Impression Sliders: On-Demand Sliders from Impression Descriptions for Parametric Design

Taiyu Homma
Meiji University
Tokyo, Japan
cs262039@meiji.ac.jp

Homei Miyashita
Meiji University
Tokyo, Japan
homei@homei.com

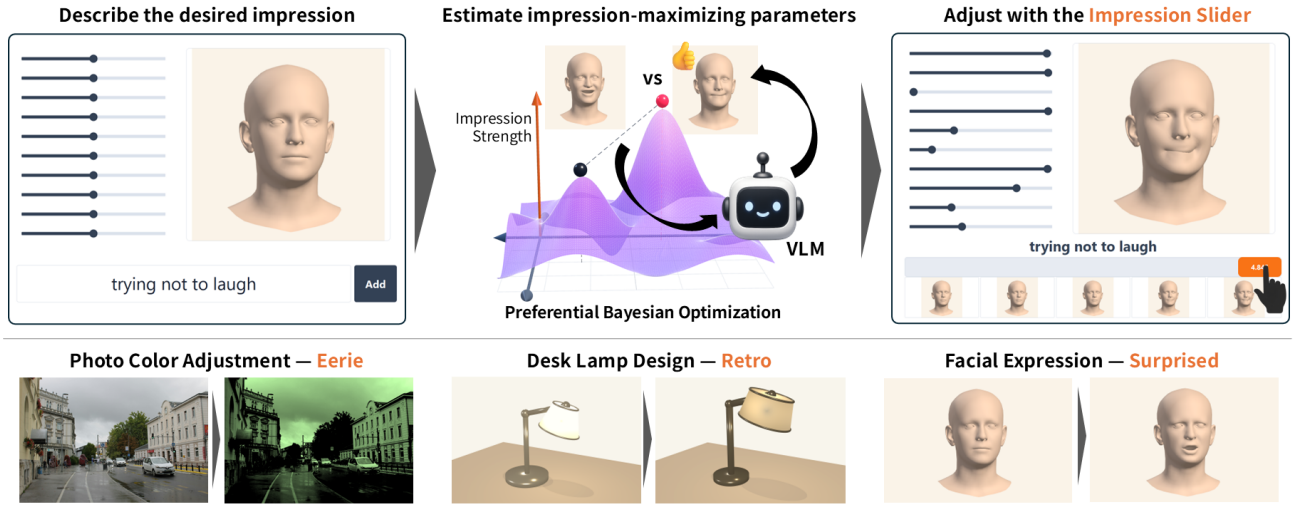


Figure 1: Overview of Impression Sliders. The user describes a desired impression (left); preferential Bayesian optimization with a VLM evaluator estimates a design that strongly conveys that impression (center); the result is presented as an Impression Slider that moves the current design toward the estimate (right). Bottom: applications to other domains.

Abstract

In parametric visual design, users often explore a parameter space, guided by a desired impression to convey to viewers—for example, seeking to make a photo look more "cinematic." However, manipulating the parameters in a direction that strengthens the intended impression is particularly difficult when the user does not understand the parameter space well. We present *Impression Sliders*, a system that turns a natural-language description of a desired impression into a slider for strengthening that impression. With a vision-language model (VLM) as the evaluator, the system uses preferential Bayesian optimization to estimate, in roughly 45 seconds, a design that strongly conveys the described impression. The resulting slider then lets users move the current design toward this estimate. We implemented the system in three domains, and a study with 19 raters indicated that the estimated designs conveyed the described impressions to viewers. An exploratory user study suggested that Impression Sliders help users adjust designs toward target impressions, particularly when the parameter space is difficult to interpret.

CCS Concepts

• **Human-centered computing** → **Interaction techniques**.

Keywords

impression-driven design, parametric design, preferential Bayesian optimization, vision-language models, design tools

ACM Reference Format:
Taiyu Homma and Homei Miyashita. 2026. Impression Sliders: On-Demand Sliders from Impression Descriptions for Parametric Design. In *The 39th Annual ACM Symposium on User Interface Software and Technology (UIST Adjunct '26), November 02–05, 2026, Detroit, MI, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3830397.3841831

1 Introduction

In parametric visual design (e.g., color grading of photographs or shape adjustment of 3D models), users often explore a parameter space, guided by a desired impression to convey to viewers. For example, a user grading a photograph may adjust the parameters, seeking to make it look more "cinematic." However, design tools typically expose individual parameters, whereas changing an impression often requires coordinated changes across multiple

This work is licensed under a Creative Commons Attribution 4.0 International License.
UIST Adjunct '26, Detroit, MI, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2855-6/26/11
https://doi.org/10.1145/3830397.3841831

parameters. Users must therefore determine for themselves how to change the parameters so that the design conveys the intended impression more strongly, which becomes particularly difficult when the parameters are numerous, their effects are hard to predict [9], or the user is unfamiliar with the parameter space.

