

John B. Cassel

The Methodological Unboundedness of Limited Discovery Processes

Abstract

Though designers must understand systems, designers work differently than scientists in studying systems. Design engagements do not discover whole systems, but take calculated risks between discovery and intervention. For this reason, design practices must cope with open systems, and unpacking the tacit guidelines behind these practices is instructive to systems methodology. This paper shows that design practice yields a methodology which applies across forms of design. Design practice teaches us to generate ideas and gather data longer, but stop when the return on design has diminished past its cost. Fortunately, we can reason about the unknown by understanding the character of the unbounded. We suppose that there might as well be an infinite number of factors, but we can reason about their concentration without knowing all of them. We demonstrate this concept on stakeholder systems, showing how design discovery informs systems methodology. Using this result, we can apply the methods of parametric design when the parameters are not yet known by establishing the concentration of every kind of factor, entailing a discovery rate of diminishing returns over discovery activities, allowing the analysis of discovery-based trade-offs. Here, we extend a framework for providing metrics to parametric design, allowing it to express the importance of discovery.

Keywords: conceptual design, discovery processes, design strategy, open systems methodology, parametric design, non-parametric design

Introduction

The earliest and most conceptual phases of a design are difficult to assess, as they happen before requirements are firmly established. This ambiguity leads to ambivalence about the contribution design might have in improving requirements, as exemplified by this quote:

... designs do not happen by accident; they must satisfy a set of pre-defined set of specifications, even if these specs sometimes get modified as the designer and client both get a better understanding of the design problem and design space. Thus, design is goal oriented. A designer's success is judged by how well his/her design meets desired goals and how well he/she has identified the alternative ways of achieving the those goals (Shah et al. 2003).

In this formulation, the development of requirements is clearly a secondary, even incidental, activity. However, designers know that one of the more powerful aspects of their practice is the choice in how they frame their activities. Capable designers choose which projects they pursue in the first place:

Competent designers act in a radically different way. They select the elements in a situation that are relevant, and choose a plan to achieve the goals. Problem solving at this level involves the seeking of opportunities, and of building up expectations. In process terms a competent designer is likely to be able to become the creator of the design situation, through strategic thinking. This is a very empowering ability, in contrast to the earlier levels of expertise in which the designer was basically just reacting to design situations as they might occur. (Dorst, 2008)

Further, it is often the recognition of an unspecified goal or constraint that is the most important aspect of a design. For example, consider the health care furniture made by Sittris (Sittris, 2009): designed for germ resistance, simple cleaning, caregiver ergonomics, bariatric patient support, and breathability, it competes in terms of health-care specific priorities, establishing its design needs through effective goal discovery. Designers know that exploring, prototyping, and other forms of discovery are critical to their practice, but the articulating a trade-off with unknown benefits is difficult.

Given the value of discovering the surrounding context, designers build their understanding of systems so that they are appropriate to their context. However, designers have to take a different approach than scientists in understanding the systems they work within. One major difference is that specific engagements do not discover whole systems, but take calculated risks between more discovery and intervention. While it is true that some designers spend whole lifetimes learning the particularities of the systems that they deal with, that knowledge grows in fits and starts, bounded by in-project constraints and out-of-project resources. Therefore, design demands that system-building methodologies assess what they have learned as they are learning it. Design practices tackle open systems, and the tacit guidelines for how they do it are instructive to both systems and design methodology.

Fortunately, we live at a time in which there are emerging concepts that deal with open systems that has useful parallels to design practice. One such concept is the non-parametric, or in other words *unboundedness*. This paper will demonstrate how this idea of the non-parametric can be applied to understanding very generic open stakeholder systems, demonstrating how the discovery function of design practice informs systems methodology.

Once we have established this design-influenced approach to dealing with open systems, we will show how to apply this finding back to design methodology, looking specifically at parametric design. Specifically, this non-parametric approach dramatically extends the way we can consider a parametric system, and may even offer to the most analytically-minded design sceptic justifications for the design process from even its fuzziest stages.

First, we will give a background to the systems of design practice and how it embodies the concept of unboundedness. Next, we develop a theory which allows us to talk about how well a given process has addressed unboundedness. Then we will apply this theory to parametric design, observing that non-parametric considerations constitute a substantial percentage of what can be said about the progress of parametric design projects. From there, we will elaborate many potential directions for future work, then conclude.

Background

First of all, there can be no doubt that capable design understands the systemic context of its interventions. Just from the perspective of futures approaches to design research, we see attention paid to all of the following system elements: trends, reinforcing paths of development (called drivers), balancing factors and resistances, tipping/leverage points, weak indicators of change (called weak signals) (Schwartz, 1991), changing environments, and the interaction between structure, function, purpose, and context (Gharajedaghi, 2011). Given this background, a provisional definition for a system is a particular selection of related elements which either calls into view or undermines particular structures, functions, purposes, behaviors, contexts, boundaries, or relationships. From this perspective, systems methodologies provide ways to bring these relations into attention and critique. This broad, nearly anarchic, definition of systems is chosen to permit the most diverse range of design practices.

