Supporting Material Writing Practice with Phraselette, a Palette of Phrases

Alex Calderwood
Computational Media
University of California
Santa Cruz, USA
acalde24@ucsc.edu

Taewook Kim
Computer Science and Communication
Northwestern University
Evanston, Illinois, USA
taewook@u.northwestern.edu

Yuqian Sun
Storytelling Lab
Midjourney
London, United Kingdom
ysun@midjourney.com

Melissa Roemmele
Midjourney
San Francisco, California, USA
mroemmele@midjourney.com

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

Max Kreminski
Midjourney
San Francisco, California, USA
mkreminski@midjourney.com

Abstract

We present a demonstration of *Phraselette*, an artistic support tool designed for compatibility with the *writerly values* of experimental poets. Following a theory of "material writing support", we introduce affordances for selecting short spans of text (on the order of a few words) to vary; constraining text generation procedures (some based on language models) to adhere to poetic intent; and searching through large spaces of potential variations for phrases that satisfy users' constraints in unexpected but evocative ways. *Phraselette* has been validated through an extended expert evaluation with 10 published poets; we found that, in contrast to the dominant prompting-based approach to interacting with language models as writing support tools, *Phraselette* is better aligned with experimental poetry practice, providing deeper support for navigating spaces of potential interpretations of poetic text.

CCS Concepts

• **Applied computing** → **Arts and humanities**; • **Computing methodologies** → *Natural language generation.*

Keywords

Creative Writing, Language Models, Poetry, Creativity Support Tool, Material, Artistic Support Tool

ACM Reference Format:
Alex Calderwood, Taewook Kim, Yuqian Sun, Melissa Roemmele, John Joon Young Chung, and Max Kreminski. 2025. Supporting Material Writing Practice with Phraselette, a Palette of Phrases. In *Designing Interactive Systems Conference (DIS '25 Companion), July 05–09, 2025, Funchal, Portugal.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3715668.3735608

1 *Phraselette*: A Contextual, Constrained Phrase Thesaurus

Phraselette is an *artistic support tool* [28] for experimental poets, meant to—rather than offload creative cognition onto machine assistance—*expand* the writer's possibility space by providing textual material only at sites where the writer requests it. It aims to integrate into creative writing practices by operating as a textual instrument [25, 37]. With use patterns analogous to a traditional thesaurus, it supports word search, and (unlike a thesaurus) does so with specificity to text context, with many avenues for creative control. Also unlike a traditional thesaurus, which supports only word-sized concepts, *Phraselette* allows search for multiword phrases. We drew design inspiration for *Phraselette* from our own personal practices of experimental writing; from literary theory's view of constraint [13, 22, 32, 36]; new media scholarship on the materiality of digital text [4, 17, 21], computational text criticism [18, 30, 33], from prior computational writing support tools [5, 6, 8, 14–16, 24, 27, 35, 38, 39]; and from critical engagement with language models [1–3, 7, 10, 11, 23, 31].

2 Background

Li et al. [28] point out that creativity support tools exert "*power over* their users" to construct a "normative ground" that structures users' "ideas, goals, and intentions". To align text generation with writers' creative activity, *Phraselette* is designed to deviate from the normative *automaticity* of many other AI-based writing support tools. We explode the unitary text generation process into many smaller, more transparent sub-processes (*phrasewells*) to support *material writing practice*: writing that engages with its own representation or the processes through which it arises [9, 17, 21].

Material writing support tools might provide writers avenues for priming their mental space or engaging with "second-generation originals and media info-bodies" as a springboard for poetic play [4]. In this view, a thesaurus might be viewed as a way for writers to

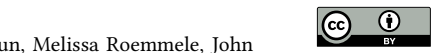

This work is licensed under a Creative Commons Attribution 4.0 International License.

