
G1.13 A Fuzzy Sets Application to Preliminary Passenger Vehicle Structure Design

Michael J. Scott, William S. Law, and Erik K. Antonsson

Engineering Design Research Laboratory
California Institute of Technology

Abstract

The *Method of Imprecision*, or M_{I} , is a formal system that uses the mathematics of fuzzy sets to model imprecision in design descriptions and performances. The M_{I} uses preference information obtained from customers and designers, as well as standard engineering analyses, to guide preliminary design decisions. The calculations of the M_{I} include statistical experimental design routines for constructing approximations when standard analysis tools are costly and function evaluations must be kept to a minimum.

G1.13.1 Overview

Unlike many fuzzy applications, which employ crisp data but model fuzzy functions, the *Method of Imprecision*, or M_{I} (Wood and Antonsson 1989), considers fuzzy inputs and outputs of the design process. Figure G1.13.1 shows the structure of the M_{I} . Design analyses can be considered crisp, but they are usually applied only after the values of design variables have been chosen, while the choice of preliminary design variable values is made on the basis of designers' informal, intuitive preferences. The M_{I} models designers' and customers' preferences as memberships of fuzzy sets, thus treating the inputs and outputs of the design process as imprecise. Engineering analysis is combined with extended fuzzy techniques to efficiently map design preferences to performances. Finally, the M_{I} aggregates the various preferences into overall preferences, which are then used to support design decisions.

G1.13.2 Motivation

One of the most critical problems in engineering design is making early decisions on a sound basis. However, the early stages of design are also the most uncertain, and obtaining *precise* information upon which to base decisions is usually impossible. The primary reason for this difficulty is that *imprecision* is an integral part of the engineering design process: not imprecision in thought or logic, but rather the intrinsic imprecision of a preliminary, incomplete description. At the concept stage, the design description is nearly completely imprecise. The design process reduces this imprecision until ultimately the final description is precise (crisp), except for tolerances, which represent the allowable limits on stochastic manufacturing variations.

Despite this evolution of imprecision, engineering design methods and computer aids nearly all utilize precise information (though some can include stochastic effects). Solid modeling CAD systems, for example, require precise geometry; there is no option to indicate an imprecise value for a dimension.

The need for a methodology to represent and manipulate imprecision is greatest in the early, preliminary phases of engineering design, where the designer is most unsure of the final dimensions and shape, materials and properties, and performance of the completed design. Additionally, those decisions with the greatest effect on overall cost are made in these early stages (Holmes 1984, Sullivan 1986, Ulrich and Pearson 1993, Whitney 1988).

The Method of Imprecision uses the mathematics of fuzzy sets as a natural representation of design imprecision. When precise information is either unavailable or prohibitively expensive, as is usually the case in preliminary design, decision support is made feasible by a fuzzy representation. In the context of

