

Modeling Design Concepts under Risk and Uncertainty using Parameterized Efficient Sets

Richard J. Malak Jr. and Christiaan J. J. Paredis

G.W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology

Copyright © 2008 SAE International

ABSTRACT

Decisions made during conceptual design can have a major impact on the success of a design project. However, the inherently imprecise nature of design is a major source of uncertainty and risk in conceptual design decisions. A single concept relates to a large set of specific design implementations, each of which has a different level of desirability based on the tradeoffs designers are willing to make. It therefore is beneficial for designers to have an understanding of the various tradeoffs they can achieve by implementing a concept. In this paper, we describe an approach to modeling design concepts under uncertainty based on a tradeoff space representation. We use the principles of decision making to develop a useful interpretation of a tradeoff space for decisions under uncertainty and to identify criteria useful for eliminating undesirable tradeoffs from consideration. We illustrate our approach to modeling and decision making on an example for the conceptual design of a gearbox.

INTRODUCTION

Deciding between competing design concepts is a critical step in a system development process and can, to a large extent, dictate the overall success of a project. Intuitively, design activities later in a process cannot compensate for the selection of a poor concept. Some researchers note that such early-process decisions are important because it becomes progressively more expensive to rectify errors [1-3]. Others argue, primarily based on studies of Toyota Motor Company and its suppliers, that improvements in early-phase decision making can lead to improved product quality and reduced design cycle times [4-6]. However, despite a clear motivation for improved conceptual design decisions, how to do this in a way that is both practically efficient and fundamentally sound remains an open question.

The main challenge is that a design concept is not a completely-defined solution to a problem; it is a general approach to implementing a function or system. Although the concept "diesel engine" implies certain aspects about the design solution, it leaves numerous implementation details unresolved. From this perspective, *a design concept is an abstraction of a large set of possible design implementations* [7]. Decision methods that rely on detailed engineering analysis, which requires a precise definition of a design, are poorly suited for the set-based nature of design concepts.

Consider again the diesel engine concept: How much does it cost? What is its reliability? How much power can it output? The answer to these questions is that it depends: different implementations will have different tradeoffs. Even if an engineering analysis model could be perfectly accurate, the imprecise, set-based nature of a design concept would lead to uncertainty in the results and expose a decision maker to considerable risk.

In this paper, we investigate an alternative to approaches that rely on explicit design definitions. Under our approach, one considers the set-based nature of a design concept explicitly by modeling the various tradeoffs achievable by implementing it. Thus, we describe a concept in terms of the possible consequences of selecting it rather than what it is in physical terms.

This investigation is part of a broader vision for conceptual design, in which designers compose system-level concepts from reusable tradeoff-space models of common functional components. For example, system engineers could model different configurations for a hybrid gas-electric power train quickly and confidently if they begin from concept tradeoff models of the main functional units (power storage, power conversion, etc.). This would enable them to incorporate detailed tradeoff considerations into their decision making without

requiring specialized domain knowledge about every component concept. Our focus in the current paper is on how to model and decide between concepts in isolation, leaving the composition of such models for future work.

In prior work, we have demonstrated an analogous approach for modeling design concepts under deterministic assumptions [8]. Although the strategy is similar, the approach we describe in the current paper is a significant departure due to the added complexity of decision making under uncertainty. In both works, we model a design concept in terms of the tradeoffs designers can achieve by implementing it. However, the deterministic notion of a tradeoff space is incompatible with decisions under uncertainty. In this paper, we contribute a new, general interpretation of what constitutes a tradeoff space and apply it in the context of concept modeling. Our interpretation is based on accepted decision theory and is meaningful independent of our modeling approach.

Another similarity is that in both works we generate a model for a design concept using only data about those implementations of it that a designer might choose rationally. Collectively, these implementations are called the *parameterized efficient set* and identifying them is complicated by the presence of uncertainty. The approach we use in the prior work is based on Pareto dominance, which is inappropriate for decisions under uncertainty. In this paper, we draw upon the decision theory literature for tests—called *stochastic dominance* criteria—appropriate for decisions under uncertainty. Although stochastic dominance can be difficult to evaluate in general, we identify a specialized criterion that can be useful in many engineering situations and apply it in our modeling approach.

We demonstrate our approach on the conceptual design of a gearbox. We consider three concepts for implementing the gearbox, each with different physical configurations. For each concept, we identify a parameterized efficient set and fit a predictive tradeoff model to it. We select the most preferred concept using these models in concert with decision preferences. We verify that our models provide accurate tradeoff predictions and yields the correct decision by comparing the results to results from a baseline method based on traditional engineering optimization techniques.

