A similarity-assisted multi-fidelity approach to conceptual design space exploration

(article starts on next page)

Computers in Industry 151 (2023) 103957

Contents lists available at ScienceDirect

Computers in Industry

journal homepage: www.sciencedirect.com/journal/computers-in-industry

A similarity-assisted multi-fidelity approach to conceptual design space exploration

Julian Martinsson Bonde [a,*], Michael Kokkolaras [b], Petter Andersson [c], Massimo Panarotto [a], Ola Isaksson [a]

[a] *Chalmers University of Technology, Gothenburg, Sweden*
[b] *McGill University, Montreal, Canada*
[c] *GKN Aerospace Engine Systems, Trollhättan, Sweden*

ARTICLE INFO

Keywords:
Engineering design
Design support
Design space exploration
Similarity metrics
Aero-engine structures

ABSTRACT

In conceptual design studies engineers typically utilize data-based surrogate models to enable rapid evaluation of design objectives that otherwise would be too computationally expensive and time-consuming to simulate. Due to the computationally expensive simulations, the data-based surrogate models are often trained using small sample sizes, resulting in low-fidelity models which can produce results that are not trustworthy. To mitigate this issue, a similarity-assisted design space exploration method is proposed. The similarity is measured between design points that have been evaluated through lower-fidelity data-based surrogate models and design points that have been evaluated using higher-fidelity physics-based simulations. This similarity information can then be used by design engineers to better understand the trustworthiness of the data produced by the low-fidelity surrogate models. Our numerical experiments demonstrate that such a similarity measurement can be used as an indicator of the trustworthiness of the lower-fidelity model predictions. Moreover, a second similarity metric is proposed for measuring the similarity of new designs to legacy designs, thus highlighting the potential to reuse knowledge, analysis models, and data. The proposed method is demonstrated by means of an aero-engine structural component conceptual design study. An open-source software tool developed to assist in data visualization is also presented.

1. Introduction

When exploring conceptual design spaces that are large and high-dimensional, Aerospace design engineers often rely on data-based surrogate models of a low fidelity. This is because models of a higher fidelity, such as finite element method simulations, are either not available or computationally expensive, which hinders rapid assessment of a large number of design points. In the latter case, Aerospace design engineers often resort to low-fidelity surrogate models built using data obtained by high-fidelity simulations with a relatively low number of sample points of the design space (Martins and Ning, 2021). The predictive capability of surrogate models can theoretically be improved by increasing the sample size. However, this is often impossible due to restrictive computational budgets and time-constrained project situations. Hence, the problem addressed in this paper is that design engineers commonly utilize underperforming surrogate models without any feedback on data trustworthiness.

In this paper, we propose a method that utilizes similarity metrics (Lin, 1998) for indicating whether low-fidelity predictions can be trusted in a specific design context. Of particular interest are high-dimensional surrogate models that have been built using a relatively small number of samples, as such setups are prone to encountering the phenomenon often referred to as *the curse of dimensionality*. This essentially means the samples needed to train a surrogate model increases exponentially with the dimensionality of the model (Wang and Shan, 2006; Keogh and Mueen, 2011). A sample size that is too small, with respect to the dimensionality, will thus result in a low surrogate model prediction accuracy.

In addition to using similarity metrics to assess the trustworthiness of surrogate model predictions an additional metric is proposed to identify similarities to already existing designs. The ability to identify and leverage similarities to existing or legacy products can be especially useful within the aircraft engine manufacturing business. When designing structural components for aircraft jet engines, the trade-off between functional performance and cost of realization is critical already in the early design phase. This is because such trade-off analysis requires in-depth knowledge about the concept definition and behavior. Structural

* Corresponding author.
E-mail address: julianm@chalmers.se (J. Martinsson Bonde).

https://doi.org/10.1016/j.compind.2023.103957
Received 18 January 2023; Received in revised form 22 May 2023; Accepted 23 May 2023
Available online 7 June 2023

aero-engine components are subject to harsh operative environments, and the weight of such components need to be minimal to reduce fuel consumption. This leads to advanced geometry optimization, which can consequently result in structures that are expensive to manufacture. To evaluate these aspects requires detailed designs, but such designs does typically not yet exist during the early design phase. Thus, there is a need to learn from existing designs, which similarity measurements can assist in addressing.

Furthermore, similarity to existing designs can be a problem as existing designs may be unnecessarily expensive and problematic to produce. Consider a scenario where a legacy product encountered technical issues during manufacturing, which caused late expensive redesigns. Any measured similarities to this historically problematic product can prompt the designers to reconsider their designs, to avoid repeating past issues. In other words, measuring the similarities to existing and legacy designs, as is done in the proposed method, can not only be used to identify "useful" similarities, but also similarities that should be avoided.

In the aircraft engine manufacturing business, a Technology Readiness Level (TRL) of 6 is often considered as necessary to start commercial product development. Conceptual work implies that known design solutions that have been used in flight (that is, have a TRL of 7 or higher) are considered mature and of low risk (Niemeyer and Whitney, 2002). However, by definition the assessed TRL of any product reduces the more the target context differs from the existing one. Nergrd and Larsson (2009) proposed to use previous designs, and knowledge gained from analysis models, to judge the quality of a TRL assessment of new designs and technologies. Thus, how similar the low-risk "proven in flight" solutions are compared to the product to be designed needs to be understood.

The aim of this study is thus to address two opportunities: (1) How similarity metrics can be used to mitigate issues typically encountered when utilizing surrogate models based on limited sample sizes, and (2) How similarity metrics can be used to highlight similarities to existing design solutions with a high degree of maturity, providing an indication of potentials for reuse of data, analysis models, and knowledge. The intended effect of addressing these opportunities is to reduce the time allocated to high-fidelity simulations, and to improve the efficiency of identifying promising design configurations.

In Section 2 the relevant theoretical background is summarized. Section 3 contains a detailed description of the proposed method. In Section 4 the results of an experiment conducted to verify the utility of the method are presented, along with an application of the proposed method to a case from the Aerospace industry. Additionally, to enable demonstrating the method fully, an open source software tool has been developed (Martinsson Bonde, 2023) to assist in measuring and visualizing similarities in design studies.

2. Theoretical background

It is common practice for design engineers to assess alternatives in the early phases of the design process by means of low-fidelity multidisciplinary analysis models. Typically, design engineers are not concerned about this fidelity compromise; they tolerate and address it in further studies of increasing detail and fidelity downstream the design process. Nevertheless, having a means to provide some sense of confidence in analysis results can be of tremendous help to design engineers while they are navigating vast design spaces with disproportionately available information. We propose a method that utilizes similarity metrics to enhance the confidence of design engineers in the above mentioned task. The literature in multi-fidelity modeling, multidisciplinary analysis (and optimization), and similarity metrics is vast; a detailed review would be out of the scope of this paper. Therefore, we review here the principles and methods that are directly relevant to our work.

2.1. Multidisciplinary assessment in design space exploration

Woodbury and Burrow (2006) define design space exploration as the computational assessment of relatively large quantities of points in the design space. How the designs are represented can vary; a common means of representation is through CAD-models that can be generated automatically. This was exemplified by Sandberg et al. (2017) who evaluated aero-engine design concept variations via the generation of solid CAD model representations.

When assessing different design concepts, it is often necessary to consider their performance in different disciplines. When developing an aero-engine structure, the final geometry needs to be aerodynamic, structurally sound, of low weight, and manufacturable. To cater to such varying and often conflicting needs, simulations can be used to generate data, which can be visualized and analyzed by cross-functional teams during design space exploration. Wall et al. (2020) demonstrated such a working procedure in their "decision arena", within which cross-functional teams make joint-decisions based on the visualization of multidisciplinary data.

In addition to data from simulations, experienced engineers often leverage knowledge gained through previous projects when making decisions regarding new designs (Smith and Duffy, 2001). For instance, a process engineer may be consulted about the manufacturability of a new design concept, in which case that engineer draws parallels to already manufactured products. If there are enough similarities between a new and an old design, analysis models may be reused in some cases (Nergrd and Larsson, 2009), thus reducing computational cost and time. For instance, Runnemalm et al. (2009) automated the process of setting up weld simulations for structural aero-engine components, which was made possible by the components' proximity in the design space. As suggested by Woodbury and Burrow (2006), exploring the design space involves moving from known points to unknown ones; the closer an unknown point is to a known point, the higher the likelihood that the analysis of the former will yield results that are similar to the ones of the latter.

Furthermore, contemporary trends of digitalization have led some to suggest that data from existing products can be reused when evaluating new designs (Cantamessa et al., 2020; Tao et al., 2019). However, data from existing products are not necessarily relevant in the context of a new design (Woodall, 2017). We argue that similarity metrics may be useful for assessing the relevancy of data, analysis models, and knowledge of existing products.

2.2. Similarity metrics in engineering design

Similarity metrics are useful in many disciplines. For example, in computer science and machine learning similarity metrics have been used in a multitude of problems, most notably for clustering and classifying data (Xu and Wunsch, 2005). Other uses of similarity metrics include spell-checking and plagiarism detection based on the *Levenshtein distance*, which measures the similarity between two strings of characters (Su et al., 2008). A more generalized similarity metric was proposed by Li et al. (2004), who claimed that *normalized compression distance* can be used to measure the similarity of two arbitrary digital objects.

In engineering design, as mentioned earlier, similarities of a new design concept to existing or legacy products can enable previous development and manufacturing knowledge to be leveraged (Li et al., 2008). Applying known design solutions in new contexts is referred to as "design reuse" (Sivaloganathan and Shahin, 1999), and can potentially reduce cost and lead-times significantly (Duffy and Ferns, 1998). As a result, methods that make use of design similarity to leverage design reuse has emerged. Such is the case in *Case-Based Reasoning* (Aamodt and Plaza, 1994) and *Design-by-Analogy* (McAdams and Wood, 2002); they both assist in the selection of solutions to design problems through use of previous solutions. In doing so, design engineers need not reinvent existing solutions.