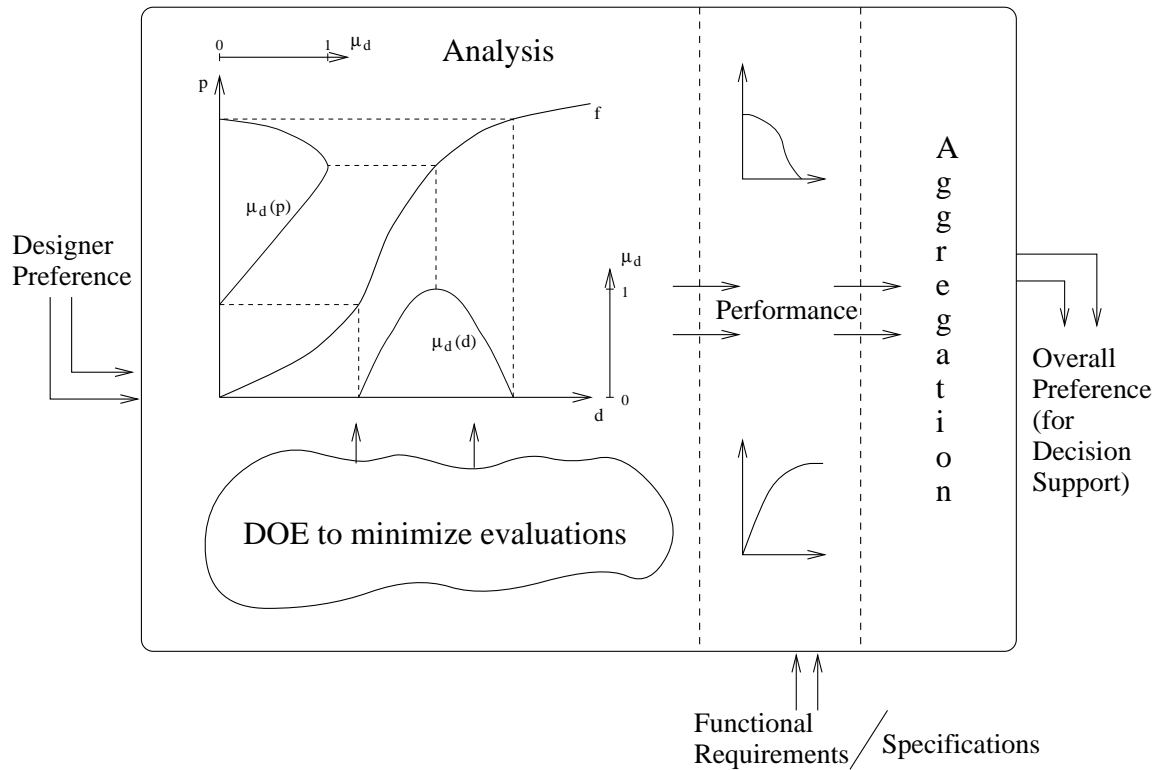


Figure G1.13.1. The Method of Imprecision

engineering design, the term *imprecision* is used to mean uncertainty in choosing among alternatives. An imprecise variable in preliminary design may potentially assume any value within a (known) range because the designer does not know, *a priori*, the final value that will emerge from the design process. A design concept can be described by a collection of (imprecise) variables; the nominal value of a length dimension is an example of an imprecise variable. Imprecision is contrasted with the stochastic uncertainty exemplified by uncontrolled variations in materials and manufacturing processes, and for which many methods, including Taguchi's method (Peace 1993, Phadke 1989), probabilistic optimization, and utility theory (Von Neumann and Morgenstern 1953, Keeney and Raiffa 1993), have been developed. These statistical and probabilistic methods, however, do not capture the imprecision inherent in design variables.

G1.13.3 Imprecision in Engineering Design

The Method of Imprecision

To represent an imprecise variable, d , a range of real numbers could be used, in the style of interval analysis (Ward and Seering 1989a and 1989b, Ward *et al* 1990). As the designer usually has a preference for certain values over others, imprecision is represented in the $M\mathcal{O}I$ by a range as well as a function, μ_d , defined on this range to describe the desirability of or preference for particular values (see Figure G1.13.2). This preference may arise objectively (*e.g.*, cost or availability of components or materials) or subjectively (*e.g.*, from experience). The $M\mathcal{O}I$ uses the formal representation of imprecise variables as ranges of values together with preferences to quantify the imprecision and perform the calculations to support preliminary design decision-making.

Design alternatives are distinguished by the values of *design variables* d_1, \dots, d_n . The whole set of design variables forms an n vector, \vec{d} , that distinguishes a particular design alternative in the *design variable space* (DVS). The set of considered values for d_i is denoted \mathcal{X}_i .

Performance variables p_1, \dots, p_q are the aspects of a design's performance that are explicitly quantified.




Figure G1.13.2. Example imprecise design variable

Each performance variable p_j is defined by a mapping f_j such that $p_j = f_j(\vec{d})$. The mappings f_j can be any calculation or procedure to evaluate the performance of a design. When the calculations are expensive or time-consuming, as in the example presented here, approximations can be constructed. The set of valid values for p_j is denoted \mathcal{Y}_j . The set of performance variables for each design alternative forms a q vector, $\vec{p} = \vec{f}(\vec{d})$, that specifies the quantified performance of a design \vec{d} . Other aspects of performance which are not quantified are not formally modeled as performance variables, and are excluded from \vec{p} . The *performance variable space (PVS)* encompasses all quantified performances $\vec{p} = \vec{f}(\vec{d})$ that are achievable by designs $\vec{d} \in \text{DVS}$.

The *functional requirement* $\mu_{p_j}(p_j)$ represents the preference that a customer has for values of the performance variable p_j :

$$\mu_{p_j}(p_j) : \mathcal{Y}_j \rightarrow [0, 1] \subset R. \quad (\text{G1.13.1})$$

$\mu_{p_j}(p_j)$ quantifies the customer's preference for values of p_j . The highest possible preference $\mu_{p_j}(p_j) = 1$ indicates complete satisfaction with that value of p_j , while $\mu_{p_j}(p_j) = 0$ represents an unacceptable value, with intermediate values showing intermediate levels of satisfaction.

