# A TECHNIQUE FOR CAPTURING STAKEHOLDER PREFERENCES AND KNOWLEDGE IN EARLY DESIGN STUDIES

**Sky Sartorius, Mirko Hornung****
****Technical University of Munich Institute of Aircraft Design**

**Keywords**: *Design process, Optimization, Conceptual Design*

## Abstract

*With the aim of accelerating early design studies and reducing design iteration cycles, a technique was developed that leverages explicitly captured* a priori *stakeholder knowledge in the form of preference maps. To put these preference maps to use, they are applied in transforming the model of the physical system into a standard-format objective function that can be used with the power of existing numerical search algorithms. The search results are a set of design points likely to be of interest for exploration in further design cycles, and various tools associated with the technique aid in understanding of the design tradeoffs and the design decision-making to follow.*

*Of the test cases used in developing the approach, the one presented here is an early exploration of a winglet retrofit for a narrow-body airliner. The technique demonstrated successfully producing insights under the uncertainty of early design studies, producing useful results for a relatively low effort compared to investing further into more mature system models and requirements. It is also an approach that enables reaping some of the benefits of numerical search algorithms much earlier in design than would normally be appropriate.*

## 1 Introduction

Computers and software play a fundamental and intrinsic role in contemporary aircraft design and trade studies. At every stage of a project, not only have computers replaced manual calculations, but often software is written even for the simplest of tasks. From day one of a new aircraft design investigation, for example, conducting initial 'back of the napkin' hand calculations has been replaced by writing formulae into a spreadsheet or a simple script. In this way, the designer not only has access to the results of these initial calculations but has simultaneously created a useful tool whereby the calculations can be repeated with different inputs and assumptions with virtually no cost. Because software is ever-present in design as a tool to be wielded by the designer, it is worth exploring techniques that have the potential to augment the mutualistic and complementary interaction between designers and their software.

Many capabilities are available to designers by the simple presence of nearly-free repetitions of a given analysis. One prominent example is automatic searching of the design space, i.e. optimization, using various search algorithms to find an optimal design. Optimization is an excellent tool that has a unique capacity for handling design problems with many variables of interest that all interact (a common situation in aeronautical design). While it is quite straightforward to produce plots and visualizations that allow a designer to fully understand a problem with only a few design variables and a handful of dependent parameters and make an informed design decision, most search algorithms are well-suited to scaling up to larger problems, which a human designer could only grasp a few dimensions at a time. Because of this, optimization also enables exploring these large design spaces, not only to find optimal solutions, but also often to enable or accelerate the discovery of feasible solutions to highly constrained problems, as well as identify and gain insight into design constraints.

This work explores specifically an approach to formally capturing *a priori* stakeholder information and leverage numerical search as a means of making use of that information.

## 1.1 Issues with Optimization in Early Design Studies

An early design study refers here to the early stages of any trade study or exploration of the design space. This could be the first rough sizing estimate made in the first days after exposure to a new set of air vehicle requirements or decision to pursue development of a new product. Alternatively, an early design study could simply be the initial investigations into some aspect of a vehicle or system that has not been previously given significant attention in the overall development. In any case, there are some characteristics of early design studies which pose a challenge to leveraging design optimization in the process.

The first such characteristic is that the requirements and/or objectives for the project are unlikely to be extremely well defined and understood. This is especially true for product design as opposed to, e.g., a traditional defense program in which the objectives and metric(s) to maximize or minimize are explicitly stated. Even then, the explicitly-stated objectives are only a best effort by the author of the requirements, and a faithful adherence to those objectives may not result in fulfilling the 'spirit' of the program goals.

Conventional wisdom regarding the analytical models used in optimization is captured well by [1]: "the underlying analysis must properly model the true physics or optimization will generate unrealistic designs." Besides uncertain requirements, another (and arguably more universal) trait of early design studies is a lack of maturity and/or fidelity of the analytical models in use. Models appropriate for conceptual aircraft design, for example, by their very nature do not capture the effects of all possible design parameters for the simple reason that these parameters are unknown at this stage. One fortunate side benefit is that these models used in early design tend to be computationally inexpensive, which helps ease what can be one of the pain points in search algorithms: long computation times.

Another reason that models in early design studies may lack maturity is simply that, especially when unique and novel design problems and/or solutions are involved (or problems or solutions novel to the particular organization), the models are in the process of being built in parallel to the design effort itself. These under-construction models are also less tested and validated, and if used in optimization have a significant chance of leading the design astray through unaccounted for responses to certain combinations of inputs.

