Surprise and reformulation as meta-cognitive processes in creative design

Kazjon Grace K.GRACE@UNCC.EDU
Mary Lou Maher M.MAHER@UNCC.EDU
Software and Information Systems, UNC Charlotte, Charlotte, NC, USA.

Abstract

Creative designs are often surprising. Designers reflect on surprising designs they themselves or others have created and are inspired to explore the causes and consequences of that surprise. We describe a meta-cognitive reflection process based on surprise and reformulation as part of a model of designing to direct a design reasoning process towards creative design. Reformulation is a critical process in designing in which the task is reframed allowing creative solutions to be sought. This process has been demonstrated empirically to be connected to the discovery of unexpected designs and the surprise those discoveries cause. This paper presents a cognitive model of this process, drawing on previous research into models of surprise to describe how they might inform reformulation. Surprise and reformulation are modeled as metacognitive processes that monitor and direct an iterative design reasoning system. Different kinds of surprise derived from perceiving unexpected designs are defined, and from them heuristics for guiding reformulation are developed.

1. Introduction

A hallmark of creative design is the ability to re-frame a problem such that the resulting solution is surprising. This ability for problem-(re)framing is an enduring challenge in the computational and cognitive modeling of design processes. Empirical evidence suggests that designers do not first analyze a problem, frame it in terms of requirements, and then proceed to synthesize solutions, but instead iteratively reinterpret, reframe, and solve design problems (Schön, 1983; Getzels & Csikszentmihalyi, 1976; Poon & Maher, 1997; Grace, Maher, Fisher, & Brady, 2014b). Schön calls this process "reflection-in-action", noting that designers seek unintended consequences early in the design process by externalizing and re-perceiving their partially formed ideas and concepts (Schön & Wiggins, 1992). Figure 1 shows the processes of professional practice as described by Schön (1987). In this model, knowing-in-action contains all the practitioner's grounded, experiential knowledge about performing design in this domain. Surprise occurs when knowing-in-action fails, triggering reflection-in-action, which is the in-the-moment process of reconceiving a design task based on unexpected perception. By contrast, reflection-on-action occurs only outside the design task, when a designer separately reflects on past processes. Reflection-on-action is about making sense of an action after it has occurred, and can lead to learning or the desire for future experimentation.

Reflection-in-action "gives rise to on-the-spot experiment. We think up and try out new actions intended to explore the newly observed phenomena, test our tentative understandings of them, or affirm the moves we have invented to change things for the better" (Schön, 1987, p28-29). This has parallels with the process of conscious attention described in Bryson (2012) as a way of directing a system towards novel stimuli. We propose that the process of reasoning about surprise is a metacognitive process that accompanies a cognitive model of creative design.

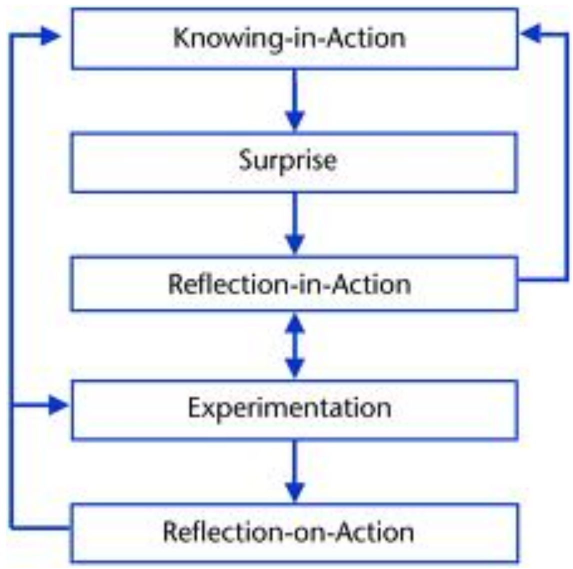

Figure 1. The two kinds of reflection and their contexts. (Schön, 1987).

Reflection-in-action occurs when reasoning about the design solution triggers the designer to transform their understanding of the current design problem, which we consider a conjunction of two processes: surprise and reformulation. Suwa, Gero, & Purcell, (2000) study this effect in practicing architects, finding that novel design requirements arose from unexpected discoveries a significant fraction of the time. They also found that unexpected discoveries begat further requirements invention, suggesting the iterative nature of the relationship between design goals and design solutions. The challenge in computationally representing the creative design process is twofold: what triggers unexpected discoveries and how are design goals invented from them?

This iteration between problem framing and problem solving observed in human cognitive processes inspired a co-evolutionary model of design, as described in Poon & Maher (1997) and shown in Figure 2. A co-evolutionary model of design operationalizes reasoning about the design solution iteratively and in parallel to reasoning about the design problem, with each guiding the reasoning of the other. In this model, generating alternative design solutions to satisfy the goals can lead to generating new goals from unexpected discoveries in the design solutions. In this paper, we elaborate on how the process of generating solutions to design problems can lead to surprise and problem reformulation. We conceptualize problem-framing as a metacognitive process triggered by surprise, as described by Schön (1987) and empirically supported by Suwa et al. (2000). This perspective raises two questions: what triggers unexpected discoveries and how do they lead to the invention of new design goals?

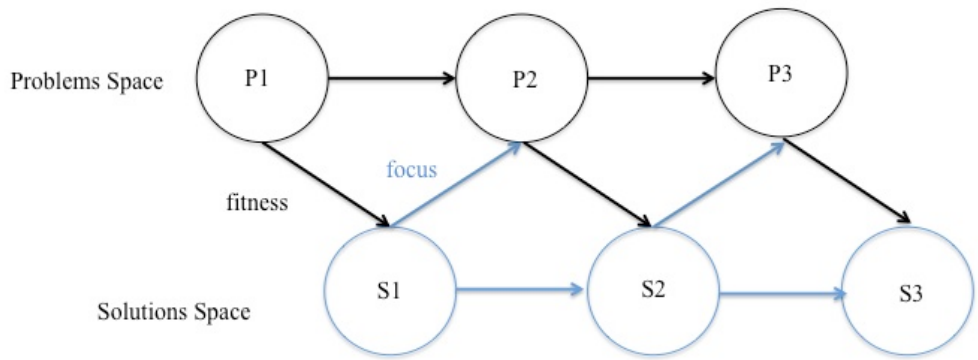

Figure 2. The co-evolutionary model of design from Poon & Maher (1997)

Reflection-in-action has clear parallels with Boden's definition of a creative artifact as one that transforms the domain (1990). According to Boden, transformational creativity changes the domain of knowledge by extending or shifting the boundaries of the conceptual space and therefore changes our expectations about new designs or products. Problem reformulation changes a designer's knowledge about the design task and the domain in which it is occurring, which makes it transformational by nature. Boden also distinguishes historically creative artifacts (h-creativity), those so recognized by a whole society, from psychologically creative artifacts (p-creativity), those recognized as creative by an individual. A transformation of the designer's domain knowledge constitutes p-creativity, with h-creativity being achieved if it changes the conceptual space for all new designs.

