MULTI-ATTRIBUTE UTILITY ANALYSIS IN SET-BASED CONCEPTUAL DESIGN

Richard J. Malak Jr.[1]
rmalak@gatech.edu

Jason Matthew Aughenbaugh[2]
jason@arlut.utexas.edu

Christiaan J. J. Paredis[1]
chris.paredis@gatech.edu

[1] The Systems Realization Laboratory
G. W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology

[2] Signal and Information Science Laboratory
Applied Research Laboratories
The University of Texas at Austin

Abstract

During conceptual design, engineers deal with incomplete product descriptions called design concepts. Engineers must compare these concepts in order to move towards the more desirable designs. However, comparisons are difficult because a single concept associates with numerous possible final design specifications, and any meaningful comparison of concepts must consider this range of possibilities. Consequently, the performance of a concept can only be characterized imprecisely. While standard multi-attribute utility theory is an accepted framework for making preference-based decisions between precisely characterized alternatives, it does not directly accommodate the analysis of imprecisely characterized alternatives. By extending uncertainty representations to model imprecision explicitly, it is possible to apply the principles of utility theory to such problems. However, this can lead to situations of indeterminacy, meaning that the decision maker is unable to identify a single concept as the most preferred. Under a set-based perspective and approach to design, a designer can work towards a single solution systematically despite indecision arising from imprecise characterizations of design concepts. Existing work in set-based design primarily focuses on feasibility conditions and single-attribute objectives, which are insufficient for most design problems. In this article, we combine the framework of multi-attribute utility theory, the perspective of set-based design, and the explicit mathematical representation of imprecision into a single approach to conceptual design. Each of the component theories are discussed, and their combined application developed. The approach is illustrated using the conceptual design of a fixed-ratio power transmission as an example. Additionally, important directions for future research are identified, with a particular focus on the process of modeling abstract design concepts.

1 Introduction

According to the paradigm of systematic design [27], the conceptual design phase takes as an input a list of requirements and objectives, and yields as an output the principle solution structures to be pursued in embodiment design. The first task of conceptual design is to distill the problem down to its core, including identifying what functions the design must perform and how these functions interact at a high level through transfers of energy, mass, and information. For example, the functions of a photocopying machine might include “acquire source image,” “move paper,” “mark a piece of paper,” and “interact with user.” The development of this function structure is not the focus of this paper, but rather the focus is on the next step in conceptual design.

Once a function structure is defined, designers seek to enumerate and then compare possible physical implementations, known as working principles or *concepts*, for each function. For example, three working principles for the function “mark a piece of paper” could be “deposit material by friction” (e.g., a pencil), “melt material onto paper” (e.g., laser jet printing), or “burn away material” (e.g., scorching the paper with a laser). Since in general there are multiple functions, each with multiple working principles, they can be combined into an overall product in many different ways, each combination forming a possible *solution concept* for the final design. In traditional systematic design, a single principal solution concept must be chosen for continued development in the embodiment design phase.

The evaluation and comparison of concepts are inherently challenging tasks. A concept is not a highly detailed product, but rather a general approach to implementing a function or system. In essence, *each design concept represents a large set of possible final designs* all based on the same concept. For example, the concept of “melt material onto paper” does not include details such as what material to melt, how much of it to melt, or how to guide the material into an appropriate mark on the paper. Because concepts are not detailed descriptions but rather are sets of alternatives defined by incomplete specifications, it is challenging to make rigorous comparisons between them. How can one decide whether it is better to deposit material by friction or melt material onto paper if one does not know exactly how each of these concepts will actually be implemented? More generally, *given a set of alternatives, how can decisions be made at the general level of conceptual design when specific design details are unknown?*

Part of the answer is to create computer aided engineering and design tools that can model the abstract characteristics of design concepts rather than the fully detailed design descriptions that most existing tools require. Before such tools for conceptual design can be developed, the fundamentally imprecise character of conceptual design must be appreciated and an effective means for comparing incomplete product descriptions must be developed. In this paper, research from several domains is brought together in order to form a framework for evaluation and comparison of alternative concepts during conceptual design. Specifically, we present a set-based approach to conceptual design using multi-attribute utility theory and imprecise probabilities.

2 Imprecision in Conceptual Design

In order to develop methods for making decisions during conceptual design, one must recognize the nature of the uncertainty that exists during this stage of the design process. In general, two aspects of uncertainty can be identified: variability and imprecision. While some authors doubt the philosophical distinction between variability and imprecision, such distinctions are useful in practice [5].

Variability, also called aleatory uncertainty (from the Latin *aleator* = dice thrower), is naturally random behavior in a physical process or property [26]. It is also known as objective uncertainty [12] and irreducible uncertainty [11]. Examples include manufacturing error, errors in communication systems, and radioactive decay. Inherent variability is best represented in stochastic terms, e.g., by a probability density function. Consequently, variability is compatible with decision approaches based on classical probability theory and expected utility maximization [35], which is the focus of most engineering research on decision making.

Imprecision, on the other hand, is due to a lack of knowledge or information [29] and sometimes is called epistemic uncertainty (from the Greek *episteme* = knowledge), reducible uncertainty [11] or subjective uncertainty [12]. Imprecision can be represented in terms of intervals if one wishes to avoid overstating what one knows to be true [19, 23]. Consequently, such representations are not immediately compatible with most engineering research on decision making. This presents a potential problem for conceptual design, where a significant proportion of the uncertainty comes from imprecision, and motivates an investigation into new approaches to decision making.

We return later to the problem of making decisions in conceptual design under imprecision. For the remainder of this section, we examine the various sources of imprecision in conceptual design. These include the structure of the design process, scarcity of relevant data, expert opinion, and the use of abstract models.

2.1 Concepts are Imprecise Design Alternatives

One can think of the exploration of design concepts as a breadth-first search of an alternative space in which the decision maker searches across high-level concepts rather than down to detailed descriptions. This approach leads to a sequential ordering of the design process from general to specific. Essentially, the guiding principle is that there is no reason to consider the detailed implementation of a specific alternative (e.g., marking paper by friction) if you can decide at a more general level that a different alternative (e.g., marking paper by melting material) is better for the given design problem (e.g., photocopier design).

This process and the inherent imprecision of design concepts are illustrated using a simple design problem in which a decision maker wishes to design a vehicle that can transport a person. We assume that through some creative ideation process (which is not the focus of this paper), the decision maker arrives at two possible vehicle concepts: a car and a bike.

Clearly, neither "car" nor "bike" is a fully detailed design specification for a vehicle. For example, both concepts must have a source of power, such as a gas engine, a diesel engine, or an electric motor. The concept "car" contains the more specific design sub-concepts of "gas car," "diesel car," and "electric car," as shown in Figure 1. Each of these sub-concepts has different characteristics. Even within a sub-concept, the characteristics of specific designs can vary significantly. For example, commercially available gasoline engines cover a large range of power and fuel efficiency. Consequently, the horsepower of the concept "gas car" is not a single number, but a set of horse powers that correspond to every member of the set of possible implementations of the "gas car" concept.

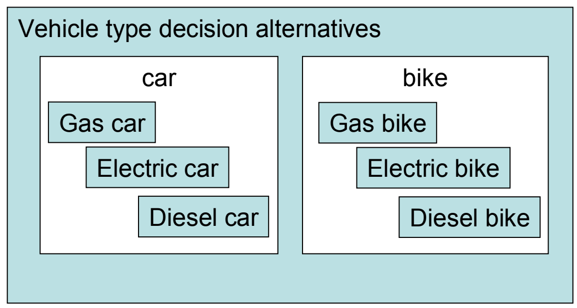

Figure 1. Sets of design alternatives

At the end of the design process, a single detailed product specification will result. However, when a decision maker is comparing design concepts, he or she lacks knowledge about the final design since there are many decisions yet to be made beyond the one the decision maker is working on. Because of this inherent lack of knowledge, the decision maker generally can only characterize the performance of concepts imprecisely.

The only way to avoid this type of imprecision early in the design process would be to completely reformulate the design process as a single optimization problem. In such a formulation, the decision maker would need to fully enumerate all possible design alternatives. For all but trivial design problems, the combinatorial explosion of fully detailed designs makes such an approach impractical. For example, in order to compare all possible car designs, a designer would have to consider the full range of possible specifications for everything from the engine and transmission to the brakes and steering. It is completely impractical to consider this is huge set of alternatives, so designers generally follow a sequential, general-to-specific systematic design approach in practice. However, such sequential approaches lead to an inherent imprecision in the decision alternatives that a designer must choose between during conceptual design.

2.2 Analysis Models Yield Imprecise Predictions

In order to compare design alternatives, engineers frequently use behavioral models to predict the performance of the alternatives in terms of attributes that are important to them such as physical behavior, cost, and reliability. Like all models, these are only abstractions and, consequently, their predictions are imprecise reflections of reality.

For example, although the laws of physics are known very precisely, one often makes significant assumptions when applying them to complex geometries, or one omits certain known—but less significant—physical phenomena from the model to reduce the complexity. For example, a model for an internal combustion engine is often abstracted into an algebraic relationship between engine speed and torque. The detailed physical phenomena (including airflow, gas-mixture combustion, friction, and inertia) are reduced into one simple algebraic relationship. This simple relationship is an idealization that may contain significant error—there are unknown or unmodeled relationships between a variety of parameters that play a role in the engine performance, such as air density, acceleration, or engine temperature. The lack of knowledge of the influence of these parameters on engine performance results in imprecision in the model's predictions. Since there is no probability distribution associated with such modeling and systematic errors, one cannot express the likelihood of occurrence for a particular error but can at best bound the size of the error, in which case the errors should be represented in terms of interval-based uncertainty.

Models introduce imprecision into design decisions throughout a design process, but the effect is most profound during conceptual design. In the preceding example, it was suggested that a simple algebraic relationship could abstractly model the complex physical relationship between engine torque and speed. This seems reasonable when the basic structure and configuration of the engine is already determined, but other approaches may be required for modeling the general concepts of "gas engine," "diesel engine," and "electric motor." In this article, we deal only with concepts for which it is reasonable to model using abstract physical relationships. The problem of characterizing general concepts remains as a topic for future investigation.