The *design preference* function $\mu_{d_i}(d_i)$ represents the preference that the designer has for values of the design variable d_i based on aspects of design performance that are not already represented by performance variables:

$$\mu_{d_i}(d_i) : \mathcal{X}_i \rightarrow [0, 1] \subset R. \quad (\text{G1.13.2})$$

As the customer has no direct basis for preferences among values of a design variable d_i , the designer must decide how values of d_i influence unquantified aspects of design performance which are not represented by performance variables.

Aggregation

In order to evaluate a design $\vec{d} \in \text{DVS}$, the various individual preferences must be combined or aggregated to give a single, overall measure. This aggregation, in practice, occurs in two stages. First, the individual design preferences μ_{d_i} are aggregated into the combined design preference μ_d and the individual functional requirements μ_{p_j} are aggregated into the combined functional requirement μ_p . Second, μ_d and μ_p are aggregated into the *overall preference* μ_o , which combines all of the preferences specified:

$$\begin{aligned} \mu_o &= \mathcal{P}_{III}(\mathcal{P}_I(\mu_{d_1}, \dots, \mu_{d_n}), \mathcal{P}_{II}(\mu_{p_1}, \dots, \mu_{p_q})) \\ &= \mathcal{P}_{III}(\mu_d, \mu_p). \end{aligned} \quad (\text{G1.13.3})$$

The *aggregation functions* \mathcal{P}_I , \mathcal{P}_{II} , and \mathcal{P}_{III} reflect the trade-off strategies for each aggregation, formalizing the designer's balancing of conflicting goals and constraints (Otto and Antonsson 1990, Scott and Antonsson 1995). As the functional requirements and design preferences can be modeled as membership functions on fuzzy sets, the aggregation of preferences is a problem of the aggregation of fuzzy sets. The choice of aggregation function is dictated by the relationship between attributes in the design problem. Consider a system of components, for example, where the failure of one component results in the failure of the entire assembly. A high preference corresponding to a long time to failure for one component cannot compensate

for a low preference corresponding to a short time to failure for another component. This requires a *non-compensating* trade-off strategy for which the aggregation function is the minimum \mathcal{P}_{\min} :

$$\mu_o(\vec{d}) = \min(\mu_{d_1}, \dots, \mu_{d_n}, \mu_{p_1}, \dots, \mu_{p_q}). \quad (\text{G1.13.4})$$

Alternatively, consider the performance variables energy storage and unit cost of an ordinary household battery. Low unit cost can partially compensate for short battery life and long battery life can partially compensate for high unit cost. These two attributes are appropriately combined with a *compensating* trade-off strategy, for which the aggregation function is the geometric weighted mean or product of powers \mathcal{P}_{Π} :

$$\mu_o(\vec{d}) = \left(\prod_{i=1}^n \mu_{d_i} \prod_{j=1}^q \mu_{p_j} \right)^{\frac{1}{n+q}} \quad (\text{G1.13.5})$$

The aggregation functions \mathcal{P}_{\min} and \mathcal{P}_{Π} are just two examples of levels of trade-off that can be modeled by the **MJ**. Intermediate levels of aggregation are also possible, and importance weightings can be specified for each attribute (see Scott and Antonsson 1995 for details). Within a single design problem, different groups of attributes may require different trade-off strategies. In general, preferences for individual attributes will need to be successively aggregated by a hierarchy of different trade-off strategies.

Mapping Design Imprecision

In implementing the Method of Imprecision, a key step is mapping design preference μ_d from the n -dimensional DVS to the q -dimensional PVS. Here computational complexity arises from two distinct sources. If the individual design preferences $\mu_{d_1}, \dots, \mu_{d_n}$ are to be combined with a non-compensating aggregation function \mathcal{P}_{\min} , the combination of design preferences can be easily computed with a method such as the Level Interval Algorithm (LIA) of Dong and Wong (1987). For aggregation functions other than \mathcal{P}_{\min} , more specialized techniques are needed. Furthermore, when the evaluation of the mappings f_j from the DVS to the PVS is expensive, it is unrealistic to exhaustively search a DVS of more than a few dimensions.

