
Machine learning based co-creative design framework

Brian Quanz, Wei Sun, Ajay Deshpande, Dhruv Shah, Jae-eun Park
IBM Research
{blquanz, sunw, ajayd}@us.ibm.com

Abstract

We propose a flexible, co-creative framework bringing together multiple machine learning techniques to assist human users to efficiently produce effective creative designs. We demonstrate its potential with a perfume bottle design case study, including human evaluation and quantitative and qualitative analyses.

Computers can assist creative processes in many ways, from facilitating idea representation and manipulation to directly collaborating through evaluating, generating, and refining ideas [1]. Directly collaborating systems are referred to as co-creative or mixed-initiative co-creative (MI-CC) systems, as both the human and computer proactively and creatively contribute to a solution (*e.g.*, a design) [2, 3, 4], with a goal of inspiring and facilitating lateral thinking, to surpass individual creative outcomes [3]. Most prior MI-CC work focuses on specific, singular applications (*e.g.*, largely video game level design) and single creative-assistant elements (*e.g.*, only suggesting new candidate designs) [3, 5, 6, 7, 8, 9, 10, 11, 12], which cannot address the diversity of user needs [1] and may limit potential and adaption for wider use. Additionally most past work relies on heuristics and evolutionary or random constrained search [3, 13, 14, 15, 16]. Only recently has machine learning (ML) been incorporated in MI-CC, in limited uses of either learning user preference [17] or learning to generate content (designs) [10]; however, we argue ML has much greater potential to assist co-creative design.

To address these shortcomings, we propose a general co-creative framework, showing how multiple ML methods can be incorporated to facilitate creativity in a co-creative system. We illustrate the framework components with an AI assistant tool for package design, and a case study of perfume bottle design (leveraging a created dataset of ~24K bottle images). Currently there is a lack of AI tools for package design, despite the huge impact package design has on product sales [18, 19, 20, 21]. Although our tool uses images as designs, the ML models are generally applicable to other domains. Further, under the idea of diagrammatic reasoning [22], *i.e.*, reasoning via visual representations, creativity systems that operate on images can potentially be applied to non-visual domains [3].

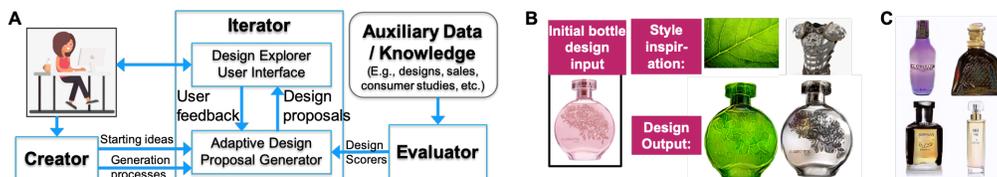


Figure 1: (A) Framework overview (B-C) Examples of design modification and generation, resp., in case study

Our framework (Figure 1A) consists of 3 key components, leveraging ML with collected design data. **Creator** comprises a set of ML tools for generating designs and variants. Generative modeling, *e.g.*, generative adversarial networks (GANs) [23], is used to capture the constraints or manifold of the design space, and generate realistic, novel designs for exploration / design proposal, with further constraint through filtering or score-optimization. Our tool includes this - Figure 1C shows random, novel designs generated by a progressive GAN [24] trained on our data. Additionally, conditional versions [25] can be used for controlled generation. Generation can be conditioned on category or user-selected design features, or parts of the design can be fixed, using image-completion style GANs

[26]. Sketch-to-image (e.g., [27]) can be used to both allow rapid prototyping and refinement with a designer-provided sketch, or to fix the design structure (e.g., using the edge map) and fill in different color palettes (part of our tool). Similarly design structure can be fixed while allowing modification of style (e.g. color and texture) through style transfer [28, 29, 30] - implemented in our tool (Figure 1B). Furthermore, creativity can be infused in generative models to expand the search space [31, 32, 33].

Evaluator leverages data to predict design modalities like novelty (e.g., [34, 35, 36]), cost, and consumer appeal, leveraging modern, highly-accurate predictive modeling [37] and transfer learning (e.g., [38, 39]) to prime modeling off the existing design set. Each evaluator provides one or more values or *scores* that characterize a design, and it is up to the user to decide its influence. In our case study we only include aesthetic modalities of shape and color. **Iterator** works with a human user to evolve design concepts and explore the design space, iteratively proposing candidate designs (including using Creator components) and learning user preference based on which designs are selected (and adapting future proposals) - leveraging again transfer learning to speed up the learning process, and also reinforcement learning [40] techniques. Evaluator modality scores feed into the user preference model as well and also help guide the user’s feedback to the Iterator.

We implemented a version of the framework as a creative assistant tool for perfume bottle design. We used a simplified Iterator interface that displays a 3x6 grid of 18 design proposals each round. Users select designs they like, clicking “SUBMIT” when finished, then the Iterator updates its model of learned user preference and produces a new set of design suggestions. We tested different strategies for design proposal, balancing between exploring the design space to inspire the user and selecting results in-line with the current, learned design goal. We evaluated the Iterator quantitatively first by plugging in an artificial user with fixed, deterministic preference (e.g., for thin bottles as measured by aspect ratio). We measured the accuracy (in terms of Area Under the Curve - AUC) of the learned preference model on a held-out test set as well as number of designs selected per round, for the different proposal methods. Results for the “thin” concept are shown in Figure 2C. We found the system was able to efficiently learn these simple fixed concepts, with a combination of proposal strategies providing efficient learning while still enabling exploration. A purely random proposal strategy (“ONLY-RAND”), however, led to a much lower rate of user-selection, which could frustrate a human user, and slightly lower learning rate.

We also performed a user study with 13 non-expert participants, assigning each a random bottle design task defined by a perfume description (such as “the smell is sweet, fruity, girly, and flirty”) - as is common in real design specifications. We demoed the design-assistive components of the tool for each participant, and had each try the Iterator for his/her task (with a combination proposal strategy). We recorded mean proposal set accuracy (batch AUC) of the learned preference model and number of proposal designs “liked” per round. We found even for this more realistic, complex case the preference was able to be learned fairly accurately and the tool quickly increased the number of proposal the users liked (Figure 2B). We also asked 5 questions about their experiences. On a scale of 1 to 5 (5 being the most): (Q1) “How useful do you think the whole suite of tools would be to help in your design process for coming up with a new, creative design?”; (Q2) “How much did it help you discover candidate designs related to the task?”; and (Q3) “How much did it help you explore the wide space of designs and stimulate your creativity?”; (free response) (Q4): “What did you like most about the tool?” and (Q5): “What would you most want to improve about the tool?”.