A final characteristic of early design studies that poses a challenge to using optimization is that there is less-than-perfect alignment between the fundamental aims of the processes. Mathematical optimization, by its nature, is focused on finding optima and, in most cases, a single optimal design solution according to an all-encompassing objective function that captures all value factors for all stakeholders. In early design studies, by contrast, because it is certain that any design point chosen will change, the designer is more interested in identifying 'interesting' design solutions and regions of the design space for further investigation and informing human-in-the-loop decision making than in finding the single *best* design point. The designer is also more interested in increasing understanding of the design space, rather than letting an optimization algorithm explore the design space behind a veil of obscurity and abstraction. Finally, especially when working on a novel design problem or solution, the designer may be more focused in this stage on building and maturing the aforementioned analytical models.

## 1.2 Technique Objectives

The aim of the technique is to capture as best as possible the preference information designers normally have *a priori* but typically only use to make design decisions *after* the design space and trade-offs are better understood. This type of preference information is normally not used in multidisciplinary design optimization (MDO). To illustrate the role this preference information plays and the effects it has on the design process, consider a simple abstracted example of an early trade study for a new aircraft design project: payload capacity of the vehicle (which the customer already specified as a requirement or preference) vs. the operational cost per unit of payload delivered (the notional value function to be minimized to yield the 'best' system), shown in Figure 1.

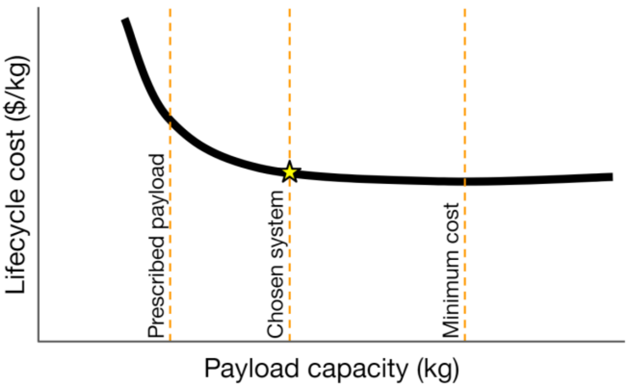


**Fig. 1. Illustration of early payload capacity trade study.**

In this situation, a significant cost improvement would be possible with a system designed with several times the payload capacity the customer originally requested. In the hypothetical case that optimization would be used for this study, the resulting design point would either have the customer-prescribed payload (if this were set as a constraint) or it would be a system many times larger than what the customer originally requested and expected. However, the design choice was neither, instead selecting a point in between.

What preference information was available to the designer before the shape of the curve was known that led to the decision? This question is the inspiration for the technique presented here. The objective of the technique is to capture this type of (often non-linear) preference information in the form of explicit semi-quantitative preference maps that can then be combined with the power of computational tools, specifically search algorithms, to provide a useful result. In this way, the aim is to accelerate the design and decision-making process and reduce design iteration cycles with automatic searching of the design space that is as informed as possible to allow the search algorithms, at every step, to be driven in the same way a human-in-the-loop designer would be.

## 2 Implementation

To develop the proposed technique, a framework (Figure 2) was built using object-oriented MATLAB and two main classes. A single top-level design space object for a given project contains a vector of design variable objects. The design variables are inputs and/or outputs to a separate system model function.

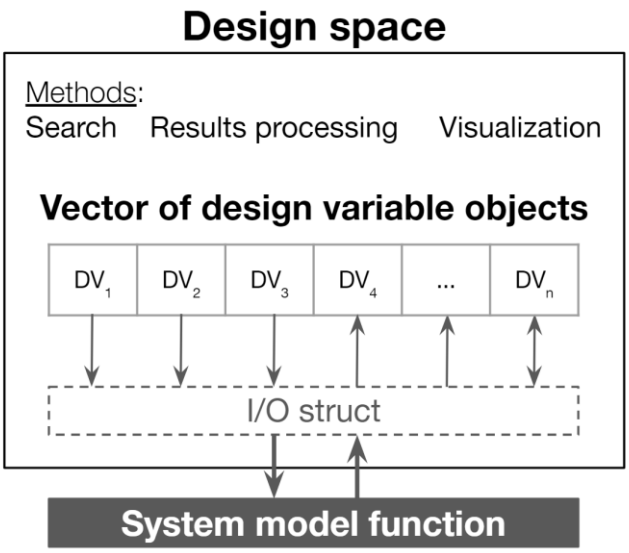


**Fig. 2. Framework architecture.**

### 2.1 System Model

The design space points to a system model function, which may have an arbitrary number of inputs and outputs. It is written in or wrapped by a MATLAB function that takes a struct as input and returns the same struct with added and updated fields for outputs. The fields of the struct may also be arbitrary; users are not restricted to pre-standardized system parameterizations, such as CPACS[2] or ADDAM[3] (though nothing precludes their use internally in the system model). Inputs and outputs can be floating point numbers, discrete numeric or non-numeric values (e.g. strings), or dimensioned variables that carry units. This flexibility allows for faster and more natural building of the system model compared to a traditional MDO setup, which normally requires that the target functions accept a single vector of continuous numeric inputs.