Case-Based Reasoning is essentially a method used for solving new problems by reusing knowledge and information from similar scenarios that has already been resolved. Aamodt and Plaza (1994) describe Case-Based Reasoning as a cyclical process: similar scenarios are retrieved and used to revise a new solution, which in turn is retained for use in future problem solving. To identify such similar scenarios similarity metrics can be used. An example of this was proposed by Akmal et al. (2014), who employed a feature-based similarity metric to augment a Case-Based Reasoning approach.

Design-by-Analogy is a concept generation approach that utilizes similarity. McAdams and Wood (2002) utilized similarity metrics to calculate the functional similarity between a new concept and existing designs. The purpose of the method is to assist design engineers in identifying design solutions to particular functions when generating new concepts. An alternative approach was suggested by Verhaegen et al. (2011), who compared new concepts to existing designs found in patent databases. In this highly automated process, candidates for Design-by-Analogy were identified by calculating their similarity.

The methods covered above are to be used primarily for assisting in the generation of new design concepts. Deploying them may enable knowledge reuse and speed-up of development times. There are also examples of similarity metrics being applied in the later stages of product development. Lupinetti et al. (2019) applied similarity metrics to compare assembly CAD models to enable knowledge reuse from previous assemblies. Bickel et al. (2020) also compared CAD models, looking at the geometrical similarity to previously produced parts as a means to reduce production times. Li and Bernstein (2017) applied similarity metrics to the process of identifying promising suppliers through the comparison of manufacturing process similarity. This list is by no means exhaustive; it is mentioned to highlight that similarity metrics can also be used to improve decision-making for processes that are downstream from design. Lupinetti et al. (2019) and Bickel et al. (2020) suggest that information created during the design phase (product geometry and assembly of CAD models) can be compared to existing designs to draw conclusion regarding new designs that will affect manufacturability. Arguably, if such similarities can be identified already during design space exploration, albeit with a reduced certainty, then downstream issues can be prevented at the early stage, thus preventing late redesigns.

Evidently, similarity metrics have multiple uses within the context of engineering design. As laid out in the introduction, the method proposed in this paper extends the use of similarity metrics in engineering design when exploring new designs. One of the mechanisms of this method is to identify similarities to existing designs, using a suitable similarity metric, to highlight potentials for reuse. This differentiates the proposed method from Design-by-Analogy, which primarily utilizes functional similarity (Nandy et al., 2021) to identify design solutions that have been used previously to achieve similar functions.

The use of other types of similarities, besides functional, enables its use for product families where the functionality is the same but the geometry is different. This is often the case for scale-based product families (Simpson et al., 2001a), in which the products can possess the same functionality but with varying dimensions. Furthermore, to our knowledge, no previous attempts have been made to utilize design similarity to augment design space exploration using surrogate models.

2.3. Early design evaluation through surrogate models

Surrogate models, sometimes referred to as “metamodels”, are approximations of higher fidelity models. These can be used as surrogates for otherwise computationally expensive models, such as finite element simulations (Simpson et al., 2001b). When referring to “surrogate models” in this paper, we exclusively refer to data-based surrogate models: surrogate models that are based on data from higher-fidelity simulations, physical measurements, or experiments.

Utilizing surrogate models has proven useful in engineering design, where optimization studies need to be conducted to find high-performing design configurations (Papalambros and Wilde, 2017). Performing optimization directly on simulation models is often impractical due to computational expenses, and thus the objective function(s) can be replaced with data-based surrogate models (Koziel et al., 2011). For this purpose, multiple surrogate modeling approaches are available, such as polynomial response surfaces, Gaussian processes, or neural networks. However, what is gained in speed can be lost in accuracy.

The choice of surrogate model, and how to configure it, depends on what needs to be predicted (Jin et al., 2001). At the same time these factors, together with the sample size of the training dataset, also impact the accuracy of the surrogate model. Typically, data from models of a higher fidelity, such as simulations, or from physical experiments, are used to “train” the surrogate model. However, such processes are often too computationally expensive or time-consuming. In such scenarios, only a few high-fidelity samples can be afforded to train the surrogate model, which compromises its accuracy.

How to select a surrogate model and optimize it for maximum accuracy is not within the scope of this paper. However, for the interested reader, it should be noted that there are multiple existing means for reducing computational cost while maintaining acceptable surrogate model accuracy. Viana et al. (2021) conducted a review of such techniques, categorizing them into four groups: (1) dimensionality reduction techniques, (2) data sampling techniques, (3) techniques for simultaneous use of multiple surrogate models, and (4) sequential sampling techniques.

The first group of techniques focuses on the reduction of dimensionality, which naturally has a direct impact on mitigating the curse of dimensionality. A common way of achieving this is through variable screening, which involves evaluating the impact each design variable has on the output, and omitting those whose impact is negligible (Koch et al., 1999). Moreover, several algorithms have been developed for the specific purpose of identifying opportunities for reducing dimensionality, that can be applied within engineering design. A recent example of this is Bird et al. (2021), who demonstrated a technique for increasing the accuracy of surrogate models used for predicting the properties of jet engine fan-blades by utilizing dimensionality reduction algorithms. In addition to minimizing the dimensionality of the problem, it is also paramount that the training dataset covers the design space appropriately. The second group of techniques focuses on experiment designs, of which there is a vast selection. Simpson et al. (2001b) and Yondo et al. (2018) both provide ample insight into this topic. The third group of techniques involves the use of multiple surrogate models simultaneously, sometimes referred to as an ensemble (Goel et al., 2007). The fourth and final group of techniques handles how the sampling is sequenced. There are many ways in which the sequence of sampling itself can be varied to improve accuracy. A technique that falls under this category is multi-fidelity surrogate modeling, which essentially involves mixing in a few samples of high-fidelity data with a larger set of low-fidelity data when training the surrogate model (Fernández-Godino et al., 2019). When adopting this approach, the high-fidelity data typically has a greater influence relative to the low-fidelity data when creating the surrogate model. The large number of low-fidelity data points combined with the few prioritized high-fidelity data points can thus yield increased accuracy.

The focus of this paper is on how to utilize similarity metrics to evaluate the trustworthiness of surrogate model predictions, and to help design engineers navigate the design space in search for interesting design regions. To the authors' knowledge this has not been previously attempted.

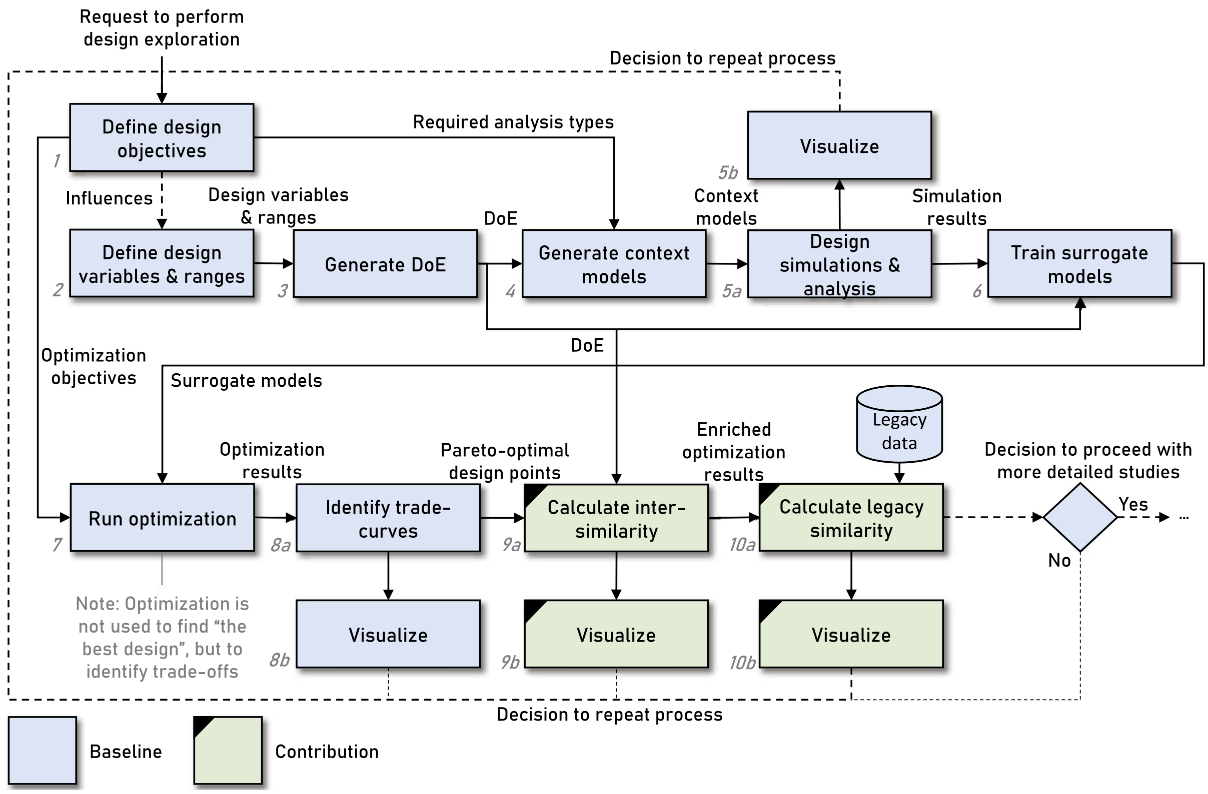

Fig. 1. The proposed method. The figure depicts a common design space exploration process derived trough consultancy with experts from a Swedish aero-engine components manufacturer. Appended to the process are two additional steps: "calculate inter-similarity" and "calculate legacy similarity", that serve as the primary contribution of this research. Additionally, two of the "visualize" activities are also highlighted as a contribution in the figure, as the method entails a degree of enrichment of traditional visualization through the use of a prototype software tool.

3. Proposed method

In this section the proposed method is presented in detail. To make the distinction easier a traditional design space exploration process is presented first, without the use of similarities. In Fig. 1, a typical design study process is visualized, along with the contributions added by the research presented in this paper. The baseline process includes all of the steps except for the calculation of inter-similarity and legacy similarity.

