Discovering the preference hypervolume: an interactive model for real world computational co-creativity

Hagg, A.

CHAPTER 1

Introduction

There is a tension between artificial intelligence (AI) systems and responsibility. An increasing number of decisions are left to AI systems that use non-transparent models. The consequences of these decisions potentially have a large impact on society. Although explainability of AI as a subfield is growing, it has a hard time to keep up with the rate of adoption, due to the myriad of economically sound possibilities offered to us by AI's speed, generalizing and generating power. At the same time, decision and model credibility issues in deep learning (DL) and optimization techniques are hard to uncover by individuals. Humans need to keep control of decisions made by AI systems. The creative process is an example where generative AI systems are creeping into a domain that is classically seen as purely human. It is in this domain where I will answer questions about how to design fruitful interactions between generative AI and human decision makers.

1.1 Creativity and Optimization

Fundamentally, every creator strives to discover their true self, while acknowledging that they will never be able to fully and concisely grasp it. The function of creativity and art is to allow reflection of both the spectator and the creator. They are necessarily interactive concepts, according to Hegel (1842), and thus inherently iterative. Creativity is a domain in which decision making inevitably takes place on human time scales, updating intuitions through reflection, which go hand in hand.

Contrarily, in real world optimization, algorithms tend to be created that solve engineering problems under the assumption that quality, constraints and preferences can be fully and concisely formalized a priori. Much of the work in optimization research tries to find global optima to such 'perfectly' formalizable problems. But the processes of problem solving and creativity are much more complex and lead to unexpected usage of such algorithms. In a 2014 interview study by Bradner et al. (2014) on the real world usage of automation in design optimization, "*participants reported consulting Pareto plots iteratively in the conceptual design phase to rapidly identify and select interesting solutions*". Yet, the employed algorithms are written with only a single design iteration in mind.

The tension between creativity (as Hegel understood it) and engineering design in practice seems to be obvious. Optimization research tends to take objectives and preferences as a given and as complete, whereas creativity is naturally open-ended. I pose that developing algorithms that embrace this discourse should be a central objective of optimization research. We should accept uncertainties as a given and build algorithms that support creative processes which help to discover preferences. The field of optimization has come a long way, tackled many issues and the advances in machine learning are opening up novel ways of dealing with the intangible: what does the engineer really want and how can we efficiently develop their understanding and support innovative thought?

1.2 Examples Drive Creative Processes

In everyday life, humans can feel insecure or even anxious about questioning their preferences and intuition. We either choose to ignore novel ideas, taking a more conservative stance towards problem solving and building upon solutions we know and understand, or rely on our intuition and trial-and-error-based techniques to develop new ideas. Those intuitions are based upon direct experience, on cross-connections we make based upon unrelated experiences, on familiarity but also on our problem solving skills, as was described in an empirical study by Raidl and Lubart (2001). We even experience a physical sensation as a reaction to these intuitions, or more precisely, introverted intuition, as defined by Jung (1923), which is paramount to discovering one's own preferences.

The artist's or creative engineer's search is not performed in a vacuum. We might have initial ideas. The ability to diverge from this starting point enables us to leave the realm of known solutions to find innovations. Yet, we need others to reflect upon ourselves and gain true insight into what our preferences are and how they relate to those of others, creating 'anchors' that allow us to more confidently explore the space of (unknown) solutions. Sartre and Elkaïm-Sartre (1946) emphasized that the Other is needed for reflection, but every artist and creator, indeed everyone, has the responsibility to choose. In a creative process, especially when *reflecting* with others, we use the Jungian ability of extraverted intuition (brainstorming or group divergent thinking, as coined by Runco (2010)) to come up with and evaluate novel solutions that might not coincide with our own intuition, or trigger our intuition in unexpected ways.

In order to discover and communicate our preferences with ourselves as well as with others we create examples that are supposed to capture abstract ideas. Without examples, communicating novel ideas can be difficult. As Sartre and Elkaïm-Sartre (1946) stated: "*No one can tell what the painting of tomorrow will be like; one cannot judge a painting until it is done*". So creative thought, which Guilford (1967) defined as the ability to perform *divergent* thinking, is about generating many examples and iteratively discovering and converging towards our preferences. It is through these examples that we can both explore and communicate our preferences.

1.3 Outline of the Dissertation

In this thesis, the ideas of Guilford about divergent thinking, Jung on intuition and Sartre on reflection by others are combined to create a Hegelian creative process. It is posed that the central object of preference discovery is a co-creative process in which the Other can be represented by a machine, as is often done in the computational creativity community. This thesis is about exploring efficient methods to enhance introverted intuition using extraverted intuition's communication lines.

Possible implementations of such processes are presented, akin to the generative-explorative creative model introduced by Ward et al. (1999), using novel algorithms that perform divergent search to feed the users' intuition with many examples of high quality solutions, allowing them to take influence interactively. In this process, the machine feeds and reflects upon human intuition, combining both what is possible and preferred. The machine model and the divergent optimization algorithms are the motor behind this co-creative process, in which machine and users co-create and interactively choose branches of an ad hoc hierarchical decomposition of the solution space.

The proposed co-creative process consists of several elements, which are treated in this thesis in the following order.

In Chapter 2, the theoretical frameworks of cognitive psychology and co-creativity are explored. A formal model for an interactive co-creative process is introduced.

