

# Nine Potential Pitfalls when Designing Human-AI Co-Creative Systems

Daniel Buschek<sup>a</sup>, Lukas Mecke<sup>b,c</sup>, Florian Lehmann<sup>a</sup> and Hai Dang<sup>a</sup>

<sup>a</sup>Research Group HCI + AI, Department of Computer Science, University of Bayreuth, Bayreuth, Germany

<sup>b</sup>Bundeswehr University Munich, Munich, Germany

<sup>c</sup>LMU Munich, Munich, Germany

## Abstract

This position paper examines *potential pitfalls* on the way towards achieving human-AI co-creation with generative models in a way that is beneficial to the users' interests. In particular, we collected a set of nine potential pitfalls, based on the literature and our own experiences as researchers working at the intersection of HCI and AI. We illustrate each pitfall with examples and suggest ideas for addressing it. Reflecting on all pitfalls, we discuss and conclude with implications for future research directions. With this collection, we hope to contribute to a critical and constructive discussion on the roles of humans and AI in co-creative interactions, with an eye on related assumptions and potential side-effects for creative practices and beyond.

## Keywords

HCI, Artificial Intelligence, Co-Creation, Design

## 1. Introduction

Ongoing advances in generative AI systems have sparked great interest in using them interactively in creative contexts and for digital content creation and manipulation: Some examples include (1) generating or modifying images with generative adversarial networks (GANs) [1, 2, 3], (2) generating controllable movements for virtual characters with recurrent neural networks, deep reinforcement learning and physics simulations [4], and (3) controllable machine capabilities for generating or summarizing when working with text [5, 6]. Such computational methods have

also entered specifically artistic domains, including visual art [7], creative writing and poetry [8, 9]. More examples can be found in a curated “ML x Art” list<sup>1</sup>.

A common vision, also present in the call for this workshop, paints a picture of creative human use of such AI as tools. In this view, these new interactive systems are hoped to realise key ideas from *creativity support tools* (CST, [10]) by leveraging AI capabilities. More specifically, this support could cast humans and AI in many different roles (for a recent overview see [11]). This includes, for example, using AI as a divergent or convergent agent, as described by Hoffmann [12], that is, to generate or evaluate (human) ideas. Related, Kantosalo and Toivonen [13] highlight alternating co-creation, with the AI “pleasing” and “provoking” the user. Moreover, Negrete-Yankelevich and Morales-Zaragoza [14] describe a related set of roles, including AI as an “apprentice”,

*Joint Proceedings of the ACM IUI 2021 Workshops, April 13-17, 2021, College Station, USA*

✉ daniel.buschek@uni-bayreuth.de (D. Buschek);

lukas.mecke@unibw.de (L. Mecke);

florian.lehmann@uni-bayreuth.de (F. Lehmann);

hai.dang@uni-bayreuth.de (H. Dang)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings  
(CEUR-WS.org)

<sup>1</sup><https://mlart.co/>, last accessed 17.12.2020

whose work is judged and selectively chosen by humans, or a leader-like role, which only leaves final configurations to the user.

Within this range of roles, the workshop call emphasises the *generative* capabilities of AI. In this paper, we thus focus on the role of AI as a generator, and the underlying goal of freeing its users to focus on a larger creative vision, while the AI takes care of more tedious steps.

With this goal in mind, this paper examines *potential pitfalls* on the way towards achieving it in practice. Our research approach is related to work on *dark patterns* in UI/UX design [15], which also examines – sometimes speculatively [16], sometimes empirically [17] – what “could go wrong”, in order to ultimately inspire directions for interaction design that are beneficial to the users’ interests. In doing so, we thus hope to contribute to a critical and constructive discussion on the roles of humans and AI in co-creative interactions, with an eye on related assumptions and potential side-effects for creative practices and beyond.

## 2. Research Approach

Our interest in collecting pitfalls is inspired by work on dark patterns [16, 17, 15]: Both pitfalls and dark patterns identify issues with user interfaces and interactions that result in experiences or outcomes which might not be in the user’s best interest. However, in contrast to what is often assumed in dark patterns, pitfalls do not imply bad intention, rather oversight or lack of information<sup>2</sup>.