3.1. The baseline design space exploration process

The baseline process typically begins with a set of design objectives derived from customer requirements. These design objectives can for instance be, for a structural component, to have a low weight and a high stiffness. The engineers then initiate the design space exploration process by defining which design variables need to be explored, and within what ranges. Which variables are considered in the design space exploration process depends on what the objective of the exploration is. Based on these definitions and ranges a Design of Experiments (DoE) is initiated, spanning the design space region of interest. Each point in this DoE is thus a design point to be evaluated. These design points are then used to generate context models, such as CAD models and finite element meshes. These models are used for various kinds of design analysis. What kinds of analysis are performed depends on the design objectives. If there is a need to create a lightweight product, then the volume and mass need to be calculated. If there is a need for the product to withstand certain mechanical loads, then a load case simulation can be performed.

With the analysis completed, the results are aggregated and coupled with their corresponding design points that were used as input. At this point the design engineers may choose to inspect the results and potentially go back and make adjustments to the initial conditions of the design study. Otherwise, the simulation results are coupled with the simulation input (the design variable values), and used to train surrogate models. This enables the prediction of design analysis results without the need for running computationally expensive simulations.

The surrogate models enables optimization to be conducted to identify trade-offs and high performing regions in the design space. The objective functions are evaluated to a degree that would be too computationally expensive for simulations. Instead, the surrogate models are used as objective functions, which produces a set of low-fidelity data. By visualizing this data, engineers can analyze the results and make the decision to either proceed with a set of design points that look promising, or to go back and redefine the initial design space region of interest, thus continuing the exploration process.

3.2. Proposed additions to design space exploration

In the proposed method, three changes are suggested to the previously described design space exploration process: (1) A stage where the similarities between design points evaluated using surrogate models, and their closest simulated neighbors are calculated. This is referred to as *inter-similarity*. (2) A stage where the similarities between both the surrogate model and simulation design points, and legacy design points are calculated. This is referred to as *legacy similarity*. (3) Enhanced visualization steps in which information of these similarities is superimposed on the traditional means of visualization, thus providing the design engineers an extra layer of decision support.

Inter-similarity

Through the baseline design space exploration process depicted in Fig. 1 two types of results are obtained. The first type of results are obtained by means of high-fidelity simulations (step 5a). The second type of results are the predictions of the surrogate models during the

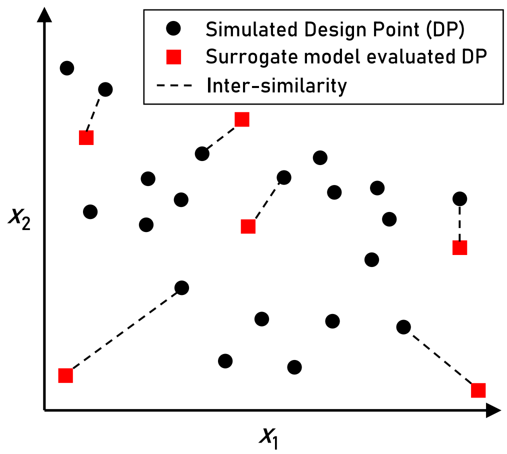

Fig. 2. Visualization of how inter-similarity is measured in a 2-dimensional design space. Inter-similarity only considers the distance to the closest simulated neighbor in the design space. A shorter dashed line represents a higher similarity.

optimization process (step 7 and 8a). Thus there are design points that have been evaluated using low-fidelity surrogate models and design points that have been evaluated through higher-fidelity simulations. Inter-similarity is defined as the distance in design space between a design point evaluated using a low-fidelity analysis (such as a surrogate model trained using higher-fidelity simulation data) to its closest neighbor evaluated at a higher fidelity (such as a physics-based simulation). This principle has been visualized in Fig. 2, and can also be formulated mathematically: Assume there exist n design points evaluated using high-fidelity analysis, and a function $S(\mathbf{x}_a, \mathbf{x}_b)$ that evaluates the similarity between two multidimensional design points, producing a lower value for a closer resemblance. The inter-similarity s_i of a design point evaluated using low-fidelity analysis $\mathbf{x}_{lf}$ can be calculated using the expression described in Eq. (1), where $\mathbf{x}_i$ represents a design point evaluated through high-fidelity analysis.

$$s_i = \min\left\{ S(\mathbf{x}_{lf}, \mathbf{x}_1), S(\mathbf{x}_{lf}, \mathbf{x}_2), \ldots, S(\mathbf{x}_{lf}, \mathbf{x}_n) \right\} \tag{1}$$

As reflected in Eq. (1), the similarity between all design points evaluated using surrogate models are compared to all design points evaluated through simulations. The closest identified similarity value is used as the inter-similarity score. Depending on what type of metric is used to calculate the inter-similarity, a higher score can either mean a high or low inter-similarity. Thus, how to interpret the inter-similarity score depends on the metric. To be consistent, a *high inter-similarity* will always mean that the design points are *close in the design space*, regardless of the similarity metric. Measuring the inter-similarity of a design point evaluated through surrogate models can have three potential benefits:

- If the inter-similarity is high in a certain design space region, then that indicates a high coverage of high-fidelity analysis in that design space region. This means that the predictions are more likely to be trustworthy. Additional simulations in this design space region are thus unlikely to yield a significantly increased surrogate model accuracy.
- If the inter-similarity is low, then trust in these predictions should also be low for any designs within this design space region. Thus, a lower inter-similarity can prompt the engineer to reconsider the trustworthiness of such results, and to consider executing additional simulations within this region of the design space, if affordable.
- Finally, as the inter-similarity can assist in locating the closest neighbor evaluated at a high fidelity, it can also help in understanding what differentiates a higher-performing design point from one that performs slightly less. This information can improve a designers understanding of the design space, and the effects of individual design variables.

Table 1
Examples of similarity metrics that can be used for different types of data.

Data type	*Similarity metric*
Numeric	Euclidean distance Cosine similarity Jaccard similarity coefficient
Text (short)	Levenshtein distance
Text (long)	Normalized compression distance
Object	Normalized compression distance
Class	Jaccard index

Legacy similarity

Legacy similarity is the similarity between design points in the design study, and designs that have been evaluated in previous projects. Naturally, this measurement is of a lower resolution than inter-similarity, since cross-generational design points may not share as many measurable similarity criteria. However, as previously established, similarities to existing designs can be highly beneficial. Thus, similarities of any magnitude could potentially be interesting to the design engineers. Furthermore, in use-cases such as the development of products within scale-based product families (Simpson et al., 2001a) multiple key design variables can be common to all products within the family, as products within the family are scaled up or down. Those key design variables can be used to measure the similarity between two products within a scale-based product family.

When calculating legacy similarity, both design points from the simulated dataset and the surrogate dataset can be evaluated. It is then calculated in a similar fashion as how inter-similarity is calculated: The similarity of each design point that has been simulated or evaluated through surrogate model, and legacy designs is calculated. The closest similarity value is used as the legacy similarity score. Once again, how the score is interpreted depends on the choice of similarity metric. To be consistent, a *high legacy similarity* will always mean that the two compared design points are *close within the design space*.

3.3. Similarity calculations

The similarity of two design points can be calculated in many different ways, depending on what type of data characterizes a design point and personal preference. Examples of similarity metrics that can be employed are listed in Table 1.

Similarity metrics such as Cosine similarity, Euclidean distance or the Jaccard similarity coefficient are well suited for numeric data, while Jaccard index can be employed if the data contains classes. For short texts, then the Levenshtein distance can be appropriate. If the data is object based or consists of long texts, then normalized compression distance can be used. In reality, these are only a few choices in an ocean of alternatives. The choice of similarity metrics should be made depending on the problem that needs to be solved, based on how well the similarity metric performs in that specific context. For numeric data, the three alternatives listed above are often chosen for their simplicity.

Neither of the similarity metrics for numeric data listed above accounts for that different features may have different value ranges. To avoid skewing the results of the similarity calculations the data thus first needs to be normalized. It should also be noted that different metrics has different response characteristics. For instance, when using Euclidean distance to measure similarity, then a score of 0 means that the two compared design points are equal (the similarity is high when the similarity score is low). On the other hand, if Jaccard similarity is used then a similarity score of 1 would mean that the two design points are equal (the similarity is high when the similarity score is high).

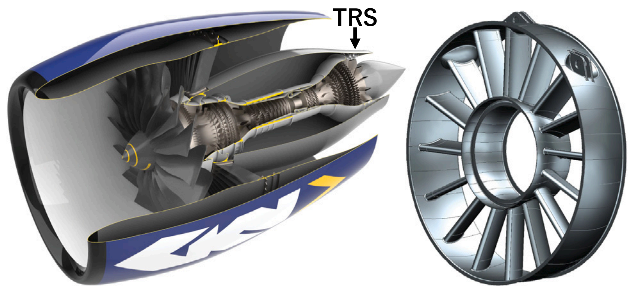

Fig. 3. The left image, provided by GKN Aerospace Engine Systems, displays the location of a Turbine Rear Structure (TRS) in a turbofan jet engine. The right image is a 3D model of a TRS.

4. Applied similarity assisted design space exploration

In this section the proposed method presented in Section 3 is demonstrated on a scenario from the Aerospace industry. In Section 4.1 the background to the scenario is presented. Before applying the proposed method to a design space exploration scenario an experiment was conducted to demonstrate that there is a correlation between inter-similarity and surrogate model prediction error. How this experiment was conducted, and what results it produced are presented in Section 4.2. To assist in visualizing and presenting the similarity data a software tool was developed. This tool is briefly described in Section 4.3. In Section 4.4 the proposed method is exemplified in a study of an aero-engine component. Finally, in Section 4.5 the results are discussed.

4.1. Scenario background

The Turbine Rear Structure (TRS) is a static component located behind the turbines of turbo-fan engines, as depicted in Fig. 3. It structurally supports the rear shaft bearing housing, while also providing mounting points such that the engine can be attached to the wing of an aircraft. Additionally, the TRS assists in directing the exhaust flow using guide vanes.

