Comparing probabilistic and entropic strategies for tracking the fragility of design space reduction decisions in set-based design

Van Houten, Joseph; Kana, Austin A.; Singer, David; Collette, Matthew

DOI
10.1016/j.oceaneng.2024.120068

Publication date
2025

Document Version
Final published version

Published in
Ocean Engineering

Citation (APA)
Van Houten, J., Kana, A. A., Singer, D., & Collette, M. (2025). Comparing probabilistic and entropic strategies for tracking the fragility of design space reduction decisions in set-based design. *Ocean Engineering*, *317*, Article 120068. https://doi.org/10.1016/j.oceaneng.2024.120068

Ocean Engineering 317 (2025) 120068

Contents lists available at ScienceDirect

Ocean Engineering

journal homepage: www.elsevier.com/locate/oceaneng

Research paper

Comparing probabilistic and entropic strategies for tracking the fragility of design space reduction decisions in set-based design ☆

Joseph Van Houten [a], *, Austin A. Kana [b], David Singer [a], Matthew Collette [a]

[a] University of Michigan, Ann Arbor, MI, USA
[b] Delft University of Technology, Delft, The Netherlands

ARTICLE INFO

Dataset link: https://github.com/Marine-Structures-Design-Lab/DesignSpace_Fragility/releases/tag/Ocean_Eng2

Keywords:
Set-based design
Space reductions
Information
Fragility
Probability
Entropy

ABSTRACT

Design space reduction decisions made in set-based design use perceptions of feasibility to eliminate unfavorable design solutions from consideration. Perceptions are formed with incomplete information, leaving them susceptible to change if new and conflicting information is made available later in the design process. This paper considers how new information originating from newly sampled design points can alter perceptions of feasibility and introduces a probabilistic and an entropic strategy for quantifying the risk of prematurely eliminating potential design solutions. Emergent designs of automated set-based design simulations gauging this risk are evaluated against ones neglecting it for an analogous design problem. The *Python*-based simulations have different disciplines randomly explore their design spaces and generate reasonable space reduction propositions, and then they give a design manager the opportunity to check the fragility of reduced design spaces before finalizing any reductions. Gathered results indicate that both the probabilistic and entropic models are able to effectively delay design decisions and help disciplines maintain a higher diversity of design solutions while designer understanding is still growing. Both models effectively delay risky space reductions and encourage a more gradual reduction of design spaces compared to simulations not including fragility checks. Furthermore, as the entropic model takes a more holistic approach by working with the history of perceptions formed in a discipline's design space rather than just the newest perceptions, space remaining and diversity results show it slightly outperforming the probabilistic model.

1. Introduction

Design decisions made within the web of interdependencies and requirements ingrained in the marine design process produce complex knowledge structures. While different methods have been proposed to characterize the knowledge generation accompanying these decisions (Braha and Reich, 2003; Hatchuel and Weil, 2009; Shields, 2017; Goodrum, 2020), each one seeks to track and better understand the emergence of (or lack thereof) design solutions. Decisions made in set-based design (SBD) focus on eliminating undesirable solutions from consideration rather than isolating and iterating on discrete solutions. For that reason, SBD decisions build up these knowledge structures more gradually, but they also leave reduced design spaces vulnerable to emergent design failures if the perceptions and information supporting them changes. Moreover, when design failures are imminent, designers using a set-based approach cannot backtrack and reopen design spaces in the same manner that iterative approaches can tweak a design because of the considerable time and resources that are already expended to keep design spaces open. Moreso than any other design approach, successful implementation of SBD depends on robust knowledge generation. Providing designers with a tool to understand the risk for new knowledge to contradict presently generated knowledge before eliminating potential design solutions from consideration would assist them in making much more informed and reliable space reduction decisions.

Using iteration to make decisions and generate knowledge is an understood reality of many complex design problems (Wynn and Eckert, 2017). In an effort to promote an efficient flow of information between iterative tasks, different studies working to enhance both the allocation of resources (Smith and Eppinger, 1997) and communicative pathways (Mihm and Loch, 2006; Parraguez et al., 2015) between said tasks have been investigated. As these strategies are improved upon to assist with iterative design decisions, they can fixate a designer's knowledge on one decision path, restricting the solutions that can be attained through others (Page, 2006). Examples of this fixation are

☆ This document is the results of the research funded by the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program. It is an extension of that presented in Van Houten et al. (2024). Copyright of the original conference paper is held with the authors.

* Corresponding author.
E-mail address: joeyvan@umich.edu (J. Van Houten).

https://doi.org/10.1016/j.oceaneng.2024.120068
Received 11 September 2024; Received in revised form 21 November 2024; Accepted 7 December 2024
Available online 18 December 2024
0029-8018/

shown in Van Houten et al. (2022) where viable solutions within a discipline's design space are significantly limited by the path chosen. In some cases, designers can lose their influence altogether if a finite set of absorbing paths constrain the knowledge structures generated from these temporal decision processes (Niese et al., 2015; Kana, 2017).

A consequence of becoming overly fixated on a particular decision path is leaving a design susceptible to emergent design failures. Dong (2017) discusses the prevalence of this problem in product development when companies introduce innovative technologies into their product's existing functional architecture. He argues that integration issues arise before the establishment of their product's physical architecture and should instead be attributed to the solution principles the design team committed to during development. He and others (Shields and Singer, 2017; Goodrum, 2020) insist that understanding emergent design failures requires a shift in viewing them from a product-centric to a knowledge-centric perspective. As Goodrum (2020) explains, a design decision is a commitment to a knowledge structure, and how those decisions affect future design activities will vary depending on how new knowledge integrates with existing knowledge. In his work, Goodrum (2020) attempts to identify when existing knowledge and new knowledge become incompatible by first mapping out the knowledge-information networks produced through previous decisions and design activities, and then by tracking various entropy-based metrics across subsequent activities. His approach shows a lot of promise, but its reactive nature limits its application to iterative design approaches embracing rework rather than convergent design approaches avoiding it.