DIS '25 Companion, Funchal, Portugal

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1486-3/25/07
https://doi.org/10.1145/3715668.3735608

Figure 1: The *Phraselette* interface. The editor (E1) allows writers to highlight text in order to create *an inlet*, a site for text revision. Users activate *phrasewells* with the side-scrolling menu (M1). Here the active phrasewells include a thesaurus (W1) which can be edited into different styles and roles with (E2) and a permanent well (W2) used for high level word constraints. Wells each provide some combination of *rephrasings* (R1, R2), *constraints* (C1, C2), *insights* (Figure 3) and *views* (V1)). A worked example is given in Section 3.2.

find relevant textual material, just as experimental collage practice views other texts as sources of inspiration. *Phraselette* thus provides opportunities for writers to find relevant textual material with the same granularity as the words given by a thesaurus.

Despite being framed in contrast to the automaticity of many existing large language model (LLM)-based applications, *Phraselette* is aimed to expose the affordances of statistical language models and the decoding algorithms that accompany them. The system lets users steer and look inside these models; it provides many suggestions, offers cues about the model's state, emphasizes the poet's mental space over the text itself, enables a plurality of customizable search criteria, and aligns these criteria with a writer's notion of constraint. We argue that certain affordances of language model technology, such as word retrieval capabilities and plain-text task specification, are, rather than analogous to or substitutive for writerly processes, able to provide a rich material work surface. Imperfect machine mirrors of the interpretations that human writers use to search for words can form a verdant ground where new ideas emerge [22, 32].

3 System Design

Phraselette is modular, supporting a variety of text generation strategies and modes of textual constraint. Disparate text generation procedures, search algorithms, models, and forms of data representation can thus be used in concert, each with different strengths and use cases that writers can tailor for their own purposes. The system connects text generation strategies, constraints, and insights into the conceptual unit of the *phrasewell* (well for short), each of which simultaneously provides new machine perspectives on text and new ways to constrain or generate alternative text.

Linking these affordances into phrasewells improves the system's modularity and extensibility. For instance, a programmer might want to introduce a WordNet-style conceptual ontology to *Phraselette*'s text representations; wells are designed with object oriented principles, so this extension can be made by implementing a new phrasewell class.

3.1 Phrasewells

Phrasewells are the main way that users interact with *Phraselette*, and can provide four distinct types of affordance:

3.1.1 Rephrasings. These are rewrite suggestions made by the Wells. Rephrasings appear within the Well UI element and are also aggregated to a window on the bottom of the screen, colored according to the Well in which they originated.

3.1.2 Constraints. Some Wells allow the writer to add constraints that reduce the search space of its own generation capabilities as well as the generation of other Wells. Constraints have two roles during rephrasing generation:

(1) To produce *advice* to the search algorithm (information that guides and directs the search)
(2) To *score* the rephrasings produced by the Well, for sorting=

Each type of Well that produces rephrasings has a unique search algorithm, so it is up to each Well to determine how to satisfy each constraint.. We refer to a Well modulated by a given constraint as receiving 'advice' from the constraint. For a complete list of currently supported phrasewell types, see Section 3.3.

3.1.3 Insights. Additional information that a Well provides. In some cases, insights are textual: such as workshop-style feedback generated according to a custom role. Insights can also be graphical; such as a word likelihood histogram. Insights function partly as mechanisms for the user to "see what the machine sees", and form a closer relationship to the generative system responsible for rephrasings.

3.1.4 Views. Some Wells provide a unique *view* on the selected text. The view may be thought of as a lens through which the document may be annotated. These Wells attach additional data to the words (tokens) for all rephrasings in the document. Wells that provide a view display it for the current selection, and to the expanded display shown when rephrasings are hovered over).

3.2 User Scenario: Working with *Phraselette*

To illustrate how *Phraselette* works in practice, we will demonstrate its use for revision. We present a user scenario: A poet, Charlie, wants to continue writing a series of poems that critiques the naivete of modernism. Charlie often works by modifying other texts, and here begins with William Carlos Williams' "The Red Wheelbarrow". Charlie opens the editor and pastes the poem (Figure 1).

