

Interactive Exploration-Exploitation Balancing for Generative Melody Composition

Yijun Zhou
The University of Tokyo
Tokyo, Japan
yjz@ui.is.s.u-tokyo.ac.jp

Masataka Goto
National Institute of Advanced Industrial Science and
Technology (AIST)
Tsukuba, Japan
m.goto@aist.go.jp

Yuki Koyama
National Institute of Advanced Industrial Science and
Technology (AIST)
Tsukuba, Japan
koyama.y@aist.go.jp

Takeo Igarashi
The University of Tokyo
Tokyo, Japan
takeo@acm.org

ABSTRACT

Recent content creation systems allow users to generate various high-quality content (e.g., images, 3D models, and melodies) by just specifying a parameter set (e.g., a latent vector of a deep generative model). The task here is to search for an appropriate parameter set that produces the desired content. To facilitate this task execution, researchers have investigated *user-in-the-loop optimization*, where the system samples candidate solutions, asks the user to provide preferential feedback on them, and iterates this procedure until finding the desired solution. In this work, we investigate a novel approach to enhance this interactive process: allowing users to control the sampling behavior. More specifically, we allow users to adjust the balance between *exploration* (i.e., favoring diverse samples) and *exploitation* (i.e., favoring focused samples) in each iteration. To evaluate how this approach affects the user experience and optimization behavior, we implement it into a melody composition system that combines a deep generative model with Bayesian optimization. Our experiments suggest that this approach could improve the user's engagement and optimization performance.

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools; User studies**; • **Information systems** → **Users and interactive retrieval**.

KEYWORDS

Generative design, human-in-the-loop machine learning, Bayesian optimization, creativity support tools

ACM Reference Format:

Yijun Zhou, Yuki Koyama, Masataka Goto, and Takeo Igarashi. 2021. Interactive Exploration-Exploitation Balancing for Generative Melody Composition. In *26th International Conference on Intelligent User Interfaces (IUI '21)*, April 14–17, 2021, College Station, TX, USA, <https://doi.org/10.1145/3397481.3450663>.

IUI '21, April 14–17, 2021, College Station, TX, USA

© 2021 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *26th International Conference on Intelligent User Interfaces (IUI '21)*, April 14–17, 2021, College Station, TX, USA, <https://doi.org/10.1145/3397481.3450663>.

'21), April 14–17, 2021, College Station, TX, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3397481.3450663>

1 INTRODUCTION

Many current content creation systems can automatically synthesize high-quality content from a set of user-specified parameters. In particular, with the recent development of deep learning techniques, the user is now able to easily and quickly obtain a high-quality image [12], 3D shape [3], or even melody [20] by just specifying a latent vector (i.e., a set of latent parameters) of a deep generative model trained for a particular domain. In such parametric content creation scenarios, the user's task is to search for an appropriate parameter set that produces the desired content. However, this task is not easy since the search space is often high-dimensional. To facilitate this task execution, researchers have investigated *user-in-the-loop optimization* [1, 5, 6, 15, 25], which is a human-AI collaborative process where the system samples candidate solutions, asks the user to provide preferential feedback on them, and iterates this procedure until finding the desired solution.

In this work, we investigate a novel approach to enhance this interactive process: allowing the user to control the sampling behavior. More specifically, we allow the user to explicitly adjust the balance between *exploration* (i.e., favoring diverse samples in the candidate set) and *exploitation* (i.e., favoring focused samples in the candidate set) in each iteration. This approach is motivated by a simple model of the creative process (which is sometimes implicitly assumed [23]): the *ideation* stage comes earlier to obtain serendipitous inspirations, and then the *refinement* stage comes later to improve the idea. That is, we take the analogy between the exploration-exploitation transition in the sampling strategy and the ideation-refinement transition in the creative process. We consider that the transition is often seamless and non-linear, and content creators purposefully control the balance in manual workflows. Thus, we expect that this additional controllability has the potential to improve user experience with the optimization process.