The TRS can be considered to be a product within a scale-based product family as each TRS has geometric similarities, are built on a common platform of knowledge, and are initially scaled up or down to suit new customer requirements. This enables comparisons between products within the product family through a set of common key design variables. When designing a TRS there are multiple key design variables that needs to be considered, including:

- The **number of vanes**, which impacts the stiffness of the structure, its weight, and the airflow.
- The **vane lean angle**, which can be adjusted to mitigate the effects of thermal expansion.
- The **inner and outer diameters** of the TRS, adjusted based on the size of the engine.
- The **thicknesses of each surface**, also impacting the stiffness of the structure, and its weight.

In Fig. 4 some of the design variables and key components of the TRS are visualized. Typically, the inner and outer flange diameters are set as constraints, as these are interfaces to adjacent parts in the engine. The inner and outer diameters are thus locked, as they are determined on an engine system level by the original equipment manufacturer. However, the number of vanes, their lean angle, and the thicknesses of a large quantity of part regions are often varied when exploring the design space. This results in a large design study dimensionality. Meanwhile, generating the geometries and simulating the conditions necessary to evaluate the design is computationally expensive. Consequently, design engineers resolve to building surrogate models based

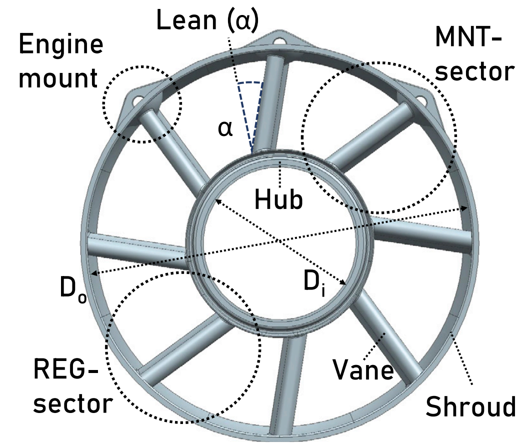

Fig. 4. Map of some of the key components and design variables of the TRS.

on a relatively small amount of simulated design points to aid in the decision-making process.

In the example used in this chapter, 8 design variables are used to determine the configuration of the TRS design: the number of vanes, the vane lean angle, and 6 thickness regions. It should be noted that these design variables control more than a single dimension on the TRS. For instance, changing the vane count will directly impact the length of the TRS (and thus also the length of the engine). Another important interaction with the vane count is the vane width: the higher the number of vanes, the lower the vane width. This is to ensure that no matter the number of vanes the airflow through the TRS will remain the same.

4.2. Experiment: Correlation study of similarity and surrogate model error

To investigate the potential of using inter-similarity as an indicator of trust in prediction results, an experiment was conducted to evaluate the correlation between inter-similarity and surrogate model prediction error. Three factors likely to influence this correlation were considered, based on what has previously been mentioned in this paper: (1) the choice of surrogate model; (2) the choice of similarity metric; (3) the sample size of the data set used to train the surrogate model. The impact of sample size on surrogate model accuracy is a known phenomenon, and will not be discussed in this paper. Additionally, it is important to remember that the choice of surrogate model is heavily dependent on what needs to be predicted. How to select an appropriate type of surrogate model has been covered extensively in literature (Simpson et al., 2001b; Yondo et al., 2018), and will not be further investigated in this experiment.

Three types of surrogate models were investigated: A second- and a third-degree polynomial response surface, and a Gaussian Process with an Radial Basis Function (RBF) kernel. For similarity metrics the normalized Euclidean distance and the Jaccard similarity coefficient were tested. The Euclidean distance was normalized by first scaling all input variables such that they spanned a range from 0 to 1, and then dividing the Euclidean distance by the square root of the dimensionality, as described by Eq. (2). The Jaccard similarity coefficient, as defined by Eq. (3), also utilized a normalized input. In these expressions, $\mathbf{x}_\mathrm{a}$ and $\mathbf{x}_\mathrm{b}$ are two normalized design points between which the similarity is measured, and k is the dimensionality of these design points.

$$\frac{\sqrt{\sum_{i=1}^{k}(\mathbf{x}_{\mathrm{a},i}-\mathbf{x}_{\mathrm{b},i})^2}}{\sqrt{k}} \tag{2}$$

$$\frac{\sum_{i=1}^{k}\min\{\mathbf{x}_{\mathrm{a},i},\mathbf{x}_{\mathrm{b},i}\}}{\sum_{i=1}^{k}\max\{\mathbf{x}_{\mathrm{a},i},\mathbf{x}_{\mathrm{b},i}\}} \tag{3}$$

Table 2
The TRS design variables and the boundaries used in the DoE for the correlation study, and for the design space exploration example. Note that the inner and outer diameters do not have a range as they are not varied in this DoE.

Design variable	*Range*	*Unit*	*Symbol*
Diameter outer	1100	mm	D_o
Diameter inner	420	mm	D_i
Vane count	[8, 18]	N/A	N
Vane lean angle	[0, 20]	°	α
Hub thickness (reg)	[1.5, 4]	mm	t_{hr}
Hub thickness (mnt)	[1.5, 4]	mm	t_{hm}
Shroud thickness (reg)	[1.5, 4]	mm	t_{sr}
Shroud thickness (mnt)	[1.5, 4]	mm	t_{sm}
Vane thickness (reg)	[1.5, 4]	mm	t_{vr}
Vane thickness (mnt)	[1.5, 4]	mm	t_{vm}

To evaluate the correlation between similarity and prediction error, the mean absolute error was plotted against ranges of inter-similarity. A correlation between error and inter-similarity should result in a slope, where data with a high inter-similarity generally results in a reduced mean error, and vice versa. The error is defined as the absolute distance between the predicted value, and the value returned by a simulation. In this test, the simulated output was the maximum deformation of a TRS under the influence of a static load, evaluated using Ansys workbench. This particular output was selected as it is computationally expensive to calculate using simulations.

First, training and test data were generated. For the training data, a set of 100 design points were evaluated. These design points were sampled using a hypercube DoE. Solid CAD geometries were then generated from these design points, using a method described in a separate paper (Martinsson Bonde et al., 2022). After generating the CAD models, each geometry was also meshed. This process, along with the geometry generation, is somewhat sensitive and expected to fail for a small fraction of the designs. A common reason for this is conflicting CAD-model constraints, which can result in invalid geometries. For the training dataset one design failed to mesh, resulting in a final sample size of $n_{train} = 99$ that were evaluated through simulation. The testing dataset had a total initial sample size of 750, also sampled using a hypercube DoE. Out of those 750 samples 19 failed to mesh, resulting in a final test sample size of $n_{test} = 731$. The dimensionality of the data was $k = 8$, and the specific boundaries used for the DoEs are described in Table 2.

To enable calculating the prediction error of the surrogate models, the design points in the test dataset were also evaluated using simulations. Thus, the difference between the simulated and the predicted value was used to determine the prediction error. The experiment process is presented in Fig. 5.

The results of the analysis visualized in Fig. 6 depicts how the choice of surrogate model affected the results. For these results, Euclidean distance was used to measure the similarity, and thus a lower inter-similarity score indicates a higher inter-similarity. It can be observed that all three surrogate models follows the expected trend of a high inter-similarity (low inter-similarity score) resulting in a lower surrogate model prediction error. In this particular instance, the second-degree polynomial response surface has the best overall performance in terms of prediction error. Yet it still displays a prominent slope, as with the other two surrogate models, where an lower inter-similarity generally results in a larger prediction error.

A second evaluation using the same data, but a different similarity metric (the Jaccard similarity coefficient) was conducted. As previously mentioned, unlike Euclidean distance, the Jaccard similarity score is high when the similarity is high. The results, as visualized in Fig. 7, show that the choice of similarity metric did not result in a significant difference in the trend. While the result is marginally different, the conclusion is the same: a high inter-similarity results in a reduced mean prediction error.

From these results it can be concluded that inter-similarity often is an indicator of surrogate model trustworthiness. At this point, one important question that might arise is *"how similar is similar enough?"*. It should first be clarified that it is not recommended for engineers to screen results based on inter-similarity as it is merely an indication of trustworthiness and of how exhaustive the high-fidelity coverage of the design space is. With that in mind, in a design study context, the exact threshold for "similar enough" will not be known. The correlation between surrogate model prediction error and inter-similarity evaluated in this section was only studied for a specific output with a particular dimensionality and sample size. In an industrial scenario, evaluating the correlation between inter-similarity and prediction error to find an appropriate threshold would likely not be feasible, as doing so would be highly computationally expensive. However, we argue that design engineers will have use for the knowledge of which design points are likely to be the most trustworthy (those with the highest inter-similarity). Having a clear indication that a design space region is not appropriately represented in the training dataset, or that a set of designs are in close proximity to a high-fidelity result can help designers assess when to deploy additional simulations, and when not to. In other words, it is not possible to say what degree of inter-similarity is enough for a prediction to be trustworthy, but having access to the similarity information can help design engineers understand their surrogate models, and their design spaces. In Section 4.3 a basic tool will be described that assists in visualizing this additional layer of information for design engineers, and in Section 4.4 the proposed method will be utilized in a design study, demonstrating how the reasoning described above can be applied in an industrial context.

4.3. Visualization tool

To make the similarity information useful, it is necessary to somehow visualize it. To that end, a prototype tool was developed that demonstrates how similarity information can be presented to a design engineer to assist in design space exploration activities. The tool is written in JavaScript (Vue.js), and can be used through a browser. A user can import a Comma-Separated Values (CSV) file into the tool to visualize the contained data using scatter plots and parallel coordinates diagrams. The intended use case is to enable the visualization of data within a CSV-file that contains: (1) the design points from a DoE, and (2) simulation and surrogate model analysis results for those design points. By pre-processing the imported data, the inter-similarity of each design point can be added as an additional column within the CSV-file. The tool can then be used to color code the plots based on inter-similarity, thus giving an indication of inter-similarity of each design point.

Scatter plots, parallel coordinate diagrams, and color coding is not new. The functionality this tool adds is the possibility to measure the distance in the design space between two design points. When looking at a scatter plot, the user can select one design point to immediately color code all other design points based on the design space distance between them and the selected design point. This does not require inter-similarity to be pre-calculated, as the calculation is done within the tool itself. Furthermore, selecting an additional point will grant the user information regarding the distance between the two points, and the differences between them (both input and output are listed). A screenshot from the design study in Section 4.4 when this functionality was utilized is available in Appendix.