We present a computational approach to discovery and surprise in design. In our model surprise is a response to an unexpected discovery made while perceiving designs. Surprise triggers meta-cognitive processes that can lead to modified evaluative and generative reasoning. We describe this cycle in terms of the Cox & Raja (2011) framework for meta-cognition, and provide both a taxonomy of the elements of surprise on which such meta-reasoning is based and a description of how different kinds of surprise can affect the design process. We then develop a set of metacognitive heuristics for reformulating design goals based on this taxonomy of surprise. These heuristics are intended to guide the search for creative designs towards potentially fertile regions of the solution space. This heuristic search approach (Newell and Simon 1976) to creativity has parallels with some of the earliest knowledge discovery systems, such as Lenat's AM (1976).

2. Surprise and creativity

Creativity is associated with new products or solutions that are both novel and valuable (or useful), as articulated in Newell, Shaw, & Simon (1959). Computational models of novelty (eg. Berlyne, 1966, 1970; Bishop, 1994; Saunders & Gero, 2001) have been developed to measure the originality of an artifact relative to what has come before. Newell and others (for example Abra, 1988) describe novelty as necessary but insufficient for creativity, forming one half of the novelty/value dyad. Two additional criteria have been offered as an extension of that dyad: surprisingness and transformational creativity. Surprise has been suggested as a critical part of computational creativity evaluation because computational models of novelty do not capture the interdependency and temporality of experiencing creativity (Macedo & Cardoso, 2001; Maher, 2010; Maher & Fisher, 2012). On the other hand, surprise has also been considered irrelevant to creativity evaluation because it is merely an observer's response to experiencing novelty (Wiggins, 2006b). Boden's (1990) transformational creativity (operationalized in Wiggins, 2006a) has been offered as an alternative by which creativity may be recognized. In both cases the addition is motived by the insufficiency of novelty as the sole accompaniment to value in the judgment of creativity. Thus far these three notions – novelty, surprise and transformativity – have been considered largely incomparable, describing different parts of what makes up creativity. There has been some abstract exploration of connections between the two – such as Boden's connection of "fundamental" novelty to transformative creativity – but no concrete unifying framework.

Previous work (Grace & Maher, 2014) unifies these three viewpoints by reconceptualizing novelty, surprise and transformativity in terms of expectation, and is from this unification that we derive our meta-cognitive view of surprise. The three perspectives on evaluating creativity are that, in addition to being valuable, 1) creative artifacts are novel, 2) creative artifacts are surprising, or 3) creative artifacts transform the design domain in which they reside. In Grace & Maher (2014) these approaches are re-conceptualized to all derive from the notion of unexpectedness, and can thus be situated within a framework illustrating their commonalities and differences.

1. Novelty occurs when an observer's expectations about the continuity of a domain are violated.
2. Surprise occurs in an observer as a response to an unexpected observation (the violation of a confident expectation).
3. Transformational p-creativity (psychological creativity) occurs when an unexpected observation leads to a significant change in a designer's conceptual knowledge about the domain, and transformational h-creativity (historical creativity) occurs when this happens collectively as a response to the same observation.

Many models of surprise involve the observation of unexpected events (Ortony & Partridge, 1987). Maher & Fisher (2012) and Grace et al. (2014b), give a definition of surprise as the violation of a confidently held expectation, a definition derived from earlier computational models both within the field of creativity (Macedo & Cardoso, 2001) and elsewhere (Ortony & Partridge, 1987; Peters, 1998; Itti & Baldi, 2004; Horvitz, et al., 2012). Models of surprise have previously looked at a variety of different kinds of expectation: predicting trends within a domain (Maher & Fisher, 2012), predicting the class of an artifact from its features (Macedo & Cardoso, 2001), the unexpected absence, modification or combination product features (Becattini et al., 2015) or the effect on the data structures of a system when exposed to a new piece of information (Baldi & Itti, 2010). The first case concerns predicting attributes over time, and involves an expectation of continuity of trends within data, the second case concerns predicting attributes relative to a classification, and is an expectation of continuity of the relationships within data, and the third case concerns the size of the change in a predictive mechanism, and is based on an expectation of continuity, but measured by the post-observation change rather than the prediction error. In each of these cases it is clear that expectation is central to the judgment of surprise, but we lack a model of the relationship of expectation to surprise and a typology of expectation as a process that leads to surprise.

This paper is motivated by the question of what happens after a designer has observed an unexpected design and become surprised. We maintain that creative designs are both surprising and valuable, and investigate how a surprising design can trigger the reformulation of a design problem. We present a model of these two processes together – surprise and reformulation – as a kind of reflection-in-action. We do this by incorporating reasoning about surprising observations as a metacognitive process that occurs in response to designs that are unexpected, and describe how that reasoning can lead to transformed design goals.

3. A meta-cognitive model of creative design based on surprise and reformulation

Consider a design system that has the goal of generating a new design that satisfies specified requirements to the greatest degree possible. Satisfying design requirements is typically complex, with many interacting variables and tradeoffs requiring the system to explore and learn about the space of possible designs in order to effectively search for valuable solutions. To this end the system adopts an iterative design process: creating a design based on expectations derived from domain knowledge, perceiving the result, analyzing its actual behavior against the requirements, and then continuing to synthesize new designs. Many more-detailed models of this process have been proposed, including those that build on the Function-Behavior-Structure (FBS) ontology (Gero, 1990), Structure-Behavior-Function (SBF) modeling (Goel 1989), Concept-Knowledge (C-K) Theory (Hatchuel, Weil, & others, 2003), or analogical thinking (Goel 2012). In these existing models, a creative (i.e novel and valuable) design may be the result of the process of searching for valuable designs, but the processes have no explicit metric for the novelty (and therefore creativity) of those designs.