The constraints of a project-based practice inform the design approach to systems. From design practice, we learn to generate ideas longer and gather data longer, being unsure

of what we know, pursuing past when returns diminish but stopping when it is impractical to gather more (Osterwalder and Pigneur, 2010). Moreover, the means of this discovery is extremely important: novice designers are sometimes frozen in “analysis paralysis”, practicing designers quickly move to see problems from the view of potential solutions (which they occasionally become fixated upon), while truly masterful designers pursue multiple lines of development simultaneously (Cross, 2004). In addition to the design process, there are particular disciplines within design that offer loose guidelines to the kinds of stakeholders and kinds of actions expected for a particular engagement. What this means is that there are concentrations of relevant stakeholders, concerns, and phenomena in the world, and that designers have both a guide for finding them and a process for knowing when they have found them.

It is worthwhile to make a few general remarks about the typical elements of design projects, to which we hope the reader can make analogies and extensions from their own work. Stakeholders are anyone who has an interest in or may alter the outcome of a design. All stakeholders can take various actions according to their own volition to bring about changes in conditions to improve or maintain their situation according to their own assessment. Sometimes, by some criteria, two situations will be incomparable even by the judgement of a particular stakeholder, one better by one criteria but worse than another, in which case they will have to make a trade-off.

Some stakeholders may undertake a design process as a part of their volition, so these stakeholders are designers. A design process is a particular strategy for selecting design activities aimed at producing some number of designs, or plans orchestrating a series of carefully configured actions in the world. These designs aim to achieve the best trade-off given all of the conditions the design could be subject to. Finding such a trade-off can be terribly complicated, as it can involve multiple stakeholders with multiple criteria, but no matter how simple or complicated it is for a given analysis, it is always the best of the choices according to each of the designers willing to pursue it, given that they need the other designers to go along with it. The activities which find designs have various costs, in time, goodwill, material resources, etc. which are then not available for the design itself.

Designers have a vast variety of design activities to choose from, and a few general categories will help us in discussing them. These include discovery activities, which elicit more information; generation activities, which produce designs as such, as well as supporting tools; evaluation activities, which assess the trade-offs of the generated designs; selection activities; which either choose particular aspects of design to undertake or abandon; and coordinative activities, in which stakeholders attempt alignment and persuasion. Although finding a better representation of design activities that best allows for significant particulars to be learned effectively is an interesting open problem, let us instead describe a simple formulation that might be usable in practice. One can index the available activities by the number of times similar/particular activities have occurred with similar/particular elements. For example, we might say we undertook a similar activity with a similar group of participants, but this is the first activity of this kind with this particular group, which has already participated in some particular other activities.

The reader will be able to decorate this simplified picture with other kinds of elements appropriate to their discipline. Now, the designer is given very few of the overall criteria, stakeholders, possible designs, or conditions, but learns about the problem through discovery activities. The fact that these elements are not known, but are the units of analysis, provides an inductive constraint that allows these parameters to be found non-parametrically in design activities, using the approach that developmental psychology applies to causal learning (Kemp et al., 2010), which we introduce next.

The Non-parametric; The Unbounded

Jordan (2010) explains that the phrase “non-parametric” is a bit of a misnomer. To be non-parametric is not to assume that there are no parameters, or elements, of the set of systems being considered. Instead, the assumption is that there is no particular set of parameters under ultimate consideration and that formulations of systems which expand flexibly to accommodate new facts is desirable. The point to take away from this is that we are no longer making judgements about a particular fact, but instead about the processes that are generating the data being observed. What this means is what any designer can tell you: while you cannot reason about factors you have not yet observed, you can reason which further design activities might do well to tell you more, based on what you know so far and what typically comes out of those processes.

Now that we can see how an assumption of unboundedness might bring design process decisions into focus, we can now unpack some ways to put this assumption to use. Suppose first we are looking to determine the some factor we do not know the choices of, such as the favourite restaurant of the residents of an unfamiliar city. If we can arrange some of the discovery activities so that they are exchangeable, which is to say so that it does not matter in which order we make our queries (Aldous, 1985; Pitman, 2006), then as we query we expect that we will continue to discover new categories, but progressively more rarely, and generally will find that new elements fit into the categories in the proportions that we have observed before. With a small number of good restaurants, we will find the popular candidates quickly, but with a larger number we will find a greater number more quickly while having further to go to discover the true proportion of preference. The concentration expresses how many categories have a significant representation. A small concentration means a fewer number of significant categories, and thus a more quickly diminishing discovery rate for new categories, while a larger number entails a larger number of categories, and thus a larger number of discovered items and a longer design process to determine the representative proportions. In all cases, we can find what restaurants people prefer even without knowing all of them.