The remainder of this paper is organized as follows. First, we survey the problem background, including prior work on concept modeling and selection. Next, we describe the proposed approach for modeling design concepts. This is followed by a description of how to formulate decisions using the tradeoff-space models. We then demonstrate our approach and verify it relative to a baseline method. Finally, we discuss the contributions, limitations and potential future work.

PROBLEM BACKGROUND

DECIDING BETWEEN DESIGN CONCEPTS

In the context of mechanical design, conceptual design involves the transformation of requirements and objectives into a definition for the basic solution structure that designers will refine in subsequent phases of a design process [9, 10]. In the broader context of systems design, conceptual design corresponds to the definition of an abstract structure for a system in terms of its functional subunits, their configuration into a system, and requirements for the design of each subunit [2, 11-13]. In either case, conceptual design entails many tasks, including functional decomposition, concept generation, and decision making, and the outcome of a conceptual design decision is an imprecise definition of the final system.

In this paper, we consider how designers can make decisions between two or more concepts. We focus on concepts for which several prior implementations exist, such that designers have access to information about these implementations or understand the principles underlying them. This includes common functional components that can be implemented using different technologies or physical configurations. Pumps, motors, speed reducers and actuators are examples of these. We anticipate that the types of models described in this paper can form a basis for composing concepts for novel systems by composing non-novel component-level concepts, and we plan to study this in future work.

A basic premise for this paper is that designers should make decisions systematically and using methods that are sound with respect to the accepted norms of decision theory. This perspective is called Decision-Based Design (DBD) in the literature [14-16]. In particular, we build upon the decision theoretic framework of Multi-Attribute Utility Theory (MAUT) [17], which is an extension of the utility theory by von Neumann and Morgenstern [18] to the case of multiple competing objectives. Using MAUT, designers can express their preferences with regard to tradeoffs under uncertainty.

A generic decision is formulated in MAUT as

$$a^* = \arg \max_{a \in \mathcal{A}} E[u(\mathbf{x}_a)], \quad (1)$$

where $\mathcal{A}$ is the set of feasible decision alternatives, $a \in \mathcal{A}$ is a specific alternative, a^* is the most preferred alternative, $\mathbf{x}_a \in \mathbb{R}^n$ is a random vector of attributes for alternative a having distribution function F_a , $u: \mathbb{R}^n \rightarrow \mathbb{R}$ is a suitably defined utility function, and $E[\cdot]$ is the expectation operator. In this context, an *attribute*, is a measure of progress toward a decision objective [17]. For example, in a choice between a diesel and a

gasoline engine, attributes may be cost, reliability and measures of technical performance (e.g., power output, max torque).

Engineers typically formulate conceptual design decisions using engineering analysis models, which map a description of a system to the relevant decision attributes. In general, this mapping is different for each concept because different concepts can involve different operating principles. One can express conceptual design decision formulated using analysis models as

$$c^* = \arg \max_{c \in C} E[u(\mathbf{f}_c(\boldsymbol{\varphi}_c))], \quad (2)$$

where $c \in C$ is a design concept, c^* is the most preferred concept, $\boldsymbol{\varphi}_c$ is vector of variables describing the concept as required for analysis and $\mathbf{f}_c(\cdot)$ represents the engineering analysis models that relate the description to a distribution in the space of decision attributes.

A significant challenge in making conceptual design decisions is that concepts typically are defined too imprecisely to support accurate analysis. In terms of the above, a parametric description, $\boldsymbol{\varphi}_c$, might include only a fraction of the design variables required to describe a finalized design implementation. Thus, a unique $\boldsymbol{\varphi}_c$ vector might correspond to a large set of design implementations, each of which could, if designers could consider them individually, have a different expected utility. Even if designers had a perfect analysis models, they would have significant uncertainty in their evaluation of a concept.

Although designers always can reach a decision using the generic formulation of Equation (2), the imprecise nature of design concepts exposes them to considerable risk. What is more, most solution methods include no mechanism for dealing with the inherently uncertain and abstract nature of design concepts. Thus, designers who use these methods may neglect this source of uncertainty and, consequently, expose themselves to even greater risk.