2.3 Data Scarcity Leads to Imprecise Statistical Models

Engineers often gather statistical data about environmental or other factors to support design decisions. Such quantitative data gives an illusion of being well-characterized, but actually it is inherently imprecise. The two most common interpretations of probability are the frequentist and subjective. The traditional formalizations under either interpretation forces all probabilities to be precise, but in actuality both types of probabilities can involve significant imprecision.

Under a *frequentist* interpretation, a probability represents the ratio of times that one outcome occurs compared to the total number of outcomes in a series of identical, repeatable, and possibly random trials. In engineering design, events are not always repeatable. Even assuming some events are essentially repeatable and data can be collected, only a limited amount of data is gathered in practice. Although in theory the relative sample frequency approaches the true relative frequency as the sample size goes to infinity, an infinite sample size is impossible to acquire in practice. Consequently, engineers will always face imprecision in their characterizations of the frequentist probabilities.

During conceptual design, it is especially challenging to collect a large sample of relevant data. For example, any prototypes of a concept will be a significant abstraction from the detailed designs that the concept represents. Even if resources were available for significant testing, extrapolation from prototype data to the abstract concept involve some degree of imprecision.

Imprecision also exists under a *subjective* interpretation of probability. Subjective probabilities are an expression of belief based on an individual's willingness to bet [10]. The process of eliciting and assessing an individual's beliefs, or willingness to bet, is resource intensive. Even assuming that precise beliefs—and hence precise probabilities—exist, it will often be impractical to fully characterize them due to constraints such as bounded rationality, time, and computational ability [16, 36, 39]. Consequently, only a partial—and therefore imprecise—characterization of subjective probabilities is available in practice .

2.4 Expert Opinions and Qualitative Judgments are Imprecise

A significant source of information in engineering design are experts who use their knowledge and experience to form judgments, beliefs, and estimates [6, 9]. Such information commonly is imprecise. A true expert understands the limitations of his or her knowledge of a particular subject matter and issues judgments accordingly. They often state their conclusions in vague linguistic terms, such as such as "unlikely," "large," or "poor," and in other cases they give imprecise ranges for a quantity, such as "the shaft is between 0.50 and 0.51 meters long" or "the reliability is at least 0.99." Imprecision also can arise in cases of conflicting expert opinion. After checking three weather services and finding temperature forecasts of 30, 31 and 33 degrees Celsius, one might conclude that the interval [30, 33] °C is a reasonable characterization of what is known about tomorrow's temperature.

3 Modeling Design Concepts

As argued in the previous section, a design concept is an imprecise definition of a product or system. The idea is to identify a promising concept by reasoning about its general properties rather than diving into specific implementation details. If a designer can conclude that a concept definitely is inferior to another concept, then no additional resources need to be expended to further explore and refine the inferior one.

In order to compare design concepts, engineers must be able to evaluate them—that is, to determine their attributes. Before attempting to predict the attributes of a concept, designers must understand what the concept includes and what it does not. Although the attributes of a design concept are imprecise, the definition of a single concept should not be vague—it should be clear whether a particular design alternative belongs to a given concept or not, even if it is unclear whether it is a good way to implement the concept. Consequently, engineers must develop formal representations for design concepts.

3.1 Approaches for Representing Design Concepts

Although it is possible for designers to reason about concepts informally using lingual definitions and qualitative ratings, the focus of this article is on comparisons that are quantitative and formal. As such, it is necessary to have a mathematical representation of a design concept.

In embodiment and detailed design, design alternatives are represented by specifying their structural characteristics, e.g., shape, size, or topological structure. At the conceptual design stage, however, the structure of a particular concept may still be too vague for structural parameters to be defined in a meaningful fashion. It is therefore common that one characterize design concepts by certain desired performance characteristics. For instance, when considering architectural concepts for a car, one may define the engine by specifying its power rating. Such representations are unambiguous, but imprecise. When defining a concept with a 175 hp engine, there is no ambiguity in the meaning of that specification, but there is imprecision in the sense that there is a large set of engines that all satisfy the performance characterization of 175 hp. The imprecision manifests itself explicitly when considering other characteristics of this engine concept; the cost or efficiency of this set of engines is almost certainly imprecise and can only be characterized by an interval of possible cost or efficiency values.

Identifying interval bounds for each attribute of a design concept is only the most basic representation. Although this approach is straightforward and often easy to achieve based on expert judgment, it can neglect dependencies that occur between the attributes. For instance, there often is a positive dependence between mass and cost because increasing the size of parts increases the material cost. In such cases, a bounding region with a more complex shape may be beneficial.

Although identifying relationships between design attributes can be difficult, it sometimes is possible to consult an abstract model that predicts attribute values based on design parameters. The model would neglect many design details and involve only a handful of design parameters compared to the quantity present in the final detailed design specification. Such models are commonplace in engineering design, and typically are algebraic in nature and ignore secondary effects (e.g., the speed-torque curve of an internal combustion engine). Using such a model, one can derive bounds on the concept attributes by bounding the model's inputs.

3.2 Identifying the Scope of a Design Concept

Identifying appropriate boundaries of a design concept is an important and often challenging task. Defining a design concept too broadly can result in a very large range of values for its defining attributes, which, as explained later in the article, can impede decision making. Conversely, too narrow a definition sometimes can mean designers are including more detail at the conceptual level than is required, thus wasting resources and effort.

One approach to defining a design concept is to consult experts who can rely on their tacit understanding of a domain to identify an appropriate range of attribute values or model inputs. This approach is common and generally useful for problems that fit within an expert's past experiences.

Another approach is to begin with a broad definition and narrow it using preference-based reasoning. This general process is illustrated in the demonstration of Section 7. Essentially, one often can refine a design concept to rule out instances that are poor solutions relative to others. For example, one likely would exclude from a "175 hp engine" design concept for a passenger car engine any instances that are unduly inefficient or costly, despite the fact that they have the correct power rating.

A third approach is to draw upon a database of prior solutions to similar problems. Using data mining and artificial intelligence techniques, one might generalize from the individuals to arrive at a characterization of a corresponding design concept. In such cases, it is necessary to distinguish between prior solutions that are part of the design concept of interest and those that are not. However, one also can turn the process around and use approaches such as data clustering to identify groups of individuals that are similar and therefore appropriate to consider as examples of the same general design concept.

In this article, we focus on the issue of making decisions involving imprecise design concepts. The determination of the best way to derive abstract models for design concepts is an equally important and parallel issue that is beyond the scope of this article. Ultimately, both issues will need to be resolved, but we start with the methods for organizing and making design decisions using imprecise concepts with the goal that these methods suggest the types of models that are needed; essentially, before creating models for design concepts, it is necessary to identify how those models will be used. Consequently, in this article we will assume that it is reasonable to use interval bounds on attribute values or model inputs to represent design concepts. The rigorous identification of these interval bounds is left for future work.

4 Decision Making in Conceptual Design

Products are generally evaluated according to multiple criteria. For example, a designer or customer of a photocopier may care about its cost, reliability, speed, accuracy, and size. Attributes such as reliability often are highly uncertain during conceptual design, so designers need a method that allows them to express their preferences under conditions of limited knowledge. Furthermore, with multiple attributes of interest, designers need a method for comparing designs across all of these attributes. A well-established normative approach to such comparisons is given in multi-attribute utility theory [18, 35].

4.1 Utility Theory as a Normative Decision Model

4.1.1 Classical Multi-Attribute Utility Theory

Under multi-attribute utility theory, a decision maker formulates his or her preferences in terms of a scalar function, called a utility function, over the domain of attribute values. This function defines the decision maker's preferred attribute values and the tradeoffs between different attributes. Uncertainty is modeled probabilistically as a distribution over the possible attribute values for a given action on the part of a decision maker. For a properly-constructed utility function, the most preferred action of the decision maker corresponds to the action that maximizes expected utility.

In the context of a design decision, actions represent the different design alternatives from which a designer can choose. Attributes are aspects of the design about which a designer cares—such as cost, reliability, or speed—and their exact values for a given design alternative might not be known with certainty. Under utility theory, designers model these uncertainties probabilistically and aggregate the possible outcomes by taking an expectation over the uncertain domain.

Mathematically, one can state a design decision according to multi-attribute decision theory as:

$$d^* = \arg\max_{d_i \in D} E_X \left[U\left(x_1(d_i), x_2(d_i), \ldots, x_N(d_i) \right) \middle| d_i \right]$$

where d_i is a particular design alternative, D is the notional set of all design alternatives, d^* is the most preferred design alternative, $x_j(d_i)$ is one of the N (uncertain) attributes corresponding to the specified design alternative, $U(\cdot)$ maps attribute values to a scalar value and $E_X[\cdot \mid d_i]$ is the expectation taken over attribute values given a particular design alternative. Note that there is no imprecision in this formulation; uncertainty about the attributes of a given decision alternative is represented using a precise probability distribution.

4.1.2 Representing Imprecision in Decision Problems

An underlying assumption of multi-attribute utility theory is that decision makers can model all uncertainty using probability theory. However, as noted in Section 2, significant imprecision exists in conceptual design and this imprecision poses challenges from a decision-making perspective. This stems, in part, from the representations with which designers can represent imprecision.

One way decision makers can represent imprecision is to use intervals defined by the upper and lower bounds on the value of a quantity. They can combine imprecise information represented this way using the rules of interval arithmetic [22] and can compare decision alternatives based on the upper and lower bounds on utility. However, this approach is effective only when *all* uncertainty is represented using intervals—standard interval analysis methods provide no means by which to incorporate probabilistic information.

To overcome this limitation, Ferson and Donald [13] have developed a formalism, called probability bounds analysis (PBA), which is based on the more general theory of imprecise probabilities [36], that can capture and compute with both inherent variability and imprecision. In PBA, one expresses uncertainty in a structure called a probability-box, or p-box. Essentially, a p-box is an imprecise cumulative distribution function (CDF). Upper and lower CDF curves are used to bound the set of distributions that the decision maker believes are consistent with the available information. For an example of the use of PBA in engineering design, see [5].Perhaps also cite Aughenbaugh et al 2006 (detc), and maybe Aughenbaugh and paredis SAE, and Schlosser/paredis SAE).

