



Universiteit
Leiden

The Netherlands

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

Hagg, A.

Citation

Hagg, A. (2021, December 7). *Discovering the preference hypervolume: an interactive model for real world computational co-creativity*. Retrieved from <https://hdl.handle.net/1887/3245521>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3245521>

Note: To cite this publication please use the final published version (if applicable).

The Creative Process

A first important step is to analyze formalisms of human creativity and the components of such processes. The aim of this chapter is to find out which components can be replaced or extended by computer algorithms. It serves as a connection between insights from psychology on the components and traits of creativity. The chapter describes requirements for systems in the field of computational co-creativity. Two main questions about the application of computational techniques on creative processes are answered: during which phases can and should we apply computational techniques and what psychological traits should be taken into account. The chapter discusses the various formalisms of creativity and co-creative processes that were described by theorists, what psychological effects can encourage and discourage creativity, what requirements computationally creative design processes have, which computer aided design processes already exist, and where they fail to meet those requirements. Finally, a formalized co-creative process is introduced, designed to extend human creativity in engineering using an AI agent.

2.1 Creativity

The ‘standard definition’ of creativity by Runco and Jaeger (2012) requires a combination of originality and effectiveness, on equal footing. Without originality or novelty, something is not considered to be creative. If we reuse or replicate old solutions to the same problems, we might rather call this replication. The advantage of replicating a solution is that it has proven to be effective. If a solution or product is ineffective, we would not see it being used, or probably at least see less use of it. This Darwinian view is certainly not a general rule, but is an accepted heuristic in engineering design.

2. THE CREATIVE PROCESS

Then again, if we come up with novel solutions, be it by transforming known solutions, transferring them from other domains or by coming up with entirely novel ideas, they certainly have to be functional and effective. This was already established by Stein (1953): “*The creative work is a novel work that is accepted as tenable or useful or satisfying by a group in some point in time*”. However, Stein stated on novelty: “*The extent to which a work is novel depends on the extent to which it deviates from the traditional or the status quo*”. Novel solutions are found at a distance from known solutions, yet only the latter are identified to be useful and represent engineers’ expert knowledge.

Novelty and (known) usefulness are in direct conflict (see Gardner (1993)). Within this conflict emerges functional creativity, combining new elements with existing knowledge. When we want to find novel solutions, we have to venture into unknown territory and deal with uncertainty. Old solutions are well-understood and there is evidence of their usefulness. A random solution we have not seen before is more novel than those that are known but there is no documented evidence for its usefulness. However, randomization certainly acts as a pressure to create novel solutions. So how do we go about in a functionally creative process? How do we deal with uncertainty, knowledge and novelty?

Runco and Jaeger (2012) acknowledged that there is an ongoing debate about how creative processes work. Simonton (2007) presented two competing theories concerning the creative process using Picasso’s production of the painting *Guernica* as an example. The ‘systematic view’ describes an expert-driven (monotonic) creative process whereas the opposing theory is a non-monotonic Darwinian blind-variation-and-selective-retention theory, which is much more unpredictable. In the latter, free association of creators is a combination of blind variation, but steered by unconscious biases called selective retention. Simonton showed that an expert panel, consisting of neutral, pro-monotonic and pro-non-monotonic experts, agreed that the process of creating *Guernica* was non-monotonic. Weisberg and Hass (2007) found that blindness, the inability to predict the outcome of a change, is a common theoretical element in both monotonic and non-monotonic processes. Again, two competing drivers of creativity can be distinguished: prediction of outcome and variation towards new solutions.

Creativity is a balancing act between blind variation to explore novel solutions and being able to predict the outcome of this variation. I pose that this balancing act can be ‘resolved’ by computational systems. If a machine could be programmed to

combine these abilities inherently difficult for humans, we could create a system that allows humans to concentrate on the task of judging the outcome with less bias than in purely human creative processes. What follows now is an overview on formalisms of functional creativity.

2.1.1 Formalisms

Formalisms of functional creativity and creative processes were subdivided into three main classes by Lubart (2001): factorial, associative and stage theories. *Factorial theories* describe creativity as an ability that consists of a set of skills and factors, or traits (Wang and Nickerson (2017)). Amabile (1988, 1996) developed a factorial theory which describes creativity as having three components: domain-relevant skills (knowledge, technical skills, special talents), creativity-relevant skills (coping with complexity, heuristics to generate novel ideas, e.g. trying counterintuitive ideas when stuck, persistence and sustained attention to a task) and task motivation. *Associative theories* define creativity as a process of combination of mental elements and associative thinking to connect those elements (Mednick (1962)). Finally, *stage theories* define creativity as a process consisting of fixed steps. This is certainly the most interesting set of theories when designing co-creative algorithms, as they describe processes that can be formalized into algorithms. For this reason, stage theories are described here in more detail.