Summary statistics of their responses to Q1-3 are given in Figure 2A. We received largely positive responses, users seeming excited about the capabilities. Users felt the sketch-to-design and style transfer components could be very helpful and time-saving. For the Iterator, users liked its ease-of-use and felt the exploration aspect was useful, providing unexpected ideas at times. However, users felt there were sometimes too many similar designs presented - a result of modeling the proposals per-image (set-based metrics like diversity will be added in future work). Having more capabilities to control the exploration was a key desire - such as the ability to fix the color theme, or take aspects of designs they liked, like shape or color only. We plan to add such capabilities as part of future work, by incorporating more Creator components mentioned, such as conditional models, in our tool.

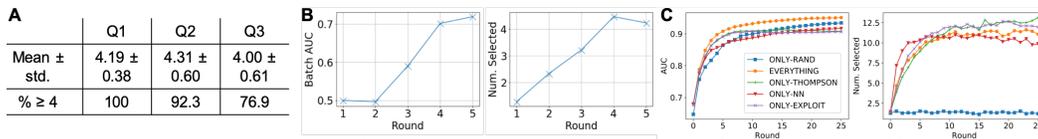


Figure 2: (A) User response stats (B-C) AUC and num. selected per round for human and artificial users, resp.

Ethical Implications

One might think a tool capable of generating creative designs could take work away from designers. However, the key point here is that it is a *co-creative* assistant - it does not replace designers, but instead works with designers, making their jobs easier and more enjoyable, and stimulating their creativity to new heights. As such it will make designers more efficient, spending less time on trivial tasks and spending more on the creative ones, making it easier to consider other data (such as cost, sales impact, or competitor data) in their designs, and helping them handle more design tasks without fatiguing their creativity. Another point is it facilitates cross-company participation in design, helping others in a company, like marketers and production teams, have a more productive collaboration with designers and possibly contribute in the design process.

Since the tool is ML-based and requires existing designs to train the model, the use of the design data could be considered an issue, if such data from one company is used to train models used in a tool offered to competitors. In some domains, many designs are public (such as publicly sold perfumes), but each company may have even more non-public designs with various data on them, such as consumer studies. However, one ethical way of providing such a beneficial tool to all designers is giving companies the choice to opt-in to a shared-data model. I.e., they could choose to keep their designs internal - using only their own and public designs for model training, and not leveraging competitor designs as well. Or they could instead choose to share some set of their internal designs to aid in general model training, and also benefit from designs supplied from other competitors.

References

- [1] T. Lubart, "How can computers be partners in the creative process: classification and commentary on the special issue," *International Journal of Human-Computer Studies*, vol. 63, no. 4-5, pp. 365–369, 2005.
- [2] D. G. Novick and S. Sutton, "What is mixed-initiative interaction," in *In Proceedings of the AAAI Spring Symposium on Computational Models for Mixed Initiative Interaction*. Citeseer, 1997.
- [3] A. Liapis, G. N. Yannakakis, C. Alexopoulos, and P. Lopes, "Can computers foster human users' creativity? theory and praxis of mixed-initiative co-creativity," *Digital Culture & Education*, vol. 8, no. 2, pp. 136–153, 2016.
- [4] G. N. Yannakakis, A. Liapis, and C. Alexopoulos, "Mixed-initiative cocreativity," in *Proceedings of the 9th Conference on the Foundations of Digital Games*. Citeseer, 2014.
- [5] A. Liapis, G. N. Yannakakis, and J. Togelius, "Sentient sketchbook: Computer-assisted game level authoring," in *Proceedings of the 8th International Conference on Foundations of Digital Games*. Society for the Advancement of Digital Games, 2013, pp. 213–220.
- [6] A. Liapis, G. Smith, and N. Shaker, "Mixed-initiative content creation," in *Procedural content generation in games*. Springer, 2016, pp. 195–214.
- [7] S. Deterding, J. Hook, R. Fiebrink, M. Gillies, J. Gow, M. Akten, G. Smith, A. Liapis, and K. Compton, "Mixed-initiative creative interfaces," in *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2017, pp. 628–635.
- [8] M. Guzdial and M. Riedl, "Game level generation from gameplay videos," in *Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2016.
- [9] M. Guzdial, N. Liao, J. Chen, S.-Y. Chen, S. Shah, V. Shah, J. Reno, G. Smith, and M. O. Riedl, "Friend, collaborator, student, manager: How design of an ai-driven game level editor affects creators," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 2019, p. 624.
- [10] M. Guzdial, N. Liao, and M. Riedl, "Co-creative level design via machine learning," *arXiv preprint arXiv:1809.09420*, 2018.
- [11] K. Gilon, J. Chan, F. Y. Ng, H. Liifshitz-Assaf, A. Kittur, and D. Shahaf, "Analogy mining for specific design needs," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 2018, p. 121.
- [12] P. Siangliulue, J. Chan, S. P. Dow, and K. Z. Gajos, "Ideahound: improving large-scale collaborative ideation with crowd-powered real-time semantic modeling," in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 2016, pp. 609–624.
- [13] A. Liapis, G. N. Yannakakis, and J. Togelius, "Enhancements to constrained novelty search: Two-population novelty search for generating game content," in *Proceedings of the 15th annual conference on Genetic and evolutionary computation*. ACM, 2013, pp. 343–350.

- [14] V. Singh and N. Gu, "Towards an integrated generative design framework," *Design studies*, vol. 33, no. 2, pp. 185–207, 2012.
- [15] H.-S. Kim and S.-B. Cho, "Application of interactive genetic algorithm to fashion design," *Engineering Applications of Artificial Intelligence*, vol. 13, pp. 635–644, 2000.
- [16] P. Von Buelow, "Suitability of genetic based exploration in the creative design process," *Digital Creativity*, vol. 19, pp. 51–61, 03 2008.
- [17] A. Liapis, G. N. Yannakakis, and J. Togelius, "Designer modeling for personalized game content creation tools," in *9th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE 2013 Workshop*. AI Access Foundation, 2013, pp. 11–16.
- [18] U. R. Orth and K. Malkewitz, "Holistic package design and consumer brand impressions," *Journal of Marketing*, vol. 72, no. 3, pp. 64–81, 2008. [Online]. Available: <https://doi.org/10.1509/jmkg.72.3.64>
- [19] PackagingWorld, "New structural design increases gatorade sales up to 50% <https://www.packworld.com/article/new-structural-design-increases-gatorade-sales-50/>," March 2006, [Online; posted 31-March-2006].
- [20] Nielsen, "Nielsen cpg innovation report: Packaging sells white paper (<http://www.nielsen.com/content/dam/corporate/us/en/reports-downloads/2016-reports/nielsen-packaging-sells-white-paper.pdf>)," December 2015.
- [21] M. Reimann, J. Zaichkowsky, C. Neuhaus, T. Bender, and B. Weber, "Aesthetic package design: A behavioral, neural, and psychological investigation," *Journal of Consumer Psychology*, vol. 20, no. 4, pp. 431 – 441, 2010, aesthetics. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S105774081000063X>
- [22] P. C.-H. Cheng, R. K. Lowe, and M. Scaife, "Cognitive science approaches to understanding diagrammatic representations," in *Thinking with diagrams*. Springer, 2001, pp. 79–94.
- [23] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [24] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of gans for improved quality, stability, and variation," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [25] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [26] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Generative image inpainting with contextual attention," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [27] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017, pp. 5967–5976.
- [28] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 2414–2423.
- [29] Y. Jing, Y. Yang, Z. Feng, J. Ye, Y. Yu, and M. Song, "Neural style transfer: A review," *IEEE transactions on visualization and computer graphics*, 2019.
- [30] X. Huang and S. Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [31] A. Elgammal, B. Liu, M. Elhoseiny, and M. Mazzone, "Can: Creative adversarial networks, generating" art" by learning about styles and deviating from style norms," *arXiv preprint arXiv:1706.07068*, 2017.
- [32] O. Sbai, M. Elhoseiny, A. Bordes, Y. LeCun, and C. Couprie, "Design: Design inspiration from generative networks," in *The European Conference on Computer Vision (ECCV) Workshops*, September 2018.
- [33] P. Das, B. Quanz, P. Chen, and J. Ahn, "Toward A neuro-inspired creative decoder," *arXiv preprint arXiv:1902.02399*, 2019. [Online]. Available: <http://arxiv.org/abs/1902.02399>
- [34] X. Ding, Y. Li, A. Belatreche, and L. P. Maguire, "An experimental evaluation of novelty detection methods," *Neurocomputing*, vol. 135, pp. 313–327, 2014.
- [35] H.-g. Wang, X. Li, and T. Zhang, "Generative adversarial network based novelty detection using minimized reconstruction error," *Frontiers of Information Technology & Electronic Engineering*, vol. 19, no. 1, pp. 116–125, 2018.
- [36] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," *arXiv preprint arXiv:1610.02136*, 2016.
- [37] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017, pp. 1800–1807.