Suppose we are looking think about the concentration of the set of features that belong to a given thing. We can similarly ask so that we will receive some set of features for each time we ask, again finding new features with diminishing returns at a rate depending upon their underlying concentration (Griffiths and Ghahramani, 2005; Thibaux and Jordan, 2007). Using these two different processes together, designers can talk about the concentration of stakeholder groups as they report varying sets of interests across discovery activities.

Given that we can now talk about how to put observations of factors in salient bundles, let us move to how to reason about relationships. There are two basic questions: does a relationship exist between the two elements at all, and if so, how does one bring about or counteract the other? We can simply make guesses to these questions and then correct them as we continue to make discoveries (Kemp et al., 2010).

With this ability to discover concentrations exchangeably, the causal effects of any particular system can be learned effectively (Goodman et al., 2008; Tenenbaum and Griffiths, 2003; Griffiths and Tenenbaum, 2007; Kemp et al., 2010). Therefore, given the loose units of analysis common to most design processes, we now can use the concentrations we observe to make trade-offs in design activity selection even for factors we have not observed.

Theory

Non-Parametric Design Trade-offs

Let us now see how the assumption of unboundedness leads to trade-offs in design discovery. We are not only interested in trade-offs in between criteria discovered at any particular point in the process, but instead in the projected generated set of criteria given a path of discovery

activities given the observed concentration. However, if it is really true that criteria are incommensurate, then discovering a new criterion could completely change the evaluation of any design choice. In order for an evaluation of discoverable value to remain coherent, we need to establish some assumptions. First of all, for every criterion, there is some margin of indifference where any stakeholder would not actually be able to discern the difference between changes in value below that margin. Following from this, there is some minimum margin of difference for all criteria, and let us say we can assume it. Finally, there is some maximum difference in value for any particular choice by any criterion. With these assumptions, we bound the value of discovering any new criteria, so that it too can fall below diminishing returns, assuring that a good trade-off still exists.

The basic trade-off we are after is between the improvement in the design that might come from being aware of additional considerations versus the cost of the discovery activities it would take to unearth those considerations.

Let us first look at the cost side of the cost/value trade-off. We can say that there is an expected cost of the minimum cost activities for discovering some new categories and relations, given the design activities undertaken so far. Every activity will have some unknown average effectiveness per kind, and if we can estimate these effectivities, then an effective path to a particular set of relations can be calculated effectively. If the effectiveness is unknown, we might choose a trial iteration in order to understand it. This is where the concentration of a given kind leaks in, as a more concentration means that the cost of finding new elements goes up faster, or in other words there is a shorter path to diminishing returns.

However, this is not actually the best way to go about minimizing the cost of the unknown. An activity might be less controlled, in that we might not know which categories and relationships it might reveal, but know that it usually is very productive. What we are looking at then is discovery effectiveness of the activity across kinds, moderated by how valuable the discovery of each kind is. Fortunately, the overall path between systems of a given cost can still be understood.

Now, let us turn to the value side of the cost/value trade-off. The value of the discovery is the further improvement in the design given a set of new categories and relations. This can be treated as the concentration of discoveries which actually cause a design change among those that are newly discovered, along with the magnitude of that change. Let us call the difference between discoveries with a different set of known or expected elements the gap, the best expected design improvement based on more activities.

The point at which to best make the trade-off is the equilibrium point. Ideally, the point of equilibrium between cost and value would be the point at which more discovery processes yield less expected value discovery than their expected costs. Unfortunately, we could always imagine some new aspect of the problem with very high impact. So, instead, we can say that we have found equilibrium if for a value gap of any particular magnitude, weighted by the likelihood of discovering it, is less than the cost of the discovery activities that cross that gap.

The implication of an equilibrium for design practice is that designers who can discover more for a given cost can have better solutions on average, as their equilibrium point contains more discovery. How might a designer get there? Of course, they may develop interpersonal and research skills that allow activities to be more effective. However, they can also attend to which activities which reveal more elements, which activities have lower costs, and which sets of factors discovered from other projects can apply to present concerns.

Application

Extending Parametric Design Non-Parametrically

The most striking form of design to contrast non-parametric design with is parametric design. To design parametrically means to design a constrained system that sets up a design space that can be explored through the variations of parameters (Gane and Haymaker, 2012; Bentley, 2007). While this approach has many advantages, it also means that only designs that can be found through the varying of parameters can be explored. Therefore, it only stands to reason that parametric assessment methods do not assess design activities aimed at discovery, and that parametric methods do not provide direct design guidance to decisions made within the formulation of requirements. This means that the trade-offs involved in the cost of discovering requirements versus the cost of not meeting undiscovered constraints and goals are currently unaddressed.