Designers can compute expected utilities using PBA. However, in the general case taking an expectation involving a p-box results in upper and lower bounds on the expected value rather than a precise value [4]. This is because each distribution inside the p-box has a particular expected value associated with it, and in general these values can be different. For example, let the p-box for a quantity X be defined by the bounding normal distributions with mean 5 and mean 6, both with standard deviation of one. Consider the first case now, $X \sim N(5,1)$ In this case, $E[X]=5$. However, the other bound of $X \sim N(6,1)$ yields $E[X]=6$. Consequently, all that can be said about the expected value of X is that it lies in the interval $[5,6]$. Although designers often can use this type of information to make decisions that are consistent with multi-attribute utility theory, there are cases that pose more of a challenge.

4.2 Challenges in Decision Making under Imprecision

An important consequence of imprecision is that it can result in *indeterminacy*: based on the available information, one may not be able to determine which decision alternative is most preferred. Strictly speaking, indeterminacy can occur when no imprecision is present. However, it is much more pronounced in cases of imprecision. For any decision in which utility theory is used to reflect preference, indeterminacy exists between two alternatives, A and B, whenever the expected utility of two alternatives is equal [35]. When imprecision exists, the expected utilities are not known precisely and become intervals, as shown in Figure 2, and indeterminacy exists whenever the expected utilities *might* be equal.

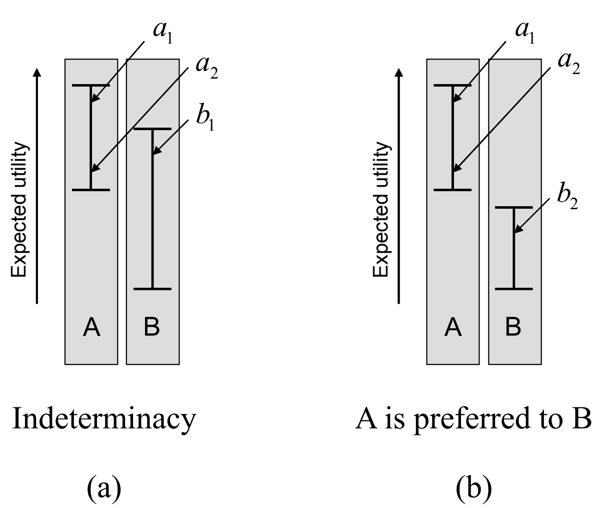

Figure 2: Two examples of comparing decision alternatives with imprecise expected utilities. Case (a) is indeterminate because of overlapping expected utility bounds. In case (b), a clear preference exists for alternative A.

For example, consider the intervals of expected utility for two alternatives (A and B) shown in Figure 2 (a). In this example, the intervals overlap. Since the true expected utility of B can lie anywhere in the given interval, the point labeled b_1 is possible. Similarly, both a_1 and a_2 are possible true values for the expected utility of A. Notice that a_1 is greater than b_1, but a_2 is less than b_1. Consequently, the available evidence leads to *indeterminacy*; the decision maker cannot determine which alternative is the most preferred, nor can the decision maker determine that he or she is definitely indifferent (because it is not necessarily true that $E[A] = E[B]$).

4.3 Common Approaches for Decision Making under Imprecision

The following is a summary of several commonly-used approaches to handling imprecision in decision problems. Each approach involves assumptions that collapse an imprecise range into a point value or presume knowledge about the relative likelihood of values within the range. Although each approach can at times be useful when making decisions, there is some question as to whether any of them can serve as a good general approach for making conceptual design decisions. An alternate approach, set-based design, is the subject of the following section.

4.3.1 Best Precise Estimate

Engineers commonly resolve imprecision by replacing an imprecise range over a quantity with their best estimate of a precise value. For example, an engineer might begin with knowledge that an external load will be between 10 and 12 kN, and assume a value of 12 kN (presumably a worst-case) in order to evaluate a precise expected utility. Once all imprecision is eliminated by making similar point assumptions, engineers can search for the decision alternative that maximizes the expected utility (under the prevailing assumptions) according to traditional multi-attribute utility theory.

Although suitable when imprecise ranges are small and the solution to the decision problem is insensitive to one's choice of precise value, general use of this approach for conceptual design decisions is questionable. Conceptual design involves high levels of imprecision, particularly with regard to the concepts themselves. Using this approach to resolve the imprecision in a design concept is tantamount to representing the entire concept using a single instance of it. Consequently, whether one design concept appears preferable to another can depend on which instances of each concept are chosen for comparison.

4.3.2 Best Estimate Probability Distribution

Rather than assuming single points, engineers often model imprecise quantities using probability distributions. The use of best-estimate probability distributions to represent uncertainty is clearly acceptable under some circumstances. For example, large data histories may be available for a quantity such as ambient temperature. Engineers can then model the ambient temperature as a random variable with a relative-frequency-based probability distribution.

Even in circumstances where data is sparse, designers often elect to model the resulting imprecision using precise probability distributions in order to use classical multi-attribute utility theory. Evidence shows that this can lead to worse design decisions compared to an approach that considers imprecision explicitly [5]. Thus, it may be advisable to avoid using a best estimate probability distribution when a significant amount of new information is likely to become available before a final decision is required. The use of precise probabilities puts the decision maker in a Catch-22 situation: in order to make a good choice about which concept to move forward with, the designer needs to have information about the expected performance of all the alternatives so that a precise probability distribution across all the design alternatives can be defined, but obtaining this information would require further study of all the alternatives, which is exactly what the decision maker is trying to avoid by making the current conceptual design decision. A set-based approach based on imprecise probabilities would allow the decision maker to eliminate some of the alternatives without having to introduce unjustified assumptions about the performance of individual alternatives.

4.3.3 Post-computation Heuristic Approaches

Rather than making assumptions that eliminate imprecision from the computation of expected utility for design concepts, one can resolve indeterminacy by applying heuristics to the imprecise expected utilities. There exist several such heuristics that can apply to conceptual design. For example, designers using a Γ-maximin strategy select the decision alternative that has the greatest lower bound on expected utility [7]. Whereas Γ-maximin is a "pessimistic" strategy, one conversely might select the concept with the greatest upper expected utility. Another strategy, known as the Hurwicz criterion, is an attempt to balance these two extremes using a weighting function [2].

Ultimately, these approaches are similar to those described above in that they aggregate an imprecise expected utility into a precise value to ensure decidability. From a procedural standpoint, the main difference is that aggregating assumptions are instituted on the results of imprecise calculations rather than on the parameters prior to computation. By applying these methods after the computations, abstract preferences such as conservativeness or optimism can be applied to the final outputs, tying them more closely to the problem being considered. Attempting to be conservative before computation requires the decision maker to understand the entire propagation of the uncertainty through the system. For example, a conservative estimate of temperature (say low) in one application could be an optimistic estimate of temperature in another application (in which a high estimate is conservative). By delaying the introduction of assumptions until after computation, the validity of the assumptions can be assessed in a more developed and appropriate context. However, these decision policies are still somewhat arbitrary in nature, and it seems unlikely that they could be applied consistently over multiple decisions, although this remains a topic for future research. As we explain in the remainder of the article, it may be possible to avoid any such arbitrary assumptions by adopting a set-based approach to design.

4.4 Need for A New Approach

Multi-attribute utility theory is a sound mathematical framework for defining and evaluating tradeoffs among multiple decision criteria and can be extended to incorporate imprecise information using the PBA formalism. However, it can lead to situations of indeterminacy when imprecision is present, as explained in Section 4.2.

By definition, indeterminacy means that there is no rational basis (given the available information) for choosing one decision over another. Common approaches for dealing with indeterminacy, such as those discussed in the preceding, involve arbitrary assumptions that serve to reduce an imprecise range of expected utility into a precise value. By assuming away the imprecision, these approaches essentially *ignore* rather than *deal with* the inherent imprecision of conceptual design.

It is not always necessary to invoke assumptions in order to resolve indeterminacy. Often it is possible for designers to support decisions rationally by reducing imprecision without eliminating it entirely. One can accomplish this by adopting a decision approach that focuses on information gathering and rational, preference-based reasoning. The next section describes such an approach, called set-based design.

5 Set-based Design

In this paper, the phrase *set-based design* is used to refer to a general approach to engineering design in which engineers reason abstractly about sets of design alternatives. The basic premise of set-based design is that it can be advantageous for engineers to delay commitment to a particular design alternative in favor of gathering information about the problem. The general effect of information gathering is to reduce imprecision to levels at which indeterminacy is resolved and designers select a single most preferred alternative. Ideas relating to set-based design exist throughout the design literature, but the focus of study has been primarily on the topics of set-valued constraints and set-based concurrent engineering. These are reviewed later in this section, along with other related literature. The main contribution in this article—the extension of set-based design to incorporate preference-based inferences—is described in Section 6. This section is an introduction to the general paradigm and existing work in set-based design.

5.1 Adopting a Set-based Approach in Conceptual Design

Conceptual design is well-suited for a set-based approach. Design concepts essentially define a set of design alternatives. Consider again the photocopy machine example from the introduction. The “melt material onto paper” concept is an abstraction that includes any number of specific implementations, each of which can involve a different material to be melted and different mechanisms for melting it. Under a set-based approach, engineers can abstract many such design alternatives into a set associated with the concept and reason about the concept based on the properties of the set. The demonstration problem of Section 7 illustrates one procedure by which designers can accomplish this.

The sets associated with concepts can be quite broad, which is a consequence of the imprecise nature of design concepts. Attributes upon which engineers base decisions—such as mass, cost, or reliability—are typically impossible to determine precisely for a concept and, as described earlier, engineers can reach a state of indeterminacy when attempting to select a single concept under conditions of imprecision. Consequently, the emphasis for set-based design approaches shifts from the typical strategy of *selecting the best* concept to that of *eliminating the inferior* ones. As described later in this section, it is clear from the literature that it is possible to eliminate some alternatives even when identifying a clear winner is not. By eliminating inferior concepts, engineers can focus their limited resources on more promising alternatives. Effectively, engineers delay commitment to a single concept while they gather information about the problem. This also may involve refining the concepts themselves by eliminating inferior members from the set associated with a concept or dividing a concept into useful sub-concepts. This process is similar to how designers traditionally think about conceptual design.