SBD is one such convergent design approach that protects against emergent design failures stemming from path fixation by having design decisions focus on eliminating undesirable regions rather than making premature commitments to hard-set characteristics. By delaying commitments and keeping variable sets open, SBD decisions create low-risk knowledge structures (Shields and Singer, 2017) and allow designers to maintain influence over a design problem while their understanding of it grows (Bernstein, 1998; Singer et al., 2009). Advantages of SBD include basing the earliest and most critical design decisions on acquired data, promoting institutional learning within the design environment, encouraging concurrence in the design and manufacturing process, and supporting a search for more globally optimal designs (Ward et al., 1995). These advantages have fueled US Navy interest in making ship design and analysis tools compatible with SBD methods (Doerry, 2012) and applying SBD to various projects such as the Ship to Shore Connector (Mebane et al., 2011), Amphibious Combat Vehicle (Burrow et al., 2014), and Small Surface Combatant (Garner et al., 2015). Despite the advantages, it is still either infrequently applied to problems in industry or generally confined to introductory design stages (Toche et al., 2020). Singer et al. (2009) claim SBD's biggest obstacle in naval design coincides with current government acquisition policies conforming to point-based methodologies. Other hurdles are summarized in McKenney and Singer (2014) and Gumina (2019) and involve having to manage misconceptions *about implementation* and lacking a regimented process *for implementation*.

The SBD implementation process is multifaceted and has disciplines individually explore areas of their design spaces to accumulate information, form perceptions of preferred and nonpreferred areas from this information, and propose space reductions from these perceptions (Bernstein, 1998). A depiction of an example design space is shown in Figure 8 of Andrews (2018) where *space reductions* refer to reducing the range of potential design solutions being left open. A Design Integration Manager (DIM) will then consider the space reductions proposed and the information supporting them to finalize a conceptually robust set of space reductions across all disciplines (Singer et al., 2009). Each of these later steps are directly tied to the information gathered at the beginning, so effective decision-making in SBD necessitates robust information. Gembarski et al. (2021) evaluates the robustness of information in decision-making by using Bayesian probabilities to model uncertainties that originate from a scarcity of information. Sypniewski (2019) takes a different approach and assesses how the inherent biases of information that has already been gathered can lead to inadequate characterization of a design space and misinformed decisions. As the robustness of information pertains to decisions made during SBD specifically, research is limited. Doerry (2015) presents a method for measuring the diversity of information in a design space to increase the likelihood of viable solutions being found later; however, this method intends to insure reduction decisions against uncertain information rather than understand the uncertainty permissible for those decisions to remain advisable. As it currently stands, there is no way for a design manager to track gathered information in SBD for the purposes of proactively gauging the risk of new knowledge to integrate with existing knowledge before committing to a space reduction decision.

The purpose of this paper is to present two new approaches for quantifying the risk of design space reduction decisions in SBD by considering the potential for new information to alter perceptions of feasibility and incite emergent design failures. In the following sections, a brief background on SBD and a design space's fragility (or the vulnerability of its perceptions to new and conflicting information) will first be provided. Next, frameworks built for assessing the fragility of design spaces and quantifying the risk of space reduction decisions from a probabilistic and entropic approach will be explained. The developed fragility frameworks differ in the extent to which they utilize previously formed perceptions to quantify risk and are intended to plug in at the very end of the space reduction process. With the frameworks established, a polynomial design problem and an autonomous SBD simulation are introduced for running experiments that compare emergent design spaces with and without these fragility checks. Results gathered from these experiments illuminate the shifting totality and diversity of potential design solutions maintained throughout the design process when fragility checks are made.

2. Set-based design

SBD is a convergent design approach that seeks a final solution through the gradual elimination of design spaces rather than cycles of rework and refinement synonymous with most iterative approaches. Bernstein (1998) describes the *ideal* way SBD should be performed with illustrative help from Fig. 1 developed by Dr. William Finch. In the early stages of SBD, disciplines individually explore areas of their design spaces and expand their ranges of potential design solutions. From a marine design perspective, these disciplines may consist of (but not be limited to) a weights division negotiating lightship and deadweight tonnage allotments along with center of gravity positioning, a stability division considering allowable beam and vertical center of gravity pairings, and a structural division contemplating various plate thickness and stiffener sizing schemes. As potential solution spaces are identified by each discipline, areas of overlap between their interdependent design spaces that may satisfy *all* requirements of the design problem become pronounced, and designers focus their exploration efforts in these more promising areas. The weights division may have its own displacement and trim requirements to satisfy, but the vertical center of gravity of a load case cannot prevent the stability division from satisfying intact or damage stability requirements, and the lightship allotment must be sufficient for the structural division to satisfy material yielding requirements within a specific safety factor. While this understanding of potential design solutions and trade-offs grows, disciplines propose space reductions that would eliminate nonpreferred areas of their design spaces. It then becomes the DIM's responsibility to consider those proposed reductions and finalize a universal set of reductions for all disciplines to abide by on the grounds of infeasibility (eliminating design solutions that cannot satisfy requirements) or dominance (eliminating design solutions not preferred by a majority of disciplines). This process of elimination is

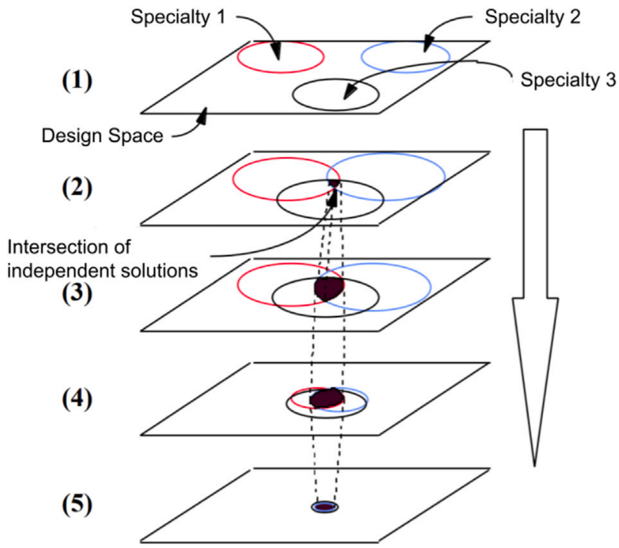

Fig. 1. Ideal convergence of the SBD process through gradual elimination of nonpreferred areas (Bernstein, 1998).

intended to continue until the disciplines have converged on a final, desirable solution satisfying all design requirements.

Through this process, a major principle of SBD is delaying decisions until the consequences of those decisions are understood (Ward et al., 1995; Singer et al., 2009). During discussions with managers utilizing "set-based concurrent engineering" at Toyota, Ward et al. (1995) learned that a critical aspect of their job is to discourage engineers from making important design decisions too soon. They believe it is necessary to delay decisions to ensure all the requirements of the customer are met while also ensuring that the design is manufacturable. Bernstein (1998) and Singer et al. (2009) discuss the benefits of delaying design decisions from the perspectives of accrued knowledge, committed costs, and stakeholder influence. They explain that knowledge of a design is gathered with time as designers run analyses to build their understanding of the characteristics and requirements driving the process. By delaying decisions through a set-based approach, designers can increase the influence maintained and decrease the costs incurred until the information and existing knowledge supporting these decisions is more robust.