Finally, the user can swap between a scatter plot view, and a parallel coordinates view. Samples can be selected in either of the view, which are then automatically highlighted in the other view. This can be useful when the user wishes to, for example, understand why some result has a low inter-similarity. By selecting the design points that reportedly have a low inter-similarity, the user can then switch to the parallel coordinates diagram to get a visualization of the design space region that lacks representation of higher-fidelity data.

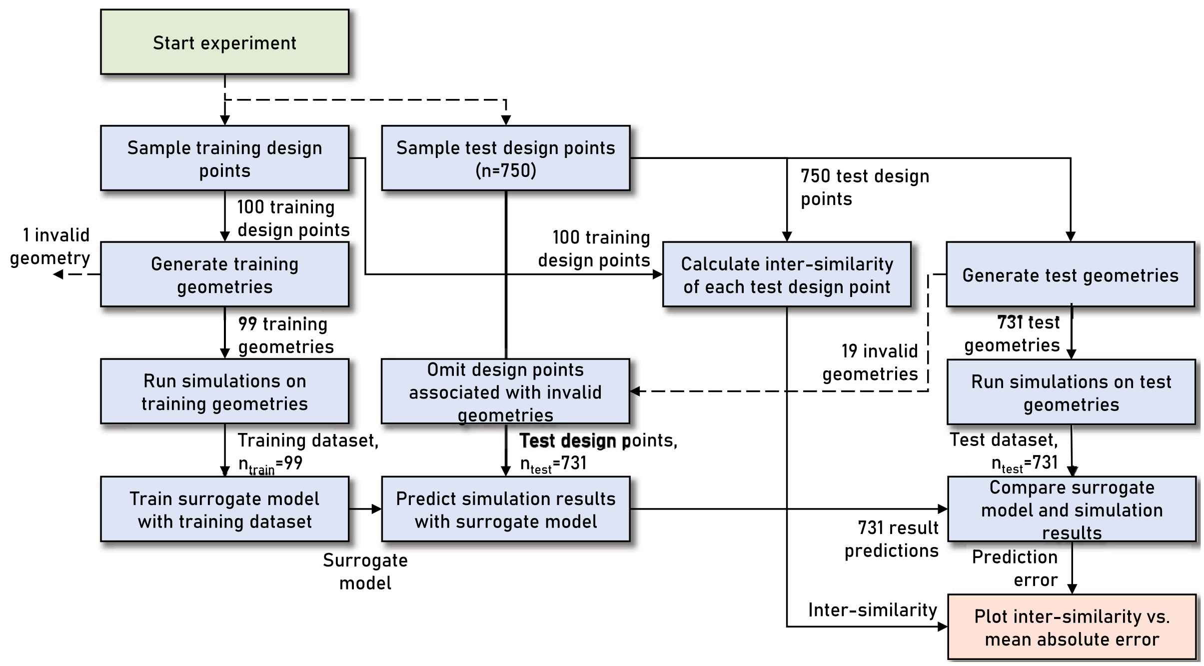

Fig. 5. The correlation study process. This schematic shows the flow of design points and their corresponding geometrical representations throughout the experiment. This sequence was run six times, with two different similarity metrics (Euclidean distance and Jaccard similarity) and three different surrogate models (2nd-degree and 3rd-degree response surfaces, and a Gaussian process with an RBF kernel).

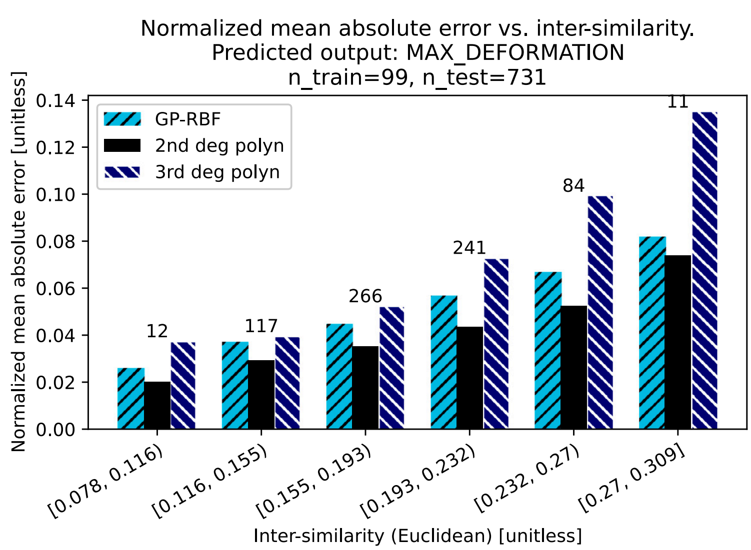

Fig. 6. Depicts the correlation of surrogate model error and inter-similarity using Euclidean distance as a similarity metric. A lower inter-similarity value indicates a closer inter-similarity. The number above the bars indicate the number of samples in that similarity range.

4.4. Application in a design study context

In this section the process flow visualized in Fig. 1 is followed for the design space exploration of a TRS component. In this exploration, the goal is to identify how the design can be adjusted through its design variables to account for stiffness and weight requirements. It is necessary for a TRS to be able to absorb large mechanical loads in cases of engine failure (such as fan-blade-out). At the same time, the design needs to be light to reduce fuel consumption. Thus, two design objectives were considered: (1) The component needs to be as stiff as possible, and (2) The component needs to be as light as possible. To measure the stiffness, the maximum deformation (δ_{max}) of the structure under a specific load was used as a proxy. A lower deformation thus means a higher stiffness. Additionally, since the entire TRS is build using the same material, the volume (V) was measured rather than the weight.

With the design objectives in place, the next step was to define appropriate design variables and ranges. For a TRS it is expected that the number of vanes, how they are leaned, and different structural wall thicknesses will have a significant impact on max deformation and volume. The considered design variables listed in Table 2 were chosen accordingly.

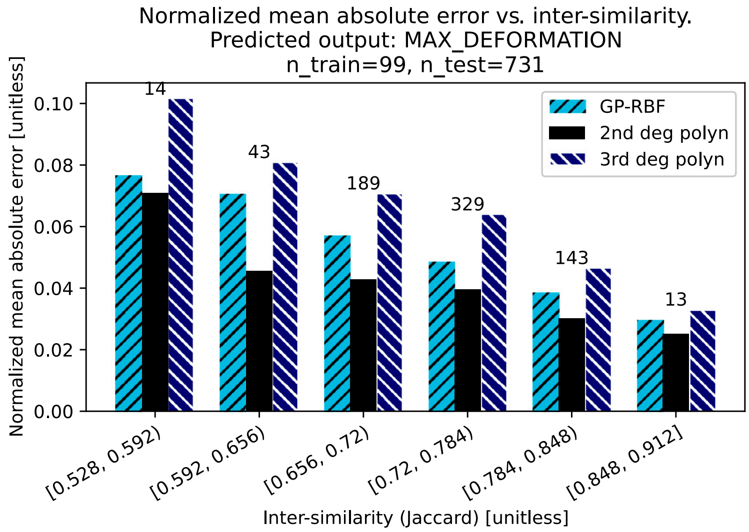

Fig. 7. Depicts the correlation of surrogate model error and inter-similarity using Jaccard similarity coefficient as a similarity metric. For this metric a high inter-similarity (maximum 1) indicates a closer inter-similarity.

A DoE with 100 samples in a hypercube configuration was instantiated, in which the selected design variables were varied within the defined ranges. These design points were then used to generate CAD-models and finite element meshes. These context models were then applied in two types of design analysis. The first analysis was a volume extraction using Siemens NX. The other analysis was a static structural load case simulation in Ansys. Both the generation of context models and the design analysis were computationally expensive. Thus, the design points from the DoE were coupled with the results from the design analysis to form a dataset that was used to train two surrogate models: one for predicting the volume, and one for predicting the stiffness. This enabled the evaluation of the stiffness and volume of a design point without running computationally expensive model generations and simulations. At this point a multi-objective optimization problem was defined (see Eq. (4)) to identify the trade-off curve between the two design objectives (low deformation and low volume) in the design space.

$$
\begin{aligned}
&\min_{\mathbf{x}} \quad \left[V(\mathbf{x}), \delta_{\max}(\mathbf{x})\right]^{\mathrm{T}} \\
&\text{where} \quad \mathbf{x} = \left[N, \alpha, t_{\mathrm{hr}}, t_{\mathrm{hm}}, t_{\mathrm{sr}}, t_{\mathrm{sm}}, t_{\mathrm{vr}}, t_{\mathrm{vm}}\right] \\
&\qquad\quad N \in \{8, 9, 10, \ldots, 18\} \\
&\qquad\quad 0^{\circ} \leq \alpha \leq 25^{\circ} \\
&\qquad\quad 1.5\ \mathrm{mm} \leq t_{\mathrm{hr}}, t_{\mathrm{hm}}, t_{\mathrm{sr}}, t_{\mathrm{sm}}, t_{\mathrm{vr}}, t_{\mathrm{vm}} \leq 4\ \mathrm{mm}
\end{aligned}
\tag{4}
$$

A genetic algorithm was utilized with a population of 100 samples. It should be noted here that the ranges for the variables used in Eq. (4) are the same as the ranges used in the initial DoE defined in Table 2. Thus, all optimization results exists within the same design space region as the simulated design points. Despite this, the results from the optimization study managed to reach lower levels of deformation and volume relative to any of the simulated samples. In Fig. 8 the results from the optimization has been plotted together with the analysis results of the simulated dataset.

Inter-similarity analysis

The optimization resulted in a clear trade-off curve. The next step was to get an indication of the trustworthiness of these new results that were generated using the surrogate models, through the use of inter-similarity. The data in Fig. 8 was thus color-coded based on inter-similarity. The inter-similarity results shows that the design points with the lowest volumes on the Pareto front are relatively far from any of the simulated design points in the design space, as they have inter-similarity values ranging from 0.25 to 0.35. This means that if a design with a very low volume and a relatively low stiffness would have been of interest, then further design points would have needed to be evaluated through simulations to increase trust in the relevant region(s) of the design space.