The classical four-stage model of the creative process by Wallas (1926) consists of the phases of preparation, where information is gathered, incubation, where ideas churn in the person’s head, illumination, where what seems to be a solution becomes apparent, and verification, where the individual checks out the apparent solution. Some stage theories show a similarity between creativity and biological evolution, for example Campbell (1960) discussed two stages of creative thought. In the first, wide variations of ideas are produced, and in the second a selection from those variations takes place. Production and selection are both relevant for creative processes. Cropley and Cropley (2010) extended the classical model of the creative process to the following phases: preparation, activation, generation, illumination, verification, communication and validation. But they stressed that this phase model “*is not an exact and concrete description of the process of the emergence of functionally creative products*”. Cropley’s model can be linked to the classical four-stage model by having the activation and generation phase subsumed

2. THE CREATIVE PROCESS

into the incubation phase. The other phases either exist in the classical model as well, or they describe parts of the transitions between phases.

Limitations Guilford (1967) criticized the four-stage model because “*it tells us almost nothing about the mental operations that actually occur*”. It was criticized in other work (see Lubart (2001)), for example in work by Goldschmidt (1991), who analyzed the sketching protocol of architectural designers. She found that the design process is much more complex than what the four-stage process allows for: “*The dialectics of sketching is the oscillation of arguments which brings about gradual transformation of images, ending when the designer judges that sufficient coherence has been achieved*”. This and other evidence from creative engineering processes leads to other, more ‘advanced’ models and more research on what actually happens during the phases.

Connection to Associative Theories Various authors described stage theories that are connected to associative theories’ mental elements. In the first stage, an external stimulus results in a search cue in the mind that leads to the activation of knowledge, both directly and through association. This search cue acts as a pulse that activates knowledge through the (connectivist model of the) brain. This effect can be translated into the context of machine creativity by the neighbourhood relations used in neural representations. In a second step, the combination of knowledge from different areas produce creative ideas. The more distant the connections that those associations activate are, the higher the originality of ideas, but at the cost of higher cognitive load (Santanen et al. (2004); Nijstad and Stroebe (2006)). Machines do not suffer from cognitive load and can be used to help humans find more distant solutions. Wang and Nickerson (2017) suggested that providing diverse and unrelated stimuli may improve idea originality.

This thesis does not limit itself to a particular stage theory but uses the model as a structural basis to develop co-creative processes. I explore the possibilities that new developments in the fields of optimization and machine learning allow for. It is important to find a ‘natural’ way to integrate machines into the creative process, using strengths of machines where humans are ill-equipped to perform certain tasks. In the following sections and subsections, requirements for co-creative processes are summarized (in boxes).

2.1.2 Iterative Processes

Cropley and Cropley (2010) acknowledged that the phases of a human creative process are not written in stone. Rather, they follow an idealized pattern that is not always directly representable in real world settings. Phases can be merged, recombined and looped. Creative processes often may not follow a linear pathway because new insights can arise while creating. For example, the activation phase might be revisited many times until a fully effective solution is developed. New insights, errors of judgment and the need for compromise due to differences in preferences of creative teams happen all the time. Looped or iterative creative processes can deal with these issues, as they allow us to reflect and communicate, and decompose the problem we are trying to solve. Therefore, looped creative processes are of special interest in this thesis.

Arieti (1976) defined the creative process as an iterative sequence of steps or phases. This looped process can be encoded using a Darwinian interpretation of creativity (see Dasgupta (2004) and Cropley and Cropley (2010)). Another example is the Geneplore model by Finke et al. (1996) in which a two-stage subdivision is presumed: in the first, generative stage (or phase), based on memory retrieval, association and analogical transfer, a preinventive structure is generated. In the second, exploratory stage, attributes and interpretations about that structure are found, out of which a meaningful creative outcome emerges. Both Arieti and Finke apply the idea of iterative search to the creative process.