To evaluate how this approach affects the user experience and optimization behavior, we implement it into an existing melody composition system [28], which combines a deep generative model [20] with a user-in-the-loop *Bayesian optimization* (BO) technique

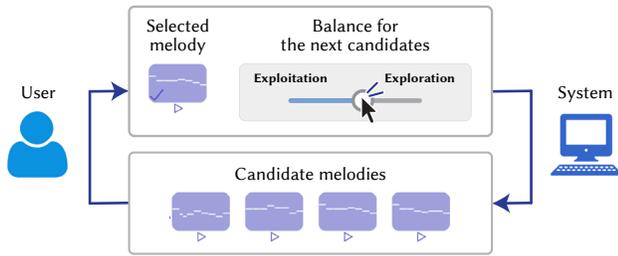


Figure 1: Interactive exploration-exploitation balancing in the search for the desired content in user-in-the-loop optimization processes. We allow the user to manually adjust the balance between *exploration* and *exploitation* in each iteration to control the variations of the system-generated candidates. Here, we focus on generative melody composition as the representative task, where the goal is to find an appropriate latent vector that generates the desired melody.

[16]. In each iteration, the user can explicitly specify the exploration-exploitation balance via a slider interface, by which the variation in the sampled candidate melodies can be controlled (Figure 1). We achieve this functionality by adapting the sampling strategy in BO determined by *acquisition functions* [22]. We conducted a simulated experiment, suggesting that the optimization performance could benefit from an appropriate balancing, and also conducted a user study with novice composers, suggesting that our approach could improve user engagement and satisfaction and better support their creativity.

2 RELATED WORK

Human-in-the-Loop Optimization. Optimization problems often involve human judgment (e.g., preference) in their objectives. In such cases, having humans in the optimization loop can be a promising solution. *Interactive evolutionary computation* [25] is one of such approaches investigated for decades. More recently, BO [22] has been extended for the human-in-the-loop use such that it accepts *preferential feedback* (i.e., given multiple options, the evaluator chooses the most preferable one) [2]. Since human-in-the-loop BO is efficient in terms of the number of necessary queries to humans, this approach has been actively investigated [1, 6, 15, 16, 28]. While BO can automatically balance exploration and exploitation by using an appropriate acquisition function when sampling new candidate solutions [22], no previous work investigates the manual approach to the balance. Our work adds user control over the acquisition function, and we investigate how this control affects the user-in-the-loop optimization process in performance and user experience.

Exploration-Exploitation Balancing in Human-Computer Interaction. The exploration-exploitation balance and similar concepts have been discussed in various human-computer interaction (HCI) contexts. For example, Koch et al. [14] presented an image recommendation system for ideation in designing mood boards, which automatically controls the exploration-exploitation balance to provide reasonable recommendations for ideation. Maudet [19] suggested to design recommendation systems so that they allow the user to manually choose from different recommendation algorithms

to switch the priority between “discovery” (exploration) and “comfort” (exploitation). Davis et al. [8] presented a human-computer co-drawing application, which offers the functionality to manually control the level of “creativity” (the degree of exploration) in the system-generated line drawings.

Creativity Support Tools for Music Composition. Designing creativity support tools has been a “grand challenge” in HCI [11, 23]. We are especially interested in designing effective tools for music content creation [10, 17, 18]. Younker et al. [27] stated that the standard process of novices to compose music consists of two steps: capturing a vague idea and refining it. Thus, we consider that it is especially important for music content creation tools to allow novice composers to explore and refine their vague ideas in a controllable way. We investigate an approach to this direction in the context of generative melody composition.

3 SYSTEM AND INTERACTION OVERVIEW

We implement the capability of interactive exploration-exploitation balancing into an existing generative melody composition system for novice composers [28]. This is a web application consisting of *composing* and *searching* modes. The composing mode offers basic functionalities to compose melodies via a piano-roll interface, and the user can toggle the searching mode when necessary. The searching mode (Figure 2) runs the user-in-the-loop BO process [16] to find an appropriate two-bar melody generated by a pre-trained variational autoencoder (VAE) called MusicVAE [20]. VAE is a variant of autoencoders, whose training is regularized so that the latent space has a smooth probabilistic distribution [13]. The latent space was reduced beforehand from 512 dimensions to 4 dimensions using another small VAE [9] so that BO works effectively [28]; refer to the previous work [28] for details. In each iteration of the user-in-the-loop-BO process, the system provides the user with four new candidate melodies sampled by BO and asks the user to select the preferable one. Optionally, the user can edit the selected candidate in the editing widget. Finally, the user specifies the exploration-exploitation balance for the next iteration via a simple slider to control the variety of the next candidates, which is the difference from the previous system [28], and then goes to the next iteration. This iteration finishes when the user feels satisfied with the found melody.