Using the interactive visualization tool described in Section 4.3, the low-fidelity design points with a low volume were selected in the scatter plot. The results were then exported to a parallel coordinates plot, which enabled a visualization of the design space region within which these design points reside. In Fig. 9, these design points are plotted on a parallel coordinates diagram. A set of filters has been applied that correspond to this design region, which is characterized by low thicknesses, few vanes and a high vane lean. By identifying the low-trust design region, it would be possible for the designers to generate additional geometries for a set of designs within this region, and analyze them through simulations. The results from those simulations could then be used to improve the performance of the surrogate models, thus increasing the trustworthiness of results from surrogate model predictions within this design space region.

Fig. 8 also shows a set of design points in the middle of the Pareto front that are in close proximity to simulated design points (an inter-similarity of around 0.15), meaning that these data are likely to be more trustworthy. Since the design objectives of this design space exploration activity were to identify stiff and light-weight structures, it was concluded that these data in the middle of the Pareto front were promising. Thus, the exploration process proceeded on to the next step: to calculate the similarity to legacy designs.

Legacy similarity analysis

When calculating legacy similarity, data available from previous design analysis and existing products are included. Since legacy design points have been tested thoroughly in the past with both simulations and in some cases physical tests, these data are of the highest fidelity available. In this design study example three hypothetical legacy designs were used, evaluated using simulations. Since these legacy designs are not based on the exact same design, they do not share all design variables. However, they are part of the same scale-based product family, and thus share some variables used for determining the size of the component, and some of its main characteristics. The common variables are: The outer diameter (D_{o}), the inner diameter (D_{i}), the

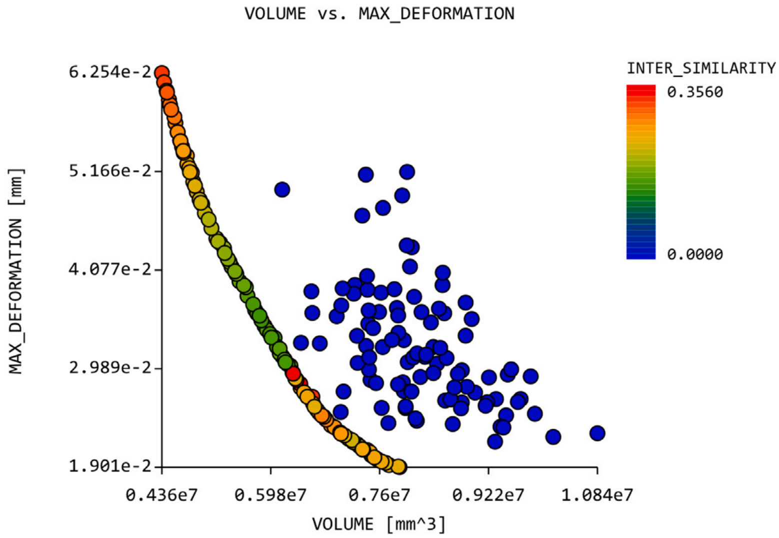

Fig. 8. Results from the optimization step color coded based on inter-similarity. Euclidean distance was used to measure inter-similarity. Thus, a lower inter-similarity indicates that the data has a close neighbor in the simulation dataset. The data points from the simulations turn up blue in the plot, as they by definition have an inter-similarity of 0 (they are similar to themselves).

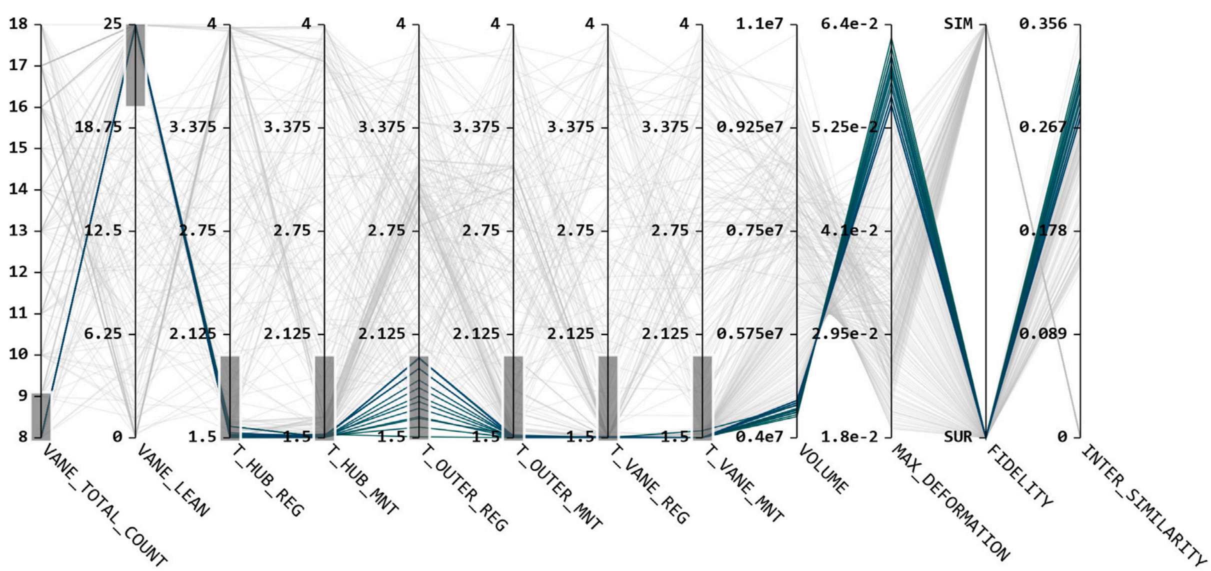

Fig. 9. Parallel coordinates diagram depicting a design space region within which the low-fidelity surrogate models report that the volume/weight of the design should be low. The boxes in the diagram represent filters, which have been used to filter out any design points not within that region, leaving only surrogate model results (as indicated by the FIDELITY-axis). To increase trust in surrogate model predictions within this region, more simulations needs to be run within this subspace.

Table 3
Legacy design points with their corresponding key design variable values.

Design name	D_o	D_i	N	α
Legacy design A	1500 mm	500 mm	18	10°
Legacy design B	800 mm	350 mm	8	20°
Legacy design C	900 mm	400 mm	12	25°

number of vanes (N), and the vane lean (α). These are the variables that were be used to determine the legacy similarity between the simulated and surrogate model evaluated design points and the legacy design points. The variables and their values for the legacy design points have been listed in Table 3.

Once again the tool was used to calculate these values and visualize the data. In Fig. 10 the three legacy data points have been plotted along with the data from the simulations and optimization. Notably, design point **B** and **C** are to the left of the Pareto front, as they were designed for smaller engines (they have a lower outer diameter). This enables them to reach lower volumes. Meanwhile, design point **A** is for a much larger engine, and thus has a relatively high volume. In the figure, one of the most trustworthy points (based on its inter-similarity) has been selected using the tool, which results in all other design points being color coded based on their similarity to the selected point. Using the color coding it was possible to identify that legacy design point **C** was the closest high fidelity design point relative to the selection. The tool reports a legacy similarity of 0.158 relative to the selection, which can be extracted from the tool by selecting both the legacy design point,

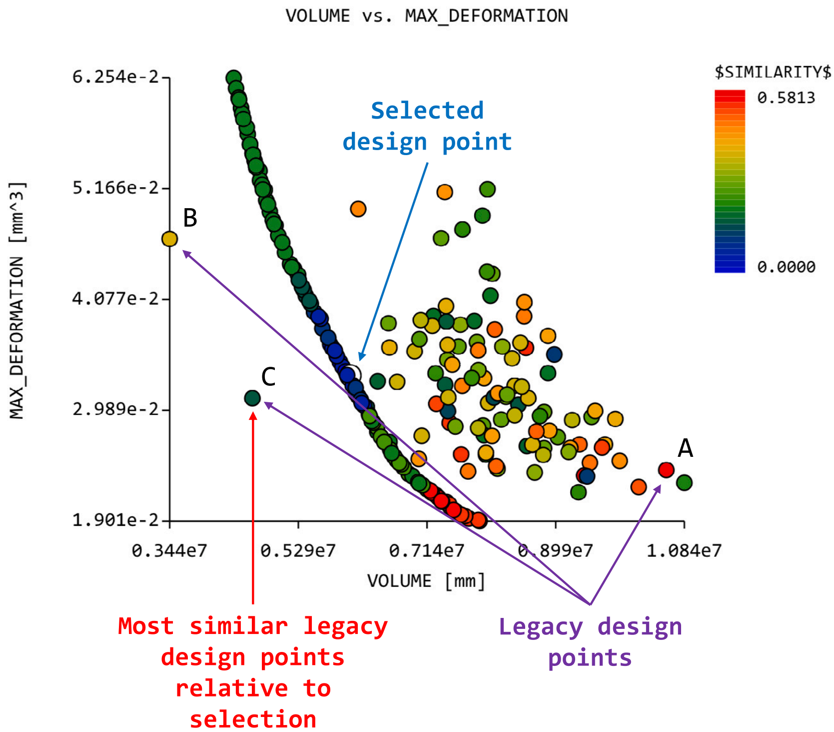

Fig. 10. Plot with data from all three levels of fidelity: surrogate model, simulated, and legacy. The legacy data has been marked in the plot. In this figure the color code is dependant on the similarity of all design points relative to the selected design point (also marked out in the figure).

and the point which you want to compare it with. This comparison also returns a lists of the differences between the two design points, which in this case reported that the legacy design point has the same number of vanes and vane lean, but has slightly smaller inner and outer diameters relative to the selected design point. Thus, the overall dimension of the legacy design point was slightly smaller, but it had the same vane configuration.