- [38] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, “Decaf: A deep convolutional activation feature for generic visual recognition,” in *International conference on machine learning (ICML)*, 2014, pp. 647–655.
- [39] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, “Cnn features off-the-shelf: an astounding baseline for recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR) workshops*, 2014, pp. 806–813.
- [40] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction (<http://incompleteideas.net/book/the-book-2nd.html>)*, 2nd ed. MIT press, 2018, [Online; accessed 5-Sept-2018].
- [41] R. Ahmed, V. Parmar, and M. Ahmed Amin, “Impact of product packaging on consumer’s buying behavior,” vol. 122, pp. 125–134, 04 2014.
- [42] G. H. Joblove and D. Greenberg, “Color spaces for computer graphics,” in *ACM SIGGRAPH Computer Graphics*, vol. 12, no. 3. ACM, 1978, pp. 20–25.
- [43] V. Losing, B. Hammer, and H. Wersing, “Incremental on-line learning: A review and comparison of state of the art algorithms,” *Neurocomputing*, vol. 275, pp. 1261–1274, 2018.
- [44] C. Robert and G. Casella, *Monte Carlo statistical methods*. Springer Science & Business Media, 2013, ch. 2.3 - Accept-Reject Methods.
- [45] C. Riquelme, G. Tucker, and J. Snoek, “Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [46] ESKO, “Esko artioscad software (<https://www.esko.com/en/products/artioscad/>),” 2019.
- [47] —, “Esko deskpack software (<https://www.esko.com/en/products/deskpack/>),” 2019.
- [48] Adobe, “Adobe illustrator software (<https://www.adobe.com/products/illustrator.html>),” 2019.
- [49] —, “Adobe photoshop software (<https://www.adobe.com/products/photoshop.html>),” 2019.
- [50] CreativeEdge, “Creative edge ic3d software (<http://www.creativeedgesoftware.com/ic3d/>),” 2019.
- [51] 99designs, “Packaging design platform (<https://99designs.com/product-packaging-design/>),” 2019.
- [52] fiverr, “Packaging design platform (<https://www.fiverr.com/categories/graphics-design/product-packaging-design/>),” 2019.
- [53] Contributors, “Creativeai.net (<http://www.creativeai.net/>),” 2019.
- [54] P. Hutchings and J. McCormack, “Using autonomous agents to improvise music compositions in real-time,” in *Computational Intelligence in Music, Sound, Art and Design*, J. Correia, V. Ciesielski, and A. Liapis, Eds. Cham: Springer International Publishing, 2017, pp. 114–127.
- [55] D. Ha and D. Eck, “A neural representation of sketch drawings,” in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=Hy6GHpkCW>
- [56] A. Kazakçi, M. Cherti, and B. Kégl, “Digits that are not: Generating new types through deep neural nets,” *CoRR*, vol. abs/1606.04345, 2016. [Online]. Available: <http://arxiv.org/abs/1606.04345>
- [57] A. Dalmia, S. Joshi, R. Singh, and V. Raykar, “Styling with Attention to Details,” *ArXiv e-prints*, Jul. 2018.
- [58] T. Nakamura and R. Goto, “Outfit Generation and Style Extraction via Bidirectional LSTM and Autoencoder,” *ArXiv e-prints*, Jun. 2018.
- [59] W. Zhou, H. Li, and Q. Tian, “Recent advance in content-based image retrieval: A literature survey,” *arXiv preprint arXiv:1706.06064*, 2017.
- [60] Google, “Google image search (<https://images.google.com/>),” 2019.
- [61] M. Wang, B. Ni, X.-S. Hua, and T.-S. Chua, “Assistive tagging: A survey of multimedia tagging with human-computer joint exploration,” *ACM Comput. Surv.*, vol. 44, no. 4, pp. 25:1–25:24, Sep. 2012. [Online]. Available: <http://doi.acm.org/10.1145/2333112.2333120>
- [62] N. Rubens, M. Elahi, M. Sugiyama, and D. Kaplan, *Active Learning in Recommender Systems*. Boston, MA: Springer US, 2015, pp. 809–846. [Online]. Available: https://doi.org/10.1007/978-1-4899-7637-6_24

1 Supplementary Material

1.1 Co-creative assistant for packaging design use-case motivation

Packaging (e.g., bag, box, or bottle) design is an integral part of product and brand identity, which influences consumer experience, buying decisions and loyalty [18]. In the age of differentiating consumer experiences, good packaging designs strive to offer unique appeal in terms of aesthetics, utility, creativity, trendiness and brand impression while keeping costs and environmental impacts at bay. There are ample examples that indicate the impact of package design on consumer impressions. For instance, according to [20], the right packaging may boost sales for an average brand by 5.5%. [21] found aesthetic packaging can boost sales for a new brand, and [41] found the preferences for packaging colors are very different between the West and Far East. The consumer industry often revamps packaging for common-use products frequently to inject new life into sales. For instance, Pepsico achieved sales growth of 35% to 50% in some countries for its Gatorade brand with new structural bottle designs [19].