Concretely, related work collected speculative dark patterns for explainability, transparency and control in intelligent interactive systems [16] by transferring dark patterns previously described for UI/UX de-

---

<sup>2</sup><https://www.merriam-webster.com/dictionary/pitfall>, last accessed 17.12.2020

sign [15]. Other work collected dark UI/UX patterns empirically by reviewing a large set of existing mobile applications [17]. Both approaches seem challenging to directly transfer to collecting pitfalls in the context of co-creative generative AI, since there are no previously defined pitfalls and no easily accessible collections (or “app stores”) of many usable applications for review.

Therefore, we followed a qualitative, speculative approach and brainstormed on potential pitfalls, or “*what could go wrong*” (cf. [18]) in interactions with co-creative AI. Here we are loosely inspired by aspects of speculative design [19], although that area typically aims to address broader issues than what we focus on here. Further inspiring “speculative futures” for human-AI co-creative systems, along with a conceptual framework, can be found in the work by Bown and Brown [20].

We particularly explore issues grounded in today’s interactions and UIs, which can be reasonably well imagined to potentially occur with the current state of the art of generative AI models. In particular, our brainstorming started from three prompts: (1) Issues arising from currently *limited capabilities of AI*, and (2) from exploring what might happen with *too much AI involvement*; plus (3) thinking *beyond use* and usage situations. Considering this approach, we see the pitfalls presented here not as a comprehensive and “definitive” list but rather as a stimulus for discussion in the research community – at the workshop and beyond.

## 3. Nine Potential Pitfalls

Table 1 shows the pitfalls we collected. In particular, we present nine pitfalls, three for each of our starting prompts, that is, for limited AI (pitfalls 1-3), too much AI involvement (pitfalls 4-6), and for aspects beyond use (pitfalls 7-9).

| Name                                    | Affected aspects                       | Problem description   | Example  | How it might have happened (examples)  | How it might be addressed (examples)   |
|---|--|---|--|--|--|
| <b>Limited AI</b>                       |  |   |  |  |  |
| 1 Invisible AI boundaries               | model, creativity, exploration         | A (generative) AI component imposes unknown restrictions on creativity and exploration.                       | An AI face image editor cannot make faces bald without also turning them male-looking.   | Model with limited generalisability beyond training data, and entangled or nonsensical (latent) dimensions w.r.t. human understanding. | UI: Show boundaries e.g. via uncertainty, samples, precision/recall [21]. AI: Improve generalisability, disentanglement; consider narrowing scope. |
| 2 Lack of expressive interaction        | usability, creativity, exploration     | The UI imposes a "bottleneck" on creative use of the AI.  | Image generator is controlled with many 1D inputs for a high D latent space [22] - vs. rich image editor tools like brushes.   | Fine-grained AI control is difficult. "Conservative" UI design focused on ensuring input stays in (training) data distribution.        | Human-centred design with target group, e.g. to inform preferable tradeoffs of UI expressiveness and model "breaking points".                      |
| 3 False sense of proficiency            | trust, reliability                     | AI suggests answers or completions that the user cannot verify or that generate a false sense of proficiency. | When prompted to complete a sentence about the population of a large city the AI delivers a reasonable number that could be correct – but might not be.                          | Language models are capable of memorizing excerpts of text and reproducing them when prompted with a similar context.                  | Learn an additional model, that can attribute generated content to an explicit source to allow for verifying correctness.                          |
| <b>Too much AI</b>                      |  |   |  |  |  |
| 4 Conflicts of territory                | usability, UX, control                 | AI overwrites what the user has manually created/edited.  | In a co-creative text editor, the user replaces terms in generated text. Later, the AI (partly) reverts these changes.   | Language model optimised for word probability and user's term was less likely.   | Keep track of user edits to protect them, ask for confirmation before changes, or to integrate this info into inference.                           |
| 5 Agony of choice                       | usability, UX, productivity            | AI provides overwhelming amount/detail of content that distracts or creates agony of choice.                  | An AI photo editor displays an excessive number of suggested variants. The resulting small previews make it hard to discern and decide.  | UI design process was focused on showing AI capabilities instead of user needs.  | Clarifying use cases and support needs, responsive / malleable UI concepts, changeable user settings.  |
| 6 Time waster                           | usability, UX, productivity            | AI interrupts user or draws attention away from the creative task itself.                                     | A co-creative music composition tool continuously shows melody completions, which keep the user busy with exploring or understanding the system instead of their ideas.          | Same as above. Also: Timing of the AI's involvement not tested with users or varying preferences between users.                        | Same as above. Attention-aware UI (e.g. AI waits to not disrupt user's focused work; or stops suggestions if user has explored it for a while).    |
| <b>Beyond use</b>                       |  |   |  |  |  |
| 7 AI bias                               | accountability, fairness, transparency | AI suggestions are biased in a certain unwanted way, w.r.t. human meaning and values.                         | An AI story generator writes gender-stereotypical protagonists (e.g. w.r.t. roles/occupations).  | AI picked up biases in the training data or created bias through its learning method. Development process unaware of biases.           | Design for easy human revision/rejection. Addressing AI bias (e.g. see [23, 24]). Learning from user feedback/actions.                             |
| 8 Conflict of Creation & Responsibility | creativity, responsibility, ownership  | A system and a user collaborate to create an output. Ownership and responsibility are unclear.                | In a co-creative text editor the AI suggests formulations that appear verbatim in the training data. Who is the owner of the resulting text?                                     | Co-creative systems operate on a continuum between user and system creation, challenging attributions of ownership.                    | Should we attribute an AI and training data providers as contributors? Do we need systems to check for (accidental) plagiarism?                    |
| 9 User and Data Privacy                 | privacy, responsibility                | Private data may be exposed through the AI system or its training data.                                       | 1) A user A works with a cloud-based AI text creator and their data is transmitted unencrypted. 2) The AI reveals (private parts of) another user B's data to A (e.g. [25, 26]). | AI models are trained on a large corpus of data and can sometimes default to replicating this data when prompted.                      | Remove private information from training sets and work with AI either encrypted or locally.  |