Eventually making these reduction decisions is challenging as design spaces cannot be understood absolutely. Different disciplines often manage large design spaces that cannot be explored completely while tolerating analyses with varying degrees of uncertainty. Moreover, it is common for changes in design requirements as well as the fidelity or underlying assumptions of analyses to be introduced throughout the design process that shift preferred and nonpreferred areas. Shields and Singer (2017) assert that space reduction decisions create low-risk knowledge structures while also acknowledging that SBD relies on considerable knowledge generation and decision-making to work effectively. In their words, "Only making decisions when the supporting knowledge is well-understood and is unlikely to change leaves stable knowledge to be further developed" (Shields and Singer, 2017).

This guideline is not only difficult to satisfy because of the limited time and effort aspects, but because designers lack context altogether over what exactly constitutes "well-understood knowledge". Each space reduction decision in SBD is supported by information that is incomplete, uncertain, and susceptible to change. When designers lack the means to account for this uncertainty of information, their reduction decisions may lead to exceedingly *fragile* design spaces, or design spaces whose perceived feasibility is vulnerable to new and conflicting information. In instances when new information does expose fragile design spaces, designers using a SBD approach cannot simply rely on backtracking and reopening design spaces either, because their timelines are limited by the considerable time and effort already spent exploring those design spaces in the first place. Providing designers with context on the robustness of their knowledge generated thus far requires them to assess the fragility of their reduced design spaces.

2.1. Fragility and space reduction decisions

To help visualize a design space's fragility, Fig. 2 has been created to mirror the red, blue, and black circles in the third layer of Bernstein's SBD process (see Fig. 1). In Fig. 2, the perceived feasible regions of each discipline are located within the circles. The green regions signify perceived feasible areas of the design space for one discipline, the yellow regions signify the same perceived feasibility for two regions, and the orange region signifies the same perceived feasibility for all three regions. Suppose the fragility is being assessed from the red discipline's perspective. One source of fragility is attributed to learning new information that alters the perceived feasible space of the red discipline itself, as depicted by the dashed red circle in Fig. 3. The pink region captures the red discipline's newly perceived feasible space, and the grey region captures its newly perceived infeasible space. If new information shifts the perceived feasible space such that the grey region contains more design solutions than the pink region, then the red discipline's originally perceived design space was very fragile.

Another source of fragility is attributed to learning new information that alters the perceived feasible space of an interdependent discipline, as depicted by the dashed blue circle in Fig. 4. Suppose the fragility is again being assessed from the red discipline's perspective. The pink region now captures newly perceived, shared feasible space, and the grey region captures newly perceived, shared infeasible space. If new information shifts the shared feasible space such that the grey region contains more design solutions than the pink region, then the red discipline's originally perceived design space was again very fragile. In both scenarios, a design space's fragility is influenced by the negative effect that new information has on its present perceptions.

While a design space's fragility directly corresponds to its vulnerability to new information, that vulnerability can be amplified by the particular space reductions that have previously been made. In both Figs. 3 and 4, the DIM may have already decided to eliminate portions of the pink region. If that is the case, disciplines would be left without newly perceived feasible space, meaning that the grey region would contain even more design solutions than the pink region. Designers want to avoid space reduction decisions that lead to exceedingly fragile design spaces, yet they must make reductions to keep the design process moving. At each space reduction cycle, every design space is susceptible to increases in design space fragility that can be further exacerbated by previous reductions. By effect, there are varying levels of risk for space reduction decisions due to the varying levels of fragility that result from prior reductions and new information.

2.2. Originating sources of new information

In the development of solutions to complex design problems, designers are compelled to explore and gain an understanding of their own discipline's design space, integrate the understanding and preferences of designers from interdependent disciplines with their own, and endure changing design requirements and maturing analyses throughout the entire process. Bearing each of these challenges in mind, three different sources of new information are worth considering when characterizing the fragility of a design space: (1) newly explored design points of one's own discipline, (2) newly explored design points of other, interdependent disciplines, and (3) new or updated design requirements or analyses. This work only considers the first originating source of new information, but the other two are important to keep in mind for future developments.

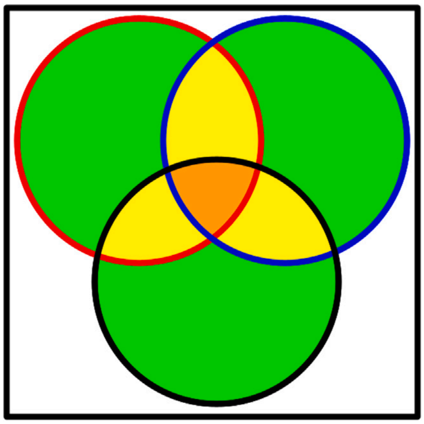

Fig. 2. Overlapping regions of perceived feasible spaces for three disciplines of a design problem (Van Houten et al., 2024).

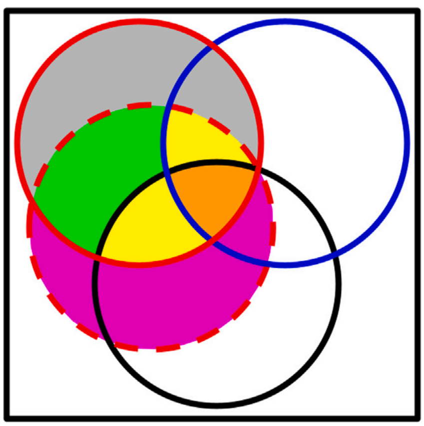

Fig. 3. Fragility attributed to design change of main discipline (Van Houten et al., 2024).

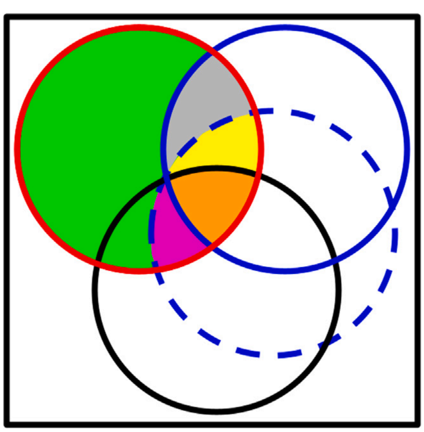

Fig. 4. Fragility attributed to design change of interdependent discipline (Van Houten et al., 2024).