5.2 Prior Research on Set-based Design

5.2.1 Inherent Imprecision and Set-based Inferences

One of the first examples of set-based design comes from Ward's Ph.D. dissertation, in which he proposes that reasoning with intervals (a special case of a set) is natural given the imprecision inherent in design [37]. For example, he notes that a final design specification actually corresponds to a set of physical artifacts that are manufactured from the specification. Consequently, he argues, to evaluate a design specification one must consider not just the nominal or ideal artifact but the set of all artifacts to which the specification can lead. Other examples of imprecision he considers include operating conditions (e.g., ranges of ambient temperature) and desired product attributes (e.g., a range of torque a motor must deliver).

Ward identifies the semantics of intervals as being important when reasoning about them and develops methods for drawing inferences in various cases. He demonstrates his ideas in the context of catalog-based design, using interval-based inferences to search the configuration space for feasible designs efficiently. In subsequent work, Finch and Ward extend the inference methods to operate with more general set-based representations [14]. However, neither the interval- nor set-based inference methods take designer tradeoff and risk preferences into account—they focus only on feasibility.

We are aware of relatively few works on set-based methods and applications in engineering design. Parunak and coauthors develop a software tool that applies set-based methods to enable collaboration among distributed designers [30]. The tool is an application of previously developed set-based ideas. Nahm and Ishikawa describe a method for applying set-based concepts using fuzzy arithmetic, design of experiments and robust design [24]. They describe methods for representing, mapping between and reducing different variable spaces (e.g., design variable space, performance variable space). Panchal and coauthors adopt set-based design ideas in their work on decentralized decision problems [28]. They describe an interval-based method by which two designers can converge to a game-theoretic equilibrium between their objectives.

The general notion of delaying decisions in a design process has also received attention in the literature. For example, Simpson and coauthors argue that it is important to maintain design freedom while building knowledge about a design [33]. Although set-based design approaches do not call for designers to delay decisions altogether, they do allow designers to maintain design freedom while building knowledge. Designers using a set-based approach do make decisions about eliminating inferior alternatives, but delay assumptions that would cause them to commit to a single alternative.

5.2.2 Success of Set-based Design Approaches in Industry

Because set-based design approaches can involve the development of multiple alternatives in parallel for longer than in other approaches, it might seem inefficient relative to an approach in which engineers select only one concept and move on. However, studies of industry suggest that set-based approaches can be advantageous, offering benefits in terms of reductions in costly design iterations, improved design quality and reduced time-to-market.

The studies of industry focus mostly on Toyota Motor Company and their suppliers, largely due to Toyota's success as an industry leader in product quality and design cycle time. In fact, Ward and coauthors argue that a set-based approach to design is a key part of the explanation for Toyota's success [38]. Sobek and coauthors examine the product development practices at Toyota from a managerial perspective and found them to be distinctly set-based [34]. They note that Toyota generates more concepts and makes a final commitment to a single concept much later than its competitors. Liker and coauthors survey hundreds of parts suppliers in the U.S. and Japanese automotive industry looking for evidence of set-based design practices [20]. They find that some companies do engage in collaboration using informal set-based approaches and identify several factors that correlate with the use of set-based approaches.

Despite providing support for the notion of set-based design, the studies of industry provide minimal guidance on how to actually perform it.

5.2.3 Managerial Methods: Set-based Concurrent Engineering

As with any general approach to design, set-based design lends itself to treatment at a managerial level. In their investigation of product development at Toyota Motor Company, Sobek and coauthors have found a relationship between set-based design and concurrent engineering [34]. They describe three main principles of what they term *set-based concurrent engineering*: (1) map the design space, (2) integrate by intersection, and (3) establish feasibility before commitment. These principles convey the importance of gathering sufficient information prior to committing to a single alternative and suggest that collaborating engineers can integrate different pieces of information by performing set-intersection operations. These form reasonable guidelines for managers, but provide little guidance to engineers on how to accomplish these objectives.

In subsequent work, Ford and Sobek investigate the impact of how long designers maintain sets of alternatives [15]. Using a computational model of a design process, they find that converging to a single design alternative too quickly or too slowly degrades project value. Although their results are from a single computational study, they demonstrate the existence of managerial tradeoffs for set-based design and indicate the need for further study.

5.3 Combining Utility Theory with Set-based Design

A key limitation of current set-based design approaches is that they lack a general means for incorporating preference information, particularly to evaluate multi-attribute tradeoffs. For example, the system for drawing interval-based inferences demonstrated by Ward [37] relies primarily on establishing feasibility with respect to imprecise targets (e.g., a mixer must stir a batter at 500 to 1000 RPM and handle loads of 25 to 75 N-m) without permitting tradeoffs among the targets. A more general approach is needed in which designers can express their preferences for tradeoffs across multiple attributes.

We propose that it is possible to overcome the limitations of multi-attribute utility theory and set-based design by combining the two approaches. Each approach has properties that make it desirable for making conceptual design decisions: utility theory allows one to express preferences for tradeoffs across multiple attributes and set-based design allows one to make progress in a design process despite imprecision and without having to invoke strong assumptions. It is possible to combine the two approaches by replacing the feasibility-oriented inferences from prior demonstrations of set-based design with inferences based on utility theory and the notions of domination described in the next section. Thus, designers can make sound, multi-attribute tradeoffs while dealing with imprecision directly—the strength of each approach compensates for the shortcomings of the other.

6 Utility-based Decisions in Set-based Design

In making set-based design decisions using utility theory concepts, one focuses on the question: *will I ever choose alternative X?* This reflects the perspective of a set-based design approach to eliminate inferior designs. There are several inferences one can draw about a set of alternatives using utility theory. This section is a brief summary of some common inference mechanisms and how they apply to conceptual design. For further information, see [31].

6.1 Domination Criteria

6.1.1 Interval Dominance

When imprecision in parameters is modeled using intervals, the calculation of expected utility will generally result in an interval of expected utility. An example of this was shown in Figure 2(a), where the intervals are overlapping and a situation of indeterminacy exists. Figure 2(b) is an illustration of a case in which the intervals do not overlap and one can reach a unique and rational decision in favor of alternative A. Intuitively, it does not matter in this case where in the given interval the true expected utility of A falls—it always will be greater than the expected utility of B, which can be no greater than its upper bound. This illustrates a situation referred to as *interval* *dominance*, one of the most basic domination conditions [40] (also called strong dominance [21]).

To evaluate interval dominance in the case of more than two alternatives, one must compare each pair of non-dominated alternatives; once an alternative is dominated by any other, it can be removed from consideration. Figure 3 is a depiction of several decision alternatives and their associated intervals of expected utility. One begins evaluating interval dominance by comparing each pair in succession—A with B, A with C, etc. In this case, alternative A dominated alternative D and one therefore can eliminate D. Because of this, one need not consider alternative D in subsequent comparisons (e.g., there is no need to compare B with D, C with D, etc.). The result of applying is a set of alternatives whose intervals of expected utility all share some region of overlap.

Interval dominance is an easy criterion to apply, but does not capture the complete relationship between alternatives. It sometimes is possible for one alternative to dominate another even when their expected utility intervals overlap. To draw such inferences, one requires a more advanced criterion that considers relationships that exist between the intervals.

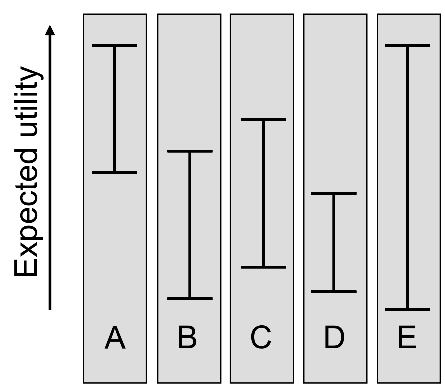

Figure 3: Many overlapping intervals of expected utlity.

6.1.2 Accounting for Shared Uncertainty

In design, there are often uncertain conditions that influence the performance of all decision alternatives in a similar fashion, factors that we define as *shared uncertainty*. When uncertainty is shared among possible actions, it means that a particular future condition or event is independent of the current decision. For example, ambient temperature is independent of the alternatives chosen—all potential final designs will have to operate over the same, but unknown, range of temperatures, such as shown in Figure 4.

As an example of uncertainty that is not shared, consider the sequential decisions of designing first a car engine and then the drive shaft. When designing the engine, the exact design of the drive shaft is unknown. However, this uncertainty is not shared by all engine alternatives, because the final design of the drive shaft will depend on the chosen engine design; the drive shaft must meet different performance requirements depending on the power of the engine, for example.

For shared uncertainties (such as temperature), the performance of alternatives should be compared assuming they are operating under the same conditions (e.g., the same temperature). A similar argument favors paired statistical testing over pooled statistical testing to remove shared systematic errors (e.g., paired t-testing). One approach that considers shared uncertainty is the *maximality* criterion for elimination [36].

Maximality is a stricter criterion than interval dominance, meaning that in general it leads to the elimination of at least as many alternatives. Maximality eliminates alternatives that are dominated at all values of the uncertain parameter by any individual other alternative. In Figure 4, alternative A dominates alternative B, because at all possible temperatures, alternative A has a higher expected utility than alternative B. Similarly, C dominates D. Thus despite the overlap in the intervals, additional eliminations can be made.

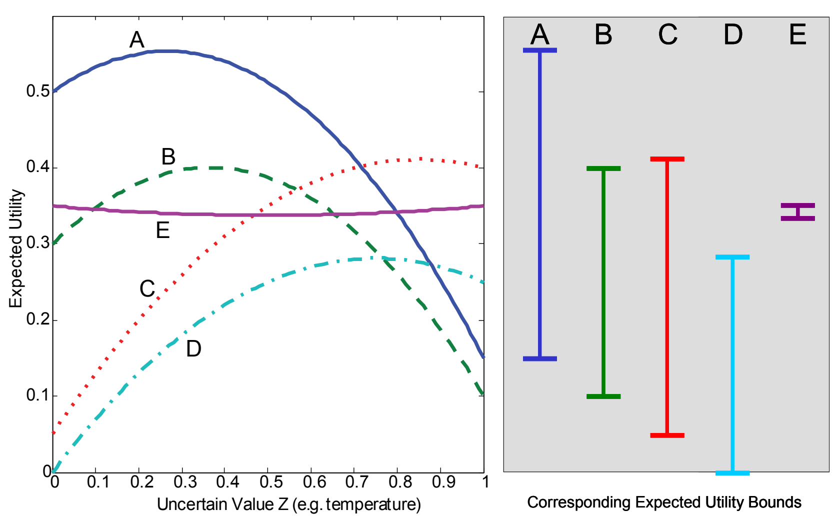

Figure 4: Expected utility varying across the range of an uncertain parameter, Z, for several design alternatives (A, B, C, D, E).