The key limitation of the LIA, that it requires monotonicity, stems from the assumption that the extreme values of f_j will occur at the corner points of $D_{\alpha_k}^d$, the n -cube which is the α -cut at α_k in the DVS. The algorithm may be improved by relaxing this assumption (Mathai and Cronin 1995). The extended problem is to find the extremal values for the performance variables:

$$\begin{aligned} p_{j \min}^{\alpha_k} &= \min\{p_j = f_j(\vec{d}) \mid \vec{d} \in D_{\alpha_k}^d\} \\ p_{j \max}^{\alpha_k} &= \max\{p_j = f_j(\vec{d}) \mid \vec{d} \in D_{\alpha_k}^d\}. \end{aligned} \quad (\text{G1.13.6})$$

Finding extrema within a subset of the DVS is a constrained optimization problem.

In choosing an optimization technique to local these extremal values, a trade-off must be made between computational cost and robustness (*i.e.*, the ability to find the correct global extremum for various starting conditions). Traditional calculus-based optimization methods converge in relatively few function evaluations but seek only local minima. Randomized search methods such as genetic algorithms offer greater robustness (Goldberg 1989) but require more function evaluations. The computational implementation employed by the **MJ** uses Powell's method, a calculus-based optimization algorithm (Adby and Dempster 1974). An important feature for a practical computational tool is a means to trade-off the number of function evaluations against accuracy and reliability. Such an adjustment enables the designer to use the same program to obtain quick estimates as well as precise evaluations. This is implemented as a user-specified fractional precision that defines termination criteria for the optimization algorithm. Our extensions to the LIA are presented in detail in Law (1996).

The second difficulty can be partially overcome by selectively approximating \vec{f} as some simple function \vec{f}^j over D_{ϵ}^d (the α -cut at infinitesimal $\alpha = \epsilon$, which represents all designs under consideration). A linear approximation is not the only choice, but higher order approximations introduce additional complexity, both in the shape of the level sets mapped onto the PVS and in the computational algorithm, that is not often justified (Law 1996).

The approach adopted here is similar to *response surface methods* (Montgomery 1991), which seek to optimize a response that is influenced by several variables. The function f_j is modeled over the search space D_{ϵ}^d . The linear approximations f'_1, \dots, f'_q are obtained using techniques adapted from statistical design

of experiments. These techniques rely on orthogonal arrays, which specify an efficient, independent set of points at which the function is evaluated. Orthogonal arrays are widely used not only for statistical design of experiments but also for the related Taguchi's method or Robust Design methodology (Peace 1993, Phadke 1989) and their direct application to engineering design is not new (Chi and Bloebaum 1995, Korngold and Gabriele 1995).

Vehicle Structure Design Example

Vehicle structure design is concerned with dozens of design variables and many performance parameters from noise and vibrational response to stiffness to manufacturing cost to style. This case study examines ten design variables (see Figures G1.13.3 and G1.13.4):

d_1	0.10–0.20	B pillar gauge (thickness of hollow rectangular cross-section)
d_2	0.10–0.20	C pillar gauge
d_3	0.07–0.13	A pillar gauge
d_4	0.10–0.20	hinge pillar gauge
d_5	0.07–0.13	roof rail gauge
d_6	0.07–0.13	rocker gauge
d_7	0.03–0.05	floor gauge (plate thickness)
d_8	0.03–0.05	roof gauge
d_9	0.15–0.25	cross-sectional area of each cross-brace (square inches)
d_{10}	-2.0–2.0	fore-aft location of B pillar (fore is positive)

(all units are inches, except where indicated otherwise)

The ranges indicated are the full ranges under consideration for this example.




Figure G1.13.3. Schematic of the automobile body

In the example under discussion, as is often the case in the automobile industry, style drives the design. Here the styling directive is that the new look is lower and sleeker: the roof should be lowered, the clearance between the frame and the ground reduced, and the window opening between the A-pillar and B-pillar should be lengthened. In addition to this small but crucial styling change, the designers are asked to make 5% improvements over the bending stiffness and torsional stiffness of last year's model without increasing the weight, so that the increase in material cost reflects only inflation. The designers thus have in hand a set of functional requirements consisting of a sketch with some vague explanation, and the targets for four relevant performance parameters:

p_1	Bending Stiffness K_B (lbs./in.)	$K_B \geq 1.05K_B^{old}$
p_2	Torsional Stiffness K_T (ft-lbs./deg.)	$K_T \geq 1.05K_T^{old}$
p_3	Weight (lbs.)	$W \leq W^{old}$
p_4	Style	subjective