Despite packaging’s critical role, current packaging design processes lack data- and AI-driven tools to facilitate and enhance these processes. Based on internal discussion with multiple consumer industries, in current design work flows for packages, marketers, product developers and designers often work in silos, wasting time in back-and-forth communication and iterative manual design refinements. Design evaluations are performed on multiple real trial prototypes with consumer testing making the entire process excessively expensive and lengthy. There is no data-driven design evaluator that allows users to estimate how a design would perform along different dimensions, such as sales, appeal to customer segments, cost, etc. Furthermore, the design process itself could benefit from assistive tools. Currently a non-designer with some ideas has to rely on a designer to create visualizations and there is no tool that allows the non-designer to make quick design tweaks and visualize her ideas. Even an expert designer may get trapped in past design biases or blind spots, and may miss potential blockbuster new designs.

We developed our ML-based co-creative framework, in part motivated by these challenges, and specifically chose the package design use-case as our target one to demonstrate the framework. The framework, and our realization of it as an AI-assistant tool, assists human designers to rapidly produce more effective and creative package designs without taking creative control away from the designer. Akin to these implicit phases in a typical design flow, our framework consists of three broad sets of tools, namely Creator, Evaluator and Iterator. Creator can take some initial inspiration seeds from a human designer (e.g., color palette, shape, inspiration images, etc.), enables rapid design space exploration by learning the latent design manifolds, and generates new designs. Evaluator enables evaluating generated designs along different dimensions, extracting different modalities of designs in the form of multi-dimensional metrics, e.g., learning relationships between design aspects and outcomes, such as aesthetics based on shape and color or novelty compared to competitors. It is able to ingest or leverage additional data or knowledge, such as consumer studies from marketing, cost constraints from product development divisions, and knowledge about shape learned from large data collections. Finally, Iterator uses an interactive UI to propose candidate designs along with their metrics to a user and gather feedback on the designer’s current preferred choices. Iterator further learns a model of the designer’s preferences for her current design goal and guides the search in the design space, facilitating the next iteration of creation and evaluation. The framework and tool are flexible in that human users with different skill levels and needs can use parts of it for different purposes. For instance, an expert designer may use Evaluator to estimate her design’s appeal and Creator or Iterator to get inspiration for out-of-box designs, while, a marketing professional may use the whole framework to tweak an existing design and visualize the results.

We demonstrate the feasibility and potential of this framework with a use case on designing perfume bottles. We show that our framework enables facilitating and enhancing the design process and grounding it in real data so both designers and non-designers can explore design possibilities and modifications for designs as desired and get quick mock-ups without long delays. Besides package design, we believe that our general framework is also applicable to broader design applications such as logo and website creation, advertisement, and others, as mentioned in the main paper. Even non-visual domains can use the same underlying machine learning models and framework components and structure, and diagrammatic reasoning is also an option.

1.2 Model Implementation for Package Design

In this section we will go over the framework again and specifically how it applies in the context of package design (and the implementation we used for our tool to demonstrate the framework). We reproduce some figures with larger size and provide more details from the main paper.

1.2.1 Overview

Our human-AI interactive design framework consists of three key components as illustrated in Figure 3. Creator collectively comprises a set of tools to enable the user to create design realizations from more primitive components as well as refine or modify designs, and also includes processes for generating designs or design variants. Evaluator consists of scorers that can characterize a design along different dimensions (like shape, color, cost, customer appeal, etc.). The outputs from Creator and Evaluator feed into Iterator - which enables the user to interactively explore the design space, by suggesting relevant candidate designs meeting her preference and current design goals along the multiple modalities, and iteratively improving design suggestions and the learned model of user preference. The end result of the system is to quickly provide the user with design inspirations and rough realizations that can be used to base final designs on.

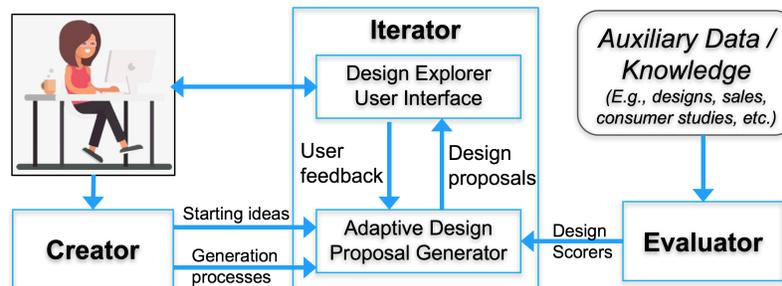


Figure 3: Overview of proposed interactive AI design framework.

1.3 Data

To show results for each component of our framework, and to evaluate the interactive system, we created a dataset of 24,126 perfume bottle product images (see Figure 13 for a sample). We first collected around 35,000 product images from the internet and filtered this to clean, bottle-only images using a classifier we trained on a small set of manually labeled images, leveraging transfer learning with a pre-trained image classification network [38, 39].

We also generated masks of the bottles to separate bottles from background by performing Canny edge detection, followed by Gaussian filter blurring and flood filling the image up to the blurred edge boundaries. This is used to determine characteristics of images like dimensions, and mask results after image operations like style transfer.

1.3.1 Creator

In many domains, a typical design process begins with a design brief which is a document that outlines the deliverables and information of the project including background, marketing message, target audience etc. It provides a designer with inspirations, directions, and constraints for the design. For example, in packaging design, a brief can indicate a primary color for a package, *e.g.*, “The handsome deep blue was chosen because it conveys a sense of tradition and strength, which is right for XYZ” (XYZ is a cologne for men).

Creator consists of several tools to suit different scenarios in this design phase. It can start from some initial inspiration seed (*e.g.*, color palette or shape), explore through the design space rapidly by leveraging learned latent design manifolds, and generate new designs.

Rapid concept exploration with sketches: Many designers start with sketches to capture their initial ideas and spend a significant amount of time to refine, detail, and explore variants of those designs. It is time-consuming but often the only way for a designer to determine if concepts are worth exploring further. One tool in Creator assists the human designer to explore through the design space

rapidly with sketches. The tool takes sketches and/or color palettes as input, and renders realistic designs. While those designs are still semi-finished, they give the designer the look and feel of different concepts, and enable him to focus more on rapid formulation of ideas.