Shaw (1989) believed that the loops and phases could be modeled mathematically. He foresaw the introduction of powerful new computational techniques that would support creative processes based on these models (see Aldous (2017)). These looped creative processes give way to the idea of computer aided design and ideation, which are central concepts used in this thesis.

Requirement

The creative process should be iterative and looped.

Before we dive into what creative machine agents can do for and with us in an iterative process, we need to establish what traits human creators need to have to be creative and what traits inhibit creativity. The mental elements that determine

2. THE CREATIVE PROCESS

the ability to be creative consist of cognitive skills and traits that are described in the next section.

2.2 Creative Cognition

Some criticism arose about models of creativity, especially on the stage models, first formulated by Guilford (1967) when he wrote about the rather abstract phase of incubation: *“It is not incubation itself that we find of great interest. It is the nature of the processes that occur during the latent period of incubation, as well as before it and after it.”*

The interest in what processes happen in the mind gave rise to creative cognition, a subfield of cognitive psychology. Creative cognition aims to advance the understanding of creativity and fundamental cognitive operations that produce creative thought, and to provide experimental evidence for such cognitive processes (see Finke et al. (1996) and Ward et al. (1999)).

To be creative, to perform a creative act, we need certain traits and abilities. These include problem definition and redefinition, divergent thinking, reorganizing of information, feature mapping, analysis, analogical thinking, selective combination, evaluation, thinking styles (being more or less detail oriented), personality traits (tolerance of ambiguity, openness, perseverance) and domain-relevant knowledge (see Lubart (2001, 2005)).

As shown in the following, some traits of our mind prevent us from being creative, which is where computers can help us.

2.2.1 Divergent Thinking

As mentioned before, novelty and usefulness, the two aspects of functional creativity, are extremes on the scale of evidence. On the one hand, old solutions are well understood and can be used with greater certainty about their usefulness. On the other hand, we do not have direct evidence for the usefulness of novel solutions. When we venture into finding novel solutions, we have to deal with this lack of evidence. So the question arises how we traverse the space of all solutions. This process has been described as divergent thinking (see Guilford (1967)): the generation of a large quantity and diversity of solutions, which may therefore

be original, in contrast to a process that converges and usually leads only to convention.

Mednick (1962) introduced an *associative theory* that describes the process of moving from one idea to another through association, e.g. through *functional proximity*. These idea chains form a structure in divergent thinking with the more original ideas, so called *remote associates*, produced late in this form of ideation, far removed from the starting point.

Requirement

Comparisons between solutions should be based on their functional expression, their *phenotype*.

Volle (2018) presented evidence that divergent thinking is connected to two modes of cognitive processes: associative and cognitive control processes. The first activate remote ideas by spontaneous *spreading of semantic activation* from close to distant associates (representatives). The control processes allow for voluntary activation, elaboration and selection of remote ideas.

Requirement

To support activation of remote ideas by semantic spreading, close as well as distant associates need to be perceived by the user.

Divergent thinking is an indicator of creative potential, yet only a moderate indicator of creative performance. Runco (2010) suggested that it is beneficial to consider a large number of ideas. They have to be produced in an efficient way in real world co-creative systems, which is important because long evaluation times are common in these settings.

Requirement

The user should be efficiently presented with a large number of ideas.

Divergent thinking is not reserved for the individual alone. Brainstorming, using Jungian extraverted intuition, is a form of group divergent thinking. Runco (2010) performed experiments with participants in brainstorming sessions. They were “told to (a) postpone judgment, (b) focus on quantity of ideas (fluency) and not quality of ideas, and (c) use each other’s ideas as springboards for one’s own

2. THE CREATIVE PROCESS

thinking.”. Runco noticed that “when in groups, social loafing is likely, and when working in a group, the most original ideas tend to be risky precisely because they are original. They are risky in the sense that an individual is taking a chance by sharing something that other people may not understand or appreciate.”

Since brainstorming with others can be restrictive to creativity, original ideas might survive longer if individuals can perform brainstorming with an ‘objective’ participant, a computer that introduces novel ideas but also reflects and varies the user’s ideas.

Requirement

The creative process should contain an objective participant to reduce the risk of restriction through misappreciation of ideas in a group of human agents.