How could we use the methods of parametric design when the parameters are not yet known? Fortunately, for every parametric framework, there is a non-parametric equivalent. We can find this non-parametric equivalent simply by looking at every kind of input parameter and establishing its concentration in the world, which entails a discovery rate of diminishing returns over various discovery activities, making discovery-based trade-offs analyzable. Now we will do exactly this: by looking at a framework for providing metrics to parametric design, we will find ways to measure improvements that come from considering the non-parametric equivalents of parametric frameworks. The Design Exploration Assessment Methodology (DEAM) is a methodology for assessing the quality of guidance provided by design processes in parametric design problems (Clevenger et al., 2013; Clevenger and Haymaker 2011, 2012). This framework provides a number of useful definitions and metrics for assessing parametric problems with given inputs. By treating these units of analysis as variables instead of as givens, we can directly extend these metrics to assess discovery.

Parametric Design Systems

Parametric design has a special terminology which expressively extend the design process. Here we look at the terminology of DEAM and make some initial qualifications regarding the use of these terms non-parametrically.

- **Stakeholders** are parties with a stake in the selection of alternatives. Parametrically, stakeholders are known while non-parametrically they must be discovered.
- **Preferences** are selections between goals made by a particular stakeholder. One parametric selection is to assign additive weights for each goal, while in the most severe parametric case we discover a preference for each specific set of goal impacts over another.
- **Goals** are declarations of the intended properties of a design solution, specified in terms of a particular *target* value. Departures from goal targets are taken to have negative impacts. Parametrically, goals are given, while non-parametrically they must be discovered.
- **Constraints** are limits placed on the properties of a design solution. Parametrically, constraints are given, while non-parametrically they must be discovered. We say that an alternative that meets all constraints is *feasible*.
- **Objectives** are the all of the goals and constraints of all stakeholders weighted by their preferences. Two factors lead to objectives not being trivially combinable: not all stakeholders legitimately have equal influence, and not all stakeholders are prepared to resolve their preferences between conflicting goals. Parametric design practitioners designate an alternative that satisfies constraints and gives the best available performance along all goals as *optimal*.
- **Variables** are decisions to be made. These decisions can be discrete or continuous, and may have dependencies. In parametric settings, it is the outcome of variables that have dependencies, while in non-parametric settings, even the existence of a variable may depend

on an earlier variable selection. Following work on design scenarios (Gane and Haymaker, 2012), let us call a particular set of variables a design scenario.

- **Options** are the potential values that can be assigned to variables, or in other words the potential outcomes of decisions. Sometimes they may be bounded, but they may also vary non-parametrically (for example, a designer may have to research manufacturing technology to best select a machine to undertake a particular operation).
- **Selections** are the options chosen for variables in choosing a particular set of options.
- An **alternative** is a complete selection of options over all variables present in the given design scenario, yielding a potential design outcome.
- An **impact** is an alternative's performance on particular goal. A simplifying assumption is that performance falls off as a percentage deviation from the goal's target. Generally, we need to tie particular levels of goal performance to particular levels of stakeholder preference.
- The **value** of an alternative is its net performance in impact across all goals as a function of stakeholder preferences. Non-parametric settings also evaluate value risk, or the potential change in the evaluation of value if more information was discovered.
- A **challenge** is the set of decisions to be made, encompassing all of the variables for which options must be selected to determine an alternative. In a non-parametric challenge, the variables to be resolved may be partially unknown and discovered in the exploration process.
- A **strategy** is a set of steps used to generate the basis for decisions regarding variables, which include no advice at all, complete algorithmic specifications, open-ended procedures, guidelines, and heuristics. In a non-parametric setting, a step-by-step procedure may be less complete than a seemingly less guided strategy, due to such procedures failing to accommodate discovery.
- An **exploration** is the sequence of variables and options considered during the course of a challenge. A design exploration in which multiple activities are available should also record these.
- A **design process** is the implementation of a strategy leading to a particular exploration to be taken in the face of a given challenge.
- The **solution** is the alternative or alternatives selected as a resolution to the challenge, ideally being the alternatives that best satisfy preferences. The best solution may be no solution, if there are no viable alternatives.