Charlie recognizes the phrase "glazed with" is doing a lot to guide the temperament of the scene. Charlie highlights it in the editor and selects *Create Inlet*. The system will now maintain this highlighted span (or *inlet*) as they work, and aim to produce revisions of the selected text. The interface displays available text transformation tools, the phrasewells, in a top-righthand-side menu (Figure 1, **M1**). Each well represents a different approach to text transformation. Charlie activates a 'reader' Well by clicking its (+) icon. It appears in the right hand column alongside an initial 'words' Well.

Charlie wants this descriptive phrase to evoke a whimsical and actively alternative, technical tone but but they also know that it should describe the behavior of water. Charlie cycles through the included reader archetypes by pressing the Well's die icon until it reads "William S. Burroughs." In addition to this reader, Charlie adds a thesaurus Well (**W1**), which they know to be less impacted by the text surrounding the highlighted inlet; they hope this will produce more contrast with the current language. To describe how this thesaurus should behave, Charlie writes "Alfred Jarry's inverted dictionary of pataphysics" to target a whimsical and parodic style (**E2**). Finally, they add a context Well to see what words would be most likely to appear in the place of "glazed with" (which they want to avoid).

Now that Charlie has assembled an assortment of Wells, each with slightly different foci, they can generate alternate rephrasings for the selection either by triggering individual Wells or by running all active Wells simultaneously.

Charlie decides to move towards more active language in order to evoke a sense of critique. In a permanent Well that deals with word-level constraints, Charlie clicks the lockbox emoji, and using the constraint selection interface for part of speech, changes to a Verb-Adverb structure (**C2**).

After Charlie runs the Wells, they each provide various forms of text or other *insights* (Section 3.1.3). Hovering over any rephrasing shows all the view data associated with it, its phonemes/speech sounds, calculated likelihood scores, and more (Figure 3.1.4). Commonly, the user will see the insights and ideas on the screen and simply make note of the new diction or come up with a different idea to prod towards. Clicking on a rephrasing will place it in the current inlet, but a common use case is to absorb the responses given by the Wells and manually make a change to the current inlet or elsewhere. For *Phraselette*, success doesn't always entail direct user acceptance of machine suggestions—noting a related word or concept often leads to revisions across the text, and may lead to knock-on effects to structure and concepts [34, 39].

Charlie sees that *Phraselette* has elevated the rephrasings "vulcanized via" and "vitrified per" (Figure 1 **R1**). They need to look up what vitrified means (perhaps consulting the dictionary Well for an approximate definition)—to convert into glass by heat. They hover over the rephrasing, noting that the context Well has satisfyingly provided a low probability for the word, and click on it, replacing "glazed with" in the text. Noting the play of "vulcanized via", they edit the stanza to: "vitrified via rain / water". They note that vitrification via water entails that the water must be *hot*, so they shift focus and highlight the (red) "wheel barrow", creating a new inlet for the phrase without removing the previous inlet. Clicking 'Run Wells' restarts the search on this inlet with the same active Wells as before, allowing them to apply the same palette of word search tools to the new context.

3.3 Types of Phrasewells

3.3.1 Thesaurus. The thesaurus phrasewell is operated by entering a description of the thesaurus or word bank from which rephrasings should be drawn. Built-in examples include descriptions like "the thesaurus James Joyce used for Ulysses", the lexicon of a romance novel, a thesaurus to produce homophones, and a thesaurus to explore hypernymy or hyponymy relationships. The thesaurus phrasewell implements constraints by algorithmically constructing LLM prompts with constraint-adherence instructions; its prompts include the Inlet's selection text, but do not consider the context surrounding this text. This is unlike other Wells, which do take into account the surrounding context. Restricting the information used to inform generation changes the bent of the rephrasings; this decision showcases that a multiplicity of diverse generation