Given our earlier definition of a creative design as one that is both surprising and valuable, a design system of this kind has the potential to produce creative designs. This occurs when a new design's perceived behavior significantly diverges from the system's expectations (making it surprising) and yet it still performs well according to the requirements (making it valuable). The system can measure this divergence from expectations (a term we are using to capture the system's knowledge of the design domain regardless of how that knowledge is structured) and thus identify unexpected designs. What a design system of this kind lacks is the meta-cognitive ability to reason about such discovered unexpectedness and thus systematically explore potentially creative regions of the design space. Such a system can recognize the creative designs of itself and others (see Grace et al., 2014a; 2014b), but it cannot go beyond a generate-and-test model and deliberately seek them out.

Similar to other models we abstract design to an iterative synthesis cycle. We focus on modeling expectation as a basis for meta-reasoning processes, which could in practice be applied to a more detailed design synthesis model. Our basic design system perceives known designs, uses their behavior to learn generalized knowledge about the space of possible designs, and then synthesizes new designs that are consistent with the design goals. We extend this with a meta-reasoning process that can direct synthesis toward creative designs.

Human designers have been observed to focus on unexpected discoveries, reflecting on them and exploring their consequences in future design actions (see Section 1). Our goal is to produce a computational model of this process that captures reflection-in-action: making an unexpected discovery, exploring it by temporarily reformulating the design task to promote it, and through that exploration making further discoveries. Such a model of design would exhibit directed exploratory behaviors, and could reason about unexpectedness and thereby deliberately pursue creative designs.

We develop a meta-cognitive model to address the problem of intentionality, and more specifically, intentionality to generate a creative design. Computational models of designing are a kind of a reasoning process, and therefore surprise, and the process of design reformulation it triggers, can be thought of as meta-reasoning that influences the design process. We adopt the three-level notion of meta-reasoning from (Cox & Raja, 2011), in which a cognitive system can be divided into the object level, the reasoning level, and the meta-reasoning level, as shown in Figure 3. We develop our meta-cognitive model of surprise using this framework in Figure 4.

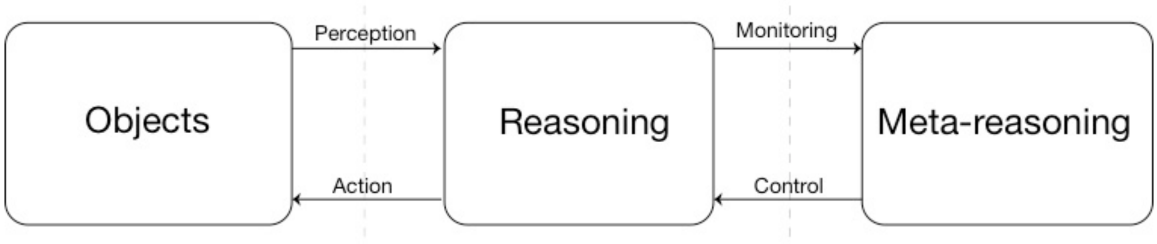

Figure 3. The three-level model of meta-reasoning (after Cox & Raja, 2011).

Following from Cox and Raja, our cognitive model of creative design has 3 levels: an objects level which comprises existing and possible designs in a design space, a reasoning level which has a generative component we call synthesis and an analytic component that we call expectation; and a meta-reasoning level that includes surprise as a process that notices when expectations are violated and reformulation that (re)directs the design goals of synthesis.

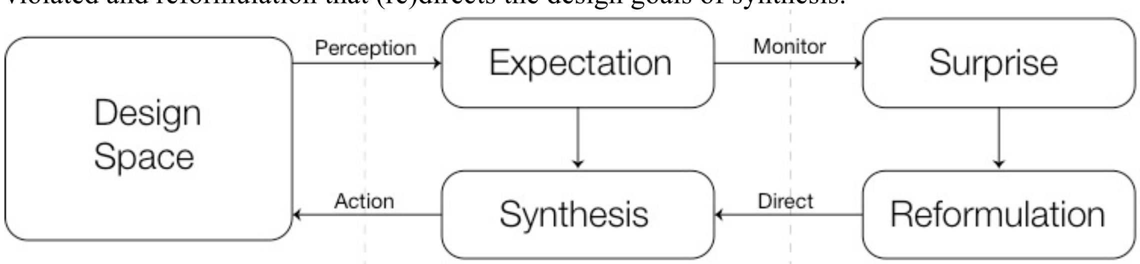

Figure 4. Surprise as a trigger for reformulation via metacognition.

The system's reasoning level contains two processes: Expectation and Synthesis. Expectation compares newly-perceived designs to trends and predictions based on learnt domain knowledge that describes the design space. Synthesis proposes and creates new designs that may be highly valuable based on the predictions made by Expectation. These two processes occur iteratively, with the system iteratively generating and analyzing designs until satisfied with the results. The reasoning process of this system thus performs a directed and exploitative search process, manipulating design variables to maximize performance. We adopt this abstract and simple representation of the complex process of design in order to focus on how unexpected discoveries can lead to new exploratory goals and intentionally surprising output.

The system's metacognition processes are triggered when the expectation process is surprised by a new design – one that was unexpected despite a confident prediction. The Surprise process identifies and characterizes this unexpected stimulus. The Reformulation process then creates or modifies the goals, requirements or constraints of the design task in order to promote exploration of what it believes to be the cause of the surprise. The specific nature of this goal is determined by reasoning about the surprising event itself, based on the taxonomy of surprise presented in Section 5. The result of this reflection is the reformulation of the design task: the modification of the requirement set based on what was found to be surprising. We refer to the surprise and reformulation processes together as reflection-in-action, as, like in Schön's model, it represents a designer re-interpreting the current design task on the basis of an unexpected observation. In this model we focus on the invention of new requirements that focus the search on an under-explored region of the space of possible designs. This re-interpretation of the design space persists until such time as the original surprising design is better understood. In the next section we develop a formal representation of the expectation process with which we can model these processes.