4 EXPLORATION-EXPLOITATION BALANCING

In BO methods, candidates are determined by an *acquisition function* [22], which evaluates the effectiveness of a candidate as the next query. We denote by $a : \mathcal{Z} \rightarrow \mathbb{R}$ an acquisition function, where \mathcal{Z} is the search space (the 4-dimensional latent space in our case). In the simple case of sampling one candidate in each iteration, it is determined as $z^{\text{next}} = \arg \max_{z \in \mathcal{Z}} a(z)$. Sampling multiple candidates in each iteration can be achieved similarly by a simple extension [21]. There are several acquisition functions that can automatically balance exploration (i.e., favor sampling from unvisited regions) and exploitation (i.e., favor sampling from high-expectation regions). One of the most popular choices is the *expected improvement* (EI), which has been used in previous human-in-the-loop BO methods [2, 6, 15, 16, 28].

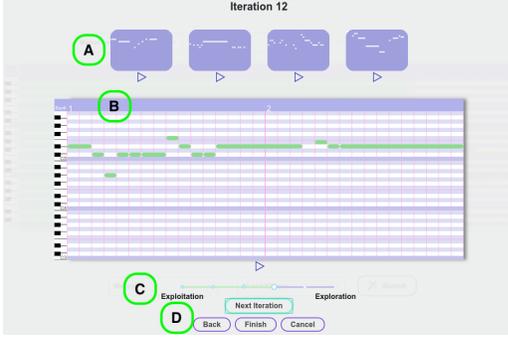


Figure 2: Interface for the search. (A) Current candidate melodies that BO sampled. (B) Editing widget. (C) Slider for specifying the exploration-exploitation balance for the next iteration. (D) Buttons to control the workflow.

In this work, the balancing mechanism needs to be controllable by the user, which cannot be achieved by EI. Thus, we propose to use another popular acquisition function called *Gaussian process upper confidence bound* (GP-UCB) [24]. GP-UCB fits our needs since it has a hyperparameter that can be set by the user to directly control the balance. The GP-UCB acquisition function at the i -th iteration is defined as follows. The posterior distribution of the objective value (i.e., user’s preference) at \mathbf{z} (which we denote by $f(\mathbf{z})$) follows a Gaussian distribution under the Gaussian process prior [22], and thus can be written as $f(\mathbf{z}) \sim \mathcal{N}(\mu_i(\mathbf{z}), \sigma_i^2(\mathbf{z}))$. Then, the GP-UCB value is calculated by

$$a_i^{\text{GP-UCB}}(\mathbf{z}) = \mu_i(\mathbf{z}) + \beta_i^{\frac{1}{2}} \sigma_i(\mathbf{z}) \quad (1)$$

where $\beta_i \geq 0$ is a hyperparameter that controls the balance; a larger value weighs exploration more, and a smaller value weighs exploitation more. In our system, the exploration-exploitation slider is mapped to β_i , where the six dots from left to right corresponds to 0.0, 0.01, 0.04, 0.16, 0.64, and 2.56.

5 SIMULATED EXPERIMENT

We conducted a simulated experiment to understand the controllability and performance of our approach, in which user response was simulated artificially. We used soft Dynamic Time Warping (s-DTW) [7] as the metric to measure the dissimilarity between two melodies, following previous work [28].

In this experiment, we compared five conditions: **Automatic** (EI), **Exploration** (GP-UCB with $\beta = 2.56$), **Intermediate** (GP-UCB with $\beta = 0.16$), **Exploitation** (GP-UCB with $\beta = 0.0$), and **Adaptive** (GP-UCB with an adaptive β). The last condition is for simulating the typical user behavior that we expected: increasing β when the current melody is far from the desired one, and *vice versa*. We set β based on the s-DTW value of the selected one, according to an empirical rule: (1) $s\text{-DTW} \leq 10.0$: $\beta = 0.0$; (2) $10.0 < s\text{-DTW} \leq 100.0$: $\beta = 0.01$; (3) $100.0 < s\text{-DTW} \leq 300.0$: $\beta = 0.04$; (4) $300.0 < s\text{-DTW} \leq 700.0$: $\beta = 0.16$; (5) $700.0 < s\text{-DTW} \leq 1500.0$: $\beta = 0.64$; (6) $1500.0 < s\text{-DTW}$: $\beta = 2.56$. These values were selected based on our observation.