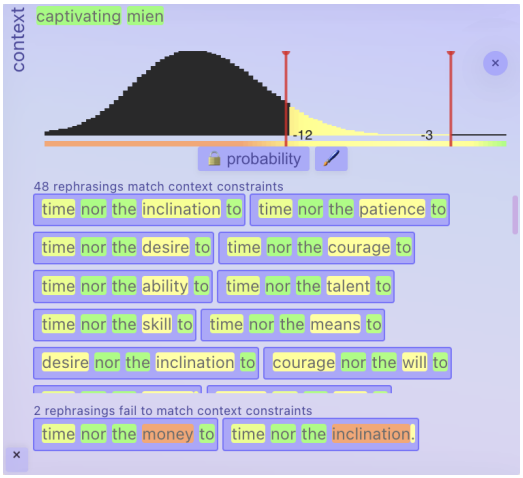

Figure 2: The context phrasewell generates within the bounds of the probability distribution set by the user.

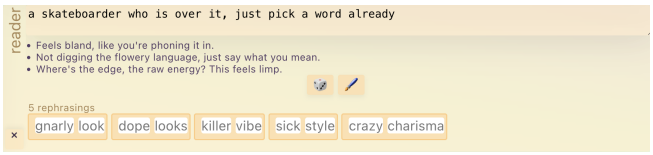

Figure 3: The reader phrasewell gives poets the ability to 'cast' an imagined persona into a role of reader to provide feedback and revision material.

strategies, each with different aims and affordances, can be used in concert.

3.3.2 Basic ('Words'). The 'Words' Well is the only phrasewell that cannot be disabled. It allows users to constrain text by word count and part of speech (POS) tags. As with other symbolic constraints, the user can specify sequences that rephrasings should contain, begin/end with, or have in order. The word count constraint provides advice to some wells by directly changing parameters to the beam-search process used in rephrasing. For others, it adds guidance to the prompt in plain English, e.g., "aim to produce between 1 and 4 words". The same is true for POS constraints, which provide advice both in the form of prompt injections and a beam search heuristic to guide the sampling process.

3.3.3 Reader. Like the thesaurus, the reader phrasewell allows users to describe the attributes of a desired reader persona in plain text. The user can rotate through predefined persona descriptions, or introduce their own. Built-in reader descriptions include surrealist writers ("Tristan Tzara, the Dadaist poet") and archetypes ("a literary critic..."). This description is then used to construct an LLM prompt which also includes the document's context and the current selection text. Upon triggering the reader Well, it first provides up to three bullet points that are generated to characterize a critique or positional reading of the current selection, given the perspective of the reader archetype in the provided description. These insights (or 'reader responses') are passed into a second prompt, which instructs an LLM to produce rephrasings in accordance to this advice. This two-step process is inspired by autodebugging [20] and the improvement of LLM writing via edit models [10]. It typically provides 5-12 rephrasings.

3.3.4 Context. This phrasewell surfaces the top 20-50 rephrasings according to a custom beam-search scoring algorithm [12, 19] which samples from among the most likely continuations to the text immediately preceding the Inlet, taking no account of the Inlet's highlighted text. The context phrasewell provides insights in the form of a histogram that shows the log-likelihood distribution of possible rephrasings according to the model. It also provides probability constraints, which can be used to limit rephrasings (across all Wells) to those of a certain likelihood level.

3.3.5 Sound. The Sound phrasewell provides constraints over the phoneme space (speech sounds) of words. Activating it introduces a phoneme-generation step to the rephrasing pipeline. Users can then constrain the sound of rephrasings to match parts of a reference pronunciation, e.g., to start with a 'k' sound or rhyme with the end of the current selection.

3.3.6 Dictionary. Like the thesaurus, this Well allows users to provide a text description to guide the system. It produces machine generated definitions to match a given style (*insights*). It takes into account the context around the current Inlet, as well as the selection text itself to produce the stylized definitions. This Well was included partly to demonstrate the breadth of possibilities provided by *Phraselette* framework. We observe that the thesaurus can often provide new coinages as well as unknown phrases; the dictionary can provide appropriately themed definition text.