Mathematically, for decision alternatives A and B, the maximality criterion is equivalent to

$$\min_{\mathbf{y}\in\mathbf{Y}} E\left[U\left(\mathrm{A},\mathbf{y}\right)-U\left(\mathrm{B},\mathbf{y}\right)\right]>0 \Rightarrow \mathrm{A}\succ\mathrm{B} \tag{1}$$

where $\mathbf{y}$ is, in general, a k-dimensional vector of values for the k imprecise uncertainties and $\mathbf{Y}$ is the set of all k-dimensional vectors falling within the domain of imprecision for the problem (typically, this is defined using the upper and lower bounds on each imprecise variable, but more complex regions are possible). This relationship indicates that alternative A is preferred to (i.e., dominates) alternative B if the minimum difference in their expected utilities over the entire imprecise region is positive. What this means is that no matter to what specific values the imprecise variables resolve, alternative A will always be preferred to alternative B if this condition holds.

Note that if the minimum difference in expected utility is less than zero, one can conclude that alternative A fails to dominate alternative B but can say nothing about whether alternative B dominates alternative A. To evaluate the latter statement, one would have to demonstrate that the maximum of the difference in expected utilities is less than zero or, equivalently, that the criterion of Equation (1) holds after exchanging A and B for one another in the expression.

6.1.3 Domination by an Individual Attribute

In general, tradeoffs exist between different attributes such that, for example, a designer might be willing to accept reduced reliability if there also are appropriate reductions in cost. However, there are times in which one can identify an alternative as being dominated without evaluating the tradeoffs completely. For example, most designers will not select a concept for which the maximum reliability is only 0.5 no matter how inexpensive it is.

The same conclusion would be obtained when considering the overall multi-attribute utility function—one's utility function would be near zero (assuming utility normalized to [0,1]) because of the unacceptability of one of the attributes in a multiplicative utility formulation [17]. Domination of this type typically happens in extreme cases that are relatively distant from the tradeoff region of interest to designers. This is analogous to how designers use constraints to bound a search problem—the constraints imply that a designer prefers some alternative within a particular region (the feasible region) to *any* alternative outside this region (the infeasible region). Inferences of this type are most useful as a means to focus the attention of designers on regions in which their preferences should be formalized more precisely. It should be noted that if the final design falls on one of these boundaries (which are essentially constraints on an optimization problem) then the designers should re-evaluate their preferences in more detail in this region in order to determine if additional tradeoffs can be made.

6.2 Strategy for Decisions in Conceptual Design

We now combine the ideas of set-based design, multi-attribute utility theory and decision-making under imprecision into a method for making conceptual design decisions. Outlined in Figure 5 is a general strategy for making conceptual design decisions that are consistent with one's overall design objectives as formulated using multi-attribute utility theory. The approach involves eliminating concepts that are dominated by others and refining the remaining concepts to enable more complete eliminations.

Initialization:	Define initial set of concepts, C^0.
Elimination:	$c_i \in C^k$ is dominated if and only if $\exists c_j \in C^k : c_j \succ c_i$, where k indicates the current step.
Refinement:	Refine the concepts by: (a) subdividing into more granular concepts such that the original concept c_i becomes new concepts $c_{i1}, c_{i2,} \ldots, c_{in}$ with $c_{ij} \subset c_i$ for $j = 1 \ldots n$ and $\bigcup_{i=1}^{n} c_{ij} = c_i$. and/or (b) gathering more information that reduces imprecision in other uncertainties about the problem (e.g., environmental factors, preferences, etc.).
Stopping Criterion:	Until a clearly superior concept is found or external pressures force commitment to one concept via arbitrary assumptions.

Figure 5: General strategy for elimination decisions in conceptual design.

The elimination step involves one or more of the domination inferences described in the preceding section. A dominated concept is one that is assured of not leading to the most preferred design alternative. In other words, no detailed design associated with the concept has an expected utility equal to the global maximum. However, it is not always possible to eliminate a concept. Sometimes designers actually are indifferent between two concepts, in which case arbitrary decision can be an effective procedure. Other times, apparent indifference, which is more appropriately termed indeterminacy, is due to imprecision that can be resolved through further analysis. When this is the case, designers can refine the problem to reduce imprecision and return later to eliminate more concepts.

Designers can take different approaches to reducing imprecision depending on the characteristics of a problem, and often a combination of approaches is warranted. One way designers can reduce imprecision is by gathering more information about a problem. This can apply to several of the sources of imprecision identified in Section 2. However, information gathering alone can be insufficient to resolve indeterminacy. Designers also can refine the concepts themselves by subdividing a concept into two or more new (sub-) concepts. Because the new concepts are subsets of their parent, their expected utility ranges will also be subsets of the parent's. Consequently, designers potentially can divide a concept such that they can eliminate portions of it and have the remaining sub-concepts be more precisely defined than their parent.

How to identify subdivisions of concepts that are beneficial to decision making remains an open research question. Some fairly abstract concepts may consist of "clusters" of solutions that designers can model more precisely as individual concepts—e.g., for the "melt material onto paper" concept, designers may group sub-concepts by the type of material to melt on the basis that different materials require different amounts of energy to melt and have different costs. Designers must also balance the desire to increase precision by dividing concepts into more granular sub-concepts with the increased computational costs of doing so. One essentially undermines the purpose of conceptual design by developing concepts to a very detailed level.

After having refined the problem, designers return to the elimination step and iterate between the two steps until they are able to reach a rationally supported decision or external considerations—such as scheduling pressures—force them to make an arbitrary choice. That the process may terminate in an arbitrary choice is not a failure of the process but a concession to the realities of engineering design.

What this process provides to designers is a way to avoid making arbitrary decisions when circumstances permit, but the freedom to invoke assumptions and control a design process as they see fit. What is more, even if eventually forced to make an arbitrary choice, designers following this process potentially can make better arbitrary choices by virtue of having first eliminated several poor choices from contention. The value of this contribution lies in its allowing designers the freedom to avoid commitment to a single decision alternative while also providing the power of rigorous inferences based on the rationality conditions associated with multi-attribute tradeoffs in utility theory.

7 Demonstration on a Vehicle Transmission Design Problem

This section is a demonstration of using utility-based inferences to perform elimination operations for set-based conceptual design. The problem involves three concepts for the transmission subsystem of a small, single-seat off-road vehicle and involves several sources of imprecision that are typical of conceptual design. The demonstration highlights how designers can delay commitment to an individual concept while they gather and process more information about the problem by taking a set-based approach. Furthermore, it illustrates how designers can use preference-based inferences to eliminate undesirable design concepts.

7.1 Problem Definition and Proposed Design Concepts

7.1.1 Problem Definition and Initial State of Information

The task is to design part of the transmission for an off-road vehicle. We presume the following information about the problem.

System and Environment

The system consists of an engine, continuously-variable transmission (CVT), a fixed-ratio transmission and a rear differential with a fixed gear ratio, arranged as depicted in Figure 6. The task is to design the fixed-ratio drive assuming the engine, CVT and rear differential have been determined previously. Although the identities of the other system components are known, some uncertainty is contributed to the design problem by the corresponding analysis models. The environment consists of an off-road race course that is of known length. Track properties that influence vehicle performance are uncertain variables. It is assumed that other vehicles on the course do not influence the vehicle. Uncertain quantities relating to the vehicle subsystems and environment are summarized in Table 1. These uncertainty models are determined by eliciting knowledge from knowledgeable engineers, consulting reference material and examining empirical data. The imprecision due to these factors is somewhat large, as is typical in conceptual design, and may be a prime candidate for reduction via information gathering at a later point in the process. For example, we could gather empirical data about the car to determine a more precise estimate of external drag.

Table 1: Summary of uncertain quantities common to all concepts.

Uncertain Quantity	***Uncertainty Model***
Total Mass	P-box ~ Normal with $\mu = [265, 288]$, $\sigma = [8, 11]$ (kg)
External Drag Coefficient	Interval: $[0.2, 0.35]$ $\left(\mathrm{N}/(\mathrm{m}/\mathrm{s})^2\right)$
Internal Drag Coefficient	Interval: $[0.00, 0.14]$ $(\mathrm{N}\cdot\mathrm{s}/\mathrm{rad})$
Course roughness Coefficient	Normal ~ $\mu = 3$, $\sigma = 0.5$

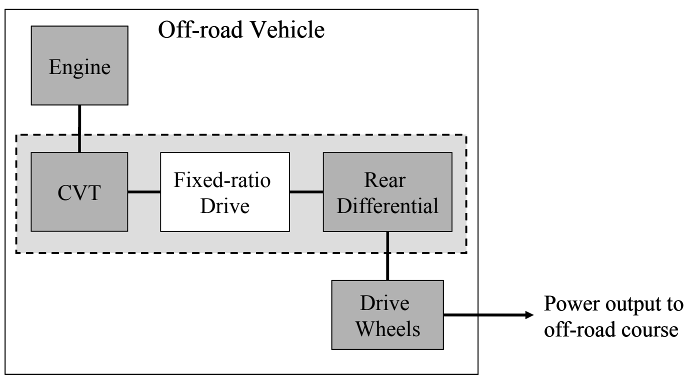

Figure 6: Configuration of and interaction among off-road vehicle components. Shaded boxes indicate previously designed components; the fixed-ratio drive is of interest in this demonstration.