**Table 1**

Overview of the collected pitfalls. Additionally, Figure 2 visualises one example for each of the categories "Limited AI", "Too much AI" and "Beyond use".

The table characterises each pitfall with a name, affected aspects (categories), a description of the problem, and a concise pitfall “vignette”: This includes an example scenario describing a system in which this issue arises, along with an illustrating diagnosis of how this might have happened in the design and development of said system, plus corresponding ideas for potential solutions or open questions. For each category of pitfalls (“limited AI”, “too much AI”, and “beyond use”) we picked one example for further illustration in Figure 2.

As an additional overview, Figure 1 locates these pitfalls within an interaction loop in human-AI co-creative systems; the loop is taken from a framework by Guzdial and Riedl [27]. In this figure we illustrate our underlying mental model of human-AI interaction. It consists of the *user* and the *AI* as potential actors collaborating on a shared *artifact*. The AI can get involved in the creation process in one of two ways: It can either be prompted to contribute through the user interface (e.g. using a predefined function to achieve an image manipulation) or it can act without a (user) prompt, e.g. to suggest edits or flag errors. We further include the training data in this model, as it provides the basis for the AI’s actions and decisions. While we located the pitfalls within this model, these locations are by no means the only possible ones. They represent our interpretations of which point in the interaction loop is most likely affected by each pitfall. As an example, a lack of expressive interaction may not only be rooted in the user interface, but can also be caused by insufficient training data to support more meaningful options.

## 4. Discussion

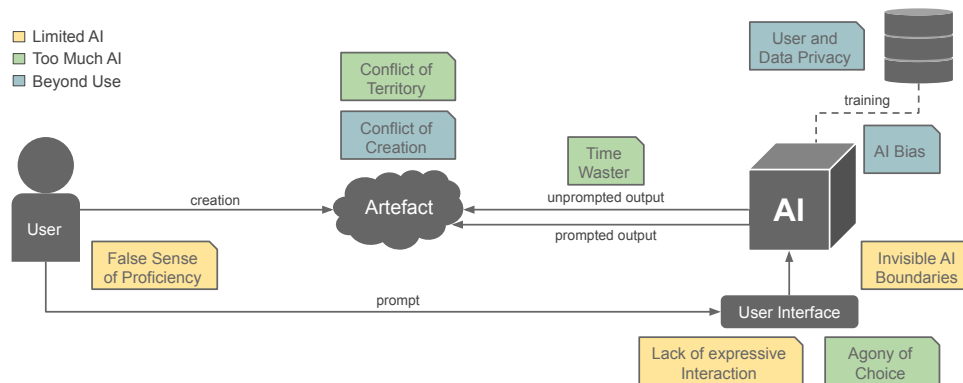
### 4.1. What are the Consequences of these Pitfalls?