## 2.2 Design Variables

The core property of the design variables in the framework is the preference map (blue line in Figure 3), which is simply a lookup table of a penalty or value as a function of a system parameter value. The ordinate of a preference map is arbitrary and can be a dollar cost, for example, or a fully abstracted points system, as long as it is consistent across all design variables.

While it is possible to build the maps with hard step changes at step changes in value (for example when wingspan breaks an ICAO code limit), experiments revealed that the technique produces slightly more useful results when the maps are built with ramps instead of steps, creating drivers for the automatic search to make ‘decisions’ similar to those that a human designer would make, but at the multidimensional scale of MDO instead of only considering the one to about five parameters that typical humans are capable of considering simultaneously (for example with a carpet plot trade study).

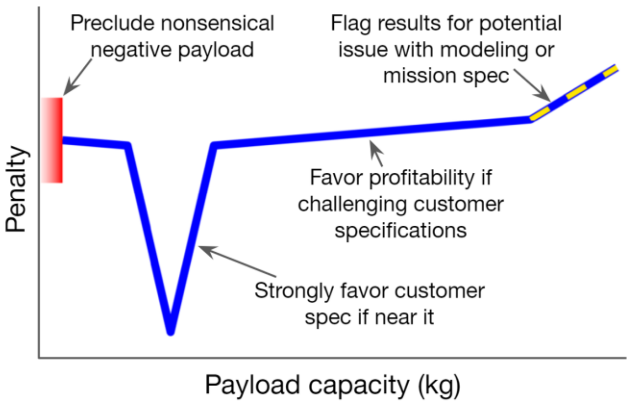


**Fig. 3. Example preference map for payload capacity trade study from Fig. 1.**

Each node and segment in a preference map can be tagged with an arbitrary amount of additional information, with the normal use case being linking to requirements, either informally or with a formal requirements tool such as IBM Rational DOORS or other requirements management schemes such as that presented in [4]. This helps make the defined design variable objects a useful source of documentation as well as reusable artifacts for future projects where some of the same drivers may be present.

There is a variety of other information captured in the design variable class. These other design variable properties either capture additional *a priori* information, such as the level of uncertainty in the scale of a particular preference map (see below), or simply help facilitate working with the design space.

## 2.3 Design Space Search

To leverage the multitude of existing off-the-shelf search algorithms, design variable preference information is transformed into a standard-format single-objective optimization problem:

$$\min_{\boldsymbol{x}} f(\boldsymbol{x}) \text{ such that } \begin{cases} g(\boldsymbol{x}) \leq 0 \\ h(\boldsymbol{x}) = 0 \\ \boldsymbol{x}_l \leq \boldsymbol{x} \leq \boldsymbol{x}_u \end{cases} \tag{1}$$

The transformation for the objective function to include the preference maps is given by the sum of design variable preferences:

$$f(\boldsymbol{x}) = \sum_{i=1}^{n} p_i\big(s(\boldsymbol{x})\big), \tag{2}$$

where $s(\boldsymbol{x})$ is the system model function that analyzes a design defined by input vector $\boldsymbol{x}$ (along with some input processing to transform $\boldsymbol{x}$ into a usable input struct for the system model function) and $p_i$ returns a preference value based on the preference map for the design parameter returned by $s$ corresponding to design variable $i$.

Upper and lower bounds are implemented as side constraints for input design variables and as nonlinear constraints for bounds on output design variables. Design variables that are both inputs and outputs and that must be consistent, for example an estimated weight input and a calculated weight output, are also enforced automatically with nonlinear constraints, removing the requirement for the system model function to converge internally.

To generate interesting *sets* of multiple design points, the technique seeks to find many unique *local* minima by searching from many different start points, $\boldsymbol{x}_{0,j}$. Note that this is contrary to the goal common in MDO to find a single, globally optimal solution, so some variation and tuning in the application of search algorithms is required. While it is possible to use a wide array of existing search algorithms without modification, because the aim here is to find local minima in an intentionally non-convex problem, the application of most algorithms can be outside the original algorithm intent, and future work includes exploring new algorithms suited and tuned for this type of application.

The other mechanism for adding variation to the resulting set of design points leverages a design variable property that is a measure of uncertainty in the preference map. Equation (2) becomes:

$$f_j(\boldsymbol{x}) = \sum_{i=1}^{n} \rho_i^{X_{i,j} \sim U(-1,1)} \cdot p_i\big(s(\boldsymbol{x})\big), \quad (3)$$

where in addition to each search using a different starting point, $\boldsymbol{x}_{0,j}$, the uncertainty factor for the design variable, $\rho_i$, is used to scale the ordinate of the preference map stochastically based on random number $X_{i,j}$ between -1 and 1. An uncertainty factor of 2, for example, means that the preference map may be scaled by between ½ and 2. A flowchart of the implementation of Equation (3) is shown in Figure 4.