PRIOR RESEARCH

A few methods from the DBD community are targeted for use on conceptual design decisions, but these have their own limitations. Chen and coauthors describe a method in which a designer models uncertainty about the final implementation of a concept explicitly using uniform probability distributions applied to the design space variables and makes decisions using a modified robust design formulation [19, 20]. Although the method is easy to apply, it constrains how designers can formalize their beliefs and preferences. Furthermore, they provide no procedure for determining the bounds of the required uniform distributions.

Wood and coauthors describe a method for abstracting the characteristics of partially-defined solutions from a database of prior design instances [21]. They perform abstraction by constructing a “design PDF” from the data and “integrating out” unneeded variables. However, their approach rests on the use of this “design PDF,” for which they provide no concrete interpretation or justification.

Mattson and Messac demonstrate an approach to selecting design concepts based on a Pareto set representation [22, 23]. They represent a concept using the Pareto set of possible implementations, which has some similarities with the approach we present here. However, they rely on dominance reasoning for making choices, which can fail to identify one concept as the most preferred (a situation called indeterminacy). Furthermore, their approach for dealing with uncertainty and risk lacks foundation in accepted decision theory.

The notion of Set-Based Design (SBD) commonly is associated with conceptual design. Under a SBD approach, designers focus on eliminating inferior implementations from a set of designs, and possibly delay a decision in favor of gathering more information [5]. Most SBD methods eliminate designs using interval- or set-based constraint propagation methods (e.g., the methods of [24-26] and the tools demonstrated in [27, 28]). Although there is evidence of analogous practices in industry [4] and that these practices are beneficial [6], SBD does not constitute a comprehensive decision method—by focusing on elimination rather than selection, SBD can lead to indeterminacy. Some research exists on extending the set-based design perspective to include eliminations based on preference information [7, 29]. Although this improves the odds of identifying a most preferred design, there are no assurances designers will avoid indeterminacy.

Several approaches to deciding between design concepts are practiced in the systems design and engineering community, the more common of which include quality function deployment [30], Pugh selection [31], and the analytic hierarchy process [32]. Hazelrigg analyzes these and other strategies, finding them each to be lacking in light of accepted results from decision theory [33]. Another systems-oriented decision approach, known as analytical target cascading, is effective at propagating design requirements down to subsystems, but assumes the system configuration is known [34, 35].

MODELING DESIGN CONCEPTS

Our strategy for modeling a design concept is to abstract information about implementations of the concept in a way that is favorable to decision making. Concepts are modeled:




Figure 1: Illustration of two types of models for computing tradeoffs: design-space-to-tradeoff-space mappings (analysis models) and mappings within the tradeoff space (tradeoff models).

- as relationships in a suitably defined tradeoff space;
- using predictive models that capture association rather than causation;
- using information about only the implementations that a designer might rationally choose; and
- in a way that is reusable across multiple decision problems.

Using the fitted models, designers can predict the most preferred tradeoff that is achievable by implementing a particular concept. They make decisions according to a two-step process of:

1. identifying the most preferred tradeoff achievable by each concept, and
2. selecting the concept that can lead to the most preferred tradeoff overall.

In prior work, we justified this general strategy and demonstrated an approach for modeling design concepts assuming designers can evaluate tradeoffs with certainty [8]. Extending the strategy to decisions formulated in MAUT is nontrivial. In this section, we explain the main characteristics of our general strategy and the extensions necessary for MAUT.

TRADEOFF SPACES

Designers can realize a particular design concept using any number of physically distinct implementations. This complicates attempts to formalize a space of design variables, and motivates us to find a modeling approach that avoids design space representations altogether.

We use the term *tradeoff space* to mean a representation of the possible tradeoffs a designer might make. A tradeoff space relates to designer objectives and is independent of the alternatives being considered. Every design implementation maps to exactly one point in a tradeoff space, and designers can compare two designs by applying their preferences to the tradeoff points for each design. Figure 1 is an illustration of the difference between using analysis models and tradeoff space models.

The tradeoff space relates closely to the concept of a decision making attribute. In the case of decisions with deterministic tradeoffs, the tradeoff space is the space of attributes. For decisions in MAUT, the tradeoff space is the space of *probability distributions* over the attribute space. This follows from the axiomatic definition MAUT as decisions between lotteries (see [17] for a development of this theory) and is evidenced by the procedures for eliciting preferences in MAUT, which involves comparisons of lotteries over attributes (as opposed to comparisons of precise values).

Strictly speaking, the notion of a “space” of probability distributions is ill-defined. However, one can specify several common distributions unambiguously by a small number of parameters. If those distributions—which includes normal, log-normal, uniform, Poisson, etc.—are good representations of the uncertainties, then designers can define a space of distributions in terms of such parameters. In this paper, we assume all uncertainties are normally distributed.