Figure G1.13.4. Finite element model of part of the automobile body

Finite element analysis is used to evaluate candidate designs for performance with respect to the quantifiable targets, and the attendant solid model will provide sketches for managers or stylists to evaluate the aesthetic impact. Since the finite element analysis is extremely costly, the M₀I applies a statistical regression (Law 1996) and arrives at the following approximations (which turn out to be linear in all but two variables):

$$K_B = 78,400 + \sum_{i=1}^8 a_{1i}(d_i - d_i^{\text{ctr}}) + 170d_9 - 240d_{10} - 630d_9^2 - 5d_9d_{10} - 88d_{10}^2 \quad (\text{G1.13.7})$$

$$K_T = 13,300 + \sum_{i=1}^8 a_{2i}(d_i - d_i^{\text{ctr}}) + 130d_9 - 38d_{10} - 620d_9^2 + 5d_9d_{10} + 4d_{10}^2 \quad (\text{G1.13.8})$$

$$W = 707 + \sum_{i=1}^9 a_{3i}(d_i - d_i^{\text{ctr}}) \quad (\text{G1.13.9})$$

where d_i^{ctr} is the value of design variable d_i at the center of the range under consideration, and the regression coefficients a_{ji} for performance variable p_j and design variable d_i are shown in Table G1.13.3.

design variable d_i	a_{1i}	a_{2i}	a_{3i}
1. B pillar gauge	10,800	20,700	567
2. C pillar gauge	25,100	4,110	460
3. A pillar gauge	33,700	8,340	390
4. hinge pillar gauge	28,500	938	276
5. roof rail gauge	161,000	15,900	260
6. rocker gauge	342,000	31,800	828
7. floor gauge	20,100	68,000	5,463
8. roof gauge	7,420	5,050	2,980
9. cross-brace area	non-linear		132
10. B pillar location	non-linear		0

Table G1.13.1. Linear regression results for performance variables.

By explicitly and formally modeling the preferences on design and performance variables as memberships in fuzzy sets, the M₀I allows for the formal resolution of issues that are commonly settled through informal negotiation. There are many such issues. Weight and stiffness will tend to increase

together, and the trade-off between these facets of performance is an example of a conflict between aspects of the design; a suitable balance point between the two depends upon the specific requirements on W , K_B , and K_T . If any of the target performances prove difficult to achieve, a negotiation between engineers and managers will ensue; aspects of this informal negotiation are captured in the formal representation of functional requirements. The incorporation of the desired styling is a prime example of a design requirement that would ordinarily be met through informal negotiation, since the engineers who effect the design will need to consult with the stylists as to the suitability of a completed design. In this illustrative example, we have ignored many complications, notably the consideration of manufacturing concerns at the design stage. Each additional complication could be addressed formally in a similar manner by the M₀I.

The first step in the formalization of the problem is the expression of more complete functional requirements through the specification of preferences $\mu_p(K_B)$, $\mu_p(K_T)$, and $\mu_p(W)$. The crisp constraints on W , K_B , and K_T are replaced by preference curves showing (more realistic) fuzzy requirements on these aspects of performance (see Figures G1.13.5 and G1.13.6). Designers are likewise polled in greater detail




Figure G1.13.5. Functional requirements on K_B and K_T




Figure G1.13.6. Functional requirement on W

as to their preferences for the various design variables. The information contained in the $\mu_p(p_j)$ already provides a starting point for trade-offs between aspects of performance using the formalisms of the M₀I. Figure G1.13.7 shows the overall preference for various values of W and K_B in a three-dimensional plot (as it is impossible to present three design variables clearly on one plot, K_T has been fixed at a nominal value); the vertical axis is the combined preference using a compensating aggregation function:

$$\mu_o = (\mu_p(K_B)\mu_p(K_T)\mu_p(W))^{\frac{1}{3}}$$

Any two or more candidate designs can be compared by examining the combined preference once the performances W , K_T , and K_B have been calculated.




Figure G1.13.7. Combined preference $\mu_p(K_B, K_T, W)$

A further complication alluded to earlier is the crucial role that styling plays in the evaluation of a design. Even if the design team is able find an “optimal” design in the sense of maximizing the combined performance shown in Figure G1.13.7, other aspects of the design’s performance, in this example: style, will need to be taken into account. Manufacturing concerns are another example of unmodeled performance that must be included. These *post hoc* evaluations are often critical to the success of the design as a whole: no matter how good the stiffness and weight numbers, a design that is perceived to be ugly or unmanufacturable will not leave the drawing board.