Observe that if all preference maps are linear, the technique simply becomes a weighted sum method for multi-objective optimization. In the case that the system model function is also smooth and convex, only the uncertainty factors, not the varying start points, lead to variation in the resulting design point set (if not all $\rho = 1$).

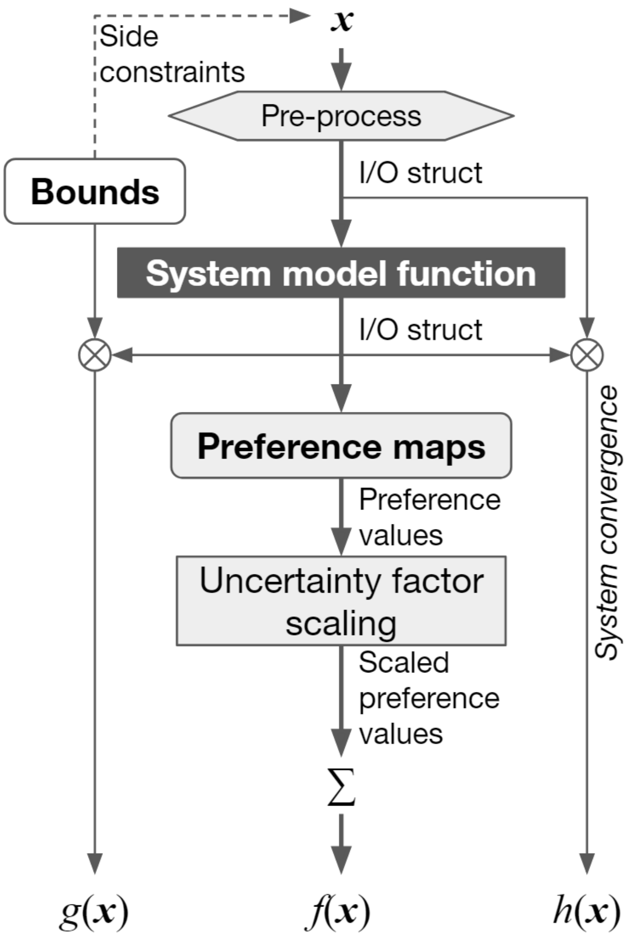

**Fig. 4. Implementation flowchart.**

## 3 Development Use Case

Different applications for aircraft generally have more or less well-defined metrics in development. Small general aviation aircraft are at one end of this spectrum, as they are often being purchased and operated by individuals for whom financial costs of aviation are of secondary importance to emotional or aesthetic factors. So, the first development test case for the technique was an initial sizing study (payload, range, etc.) of a personal use, one-off home-built experimental general aviation airplane.

While a personal homebuilt experimental aircraft design represents a case with early, uncertain, and fungible requirements, at the other end of the spectrum are systems like airliners, where there is an existing and well-understood business model and a small set of clear, overarching value functions to maximize or minimize (costs and profits), often supported by market research. The second development test case (and the case used here for a deeper illustration of the preference mapping technique) was a winglet retrofit study with a system model that is very simple and does not go all the way to capturing a top-level figure of merit like direct operating cost. The design spaces for both test cases are relatively small and simple, but the uncertainty that the tool addresses is quite different for each.

### 3.1 Winglet System Model

The system model function for this use case analyzes a potential winglet retrofit on a narrow-body airliner. Though the airliner industry has well-defined objectives and many well-established models, the modeling used for this test case is intentionally simple to mimic the lack of such fidelity in early design studies. The primary analysis for lift, induced drag, and bending moments is based on the vortex lattice method as implemented in Athena Vortex Lattice (AVL) software [5], with simple models or surrogates used to yield other parameters for use with preference maps (weight and parasite drag, for example).

For this study, the modifications are limited to the tip section of the wing that is entirely outboard of the aileron and slat (Figure 5). As the analysis is subsonic, the leading edge sweep of both the winglet (if present) and the tip section are enforced to be at least that of the main wing. The main parameters to adjust are the tip panel geometry, the winglet size, and the winglet twist and incidence. Weight change is based on a simple areal weight parameter, and parasite drag change is based on a constant section drag coefficient.

### 3.2 Design Variable and Preference Maps

A total of sixteen design variables were defined for the winglet study. Six of those are shown in Figure 6. The preference maps were scaled based on cruise drag, i.e., one percent of cruise drag reduction was thought of as a unit of 'currency' when building the preference maps (and hence the linear preference map for cruise drag reduction).