While Table 1 lists concrete example problems, here we reflect more broadly on the consequences of such pitfalls for co-creative generative systems. In particular, we see two broad directions – overt and covert consequences.

First, users might be annoyed, distracted, or otherwise put off by bad user experiences due to these pitfalls. For example, cases where the AI directly overwrites the user (pitfall 4), or distracts the user from their productive task (pitfall 6) might be particularly harmful in this regard. Observing AI failures might lead to *algorithm aversion*, as described by Dietvorst et al. [29]. In these cases, users might avoid future use of such systems.

In contrast, users might also be affected negatively *without* noticing it. For example, this might be the case if the AI implies invisible boundaries (pitfall 1) that hinder creative exploration. Similarly, “silent” issues might result from the generative AI introducing incorrect information (pitfall 3), distractions (pitfall 6), or biases and legal issues (pitfalls 7-9). Users might only (much later) stumble across issues in downstream processes, evaluations or reflections. If such issues then affect evaluation of the user’s creative work (e.g. due to false information, pitfall 3), this might result in *algorithm anxiety*, described by Jhaver et al. [30].

Overall, the pitfalls might thus result in a range of possible consequences, from bad user experiences, negative impacts on creative work, abandonment of tools, to broader issues, including privacy related and legal ones.



**Figure 1:** Visualisation of our underlying mental model of the interaction loop in human-AI co-creative systems. We place our identified pitfalls (see Table 1) in this loop based on the position where they most likely occur.

## 4.2. How can the Pitfalls Inform Research and Design of Co-Creative Generative Systems?

Put briefly, this position paper describes what could go wrong in order to stimulate discussions of how to get it right. More concretely, here we describe three potential uses.

### 4.2.1. Raising Awareness of Design Considerations

The described pitfalls can help researchers and designers to think about a wide range of concrete aspects of interaction and UI design for co-creative generative systems (e.g. temporal and spatial integration of AI actions in UIs). In this way, they may *raise awareness for making design choices explicit* that might have otherwise not been prominently considered. These design choices could then also be considered in light of relevant frameworks, such as Horvitz’ mixed initiative principles [31] or the co-creative framework described by Guzdial and Riedl [27] (cf. Figure 1).

### 4.2.2. Informing Comparisons and Baselines

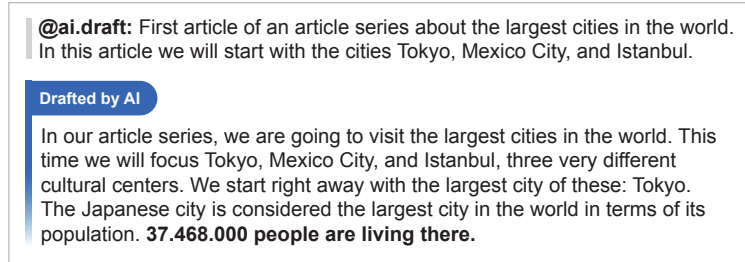
Moreover, the problematic systems described in the pitfalls in Table 1 might *inspire informative baseline systems for comparison* with (hopefully) better solutions. For example, a typical HCI user study on an AI photo editor might compare an AI vs non-AI version. However, as illustrated with the example for pitfall 5 (Agony of Choice), another insightful evaluation might further use a baseline that involves AI “even more” than the intended design solution to be evaluated.