4 User Study

We conducted an *extended user study* [29] with 10 recently published poets over the course of two weeks of use, aiming to understand how such a writing tool can be integrated within real world art-making practice [8]. Our qualitative coding of interview transcripts found that participants on the whole:

- Expressed positivity about the tool after using it for two weeks (Figure 4). One participant "a dependency growing in a good way".
- Used the system for revision, conceptual discovery, and translation.
- Described it in terms of play. Described moments where useful words and phrases appeared as highly rewarding.
- Described their role as curator, editor, or workshop participant.
- Found the multiplicity of options useful.
- Compared *Phraselette* postively to their impression of other AI systems, feeling higher levels of creative control and ownership over the generated text than they expected.

One user's description of *Phraselette* as a conceptual "chisel" suggests a view of the system as a defamiliarizing force [8], a "loser of poetry" and its meaning [22]. Broadly, our findings also support Li et al.'s theory of "normative ground" [28]; Roemmele's findings around "inspiration through observation" of machine-generated text [34]; and Lawton et al.'s findings regarding when users perceive co-creative systems as collaborators rather than tools [26].

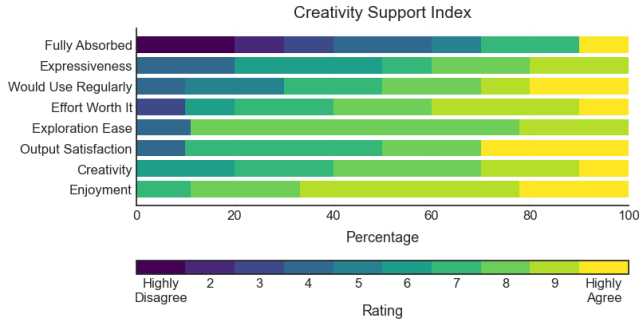

Figure 4: Creativity Support Index responses indicate that participants enjoyed using *Phraselette* and valued its output. In interviews, many mentioned that they were planning to pursue publication of work authored using the tool.

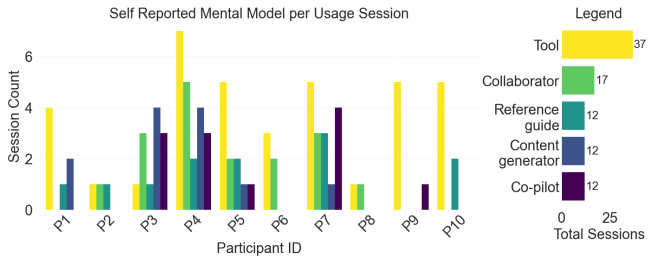

Figure 5: In response to "I used the system as a _____", writers most often answered 'Tool' after each prolonged interaction with Phraselette.

In particular, *Phraselette*'s placing of initiative firmly within the hands of the user may contribute to broad user agreement that the system is best viewed as a tool, rather than a collaborator or co-pilot (Figure 5).

References

[1] Dhruv Agarwal, Mor Naaman, and Aditya Vashistha. 2024. AI suggestions homogenize writing toward western styles and diminish cultural nuances. *arXiv preprint arXiv:2409.11360* (2024).

[2] Barrett R Anderson, Jash Hemant Shah, and Max Kreminski. 2024. Homogenization effects of large language models on human creative ideation. In *Proceedings of the 16th Conference on Creativity & Cognition*. 413–425.

[3] 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*. 128–138.

[4] Giselle Beiguelman. 2006. Nomadic Poetry. In *New Media Poetics: Contexts, Technotexts, and Theories*. MIT Press.

[5] Kyle Booten. 2023. Lotus Chorus Workshop: Designing for Cognitive Overload. In *Proceedings of the Eleventh Conference on Computation, Communication, Aesthetics & X*. i2ADS — Research Institute in Art, Design and Society. https://doi.org/10.34626/XCOAX.2023.11TH.185

[6] Kyle Booten and Katy Ilonka Gero. 2021. Poetry Machines: Eliciting Designs for Interactive Writing Tools from Poets. In *Creativity and Cognition*. ACM. https://doi.org/10.1145/3450741.3466813

[7] Daniel Buschek. 2024. Collage is the New Writing: Exploring the Fragmentation of Text and User Interfaces in AI Tools. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 2719–2737.