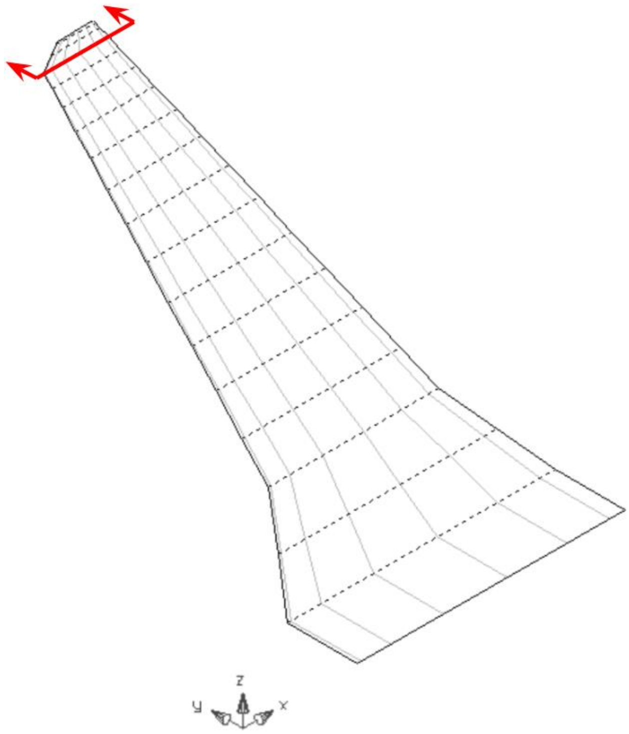


**Fig. 5. Narrow-body airliner half-wing AVL model with tip modification region marked.**

Because the system model contained no structural analysis, preference maps for wing root bending moment and winglet structural aspect ratio were used as surrogates, with any root bending moment increase greater than 10% flagged as likely requiring a more major structural redesign effort. The wing span preference map captures the significant penalty for exceeding the ICAO Code C maximum of 36 meters, with a hard bound at the Code D limit. The $x_{spar}$ preference maps drive favorable design adjustments in the case that a main wing spar is close to aligning with the thick part of the winglet root section. The winglet span preference map captures that if the best design has only a very tiny winglet, it is desirable to simplify and not have any winglet at all. Uncertainty factors typically between two and three were set for design variables other than cruise drag change.