### 4.2.3. Making the Criteria for Successful Design Explicit

Evaluating technical systems for creative use is challenging [32], for example, since creativity and quality criteria are often hard to operationalise, and may require interdisciplinary knowledge. Additionally involving AI can be expected to complicate evaluations further. Here, our pitfalls and examples may provide helpful concrete starting points, as a *thinking prompt towards developing evaluations* that satisfy both HCI and AI inter-



- (a) An AI photo editor displays an excessive number of suggestions. Due to the number of options and the small previews (making it hard to see what each option achieves) the user is left in an agony of choice.
- (b) Example for an AI leaking sensitive information from the training dataset (based on [25]), either as a suggestion or as a response to a primer (enabling active attacks). Such an attack has been demonstrated by Carlini et al. [26].



- (c) A text editing tool could offer intelligent features, e.g. drafting paragraphs or completing a sentence. Yet, the AI might not have the capability to refer to sources – to the human it remains unclear if the claims in a text are true. This leads to a false sense of proficiency. Here, the AI drafted a paragraph with claims about Tokyo’s approximate population (bold). However, it refers to the metropolitan area, not the city proper. The interface in the figure is inspired by Yang et al. [28].

**Figure 2:** Collection of visual examples for the pitfalls shown in Table 1. Here we show potential interfaces and situations in which selected pitfalls may occur, leading to (a) agony of choice, (b) a breach of privacy or (c) a false sense of proficiency.

ests. For instance, readers and workshop participants (with different backgrounds) could think about how they would improve the design – and evaluate it – for a concrete problematic example system in Table 1; and in particular how they might then make explicit and formulate their criteria in these cases.

### 4.3. Will the Pitfalls Vanish with Better AI?

One may ask if the illustrated issues might simply vanish in future systems that can build on better AI capabilities. Based on our

considerations here, we do not expect this to be case: Co-creative systems involving both human and AI actions are not only limited by AI capabilities. We also have to expect problems arising from interaction and UI design as well as from integration into creative human practices. For example, a lack of expressiveness in interactions (pitfall 2) can still cause problems for creative human use, even in a system with a powerful, “perfect” generative model under the hood.

In summary, the pitfalls highlight that human-AI co-creative systems sit at the intersection of HCI and AI, and that successful



designs need to consider human-centred aspects in the process. Our pitfalls reflect this in their mix of issues relating to interaction, UI and AI. We thus aim to motivate interdisciplinary work on such systems, also regarding research and design methodology.

## 5. Conclusion

One vision of interactive use of AI tools in co-creative settings focuses on the role of the AI as a generator that augments what people can achieve in creative tasks. This paper examined *potential pitfalls* on the way towards achieving this vision in practice, starting from three speculation prompts: Issues arising from (1) limited AI, (2) too much AI involvement, and (3) thinking beyond use and usage situations.

Concretely, we collected a set of nine potential pitfalls (Table 1) and discussed possible consequences and takeaways for researchers and designers along with illustrating examples. With this collection, we hope to contribute to a critical and constructive discussion on the roles of humans and AI in co-creative interactions, with an eye on related assumptions and potential side-effects for creative practices and beyond.

## Acknowledgments

This project is funded by the Bavarian State Ministry of Science and the Arts and coordinated by the Bavarian Research Institute for Digital Transformation (bidt).

## References

- [1] D. Bau, J.-Y. Zhu, H. Strobel, A. Lapedriza, B. Zhou, A. Torralba, Understanding the role of individual units in a deep neural network, Proceedings of the National Academy of Sciences (2020). URL: <https://www.pnas.org/content/early/2020/08/31/1907375117>. doi:10.1073/pnas.1907375117.
- [2] E. Härkönen, A. Hertzmann, J. Lehtinen, S. Paris, GANSpace: Discovering Interpretable GAN Controls (2020). URL: <https://arxiv.org/abs/2004.02546v1>.
- [3] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, T. Aila, Analyzing and improving the image quality of stylegan, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020.
- [4] S. Park, H. Ryu, S. Lee, S. Lee, J. Lee, Learning predict-and-simulate policies from unorganized human motion data, ACM Trans. Graph. 38 (2019). URL: <https://doi.org/10.1145/3355089.3356501>. doi:10.1145/3355089.3356501.
- [5] S. Dathathri, A. Madotto, J. Lan, J. Hung, E. Frank, P. Molino, J. Yosinski, R. Liu, Plug and Play Language Models: A Simple Approach to Controlled Text Generation, arXiv:1912.02164 [cs] (2020). URL: <http://arxiv.org/abs/1912.02164>, arXiv: 1912.02164.
- [6] S. Gehrmann, H. Strobel, R. Krüger, H. Pfister, A. M. Rush, Visual interaction with deep learning models through collaborative semantic inference, IEEE Transactions on Visualization and Computer Graphics 26 (2020) 884–894. doi:10.1109/TVCG.2019.2934595.
- [7] M. Akten, R. Fiebrink, M. Grierson, Learning to see: You are what you see, in: ACM SIGGRAPH 2019 Art Gallery, SIGGRAPH '19, Association for Computing Machinery, New York, NY, USA, 2019. URL: <https://doi.org/10.1145/3306211.3320143>. doi:10.1145/3306211.3320143.