<u>Preferences</u>

The setting for this example is a race with prize money awarded to the winner. We presume that designers care about three attributes: prize money won, system reliability, and the cost of constructing the system. The overall design objective is to find the most preferred design, which is the one that maximizes expected utility and is identified mathematically as

$$d^* = \arg\max_{d \in D} E\left[U(d)\right],$$

where

$$U(d) = R(d)W(d) - C(d),$$

$U(d)$ is the utility (in dollars) for a specific design alternative, d, from the notional set of all alternatives (denoted D), $R(d)$ is the reliability of the transmission, $W(d)$ is the prize money as predicted by a model based on historical race data, and $C(d)$ is the combined cost (in dollars) of constructing the fixed-ratio drive and of retrofitting the existing subsystems to interface with the drive.

The retrofitting cost exists because we presume the other subsystems are preexisting and modifications of them come at a cost. For the fixed-ratio drive to interface with the other system components, its input shaft must align with the output shaft of the CVT and its output shaft must align with the input shaft of the differential. An alternative design approach would be to constrain the fixed-ratio drive design to match the preexisting shaft-to-shaft distance. However, this approach is incapable of capitalizing on potentially beneficial tradeoffs—e.g., one might pay a little to retrofit the preexisting system but gain much in terms of performance. Thus, domain expertise is used to develop an imprecise cost model in which small modifications cost very little but the cost grows quickly for larger modifications, with the modification magnitude taken as the required change in shaft-to-shaft distance.

7.1.2 Proposed Design Concepts

Three concepts are proposed as potential solutions to the design problem: a belt drive, a chain drive and a gearbox. They are described below.

<u>Belt Drive</u>: The belt drive concept consists of a flat, multi-ply belt running in a single stage between input and output pulleys.

<u>Chain Drive</u>: The chain drive concept consists of a single-stage from input to output sprocket. It uses a standard double-strand roller chain.

<u>Gearbox</u>: The gearbox concept consists of a train of three gears: an input pinion, an idler and an output gear. Gears are constructed from non-exotic materials (e.g., carbon steel).

7.2 Set-based Solution Procedure

The solution procedure is based on the general approach outlined in Figure 5, with an interweaving of concept refinement and elimination operations. The steps are outlined below and detailed in the following sections.

<u>Initial Concept Modeling and Refinement</u>. Each design concept is modeled using bounds on two attributes (reliability and cost) and a single parameter (transmission ratio) which is used to predict an imprecise range for the third attribute (prize money). Preference-based inferences are used to refine each concept by eliminating dominated ranges of the transmission ratio parameter.

<u>Initial Comparison of Concepts</u>. Each design concept is compared using the domination criteria described in Section 6.1. Interval dominance is attempted first, followed by maximality if interval dominance fails to identify a single non-dominated concept.

Secondary Concept Modeling and Refinement. Each design concept is modeled with more precision, although still with imprecise abstract relationships. All attributes are predicted from parameters specific to each concept. The parameters are modeled as imprecise ranges. Preference-based inferences are used to refine each concept by eliminating dominated regions of the parameter space.

Secondary Comparison of Concepts. Same as initial comparison step, except uses the more precise concept models.

7.2.1 Initial Concept Modeling and Refinement

One of the first considerations when designing a transmission is to determine the drive ratio. Under a classical approach, a designer might select a precise drive ratio and evaluate each concept for that ratio. However, a set-based approach permits designers to carry forward a range of drive ratios that may lead to the most preferred solution. This allows designers to make tradeoffs later that are difficult to evaluate under the imprecise information available in early-stage conceptual design.

To eliminate dominated drive ratios, we must first establish the relationship between this parameter and the attributes about which designers care. There exists no direct and unique relationship between drive ratio and reliability—for the concepts being considered, one can design drives with a wide range of reliability values for a given ratio. The same can be said for cost. Thus, at this point we model each of these attributes as intervals that are independent of the drive ratio and of each other. Appropriate bounds for the intervals are determined using domain expertise.

It is possible to model the relationship between drive ratio and the prize money attribute using simple algebraic relationships. Although the relationships are simple, there are several of them and a full exposition is omitted in the interest of brevity (readers can consult [8] for details). The key features of this model are that it is imprecise, it captures the general behavior of the attribute as a function of the parameter, and it is common to all design concepts. By propagating the imprecise range of the drive ratio through the model, we can predict the imprecise range of the prize money attribute. Combining this information with the corresponding intervals for cost and reliability, we can compute bounds for the expected utility of a concept.

Using this model, we refine the range of drive ratio values by using the interval dominance inference mechanism to identify drive ratio values that can be eliminated (i.e., no transmission with these drive ratios would maximize expected utility). After eliminating some drive ratios, we may be able to refine our bounds on cost and reliability. For example, ruling out extremely high ratios may lead us to reduce our upper bound on cost. Given this increased precision, we can try to eliminate more of the drive ratio range.

The elimination process can iterate but quickly converges for each concept to a drive ratio range that cannot be reduced under the given level of imprecision. This progression is evident from the data associated with the gearbox concept, displayed in Table 2. A similar procedure leads to refined bounds for the other concepts. The results are summarized in Table 3.

Table 2: Progression of characterization of gearbox concept using expert opinion and dominance inferences using an abstract model (new information at each step in bold).

Quantity	*Initial Bounds based on expertise*	*After elimination on drive ratio*	*Revised bounds based on expertise and elimination results*	*After further elimination on drive ratio*
Drive Ratio	Any positive real	**[0.875, 3.25]**	[0.75, 3]	**[0.85, 3.075]**
Reliability	**[0.95, 1]**	[0.95, 1]	**[0.975, 1]**	[0.975, 1]
Cost	**[250, 1000]**	[250, 1000]	**[250, 550]**	[250, 550]
Expected Utility	Unknown	**[-997, 1736]**	[-997, 1739]	**[-547, 1739]**

Table 3: Definitions of each concept after initial refinement stage.

Quantity	***Refined Concept Definitions***		
	Belt Drive	*Chain Drive*	*Gearbox*
Drive Ratio	[0.80, 3.00]	[0.75, 3.25]	[0.85, 3.025]
Reliability	[0.985, 1]	[0.98, 1]	[0.975, 1]
Cost	[150, 700]	[150, 650]	[150, 450]
Expected Utility	[-700, 1839]	[-950, 1839]	[-547, 1839]

Table 4: Minimum differences in expected utility for pair-wise comparisons of concepts.

	Belt Drive	*Chain Drive*	*Gearbox*
Belt Drive	n/a	-2484	-2225
Chain Drive	-2108	n/a	-1379
Gearbox	-350	-2434	n/a

7.2.2 Initial Comparison of Concepts

At this point, there is significant overlap in the imprecise range of expected utility for the three concepts. This is evidenced in Table 3; all three concepts have nearly identical upper expected utilities and their intervals overlap significantly. Consequently, it is impossible to eliminate any of the concepts using the interval dominance criterion.

Despite being a stronger form of inference, the maximality criterion also leads to no concepts being eliminated at this stage. This is evidenced by the data of Table 4. Each table cell holds the minimum difference in expected utility between the column concept and row concept over the domain of imprecision (i.e., the results of the minimization in Equation (1), with the column as alternative A and the row as alternative B). All comparisons result in a minimum value that is less than zero, which means no eliminations are possible and the decision problem requires refinement.

It is possible that all concepts are actually preferred equally, but it also is possible that we cannot eliminate any concepts at this point because the models we are using fail to capture the relationships that the maximality criterion exploits. Although it is possible to make an arbitrary assumption at this point and select a single concept—e.g., choose the gearbox on the basis that it has the highest lower bound (i.e., it is the Γ-maximin solution)—we choose to delay such assumptions while we work to reduce imprecision in the problem. Thus, in the next step we develop better analysis models for the concepts.

7.2.3 Secondary Concept Modeling and Refinement

Thus far, we have modeled two of the design attributes (cost and reliability) using imprecise ranges based on expert opinion and the third (prize money) as a function of a single design parameter using an imprecise algebraic model. One way to reduce imprecision in a design concepts is to model all three attributes using design parameters appropriate for the concept. Doing this allows us to account for dependencies that may exist between the attributes.

Developing the Analysis Models for Each Concept

Figure 7 is a depiction of the relationship between the generic transmission model and the specialized design concept models. Each of the new concept models inherits the drive ratio parameter as well as having three unique parameters of its own. The models are algebraic and based on those found in standard references relating to the particular concepts (e.g., see [1, 25, 32]).

Other parameterizations of the design concepts are possible in general. These are chosen because they are common parameters for specifying these kinds of design concepts at this level of abstraction. Note that several other parameters for each concept are abstracted away (e.g., lubrication types, numbers and types of bearings required, shaft and enclosure considerations, etc.). Thus, the models are quite imprecise despite the presence of geometric parameters.

Each model includes many uncertainties, which are summarized in Table 5. The focus of this demonstration is on the *process*, not the specific numbers being used. However, it is informative to consider the general approach by which we arrived at some of these models, which were determined using a combination of domain knowledge, data from reference materials and physical modeling. For example, the tensile strength of chain is modeled as a function of chain pitch length by fitting a relationship to catalog data using regression methods. The model is associated with a p-box representing the modeling uncertainty. Cost models were similarly fit to catalog data but the uncertainty associated with them was greater and less justifiable as random, so we deemed an interval representation more appropriate. The application factors for each concept were modeled as intervals using standard reference tables. These tables are themselves imprecise, involving vague linguistic descriptions of application situations that index to precise factors.

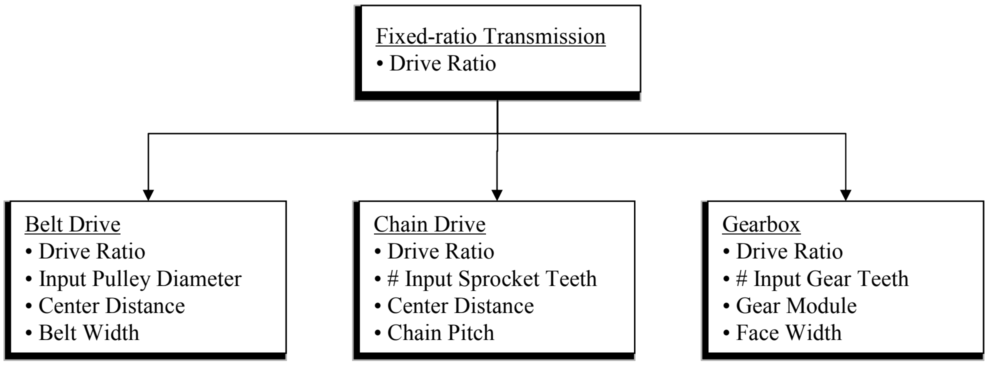

Figure 7: Specialization of generic transmission model into concept-specific models. Bulleted items indicate model parameters.

