/10.1145/3526113.3545616>
- [80] M.Q. Patton. 2002. *Qualitative Research & Evaluation Methods*. SAGE Publications.
- [81] J. Scott Penberthy and Daniel S. Weld. 1992. UCPOP: A Sound, Complete, Partial Order Planner for ADL. In *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning* (Cambridge, MA) (KR '92). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 103–114.
- [82] Zhenhui Peng, Qingyu Guo, Ka Wing Tsang, and Xiaojuan Ma. 2020. Exploring the Effects of Technological Writing Assistance for Support Providers in Online Mental Health Community. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3313831.3376695>
- [83] Rafael Pérez Y Pérez and Mike Sharples. 2001. MEXICA: A computer model of a cognitive account of creative writing. *Journal of Experimental & Theoretical Artificial Intelligence* 13, 2 (2001), 119–139. <https://doi.org/10.1080/09528130010029820>
- [84] Andrei Popescu-Belis, Alex Atrio, Valentin Minder, Aris Xanthos, Gabriel Luthier, Simon Mattei, and Antonio Rodriguez. 2022. Constrained Language Models for Interactive Poem Generation. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France, 3519–3529.
- [85] Mark Riedl. 2009. Vignette-Based Story Planning: Creativity Through Exploration and Retrieval. *Proceedings of the International Joint Workshop on Computational Creativity 2008* (01 2009).
- [86] Mark O. Riedl. 2010. Scalable Personalization of Interactive Experiences through Creative Automation. *Comput. Entertain.* 8, 4, Article 26 (dec 2010), 3 pages. <https://doi.org/10.1145/1921141.1921146>
- [87] Mark O. Riedl and R. Michael Young. 2010. Narrative Planning: Balancing Plot and Character. *J. Artif. Int. Res.* 39, 1 (sep 2010), 217–268.
- [88] Duncan Robertson. 1967. The Dichotomy of Form and Content. *College English* 28, 4 (1967), 273–279.
- [89] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. arXiv:2112.10752 [cs.CV]
- [90] Nikhil Singh, Guillermo Bernal, Daria Savchenko, and Elena L. Glassman. 2023. Where to Hide a Stolen Elephant: Leaps in Creative Writing with Multimodal Machine Intelligence. *ACM Trans. Comput.-Hum. Interact.* 30, 5, Article 68 (sep 2023), 57 pages. <https://doi.org/10.1145/3511599>
- [91] Reid Swanson and Andrew S. Gordon. 2012. Say Anything: Using Textual Case-Based Reasoning to Enable Open-Domain Interactive Storytelling. *ACM Trans. Interact. Intell. Syst.* 2, 3, Article 16 (sep 2012), 35 pages. <https://doi.org/10.1145/2362394.2362398>
- [92] David Sweet. 2021. Before and after photography. *Journal of Contemporary Painting* 7, 1-2 (2021), 163–175. https://doi.org/10.1386/jcp_00029_1
- [93] David Thue, Vadim Bulitko, Marcia Spetch, and Eric Wasylishen. 2007. Interactive Storytelling: A Player Modelling Approach. In *Proceedings of the Third AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (Stanford, California) (AIIDE'07). AAAI Press, 43–48.
- [94] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiannan Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288 [cs.CL]
- [95] S.R. Turner. 1992. *MINSTREL, a Computer Model of Creativity and Storytelling*. Number pt. 1 in CSD (Series). UCLA Computer Science Department.
- [96] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc.
- [97] Nick Walton. 2019. AI Dungeon 2. <https://aidungeon.cc/>
- [98] Stephen G. Ware, R. Michael Young, Brent Harrison, and David L. Roberts. 2014. A Computational Model of Plan-Based Narrative Conflict at the Fabula Level. *IEEE Transactions on Computational Intelligence and AI in Games* 6, 3 (2014), 271–288. <https://doi.org/10.1109/TCIAIG.2013.2277051>
- [99] Kento Watanabe, Yuichiroh Matsubayashi, Kentaro Inui, Tomoyasu Nakano, Satoru Fukayama, and Masataka Goto. 2017. LyriSys: An Interactive Support System for Writing Lyrics Based on Topic Transition. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces* (Limassol, Cyprus) (IUI '17). Association for Computing Machinery, New York, NY, USA, 559–563. <https://doi.org/10.1145/3025171.3025194>
- [100] Hong Yu and Mark O. Riedl. 2012. A Sequential Recommendation Approach for Interactive Personalized Story Generation. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1* (Valencia, Spain) (AAMAS '12). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 71–78.
- [101] Hong Yu and Mark O. Riedl. 2013. Data-Driven Personalized Drama Management. In *Proceedings of the Ninth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (Boston, MA, USA) (AIIDE'13). AAAI Press, 191–197.
- [102] Hong Yu and Mark O. Riedl. 2014. Personalized Interactive Narratives via Sequential Recommendation of Plot Points. *IEEE Transactions on Computational Intelligence and AI in Games* 6, 2 (2014), 174–187. <https://doi.org/10.1109/TCIAIG.2013.2282771>
- [103] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: Story Writing With Large Language Models. In *17th International Conference on Intelligent User Interfaces* (Helsinki, Finland) (IUI '22). Association for Computing Machinery, New York, NY, USA, 841–852. <https://doi.org/10.1145/3490099.3511105>