To automatically generate solutions, evolutionary techniques are used, which are widely applied in divergent optimization settings. Based on natural phenomena such as behavior encoding, mutation, and heredity, they provide a natural means to search through large and high-dimensional search spaces. The biological and ecologic concepts of the extended phenotype and diversity are applied to evolutionary encodings in Chapter 3 to answer research question I: “Are solutions best compared using their genomes or their expressed phenotype or behavior?”. A case is made to use a phenotypic search method called quality diversity (QD), a divergent search method which will be explained in Chapter 3, which is compared to state of the art multi-solution optimization algorithms. These algorithms can fulfill the divergent and convergent component from the AI’s side. The comparison answers research question II: “What multi-solution optimization method produces the highest phenotypic diversity?”.

The chapter also discusses what types of diversifying features can be used in phenotypic search to compare and produce a diversity of solutions. Research question III: “Can we produce more diverse solution sets when learning phenotypic niching from data instead of using predefined features?” is used to determine, how we can use generative model (GM) to increase solution diversity. GM can be used both to increase diversity of solution sets generated with predefined encodings as well as to provide the encodings altogether. Research question IV: “What are the limitations of generative models in terms of the possible diversity of the solutions they create?” is answered to decide how to use GM in a co-creative setting. See Table 1.1 for a list of publications Chapter 3 consists of.

Due to the fact that the QD method is generally computationally intense when applied to real world engineering processes, efficiency aspects that arise when applying the co-creative model to real world problem domains are discussed in

Table 1.1: Publications for Chapter 3

Hagg, A., M. Preuss, A. Asteroth, and T. Bäck (2020). An Analysis of Phenotypic Diversity in Multi-Solution Optimization. In Proceedings of the 9th International Conference on Bioinspired Optimisation Methods and Their Applications - BIOMA 2020.
Hagg, A., S. Berns, A. Asteroth, S. Colton, and T. Bäck (2021). Expressivity of Parameterized and Data-driven Representations in Quality Diversity Search. In Proceedings of the Genetic and Evolutionary Computation Conference - GECCO 2021.
Hagg, A. (2021). Phenotypic Niching using Quality Diversity Algorithms (accepted). In M. Epitropakis, X. Li, M. Preuss, and J. Fieldsend (Eds.), Metaheuristics for Finding Multiple Solutions. Springer Press.

Chapter 4. For a fluid dynamics domain, behavioral features of fluids are expensive to calculate. In order to increase efficiency, research question V: "Can we model behavioral features in a surrogate-assisted way by sampling based on optimality alone?" is answered. Neural encodings are used in a number of generative approaches, which is why in a robotics domain research question VI: "Can we model neural encodings' behavior ex situ by sampling their outputs and using a behavior-kernel?", is answered. See Table 1.2 for a list of publications Chapter 4 consists of.

Table 1.2: Publications for Chapter 4

Hagg, A. (2017). Hierarchical surrogate modeling for illumination algorithms. In Proceedings of the Genetic and Evolutionary Computation Conference - GECCO 2017.
Hagg, A., M. Zaefferer, J. Stork, and A. Gaier (2019). Prediction of neural network performance by phenotypic modeling. In Proceedings of the Genetic and Evolutionary Computation Conference - GECCO 2019.
Hagg, A., D. Wilde, A. Asteroth, and T. Bäck (2020). Designing air flow with surrogate-assisted phenotypic niching. In Proceedings of the 16th International Conference on Parallel Problem Solving from Nature - PPSN 2020.

In Chapter 5 the user selection is modeled and incorporated into the AI agent's learning and adaptation phase. In order to achieve this, the following research questions are answered: “Can QD results be summarized using representatives?”, to communicate results to the human, “Can QD be influenced by the user by their selected representatives across domains?” and “Can selected prototypical genomes be modeled?”,to allow humans to take influence, and finally: “Can we constrain parameters by penalizing QD’s objective?” and “Can we constrain phenotypes by penalizing QD’s objective?”, to finally produce the full co-creative process. See Table 1.3 for a list of publications Chapter 5 consists of.

Table 1.3: Publications for Chapter 5

Hagg, A., A. Asteroth, and T. Bäck (2018). Prototype discovery using quality-diversity. In Proceedings of the 16th International Conference on Parallel Problem Solving from Nature - PPSN 2018, pp. 500-511. Springer.
Hagg, A., A. Asteroth, and T. Bäck (2019). Modeling User Selection in Quality Diversity. In Proceedings of the Genetic and Evolutionary Computation Conference - GECCO 2019.
Hagg, A., A. Asteroth, and T. Bäck (2020). A Deep Dive Into Exploring the Preference Hypervolume. In Proceedings of the International Conference on Computational Creativity - ICCC 2020.

In Chapter 6 I summarize the findings in this thesis. The appendices contain background information on basic methods that are used in this work. They discuss evolutionary algorithms, a dimensionality reduction method, and a widely used statistical modeling technique.

We tend to portray human creativity as having few boundaries. It might seem folly for a computer scientist to venture into such a domain where we have biases about human capabilities and are not capable of quantifying results in an objective manner without involving the human we aim to partially replace. This dissertation is an effort to connect the field of computer science with the fields of philosophy, psychology, and biology.

The ethical, social and political consequences of replacing humans by algorithms are part of a wider debate about how we organize society in the light of artificial intelligence’s deep impact. But instead of replacing the human, this thesis explicitly chooses to embrace human-computer interaction in creative design, putting the human back into the loop of algorithmic design in generative AI and optimization.