[8] Alex Calderwood, Vivian Qiu, Katy Ilonka Gero, and Lydia B Chilton. 2018. How Novelists Use Generative Language Models: An Exploratory User Study. In *23rd International Conference on Intelligent User Interfaces*. ACM.

[9] John Cayley. 2006. Time Code Language: New Media Poetics and Programmed Signification. In *New Media Poetics Contexts*. https://direct.mit.edu/books/book/2676/chapter-abstract/72465/Time-Code-Language-New-Media-Poetics-and?redirectedFrom=PDF

[10] Tuhin Chakrabarty, Philippe Laban, and Chien-Sheng Wu. 2024. Can AI writing be salvaged? Mitigating Idiosyncrasies and Improving Human-AI Alignment in the Writing Process through Edits. *arXiv preprint arXiv:2409.14509* (2024).

[11] Anil R Doshi and Oliver P Hauser. 2024. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances* 10, 28 (2024), eadn5290.

[12] Markus Freitag and Yaser Al-Onaizan. 2017. Beam Search Strategies for Neural Machine Translation. In *Proceedings of the First Workshop on Neural Machine Translation*. 56–60. https://doi.org/10.18653/v1/W17-3207 arXiv:1702.01806 [cs].

[13] Paul Fussell. 1979. Poetic Meter and Poetic Form. *University of Pennsylvania* (1979).

[14] Richard P. Gabriel, Jilin Chen, and Jeffrey Nichols. 2015. InkWell: A Creative Writer's Creative Assistant. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*. ACM, Glasgow United Kingdom, 93–102. https://doi.org/10.1145/2757226.2757229

[15] 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*. ACM, San Diego CA USA, 597–603. https://doi.org/10.1145/3325480.3326573

[16] 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*. ACM, Glasgow Scotland Uk, 1–12. https://doi.org/10.1145/3290605.3300526

[17] N. Katherine Hayles. 2006. The Time of Digital Poetry: From Object to Event. In *New Media Poetics*, Adalaide Morris and Thomas Swiss (Eds.). The MIT Press, 181–210. https://doi.org/10.7551/mitpress/5002.003.0013

[18] Daniel C Howe. 2013. Reading, writing, resisting: literary appropriation in *The Readers Project*. In *Proceedings of the 19th International Symposium on Electronic Art*. 178–181.

[19] Katnoria. 2020. Visualising Beam Search and Other Decoding Algorithms for Natural Language Generation. https://medium.com/voice-tech-podcast/visualising-beam-search-and-other-decoding-algorithms-for-natural-language-generation-fbba7cba2c5b

[20] Jack Kelly, Alex Calderwood, Noah Wardrip-Fruin, and Michael Mateas. 2023. There and Back Again: Extracting Formal Domains for Controllable Neurosymbolic Story Authoring. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 19, 1 (Oct. 2023), 64–74. https://doi.org/10.1609/aiide.v19i1.27502

[21] Matthew Kirschenbaum. 2001. Materiality and Matter and Stuff: What Electronic Texts Are Made Of. (2001). https://electronicbookreview.com/essay/materiality-and-matter-and-stuff-what-electronic-texts-are-made-of/

[22] Max Kreminski. 2024. Computational Poetry is Lost Poetry. In *Proceedings of the Halfway to the Future Symposium*. ACM, Santa Cruz CA USA. https://doi.org/10.1145/3686169.3686179

[23] Max Kreminski and Chris Martens. 2022. Unmet creativity support needs in computationally supported creative writing. In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*. 74–82.

[24] Max Kreminski and Michael Mateas. 2021. Reflective Creators. In *International Conference on Computational Creativity*.