To observe how new information originating from newly explored design points of one's own discipline can impact perceptions of design space behavior, consider Fig. 5. With the information from design points presently available in Fig. 5(a), clear regions of feasibility have been formed for the discipline; designers of this discipline are perceiving smaller values of *Variable 1* to be feasible and larger values of *Variable 1* to be infeasible. However, those perceptions shift in Fig. 5(b) when new information originating from newly tested design points becomes available. Larger values of *Variable 1* are still perceived as infeasible, but designers have also learned they may have less area to work with for smaller values of *Variable 1* than they previously thought. Before learning this new information, suppose the decision is made to eliminate some of the smallest values of *Variable 1* because (in contrast to this discipline) other disciplines prefer large values of *Variable 1* to small values. Designers of this discipline may be inclined to approve the space reduction thinking they still have plenty of feasible space with which to work. Later, they would regret to learn that the space reduction decision has limited far more feasible solutions remaining for them than they originally anticipated.

The intent of a fragility framework will be to protect design spaces against scenarios like the one described. DIMs may be capable of taking proposed space reductions from disciplines and carefully assessing the impact those reductions would have on other disciplines with the information *at hand*, but they lack a tool for understanding the consequences of those reductions if the perceptions formed from that information *changes*.

3. Fragility framework

Traditionally in SBD, the space reduction decision process ends with the set of reductions finalized by the DIM. At this point, designers have explored their own design spaces to form perceptions and propose space reductions, and the DIM has merged them together with the information available through infeasibility- or dominance-based decisions. As discussed though, this process, which only considers present information, leaves reduced design spaces vulnerable to new information.

The intent of a fragility framework is to gauge the vulnerabilities of each discipline's design space to new information before committing to any space reductions. To accomplish this goal, a developed framework will require components that address various complexities inherent to the space reduction process. Table 1 summarizes those space reduction complexities and corresponding fragility framework requirements. In this work, a Probabilistic Fragility Model (PFM) and an Entropic Fragility Model (EFM) are introduced for fragility assessment. The PFM is replicated from Van Houten et al. (2024), while the EFM is a new model introduced in this work to overcome some of the PFM's shortcomings. Both frameworks are still a work in progress and do not address every framework requirement outlined in the table. Still, they address many complexities inherent to SBD's space reduction process and have the potential to be expanded further in future work. Both frameworks have identical beginning and ending steps for how they form their initial perceptions of feasibility and assess the risk of changing perceptions on a design space; their differences lie in how they leverage the history of those formed perceptions in the middle.

3.1. Forming initial perceptions of feasibility

To form initial perceptions of feasibility throughout a design space, designers need to extrapolate data gathered from explored areas of their design spaces to unexplored areas. To that end, the first fragility framework step starts off by calculating a *pass–fail* amount for each explored point in a design space. Calculation of the pass–fail amount is intended to give designers an idea of how much an explored design solution passes or fails a discipline's requirements based on the value of that design solution's output value within the objective space.

For design solutions meeting all requirements, a *pass* amount is calculated as the minimum normalized difference of the calculated output value to the nearest requirement threshold. Eq. (1) shows an example calculation of the pass amount (PA) involving three objective space requirements ($y_1 > 0.2$, $0.3 < y_2 < 0.6$, $y_1 + y_2 < 0.8$), three calculated output values ($y_1 = 0.4$, $y_2 = 0.35$, $y_1 + y_2 = 0.75$), and three ranges of calculated values from explored points ($y_1 \in [0.05, 1.2]$,

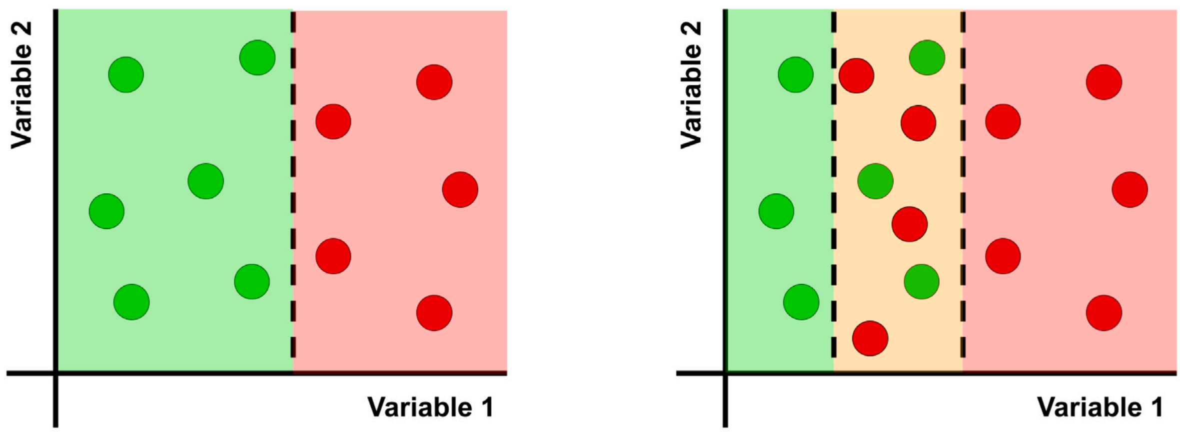

(a) Perceptions of feasibility before sampling new points   (b) Perceptions of feasibility after sampling new points

Fig. 5. Comparison of perceptions of feasibility before and after sampling new points within a design space. Green points represent tested designs that are feasible, red points represent tested designs that are infeasible, green regions represent perceived feasible spaces, red regions represent perceived infeasible spaces, and yellow regions represent spaces of mixed feasibility (Van Houten et al., 2024).

Table 1
Complexities that exist when making space reduction decisions with uncertain information and the corresponding fragility framework requirements addressing these complexities (shaded green rows are addressed in this paper's frameworks) (Van Houten et al., 2024).