2.2.2 Priming

A lack of creativity when generating ideas was found to be connected to recently activated knowledge and examples. The hypothesis *“that a broad or narrow scope of perceptual attention engenders an analogously broad or narrow focus of conceptual attention, which in turn bolsters or undermines creative generation”* was tested by Friedman et al. (2003). Their experiments provided support for *“the attentional priming hypothesis, suggesting that situationally induced variations in the scope of perceptual attention (and simple cues associated with such variations) may correspondingly expand or constrict the focus of conceptual attention within the semantic network, thereby improving or diminishing creativity”*. Presenting a diverse set of examples will therefore increase creativity in the human.

Requirement

A diverse set of examples primes the user for an increase of creativity.

2.2.3 Convergent Thinking

Divergent thinking is necessary to find novel solutions, yet their usefulness is not a primary goal in the process. Solutions might not be as useful as they could be. In order to innovate, ideas need to be optimized towards reaching goals to be

useful. As stated by Müller-Wienbergen et al. (2011): “*The human mind is prone to reproduce what it is used to. [...] Hence, IT systems supporting creative work have to support creative individuals by extending their personal knowledge while, at the same time, preventing them from merely walking down beaten tracks*”. They “*propose a design theory for IT systems that support both convergent and divergent thinking*”.

Both convergent and divergent thinking are involved in functional creativity. A computer aided ideation process should therefore be designed to help with both divergent as well as convergent thinking.

Requirement

The artificial agent should perform both divergent as well as convergent thinking.

2.2.4 Reorganization of Information

The process of reorganizing information is part of creative thinking as well. Reorganization, according to Lubart (2001), consists of “*problem construction, information encoding [...], category search [...], specification of best fitting categories, combination and reorganization of category information to find new solutions*”. Mumford et al. (1991) postulated that “*information is stored and interpreted in categorical structures [...]. These categorical structures or schemata represent an organized interrelated set of discrete pieces of information, where organization is derived from the central features of the category. Information about solutions is reorganized alongside feature categories*”. The organization of divergent search results along meaningful features is an inherently important feature to communicate them to the user.

Requirement

Solution candidates have to be presented along meaningful feature dimensions.

2. THE CREATIVE PROCESS

2.2.5 Dissonance

A link was found between cognitive dissonance and creativity. Dissonance causes unease and discomfort, but is shown to be a possible precursor to creativity, simulating human cognition. As Runco and Pritzker (2020) described: “*Since dissonance is discomforting, it often functions as a source of arousal. The dissonance provides most people with a directional goal in terms of which they can either search for memories to support the desired conclusions or search out new knowledge to create a new explanation. The latter is actually how cognitive dissonance leads to the motivated reasoning resulting in creativity.*”

So how can we trigger dissonance, possibly even in the minds of ‘less creative’ people? I pose that by having a computer generate novel but functional solutions, the minds of engineers can be triggered, even if they seek to avoid these triggers, by codifying the induction of cognitive dissonance within the creative process.

Requirement

A machine should ideally trigger cognitive dissonance in the user by generating novel and functional solutions.

2.2.6 Resonance: Experience, Intuition and Bias

Psychologists have observed an interesting effect of experience when they tried to correlate it with creativity. On the one side, a layman without any experience usually has no way to recognize problems or novel solutions. One needs to be somewhat of an expert to be able to perform activation, due to the knowledge base necessary to be manipulated to yield effective novelty. However, weathered experts can have a “*vested interest in maintaining the status quo*”, as Cropley and Cropley (2010) stated. Even without such a conscious interest, functional fixity and confirmation bias can lead to an unwillingness or inability to recognize new problems and solutions. There is quite some evidence that the “*relationship between level of preexisting knowledge and creativity is U-shaped: Both very high (great expertise) and very low (ignorance) levels of preexisting knowledge may inhibit Activation*” (see Cropley and Cropley (2010), Mumford and Gustafson (1988), and Martinsen (1995)).

Humans thus tend to generate solutions that are based upon those they have encountered before. This makes the phase particularly hard for those that lack that experience. They might tend to perform domain substitution, using solutions from other domains, but only after they have understood the problem at least to a certain extent. Experts are more able to generate solutions, but here may arise a problem again, as too much experience could lead to a more conservative stance on design.

So with this in mind, computer aided creative processes should be designed to support dealing with both extremes of the scale of expertise. A need for creative processes that can support both novice as well as experts exists. As Bonnardel and Zenasni (2010) emphasized: “*novices and expert designers differ in their creative process and the kind of help that a computer system could best provide depends on the user’s level of expertise*”.

Requirement

The creative process should be designed to support both extremes of the scale of expertise.