Prior work has supported impression-guided exploration by estimating impression–parameter correspondences from human evaluation. However, supporting exploration guided by an impression described during design work remains difficult. AttribIt [3] learns adjective–shape correspondences from crowdsourced comparisons, but is limited to adjectives learned in advance. Koyama et al. [5] estimate a goodness function for an arbitrary user-described criterion, but require crowdsourced evaluation that takes tens of minutes. Design Adjectives [13] learns an individual user's intent from user-scored examples, but requires multiple examples. CLIP [12] has been used to guide generative models toward described concepts [11], but evaluates each design independently, making it difficult to assess impressions whose meaning depends on comparison with other candidates, such as "compact."

We present *Impression Sliders*, a system that turns an impression described while designing into an on-demand control for adjusting a parametric design. Using a vision-language model (VLM) as the evaluator, preferential Bayesian optimization (PBO) estimates, in roughly 45 seconds, a design that strongly conveys the impression to viewers. The estimate is presented as an *Impression Slider*, which lets users move the current design toward this estimate. Rather than adjusting parameters individually, users explore designs by manipulating Impression Sliders, creating new ones as needed.

2 System

Given a natural-language target impression such as "retro," the system models a latent utility $f(x)$—how strongly the VLM judges a design x to convey the impression—with a Gaussian process and estimates its maximizer via PBO (Figure 1, center). At each iteration, candidates selected by the qEUBO acquisition function [1] are rendered and compared pairwise by a VLM (GPT-5.4 mini [10]); prior work reports task-dependent agreement between GPT-4V evaluations and human judgments [14]. The comparison outcomes update the posterior via PairwiseGP [2, 4]. Our implementation, built on BoTorch [2], initializes with all-pairs comparisons among 10 Sobol samples, then runs 8 iterations of qEUBO ($q = 7$), with the VLM comparing all 21 pairs among the candidates in parallel (randomized A/B order); the iteration count was chosen based on a preliminary simulation to balance waiting time against estimation quality. Slider generation averaged 43.8 s (photo) and 48.3 s (face) in the exploratory user study.

The estimated design defines the positive end of an *Impression Slider* (Figure 1, right). Dragging toward the positive end linearly interpolates from the current parameters toward the estimated design; dragging toward the negative end interpolates back toward the parameters at the slider's creation (the optimization does not provide a reliable weakening direction). On release, the handle recenters while the parameters are retained, so subsequent drags start from the current design.

Table 1: Mean impression-strength ratings on a 7-point scale ($N = 19$), with per-description Friedman and Bonferroni-corrected Wilcoxon tests: [a] vs. no edit; [b] vs. CLIP ($p < .05$).

Domain	Description	No edit	CLIP	Ours
Photo	eerie	2.16	4.37[a]	**5.12**[a]
Photo	warm	2.21	2.96	**5.54**[a,b]
Face	trying not to laugh	1.58	3.39[a]	**4.82**[a,b]
Face	surprised	1.47	3.88[a]	**4.19**[a]
Lamp	retro	4.68	4.11	**5.67**[b]
Lamp	compact	3.89	1.82[a]	**5.96**[a,b]

We implemented the system in three domains (Figure 1): photo adjustment (seven parameters), facial expressions (the first 10 expression PCA coefficients of the FLAME head model [8]), and desk lamp design (shape, shade, and material parameters).

3 Evaluation

Human evaluation of estimated designs. For six impression descriptions (two per domain), we compared designs obtained under three conditions: (1) no edit, (2) Bayesian optimization of the CLIP [12] score for the description (qEI, $q = 7$, 12 iterations), and (3) our method. Descriptions were presented in Japanese (reported here in English translation); CLIP received the English translations, matching its English-centric training data. Nineteen raters (12 male, 7 female; ages 17–60, median 20) rated, in randomized order, how strongly they felt the described impression from each design on a 7-point scale. CLIP and our method were each run three times per description; for each rater, ratings were averaged over the three runs. Our method achieved the highest mean ratings for all six descriptions (Table 1). For "compact," the margin over CLIP was the largest, consistent with relative judgments being suited to comparative descriptions.

Exploratory user study. In the photo adjustment and facial expression domains, we conducted a within-subjects comparison between a baseline condition with conventional parameter sliders only and an Impression Sliders condition that added them ($N = 6$, ages 21–23, order counterbalanced). No participant routinely used photo adjustment tools, and none had edited facial expressions before. The target descriptions were the same across conditions: horror-film-like eeriness for photos and trying not to laugh for faces. We recorded interaction logs, 7-point questionnaire responses, and semi-structured interviews. All participants created Impression Sliders (1–4 per task, 28 in total), often using descriptions beyond the task statement, such as "like Resident Evil."