Space reduction complexity	Framework requirement
Space reductions are focused on eliminating undesirable solutions from a ranging design space. The desirability of solutions are rooted in perceptions of feasibility formed by running *discrete* design points through the analyses established by each discipline.	The framework needs to form initial perceptions of feasibility with presently available information. A technique for converting information from explored points and their output values into perceptions of feasibility *throughout* each discipline's design space is required.
Perceptions of feasibility are uncertain because they are formed with incomplete information within a discipline's design space. Information from newly analyzed design points *within a design space* could alter perceptions.	Formed perceptions of feasibility for unexplored areas of the design space are not definitive. The framework should account for the possibility of new design points being tested with feasibility that is contradictory to expectations.
The number of ways new information can alter perceptions of feasibility within a design space is *unbounded and unknown* until the information is made available. The risk of a space reduction in context of itself is unlimited.	Comparing the fragility of a reduced design space to a non-reduced design space and determining what new information a discipline *can* handle rather than it would *have to* handle will narrow the DIM's scope and allow space reduction risk to be quantified.
A design space may be fragile when considering all input variables together (i.e. x_1, x_2, x_3) and when considering various *combinations* of input variables (i.e. x_1, x_2).	The framework cannot only measure the fragility of a design space as a whole. It must be flexible enough to also identify component-based fragilities.
Perceptions of feasibility are uncertain because of the *interdependencies* that exist through shared variables between disciplines. Vulnerabilities of one design space to new information could directly or indirectly amplify the vulnerabilities of other design spaces.	The framework must include a cross-discipline component that ties the individualistic fragilities of each discipline together such that the vulnerabilities tracked across interconnected design spaces are representative of their dependencies on each other.
Perceptions of feasibility are uncertain because they are formed with output information that is susceptible to change. New information originating from *changes to design requirements or analyses* could alter perceptions.	The location of calculated output values within the objective space must not be treated as definitive. Instead, the framework should account for the possibility of output values and requirements shifting in relation to each other.

$y_2 \in [0, 0.9]$, $y_1 + y_2 \in [0.05, 2.1]$). Infeasible design solutions will have a pass amount of zero.

$$PA = \min\left(\frac{|0.4 - 0.2|}{1.2 - 0.05}, \frac{|0.35 - 0.3|}{0.9 - 0}, \frac{|0.75 - 0.8|}{2.1 - 0.05}\right) \tag{1}$$

For design solutions failing at least one requirement, a *fail* amount is calculated as the normalized root mean square difference of the calculated output value to all of the requirements. Eq. (2) shows an example calculation of the fail amount (FA) involving three objective space requirements ($y_1 > 0.2$, $0.3 < y_2 < 0.6$, $y_1 + y_2 < 0.8$), three different output values ($y_1 = 0.1$, $y_2 = 0.25$, $y_1 + y_2 = 0.35$), and three ranges of calculated values from explored points ($y_1 \in [0.05, 1.2]$, $y_2 \in [0, 0.9]$, $y_1 + y_2 \in [0.05, 2.1]$). Feasible design solutions will have a fail amount of zero.

$$FA = \sqrt{\frac{\left(\frac{0.1-0.2}{1.2-0.05}\right)^2 + \left(\frac{0.25-0.3}{0.9-0}\right)^2 + \left(\frac{0}{2.1-0.05}\right)^2}{3}} \tag{2}$$

Once the pass and fail amounts are calculated for each explored design solution, the pass–fail amount is simply taken as pass amount minus the fail amount. As either the pass or fail amount of each explored solution will always be zero, subtracting the fail amount from the pass amount is solely done for sign convention purposes. Potential design solutions meeting all objective space requirements will have positive pass–fail amounts, while potential design solutions failing at least one objective space requirement will have negative pass–fail amounts.

Data from each explored design solution's unique input values and calculated pass–fail amount are then used to train a Gaussian Process Regressor (GPR) with a Radial Basis Function (RBF) kernel. The GPR is chosen because both the mean and standard deviations of predictions are used in fragility assessment, and it is paired with the RBF kernel to handle non-linearities in feasible behavior and instill more confidence in predicted behavior for areas closer to explored points. The GPR is allotted a very small noise value for training such that the learned model fits directly through each explored design solution. Finally, remaining (unexplored) design solutions are discretized within the discipline's design space, and the trained GPR is used to predict their mean pass–fail amounts, which are normalized between −1 and +1. Positive pass–fail predictions indicate that an unexplored design solution is perceived as feasible, while negative predictions indicate that an unexplored design solution is infeasible.

Fig. 6 depicts this process for the remaining areas of a design space involving two input variables (x_1 and x_2). On the left-hand side of the

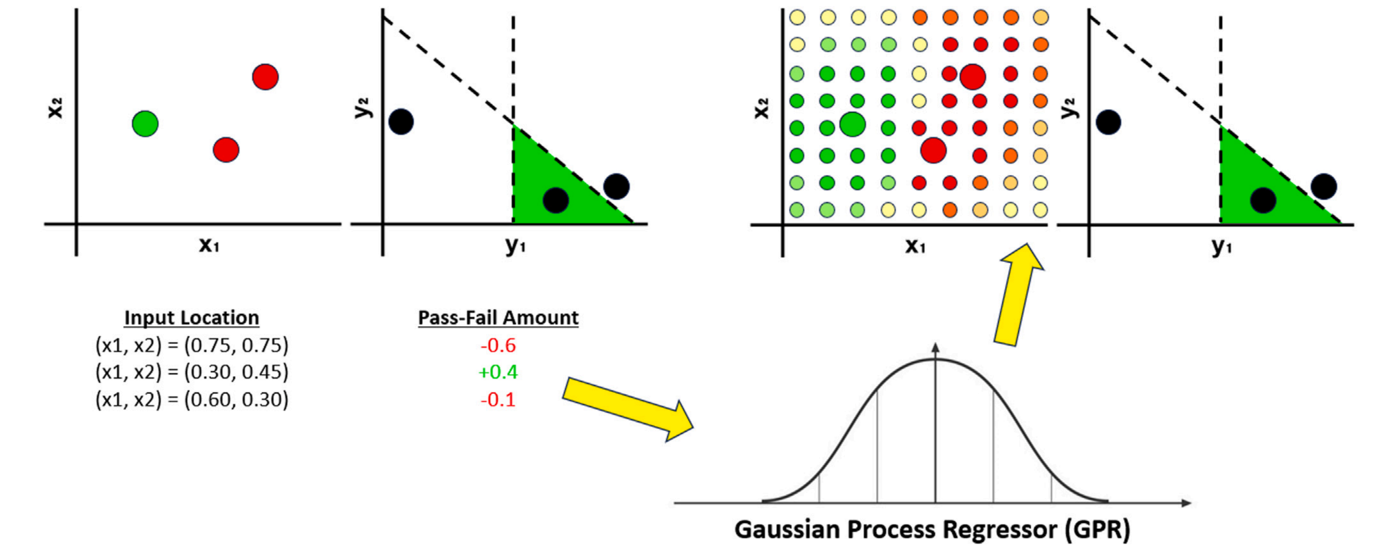

Fig. 6. Forming perceptions of feasibility for unexplored areas of the design space (Van Houten et al., 2024).

figure, pass–fail amounts are formed for three explored (larger) points. Data from those explored points train a GPR, and then the trained GPR forms predictions for the unexplored (smaller) points of the design space.