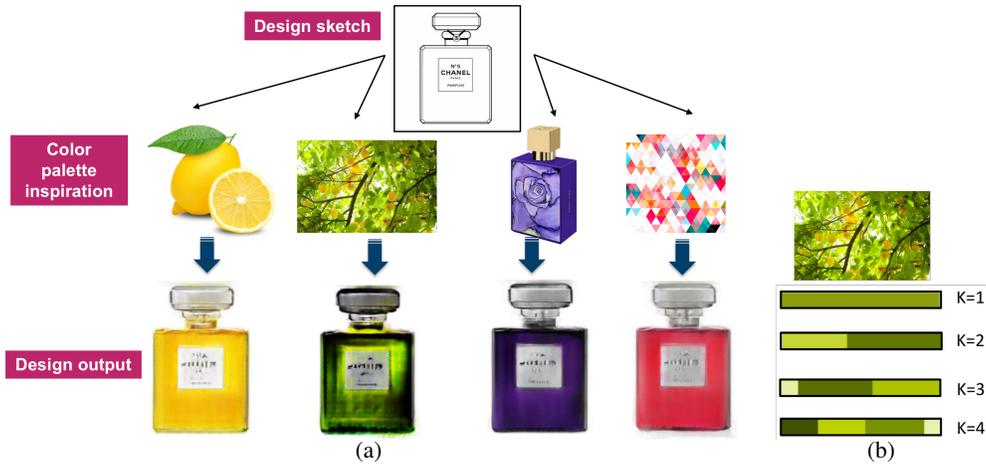


Figure 4: (a) An example of Creator on perfume bottle designs based on inspiration images. (b): Color palettes generated from an inspiration image for different number of clusters. These feed into the conditional GAN to guide the bottle sketch completion

The framework behind this tool is adapted from Pix2Pix [27], which uses a conditional generative adversarial network (cGAN) to learn a mapping from an input image (corresponding to the image edge map) to an output image (corresponding to the complete image). Differentiating our approach, we make a key modification to the model for our application: our tool takes additional images as input from which a color palette is extracted and fed into the generator to influence the color in the generated result (see Fig 4(a) for an example of some results). This additional input is crucial for package design as clients often specify a color palette or provide color inspirations in the design brief. In addition this approach could be used to enable constrained generation for exploration while fixing the structure, but exploring different sample colored palettes - and could also be extended to have multiple color palettes corresponding to different parts of the design.

To obtain the color palette, we first run K-means to cluster the pixel intensities of an RGB image, then create a histogram which depicts the dominant colors represented by each centroid, where the width for each color is proportional to the number of pixels in that cluster. Fig 4(b) shows various color palettes generated from the same image with respect to different cluster sizes. Each centroid represents a dominant color in the inspiration image. Pixels that belong to a given cluster will be more similar in color than pixels belonging to a separate cluster, and the number of dominant colors is determined by the cluster size which is an input parameter of K-means. The color palette of each training bottle image is fed into the conditional GAN model as part of the conditional input, along with the edge map of the image, during training, and when applying the model with a bottle sketch the extracted color palette from inspiration images are used.

Figure 5 shows two more examples to showcase the robustness of this creation tool. Instead of a sketch of a bottle, the “inspirational shapes” are given by an assembly of an apple and a glass cap in the first example, and a lemon slice clipart and a rectangle in the second. These examples demonstrate that even without a formal sketch, the tool is capable of producing realistic-looking designs from quick and rough shape input given by a designer.

Style transfer and flanker design: Borrowing a perfumery concept, a flanker refers to a newly created perfume that shares some attributes (e.g., packaging or notes) of an already existing perfume. These attributes may be the name, packaging or notes of the existing fragrance. For example, Dior’s 1985 fragrance Poison was followed by Tendre Poison (1994), Hypnotic Poison (1998), Pure Poison (2004), Midnight Poison (2007) and Poison Girl (2016), all with roughly the same packaging in different styles. In terms of flankers’ packaging, they tend to share very similar bottle designs, but differ in terms of colors, texture and potentially material. For this task, we utilize the approach of neural style transfer [28] to copy the style from a style image and apply it to a content image representing the design we want to modify the style of. When it is applied to flanker package design,



Figure 5: Examples of Creator with abstract shape input.



Figure 6: Flankers for Dior Poison. Image source: <https://cleopatrasboudoir.blogspot.com/2013/05/poison-by-christian-dior-c1985.html> (We can include this image in the supplementary material)

the content image can be the an existing design (e.g., perfume bottle), while the style could be an inspiration image from the design brief. Fig 7 depicts examples of style transfer with our data. For example, the last column is an example of a flanker design where we transfer the style of shiny armor to a perfume bottle, creating a “masculine” version of the existing bottle. Style transfer can also be used to generate design variants to help explore the design space and provide creative design proposals to the human user in the co-creative system. The results of style transfer can even be varied based on changing the settings (as illustrated in Figure 8), adding another dimension of generated design variation.

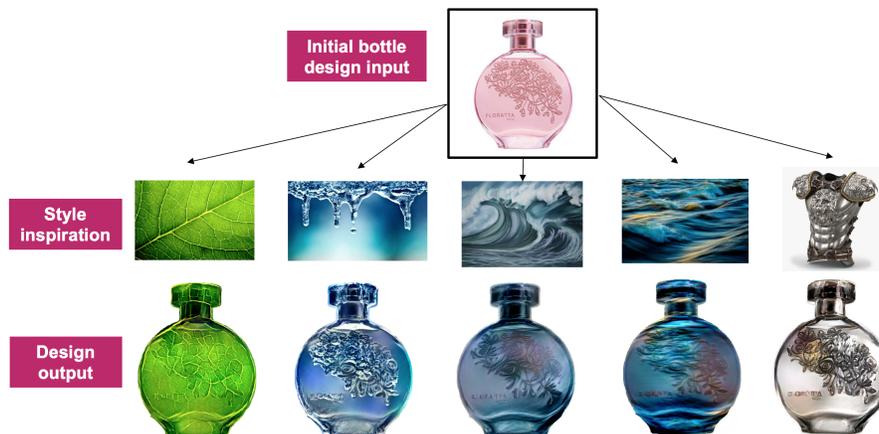


Figure 7: Examples of flanker bottle designs with style transfer.

Design manifold exploration with generative modeling: The space of design specifications, such as images depicting a design, is huge and difficult to explore directly. However, designs exist on some lower dimension manifold encompassing the inherent design characteristics and constraints, such as symmetry, specific cap and body shapes and locations, etc. By learning to model and generate designs in this low-dimension space we can facilitate design space exploration, e.g., to find similar designs, interpolate between designs, and even generate completely new designs. Our system uses



Figure 8: Examples of flanker bottle designs with style transfer for different packaging type: wallets. Each result is generated with different settings for the style transfer, generating how style transfer can contribute to expanding the search space of designs, even without vaying style images.

generative neural network models to accomplish this. For example, we trained a progressively grown GAN (PGAN) [24] on our collection of bottle images. An example of 5 random, non-cherry-picked designs generated by this trained network are shown in Figure 9. We found it is able to generate fairly realistic, novel designs, quite dissimilar from “neighboring” training-data designs in the embedding spaces. Along with the other Creator components, and conditional GANs, this can be used as a key component for design space exploration described in Section 1.3.3 - for generating new and nearby designs to help flesh out and fully explore the design space.



Figure 9: Examples of random, new designs generated.

1.3.2 Evaluator

Evaluator consists of a number of scorer systems each evaluating the design along a given dimension, or *modality*, trained using relevant data for that modality. For example, a scorer may be a predictor estimating the cost of a design based on design specifications and associated manufacturing cost details, or predicting the appeal to different demographics based on consumer studies and sales data, or quantifying novelty from an existing database of designs. Each scorer takes a design specification (e.g., an image representing the bottle design), and maps it to one or more values representing the design along its modality, e.g., a single estimated cost value, or a customer preference profile vector, or a shape representation vector, etc.