Impression Sliders appeared to give participants a direction toward the target impressions. Ratings for "I knew in which direction to move to approach the goal" were higher with Impression Sliders in both domains (photo: 5.00→6.17; face: 4.17→5.50). The interviews were consistent with this pattern: participants reported that they could not imagine how to approach the target impression with parameter sliders alone, whereas Impression Sliders gave them an initial direction. Participants also discovered designs they said they

would not have thought of using conventional sliders alone, such as asymmetric mouth corners for the trying-not-to-laugh expression.

Satisfaction with the final designs, however, differed by domain. Shown the two designs side by side, all six participants preferred the design from the Impression Sliders condition for facial expressions, whereas only three did so for photo adjustment. This may be because the photo parameters in our study were more interpretable than the facial-expression parameters, allowing participants to reach satisfying designs on their own. Impression Sliders may therefore be especially helpful when users have difficulty interpreting the parameter space.

4 Limitations and Future Work

What conveys an impression depends on viewers' cultural backgrounds, yet a single VLM currently makes this judgment, so the estimates may not reflect the intended audience. To keep latency low, PBO is limited to 8 iterations, which can leave the estimate under-optimized; more sample-efficient methods [6, 7] could improve estimates under the same budget. Finally, the impression's perceived strength may not change monotonically along the slider's path. Future systems could use additional comparisons to construct perceptually meaningful paths or support compound operations, such as strengthening one impression while preserving another.

References

[1] Raul Astudillo, Zhiyuan Jerry Lin, Eytan Bakshy, and Peter Frazier. 2023. qEUBO: A Decision-Theoretic Acquisition Function for Preferential Bayesian Optimization. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 206)*, Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (Eds.). PMLR, 1093–1114. https://proceedings.mlr.press/v206/astudillo23a.html

[2] Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. 2020. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 21524–21538. https://proceedings.neurips.cc/paper_files/paper/2020/file/f5b1b89d98b7286673128a5fb112cb9a-Paper.pdf

[3] Siddhartha Chaudhuri, Evangelos Kalogerakis, Stephen Giguere, and Thomas Funkhouser. 2013. Attribit: content creation with semantic attributes. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) *(UIST '13)*. Association for Computing Machinery, New York, NY, USA, 193–202. doi:10.1145/2501988.2502008

[4] Wei Chu and Zoubin Ghahramani. 2005. Preference learning with Gaussian processes. In *Proceedings of the 22nd International Conference on Machine Learning* (Bonn, Germany) *(ICML '05)*. Association for Computing Machinery, New York, NY, USA, 137–144. doi:10.1145/1102351.1102369

[5] Yuki Koyama, Daisuke Sakamoto, and Takeo Igarashi. 2014. Crowd-powered parameter analysis for visual design exploration. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) *(UIST '14)*. Association for Computing Machinery, New York, NY, USA, 65–74. doi:10.1145/2642918.2647386

[6] Yuki Koyama, Issei Sato, and Masataka Goto. 2020. Sequential gallery for interactive visual design optimization. *ACM Trans. Graph.* 39, 4, Article 88 (Aug. 2020), 12 pages. doi:10.1145/3386569.3392444

[7] 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, Article 48 (July 2017), 11 pages. doi:10.1145/3072959.3073598

[8] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Trans. Graph.* 36, 6, Article 194 (Nov. 2017), 17 pages. doi:10.1145/3130800.3130813

[9] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. 1997. Design galleries: a general approach to setting parameters for computer graphics and animation. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., USA, 389–400. doi:10.1145/258734.258887

[10] OpenAI. 2026. GPT-5.4 mini Model. `https://developers.openai.com/api/docs/models/gpt-5.4-mini`. Accessed: 2026-07-05.

[11] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. 2021. StyleCLIP: Text-Driven Manipulation of StyleGAN Imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 2085–2094.

[12] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 8748–8763. https://proceedings.mlr.press/v139/radford21a.html

[13] Evan Shimizu, Matthew Fisher, Sylvain Paris, James McCann, and Kayvon Fatahalian. 2020. Design Adjectives: A Framework for Interactive Model-Guided Exploration of Parameterized Design Spaces. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '20)*. Association for Computing Machinery, New York, NY, USA, 261–278. doi:10.1145/3379337.3415866

[14] Xinlu Zhang, Yujie Lu, Weizhi Wang, An Yan, Jun Yan, Lianke Qin, Heng Wang, Xifeng Yan, William Yang Wang, and Linda Ruth Petzold. 2023. GPT-4V(ision) as a Generalist Evaluator for Vision-Language Tasks. arXiv:2311.01361 [cs.CV] https://arxiv.org/abs/2311.01361