3.2. Considering the consequences of incorrect perceptions - Regret and windfall

In the next step, designers need to consider the consequences of their formed perceptions of feasibility being incorrect. This requirement leads to the introduction of regret and windfall in a design space. Suppose the sampled design space in Fig. 7 is considering the space reduction depicted by the black box. The space reduction would eliminate portions of the design space perceived as feasible (top-left) as well as portions of the design space perceived as infeasible (top-right). Now suppose new information comes along that throws off those perceptions of feasibility as depicted by the left-hand design space in Fig. 8. This new information would cause designers to regret the space reduction if they are left *with infeasible* space that was expected to be feasible or left *without feasible* space that was expected to be infeasible (instances of regret). In contrast, the new information would benefit designers if they are left *with feasible* space that was expected to be infeasible or left *without infeasible* space that was expected to be feasible (instances of windfall).

Before committing to a space reduction, various fragility models must assess these *potentials* for windfall and regret for the reduced design space in context of the non-reduced design space. This logic allows designers to consider the consequences and quantify the risk of moving forward with a space reduction compared to delaying the space reduction.

3.2.1. Probabilistic fragility model approach

The main idea behind the PFM is to characterize a discipline's present understanding of a design space with straightforward probabilities of feasibility and infeasibility and then to quantify its vulnerability based on how likely those perceptions are to change. In other words, the PFM primarily works with the complementary probabilities for a discipline's unexplored design space at a particular moment in time.

The complementary probabilities of feasibility are attained by leveraging the means and standard deviations of the trained GPR's predicted pass–fail amounts. Each unexplored point in the design space will have a predicted pass–fail amount and a standard deviation associated with that prediction. With those two pieces of information, a normal probability distribution centered around the pass–fail prediction is formed.

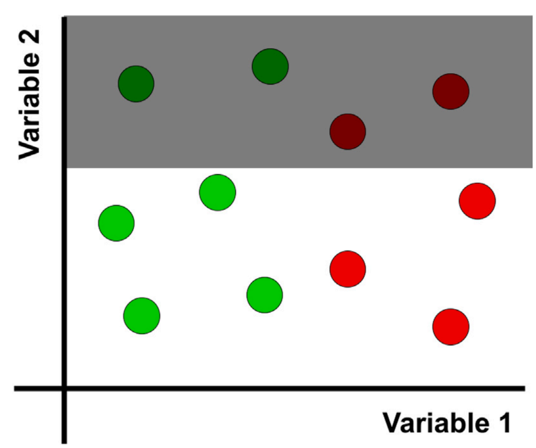

Fig. 7. Design space considering a proposed space reduction (signified by the black box) (Van Houten et al., 2024).

From there, the complementary probability of feasibility or infeasibility for the unexplored point is determined by calculating the portion of the probability distribution lying on the opposite side of zero as the predicted pass–fail amount.

Whether the complementary probabilities contribute to a design space's potential for regret or windfall depends on its presently perceived feasibility and where the point falls in relation to the area of the design space that would be eliminated. Accordingly, the unexplored points' complementary probability of feasibility is added to the proper sum of either regret or windfall potentials for the reduced and non-reduced design spaces. Fig. 9 depicts how these potentials may end up looking between the reduced and non-reduced design spaces for a proposed space reduction. The summed regret and windfall potentials will eventually be used to quantify *added* regret and windfall potentials accompanying a space reduction in Eqs. (5) and (6) of Section 3.3. Because the potentials for regret and windfall of the PFM will be compared to those of the EFM, the complementary probabilities are also normalized before summing.

3.2.2. Entropic fragility model approach

While the PFM may prove to be a straightforward and effective strategy, it will likely be hindered to some extent by temporal biases. Each time a fragility check is performed with the PFM, it only considers

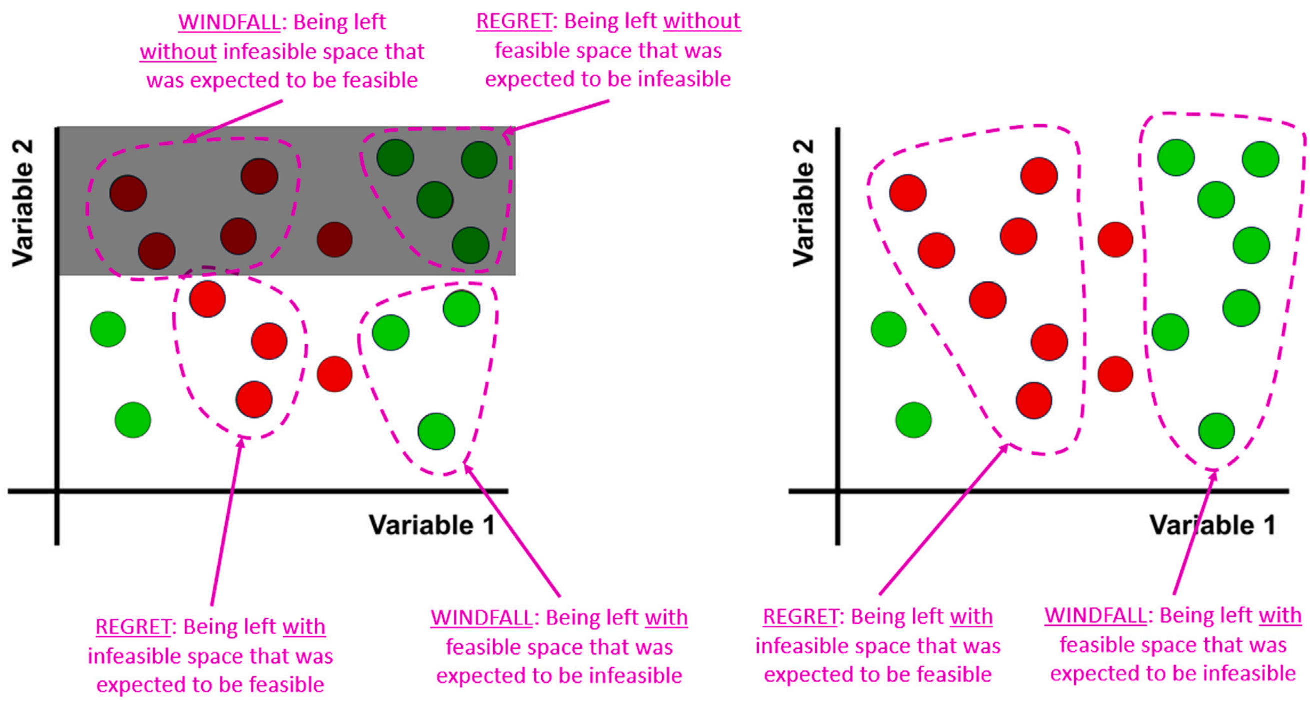

Fig. 8. Instances of regret and windfall for the reduced design space (left) and non-reduced design space (right) (Van Houten et al., 2024).