In this case the engineers have a simple measure that can guide them in their work. The stylists have expressed a preference for a wide window opening; this is interpreted as a preference for the nominal location of the B-pillar in the original solid model, as shown in Figure G1.13.8. The B-pillar need not be placed precisely in that location, but deviations from there carry a style penalty. A similar preference might be placed on the height of the roof. These preferences are *not* meant to replace the stylist’s judgment in his final analysis of the aesthetics of the design; they are a tool to guide the structural designer to avoid expensive redesigns. Since this preference is expressed on a design variable (B-pillar location), the M_{QI} treats it as a design preference. The location of the B-pillar will affect the stiffness; this will appear in the finite element calculation. There may well be other preferences for B-pillar location or for other design variables, for example, manufacturing concerns. The ostensibly continuous variables representing pillar gauge can be made effectively discrete if the designer expresses non-zero preferences for only those values that are available off the shelf. The computations of the M_{QI} take into account all of this information, with the possibility to assign weights and hierarchies. The overall preference for a design, calculated by the M_{QI}, thus contains the analysis that has been performed (in this case, the finite element analysis) as well as preferences for the aspects of the design, such as style, that are not calculated by the analysis tools.

G1.13.4 Development

A full finite element model for vehicle structure analysis takes several minutes on a Cray C90, while the approximate FEM takes about a minute on a Sun Sparc10 Workstation. The M_{QI} with its statistical extensions has been implemented in C. This implementation, some details of which are outlined in Law (1996), is of insignificant computational time compared to the simplified FEM.

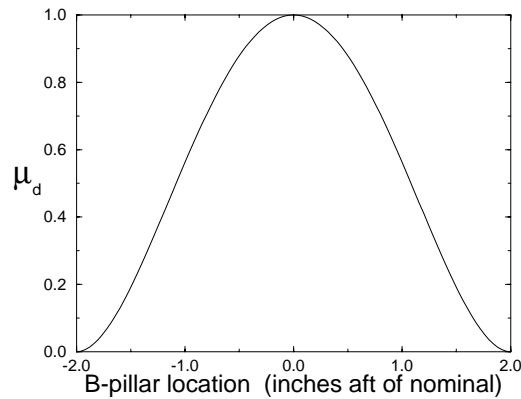


Figure G1.13.8. Stylists' preference for B-pillar location

G1.13.5 Results

For the example of the preliminary design of the automobile structure, results were calculated both with and without the incorporation of the styling preference on the location of the B-pillar. The various aspects of performance were determined to trade-off in a compensating fashion.

Both stiffnesses, and the preferences for them, increase with all design variables except for d_{10} , B-pillar location. Weight increases with these same variables, and so preference for weight decreases. For most of the variables, the peak performance is to be found at one or the other extreme of the design range:

d_1	0.18
d_2	0.10
d_3	0.07
d_4	0.10
d_5	0.13
d_6	0.12
d_7	0.03
d_8	0.03
d_9	0.15

The most preferred location for the B-pillar, considering only stiffnesses and weight, is $d_{10} = -2$, and the differences in overall preference among different locations are slight. When the styling preference is taken into account, the most preferred B-pillar location moves closer to the nominal value ($d_{10} = -0.03$), and the differences are greater.

It is not possible to represent the effects of all design variables on a single graph, but the effects of one or two design variables (with others held fixed at some nominal values) are easily shown. Figure G1.13.9 shows two overall preferences μ_o plotted against various B-pillar locations, with the other design variables held fixed. The optimal value for weight and stiffness alone, with no style considerations, is for the B-pillar location to be near -2, but it is seen that the variation in overall preference is small over the entire range. If style is also taken into account, the overall preference is strongly biased towards the nominal B-pillar location. As would be expected, this bias is reduced if the importance weighting for the style consideration is reduced.

Thus issues that would be considered informally in a typical design process, are explicitly calculated in the context of the M_oJ.

G1.13.6 Comparison with other methods

While internally the M_oJ makes some use of techniques from classical optimization and experimental design, its primary contribution is the use of fuzzy mathematics to model the inherent imprecision in the design process that is ignored by these other approaches. There is thus no computational benchmark available for comparing M_oJ calculations. The method's usefulness is demonstrated when the formal, computable, design




Figure G1.13.9. Overall preference for different B-pillar locations, without and with style considerations

support it offers eliminates the need for time-consuming, informal activities and negotiations to settle design questions.