Another facet of parametric problems that is not codified in the DEAM formulation is the concern for *conditions*, which refers to factors of the challenge which vary parametrically, but are not controlled by the user. Other work refers to these as uncertainties (Johansson and Krus, 2005), which provides another way to look at these non-designed parameters. Typically, constraints must hold in all conditions considered within range in the challenge formulation, and the value of goals might be assessed by their expected value within the conditions. In addition to these terms, it is difficult to talk about non-parametric settings without some additional vocabulary. Among these we find:

- **Activities** are the various procedures undertaken in the course of a design exploration, the purposes of which include discovery, evaluation, and analysis. An open design process not only selects options to analyze, but also selects activities to undertake. Although activities themselves are subject to discovery, working from a relatively stationary set of alternatives allows for learning which activities are most effective.
- A **design scenario** is set of variables which, when resolved, lead to an alternative. By definition, the choice between scenarios itself is a variable. For example, in a structural application, we might choose between particular shapes of beams, which lead to different parameters the values of which specify an exact shape.
- **Challenge scenarios** are the differing sets of conditions in which alternatives are evaluated. These allow for overall feasibility, reliability, robustness, optimal average performance, and other nuances when considering value. Challenge scenarios also must be discovered. This

term reflects the more typical use of the word scenarios, but for our purposes when we say scenarios we will mean design scenarios by default. We say that an alternative that remains feasible under all conditions it might plausibly encounter is *robustly feasible*, if it remains optimal across conditions it is *robustly optimal*, and more informally if the alternative has consistently good performance across all plausible conditions it is *robust*.

- **Functions** are what designs are said to do. What function means in parametric design is insightful, but unfortunately, clouded by product/process and objective/subjective schools of thought (Erden et al., 2008). Fortunately, our purposes allow us to be inclusive regarding approaches to functions, as what we are after is to develop all potential linkages between objectives and the system structure.

Non-parametric Assessment

With this terminology in place, we can describe various metrics for assessing how well particular strategies yielded effective explorations of particular challenges. These include metrics for assessing the difficulty of the challenge, the coverage of strategies, and the value of the alternatives discovered. Following a similar convention as papers on DEAM, we will use italics to delineate metrics. We will first look uniformly at each DEAM metric, and then consider what other metrics might be considered for non-parametric design evaluation, leading to a non-parametrically extended DEAM (say NDEAM).

With this terminology in place, we can describe various metrics for assessing how well particular strategies yielded effective explorations of particular challenges. These include metrics for assessing the difficulty of the challenge, the coverage of strategies, and the value of the alternatives discovered. Following a similar convention as papers on DEAM, we will use italics to delineate metrics. We will first look at each DEAM metric, and then consider what other metrics might be considered for non-parametric design evaluation.

Challenge Assessment

Challenge assessment metrics look at the variety of difficulties imposed by the challenge in order to give various strategies a comparative baseline. There are two key reasons for this. First of all, challenge metrics assess the particular character of the difficulty, implying that particular strategies may be more appropriate for the challenges presented. Secondly, uniformly less sophisticated strategies may do equally well on uniformly less challenging problems, solving a challenge effectively with less resources.

The *Objective Space Size* is taken to be the total number of goals and constraints considered within the scope of a given exploration. The non-parametric equivalents of this metric would consider the growth of the objective space. For example, the *Objective Concentration* measures the likelihood of discovering a new goal in a particular activity, which is expected to diminish in rough proportion to how many times that activity was undertaken. This can further be decomposed by examining the *Stakeholder Concentration* and the *Goal per Stakeholder Concentration*. Equivalent discovery concentrations can be proposed for constraints (we might expect changes in constraints for areas with scientific or regulatory uncertainty). These rates can be measured in terms of sample count, overall time, or cost. These concentrations reflect the intuition that a challenge becomes more difficult if high-impact goals and constraints are costly to discover.

Alternative Space Independence is the number of first-order interactions among design variables divided by the total number of variables. It represents the degree to which coupling between the effect of the variables determines impact performance. A metric for the discovering of coupling is the *Coupling Concentration*, or the relative commonality of interactions. The designer might have some domain knowledge about discovering couplings, such as knowing that certain electrical components cause noise that disrupts the functionality of other components.

Impact Space Complexity is the number of variables found to result in performance trade-offs (divergent impacts) divided by the total number of variables, representing the percentage of variables with competing objectives. The *Impact Complexity Concentration* is the concentration given by the number of trade-offs per discovered variable and goal.

Another source of additional metrics in open formulations is the additional conditional structure of scenarios, both design scenarios and challenge scenarios. The *Conditional Variable Size* is the number of variables that only come into play depending on an option selected earlier in the design process.

Just as different design scenarios lead to different variables and occasionally even different trade-offs between goals, challenge scenarios lead to different conditions and occasionally suggest further constraints. The *Challenge Scenario Size* is the number of different challenge scenarios, while the *Challenge Scenario Depth* corresponds to the number of factors that need to be resolved in order to specify a particular set of conditions. There are a variety of other complexities in discovering and applying challenge scenarios that are assessed through the disciplines of strategic foresight (Schwartz, 1991; van der Heijden, 1997) and strategic forecasting. These matters are beyond the scope of this paper. However, these areas also involve the application of activities to discovering the sorts of parameters useful here, and are subject to similar assessments of size, concentration, complexity, and cost.