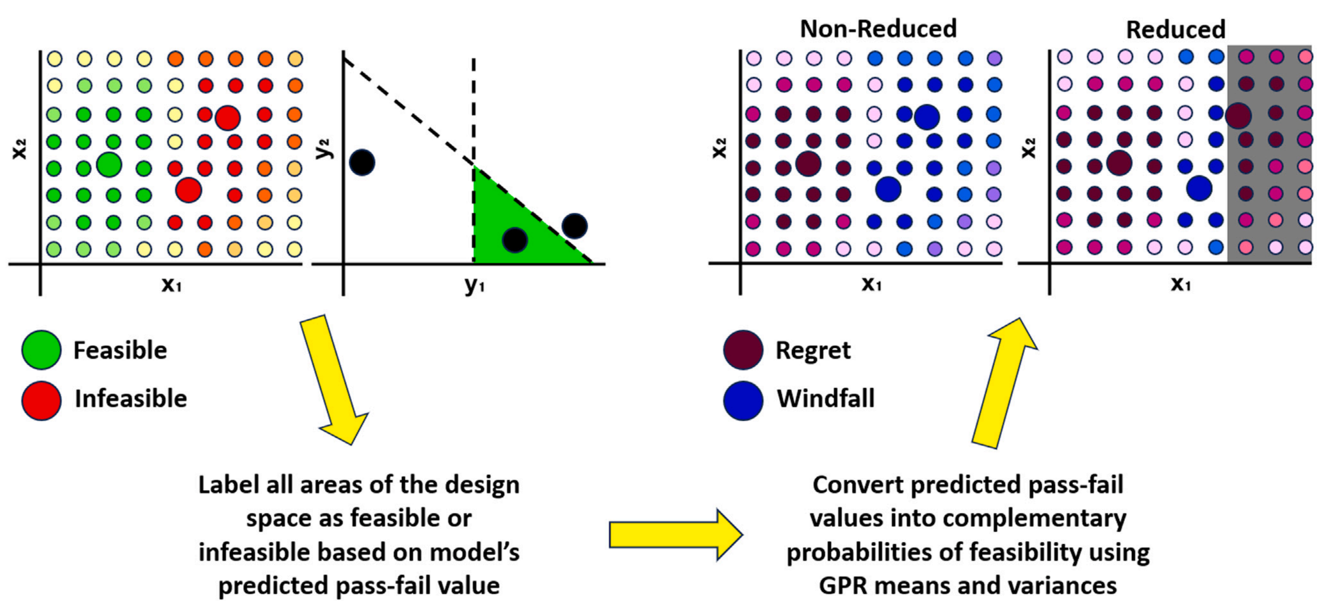

Fig. 9. Converting perceptions of feasibility into potentials for regret and windfall for a non-reduced and reduced design space (Van Houten et al., 2024).

perceptions of feasibility formed at that present moment in time. Those present perceptions should be the most accurate ones formed thus far because they are derived from the either the same or more explored points than past perceptions. However, those present perceptions could still be fluctuating significantly compared to previous ones, indicating that future perceptions are likely to do the same. The PFM disregards any sort of fluctuating perceptions in its risk assessment, limiting the effectiveness of this approach when present perceptions especially contrast actual design space feasibilities. To overcome this temporal handicap, the EFM is developed.

The main idea behind the EFM is to characterize a discipline's present uncertainty of a design space from its *history* of predicted pass–fail values. Rather than working solely with present probabilities, the EFM uses the breadth of information acquired up to a certain point in time to convert those probabilities into information entropies.

To understand the EFM approach, it helps to first understand what exactly is information entropy. Information entropy was first introduced by Shannon (1948) as a means to quantify informational content from a set of events with various probabilities. For a pair of events having equal probability, an observer would have high uncertainty about a future event's outcome, corresponding to a higher entropy. Contrarily, if the events have more of an 80%–20% probability split, an observer is more certain of the outcome of a future event, corresponding to a lower entropy.

Since Shannon's introduction of information entropy, there have been plenty of different extensions introduced. One such extension called Generalized Cumulative Residual Entropy (GCRE) works with the *cumulative* probability distributions of events (Rao et al., 2004). An advantage of GCRE over Shannon Entropy (SE) is that GCRE naturally incorporates the magnitudes of potential outcomes into an observer's uncertainty quantification. Goodrum (2020) uses the same logic behind GCRE to introduce an entropy measure working with a time–history of outcomes called Target Value Entropy (TVE). Eq. (3) shows how TVE is calculated for a time–history of outcomes ($V = \{V_0, V_1, \dots, V_n\}$). As more points are explored within a design space and new pass–fail predictions are formed over time, TVE can give designers an idea of how uncertain they are for those predictions to keep fluctuating. High TVE values would indicate pass–fail predictions are continuing to fluctuate and the design space is still quite fragile, and vice versa for low TVE values.

$$TVE(V) = -\int_{-\infty}^{\infty} p(V > v)\log_2 p(V > v)dv \tag{3}$$

After every exploration cycle, new pass–fail predictions are made for every unexplored design point remaining within each discipline's design space. These predictions will also have a history of standard deviations associated with them that will presumably decrease as more points are explored. To calculate the TVE for each unexplored point, a probability needs to be assigned to each of its pass–fail predictions. Those probabilities are formed by inverting the standard deviations of the predictions, summing the inverted standard deviations together, and then assigning the probability as the fraction of the inverted standard deviation to the sum. If an unexplored point has predicted pass–fail amounts of −0.1, 0.0, and 0.2 with respective standard deviations (σ) of 0.1, 0.2, and 0.4, those pass–fail amounts would be assigned respective probabilities of about 0.57, 0.29, and 0.14 (where Eq. (4) provides an example calculation for the first probability, p_1).

$$p_1 = \frac{\sigma_1^{-1}}{\sum_{i=1}^{n} \sigma_i^{-1}} = \frac{0.1^{-1}}{0.1^{-1} + 0.2^{-1} + 0.4^{-1}} = 0.57 \tag{4}$$

To be able to calculate TVE after every exploration cycle, including the first one, pass–fail predictions also need to be made for the special zeroth time case. As no points have been explored yet, a designer would have no idea of what areas of the design space are passing or failing; for that reason, each unexplored point is assigned a pass–fail amount in the middle at 0.0 with a standard deviation of $\frac{1}{\sqrt{3}}$, which is reflective of a uniform distribution ranging between −1 and +1.