4. A symbolic representation of expectation and surprise

We now introduce a symbolic representation of the expectation processes that occur at the reasoning level of our model. This representation is the basis of the metacognitive surprise and reformulation processes of our model, which reason about the surprising design and the expectation that triggered it. The surprise process in our model occurs as a reaction to a design for which a confident expectation had an unexpected result. Our metacognitive model reasons about this expectation and the design that violated it in order to transform the design goals. We develop a formalism for these interactions below.

At the object level of our model according to Cox and Raja's framework, a design space $\mathbf{D}$ defines possible designs, and additionally a set of known designs $\mathbf{O}$. A design $\mathbf{d} \in \mathbf{O}$ is a member of that set, having been previously observed (and possibly also created) by the system. The set of known designs typically contains those by many different creators, and it is through this set that the system is influenced by societal trends and the solutions of others. Each design consists of a set of intrinsic attributes $\mathbf{O_{ai}}$. These attributes are called *intrinsic* as they arise from the representation of the object within the design space. This model assumes the set of attributes within the design space is fixed, and that each object contains the same set of attributes. The designs are then perceived by the system, which creates a set of *extrinsic* attributes $\mathbf{O_{ae}}$, so called as they arise from the system's interpretation of the design and not from the representation used in the design space. We use $\mathbf{O_a}$ to refer to the set $\mathbf{O_{ai}} \cup \mathbf{O_{ae}}$, with all attributes $\mathbf{a} \in \mathbf{O_a}$ being either intrinsic or extrinsic in nature.

An expectation model $\mathbf{E}$ consists of a 5-tuple $\{\mathbf{Pr,Pd,M,C,U}\}$. Expectation models are specific to a particular design space, and can be applied to any object within that space. The predictor set $\mathbf{Pr} \subseteq \mathbf{O_a}$, comprises the object attributes on which expectations are based. The predicted set $\mathbf{P_d} \subseteq \mathbf{O_a}$ comprises the object attributes about which expectations are made. $\mathbf{P_r}$ and $\mathbf{P_d}$ are mutually exclusive: $\nexists \mathbf{a}, \mathbf{a} \in \mathbf{P_r}, \mathbf{a} \in \mathbf{P_d}$. The model $\mathbf{M}$ is a mapping which takes values for each attribute in $\mathbf{P_r}$ and produces a set of values for each attribute in $\mathbf{P_d}$. The confidence measure

$\mathbf{C(M(P_r))}$ returns a scalar value in the interval [0,1] that is the predictive model's degree of a priori confidence in a given prediction. Confidence measures vary based on the specifics of the predictive model, but often incorporate notions like sparsity and noisiness in the relevant region of the design space. The unexpectedness measure $\mathbf{U(M(P_r),P_d)}$ returns a scalar value in the interval [0..1] that is the deviation between the expected values of the predicted set and the actual values observed in a design for any instance of the expectation. When both unexpectedness and confidence are high, the meta-reasoning process we call surprise determines that the design is surprising, as described in Section 5.1.

The synthesis process takes a set of expectations **X**, a set of design goals **G** and the set of known designs **O** and produces a new design $\mathbf{d'}$. A design goal $\mathbf{g} \in \mathbf{G}$ is a function that takes a design as input and returns a scalar value in the range [0..1] that indicates the degree to which that design satisfies the goal. The attributes of $\mathbf{d'}$ are chosen to maximize the goals in **G** based on the predictions of **X**. Beyond this notion of goal-driven generative design we leave the mechanics of the synthesis process deliberately open, since it is not the focus of this paper.

Surprise occurs as a reaction to a design $\mathbf{d'}$ for which one or more expectations $\mathbf{x} \in \mathbf{X}$ have both a high unexpectedness and a high confidence. The higher these two quantities, the more surprising the design. The surprise process produces a representation of the kind of unexpectedness exhibited by each surprise-triggering expectation $\mathbf{x} \in \mathbf{X}$ in response to the new design. The confidence and unexpectedness thresholds required for an expectation failure to be labeled "surprising" are specific to the design space and the implementation details. We define $\mathbf{S_d}$ as a surprise description process which takes the set of expectations **X**, the design $\mathbf{d'}$ and the vectors of unexpectedness and confidence that resulted from applying each expectation in **X** to $\mathbf{d'}$ and produces a set of surprise descriptions **R** that represent the model's knowledge about the kind of surprise it is experiencing: $\mathbf{R= S_d(X, d').}$ These surprise descriptions, (e.g. "this mobile device violated a trend-based perceptual prediction about phone size"), are based on the taxonomy of surprising design observations described in Section 5.1.

Reformulation uses the model's knowledge about the surprising design it is reflecting on to transform the set of design goals **G**. This process takes the set of surprise descriptions in **R**, as well as the surprising design $\mathbf{d'}$ and produces a new set of design goals $\mathbf{G'}$. This transformed set of goals is intended to generate designs that focus on the surprisingness of $\mathbf{d'}$. This can be accomplished by using the knowledge in **R** to promote exploration of parts of the design space of $\mathbf{d'}$ that caused confident expectations to fail.

In Section 5 we describe how the surprise and reformulation processes operate given the formalization presented here.

5. Surprise and reformulation as a model of reflection-in-action

Cognitive studies have suggested that the perception of surprising designs influences a designer's future design actions. Not all designs that cause surprise do so in the same way, and in order to describe how the perception of such events influence design we must first establish how such events can differ. These differences can then be used to model the "surprise description" process in the formalization described in Section 4. With this representation of the kind of surprise experienced in response to an unexpected design we can then describe how reformulation would transform the design goals of a system. In this section we describe several kinds of unexpectedness in terms of our model of expectation, and present some heuristics for reformulation based on that taxonomy.

In order to illustrate the kinds of surprise described in this section we use a dataset from previous research into evaluating creativity in mobile device designs. This research examined ~5,000 mobile devices from 1989 to 2014, comparing them on the basis of 13 real-valued attributes derived from their technical specifications, as well as their date of release. The research was performed by building regression-based models of the expected relationships between those attributes, and hierarchical clustering models of how the devices could be divided into successively more specific categories based on those attributes. The models and their results are described in detail in Grace et al., (2014a, 2014b).