G1.13.7 Conclusions

The Method of Imprecision (MI), briefly described here, forms the basis for formally representing and manipulating imprecise information in engineering design while simultaneously incorporating engineering designers' experience and judgment into design decisions. The underlying methodology builds upon fuzzy mathematics by mapping fuzzy sets (design information) through crisp functions (engineering analyses), and utilizing a family of idempotent mixed connectives for aggregation. When the engineering analysis is expensive, traditional computation of the extension principle is made more efficient by incorporating design of experiments methods. The designer can adjust the precision of these computations in order to accommodate different stages of the design process (from the most imprecise preliminary stage to the precise end result).

The method has been demonstrated in a range of applications, including the passenger vehicle structure design problem, one iteration of which was presented here. These results illustrate the ability of the method to incorporate incommensurate design information, along with subjective information (*e.g.*, styling) into a formal method for engineering design decision-making.

References

- Aday P and Dempster M 1974 *Introduction to Optimization Methods* (London: Chapman and Hall)
- Chi H W and Bloebaum C L 1995 Mixed variable optimization using Taguchi's orthogonal arrays *Advances in Design Automation - 1995* vol DE-82 501–508 (New York: ASME)
- Dong W M and Wong F S 1987 Fuzzy weighted averages and implementation of the extension principle *Fuzzy Sets and Systems* **21(2)** 183–199
- Goldberg D E 1989 *Genetic Algorithms in Search, Optimization, and Machine Learning* (Reading, Mass.: Addison-Wesley)
- Holmes M F 1984 Machine dynamics, The need for greater productivity *Research Needs in Mechanical Systems* ed. K N Reid 140–159 (New York: ASME)
- Keeney R and Raiffa H 1993 *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (Cambridge, England: Cambridge University Press)
- Korngold J C and Gabriele G A 1995 Multidisciplinary analysis and optimization of discrete problems using response surface methods *Advances in Design Automation - 1995* vol DE-82 173–180 (New York: ASME)
- Law W S 1996 *Evaluating Imprecision in Engineering Design* PhD thesis (Pasadena, Cal.: California Institute of Technology)
- Mathai T and Cronin D 1995 Application of fuzzy sets to car body structure NVH design. *Proceedings of IX International Conference on Vehicle Structural Mechanics and CAE* 107–115
- Montgomery D C 1991 *Design and Analysis of Experiments* (New York: Wiley)
- Otto K N and Antonsson E K 1991 Trade-Off Strategies in Engineering Design *Research in Engineering Design* **3(2)** 87–104

-
- Peace G S 1993 *Taguchi Methods: a Hands-On Approach* (Reading, Mass.: Addison-Wesley)
- Phadke M 1989 *Quality Engineering Using Robust Design* (Englewood Cliffs, N.J.: Prentice Hall)
- Scott M J and Antonsson E K 1995 Aggregation Functions for Engineering Design Trade-offs *9th International Conference on Design Theory and Methodology* ed. A Ward 389–396
- Sullivan L P 1986 Quality function deployment. *Quality Progress* **June** 39–50
- Ulrich K T and Pearson S A 1993 Does product design really determine 80% of manufacturing cost? Working Paper MSA 3601-93 (Cambridge, Mass.: MIT, Sloan School of Management)
- Von Neumann J and Morgenstern O 1953 *Theory of Games and Economic Behavior* (Princeton, N.J.: Princeton University Press)
- Ward A C, Lozano-Pérez T and Seering W P 1990 Extending the constraint propagation of intervals *Artificial Intelligence in Engineering Design and Manufacturing* **4(1)** 47–54
- Ward A C and Seering W P 1989 The performance of a mechanical design compiler memo 1084 (Cambridge, Mass.: MIT Artificial Intelligence Laboratory)
- Ward A C and Seering W P 1989 Quantitative inference in a mechanical design compiler memo 1062 (Cambridge, Mass.: MIT Artificial Intelligence Laboratory)
- Whitney D E 1988 Manufacturing by design *Harvard Business Review* **66(4)** 83–91
- Wood K L and Antonsson E K 1989 Computations with Imprecise Parameters in Engineering Design: Background and Theory *ASME Journal of Mechanisms, Transmissions, and Automation in Design* **111(4)** 616–625