After establishing that legacy design point **C** was relatively similar to the selected design point on the Pareto front, two important conclusions could be made. Firstly, since the selected design point and the legacy design point **C** also have similar outputs (volume and max deformation) that information helps verify the trustworthiness of the surrogate model results. Secondly, the similarity between the two design points may be close enough to enable some degree of knowledge, model, or data reuse. Using the information presented in the tool (the differences between the two design points) a team of cross-functional engineers could decide on whether or not knowledge, analysis models and/or data could be reused. The potential of reuse can then influence whether more detailed studies should be conducted on the design points of interest.

4.5. Discussion

As stated in the introduction, the aim of the proposed method is to reduce the time allocated to high-fidelity simulations, and to improve the efficiency of obtaining promising design configurations. Thus, the first question that needs to be answered is: *does the proposed method achieve this goal?* In the design study, a situation was exemplified where the surrogate model had produced trustworthy or non-trustworthy results depending on the desired outcome of the design space exploration activity. If the designers had been looking for low-volume designs where high deformation was acceptable, then the inter-similarity metric clearly indicated that the data produced by the surrogate models was not to be trusted. In other words: it was too far from anything that had been previously simulated. It is possible that this could have been avoided by tweaking or choosing another design of experiments. However, there is no way of knowing beforehand where in the design space the Pareto-optimal design points will be situated. An experienced engineer may be able to predict which design variables are the most important for individual objectives, though such predictions are difficult in trade-off scenarios, and less useful in multi-objective design optimization scenarios where there are more than two objectives that needs to be fulfilled. In the example scenario, the inter-similarity measurement served as an indication to the designers not to trust these data. A designer who encounters this kind of result needs to run more simulations, or understand that the surrogate model results may be far from their true values. On the other hand, the more balanced alternatives in the middle of the Pareto front were found to be relatively close to previously simulated design points. In this scenario, running more simulations in this area would likely not result in improved or more accurate results. Thus, the inter-similarity has saved the designer precious computational resources by indicating that the results are likely trustworthy.

A second question that needs to be addressed is: *in what scenarios is the proposed method useful?* Inter-similarity provides design engineers with an additional layer of information which can assist in understanding the trustworthiness of surrogate model predictions, and the coverage of high-fidelity analysis within the design space. This additional layer is created using already existing information (the design points). Since no new information is necessary to create this layer of information, it can be included in any study that utilizes surrogate models. In Section 4 aero-engine component design is used as to demonstrate the proposed method. However, the use of inter-similarity is generally applicable to broader model-based design analysis and optimization applications. Arguably, inter-similarity can be useful in any design space exploration scenario where compromises needs to be made on surrogate model accuracy. Calculating inter-similarity is inexpensive from a computational perspective, and can thus be done without penalty. It is also easy to automate this process. Legacy similarity, on the other hand, has a slightly narrower use case. It is well-suited for scenarios where previous designs and existing products have been similar enough to the new designs to enable measurable comparisons. As previously mentioned, a good example is when developing new designs within a scale-based product family, where new products are scaled up/down versions of existing solutions. Having the possibility to measure the similarity between new designs and existing solutions can be useful in three ways: (1) to verify simulation results based on test or experiment data from higher fidelity analysis run on an existing solution; (2) to identify the potential to reuse knowledge, analysis models, and data from previous designs and products; (3) to avoid or reconsider designs that are similar to existing products that encountered problems during their life cycles (e.g., manufacturability or reliability problems).

One final question: *does utilizing inter-similarity and legacy similarity make the job of the design engineer easier?* If the effort to calculate inter-similarity and legacy similarity is too great, then it may outweigh the potential benefits of utilizing those metrics. In the case of inter-similarity, it can easily be calculated automatically either using a script run after the optimization stage (as was done in the design study presented in this paper), or in real-time while using the tool if the dataset is small enough. By utilizing vector operations, calculating inter-similarity can be done in less than a second for datasets with a sample size of 1000. Increasing the sample size will naturally also increase the calculation time, but it will likely remain insignificant if implemented in an efficient way. Legacy similarity, however, provides a greater challenge as it requires the engineers to have access to downstream data from existing products in an appropriate format. This has proven difficult historically (Andersson and Isaksson, 2008), and would require an initial effort to configure a database that can store the necessary information, or some other mechanism that is able to aggregate the necessary information from existing databases. However, granted that the design engineers have access to downstream data in an appropriate format, and that legacy designs are comparable to new designs, then calculating and utilizing legacy data would be both attainable and useful.

5. Conclusion

A method was proposed to assist design engineers in evaluating the trustworthiness of surrogate model results during design space exploration. It considers similarity to design instances assessed with a high confidence as a means to improve trust in low-fidelity predictions. Such similarity was referred to as "inter-similarity": the closest distance in the design space from a surrogate model evaluated point, to a point analyzed with simulations. The intended benefit of utilizing this measurement is to reduce the need for running computationally expensive simulations, and to assist design engineers in navigating complex design spaces.

An experiment was conducted where the correlation between inter-similarity and surrogate model prediction error was evaluated. The results of this experiment suggested that inter-similarity is a useful indicator of trust for surrogate models with high dimensionality that have been trained with a low sample size. Scenarios in which such "low-quality" surrogate models are used are common in industry. This is because simulations are typically used to evaluate designs, but are often too computationally expensive to be used for design space exploration. Surrogate models are then used as a means to reduce the need for simulations, thus reducing computational expenses for the price of diminished accuracy.

The method also utilizes similarity to existing products and designs analyzed with a high fidelity. This is referred to as "legacy similarity": the distance in the design space from a new design point, to a point that has been evaluated with high fidelity in previous endeavors. The purpose of this metric is to provide a means to indicate relevancy of existing knowledge, analysis models, and data. At the same time, it can also be used to verify analysis results by indicating that a similar existing product/legacy design has been evaluated and produced similar results. This, together with inter-similarity, was exemplified in a design study conducted on a static aero-engine component. In this design study it was demonstrated how both of these metrics (inter-similarity and legacy similarity) can be used to improve the design space exploration process. To facilitate the use of these metrics in the design study and in future industrial applications a prototype software tool was developed to visualize the similarity information.

The relevance and utility of an inter-similarity metric in practice was demonstrated through an experiment, and a design space exploration study. It was shown that this metric can be used together with other metrics in early design studies to assist design engineers. Furthermore, inter-similarity is relatively easy to calculate and deploy, meaning that it can be integrated into new development projects without significant effort. Legacy similarity, however, has more challenging prerequisites: downstream data such as data from physical tests, or from high-fidelity simulations, need to be accessible in a useful format. In practice, this is not as straightforward as it might appear. Such data are created, stored, and used for other purposes, and it would require a larger effort to set up such conditions within many companies. If the prerequisites can be resolved then legacy similarity can be utilized in a similar fashion as inter-similarity in situations where previous generations of products are comparable to new designs, such as in product families. It should however be noted that inter-similarity and legacy similarity are independent from each other, and does not require the other to function as intended. However, when utilized together, as in the proposed method, both computational expenses and development lead times can potentially be reduced.

CRediT authorship contribution statement

Julian Martinsson Bonde: Writing – original draft, Writing – review & editing, Software, Methodology, Conceptualization. **Michael Kokkolaras:** Supervision, Writing – original draft, Writing – review & editing, Conceptualization. **Petter Andersson:** Supervision, Project administration, Conceptualization. **Massimo Panarotto:** Supervision, Conceptualization. **Ola Isaksson:** Supervision, Project administration, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

This research was supported by the Swedish Governmental Agency for Innovation Systems (VINNOVA) through the DIFAM project, grant number [2019-02756]. Additionally, the research was also supported through funding from the Clean Sky 2 Joint Undertaking under the European Union's Horizon 2020 research and innovation programme via the DIAS project, grant number [887174].

Appendix. Screenshot from software visualization tool

See the figure.

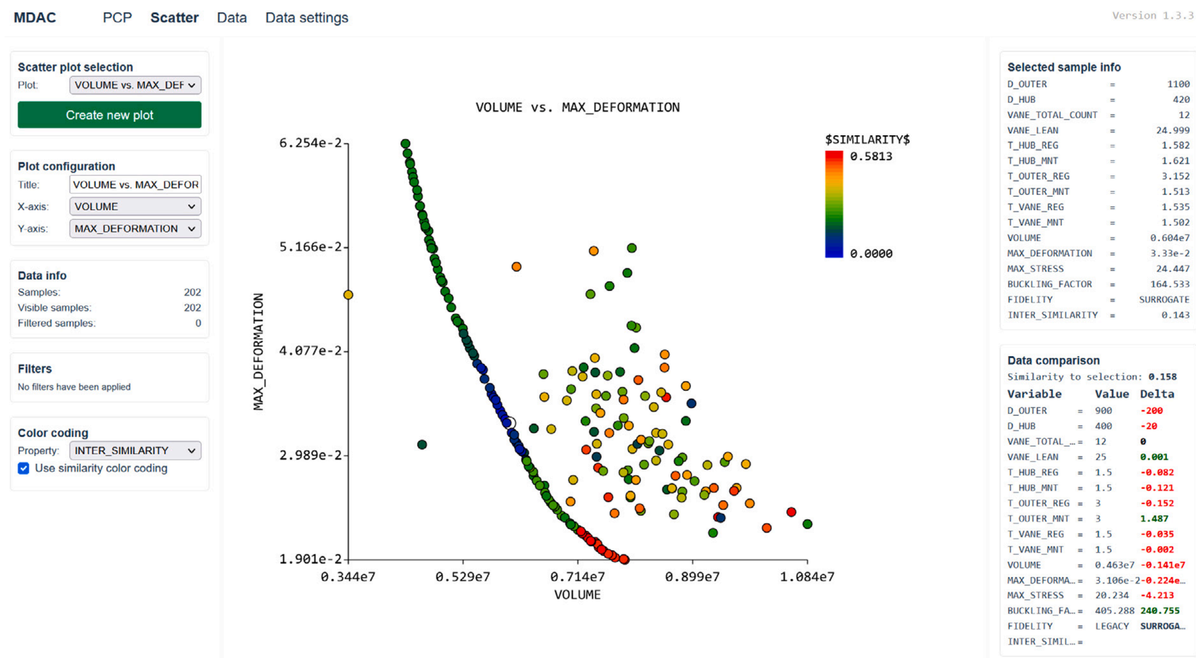

References

Aamodt, A., Plaza, E., 1994. Case-based reasoning: foundational issues, methodological variations, and system approaches. AI Commun. 7, 39–59. http://dx.doi.org/10.3233/AIC-1994-7104.

