

A Deep Dive Into Exploring the Preference Hypervolume

Alexander Hagg^{1,2}, Alexander Asteroth¹, Thomas Bäck²

¹Bonn-Rhein-Sieg University of Applied Sciences, ²Leiden Institute of Advanced Computer Science

¹Germany, ²The Netherlands

alexander.hagg@h-brs.de, alexander.asteroth@h-brs.de, t.h.w.baeck@liacs.leidenuniv.nl

Abstract

Computers can help us to trigger our intuition about how to solve a problem. But how does a computer take into account what a user wants and update these triggers? User preferences are hard to model as they are by nature vague, depend on the user's background and are not always deterministic, changing depending on the context and process under which they were established. We pose that the process of preference discovery should be the object of interest in computer aided design or ideation. The process should be transparent, informative, interactive and intuitive. We formulate Hyper-Pref, a cyclic co-creative process between human and computer, which triggers the user's intuition about what is possible and is updated according to what the user wants based on their decisions. We combine quality diversity algorithms, a divergent optimization method that can produce many, diverse solutions, with variational autoencoders to both model that diversity as well as the user's preferences, discovering the preference hypervolume within large search spaces.

Introduction

Although we will never be able to fully and concisely grasp it, Hegel describes that in its core, the creative process strives to discover one's true self. Art has the task to reflect upon the spectator and is necessarily interactive (Hegel 1842).

We experience *intuitions* about what we want but are often not able to formalize our preferences. They are based upon direct experience, on cross-connections we make based upon unrelated experiences, on familiarity but also on our problem solving skills (Raidl and Lubart 2001). We experience a physical sensation as a reaction to these intuitions, as defined by C.G. Jung, or more precisely, introverted intuitions, which are paramount to discovering one's own preferences (Jung 1923). Research that provided evidence that intuition and creativity are positively related (Raidl and Lubart 2001) however shows that the debate on creativity and intuition has not been and might never be settled.

We use the definition of intuition as a sensation that arises when a perceived pattern is unconsciously matched to another formerly perceived one (Rosenblatt and Thickstun 1994). One source that can act as a stimulus to intuition is external (Raidl and Lubart 2001). The artist's or creative engineer's vigorous search cannot be performed in a vacuum. We might have an initial idea and the ability to perform a

divergent search, yet need others to reflect upon themselves and gain true insight into what they are and want. Sartre emphasizes that the Other is needed for reflection, but the artist, every creator, or indeed everyone, has the responsibility to choose (Sartre and Elkaïm-Sartre 1946). In a creative process, especially when *reflecting* with others, we might use the Jungian ability of extraverted intuition, or brainstorming, to come up with and reflect upon novel solutions that do not coincide with our own intuition (Jung 1923).

In order to communicate and reflect upon our preferences with ourselves as with others we need to create examples that try to capture abstract ideas. Creativity, which Guilford defined as the ability to perform *divergent* thinking, is about generating many examples that adhere to one's preferences (Guilford 1967). It is through these examples that we can both explore and communicate our preferences.

We combine the ideas of Guilford about divergent thinking, Jung on intuition and Sartre on reflection by others and pose that the central object of preference discovery in computational creativity should be the creative process that includes all three aspects, whereby the Other is represented by a computer and by other creators. We present an example of such a process, akin to the generative-explorative creative model (Ward, Smith, and Finke 1999), using quality diversity algorithms to perform divergent search to trigger the intuition of the users, allowing them to take influence interactively. The computer feeds and reflects upon this interaction and merges the intuitions of what is possible and preferred in a model. The model and the quality diversity algorithm are the motor behind the creative process, in which computer and users co-create a common understanding.

Related Work

Evolutionary computation has often given us unexpected solutions to engineering problems (Lehman, Clune, and Mišević 2018). Novelty search (Lehman and Stanley 2011) took the idea of divergent search to a new level by abandoning the objective function altogether, its only goal to find a set of novel solutions. Reintroducing the objective to this purely divergent search method gave way to quality diversity (QD) algorithms like MAP-Elites (Cully et al. 2015). As in multimodal optimization (Preuss 2015), it finds a diverse set of high quality optimizers, but instead of performing niching in the search space directly, it does so in phenotypic

or behavioral space. First applied to robotics, QD finds a large number of high-performing robot controller morphologies by only comparing fitness between similar solutions, in terms of their morphology or behavior. QD keeps track of solutions in an archive of niches and finds a subset of regions in genetic space, called the *elite hypervolume* (Vassiliades and Mouret 2018), or *prototypes* (Hagg, Asteroth, and Bäck 2018). Similarly, we can describe the volume that contains the preferred solutions the *preference hypervolume*.