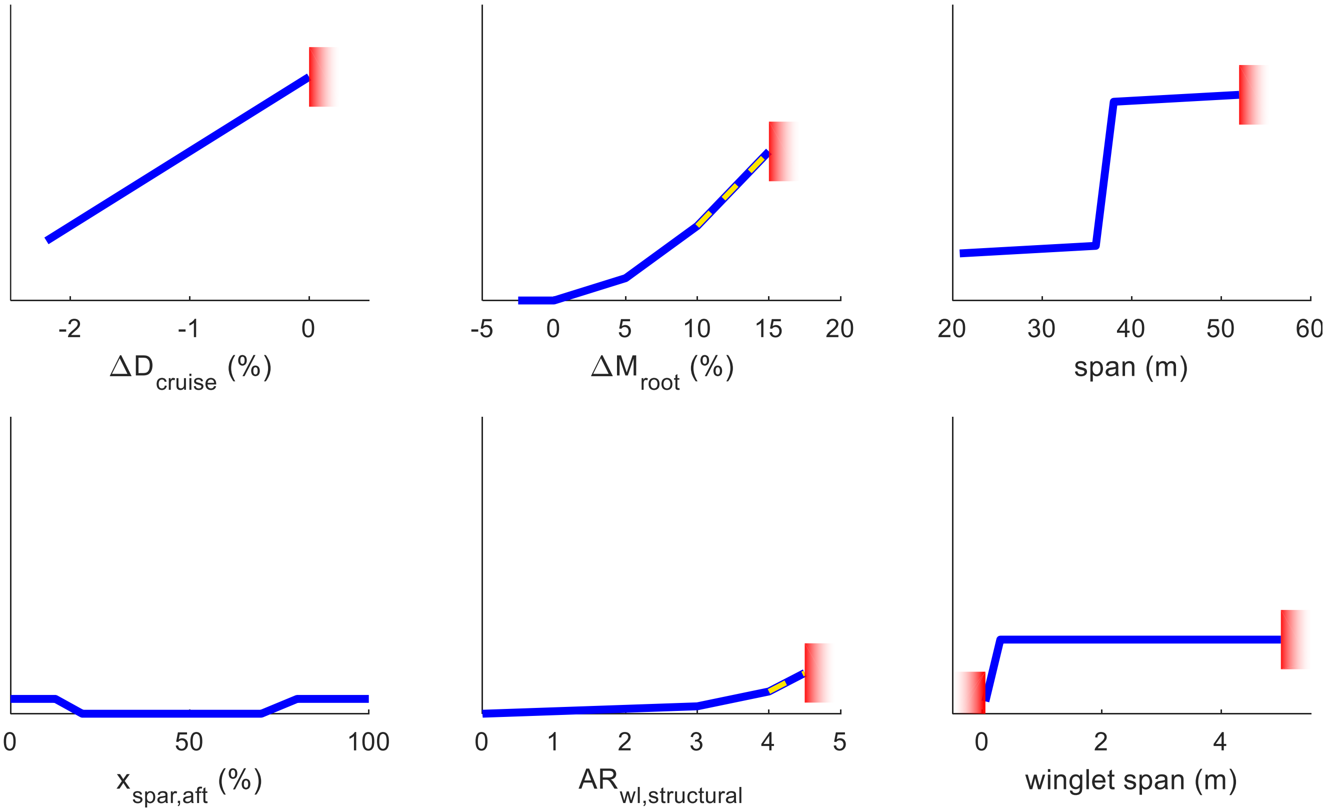


**Fig. 6. Selected preference maps used in winglet study. All ordinate axes are the same penalty scale.**

### 3.3 Use Case Results

The design space class, in addition to coordinating search to generate interesting sets of design points, also provides tools to process and visualize those results to help facilitate making decisions and moving the design process forward.

After a search of the design space, the result is a large set of design points. Some processing methods are available to make the set more manageable. In the (rare in tests so far) case that there are points that are completely dominated by another (better preference map value for every variable), they can be removed, along with searches that failed for one reason or another. Clustering algorithms were also explored, with the most useful being those that preserve outliers, since those are often design points of interest.

Visualizing a multi-dimensional design space is quite useful for providing a deeper understanding of the system for stakeholders. Sometimes the density of resulting points can be valuable, in which case a corner plot (implemented by [6]) is useful. There are also tools built in to support simply creating a spreadsheet data table of the various designs and their attributes. However, thus far the most commonly valuable tool, inspired by and implemented similarly to the multi-objective genetic algorithm results visualization tool in [7], is the scatter matrix augmented with data brushing. This visualization of the results for the winglet study is shown in Figure 7, where the marker color corresponds to the sum of preferences before uncertainty factor scaling and flagged design points are denoted by smaller marker size.