Table 5: Concept-specific uncertainties.

Belt Drive	
Cost model error	Interval: [-45, 45] ($)
Application factor	Interval: [1.85, 1.95]
Linear density, per 0.152 m belt width	P-box ~ $Normal([0.65,0.7],[0.025,0.05])$ (kg/m-length)
Coefficient of friction, belt on steel	PDF ~ $Normal(0.3,0.033)$
Tensile Strength, belt	P-box ~ $Normal([140,150],[5,8])$ (N/mm-width)
Fatigue constant, belt	[9, 14]
Chain Drive	
Cost model error	Interval: [-45, 45] ($)
Application factor	Interval: [1.65, 1.75]
Linear density model error, chain	Interval: [0.0374, 0.1364] (kg/m)
Tensile strength model error, chain	P-box ~ $Normal([-0.05,0.05],[0.005,0.01])$ (N)
Galling model constant	PDF ~ $Normal(0.3226,0.01)$
Gearobx	
Cost model error	Interval: [-25, 25] ($)
Bending strength geometry factor	Interval: [0.27, 0.37]
Gear quality factor	P-box ~ $Normal([8,10],[0.5,0.55])$
Bending strength, gear material	P-box ~ $Normal([170,230],[20,40])$ (MN/m^2)
Contact strength, gear material	P-box ~ $Normal([590,660],[20,40])$ (M N/m^2)
Application factor	Interval: [1.65.1.85]

Although some parameters are modeled explicitly as a function of the design parameters— e.g., linear density of the belt, tensile strength of the chain, cost—others that could be are not. For instance, the bending strength geometry factor is a function of , among other things, the number of teeth on each of two gears in mesh. Although it is possible to gauge a more precise value for this factor given particular design parameter values, a modeling decision was made to represent this uncertainty using an interval.

This is perfectly valid, since any design alternative that dominates another with the interval representation would also dominate using a more precise model (assuming the interval subsumes the values of the more precise model). It always is possible to move to a more precise model later if we suspect that reducing this imprecision could lead to the elimination of more concepts [3].

Refining the Parameter Ranges for Each Concept

Given the newly defined models for our concepts, it is possible to refine the sets defined by our concepts further through an elimination procedure analogous to that used in Section 7.2.1. To do this, we use the maximality criterion in concert with a reference design to serve as a baseline for comparison. For each concept, a promising instance is chosen based on designer intuition (defined by precise values for the concept's design parameters). This *reference design* serves as a baseline against which all other members of the concept set are measured and is used here because efficient numerical methods for applying maximality over a continuous range is a current research topic. If the reference design dominates some subset of the concept set, then it is reasonable to remove that subset from the concept definition. The result of this refinement step is more precise bounds on the design parameters for each concept, such that the bounds encompass the non-dominated instances within the original concept set.

The resulting reduced bounds are summarized in Table 6. One can see that for each concept the bounds on transmission ratio are tighter here than they were in Table 3. The same also is true for the bounds on expected utility. This is because the more detailed analysis models account for the general relationship between the design attributes that previously were modeled as being independent.

Given the new, more precise, design concept definitions it is worth checking again to see if we can eliminate any of the concepts.

Table 6: Design concept set definitions after refinement based on comparisons with a reference design.

Belt drive		*Chain Drive*		*Gearbox*	
Ratio	[1.1, 2.2]	Ratio	[1.1, 2.2]	Ratio	[1.3, 2.2]
Input pulley diam.	[0.05, 0.12] (m)	# teeth, input	[16, 26]	# teeth, input	[15, 20]
Belt width	[76, 203] (mm)	Chain pitch	[9.5, 15.9] (mm)	Face width	[6, 10] (mm)
Center dist.	[0.200, 0.500] (m)	Center dist.	[20, 35] (pitches)	Module	[4.15, 4.25] (mm)
Expected Utility	[-625, 1805]	Expected Utility	[72, 1831]	Expected Utility	[1281, 1809]

7.2.4 Check for Dominated Concepts

By observation of Table 6, one cannot eliminate any of the concepts using interval bounds rationale. As in Section **Error! Reference source not found.**, we turn again to pair-wise comparisons of the design concepts and the maximality criterion for elimination. Using maximality, we are able to conclude that one of the concepts—the belt drive—can be eliminated by virtue of its being dominated by the gearbox concept. The comparison data appear in Table 7. Recall that the condition for dominance is that the minimum difference in expected utility is greater than zero.

Table 7: Minimum differences in expected utility for pair-wise comparisons of the concepts.

	Belt Drive	*Chain Drive*	*Gearbox*
Belt Drive	n/a	-1029	611
Chain Drive	-1825	n/a	-326
Gearbox	-1809	-2836	n/a

That the belt drive should be dominated is not obvious by inspecting the ranges of expected utility. As evidenced by the data, there are instances of the belt drive concept that have relatively high expected utility. What the domination results indicate is that the uncertainty shared between the concepts forces a relationship between their expected utilities such that for a given point in the shared uncertainty space there exists at least one instance of the gearbox concept that has a higher expected utility than any instance of the belt drive concept. The reason for this is that the gearbox typically is a more compact solution and therefore does not incur as large of retrofitting costs as the belt drive. This effect also explains why the gearbox nearly dominates the chain drive concept.

At this point, it is possible to continue the process with the two remaining concepts. However, our demonstration of the procedure will end here. Because there is indeterminacy about which of the two remaining concepts is more preferred, one must eliminate additional imprecision in order to eliminate one of them. As noted previously, there are many remaining sources of imprecision in this problem that are prime candidates for elimination.

8 Discussion

8.1 Delaying Commitment while Making Rational Decisions

The set-based approach to design demonstrated in the preceding section allows designers to make conceptual design decisions without having to commit to a single concept in a single step. The multi-step design is advantageous because it allows designers to use their limited resources more efficiently. At each step, designers focus their development effort on only the design concepts that may lead to the most preferred design solution (i.e., the non-dominated concepts).

There are multiple examples of this in the demonstration. Eliminating particular ranges of transmission ratio in Section 7.2.1 allowed us to focus on single-stage belt and chain drives and a relatively simple gear train. Had we been unable to eliminate large ratios, the concepts would have required broader definitions to include such scenarios. Thus, we made a decision that focused our attention on concepts that can lead to the design solution we prefer most. This happens again in Section 7.2.4 when we are able to eliminate the belt drive concept; if the problem were carried forward, future design effort would focus on the gearbox and chain drive concepts.

Although existing descriptions of set-based design include the notion of delaying commitment while awaiting more information, the combination of set-based design and multi-attribute utility described here is unique and powerful. This combination allows designers to eliminate concepts according to their overall preference structure, taking into account the various tradeoffs they are willing to make. This is illustrated at several points in the demonstration—for example, we use utility-based comparisons to a reference design during a refinement step of Section 7.2.3 to eliminate dominated sub-concepts from each particular concept set.

Prior work on set-based design describes eliminations based only on feasibility information. Thus, one is more limited in the inferences that can be drawn. For example, applying feasibility-based inferences to refine the concepts in Section 7.2.4 would lead one to eliminate physically impossible concepts (e.g., gear trains that do not mesh, chain drives that have incompatible chains and sprockets, etc.) but would leave unevaluated any of the tradeoffs about which we have a preference. Designers would be unable to eliminate sub-concepts despite the fact that they never would choose them over other sub-concepts.

8.2 Handling Concepts that are More Abstract

In this example, the concepts (belt drive, gearbox, chain drive) were sufficiently embodied that we could use models to predict attributes of interest with some degree of confidence. Specifically, the concepts had definite forms, but uncertain dimensions. As such, abstract models of cost and reliability could be used to evaluate imprecise inputs. In more general applications, the structure of the concepts themselves may be imprecise. An open and significant question is *how to model the performance of abstract concepts during conceptual design*.

This question is fundamentally different from those traditionally addressed in computer aided design, which has focused on creating a mathematical and computer-interpretable model of a defined form, and then exploring its properties. A general concept could include several different forms as detailed instances. For example, one might consider two different transmission concepts, one being "mechanical transmission" and the other being "hydrodynamic transmission." The mechanical transmission concept would include all of those explored here.

A related issue is expanding the concept representations to be more expressive. We use simple representations in our demonstration, relying on involving interval bounds on model parameters to define our concepts and, using these bounds and models, we are able to predict a corresponding range of expected utility. Although the bounds we use are capable of encompassing the non-dominated set for each design concept, they are less than ideal because they also admit some sub-concepts that are dominated. This occurs because there are relationships between the parameters that manifest in the attribute space. In the belt drive concept, for example, the minimum width a pulley could have and still be in the non-dominated set depends on the pulley diameter (both parameters relate to whether the belt will slip, which relates to its reliability).

The main consequence of including dominated sub-concepts in the representation of a design concept is that predictions of expected utility are less precise. The predicted range of expected utility includes the lower utility values associated with the dominated sub-concepts. This reduces one's ability to eliminate concepts, which makes the overall process less efficient. Because our focus is on the general procedure for set-based conceptual design, the simple interval representation is adequate for our demonstration.

The predictive modeling of concepts is an important area for future work. The initial emphasis of this research is not in developing tools, but identifying the fundamental principles and approaches for modeling imprecise concepts.

8.3 Higher-level Management of the Set-based Design Process

The example problem presented in this article is a clear demonstration of gradual refinement of a concept, a process in which the decision maker does not commit to a single alternative at each step but rather uses utility theory to make set-based, elimination decisions. This example focused on a significant step in set-based design, but it did not address the managerial aspects of set-based design. For example, open questions include *how can a decision maker efficiently partition the design space into concepts*, and *when should the decision maker terminate the set-based process and choose a single solution*?

The partitioning question is crucial to the efficiency of a conceptual design process. Every concept designers identify requires some amount of modeling and characterization. Breaking a design space into too many concepts can lead to an excessive modeling workload. On the other hand, concepts defined too broadly will be very imprecise and it may not be possible to eliminate any of them. An effective procedure will balance the desire to eliminate as much of the design space as possible with the modeling and computational costs of doing so.