A computer aided ideation process using QD can be developed that is based on an a posteriori articulation of preference (Hagg, Asteroth, and Bäck 2018), or “design by shopping” (Balling 1999). By using a preference model based on genetic similarity to selected solutions, and incorporating a factor in the objective function that rewards solutions that are closer to the selected ones, new solutions generated by the system are similar to the user’s selection (Hagg, Asteroth, and Bäck 2019). This approach depends on whether genomes that are closer together also are close in their expressed form, which cannot always be guaranteed. When a user prefers a solution, they would certainly expect the updated solutions to be similar in terms of their expressed morphology or behavior, not their encoding.

In reinforcement learning, learning a neural network model from human preferences makes it possible to find robust controllers without an explicit objective and instead showing the user pairs of examples and letting them pick which one they prefer (Christiano et al. 2017). Showing the user a rich set of high performing designs and having the user select designs is not new (Stump et al. 2003). Recent work shows that using generative or latent models that are trained on a diverse set of solutions (Fernandes, Correia, and Machado 2020) allow the user to easily search in the latent space created by the model. This allows interactive evolution of the latent model, for example to interactively recreate images (Bontrager et al. 2018). However, the latent space is per definition an interpolative space which seems to be less suited for ideation processes than an extrapolative space.

HyperPref

We introduce HyperPref, an implementation of the idea of integrating divergent thinking, intuition and reflection into an interactive co-creative process (see Fig. 1). The central process consists of two alternating steps: I) the computer initiates the process by producing a diverse set of high quality solutions and II) the users select the solutions they prefer, after which the computer updates the set of solutions that are preferred as well as high performing.

An initial pool of random solutions is generated and evaluated using a user-defined objective, which can be an optimality criterion or a more general criterion about the appearance of solutions, throwing a wider net for more “free” thinking. The genomes are then expressed into their phenotypes. A latent model, a variational autoencoder (VAE), is trained to compress the phenotypes into a low-dimensional description. This allows us to determine how similar solutions are and perform phenotypic niching despite of the high dimensionality of the phenotypes. QD creates such a




Figure 1: Discovering the preference hypervolume with HyperPref. The co-creative process (gray box) consists of computer-generated solution sets (step I) that are influenced by the users’ selection (step II). A latent model is trained on the phenotypes (1,2) of a random set of solutions. The model predicts phenotypic features (3) while creating diverse solutions using quality diversity (4), enhancing the intuition of the user. The user can now select/deselect solutions. The selection and a snapshot of the latent model form a constraint model. In the next iteration, the selected solutions are extracted and a new population is created by adding small mutations to the selection (5). The objective function is adjusted with a constraint penalty and the process resumes at (1). The phenotypic latent model (not the constraint model) is updated (2) and an intuition about what is high performing and within the user’s selection is expressed (4).

latent niching archive, consisting of high-performing solutions (according to the user-defined objective function) and triggers a first intuition of what good solutions can look like.

Triggering their intuition, users select their preferred solutions from the archive, and a snapshot of the latent model and preferences is saved. By using a similarity metric based on the latent distance of new candidate solutions to the preferred and non-preferred solutions we can determine whether a new candidate solution is likely to be part of the preference hypervolume or not. The preferred solutions are used to create a new set of initial solutions by perturbing the original solutions and adding them to the set of originally selected ones, adding possibly new innovations into the data set. The preference hypervolume usually consists of disconnected regions in the search space. By increasing the mutation strength (the σ of the normal distribution from which the amount of perturbation is chosen), we allow finding shapes between phenotypic clusters.

Note that in contrast to other work, we do not directly use the latent space for the search, only to compare phenotypic similarity of solutions. Although the VAE would allow this, it would constrain the search to the interpolated space between selected solutions. This would only be really sensible when the first latent model from which the users select solutions was trained on all feasible and relatively high performing solutions. Although QD is a strong mechanism to find diverse solutions, it does not guarantee that all solutions are found. We can also assume that the latent model will not be able to model all variations within the solution set,




Figure 2: Shapes are encoded with 16 polar control points that get transformed into the final shape using locally interpolating splines (Catmull-Rom). The shape is then discretized to serve as the input to the latent model.

especially not with low-dimensional latent space, which are necessary for QD to remain feasible. In conclusion, not limiting the search to the latent space will allow more innovative solutions to be found once the users constrain the search to their preferences, considering solutions that would not be considered by the first model.