Figure 5 shows three mobile devices that were evaluated as surprising designs using models of expectation. The Kyman mobile device was highly unexpected because its form factor is unlike a mobile phone or tablet: it is a retail barcode scanner that runs a mobile OS. The iPhone 4 was highly unexpected for its display pixel density and storage capacity relative to its release date, both factors on which Apple's marketing for the device was based. The ICD Vega was unexpected for its large size: it has a 15" display and was marketed as a "tablet for the home" and an alternative to a television.

Datalogic Mobile Kyman Apple iPhone 4 ICD Vega

Figure 5. Three highly surprising mobile devices (after Grace et al 2014a).

5.1 Surprise: reasoning about unexpectedness

Surprise is a process that takes an unexpected design $\mathbf{d}' \in \mathbf{O}$ and the expectation $\mathbf{x} \in \mathbf{X}$ that found that design to be unexpected. Surprise then reasons about $\mathbf{x}$ and produce a description of it that can be used, along with $\mathbf{d}'$, by reformulation. Using our model of expectation in design outlined in Section 4 we have identified seven dichotomies that describe how an expectation that triggered surprise can influence reformulation:

1. Conditional vs unconditional,
2. Holistic vs reductionist,
3. Impact- vs accuracy-measured unexpectedness,
4. Scope-restricted vs scope-complete,
5. Structured vs unstructured predictions,
6. Trend-based vs trend-agnostic, and
7. Epistemic vs perceptual predictions.

Conditional vs unconditional: The simplest models of expectation are *unconditional*, in that they make predictions about a design that are not dependent on any other attributes of that design. Formally we specify an unconditional expectation as one where $\mathbf{P_r} = \emptyset$, meaning that there is no "predictor" and the expectation applies to all members of the design space. Based on our dataset it is reasonable to expect that the thickness of all smartphones will be normally distributed about 11mm with a standard deviation of 1.5mm. This is an unconditional expectation (assuming "smartphones" is the design space, not a label applied within it) as this prediction is not contingent on any other attributes or categories – i.e. $\mathbf{P_r}$ is empty. By contrast, the expectation that all devices with a CPU speed above 500Mhz will have been released after 2005 is a conditional expectation – CPU speed is in $\mathbf{P_r}$, while release date is in $\mathbf{P_d}$.

Holistic vs reductionist: Expectation models can be distinguished by whether they make *holistic* predictions about $\mathbf{d}'$ as a whole or *reductionist* predictions about some of its attributes. For example, expecting all new mobile devices to be proximal to a set of archetypal devices in attribute-space is a holistic expectation. Formally we say that a holistic expectation is one for which $\mathbf{P_r} \cup \mathbf{P_d} \supseteq \mathbf{O_{ai}}$, or that all intrinsic attributes of any design in the space are either predictors or predicted. As defined previously an unconditional expectation has an empty $\mathbf{P_r}$, meaning that in an unconditional holistic expectation $\mathbf{P_d} \supseteq \mathbf{O_{ai}}$. Extrinsic attributes may be present in either the predictor, predicted, or both. Based on our dataset, the first phablet was an example of surprise caused by a holistic expectation as existing devices were predicted (based on all their attributes) to fit either the category of "smart phone" or "tablet", but the first phablet created a new category that was a mixture of both. This category label is an example of an extrinsic attribute in $\mathbf{P_r}$ by which all the intrinsic attributes of a design are predicted. An example of a surprise caused by a reductionist expectation is the iPhone 4 because its display pixel density and storage capacity were very high relative to other designs at that time.

Impact- vs accuracy-measured unexpectedness: There are two approaches by which to construct a measure of unexpectedness (Baldi & Itti, 2010). The first is the degree to which an expectation's a priori predictions of a design are wrong, which we call "accuracy-measured unexpectedness". The second is the degree to which the expectation model must be updated when adapting to fit a new design, which we call "impact-measured unexpectedness". As an example of the impact measure, consider an expectation model that clusters phones into different niches, then measures the degree to which a new phone perturbs its clustering. When such a model encounters a phone that serves as the "missing link" between two previously-separate groups, for example a group of mobile devices we now call phablets, that device would be labeled highly unexpected using the impact measure as it affects the clustering of a every phone in each group as well as new phones. Impact-based measures of unexpectedness have been connected to Boden's notion of transformational creativity (Grace & Maher, 2014). For an example of accuracy measurement of unexpectedness, consider measuring how unexpected a 6mm-thick mobile device is based on the expectation that all mobile devices are distribution around 11mm thickness. Unexpectedness is measured based on a calculation of how rarely the expectation model would expect to see a phone that thin, a measure based on the accuracy of the model.

Scope-restricted vs scope-complete: Some expectations are specific to some part of the design space, in that they will only ever make confident predictions about some portion of known designs. We refer to these as *scope restricted* expectations, and to their inverse as *scope unrestricted*. Formally we say that an expectation model is scope restricted when the expectation only applies to a subset of designs: there exist regions of the space of possible designs for which $\mathbf{C(M(P_r))}$ is 0, and that region contains at least one known design. For example, based on our database the expectation process finds that smart phones above \$500 will have above-average storage space. This expectation is scope restricted because it makes no prediction about devices with a price attribute (a member of this expectation's $\mathbf{P_r}$) below \$500, and its confidence function will always return 0 for them. The exclusion of at least one known design ensures that prediction is restricted meaningfully, as some predictive models may have hypothetical restrictions that bear no impact on the explored parts of the design space, and distinguishing those serves little purpose.

Structured vs unstructured predictions: A surprising design can be described based on the structure of the attributes involved in the expectation that triggered it. Attributes (intrinsic or extrinsic) are called *structured* when there are known relationships between them that stem from the nature of the representation involved. This structure can be spatial (such as between the pixels of an image), temporal (such as between the notes of a song), or typological (such as between the branches of a phylogenetic tree). Either or both $\mathbf{P_r}$ and $\mathbf{P_d}$ can contain structured attributes.