Akmal, S., Shih, L.H., Batres, R., 2014. Ontology-based similarity for product information retrieval. Comput. Ind. 65 (1), 91–107. http://dx.doi.org/10.1016/j.compind.2013.07.011.

Andersson, P., Isaksson, O., 2008. Manufacturing system to support design concept and reuse of manufacturing experience. In: Mitsuishi, M., Ueda, K., Kimura, F. (Eds.), Manufacturing Systems and Technologies for the New Frontier. Springer London, London, pp. 137–140. http://dx.doi.org/10.1007/978-1-84800-267-8_28.

Bickel, S., Sauer, C., Schleich, B., Wartzack, S., 2020. Comparing CAD part models for geometrical similarity: A concept using machine learning algorithms. Procedia CIRP 96, 133–138. http://dx.doi.org/10.1016/j.procir.2021.01.065.

Bird, G.D., Gorrell, S.E., Salmon, J.L., 2021. Dimensionality-reduction-based surrogate models for real-time design space exploration of a jet engine compressor blade. Aerosp. Sci. Technol. 118, 107077. http://dx.doi.org/10.1016/j.ast.2021.107077.

Cantamessa, M., Montagna, F., Altavilla, S., Casagrande-Seretti, A., 2020. Data-driven design: the new challenges of digitalization on product design and development. Des. Sci. 6, e27. http://dx.doi.org/10.1017/dsj.2020.25.

Duffy, A.H.B., Ferns, A.F., 1998. An analysis of design reuse benefits. In: Lindemann, U., Birkhoffer, H., Meerkamm, H., Vajna, S. (Eds.), Proceedings of the 12th International Conference on Engineering Design. ICED '99, Design Society, pp. 799–804.

Fernández-Godino, M.G., Park, C., Kim, N.H., Haftka, R.T., 2019. Review of multi-fidelity models. AIAA J. 57 (5), 2039–2054. http://dx.doi.org/10.2514/1.j057750.

Goel, T., Haftka, R.T., Shyy, W., Queipo, N.V., 2007. Ensemble of surrogates. Struct. Multidiscip. Optim. 33 (3), 199–216. http://dx.doi.org/10.1007/s00158-006-0051-9.

Jin, R., Chen, W., Simpson, T.W., 2001. Comparative studies of metamodelling techniques under multiple modelling criteria. Struct. Multidiscip. Optim. 23 (1), 1–13. http://dx.doi.org/10.1007/s00158-001-0160-4.

Keogh, E., Mueen, A., 2011. The curse of dimensionality. In: Encyclopedia of Machine Learning. Springer, pp. 257–258.

Koch, P.N., Simpson, T.W., Allen, J.K., Mistree, F., 1999. Statistical approximations for multidisciplinary design optimization: the problem of size. J. Aircr. 36 (1), 275–286. http://dx.doi.org/10.2514/2.2435.

Koziel, S., Ciaurri, D.E., Leifsson, L., 2011. Surrogate-Based Methods. In: Koziel, S., Yang, X.-S. (Eds.), Computational Optimization, Methods and Algorithms. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 33–59. http://dx.doi.org/10.1007/978-3-642-20859-1_3.

Li, K., Bernstein, W.Z., 2017. Developing a capability-based similarity metric for manufacturing processes. In: International Manufacturing Science and Engineering Conference. Vol. 3.

Li, M., Chen, X., Li, X., Ma, B., Vitanyi, P.M.B., 2004. The similarity metric. IEEE Trans. Inform. Theory 50 (12), 3250–3264. http://dx.doi.org/10.1109/TIT.2004.838101.

Li, Z., Raskin, V., Ramani, K., 2008. Developing engineering ontology for information retrieval. J. Comput. Inf. Sci. Eng. 8 (1), http://dx.doi.org/10.1115/1.2830851.

Lin, D., 1998. An information-theoretic definition of similarity. In: Icml, Vol. 98, No. 1998. pp. 296–304.

Lupinetti, K., Giannini, F., Monti, M., Pernot, J.-P., 2019. Content-based multi-criteria similarity assessment of CAD assembly models. Comput. Ind. 112, 103111. http://dx.doi.org/10.1016/j.compind.2019.07.001.

Martins, J.R.R.A., Ning, A., 2021. Engineering Design Optimization. Cambridge University Press, http://dx.doi.org/10.1017/9781108980647.

Martinsson Bonde, J., 2023. Multidisciplinary analysis client (MDAC) (version v1.9.0). http://dx.doi.org/10.5281/zenodo.7973486.

Martinsson Bonde, J., Brahma, A., Panarotto, M., Isaksson, O., Wärmefjord, K., Söderberg, R., Kipouros, T., Clarkson, P.J., Kressin, J., Andersson, P., 2022. Assessment of weld manufacturability of alternative jet engine structural components through digital experiments. In: ISABE 2022. Ottawa, Canada, http://dx.doi.org/10.5281/zenodo.7973381.

McAdams, D.A., Wood, K.L., 2002. A quantitative similarity metric for design-by-analogy. J. Mech. Des. 124 (2), 173–182. http://dx.doi.org/10.1115/1.1475317.

Nandy, A., Dong, A., Goucher-Lambert, K., 2021. Evaluating quantitative measures for assessing functional similarity in engineering design. J. Mech. Des. 1–42, http://dx.doi.org/10.1115/1.4052302.

Nergård, H., Larsson, T., 2009. Challenges for experience feedback in engineering design. In: Proceedings of the ASME 2009. In: International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, vol. 2, San Diego, California, USA, pp. 835–842. http://dx.doi.org/10.1115/DETC2009-87305.

Niemeyer, J.K., Whitney, D.E., 2002. Risk reduction of jet engine product development using technology readiness metrics. In: Proceedings of the ASME Design Engineering Technical Conference. Vol. 4, pp. 3–13. http://dx.doi.org/10.1115/DETC2002/DTM-34000.

Papalambros, P.Y., Wilde, D.J., 2017. Principles of Optimal Design: Modeling and Computation, third ed. Cambridge University Press, http://dx.doi.org/10.1017/9781316451038.

Runnemalm, H., Tersing, H., Isaksson, O., 2009. Virtual manufacturing of light weight aero engine components. In: ISABE 2009. Montreal, pp. 170–176.

Sandberg, M., Tyapin, I., Kokkolaras, M., Lundbladh, A., Isaksson, O., 2017. A knowledge-based master model approach exemplified with jet engine structural design. Comput. Ind. 85, 31–38. http://dx.doi.org/10.1016/j.compind.2016.12.003.

Simpson, T.W., Maier, J.R., Mistree, F., 2001a. Product platform design: method and application. Res. Eng. Des. 13 (1), 2–22. http://dx.doi.org/10.1007/s001630100002.

Simpson, T.W., Peplinski, J.D., Koch, P.N., Allen, J.K., 2001b. Metamodels for computer-based engineering design: Survey and recommendations. Eng. Comput. 17 (2), 129–150. http://dx.doi.org/10.1007/PL00007198.

Sivaloganathan, S., Shahin, T.M.M., 1999. Design reuse: An overview. Proc. Inst. Mech. Eng. B 213 (7), 641–654. http://dx.doi.org/10.1243/0954405991517092.

Smith, J.S., Duffy, A.H.B., 2001. Re-using knowledge—why, what, and where. In: Proceedings of International Conference on Engineering Design. Glasgow, pp. 227–234.

Su, Z., Ahn, B.-R., Eom, K.-Y., Kang, M.-K., Kim, J.-P., Kim, M.-K., 2008. Plagiarism detection using the levenshtein distance and smith-waterman algorithm. In: 2008 3rd International Conference on Innovative Computing Information and Control. p. 569. http://dx.doi.org/10.1109/ICICIC.2008.422.

Tao, F., Sui, F., Liu, A., Qi, Q., Zhang, M., Song, B., Guo, Z., Lu, S.C., Nee, A.Y., 2019. Digital twin-driven product design framework. Int. J. Prod. Res. 57 (12), 3935–3953. http://dx.doi.org/10.1080/00207543.2018.1443229.

Verhaegen, P.-A., D'hondt, J., Vandevenne, D., Dewulf, S., Duflou, J.R., 2011. Identifying candidates for design-by-analogy. Comput. Ind. 62 (4), 446–459. http://dx.doi.org/10.1016/j.compind.2010.12.007.

Viana, F.A.C., Gogu, C., Goel, T., 2021. Surrogate modeling: tricks that endured the test of time and some recent developments. Struct. Multidiscip. Optim. 64 (5), 2881–2908. http://dx.doi.org/10.1007/s00158-021-03001-2.

Wall, J., Bertoni, M., Larsson, T., 2020. The model-driven decision arena: augmented decision-making for product-service systems design. Systems 8 (2), 22. http://dx.doi.org/10.3390/systems8020022.

Wang, G.G., Shan, S., 2006. Review of Metamodeling Techniques in Support of Engineering Design Optimization. J. Mech. Des. 129 (4), 370–380. http://dx.doi.org/10.1115/1.2429697.

Woodall, P., 2017. The data repurposing challenge: new pressures from data analytics. J. Data Inf. Qual. 8 (3–4), 1–4. http://dx.doi.org/10.1145/3022698.

Woodbury, R.F., Burrow, A.L., 2006. Whither design space? Artif. Intell. Eng. Des. Anal. Manuf. 20 (2), 63–82. http://dx.doi.org/10.1017/S0890060406060057.

Xu, R., Wunsch, D., 2005. Survey of clustering algorithms. IEEE Trans. Neural Netw. 16 (3), 645–678. http://dx.doi.org/10.1109/TNN.2005.845141.

Yondo, R., Andrés, E., Valero, E., 2018. A review on design of experiments and surrogate models in aircraft real-time and many-query aerodynamic analyses. Prog. Aerosp. Sci. 96, 23–61. http://dx.doi.org/10.1016/j.paerosci.2017.11.003.