2.3 Computational Creativity

After describing how creativity can be decomposed in interacting phases, what cognitive traits encourage or discourage creativity, and how these traits should be supported by computers, in this section computational creativity and computer aided ideation are discussed.

Computational creativity, according to Colton et al. (2012), is the “*philosophy, science and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviours that unbiased observers would deem to be creative*”. This is a more strict interpretation of computational creativity. As the Association for Computational Creativity (2020) defined: “*The goal of computational creativity is to model, simulate or replicate creativity using a computer, to achieve one of several ends: to construct a program or computer capable of human-level creativity, to better understand human creativity and to formulate an algorithmic perspective on creative behavior in humans, and to design programs that can enhance human creativity without necessarily being creative themselves*”.

2. THE CREATIVE PROCESS

The difference between computationally creative systems and creative support tools lies in the degree of *responsibility* of the system we create (see Colton et al. (2012)). In the former case, the system itself should be deemed creative, in the latter case it only assists the user in being creative. Maher (2012) introduced a new perspective on creativity as being ascribed “*to a computational agent, an individual person, collectives of people and agents and/or their interaction*”. A computer aided design system was first introduced by Kamensky and Liu (1963). It featured an early interactive creative process involving humans and a computer.

Lubart (2005) proposed four categories of human-computer interaction to promote creativity:

- computers may facilitate the management of creative work
- computers may facilitate communication between individuals collaborating on creative projects
- the use of creativity enhancement techniques
- the creative act through integrated human-computer cooperation during idea production

The third approach is dominant in most work (see Wang and Nickerson (2017)) but this thesis takes specific interest in defining creative support systems in which the computer agent and the human user or users cooperate on a somewhat equal footing. The system should be able to be used as a creativity support system for individuals as well as for groups, taking advantage of both introverted as well as extraverted intuition and incorporating the artificial agent as a colleague, or rather an inspiring subordinate.

Requirement

The creative process should incorporate an artificial agent as an inspiring subordinate, and allow one or multiple users to interact.

As was established by Lubart (2005), some of the creative artificial intelligence programs that were deemed a ‘failure’ because they were not fully autonomous and always involved some level of human interference, were actually examples of successful human-computer interactions that facilitate creativity. He stated that it is possible to conceive of computers as real partners. He continued to say that “*The most ambitious vision of human-computer interaction for creativity involves a*

real partnership, in which humans and computers work hand in hand". Machines *can themselves be creative, or contribute new ideas in a dialogue with humans*". Although this position on machine creativity can be criticized, because in the end we humans judge whether something is creative or not, or at least define the metrics that can measure creativity, the weaker proposition of machines contributing and communicating new ideas leads to co-creative processes. Requirements for an iterative co-creative process are described in the following section.

2.3.1 Requirements of Co-creative Systems

The requirements that were defined in Sections 2.1.2 – 2.3 can be summarized in three general requirements.

- The creative **process** should be evolutionary, iterative and interactive, containing (at least one) objective participant as well as one or more users.
- The **search method** used by the objective participant should generate a diverse set of novel and functional solutions to trigger both dissonance as well as resonance with novice and expert users.
- Solution **similarity** should be based on their functional expression, the way solutions solve a problem, and presented along meaningful feature dimensions.

Research on functional co-creativity has produced a large set of (possible) requirements. Two complementary works have had a large impact on co-creative design environments. Parmee and Bonham (2000) defined requirements for techniques that support innovative and creative design activity in an interactive design team and evolutionary search environment. Müller-Wienbergen et al. (2011) stressed the importance of knowledge in creative work and propose that creativity support systems should support knowledge acquisition.

The full list of requirements is based on the previous sections and requirements extracted from Parmee and Bonham (2000) and Müller-Wienbergen et al. (2011). The list is summarized in Tables 2.1, 2.2, 2.3, and 2.4.