Our definition for a tradeoff space makes intuitive sense when one considers the kinds of tradeoffs that occur when uncertainty is involved. For example, when product quality is a concern it is common for designers to trade mean performance to achieve reduced variability in that same attribute. One can visualize this tradeoff in the mean-variance space of that attribute. Figure 2 is an illustration of the relationship between the attribute space and the tradeoff space for MAUT decision problems.

PREDICTIVE MODELING

There is growing recognition that, fundamentally, engineering design activities are part of an overall business enterprise, and that designers must make decisions in this context [36-39]. Some authors even advocate a decision approach based on explicit modeling and simulation of enterprise activities [40, 41]. However, we are skeptical about whether the gains would outweigh the costs associated with such a

Attribute Space

Tradeoff Space




Figure 2: Illustration of distributions in an attribute space mapping to points in a tradeoff space defined in terms of parameters of the distributions.

modeling effort. Instead, we adopt a strategy based on predictive modeling.

A predictive model is something with which users can draw inferences about the unknown value of one variable given the values of other variables [42, 43]. The power of predictive models is that they permit useful inference in cases where the underlying causal mechanism is unknown or too complex to formulate analytically. Although currently more common in the context of marketing and finance, such models also are useful for engineering design, where designers can use predictive modeling techniques to infer relationships too complex to construct directly from engineering principles. Statistical cost estimation is a common example [44-46].

Our strategy is to model a concept by fitting a predictive model to tradeoff space data about prior implementations of the concept. The fitting can be achieved using standard regression analysis or interpolation methods. The motivation for using information about prior implementations is that it will include enterprise considerations implicitly. Although this is a departure from the traditional approach of engineering analysis, the idea of generating tradeoff-space predictive models has been demonstrated previously in the contexts of power transmission [47], environmentally benign design [48], and motorized scooters [49].

To express the model and its role mathematically, let θ denote a particular tradeoff in the m -dimensional space of tradeoffs, Θ . A predictive model, $\tau(\cdot)$, is formulated to use $p < m$ dimensions of the tradeoff space to infer the remaining $m - p$ dimensions. Thus, one can express a complete vector in the tradeoff space as

$$\theta = [\theta', \tau(\theta')],$$

where θ' denotes the p -dimensional vector of tradeoff dimensions used as independent variables in the predictive model. The selection of which dimensions to

use as independent variables is arbitrary, and one should structure the model so it yields the best predictions possible.

STOCHASTIC DOMINANCE AND EFFICIENCY

Designer Rationality, Efficiency and Dominance

To model all feasible tradeoffs that correspond to a design concept would be unproductive, since a *rational designer never would choose most of these implementations*. Therefore, it is beneficial to distinguish between feasible and “good” tradeoffs. Two decision theoretic concepts—efficiency and dominance—are helpful for making this distinction.

Our strategy is to model only the tradeoffs that are *efficient*, which means that designers cannot improve in any tradeoff criteria without sacrificing another. It is common to identify the efficient tradeoffs by performing dominance tests. An alternative, A , is said to be *dominate* another, B , when decision makers can guarantee that A is preferred to B . A dominated design cannot be efficient, so one can identify the efficient set using dominance tests.

For deterministic decision problems, the set of efficient alternatives is known commonly as the Pareto set [17, 50]. Pareto sets have been used in a number of frameworks for decision making in design [51-55], including concept selection [8, 22]. However, the notion of Pareto dominance is tied inherently to deterministic decision making. Some authors have attempted to extend Pareto dominance to decisions under uncertainty (e.g., [23]), but the appropriate approach, in light of the norms of decision theory, requires an expanded interpretation of dominance.

Stochastic Dominance

To identify the efficient set for decisions formulated in MAUT, designers must appeal to the general notion of *stochastic dominance* [56, 57]. To date, research into stochastic dominance criteria is limited mostly to the economics and operations research literature; we are unaware of prior applications in the context of engineering design.

A stochastic dominance test involves comparing distribution functions that are defined over an attribute space—i.e., it is a tradeoff-space evaluation. It also requires assumptions about the mathematical structure of the corresponding utility function. This is similar to the monotonicity assumption associated with the deterministic concept of Pareto dominance, but the useful assumptions are more varied in the stochastic case owing to the greater complexity of preferences for decisions under uncertainty. Researchers have identified a number of stochastic dominance tests.