The computer updates the latent model, which now describes the similarities within the preference hypervolume. By adjusting the objective function, adding the constraint model, the intuition is updated. This process can continue until the users are satisfied. Novelty search (as opposed to QD) and autoencoders have been combined (Liapis, Yanakakis, and Togelius 2013), but without involving an explicit external objective. Our generator searches for quality as well, and we use the autoencoder not as a way to enhance novelty but rather to capture the user’s choice.

Demonstration

HyperPref is demonstrated on a 2D shape domain, consisting of local interpolating splines (Catmull and Rom 1974). The splines are encoded by a polar coordinate based genome (see Fig. 2). By controlling the radius r and angle θ , a large variety of convex and concave shapes can be created.

We simulate two use cases. In an artistic case, the users are looking to design a ninja star, starting from *centrally symmetric* shapes. The second case, which is closer to creative engineering, starts out with *unbalanced* shapes with the objective to find wing profiles. The first objective prefers solutions that are point symmetric through the center point of the shape. The shape is sampled at $n = 100$ equidistant locations on its circumference, after which the symmetry metric is calculated. The metric is based on the symmetry error E_s , the sum of Euclidean distances of all $n/2$ opposing sampling locations to the center: $f_P(\mathbf{x}) = \frac{1}{1+E_s(\mathbf{x})}$, $E_s(\mathbf{x}) = \sum_{j=1}^{n/2} \|\mathbf{x}_j - \mathbf{x}_{j+n/2}\|$. The second objective maximizes the distance between the center of mass and the center of the bounding box around the shape.

For simplicity, we use a 2D latent space which only captures the similarity of solutions based on the largest phenotypic variance. The shape genome consists of 16 genes. QD produces 32 new child solutions for 1024 generations, by using a normally distributed mutation operator with $\sigma = 10\%$ of the parameters’ range as a generator of diversity. The archive holds 20×20 solutions. The convolutional VAE is trained on a GPU with 128 by 128 pixel representations of




Figure 3: Top: starting with centrally symmetric shapes to design ninja stars. Bottom: starting with unbalanced shapes to design wing profiles.

the shapes and the ADAM (Kingma and Ba 2014) training method. The encoder consists of two convolutional and non-linear (reLu) layers, eight filters of size three with a stride of one. Training is performed with a learning rate of 0.001 and maximizes the evidence lower bound, thereby minimizing the Kullback-Leibler divergence between the original and the latent distribution. The solution set is updated using as many perturbed ($\sigma = 10\%$) versions of the selected shapes as the 64 initial solutions. The constraint penalty, which is multiplied with the original fitness function, is based on the *user selection drift* (Hagg, Asteroth, and Bäck 2019), with the minimal distance s of a candidate solution x to a selected solution and $\bar{s}$ to a deselected solution:

$$p(x) = \begin{cases} 1, & \text{if } \frac{s}{(s+\bar{s})} < 0.5 \\ 1 - 2 \cdot \left(\frac{s}{(s+\bar{s})} - 0.5 \right), & \text{otherwise} \end{cases}$$

Experiments Fig. 3 shows the initial computer-generated solution set on the left. The diversity of the sets is clearly visible. We then simulate a group of creators that all have a different preferred shape (shown in the center). After selection, the computer updates the set, reflecting the combination of the creators’ choice and the general objective. This process of user selection is repeated once more. The resulting sets of ninja stars or wing profiles to look like is shown in the second preference hypervolume on the right.

Discussion The primary solution set contains a large and diverse number of shapes, offering users inspiration and feeding their intuition about how shapes could look like. With a specific goal in mind, namely designing ninja stars or wing profiles, users and computers can co-create in an intuitive creative process. The process offers reflection, by combining preferred shapes and zooming in on the preference hypervolume. Only two steps are necessary to create shapes that are close to what one could and would expect from such a creative process.

Conclusion

We showed how to combine the divergent search of quality diversity to trigger the user intuition about what solutions are possible and high performing, allowing creators to select shapes they prefer by shopping for designs, and then having the computer reflect upon that selection, incorporating the preferences through a constraint model and discovering the preference hypervolume in an intuitive, co-creative manner. The features upon which the initial set is based are generated by a variational autoencoder, trained on the phenotypical expression of the solutions, rather than hand-crafted features or genetic similarity. The constraint model is based upon a snapshot of that model in combination with the set of selected and non-selected solutions.