Termination of the set-based process is a problem for which there may exist no rigorous answer. Ideally, the process will terminate when the concepts are defined with sufficient precision that a clear winner becomes evident. In practice, the situation is more complex. As noted in the literature review, there is evidence that the value of performing set-based design can be limited if one selects down to a single alternative too quickly or too slowly [15]. External concerns such as market pressures or contractual deadlines may force designers to select a single concept despite the existence of indeterminacy. In such cases, the value of set-based design lies in eliminating the clearly undesirable concepts prior to the selection deadline. In other cases, there may never be a unique most-preferred concept. History has shown that often there are multiple engineering solutions to the same problem. A gearbox and a chain drive may essentially be equivalent with respect to certain preferences. The challenge then is identifying that this is the case, so that designers can confidently make an arbitrary selection and move on to other phases of design.

9 Summary

Conceptual design decisions involve significant levels of imprecision, owing to many sources. Although multi-attribute utility theory is an effective framework for formalizing and making decisions in engineering design, the imprecision inherent in conceptual design makes it difficult for designers to identify the concept that leads to the most preferred design. Set-based design is a general approach to making design decisions in which one focuses on eliminating demonstrably inferior alternatives rather than on committing in one step to a single alternative. However, prior work on set-based design has been limited to feasibility-oriented reasoning. In this article, we describe a general approach to making conceptual design decisions that combines the formal tradeoff analysis of multi-attribute utility theory with the elimination-based perspective of set-based design. The result is a general procedure by which designers can make conceptual design decisions without having to commit to a single alternative if current information does not allow for rational support of such a decision.

Many open questions remain about set-based design and how to implement it efficiently. It requires a new perspective on concept modeling and computer-aided design tools before its full potential can be realized. Guidelines for applying it effectively also are needed. However, despite the open issues, the set-based nature of conceptual design leads to the conclusion that there is a clear potential for set-based approaches to conceptual design.

Acknowledgement

This research is sponsored in part by NSF grant DMI-0522116. Jason was supported by a National Science Foundation Graduate Research Fellowship before joining the Applied Research Laboratories at the University of Texas at Austin. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation. Additional support is provided by the G.W. Woodruff School of Mechanical Engineering at the Georgia Institute of Technology. The authors thank Morgan Bruns, Steve Rekuc and Jay Ling for their contributions to this work.

References

[1] American Chain Association (2006) *Standard Handbook of Chains: Chains for Power Transmission and Material Handling* 2 ed, Dekker Mechanical Engineering, Taylor & Francis, Boca Raton, FL.

[2] Arrow K and Hurwicz L (1972) An Optimality Criterion for Decision-Making Under Ignorance. In Carter CF and Ford JL (eds) Uncertainty and Expectation in Economics: Essays in honour of G. L. S. Shackle. Blackwell, Oxford.

[3] Aughenbaugh JM, Ling J and Paredis CJJ (2005) The Use of Imprecise Probabilities ot Estimate the Value of Information in Design, in *Proceedings of*

the 2005 ASME International Mechanical Engineering Congress and RD&D Expo, Orlando, FL, 5-11 November.Paper No. IMECE2005-81181.

[4] Aughenbaugh JM (2006) Managing Uncertainty in Engineering Design Using Imprecise Probabilities and Principles of Information Economics. Ph.D. thesis, Georgia Institute of Technology, Atlanta, GA.

[5] Aughenbaugh JM and Paredis CJJ (2006) *The Value of Using Imprecise Probabilities in Engineering Design*. Journal of Mechanical Design. 128: 969-979.

[6] Ayyub B (2001) *Elicitation of Expert Opinions for Uncertainty and Risks*, CRC Press, New York.

[7] Berger JO (1985) *Statistical Decision Theory and Bayesian Analysis* 2nd ed, Springer.

[8] Bruns M (2006) Propagation of Imprecise Probabilities through Black Box Models. MS thesis, Georgia Institute of Technology, Atlanta, GA.

[9] Cooke RM (1991) *Experts in Uncertainty: Opinion and Subjective probability in Science*, Oxford University Press, Oxford.

[10] de Finetti B (1964) Foresight. Its Logical Laws, Its Subjective Sources (translated). In Kyburg HE and Smokler HE (eds) Studies in Subjective Probability. Wiley, New York.

[11] Der Kiureghian A (1989) *Measures of Structural Safety under Imperfect States of Knowledge*. Journal of Structural Engineering. 115: 1119-1139.

[12] Ferson S and Ginzburg L (1996) *Different Methods are Needed to Propagate Ignorance and Variability*. Reliability Engineering and System Safety. 54.

[13] Ferson S and Donald S (1998) Probability Bounds Analysis, in *Proceedings of the International Conference on Probabilistic Safety Assessment and Management (PSAM4)*, New York, NYSpringer-Verlag.

[14] Finch WW and Ward AC (1997) A Set-based System for Eliminating Infeasible Designs in Engineering Problems Dominated by Uncertainty, in *Proceedings of the 1997 ASME Design Engineering Technical Conferences*, Sacramento, CA.Paper No. DETC97/DTM-3886. American Society of Mechanical Engineers.

[15] Ford DN and Sobek II DK (2005) *Adapting Real Options to New Product Development by Modeling the Second Toyota Paradox*. IEEE Transactions on Engineering Management. 52: 175-185.

[16] Groen FJ and Mosleh A (2005) *Foundations of probabilistic inference with uncertain evidence*. International Journal of Approximate Reasoning. 39: 49-83.

[17] Keeney RL (1982) *Decision Analysis: An Overview*. Operations Research. 30: 803-838.

[18] Keeney RL and Raiffa H (1993) *Decisions with Multiple Objectives* (1976 orig) 2 ed, Combridge University Press, Cambridge, UK.

[19] Kreinovich V, Ferson S, Ginzburg L, Schulte H, Barry M and Nguyen H (1999) From Interval Methods of Representing Uncertainty to a General Description of Uncertainty, in *Proceedings of the International Conference on Information Technology*, Bhubaneswar, India, December 20-22, 1999Mohanty H and Baral C (eds) McGraw-Hill. 161-166.

[20] Liker JK, Sobek II DK, Ward AC and Cristiano JJ (1996) *Involving Suppliers in Product Development in the United States and Japan: Evidence for Set-based Concurrent Engineering*. IEEE Transactions on Engineering Management. 43: 165-178.

[21] Luce RD and Raiffa H (1957) *Games and Decisions*, Wiley, New York.

[22] Moore RE (1966) *Interval Analysis*, Prentice Hall, Englewood Cliffs, NJ.

[23] Muhanna RL and Mullen RL (2004) Interval Methods for Reliable Computing. In Nikolaidis E, Ghiocel DM and Singhal S (eds) Engineering Design Reliability Handbook. CRC Press, New York, NY.

[24] Nahm Y-E and Ishikawa H (2005) *Novel Space-based Design Methodology for Preliminary Engineering Design*. International Journal of Advanced Manufacturing Technology.

[25] Norton RL (2000) *Machine Design: An Integrated Approach* 2 ed, Prentice Hall, Upper Saddle River, NJ.

[26] Oberkampf WL, DeLand SM, Rutherford BM, Diegert KV and Alvin KF (2002) *Error and Uncertainty in Modeling and Simulation*. Reliability Engineering and System Safety. 75: 333-357.

[27] Pahl G and Beitz W (1996) *Engineering Design: A Systematic Approach* 2nd ed, Wallace K, Blessing L and Baurt F (trans), Wallace K (ed), Springer-Verlag, London.

[28] Panchal JH, Fernandez MG, Allen JK, Paredis CJJ and Mistree F (2005) An Interval-based Focalization Method for Decision-Making in Decentralized Multi-Functional Design, in *Proceedings of the ASME 2005 Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Long Beach, CA.Paper No. DETC2005-85322.

[29] Parry GW (1996) *The Characterization of Uncertainty in Probabilistic Risk Assessment of Complex Systems*. Reliability Engineering and System Safety. 54: 119-126.

[30] Parunak HVD, Ward AC, Fleischer M and Sauter J (1997) A Marketplace of Design Agents for Distributed Concurrent Set-based Design, in *Proceedings of the ISPE/CE97: Fourth ISPE International Conference on Concurrent Engineering: Research and Applications*, Troy, MI, 20-22 Aug. 1997. 287-293.

[31] Rekuc SJ, Aughenbaugh JM, Bruns M and Paredis CJJ (2006) Eliminating Design Alternatives Based on Imprecise Information, in *Proceedings of the Society of Automotive Engineering World Congress* Paper No. 2006-01-0272.

[32] Shigley JE, Mischke CR and Brown TH (eds) (2004) Standard Handbook of Machine Design. McGraw-Hill, New York.

[33] Simpson T, Lautenschlager U and Mistree F (1998) Mass Customization in the Age of Information: The Case for Open Engineering Systems. In Porter AL and Read WH (eds) The Information Revolution: Current and Future Consequences. Ablex Publishing Corporation, Greenwich, Connecticut.

[34] Sobek II DK, Ward AC and Liker JK (1999) *Toyota's Principles of Set-Based Concurrent Engineering*. Sloan Management Review. 40: 67-83.

[35] von Neumann J and Morgenstern O (1980) *Theory of Games and Economic Behavior* 3 ed, Princeton University Press, Princeton, NJ.

[36] Walley J (1991) *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, New York.

[37] Ward AC (1989) A Theory of Quantitative Inference Applied to a Mechanical Design Compiler. PhD thesis, MIT, Boston.

[38] Ward AC, Liker JK, Cristiano JJ and Sobek II DK (1995) *The Second Toyota Paradox: How Delaying Decisions Can Make Better Cars Faster*. Sloan Management Review. 36: 43-61.

[39] Weber M (1987) *Decision making with incomplete information*. European Journal of Operational Research. 28: 44.

[40] Zaffalon M, Wesnes K and Petrini O (2003) *Reliable Diagnoses of Dementia by the Naive Credal Classifier Inferred from Incomplete Cognitive Data*. Artificial Intelligence in Medicine. 29: 61-79.