[25] Max Kreminski and Michael Mateas. 2021. Toward Narrative Instruments. In *Interactive Storytelling*, Alex Mitchell and Mirjam Vosmeer (Eds.). Vol. 13138. Springer International Publishing, Cham, 499–508. https://doi.org/10.1007/978-3-030-92300-6_50 Series Title: Lecture Notes in Computer Science.

[26] Tomas Lawton, Kazjon Grace, and Francisco J Ibarrola. 2023. When is a tool a tool? User perceptions of system agency in human–AI co-creative drawing. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 1978–1996.

[27] Mina Lee, Katy Ilonka Gero, John Joon Young Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thiemo Wambsganss, David Zhou, Emad A. Alghamdi, Tal August, Avinash Bhat, Madiha Zahrah Choksi, Senjuti Dutta, Jin L. C. Guo, Md Naimul Hoque, Yewon Kim, Simon Knight, Seyed Parsa Neshaei, Agnia Sergeyuk, Antonette Shibani, Disha Shrivastava, Lila Shroff, Jessi Stark, Sarah Sterman, Sitong Wang, Antoine Bosselut, Daniel Buschek, Joseph Chee Chang, Sherol Chen, Max Kreminski, Joonsuk Park, Roy Pea, Eugenia H. Rho, Shannon Zejiang Shen, and Pao Siangliulue. 2024. A Design Space for Intelligent and Interactive Writing Assistants. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3613904.3642697

[28] Jingyi Li, Eric Rawn, Jacob Ritchie, Jasper Tran O'Leary, and Sean Follmer. 2023. Beyond the Artifact: Power as a Lens for Creativity Support Tools. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*.

ACM, San Francisco CA USA, 1–15. https://doi.org/10.1145/3586183.3606831

[29] Tao Long, Katy Ilonka Gero, and Lydia B Chilton. 2024. Not Just Novelty: A Longitudinal Study on Utility and Customization of an AI Workflow. In *Designing Interactive Systems Conference*. ACM, IT University of Copenhagen Denmark, 782–803. https://doi.org/10.1145/3643834.3661587

[30] Franco Moretti. 2005. *Graphs, maps, trees: abstract models for a literary history*. Verso.

[31] Vishakh Padmakumar and He He. 2023. Does Writing with Language Models Reduce Content Diversity? *arXiv preprint arXiv:2309.05196* (2023).

[32] Allison Parrish. 2015. Exploring (Semantic) Space With (Literal) Robots. http://opentranscripts.org/transcript/semantic-space-literal-robots/

[33] Stephen Ramsay. 2011. *Reading Machines: Toward and Algorithmic Criticism*. University of Illinois Press.

[34] Melissa Roemmele. 2021. Inspiration through observation: Demonstrating the influence of automatically generated text on creative writing. *arXiv preprint arXiv:2107.04007* (2021).

[35] Allen Roush, Sanjay Basu, Akshay Moorthy, and Dmitry Dubovoy. 2022. Most Language Models can be Poets too: An AI Writing Assistant and Constrained Text Generation Studio. In *Proceedings of the Second Workshop on When Creative AI Meets Conversational AI*, Xianchao Wu, Peiying Ruan, Sheng Li, and Yi Dong (Eds.). Association for Computational Linguistics, Gyeongju, Republic of Korea, 9–15. https://aclanthology.org/2022.cai-1.2/

[36] Jean-Jacques Thomas. 1988. Oulipo: A Primer of Potential Literature.

[37] Noah Wardrip-Fruin. 2007. Playable media and textual instruments. *The aesthetics of net literature: writing, reading and playing in programmable media* (2007), 211–80.

[38] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: story writing with large language models. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*. 841–852.

[39] David Zhou and Sarah Sterman. 2024. Ai.llude: Investigating Rewriting AI-Generated Text to Support Creative Expression. In *Proceedings of the 16th Conference on Creativity & Cognition*. 241–254.