2. THE CREATIVE PROCESS

Table 2.1: Process requirements for co-creative processes I

ID	Requirement
A1	The process should be iterative
A2	The process should be interactive and contain (at least one) objective participant
A3	The initial solution set presented to user should need as little guidance from the user as possible
A4	The process should capture design knowledge through designer interaction
A5	The process should be adaptable along the resonance-dissonance dimension
A6	The process should contain evolutionary components

Table 2.2: Knowledge requirements for co-creative processes II

ID	Requirement
B1	Solutions should be categorized
B2	Knowledge should be organized hierarchically
B3	User can filter knowledge base dynamically
B4	Process should provide diverse perspectives on existing knowledge

Table 2.3: Search requirements for co-creative processes III

ID	Requirement
C1	Search should efficiently generate a diverse set of novel, functional solutions
C2	Search should efficiently sample complex design spaces
C3	Search should identify regions of design feasibility and optimality
C4	Search should support addition, removal and/or variation of constraints, objectives and variable parameter bounds
C5	Solution similarity should be based on their functional expression and presented along meaningful feature dimensions

Table 2.4: Communication requirements for co-creative processes III

ID	Requirement
D1	Communication should be evocative through features such as visualization and abstraction

An overview of the state of the art in co-creative systems follows.

2.3.2 Related Co-creative Systems

In this section I will discuss some of the many co-creative support systems developed in the last few decades. It is by no means intended to be complete, but provides a rough overview. The systems described here are generative, presenting a set of solutions, and form an active participant in an interactive co-creative process involving AI and users.

A priori generation, parameter range driven selection. A typical work in interactive design, Parmee and Bonham (2000) and Parmee (2002) introduced an environment in which, by iterative adjustment of rule-based preferences (parameter, constraint and objective weight ranges), an evolutionary algorithm is used to rapidly discover design requirements. Bradner et al. (2014) found that *some* architecture firms use parametric models, associated parameters and simulations as a form of design documentation. Having a parameterized representation at some point in the design process allows designers to tune parameters a posteriori. This process also allows classical parameter-based optimization to take place. Yet, their finding that “*parameters tell the design story*”, which is lifted into the article’s title, is not concisely defined. Moreover, the finding could be rephrased to “*design parameters summarize the history of the design process and the encoding replaces the final object with a collection of objects that are connected through their parameters*”. The question remains whether it is parameters that tell the design story or whether, alternatively, phenotypes and how they are related to each other through high level features give more insight into what solutions could look like, providing more understanding while allowing novel solutions that do not ‘fit’ the parameter space to be protected.

A priori generation, example driven selection. Opposite to a focus on finding parameter ranges, the ‘design by shopping’ paradigm by Balling (1999) and Stump et al. (2003) introduced an alternative way for users to express preferences in a process of *a posteriori articulation of preference*. A Pareto front of optima is created by a multi-objective optimization algorithm, after which engineers choose a solution to their liking. In a study by Bradner et al. (2014) the authors

2. THE CREATIVE PROCESS

reported that many participants consult “*Pareto plots iteratively in the conceptual design phase to rapidly identify and select interesting solutions*”. Pareto plots offer alternative trade-off solutions, which implies that the dimensions of variation follow some style of objective. Herring et al. (2009) showed that examples are a cornerstone of creative practice and are utilized for many reasons throughout the design process.

Secretan et al. (2011) performed an online experiment in which they had a large crowd of users improve the quality and quantity of generated artifacts, using an underlying non-linear evolutionary encoding. Gerber et al. (2012) generated design alternatives according to user defined parameter ranges and introduced automatic quality analysis of these alternatives. Woolley and Stanley (2014) introduced a collaboration between *novelty search*, an optimization method without an objective other than finding novel solutions, and, through interaction, human users. They found that this combination, in which humans select robot behaviors from a computer-generated set, leads to a faster convergence on solutions to some deceptive problems. However, this method still aims to find a single solution to a task. Although it does not emphasize computational creativity, it has promising components. Similarly, Sørensen et al. (2016) showed that users can guide evolution with respect to their preferences. They performed an experiment where users iteratively selected Super Mario behaviors from a set of generated candidates. The participants were shown to be able to evolve diverse behaviors.

In reinforcement learning, creating a neural network model from human preferences makes it possible to find robust controllers without an explicit objective. Christiano et al. (2017) created a system that presents the user with pairs of examples, letting them pick which one they prefer. Bontrager et al. (2018) used evolutionary techniques to search the latent space of generative models and presented these results to participants in an experiment. The participants then selected those results that resemble a predefined target. The experiment showed that the combination of techniques allows a user to evolve complex objects to a specific goal by selection.