For our proof of concept we focused on two key, universal design modalities - shape and color. To represent the shape modality, we use high-layer outputs of a pre-trained image classification neural net trained on a large image data set (ImageNet) as these have been shown to provide general image features [38, 39], and style transfer work suggests high-level layer outputs capture image structure (i.e., shape) well [28]. Specifically we use the pre-trained Xception network’s global pool output layer [37] - a 2048-D vector.

To represent the color modality, we use a simple a 4-D descriptor, which we found to be effective at capturing dominant color themes. We transform each pixel to the HSV space [42] and convert the radial hue value to its 2-D coordinate on the unit circle. We then take the average of these pixel vectors across the image, resulting in a simple 4-D descriptor.

1.3.3 Iterator

The Iterator component enables interactive exploration of the design space by iteratively proposing a set of candidate designs to a user, collecting feedback, and updating a learned model of the user’s current preference to refine the proposals for the next round. It leverages a set of different proposal methods, trading off between exploring in different ways and exploiting current learned knowledge of the user’s preference. Exploration can help inspire and lead the user to creative designs, while exploitation is necessary to rapidly propose relevant results. The overall process is given in Algorithm 1.

```
input :# of candidates per proposal method
output :Set of selected candidate designs
Generate random set of candidate designs,  $C$ ;
while User does not terminate do
    Show user set of candidate designs  $C$ ;
    Retrieve feedback from user ;
    Update user preference models from feedback;
    (Optional) update generator to align with user preference;
    Empty candidate design set  $C$ ;
    for each proposal method do
        | Generate proposals from method and add to  $C$ ;
    end
end
```

Algorithm 1: Interactive Design Exploration

Iterator consists of 5 components. First is an interface to present proposed designs and obtain user feedback. In our proof of concept, the feedback is just a simple “like” or not and the interface is a graphical user interface visually displaying candidate designs (Figure 13). In general the interface can provide more details per design and the feedback can be more nuanced, including ratings along different modalities. The second component is a generator, for generating candidate designs. This can be on-the-fly generation using the Creator (e.g., generative models, style transfer, etc., with random inputs). Alternatively, candidate designs can be generated and collected in advance and simply sampled from (i.e., a pre-built database of designs). The generation can also be biased using the user preference model or design starting point inputs. The third component is the set of Evaluators (see Section 1.3.2) - providing embeddings of the design along different modalities. These encode key design dimensions users may care about, so the user preference models are based off of them, and they can also be incorporated in the UI to inform and guide the users.

The fourth component is a trainable user preference model, i.e., a predictive model of the user’s design preference, in the current design session. For the case of “liking” a design or not, this is simply a probabilistic classification model. This model is updated after each round of proposals and feedback from the user, via incremental learning [43], or simply re-training, which is typically quick enough due to the small sample size. This model can be based on both the full-fledged design representations (e.g., the design images in our package design use case) and the Evaluator modality outputs, or for simplicity it could be restricted to be based off of only the modality outputs (although this could limit its potential accuracy and ability to account for modalities not captured in the given set).

Finally, the fifth component is a set of proposal methods. These involve sampling from or searching over the design space (which can be represented by the Evaluator modality embeddings for simplicity and efficiency), and vary over providing design space exploration, enabling rapid user-preference learning, and exploiting learned user preference. The following proposal methods are used.

RAND: This method randomly samples from the generator inputs/database of designs. We also implemented a variant that samples from the distribution imposed by the user preference model via rejection sampling [44].

EXPLOIT: A larger set of designs (e.g., a few hundred) are sampled via RAND and the top scoring designs according to the user preference model are taken for the proposals.

THOMPSON: Based on a reinforcement learning perspective [40], for each candidate proposal it samples a user preference model from an estimated distribution over preference models and takes

the top-scoring design from a RAND sample. To capture the distribution over models, bootstrap sampling is used (modeled after the approach in [45]). A set of k different models are maintained and each has probability p of being updated with a labeled design after each round - so that each model is incrementally trained on a different, random subsample of data.

NN: This proposal method samples designs in nearby neighborhoods of selected designs in the previous round - in different modality embeddings. This enables quickly seeing similar designs and design variations to ones the user liked, to inspire the user and rapidly expand the set of relevant designs proposed.

1.4 Experiment

We focus on most-thoroughly evaluating the Iterator component of our proposed framework, as it includes the other components and lends itself to quantitative and qualitative study. We implemented and tested the Iterator approach for simple binary selection feedback. We focus on the case of perfume bottle designs as represented by bottle product images (data described in Section 1.3). Screen shots of the interface are shown in Figures 12 and 13. When a user clicks to select a design, it is highlighted in green. After finishing selections in the current round, the user clicks the “SUBMIT” button to submit the selections, and clicks the “END” button when finished exploring designs. We focus on the representative approach of using a pre-generated database of designs, as described in Section 1.3.3, for simplicity.

1.4.1 Setup

We test the ability of the system to aid in discovering and learning a number of design concepts. We fix each design concept in advance then interact with the AI system through the interface to see how quickly it can learn the concept and propose relevant designs. For each experiment run we plug in a program to simulate human operation according to fixed scoring rules for selecting designs.

Design concept tasks: In order to avoid overly subjective results, we use 4 simple design concepts, scored automatically by fixed rules. The first two tasks are **RED** and **BLUE** - which represent the concept of dominant red or blue colored bottles, respectively. These are scored by converting each RGB pixel value to the difference between the red / blue channel and the maximum value of the other channels, and averaging over all pixels. If this score is then greater than an outlier threshold, it is classified as a positive case. Similarly the next two concepts, **FAT** and **THIN**, are also based on a threshold rule - in this case on the aspect ratios (width divided by height) of the bottles, which are determined by extracting the width and height of each bottle using masks derived as explained in Section 1.3. **THIN** corresponds to very tall, skinny bottles, and **FAT** to very short, fat bottles.

Additionally, to better simulate human interaction of sometimes selecting cases that seem close to the target concept - for each concept we set another, smaller threshold that includes roughly 2000 additional bottles. We then assign a probability using a sigmoid function of the bottle’s score normalized between the two thresholds, and randomly label these bottles according to the probabilities in each run. The final percentages of positively labeled cases are shown for each automated task in Table 1 - where “Always” indicates the percentage of bottles that are always labeled positive, according to the first, tighter threshold, and “Average” the average number labeled positive across all runs - including the randomly labeled ones. We also illustrate the results of manually using the tool for a fifth design concept of “circle-shaped bodies”.

Task	RED	BLUE	FAT	THIN
Average	9.2%	2.1%	9.5%	7.5%
Always	5.5%	0.5%	7%	3.6%

Table 1: Percentage of positive-labeled cases per automated task.