Trend-based vs trend-agnostic: Surprise often occurs from expectation models that capture *trends*. Trend-based models contain within their $\mathbf{P_r}$ an attribute that is based on the time the design was released. This "release date" or "age" attribute permits prediction of changes in the design space over time, and the effect of that drift on the attributes of objects. Based on our database one could assume a trend, present since 2005, of mobile device speeds doubling every 18 months, a well known prediction about computing power in general referred to as "Moore's Law" that has only applied to the mobile device domain for the last decade. This is a very different kind of temporal dependency from that mentioned in structured predictions above: a piece of music has temporal relationships between its attributes (as its representation will contain a temporal dimension), while trend-based expectations capture temporal relationships between design objects that are observed over time.

Epistemic vs perceptual predictions: The extrinsic or intrinsic nature of the attributes involved in an expectation model (see Section 4) can influence surprise-triggered reflection-in-action. We distinguish expectation models that are primarily concerned with the attributes intrinsic to designs, which we call *perceptual expectations* from those primarily concerned with the system's conceptual knowledge about expectations, which we call *epistemic expectations*. This parallels the difference between perceptual and epistemic curiosity outlined in Berlyne (1966). Based on our dataset one could construct the perceptual expectation that wider phones will also be taller, and the epistemic expectation that phablets (an extrinsic label applied to devices and not part of their intrinsic representation) will have larger screens than smart phones. Formally, a perceptual expectation is one in which neither $\mathbf{P_r}$ nor $\mathbf{P_d}$ contain any extrinsic attributes, while an epistemic expectation is one in which either $\mathbf{P_r}$ or $\mathbf{P_d}$ contains only extrinsic attributes. This definition leaves open the possibility of hybrid perceptual-epistemic expectations for which neither of those statements holds.

5.2 Reformulation heuristics

Using the seven characteristics of surprise identified in Section 5.1 we have developed five heuristics for directing the synthesis of new designs that respond to a specifically design that is surprising.

Reformulation Heuristic 1: Induce a goal to explore the design space proximal to the design perceived as surprising. This heuristic is intuitively the simplest response to discovering something unexpected: search the region nearby. It applies primarily to surprise from unconditional, accuracy-measured expectations.

When a system is surprised a design that is highly unexpected according to a confident unconditional expectation, this typically indicates a shift in the perceived boundaries of the design space. An unconditional expectation by definition predicts for all objects, meaning the knowledge it encapsulates describes the limits on what is known to be possible within the space. The system was confident in its understanding of the previous boundaries of the space, and so a surprise may indicate a paradigm shift of which the system was unaware. This kind of expectation model is often referred to as a model of novelty – unconditional predictions of what is or is not considered probable for a particular class of object. In response to observing a design outside those boundaries a design system can induce a simple goal: extend the boundary of the attribute that had an unexpected value and synthesize designs with this new value range. This heuristic can apply to holistic expectations, resulting in exploring a region of the space defined for all attributes, or to reductionist expectations, resulting in a goal to explore constrained ranges of some attributes in conjunction with existing ranges for others. For example, if a system expects phones to be less than 13 mm in thickness and encounters a surprisingly thick phone that is valuable, this heuristic would change the goals to explore thicker devices (leaving, in this reductionist case, the desired ranges of other attributes unchanged).

When surprise arises from an accuracy-measured expectation, the only information available is the degree to which the new design diverged from that expectation. In this case the "proximal exploration" kind of reformulation described by this heuristic is particularly applicable For example, if a new mobile device is introduced that has a landscape orientation when most of the other phones have a portrait orientation, then a goal to explore the synthesis of landscape phones would lead to highly valued phones with a landscape orientation. By contrast, impact-measured expectations provide additional information about how the conceptual structure has changed, making other reformulation strategies possible.

Reformulation Heuristic 2: *Induce a goal that relaxes a constraint on an attribute that bounds the design space, in response to that constraint being violated by a surprising design.* This heuristic is a special case of the first heuristic in which the expectation model is defined in only one dimension and the unexpected design lies outside the previously defined boundaries of the space of possible designs. It applies to specific kinds of reductionist expectations.

This heuristic applies when a reductionist expectation that predicts the value of a single attribute is violated by a newly observed design that has a value for that attribute which would a priori have been considered infinitesimally likely. In that case the design system can relax the boundaries of the design space along that dimension to include the new design, and then focus its synthesis on that newly allowed region. This can be considered exploring the space "near" that design, but in this case that proximity is defined only in terms of a single attribute – new designs created by synthesis may bear no relation to the originating design in any other dimension. For example, in the early 2000s several mobile device manufacturers experimented with "ultralight notebooks", effectively laptops that ran mobile device operating systems, leading to their inclusion in our database. This was highly surprising as the width of these devices was higher than the largest of tablets by a significant margin. While this particular product category was short-lived at the time (although it is being revisited by the mobile-desktop OS convergence visible in, for example, Windows 8), a design system presented with those designs could relax its width constraint and explore very large devices.

Reformulation Heuristic 3: Create a new category from the surprising design and induce a goal to synthesize new designs that help refine that category. This heuristic is primarily relevant to surprise from holistic and impact-measured expectations, and it can also be influenced by scope-restricted expectations.

When a system experiences holistic surprise, particularly when the relevant expectation is using extrinsic attribute(s) to predict the nature of designs, it can create a new category of designs, assign to it the surprising design, and then focus synthesis upon it. This overlaps with the first reformulation heuristic in that designs proximal to the originating "phablet" will be explored, but differs in that the boundaries of the new category and its interactions with other categories can also be explored. For example, given the rise of "phablets" as popular (i.e. valuable) mobile devices, a goal could be created to explore the boundaries of the new category "phablet": what attribute ranges are typical, where are its borders with the categories "phone" and "tablet", and how does it interact with other labels and categories that are known?

This heuristic also applies to the surprise that results from a high impact design: A high impact new design will cause a change in knowledge structure. That change may already include a new conceptual category in which this design is placed, or it may contain information on which a new category can be based. As an example of this heuristic being applied in our dataset, a new kind of phone was introduced in the mid-2000s that had a relatively large touch screen and no physical keyboard, unlike the dominant smartphones at that time. This lead to a re-categorization of smartphones that involved reduced importance on Blackberry-style keyboards and increased importance on iPhone-style large touch screens. The exploration of this new category led to many highly valued designs and eventually the market dominance of the latter category.

Scope-restricted expectations can also influence this heuristic, as the regions for which they make confident predictions are a candidate for bounding the new category. With only one archetypal design assigned to the new category – the unexpected design under consideration – the boundaries of the new category must be refined through exploration, and the information provided by existing scope-restricted expectations makes a good initial candidate.