- [8] K. I. Gero, L. B. Chilton, Metaphoria: An algorithmic companion for metaphor creation, in: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 1–12. URL: <https://doi.org/10.1145/3290605.3300526>. doi:10.1145/3290605.3300526.
- [9] M. Ghazvininejad, X. Shi, J. Priyadarshi, K. Knight, Hafez: an Interactive Poetry Generation System, in: Proceedings of ACL 2017, System Demonstrations, Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 43–48. URL: <https://www.aclweb.org/anthology/P17-4008>.
- [10] J. Frich, L. MacDonald Vermeulen, C. Remy, M. M. Biskjaer, P. Dalsgaard, Mapping the landscape of creativity support tools in hci, in: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 1–18. URL: <https://doi.org/10.1145/3290605.3300619>. doi:10.1145/3290605.3300619.
- [11] A. Kantosalo, A. Jordanous, Role-Based Perceptions of Computer Participants in Human-Computer Co-Creativity, in: 7th Computational Creativity Symposium at AISB 2020, 2020. URL: <https://research.aalto.fi/en/publications/role-based-perceptions-of-computer-participants-in-human-computer>.
- [12] O. Hoffmann, On modeling human-computer co-creativity, in: S. Kunifuji, G. A. Papadopoulos, A. M. Skulimowski, J. Kacprzyk (Eds.), Knowledge, Information and Creativity Support Systems, Springer International Publishing, Cham, 2016, pp. 37–48.
- [13] A. Kantosalo, H. Toivonen, Modes for Creative Human-Computer Collaboration: Alternating and Task-Divided Co-Creativity, in: Proceedings of the Seventh International Conference on Computational Creativity (ICCC 2016), Sony CSL, 2016, pp. 77–84. URL: <https://researchportal.helsinki.fi/en/publications/modes-for-creative-human-computer-collaboration-alternating-and-t>.
- [14] S. Negrete-Yankelevich, N. Morales-Zaragoza, The apprentice framework: planning and assessing creativity, Fifth International Conference on Computational Creativity, 2014.
- [15] C. M. Gray, Y. Kou, B. Battles, J. Hoggatt, A. L. Toombs, The dark (patterns) side of ux design, in: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 1–14. URL: <https://doi.org/10.1145/3173574.3174108>. doi:10.1145/3173574.3174108.
- [16] M. Chromik, M. Eiband, S. T. Völkel, D. Buschek, Dark patterns of explainability, transparency, and user control for intelligent systems, in: C. Tractner, D. Parra, N. Riche (Eds.), Joint Proceedings of the ACM IUI 2019 Workshops co-located with the 24th ACM Conference on Intelligent User Interfaces (ACM IUI 2019), Los Angeles, USA, March 20, 2019, volume 2327 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2019. URL: <http://ceur-ws.org/Vol-2327/IUI19WS-ExSS2019-7.pdf>.