The resulting creative process, which continuously visualizes and updates the creators' intuition, was shown in a simple 2D shape domain. The updates can be fast, depending on the GPU used for training the VAE and the number of QD updates and VAE prediction speed. The current bandwidth of GPUs is such that the method is close to being on-line.

We recognize that optimizations and variations of the introduced process exist. We used a 2D latent space for the purposes of simplicity and visualization, but higher-dimensional latent spaces are more accurate in measuring similarity in detailed shapes. An interesting research path will be to analyze the differences between searching the VAEs latent space, interpolating between selected shapes, and searching genetic space directly, allowing extrapolation away from the modeled surface, which seems to be more fitting for a creative process. Often times we only find innovative solutions during the creative process, and we certainly hope that unexpected, novel solutions are discovered once we made the first few design decisions.

We took a short but deep dive into exploring the preference hypervolume, combining quality diversity with latent models and interactive user selection into a co-creative process that shows what we can expect in the near future when creators work together. We put the human into the loop by feeding and reflecting upon their intuitions, leading the creative process by example: *No one can tell what the painting of tomorrow will be like; one cannot judge a painting until it is done* (Sartre and Elkaïm-Sartre 1946).

References

Balling, R. 1999. Design by shopping: A new paradigm? In *Proceedings of the Third World Congress of structural and multidisciplinary optimization (WCSMO-3)*, volume 1, 295–297. International Soc. for Structural and Multidisciplinary Optimization Berlin.

Bontrager, P.; Lin, W.; Togelius, J.; and Risi, S. 2018. Deep interactive evolution. *Lecture Notes in Computer Science* 267–282.

Catmull, E., and Rom, R. 1974. A class of local interpolating splines. In *Computer aided geometric design*. Elsevier. 317–326.

Christiano, P. F.; Leike, J.; Brown, T.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning

from human preferences. In *Advances in Neural Information Processing Systems*, 4299–4307.

Cully, A.; Clune, J.; Tarapore, D.; and Mouret, J.-B. 2015. Robots that can adapt like animals. *Nature* 521(7553):503–507.

Fernandes, P.; Correia, J.; and Machado, P. 2020. Evolutionary latent space exploration of generative adversarial networks. In *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*, 595–609. Springer.

Guilford, J. P. 1967. The nature of human intelligence.

Hagg, A.; Asteroth, A.; and Bäck, T. 2018. Prototype Discovery using Quality-Diversity. In *Parallel Problem Solving From Nature (PPSN)*.

Hagg, A.; Asteroth, A.; and Bäck, T. 2019. Modeling User Selection in Quality Diversity.

Hegel, G. W. F. 1842. *Vorlesungen über die Ästhetik*, volume 1. Duncker und Humblot.

Jung, C. 1923. Psychological types: or the psychology of individuation.

Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Lehman, J., and Stanley, K. O. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19(2):189–223.

Lehman, J.; Clune, J.; and Misevic, D. 2018. The surprising creativity of digital evolution. In *Artificial Life Conference Proceedings*, 55–56. MIT Press.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2013. Designer modeling for personalized game content creation tools. In *Ninth Artificial Intelligence and Interactive Digital Entertainment Conference*.

Preuss, M. 2015. *Multimodal Optimization by Means of Evolutionary Algorithms*.

Raidl, M.-H., and Lubart, T. I. 2001. An empirical study of intuition and creativity. *Imagination, Cognition and Personality* 20(3):217–230.

Rosenblatt, A. D., and Thickstun, J. T. 1994. Intuition and consciousness. *The Psychoanalytic Quarterly* 63(4):696–714.

Sartre, J.-P., and Elkaïm-Sartre, A. 1946. L'existentialisme est un humanisme.

Stump, G. M.; Yukish, M.; Simpson, T. W.; and Harris, E. N. 2003. Design space visualization and its application to a design by shopping paradigm. In *ASME 2003 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, 795–804. American Society of Mechanical Engineers Digital Collection.

Vassiliades, V., and Mouret, J.-B. 2018. Discovering the elite hypervolume by leveraging interspecies correlation. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 149–156.

Ward, T. B.; Smith, S. M.; and Finke, R. A. 1999. Creative cognition. *Handbook of creativity* 189:212.