Reformulation Heuristic 4: Induce a goal to focus on a particular relationship between attributes found surprising in an observed design. This heuristic is particularly relevant to conditional, epistemic and structured expectations, each of which define different kinds of relationships between attributes.

Conditional expectations concern a relationship that is predicted to exist between two or more attributes. When this is violated, the resulting surprise concerns a design that breaks that predicted relationship in some way. When applied to conditional expectations this heuristic explores that relationship, creating a goal to seek further designs that violate it, regardless of whether they do so in a way that resembles how the originally surprising design did it. For example, a design system could, from our dataset, reasonably expect in 2009 that any mobile phone released would have a screen aspect ratio of approximately 1.5:1. In 2010 several designs were released with aspect ratios of 16:9, including the iPhone 5. This violated an existing relationship (screen height vs screen width), which could prompt reformulation to explore other possible values outside of the previously-predicted ratio, including but not limited to the 16:9 ratio that was found surprising.

When surprise is triggered based on a prediction involving a *structured* attribute, that structure represents a relationship that can be exploited in the resulting exploratory goals. This involves leveraging the notion of proximity within the set of attributes, as defined by the structure of those attributes. This relationship can then be used to designs that have the same features as the surprising object among attributes nearby in space, time or an abstract hierarchy. In some cases a surprising element could be repeated within an object, such as a sequence within a piece of music. Proximity-based exploration permits a design system to explore this possibility as a result of encountering or producing a surprising object. For example, if a design system that produces mobile phones used as part of its design representation images of the devices (which is not part of the representation in our previous work, but could be) then it could be surprised by a phone with beveled corners. Given a machine vision algorithm that could detect the typical visual features of a phone, one with differently styled corners would be surprising. Reformulation could use the structure of the attribute representation (here a 2D structure of pixels) to explore adding the same surprising feature (beveled corners) nearby – such as on buttons or display edges.

The *perceptual* or *epistemic* nature of an expectation has consequences for this reformulation heuristic. In addition to the categories described in heuristic 3, the extrinsic attributes found in epistemic expectations can concern concepts that describe a relationship between particular attributes. Surprise arising from such expectations can indicate that the system's conceptual knowledge has become outdated, prompting exploration of alternate possible ways of describing designs within the space. Such alternate knowledge structures would then prompt the creation of designs that help clarify those descriptors. For example, several recently released smartphones are utilizing curved displays in a variety of ways, such as to create a curved phone or to add a secondary display to a phone's beveled edge. This feature is an abstract extrinsic attribute, not a well-defined category: these devices belong to different market segments and their other attributes are very different, but they share this one feature. This heuristic could be used to prompt synthesis to focus on generating devices that possess this particular descriptor, which may create designs that are otherwise unlike the original surprising "curved display" devices.

Reformulation Heuristic 5: Induce a goal to explore a specific domain trend given a new, surprising design that either opposes it completely or is sufficiently "ahead" of it. This heuristic is relevant when surprise from trend-based expectation is violated for a design – either by moving in the opposite direction to that trend or by moving in the same direction so far as to be unexpected.

Surprise arising from trend-based expectations contains additional information on which exploratory design goals can be based. Trends have a direction – the dominant vector of attribute change over time on which the prediction is based. A design perceived as surprising relative to that trend can thus be classified as being "with" the trend (i.e. further in the direction that the trend is heading and thus "ahead of its time" if the trend continues) or "against" it (i.e. in the direction opposite the trend's heading). Ahead-of-their-time designs can prompt the exploration of future designs by projecting trend-based models of other attributes into the future proportionally to the surprising object's unexpectedness. Trend-opposed designs are typically can be explored for the possibility of a novel "niche" based on rejecting the mainstream or returning to what was once the norm. For example, several Android devices retain a large slide-out physical keyboard in addition to their full-sized touch screen. These devices opposed the domain trend of a move to touch-only input, but are popular among the visually impaired, for whom the tactile feedback is invaluable.

Trend-based surprise can also be differentiated by whether the discovered object occurs at the beginning or end of a domain trend, which, in the case of a regression-based expectation model, can be calculated using the second derivative of the trend. Ahead-of-their time objects that occur at the beginning of a new trend are particularly good candidates for exploratory design, as they may be the progenitors of a domain transformation. For example, consider a design system surprised by a very large phone at a time when the domain is trending towards both larger (but not so large as to explain the new observation) and thinner devices. Assuming the surprising phone is about six months ahead of the currently expected trend, the system could create a goal to explore designs that are that large, but also six months ahead of the trend towards thinness.

6. Discussion

We have presented a metacognitive model of a design system that can reformulate its goals when surprised by designs it observes or creates. In this model transformational creativity occurs when knowledge about the domain is updated to correctly predict a new design. This knowledge is expressed in terms of a set of predictive models that make expectations about newly observed designs. We extend our past models of creativity recognition to include meta-cognitive processes of surprise and reformulation. These processes can change design goals by reasoning about the cause of the surprise.

Our metacognitive model of design makes the process of transformational creativity intentional, by setting goals for new designs that will require an adaptation of domain knowledge. This contrasts with existing models of design synthesis that can lead to creative designs (through search-based processes, including analogical reasoning), but not explicitly recognize or seek out the novelty that is an essential aspect of that creativity. Our model is surprised by unexpected designs and uses that knowledge to reformulate design goals and bias the system towards exploring the unexpected with the intention to synthesize further unexpected and valuable designs. We claim that this model operationalizes effects such as reflection-in-action (Schön, 1983) and the s-invention of design requirements (Suwa et al., 2000).

A necessary precondition to a design system with the capacity to surprise itself is a perception process (and resultant object representation) that differs from the one used in generation. This could include observations of actual or simulated design behavior according to performance metrics as well as the learning of separate rule sets for *describing* the space of known designs and *traversing* that space, as proposed in Wiggins (2006a).