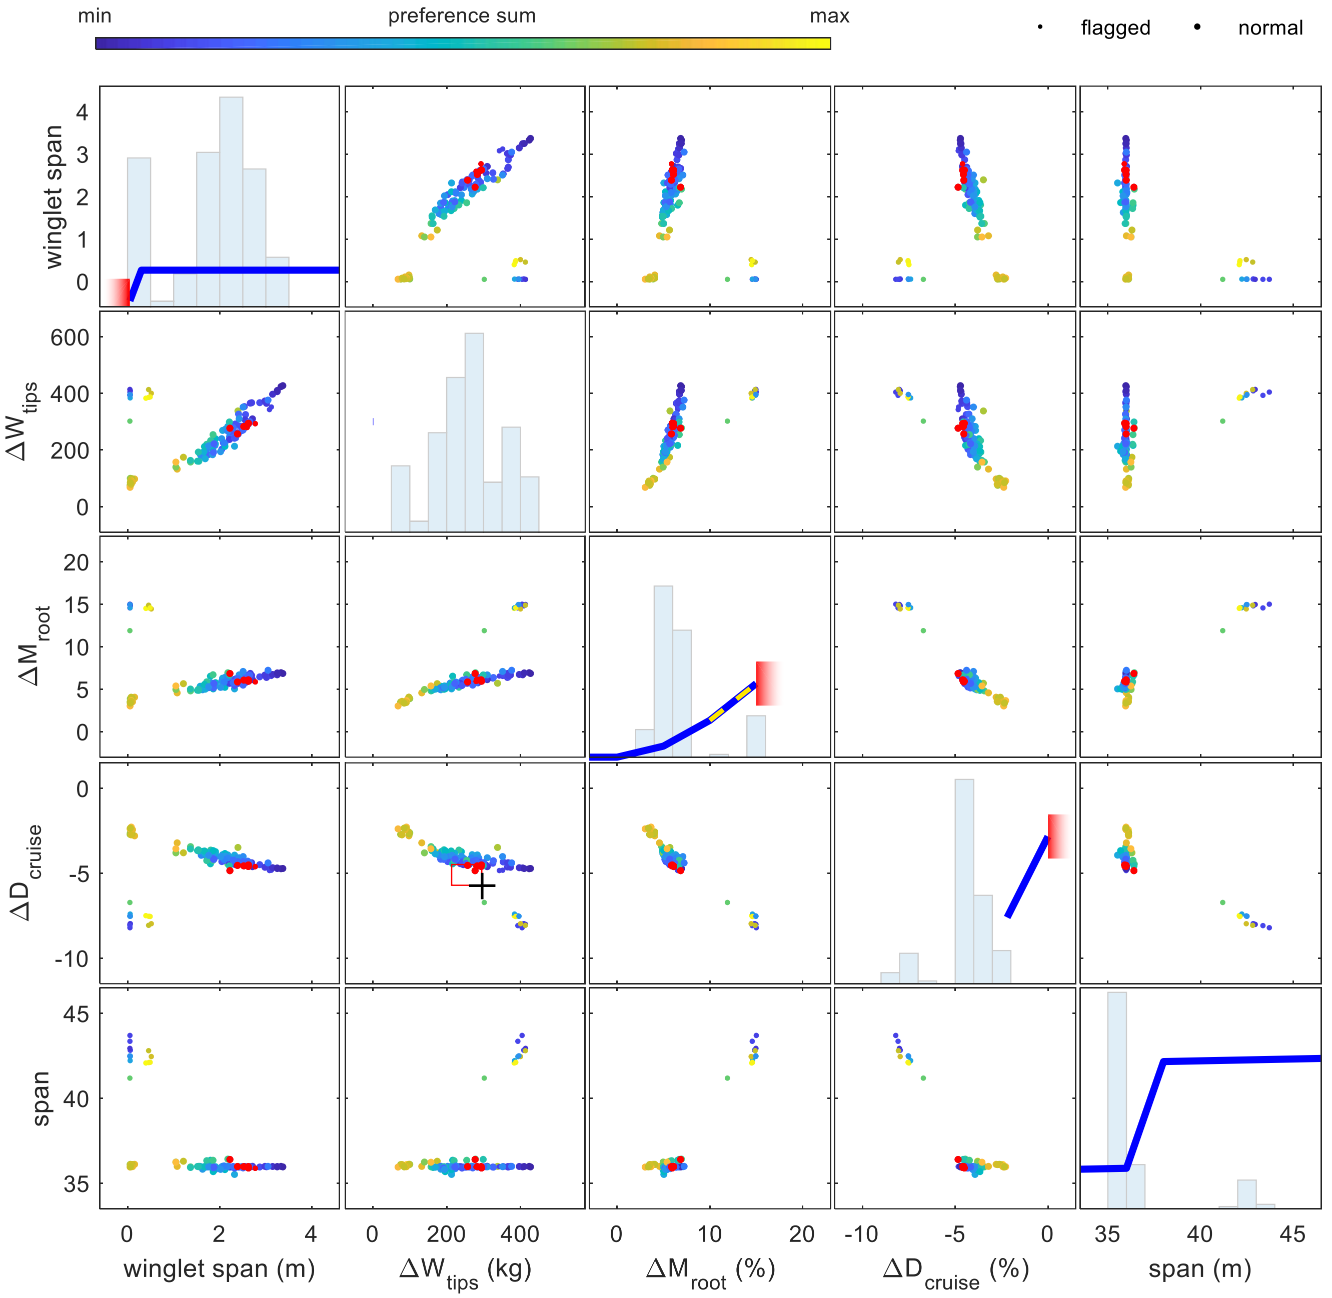


**Fig. 7. Winglet study scatter matrix results visualization with example data brushing (red). The main diagonal shows preference maps for reference, along with underlaid histograms to indicate distribution and density of design points.**

The results reveal a few large clusters of design points in interesting regions of the design space (with some variation within clusters as well). Representative wing geometries for the three clusters is shown in Figure 8. Figure 8(a) is representative of the large cluster of results that feature a prominent winglet, 8(b) depicts a simple span extension up to but not violating the ICAO Code C limit, and 8(c) shows a design with very high bending moments (flagged ahead of time as likely being unmanageable) but very good drag performance.