In open challenges, the scale of the exploration itself becomes important. In particular, the challenge a particular design process might mitigate is ongoing in real time, such that impacts can be expected to befall stakeholders without timely intervention. In addition, there may be time and resource constraints on the project itself. These challenge-specific constraints on exploration may not be given. The *Real Impact Rate* is the rate at which impacts occur.

Strategy Assessment

Strategy assessment metrics consider the coverage and efficiency of strategies independent of outcome. However, independence of outcome is not the same as being independent of the challenge. In open problems, the assessment of the challenge is made throughout the execution of the strategy, such that the degree to which the size of the challenge was assessed should also provide metrics for the assessment of the strategy.

Objective Space Quality is the extent to which the goals analyzed match the goals proposed, as measured by the overall percentage of goal coverage. In a conceptual design space, this would also imply either developing or explicitly ruling out generic guidelines for appropriate goals, including fitness of purpose, total cost of ownership, manufacturability, ecological impact, regulatory compliance, material and energy efficiency, and similar cross-domain concerns. Therefore, *Objective Look-ahead Effectiveness* is the degree to which objectives assessed prior to establishing stakeholder preference come to be established as objectives by stakeholders.

Alternative Space Sampling is the number of alternatives divided by the number of alternatives required for a “significant sampling” of the entire alternative space. Let us refine this metric a bit, in that what we are not interested in approximating the value of all alternatives, but rather reducing the risk that an unevaluated alternative is of higher value. We can rephrase *the Alternative Space Sampling* as the *Alternative Discovery Gap*, or a confidence-weighted expectation of how much new alternatives might improve over the best known alternatives minus what it is expected to cost to find those alternatives. In an open problem with multiple discovery and evaluation activities, we would like to reduce the risk of not discovering an alternative. The *Variable Discovery Rate* is the rate at which new variables are discovered. This leads to a metric for the confidence that we will not discover variables

that contain alternatives more valuable than the cost of discovery activities, called the *Variable Discovery Gap*. Notice that the *Variable Discovery Gap* is bounded by the *Alternative Discovery Gap*, in that we cannot assess what we know versus what we could discover if we do not know the value of what we know.

Alternative Space Flexibility is the average number of option changes between any two alternatives divided by the number of variables modeled. This metric indicates the variety of alternatives that were explored in a given exploration. In a non-parametric challenge, we are interested in the average number of different variables modeled, divided by the average number of variables modeled per alternative. Given that different sets of variables correspond to different scenarios, we can call this metric the *Scenario Space Flexibility*.

Given the DEAM metrics, let us now consider some strategy metrics that appear when we expand to look at open problems. First, we'll look at metrics that assess the quality of the activities undertaken, and then we will consider to assess how well a given design process characterizes the underlying challenge. The most simple metric for assessing activities is the *Average Activity Cost*. If a given set of activities can produce the same discoveries and evaluations as another set for fewer resources, then those activities are superior. Clearly, the total activity cost is contingent on the determination that activities do not need to be undertaken. Working through an idealized model or interviewing certain stakeholders may reveal that certain options lead to infeasible or terrible alternatives. In the event that some activities render others unnecessary, we reduce the activities that we need to undertake.

The *Activity Reduction Effectiveness* is the degree to which the outcomes from one activity can prevent the need to undertake other activities, measured as one minus the cost of the undertaken activities divided by the cost of downstream activities that would need to be undertaken to make a similar assessment. Even more strongly, when activities are undertaken prematurely, their results can be entirely undermined by the findings from other activities, rendering them ineffective. Therefore, this metric can yield negative values.

There is a downside to activity reduction, namely that the approximation or idealization employed incorrectly analyzes the true behavior of the design, eliminating good designs spuriously. This *Activity Reduction Risk* looks at the potential error caused in both neglecting valuable alternatives and in pursuing spurious alternatives due to mistakes.

We can only assess if the activities of an exploration have found valuable alternatives that merit their cost if the challenge has been sufficiently articulated. The metrics we have already seen for assessing this include the *Alternative Discovery Gap* and the *Variable Discovery Gap*. The Stakeholder Concentration and Goal per Stakeholder Concentration give rise to the Preference Discovery Gap, or the expected difference between the value caused by newly discovered stakeholders (along with their criteria, preferences, and goals) versus the cost to discover those stakeholders. The *Challenge Discovery Gap* similarly looks at assessing the potential for undiscovered challenge spaces in invalidating the value of alternatives addressing the overall objectives. In short, for every parametric input, there is an underlying non-parametric concentration which generates a discovery rate, leading to a decision-theoretic assessment of the loss caused by neglecting undiscovered inputs, which we call the discovery gap.