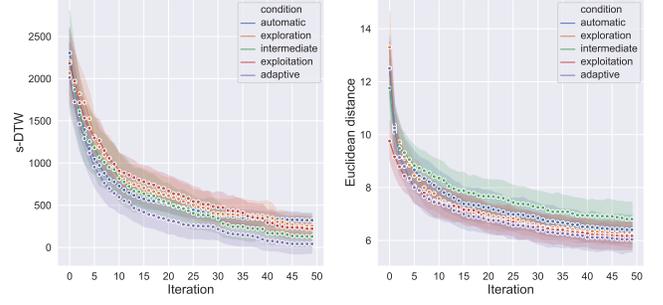


Figure 3: Residuals in the simulated experiment. Each plot shows the mean of 100 trials for each condition with the colored region showing the standard deviation. The left shows s-DTW, and the right shows the Euclidean distance in the search space.

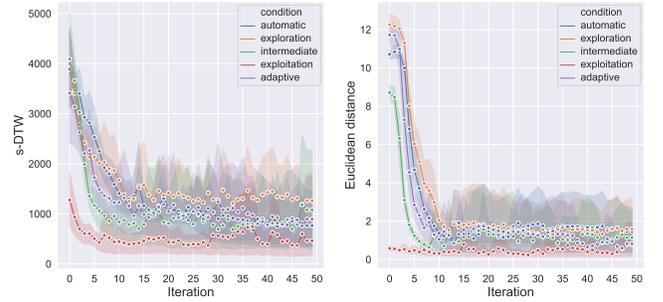


Figure 4: Candidate variations in the simulated experiment. Each plot shows the average distance of every pair of candidates in each iteration of 100 trials for each condition with the colored region showing the standard deviation. The left shows s-DTW, and the right shows the Euclidean distance in the search space.

Each experiment trial was conducted as follows. (1) Generate a desired melody by randomly sampling on the search space; (2) In the t -th iteration ($t = 1, \dots, 50$), select the one with the smallest s-DTW to the desired melody as the preferable one from the candidate melodies. We recorded both s-DTW values and Euclidean distances in the search space as residuals through iterations. We performed 100 trials for each condition.

As a result (Figure 3), we observe that the s-DTW values decreased as iterations went on in every condition. Among them, **Adaptive** could approach to the desired melody the most quickly. We also calculated the average s-DTW and Euclidean distance among every possible pair of candidates in each iteration (Figure 4), which indicates the variety in the candidates. We observe that, in every condition, the values decreased quickly in the first 10 iterations and then became relatively stable, where **Exploitation** was the smallest, **Exploration** was the largest, and the others stayed intermediate. From these observations, we conclude that GP-UCB can effectively control the variation of candidates, and the **Adaptive** condition, which simulates the user behavior in our approach, can work efficiently.

6 USER STUDY

We conducted a user study to evaluate how effectively our **manual** method (i.e., GP-UCB with user-controllable β s) can improve user experience and satisfaction in comparison with the **auto** method (i.e., EI), which is considered a baseline. We followed a within-subject design with counter-balancing.

Participants. We recruited twelve participants, consisting of seven females and five males, aged from 23 to 37 (mean: 26). They had diverse musical backgrounds; three had no formal music training, four had years of piano training at their early ages, and the others had played other instruments such as accordion, ukulele, and flute. To ensure that they were novices in composition, we recruited them in accordance with the following criteria: (1) wish to compose music; (2) know little/no music theory and have little/no formal training on composition; and (3) have not or hardly used composing software and can control the note sheet. Each participant received a 10-USD gift card as a reward.

Procedure and Task. We conducted the user study in a remote way without face-to-face contact. Each participant was first tutored by the experimenter about the usage and the features of the system (5 minutes). Then, they were asked to finish two tasks. In each task, they were asked to search for a two-bar melody with either the **manual** or **auto** method. In the first iteration, they were asked to select a two-bar melody they preferred from randomly sampled candidates. Then, they were asked to search for their desired two-bar melody based on the selected one in the first iteration. They were also asked to perform at least 25 iterations and finish each task within 10 minutes. Finally, they filled post-study questionnaires and completed a semi-structured interview (10 minutes).