A posteriori generation, example driven selection. In an example of a posteriori user selection, Liapis et al. (2013b) introduced a constrained novelty search to generate alternatives to the user’s design. A distance metric based on the Hamming distance between expressed game level maps was used to generate novel

alternatives. The authors then take the idea introduced by Schmidhuber (2007) that beauty can be expressed as the compressibility of objects. This quantitative method measures ‘elegance’ or regularity of an observation. Liapis et al. extended the approach using an autoencoder that compresses the set of solutions which is produced in an exploratory phase. They defined an ‘interestingness’ metric that measures “*where there is potential to improve the compression ratio, or in other words potential [...] to learn about this type of environment*”. This springs from the idea that highly compressible data sets might be elegant but also empty of information. According to the authors, adding a new data point to a solution set should be dependent on how much it improves the compressibility by an autoencoder, uncovering elegance and regularity in the full data set.

The interestingness metric, based on the distance of a solution to known solutions in the autoencoder’s latent space, was used by Liapis et al. (2013) to inform the exploratory step in the next iteration, producing newer, more interesting solutions. This approach was then used by Yannakakis et al. (2014) in an iterative mixed-initiative co-creative process. Alternatives to human design updates are suggested by the computer in real-time. Fitness dimension weights are adjusted to reflect the user’s selection (see Liapis et al. (2014)). However, the mere suggestion of alternatives to the user’s ideas might not be sufficient to trigger cognitive dissonance that might truly inspire the person. Difficult problem domains can have unexpected solutions, and relying on human preferences alone is not enough to reduce the effect of the u-shaped creativity curve. Furthermore, the initial set of solutions might have a large impact on the user’s cognitive priming.

2.4 Process Model

A model for computational co-creativity is defined here. It is built around the reflective and provoking interaction between an objective participant, an AI agent, and one or more users.

The users formalize the representation and quality of solutions in closed form. The agent then performs divergent-convergent search to create a large and diverse set of useful solutions. The set is presented to the users to expand the way they are a priori primed, trigger cognitive dissonance in case they find unexpected solutions, and help discover their preferences. The agent is an active but objective partner

2. THE CREATIVE PROCESS

in the creative process, as it does not have preferences of its own, but rather tries to create as many solutions as possible. This brainstorming aid can make an inexperienced user understand what solutions might look like and inspire an experienced user to think outside of the box.

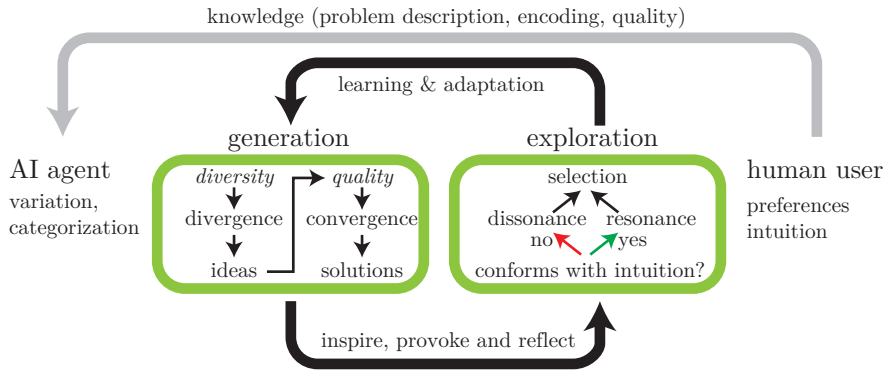


Figure 2.1: Model of the interaction between an objective, creative agent and a human user.

The process model is shown in Fig. 2.1. It is somewhat similar to the generative-explorative models employed by Finke et al. (1996) and Ward et al. (1999), as it consists of the two stages of generation and exploration. In the model introduced here, the generation stage is performed by the AI agent, which, based upon the problem representation defined by the user, creates ideas and turns them into categorized solutions using a diverge-converge operation. This operation is based on evolutionary optimization because of the amount of evidence that they foster diversity (see Hagg et al. (2017) and Lehman et al. (2020)).

The human user then explores these solutions, based upon their internal knowledge and intuition, being inspired by the dissonance caused by unexpected solutions. The user expresses their preferences through selection of solutions. The AI is responsible to inspire and provoke the human but also to reflect the user's preferences. It does so by collecting relevant information through learning and adaptation.

This by no means makes the AI component a fully creative, independent agent, as the knowledge and activation, as well as the exploratory phases, are still the responsibility of the user. They need to define an encoding of a solution and can influence the AI agent's search by adding objectives. However, the AI agent is responsible for generation (divergence), illumination (convergence), verification

(quality) and communication (presentation) of solutions to the user. By adding random variation and evolutionary search inside the agent, ideas (novelty) are added in an objective manner. The agent is also responsible to find high quality solutions and learn the user’s preferences. The human is responsible to select solutions and determine when to end the process.