Exploration Assessment

Exploration metrics examine the value of the solution, or how well the particular application of the strategy met the challenge. The exploration of an open challenge may discover several distinctly satisfiable objectives, and therefore should be assessed for the full variety of solutions produced. The *Value Space Average* is the mean value of the set of alternatives analyzed, while *Value Space Range* measures the dispersion, the standard deviation serving well for this purpose. If the design process discovers distinctly achievable sets of objectives,

each of those alternatives is assessed only with respect to the value of alternatives according to the set of objectives that were being pursued. One can distinguish separable set of objectives by a subset of the stakeholders are indifferent to that alternative. Therefore, open problems allow *Value Space Averages* and *Value Space Ranges*. *Value Space Maximum* is the top value of the alternatives generated in a design exploration. Given that open challenges might provide multiple preferred solutions, we instead assess the *Value Space Maxima* of the exploration. The *Value Space Iterations* metric is the number of alternatives generated before the highest value is reached. Nominally, reducing this number is the efficiency of the exploration.

In an open exploration, there are two important milestones for efficiency. The most direct analog to Value Space Iterations is Cost to Value Maxima, which reflects the actual cost of activities to reach a peak value alternative. However, the more important efficiency metric might be the previously developed *Overall Cost to Cost/Risk Equilibrium*, which is the cost to determine that further activities would not likely be effective in discovering further objectives or valuable alternatives. In other words, this is the effectiveness of the design strategy in correctly determining the exploration is completed. It may be that a given exploration did not find this equilibrium, at which point we can still assess a *Cost to Cost/Risk Gap Margin*, though it may be difficult to compare explorations that stopped with different margins.

Future Work

As this work presents a new concept to design methodology, there is a great deal that can be done to develop and apply the approach highlighted here and this should only be taken as a broad and incomplete overview of those opportunities. Further work currently being developed investigates that, once given a variety of discovery metrics, how do we take advantage of these measures to improve design processes? The ability to catalog methods by their capabilities for effective discovery, evaluation, generation, persuasion, and where necessary, termination could be greatly facilitated by such tools. The most obvious direction for future work in this regard is to create a platform that can assess the NDEAM metrics given above. Developers of computational product models (Wolkl and Shea, 2009) identify four major categories of computer support for design: product models, editors, operational methods, and data/knowledge bases. This work implies that for every kind of model, there should be an operational method capable of assessing non-parametric metrics for it, so there is a rich body of extensions possible for existing approaches.

Next, there are features of technical design, including reliability, testability, and maintainability, that should be more directly incorporated into and subject to non-parametric analysis. For example, particular kinds of checklists and other theory-level guides and heuristics may demonstrate their value in determining design activities in a broad range of domains, such as conceptual guidelines for improving robustness (Jugulum and Frey, 2007). It would also help to explore potential interfaces between nonparametric discovery and experimental design. In a similar vein, it might well be profitable to catalog existing methods by both their capabilities for discovery and persuasion and their cost. Another view into applying design metrics to design methods would be to analyze case-studies of previous projects for their performance according to discovery metrics. One compelling possibility is if this work could be extended to disciplines far from material artifacts, such as cultural policy.

Another pending concern is that this work does not deal directly with domains that are subject to change and so there may be an extension of this work to a non-parametric control theory for handling initially unknown ongoing rates of change among initially unknown parameters. It may also be possible to put a non-parametric lens on theories as such (Ullman et al., 2010).

One final, very speculative application is constructing design education curriculum for methods. A good portfolio of methods will strategically discover and persuade, directly developing the ability to discover in a wide range of design situations at low cost. The development and strategic deployment of discovery abilities may be considerations for design practices when putting together teams and developing training materials.

Conclusion

Overall, we have seen how design practice informs the approach for analyzing systems, which in turn gives new tools to a design methodology. Design practices, through necessity, already employ discovery activities with low costs. These design activities are heuristics for improving a particular trade-off, namely the expected risk due of undiscovered factors versus discovery activity cost. We introduced how the non-parametric is a powerful conceptual tool for talking about the discovery of the unknown. After that, we found that this concept allowed for trade-offs applicable to all discovery. We then saw that these findings applied back to parametric design, allowing for the assessment of the discovery of every factor in parametric models. Finally, we showed that this connection between systems and design opens up a wide range of future work for both areas.

With the activities of discovery, creation, critique, and rhetoric now given an expression of trade-offs that may appeal to analytic tendencies, the content of those activities is given over to the traits of the designer's creative powers: ineffable, tacit, profound, and deep, happy stowaways on the planned and provisioned excursion of discovery, with justification and expression enjoying each other's company as they venture into the horizon of the unknown.