Questionnaires. To evaluate user experience of the creative process, we used Creativity Support Index (CSI) [4], which is a standardized psychometric tool for assessing the perceived creativity support of a tool. It consists of six factors: **enjoyment**, **exploration**, **expressiveness**, **immersion**, **results worth efforts**, and **collaboration**. Even though our system does not involve **collaboration**, we included it as suggested by the original paper [4]. According to the original paper [4], we chose the **weighted score** for all the paired t -test, which was calculated by multiplying a participant's factor score by the factor count. The goal is to make the weighted score more sensitive to the factors that are more important considered by the participant. Our questionnaires also included a question that asks the **satisfaction** score on their melody compositions in a ten-point scale as CSI. Finally, our questionnaires ask the participants to indicate their **preference** over the **manual** and **auto** methods.

Semi-Structured Interview. We conducted a semi-structured interview for each participant, focusing on the overall experience, perceived advantages and disadvantages, and the possible improvement. In particular, we asked their feeling and experience of using the exploration-exploitation slider, and how this interaction influenced their experience.

Questionnaires Results. In the results of CSI (Table 1), the **manual** method outperformed the **auto** method in all factors, and there is statistical significance in **enjoyment** ($p < 0.01$), **exploration**

($p < 0.05$), **expressiveness** ($p < 0.01$), **results worth efforts** ($p < 0.01$), and the overall CSI score ($p < 0.01$). We can also observe from the counts that the participants considered **exploration** and **expressiveness** as important factors in the composition task. The mean **satisfaction** scores of the **manual** and **auto** methods were 8.25 (SD: 0.87) and 6.83 (SD: 2.08), respectively, which shows statistical significance ($p = 0.041$). As for the **preference**, eleven out of twelve preferred the **manual** method over the **auto** method.

Interview Results. Nine participants indicated positive feelings of control and engagement when using the **manual** method as the reason they preferred the **manual** method: “It makes me feel that I can express my idea to the system as the slider gives me more options and control” (P9); “It gives me the opportunity to indicate my own preference, instead of only relying on the machine’s result” (P12); and “It’s easier to indicate my desired exploration direction with the slider, which makes the experience more engaging” (P3). Also, two participants mentioned that the **manual** method inspired their creativity: “It’s more interesting for me to be able to control the system. The results can be tuned to my idea and which makes me more creative and inspires me with more ideas” (P6); and “It’s a great and brand-new experience. During the iterations, I sometimes changed my mind on how the desired melody would be, and the manual method made it possible for me to tell the system” (P5).

7 DISCUSSION

The dimensionality of the search space in our system is four, which we set following previous work [28]. However, this dimensionality is relatively lower than those of typical deep generative models. Investigating how the exploration-exploitation balancing works in higher-dimensional spaces is an important future work. As for the user interface, the observability is currently limited when the user interacts with the exploration-exploitation slider; a mechanism to show the impact of the slider value on music generation would help the user understand the interaction flow. While we focused on melody composition as the representative task, parametric design scenarios are very common in various domains, such as photo editing [15, 16] and 3D modeling [26]. We expect that our findings can be applied to those scenarios, but further investigation is necessary. Our simulated experiment suggested the effectiveness of the adaptive balancing approach. That is, the performance of BO can be improved if the exploration-exploitation balance is properly controlled. This finding is interesting in both user-in-the-loop and non-user-in-the-loop BO settings, and further analysis is important.

8 CONCLUSION

We found that the exploration-exploitation balancing was effective in generative melody composition, a user-in-the-loop optimization problem. Our user study suggested that our approach could enhance the user’s control and engagement, inspire their creativity, and help them express themselves better through the manual control of the balance. We also found that participants could be more satisfied with the found results with this manual approach. We hope our work inspires researchers to investigate deeper into designing creativity support tools by incorporating this approach.

Table 1: The results of the creativity support index (CSI) in the user study. Overall, the manual method (ours) obtained higher scores than the auto method (baseline).