System setup and procedure:

Our interface provides 18 proposals to the user each round in a 3x6 grid. We compare 5 different proposal methods. The first 4 correspond to generating proposals using each of the 4 proposal methods of Section 1.3.3: RAND, EXPLOIT, THOMPSON (with $k = 5$ and $p = 0.75$), and NN. The fifth method is a combination of these, labeled “EVERYTHING”, with the number drawn from each method set to: 4, 1, 9, and 4, respectively - and using the rejection sampling approach for RAND. We run the system for 26 rounds for each method, repeated 90 times with different random initialization

and labeling, and report results averaged across the 90 runs. For each run we randomly hold out 2000 designs (with stratified sampling) to use as test cases for measuring test accuracy of the user preference models. For the accuracy metric, we report Area Under the Curve (AUC) as the data is imbalanced.

For the user preference model class, we use a support vector machine classifier per embedding with an RBF kernel, fixing the error weight C to 100 and setting the kernel width σ^2 to half the 10th percentile of all pairwise distances for that embedding. The per-embedding classifiers are then combined via a logistic regression classifier.

1.4.2 Results

For simplicity, because results were similar, we only present the results of the automated tasks “RED” and “THIN” in Figures 10 and 11, which depict test AUC per iteration and number of user-selected designs per iteration, respectively, averaged over the 90 random runs. Generally all methods can achieve high AUC and learn the target concept fairly quickly - but the combination method (EVERYTHING) achieves among the top AUC per round in most cases, standing out in particular for the THIN task which may be the most complex task. We note that for the RED task those methods that most quickly achieve high AUC see subsequent degradation in model accuracy - we suspect this may be due to a resulting early imbalance of positive cases seen in the training data and propose to investigate rebalancing for training in future work.

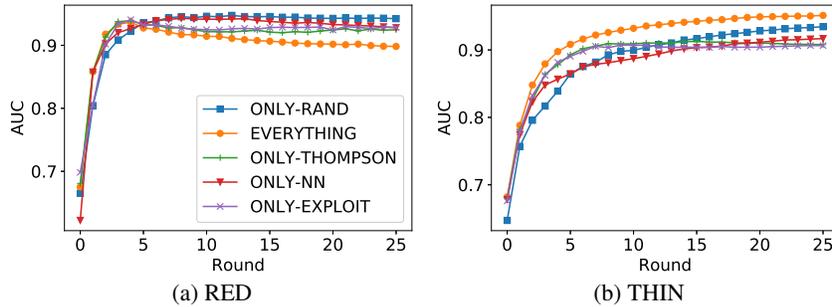


Figure 10: Test AUC per round comparison for two of automated tasks with interactive explorer system.

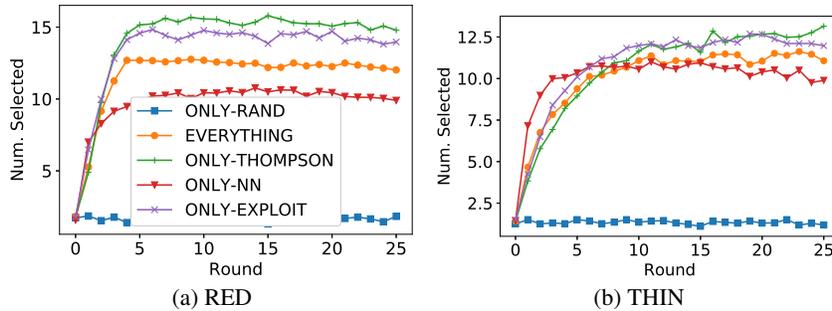


Figure 11: Number selected per round comparison for two of automated tasks with interactive explorer system.

In terms of number of user-selected proposals per iteration, we see that different methods yield the fastest increase and most total selections for different tasks, except for RAND which is never able to generate a significant number of relevant proposals. Most methods quickly provide more and more relevant proposals, while still allowing some exploration. Note also that, unlike reinforcement learning, the goal here is not just to get the most selections, but also to incorporate a significant amount of exploration to try and inspire the user to discover new creative ideas or directions. For this reason we feel if there are not some rejections each round then not enough exploration is included, and the combined approach, EVERYTHING, seems to typically provide a large number of selected designs, without overdoing it.

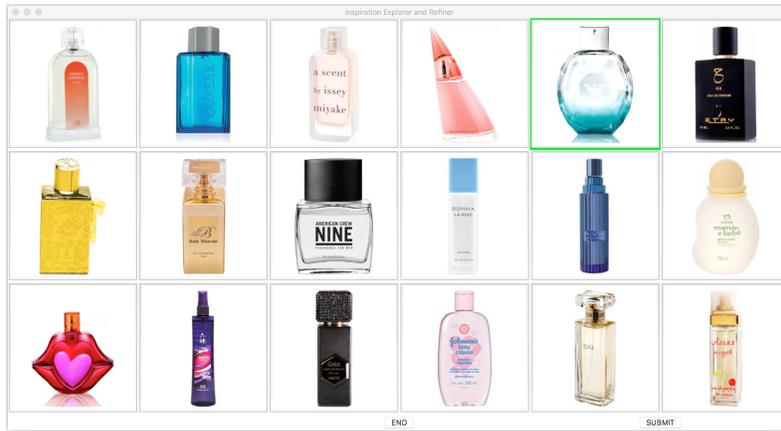


Figure 12: Starting sample of designs

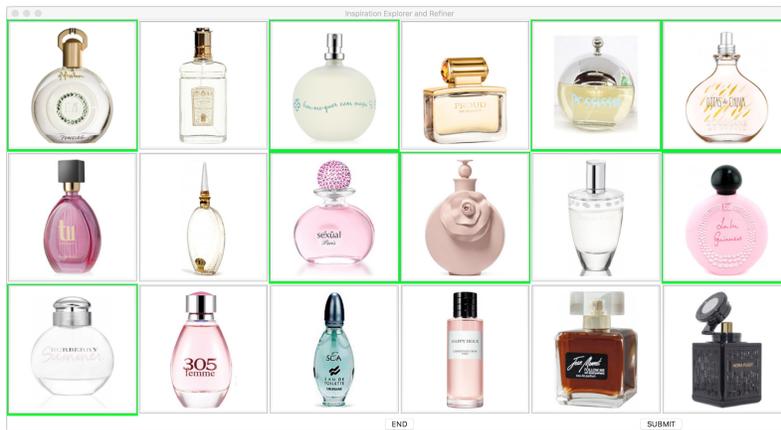


Figure 13: User interface - at the 5th round, selections in green.

1.5 User Study

We performed a user study of the proposed AI design assistant with 13 volunteers (7 male, 6 female), to evaluate its use in a simulated real, initial design process. These were non-expert volunteers since a key use-case of our tool is to assist novice designers or even non-designer business users, such as marketers, needing some quick seed ideas. First asking them to consider how they would design a perfume bottle on their own, we then explained and showed them each component of the AI design framework. We then randomly selected one or more design tasks for each and asked them to try out the interactive Iterator component for the design task, and provide feedback.