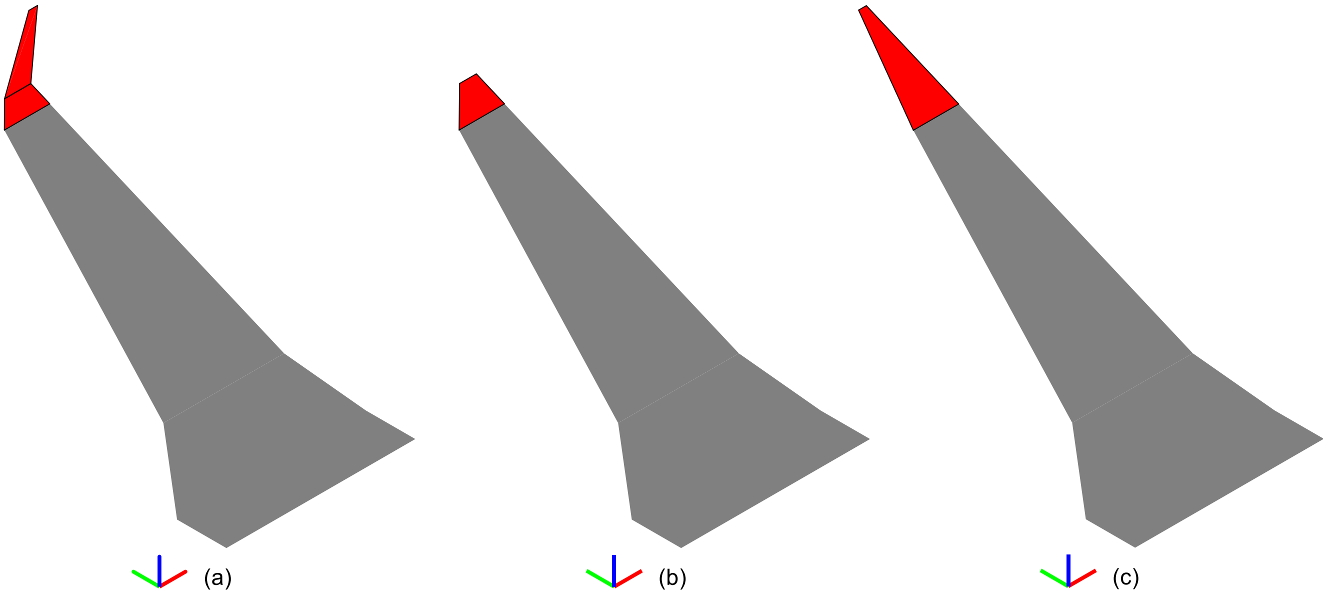


**Fig. 8. Isometric views of representative unique wing geometries from each of the major results clusters (triads are one meter for scale).**

## 4 Conclusions

The results shown above are not the final step of the process. Rather, the important next step is for the designer and other stakeholders to use the results to make human-in-the-loop design decisions that narrow down the design space and allow proceeding with further design iterations. The technique enables this by generating qualitatively distinct alternatives and presenting only those that are likely to be of interest for further exploration. In addition, through tools such as the scatter matrix with data brushing, the technique enables exploration and understanding of the quantitative tradeoffs between designs. These sets of information yield the insights needed for quick, informed design decision-making and accelerated design cycles.

This is made possible in part by using the power of numerical search and optimization much earlier, when normally the system model and/or requirements are too immature to be appropriate for use with optimization. With the technique, however, the process can proceed under significant but explicitly captured uncertainty. As an example, when the winglet study use case was set up as a weighted sum multi-objective optimization, the one design point returned was a high-aspect ratio, high-span, and very high bending moment design similar to Figure 8(c).

Another key benefit of the approach is the relatively low amount of effort required to build preference maps in comparison to refining requirements and system models (along with additional benefit of the forced critical thinking exercise that comes with building preference maps). The resulting winglet geometries in the use case presented here, for example, are quite reasonable and realistic without having any sort of highly informed analysis for structures or weights. The preference maps used as surrogates were built using 100% *a priori* knowledge and therefore did not require any research and development to generate. The technique provides an alternative path to putting effort into market research to lock in requirements or building more and more in-depth system models, and therefore the technique can be a significant improvement in return on invested effort by producing similar design results using low-effort preference maps in lieu of high-effort requirements and models.

Initial experience using the technique with the two simple test cases has shown enough promise to warrant further development and exploration of the approach. The architecture of the tool makes it a useful framework for other design tasks, such as Monte Carlo probabilistic design, more conventional optimization setup (or just simply design convergence setup), and generating visual artifacts useful for inclusion in reports, design reviews, etc. Future development specific to the technique presented here will focus on more use cases to explore inclusion of a wider variety of stakeholders (including non-technical) in the workflow, scalability to slightly higher-dimensional search spaces, and applicability to later stages of design. Regarding the latter, it is anticipated that the usefulness of the technique will diminish at later stages of design when uncertainty is significantly reduced, requirements are solidified, and system models much more closely capture true physics.

## References


[1] Vanderplaats, G. *Multidiscipline Design Optimization*. 1st Edition, VR&D, 2007.

[2] Nagel B, Böhnke D, Gollnick V, Schmollgruber P, Rizzi A, La Rocca G, and Alonso J. Can we establish a common language? *28th International Congress of the Aeronautical Sciences*, Brisbane, Australia, 2012.

[3] Herbst S and Hornung M. ADDAM: an object oriented data model for an aircraft design environment in MATLAB. *AIAA Modeling and Simulation Technologies Conference*, Dallas, TX USA, 2015.

[4] Glas M and Sartorius S. Towards a continuous build-up process of a reusable requirements-based system model. *IEEE Aerospace Conference*, Big Sky, MT, USA, 2012.

[5] Drela M and Youngren H. Athena Vortex Lattice (AVL) ver. 3.35. *Computer software*. 2014. Retrieved from web.mit.edu/drela/Public/web/avl.

[6] Adler W. cornerplot ver. 1.5. *Computer software*. 2018. Retrieved from github.com/wtadler/cornerplot.

[7] Sartorius S. Optimization and design space visualization applied to early conceptual design of compound helicopters. *52nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Denver, CO USA, 2011.


## Contact Author Email Address

Sartorius@TUM.de

## Copyright Statement