John B. Cassel

Wolfram Research,
Champaign, IL, USA

Email address: john.benjamin.cassel@gmail.com

References

- Aldous, D. (1985). Exchangeability and related topics. *Ecole de probabilités de Saint-Flour, XIII-1983*, 1–198.
- Bentley, D. (2007). *Architectural geometry*. Exton, UK: Bentley Institute Press.
- Clevenger, C.M. and Haymaker, J.R. (2011). Metrics to assess design guidance. *Design Studies*, 32, 431–456.
- Clevenger, C. M. and Haymaker, J.R. (2012). The value of design strategies applied to energy efficiency. *Smart and Sustainable Built Environment*, 1, 222–240.
- Clevenger, C.M., Haymaker, J.R., & Ehrich, A. (2013). Design exploration assessment methodology: Testing the guidance of design processes. *Journal of Engineering Design*, 24 (3), 165-184.
- Cross, N. (2004). Expertise in design: An overview. *Design Studies*, 25, 437–441.
- Dorst, K. (2008). Design research: A revolution-waiting-to-happen. *Design Studies*, 29, 4–11.
- Erden, M. S., Komoto, H., van Beek, T. J., D’Amelio, V., Echavarría, E., and Tomiyama, T. (2008). A review of function modeling: Approaches and applications. *Artificial Intelligence for Engineering Design, Analysis, and Manufacturing*, 22, 147–169.
- Gane, V. and Haymaker, J. (2012). Design scenarios: Enabling transparent parametric design spaces. *Advanced Engineering Informatics*, 26, 618–640.
- Gharajedaghi, J. (2011). *Systems thinking: Managing chaos and complexity: A platform for designing business architecture*. Elsevier.
- Goodman, N. D., Ullman, T. D., and Tenenbaum, J. B. (2008). Learning a theory of causality. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, 2188–2193.
- Griffiths, T. and Ghahramani, Z. (2005). Infinite latent feature models and the Indian buffet process. *Advances in Neural Information Processing Systems*, 18 (17).
- Griffiths, T. L. and Tenenbaum, J. B. (2007). Two proposals for causal grammars. In Gopnik, A. and Schulz, L. (eds.), *Causal learning: Psychology, philosophy, and computation* (323–346). Oxford, UK: Oxford University Press.
- Johansson, B. and Krus, P. (2005). Probabilistic analysis and design optimization of Modelica models. *Proceedings of the 4th International Modelica Conference, March 7-8, 2005*, 247–254.
- Jordan, M. I. (2010). Bayesian nonparametric learning: Expressive priors for intelligent systems. In Dechter, R., Geffner, H. and Halpern, J.Y. (eds.), *Heuristics, Probability, and Causality: A Tribute to Judea Pearl*, Chapter 10. London, UK: College Publications.
- Jugulum, R. and Frey, D. D. (2007). Toward a taxonomy of concept designs for improved robustness. *Journal of Engineering Design*, 18, 139–156.
- Kemp, C., Goodman, N. D., and Tenenbaum, J. B. (2010). Learning to learn causal models. *Cognitive Science*, 34 (7), 1185–1243.
- Osterwalder, A. and Pigneur, Y. (2010). *Business Model Generation*. Hoboken, NJ: John Wiley and Sons.
- Pitman, J. (2006). *Combinatorial stochastic processes*, 32, 1875. Springer.
- Schwartz, P. (1991). *The art of the long view: Planning for the future in an uncertain world*. New York: Currency Doubleday.
- Shah, J. J., Vargas-Hernandez, N., and Smith, S. M. (2003). Metrics for measuring ideation effectiveness. *Design Studies*, 24, 111-134.
- Sittris BA. (2009). Sittris product brochure. Retrieved May 16, 2011, from http://www.sittris.com/brochures/Sittris_Company_Brochure_2009.pdf.
- Tenenbaum, J. B. and Griffiths, T. L. (2003). Theory-based causal inference. In Becker, S., Thrun, S., and Obermayer, K. (eds.), *Advances in Neural Information Processing Systems 15*, 35–42. Cambridge, MA: MIT Press.

- Thibaux, R. and Jordan, M. I. (2007). Hierarchical beta processes and the Indian buffet process. In Meila, M. and Shen, X. (eds.), *Proceedings of the 11th International Conference on Artificial Intelligence and Statistics*. Madison, WI: Omnipress.
- Ullman, T. D., Goodman, N. D., and Tenenbaum, J. B. (2010). Theory acquisition as stochastic search. *Proceedings of the Thirty-Second Annual Conference of the Cognitive Science Society*, 2840– 2845, Portland, OR.
- van der Heijden, K. (1997). Scenarios, strategy, and the strategy process. Global Business Network. Retrieved April 14, 2011, from http://www.gbn.com/consulting/article_details.php?id=31.
- Wolkl, S. and Shea, K. (2009). A computational product model for conceptual design using SysML. *Proceedings of the ACSM 2009 International Design Engineering Technical Conferences (IDETC/CIE 2009)*, San Diego, CA.