Factor	Counts	Manual method (ours)		Auto method (baseline)		<i>p</i> -value
		Score (SD)	Weighted score (SD)	Score (SD)	Weighted score (SD)	
Enjoyment	3.00	17.17 (1.70)	51.50 (5.10)	11.17 (3.79)	33.50 (11.36)	5e-5
Exploration	3.83	15.50 (3.63)	59.42 (13.92)	11.33 (3.92)	43.44 (15.01)	.013
Expressiveness	3.33	15.67 (2.06)	52.22 (6.87)	10.58 (4.14)	35.28 (13.81)	.001
Immersion	2.67	15.42 (2.19)	41.45 (6.01)	14.33 (1.72)	37.58 (4.21)	.192
Results Worth Efforts	1.50	15.42 (1.38)	23.18 (2.16)	12.17 (3.52)	17.73 (5.19)	.007
Collaboration	0.67	10.67 (0.89)	7.03 (0.55)	10.58 (1.44)	6.97 (0.92)	.866
CSI			78.16 (8.63)		57.89 (15.69)	.001

ACKNOWLEDGMENTS

This work was supported in part by JST CREST Grant Number JPMJCR17A1, Japan, and JST ACCEL Grant Number JPMJAC1602, Japan.

REFERENCES

- [1] Eric Brochu, Tyson Brochu, and Nando de Freitas. 2010. A Bayesian Interactive Optimization Approach to Procedural Animation Design. In *Proc. SCA '10*. 103–112. <https://doi.org/10.2312/SCA/SCA10/103-112>
- [2] Eric Brochu, Nando de Freitas, and Abhijeet Ghosh. 2007. Active Preference Learning with Discrete Choice Data. In *Proc. NIPS '07*. 409–416. <http://papers.nips.cc/paper/3219-active-preference-learning-with-discrete-choice-data>
- [3] Zhiqin Chen and Hao Zhang. 2019. Learning implicit fields for generative shape modeling. In *Proc. CVPR 2019*. 5939–5948. <https://doi.org/10.1109/CVPR.2019.00609>
- [4] Erin Cherry and Celine Latulipe. 2014. Quantifying the Creativity Support of Digital Tools through the Creativity Support Index. *ACM Trans. Comput.-Hum. Interact.* 21, 4, Article 21 (June 2014), 25 pages. <https://doi.org/10.1145/2617588>
- [5] Chia-Hsing Chiu, Yuki Koyama, Yu-Chi Lai, Takeo Igarashi, and Yonghao Yue. 2020. Human-in-the-Loop Differential Subspace Search in High-Dimensional Latent Space. *ACM Trans. Graph.* 39, 4 (2020), 85:1–85:15. <https://doi.org/10.1145/3386569.3392409>
- [6] Toby Chong, I-Chao Shen, Issei Sato, and Takeo Igarashi. 2021. Interactive Optimization of Generative Image Modelling using Sequential Subspace Search and Content-based Guidance. *Comput. Graph. Forum* (2021). <https://doi.org/10.1111/cgf.14188>
- [7] Marco Cuturi and Mathieu Blondel. 2017. Soft-DTW: a Differentiable Loss Function for Time-Series. In *Proc. ICML '17*. 894–903. <http://proceedings.mlr.press/v70/cuturi17a.html>
- [8] Nicholas Davis, Chih-Pin Hsiao, Kunwar Yashraj Singh, Lisa Li, and Brian Magerko. 2016. Empirically Studying Participatory Sense-Making in Abstract Drawing with a Co-Creative Cognitive Agent. In *Proc. IUI '16*. 196–207. <https://doi.org/10.1145/2856767.2856795>
- [9] Monica Dinulescu, Jesse Engel, and Adam Roberts. 2019. MidiMe: Personalizing a MusicVAE model with user data. In *Proc. NeurIPS Workshop on Machine Learning for Creativity and Design '19*.
- [10] Morwaread M. Farbood, Egon Pasztor, and Kevin Jennings. 2004. Hyperscore: A Graphical Sketchpad for Novice Composers. *IEEE Comput. Graph. Appl.* 24, 1 (Jan. 2004), 50–54. <https://doi.org/10.1109/MCG.2004.1255809>
- [11] Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, and Peter Dalsgaard. 2019. Mapping the Landscape of Creativity Support Tools in HCI. In *Proc. CHI '19*. 389:1–389:18. <https://doi.org/10.1145/3290605.3300619>
- [12] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *Proc. ICLR 2018*. <https://openreview.net/forum?id=Hk99zCeAb>
- [13] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *Proc. ICLR '14*. <http://arxiv.org/abs/1312.6114>
