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

Hagg, A.

Summary

Creative processes in engineering are time-consuming. For quite some time, we have striven to develop algorithms that help us alleviate this task and extend our ability to think creatively. Automated generative systems are creeping into this domain. In this thesis, questions are answered about how to design fruitful interactions between generative systems and human decision makers. Instead of replacing the human, this thesis embraces human-computer interaction in creative design, putting the human back into the loop of algorithmic design in generative systems and optimization. This thesis mostly revolves around shape optimization in expensive settings, for example when creating shapes that are used in fluid dynamics. It is here where a large potential exists to increase the efficiency and efficacy of creative engineering.

In Chapter 2, common theories on creativity and the creative process are described. Aspects from creative cognition that can hinder or facilitate the ability to find innovate solutions are taken as an inspiration to design a computational co-creative process model. Requirements are developed which allow this thesis to then explore efficient methods to enhance engineers' capabilities independent of their experience level.

The ability to generate diverse solutions in an automated manner depends on how we measure similarity and how diversity can be maintained. Chapter 3 discusses a key concept of evolutionary computation, that of the separation of search space (the genome) and the actual solution space (the phenotype). By measuring similarity and diversity in high-dimensional phenotypic space, more diverse solution sets can be developed. When comparing multisolution optimization paradigms, it is shown that quality diversity (QD) optimization, specifically the Voronoi-Elites (VE) algorithm, produces the most diverse solution sets. QD performs niching in a low-dimensional space that describes factors of phenotypic variation. By making use

of a generative model (GM), e.g. variational autoencoders, the AutoVE algorithm combines VE with such an autoencoder. It is shown that we can automatically learn phenotypic features in an on-line fashion instead of using a manual a priori definition. Limitations of GM are shown as well. Evidence is gathered that show that GM should be used as a niching, not as a search space in divergent search.

Efficiency is a key ingredient to co-creative systems in fluid dynamics and other computationally expensive problems. In Chapter 4, surrogate-assisted methods are developed that can predict phenotypic features and behavior. The ability to model behavioral features through surrogate-assisted phenotypic niching (SPHEN) by sampling training examples based on optimality alone, allows us to not only create optimal solution sets efficiently but also to predict their diversity in an on-line fashion. The ability to model neural encodings' behavior using phenotypic distance (PHD) allows the use of more complex encodings in optimization. Both are key to developing flexible, efficient QD algorithms.

Finally, the preference hypervolume, which contains all preferred (phenotypic) solutions is discussed in Chapter 5. This hypervolume is what a co-creative process aims to find and describe and is the central object of this research, in which we use the design by shopping paradigm to allow a user to select their preferences from generated examples. The ability to efficiently generate diverse solution sets is helpful but the vast number of solutions and complex aspects of diversity needs to be presented to a user. This is accomplished by compressing the solution set into representative genetic prototypes. Users then are able to effectively select what examples they prefer and which ones they dislike. This selection is turned into the user selection drift metric that can be used as a penalty to any objective function. The underlying selection is modeled using the user decision hypersurface model (UDHM) preference model. This forces algorithms like QD to start producing artifacts that are similar to the user's preferred solutions. The UDHM is integrated with QD in the prototype discovery using quality diversity (PRODUQD) algorithm. The chapter then combines the idea of focussing on phenotypes, not genomes (Chapter 3), to produce the phenotypic drift metric and penalty. The final algorithm, HyperPref, implements the requirements of a co-creative process as defined in Chapter 2.