Our future work in this area will investigate two kinds of system based on the model presented here. The first is an "innovation analytics" system, in which a computational system performs the expectation, surprise and reformulation processes by perceiving and monitoring the data in a highly populated design space and humans perform synthesis (as well as their own versions of the other processes, which will inevitably influence their creations). This mixed-initiative human-computer collaboration would provide suggestions and descriptions of trends and changes in the domain to the designers, but those designers would be able to follow or ignore those suggestions at will. For complex real-world design domains like the mobile phones example presented here this collaboration seems the most feasible path to a complete implementation of the model described here. We will also pursue a complete implementation of the model as a "computational creativity" system, in which both the metacognitive model of surprise and reformulation and the design synthesis are computational, allowing us to further explore cognitive models of reflection-in-action.

References

Abra, J. (1988). *Assaulting Parnassus: Theoretical views of creativity*. Lanham, MD: University Press of America.

Baldi, P., & Itti, L. (2010). Of bits and wows: a Bayesian theory of surprise with applications to attention. *Neural Networks*, *23*(5), 649–666.

Becattini, N., Borgianni, Y., Cascini, G., Rotini, F., & others. (2015). An Investigation on Factors Triggering Surprise. In *DS79: Proceedings of The Third International Conference on Design Creativity, Indian Institute of Science, Bangalore*.

Berlyne, D. E. (1966). Exploration and curiosity. *Science*, *153*, 25–33.

Berlyne, D. E. (1970). Novelty, complexity, and hedonic value. *Perception & Psychophysics*, *8*(5), 279–286.

Bishop, C. M. (1994). Novelty detection and neural network validation. In *Vision, Image and Signal Processing, IEE Proceedings-* (Vol. 141, pp. 217–222). IET.

Boden, M. A. (1990). *The creative mind: Myths and mechanisms*. Basic Books.

Bryson, J. (2012). A role for consciousness in action selection. *International Journal of Machine Consciousness* 4(2), 471–482.

Cox, M. T., & Raja, A. (2011). *Metareasoning: Thinking about thinking*. MIT Press.

Gero, J. S. (1990). Design prototypes: a knowledge representation schema for design. *AI Magazine*, *11*(4), 26.

Getzels, J. W., & Csikszentmihalyi, M. (1976). *The creative vision*, John Wiley & Sons, New York.

Goel, A. & Chandrasekharan, B. (1989). Functional representation of designs and re-design problem solving. *Proceedings of the 11th Joint Conference on Artificial Intelligence*, Morgan Kaufman, USA, pp. 1388-1392.

Goel, A. (2012). Analogical thinking, systems thinking, visual thinking and meta thinking: Four fundamental processes of design creativity. In *Proc. Design Society's SIG Design Creativity Workshop*, College Station, Texas.

Grace, K., & Maher, M. L. (2014). What to expect when you're expecting: The role of unexpectedness in computationally evaluating creativity. In *Proceedings of the 5th International Conference on Computational Creativity*. Ljubjana, Slovenia.

Grace, K., Maher, M. L., Fisher, D., & Brady, K. (2014a). Modeling expectation for evaluating surprise in design creativity. In *Design Computing and Cognition 2014 (to appear)*.`

Grace, K., Maher, M. L., Fisher, D. H., & Brady, K. (2014b). Data-intensive prediction of creative designs using on novelty, value and surprise. *International Journal of Design Creativity and Innovation, (ahead of Print)*.

Hatchuel, A., Weil, B., & others. (2003). A new approach of innovative Design: an introduction to CK theory. In *DS 31: Proceedings of ICED 03, the 14th International Conference on Engineering Design, Stockholm*.

Horvitz, E. J., Apacible, J., Sarin, R., & Liao, L. (2012). Prediction, expectation, and surprise: Methods, designs, and study of a deployed traffic forecasting service. *arXiv Preprint arXiv:1207.1352*.

Itti, L., & Baldi, P. (2004). A surprising theory of attention. In *IEEE Workshop on Applied Imagery and Pattern Recognition* (pp. 45–54).

Lenat, D. B. (1977). The ubiquity of discovery. *Artificial Intelligence,* 9(3), 257–285.

Macedo, L., & Cardoso, A. (2001). Modelling Forms of Surprise in Artificial Agents. In *23rd Annual Conference of the Cognitive Science Society*.

Maher, M. L. (2010). Evaluating creativity in humans, computers, and collectively intelligent systems. In *Proceedings of the 1st DESIRE Network Conference on Creativity and Innovation in Design* (pp. 22–28).

Maher, M. L., & Fisher, D. H. (2012). Using AI to evaluate creative designs. In *2nd International conference on design creativity (ICDC)* (pp. 17–19).

Newell, A., Shaw, J., & Simon, H. A. (1959). *The processes of creative thinking*. Rand Corporation.

Newell, A. & Simon, H. A. (1976). Computer science as empirical enquiry: symbols and search. *Communications of the ACM*, 19(3), 113–126.

Ortony, A., & Partridge, D. (1987). Surprisingness and expectation failure: what's the difference? In *Proceedings of the 10th international joint conference on Artificial intelligence-Volume 1* (pp. 106–108).

Peters, M. (1998). Towards artificial forms of intelligence, creativity, and surprise. In *Proceedings of the twentieth annual conference of the cognitive science society* (pp. 836–841).

Poon, J., & Maher, M. L. (1997). Co-evolution and emergence in design. *Artificial Intelligence in Engineering*, *11*(3), 319–327.

Saunders, R., & Gero, J. S. (2001). Artificial creativity: A synthetic approach to the study of creative behaviour. *Computational and Cognitive Models of Creative Design V,* Key Centre of Design Computing and Cognition, University of Sydney, Sydney, 113–139.

Schön, D. A. (1983). *The reflective practitioner: How professionals think in action* (Vol. 5126). Basic Books, New York.

Schön, D. A. (1987). *Educating the reflective practitioner: Toward a new design for teaching and learning in the professions*. Jossey-Bass, San Francisco.

Schön, D. A., & Wiggins, G. (1992). Kinds of seeing and their functions in designing. *Design Studies*, *13*(2), 135–156.

Suwa, M., Gero, J., & Purcell, T. (2000). Unexpected discoveries and S-invention of design requirements: important vehicles for a design process. *Design Studies*, *21*(6), 539–567.

Wiggins, G. A. (2006a). A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems*, *19*(7), 449–458.

Wiggins, G. A. (2006b). Searching for computational creativity. *New Generation Computing*, *24*(3), 209–222.