- [17] L. Di Geronimo, L. Braz, E. Fregnan, F. Palomba, A. Bacchelli, Ui dark patterns and where to find them: A study on mobile applications and user perception, in: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 1–14. URL: <https://doi.org/10.1145/3313831.3376600>. doi:10.1145/3313831.3376600.
- [18] L. Colusso, C. L. Bennett, P. Gabriel, D. K. Rosner, Design and diversity? speculations on what could go wrong, in: *Proceedings of the 2019 on Designing Interactive Systems Conference*, DIS '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 1405–1413. URL: <https://doi.org/10.1145/3322276.3323690>. doi:10.1145/3322276.3323690.
- [19] A. Dunne, F. Raby, *Speculative Everything: Design, Fiction, and Social Dreaming*, MIT Press, 2013. Google-Books-ID: 9gQyAgAAQBAJ.
- [20] O. Bown, A. R. Brown, *Interaction Design for Metacreative Systems*, Springer International Publishing, Cham, 2018, pp. 67–87. URL: [https://doi.org/10.1007/978-3-319-73356-2\\_5](https://doi.org/10.1007/978-3-319-73356-2_5). doi:10.1007/978-3-319-73356-2\_5.
- [21] T. Kynkäänniemi, T. Karras, S. Laine, J. Lehtinen, T. Aila, Improved precision and recall metric for assessing generative models, in: *Advances in Neural Information Processing Systems*, 2019, pp. 3927–3936.
- [22] E. Härkönen, A. Hertzmann, J. Lehtinen, S. Paris, Ganspace: Discovering interpretable gan controls, arXiv preprint arXiv:2004.02546 (2020).
- [23] J. Buolamwini, T. Gebru, Gender shades: Intersectional accuracy disparities in commercial gender classification, in: S. A. Friedler, C. Wilson (Eds.), *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, PMLR, New York, NY, USA, 2018, pp. 77–91. URL: <http://proceedings.mlr.press/v81/buolamwini18a.html>.
- [24] D. S. Shah, H. A. Schwartz, D. Hovy, Predictive biases in natural language processing models: A conceptual framework and overview, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Online, 2020, pp. 5248–5264. URL: <https://www.aclweb.org/anthology/2020.acl-main.468>. doi:10.18653/v1/2020.acl-main.468.
- [25] N. Carlini, C. Liu, Ú. Erlingsson, J. Kos, D. Song, The secret sharer: Evaluating and testing unintended memorization in neural networks, in: *28th USENIX Security Symposium (USENIX Security 19)*, USENIX Association, Santa Clara, CA, 2019, pp. 267–284. URL: <https://www.usenix.org/conference/usenixsecurity19/presentation/carlini>.
- [26] N. Carlini, F. Tramer, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, D. Song, U. Erlingsson, A. Oprea, C. Raffel, Extracting training data from large language models, 2020. arXiv:2012.07805.
- [27] M. Guzdial, M. Riedl, An Interaction Framework for Studying Co-Creative AI, arXiv:1903.09709 [cs] (2019). URL: <http://arxiv.org/abs/1903.09709>, arXiv:1903.09709.

- [28] Q. Yang, J. Cranshaw, S. Amershi, S. T. Iqbal, J. Teevan, Sketching nlp: A case study of exploring the right things to design with language intelligence, CHI '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 1–12. URL: <https://doi.org/10.1145/3290605.3300415>. doi:10.1145/3290605.3300415.
- [29] B. J. Dietvorst, J. P. Simmons, C. Massey, Algorithm aversion: people erroneously avoid algorithms after seeing them err, *Journal of Experimental Psychology. General* 144 (2015) 114–126. doi:10.1037/xge0000033.
- [30] S. Jhaver, Y. Karpfen, J. Antin, Algorithmic anxiety and coping strategies of airbnb hosts, in: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 1–12. URL: <https://doi.org/10.1145/3173574.3173995>. doi:10.1145/3173574.3173995.
- [31] E. Horvitz, Principles of mixed-initiative user interfaces, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '99, Association for Computing Machinery, New York, NY, USA, 1999, p. 159–166. URL: <https://doi.org/10.1145/302979.303030>. doi:10.1145/302979.303030.
- [32] C. Lamb, D. G. Brown, C. L. A. Clarke, Evaluating computational creativity: An interdisciplinary tutorial, *ACM Comput. Surv.* 51 (2018). URL: <https://doi.org/10.1145/3167476>. doi:10.1145/3167476.