Specifically, we asked the following 5 questions about their experience compared to without using the framework. On a scale of 1 to 5 (5 being the most, and 1 being not at all): **(Q1)** “How useful do you think the whole suite of tools would be to help in your design process for coming up with a new, creative design?”; **(Q2)** “How much did it help you discover candidate designs related to the task?”; and **(Q3)** “How much did it help you explore the wide space of designs and stimulate your creativity?”. Questions 4 and 5 were free response questions: **(Q4)**: “What did you like most about the tool?” and **(Q5)**: “What would you most want to improve about the tool?”.

There were 6 design tasks given as short text descriptions, which mirrored the ambiguous briefings usually provided to designers, meant to capture abstract feelings or concepts. An example of 2 design tasks were “The smell is sweet, fruity, girly, and flirty” and “Funky, quirky, unique, - they want to stand out from the crowd”.

1.5.1 Results

Overall we received positive responses¹, users seemingly pleasantly surprised and excited about the capabilities offered and the ability of the Iterator to quickly suggest related designs. The mean and standard deviation of scores for the first 3 questions are given in Table 2, and percentage of ratings ≥ 4 in Table 3, suggesting users did find the proposed system useful.

Q1	Q2	Q3
4.19 ± 0.38	4.31 ± 0.60	4.00 ± 0.61

Table 2: Mean \pm std. dev. of user ratings per question

Q1	Q2	Q3
100	92.3	76.9

Table 3: Percent of user ratings ≥ 4 per question

Since the ground truth concepts were only in the users’ minds, we could not hold out a test set to evaluate performance, so instead we report average AUC, log-loss, and number selected per round across the batch of 18 proposal, averaged across 18 runs (some users tried multiple design tasks), in Figure 14. We show results for the first 5 rounds, since we only enforced users to go through 5 rounds. As in the automated study, we see the system is able to quickly learn the user’s preference and increase selections.

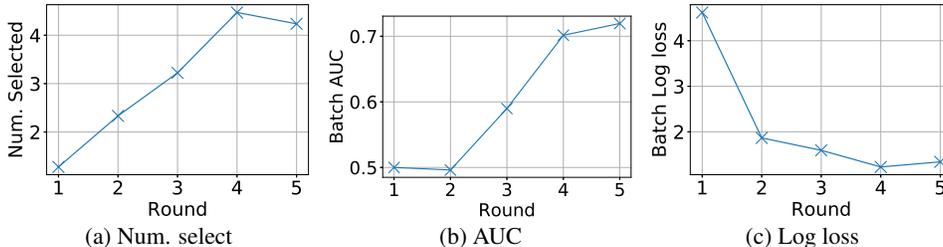


Figure 14: Number selected, AUC, and log-loss per round.

From analyzing the detailed responses, some themes and insights emerged. Users felt the sketch-to-design and style transfer components could be very helpful and save a lot of time. For the Iterator, users liked how easy-to-use it was - the aspect of being able to see many different designs and simply select what they liked. For the most part users felt it could hone-in on their design concepts, and felt the exploration aspect was useful and in some cases provided interesting or unexpected ideas. However, users sometimes felt there were too many similar designs presented and wanted to see more variation around their target concept, as well as an ability to see past selections and use them to derive further variations. Having more capabilities to control the exploration was a key desire - such as the ability to fix the color theme, or take aspects of designs they liked, like shape or color only, or vary certain aspects of selected designs.

1.6 Additional related work

In addition to co-creativity, the presented framework is related to and leverages several areas of work, which we briefly summarize here. A number of packaging design tools exist that allow a user to create 2D or 3D forms of structural and graphic designs. These include ESKO ArtiosCAD [46], ESKO DeskPack [47], Adobe Illustrator [48], Adobe Photoshop [49] and Creative Edge iC3D [50]. These typically require extensive training and hands-on learning. There also exist commercial platforms that match designers according to users’ needs [51, 52].

Recent advances in AI techniques such as Generative Adversarial Networks (GANs), neural style transfers and variational autoencoders (VAEs) have enabled explorations of creativity [53] in domains ranging from music [54] and arts [55, 56] to fashion [57, 58]. A key challenge that remains to be addressed is not generating random designs, but how we enable a human in-the-loop to create a useful design with active control, which we attempt with our framework.

¹Complete responses, question and task details and results, and code will be made available upon publication.

The design space exploration aspect of our framework, i.e., trying to find designs matching user intent, is somewhat similar to visual/semantic search [59, 60]. However, unlike in search, the designer may not know or be able to articulate the target design up-front, and may further want to build on top of initial concepts. Our system instead allows a user to interact with it and learns about the user’s preference based on explicit feedback, and also presents creative suggestions for inspiration that may be outside of what the user originally has in mind.

The interactive learning and design proposal components of our framework is related to reinforcement learning (RL) [40]. However, unlike in RL, the user may be focused on creative exploration around a design goal as opposed to choosing a single best set of designs, and our goal is to help the user discover and come across new design ideas. Similarly with respect to Iterator, active learning research [61, 62] could be used to learn an underlying design concept from as few as possible user labels. However, active learning does not have the concept of exploiting which is essential for the human user to get an immediate benefit from using the system.

Related to our Iterator, there is a large body of work on evolutionary computation to evolve designs [14], including interactive genetic algorithms, e.g., for clothing [15] and bridge design [16]. However, unlike our Iterator, these approaches are bound by the random populations that evolve (i.e., the user simply evaluates evolved examples from a given population), require structured specifications and generative procedures to enable the evolutionary generation (which can be difficult to obtain or result in less-realistic design realizations), and do not learn a re-usable user preference model. We take a different, learning-based approach with our iterator that is also capable of efficiently exploring existing designs and learning an explicit user preference model which has additional applications, such as filtering and scoring new designs directly. Incorporating evolutionary computation approaches in our framework for further design evolution could be an interesting area of future work.

1.7 Conclusion

We proposed an interactive AI framework for design, with a realization and implemented application for package design, to speed-up the design process, make it easier for both designers and non-expert users, and enhance creativity and exploration of the design space. It consists of 3 main components: Creator, for quickly creating design realizations and variations and automatically generating designs; Evaluator, for evaluating designs along different modalities of interest, like customer impact and aesthetics, and to help guide design exploration; and Iterator, for interactively exploring the design space and learning the user’s current design preference. We constructed a perfume bottle data set to carry out a proof of concept. We showed results for a number of Creator components including design generation, sketch and color to design realization, and style transfer. Additionally we implemented and tested a complete realization of the Iterator system and demonstrated it is capable of quickly learning the user preference and generating relevant design proposals.

There are several areas of future work. One is to test incorporating more modalities, e.g., sales and consumer impact, design cost, etc. Another important area of work is developing new metrics to evaluate the success of the proposals that specifically incorporate measurements of exploration, creativity, and surprise. Finally, an important future step is enhancing the framework based on the feedback, e.g., adding the capability for user-controlled design variation in Iterator, and integrating the different components into one tool.