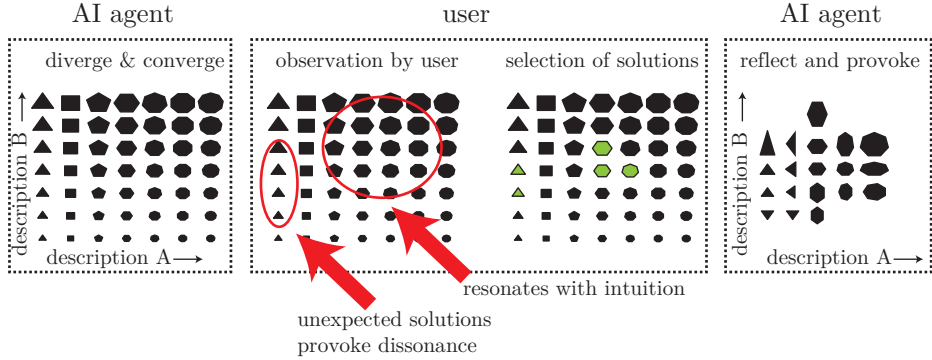


Figure 2.2: Example of a dialogue between an objective, creative agent and a human user. First, the AI agent uses divergent and convergent techniques to generate a diverse set of solutions along meaningful feature dimensions. The human agent is then confronted with this set. Some solutions will conform to their intuition, but unexpected solutions provoke dissonance. The human agent then selects the solutions according to their preferences. The AI agent in its turn reflects upon this selection and provokes the user again by finding variations on the selected solutions.

Fig. 2.2 shows an example of a partial dialogue between the AI agent and the user. The solutions presented by the agent are categorized and visualized in an ordered fashion, along descriptive feature dimensions. After the user’s selection, the AI agent creates a model of the user’s selection and returns a new solution set that is in line with the user’s selection.

The process is iterative (req. A1), consists of interaction (req. A2) between user(s) and an objective AI agent (req. A3), retains information about relevance and features of (intermediate) solutions (req. B1), efficiently generates a diverse set of useful ideas (req. C1, C2, C3, C5), retains information about user(s) preferences (req. A4), generates variations of solutions (req. B4), communicates solutions to the user(s), allows the user to influence the representation of the solution set (req. B3), and allows a user guided, hierarchical decomposition of the solution space (req. B2). Finally, the process is adaptable along dimension resonance and dissonance

2. THE CREATIVE PROCESS

(req. A5). The scope of this thesis does not include req. C4, the addition, removal and or variation of constraints, objectives and parameter bounds.

2.5 Chapter Summary

This chapter discussed the creative process and connected it to cognitive effects that need to be taken into account when designing a co-creative process where AI and human agents work together. Requirements were extracted for co-creative systems and examples for such systems that can be found in literature were described.

A process model for computational co-creativity was defined that is built around divergent search. The process in itself is based on decades of research in cognitive psychology, creativity, and (evolutionary) optimization. The contribution is inspired by the ability and interest in novelty- and diversity-based search. Only in the past decade, methods have been developed that emphasize diversity as well as quality, and the understanding of very high-dimensional objects and spaces.

This work tries to provide a deeper connection between cognitive psychology and (divergent) optimization. An emphasis on high functional diversity of solutions is made to increase the probability of both dissonance and resonance. The use of an artificial, objective participant allows us to question the user's intuition, advancing their understanding of design problems. The objective participant can help to combat group pressure between peers and presents a large number of functional ideas to prime the user and trigger dissonance.

By using phenotypic features to structure the solution set shown to the user, the process allows the user as well as the AI agent to learn a model of the functional solution space over meaningful differences and better understand the user's preferences about *how* artifacts solve a problem. The process is capable of handling high-dimensional artifacts in complex and real world engineering settings.

The objective participant feeds and reflects upon the interaction with the user and merges the intuitions of what is possible and preferred in a model. The model and the divergent search algorithm are the motor behind the creative process, in which computer and users co-create a common understanding.

2.5 Chapter Summary

The following chapters of this dissertation will describe how to efficiently create a diverse solution set while taking into account the user's preferences.