- [14] Janin Koch, Andrés Lucero, Lena Hegemann, and Antti Oulasvirta. 2019. May AI? Design Ideation with Cooperative Contextual Bandits. In *Proc. CHI '19*. 633:1–633:12. <https://doi.org/10.1145/3290605.3300863>
- [15] Yuki Koyama, Issei Sato, and Masataka Goto. 2020. Sequential Gallery for Interactive Visual Design Optimization. *ACM Trans. Graph.* 39, 4 (July 2020), 88:1–88:12. <https://doi.org/10.1145/3386569.3392444>
- [16] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential Line Search for Efficient Visual Design Optimization by Crowds. *ACM Trans. Graph.* 36, 4 (July 2017), 48:1–48:11. <https://doi.org/10.1145/3072959.3073598>
- [17] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J. Cai. 2020. Novice-AI Music Co-Creation via AI-Steering Tools for Deep Generative Models. In *Proc. CHI '20*. 1–13. <https://doi.org/10.1145/3313831.3376739>
- [18] Charles Martin, Henry Gardner, Ben Swift, and Michael Martin. 2016. Intelligent Agents and Networked Buttons Improve Free-Improvised Ensemble Music-Making on Touch-Screens. In *Proc. CHI '16*. 2295–2306. <https://doi.org/10.1145/2858036.2858269>
- [19] Nolwenn Maudet. 2019. Dead Angles of Personalization: Integrating Curation Algorithms in the Fabric of Design. In *Proc. DIS '19*. 1439–1448. <https://doi.org/10.1145/3322276.3322322>
- [20] Adam Roberts, Jesse Engel, Colin Raffel, Curtis Hawthorne, and Douglas Eck. 2018. A Hierarchical Latent Vector Model for Learning Long-Term Structure in Music. In *Proc. ICML '18*. 4364–4373.
- [21] Matthias Schonlau, William J. Welch, and Donald R. Jones. 1998. *Global versus local search in constrained optimization of computer models*. Lecture Notes–Monograph Series, Vol. 34. Institute of Mathematical Statistics, Hayward, CA, 11–25. <https://doi.org/10.1214/lnms/1215456182>
- [22] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. 2016. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proc. IEEE* 104, 1 (January 2016), 148–175. <https://doi.org/10.1109/JPROC.2015.2494218>
- [23] Ben Shneiderman. 2007. Creativity Support Tools: Accelerating Discovery and Innovation. *Commun. ACM* 50, 12 (Dec. 2007), 20–32. <https://doi.org/10.1145/1323688.1323689>
- [24] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. 2012. Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting. *IEEE Trans. Inf. Theory* 58, 5 (May 2012), 3250–3265. <https://doi.org/10.1109/TIT.2011.2182033>
- [25] Hideyuki Takagi. 2001. Interactive Evolutionary Computation: Fusion of the Capabilities of EC Optimization and Human Evaluation. *Proc. IEEE* 89, 9 (Sep. 2001), 1275–1296. <https://doi.org/10.1109/5.949485>
- [26] Jerry O. Talton, Daniel Gibson, Lingfeng Yang, Pat Hanrahan, and Vladlen Koltun. 2009. Exploratory Modeling with Collaborative Design Spaces. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 167:1–167:10. <https://doi.org/10.1145/1618452.1618513>
- [27] Betty Anne Younker and William H. Smith. 1996. Comparing and Modeling Musical Thought Processes of Expert and Novice Composers. *Bulletin of the Council for Research in Music Education* 128 (1996), 25–36. <http://www.jstor.org/stable/40318786>
- [28] Yijun Zhou, Yuki Koyama, Masataka Goto, and Takeo Igarashi. 2020. Generative Melody Composition with Human-in-the-Loop Bayesian Optimization. In *Proceedings of the 2020 Joint Conference on AI Music Creativity (CSMC-MuMe 2020)*. <https://arxiv.org/abs/2010.03190>