With the history of pass–fail predictions and standard deviations gathered for each unexplored point, the TVE of those predictions can now be calculated. Whether those TVE values contribute to a design space's potential for regret or windfall follows the same exact logic as is done for the complementary probabilities calculated in the PFM. Depending on an unexplored point's *presently* perceived feasibility and where it falls in relation to the area of the design space that would be eliminated, its normalized TVE is added to the proper sum of regret or windfall potentials for the reduced and non-reduced design spaces. And in a similar fashion, the summed windfall and regret potentials will be used to quantify *added* regret and windfall potentials in the following subsection.

3.3. Assessing risk from regret and windfall

Once the potentials for regret and windfall are calculated for all the unexplored areas of a design space with the PFM or EFM, the risk of committing to a space reduction can be quantified. Gathering these potentials for both the reduced and non-reduced design spaces permits the DIM to gauge the risk of approving a space reduction in context of leaving a design space untouched.

Eqs. (5) and (6) introduce an added potential for regret metric (Δ_{reg}) and an added potential for windfall metric (Δ_{wind}) based on the summation of the regret and windfall potentials (ϕ_{reg} and ϕ_{wind}) of each unexplored point (x) in the reduced and non-reduced design spaces. Each metric tracks how opportunities for regret and windfall would shift if the DIM moves forward with a certain space reduction. When these metrics are calculated, they are combined in Eq. (7) to create an endured added risk metric ($Risk$) that weighs the potential for new information to be introduced that would hurt designers against new information that would benefit them.

$$\Delta_{reg} = \frac{\sum_{i=1}^{n} \phi_{reg,red}(x)}{\sum_{i=1}^{n} \phi_{reg,nonred}(x)} - 1 \tag{5}$$

$$\Delta_{wind} = \frac{\sum_{i=1}^{n} \phi_{wind,red}(x)}{\sum_{i=1}^{n} \phi_{wind,nonred}(x)} - 1 \tag{6}$$

$$Risk = \Delta_{reg} - \Delta_{wind} \tag{7}$$

At last, the risk of moving forward with a space reduction decision is quantified in context of forgoing the space reduction, and the DIM can decide if they are willing to accept that risk at a particular moment in time of the design process. Low risk values align with space reductions that see greater potentials for windfall and lesser potentials for regret in a reduced design space, while high risk values align with the opposite. Both the endured added risk and the added risk that a DIM is willing to accept will fluctuate as a design problem progresses. To next observe how emergent design spaces are impacted with this new space reduction risk consideration in mind, the fragility frameworks need a design problem and a SBD simulation into which they can be inserted.

4. Introducing an analogous design problem and SBD simulation

Investigating the impacts that a SBD process including fragility checks has on emergent design spaces requires a design problem. A more complex problem could be adopted from other papers or case studies, but using a simple problem for this initial assessment keeps focus directed on the frameworks and makes interpretation of results in their early stages of development more straightforward.

As shown in Fig. 10, the design problem created for this work involves three different disciplines having some shared input variables and unique output variables. The input variables are analogous to the different ship characteristics that a discipline has influence over, while the output variables are analogous to the different ship performance characteristics with which a discipline is concerned. This design problem involves polynomial mathematical equations (shown in Eqs. (8)–(12)) which are meant to be analogous to the different parametric models or design programs with which marine design disciplines may be working. Each discipline also has requirements that must be satisfied for their input and output variables. The bounds on all the input variables are normalized between 0 and 1. The bounds on the output variables are unique and described as follows: $0 \leq y_1 \leq 0.4$ or $1.2 \leq y_1 \leq 1.6$, $0.5 \leq y_2 \leq 0.7$, $0.2 \leq y_3 \leq 0.5$, $0 \leq y_4 \leq 0.5$, $0.8 \leq y_5 \leq 1.6$. For a more detailed explanation of the design problem and visualizations of each discipline's feasible design spaces, see Van Houten et al. (2024).

Discipline 1:

$$y_1 = 0.8x_1^2 + 2x_2^2 - x_3 \tag{8}$$

Discipline 2:

$$y_2 = 1.25x_5 - 12.5x_3^3 + 6.25x_3^2 \tag{9}$$

$$y_3 = \left(x_4^3 + x_5\right)^2 \tag{10}$$

Discipline 3:

$$y_4 = 2x_5 + 0.2\sin\left(25x_6\right) - x_1^{\frac{1}{5}} \tag{11}$$

$$y_5 = x_1^{\frac{1}{3}} - \cos\left(3x_5\right) \tag{12}$$

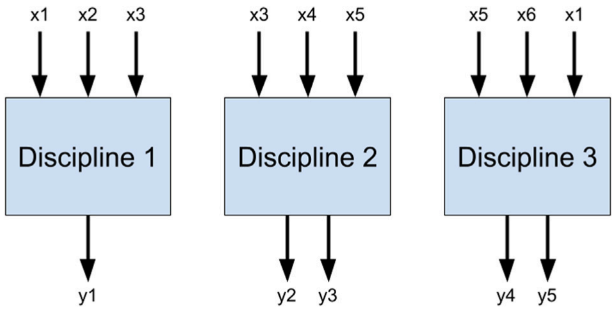

Fig. 10. Input and output variables for three disciplines of the polynomial design problem (Van Houten et al., 2024).

To fairly evaluate the fragility of emergent design spaces for this design problem with and without the fragility frameworks, a SBD simulation that proposes reasonable space reductions is also needed. With the simulation, experiments can be run that compare emergent design spaces when there are fragility checks being made with the PFM or EFM compared to when there are no fragility checks. The simulation created is coded in *Python* and is designed to operate autonomously for a couple of reasons. For one, automating the simulation removes the impact that human inconsistencies would have on an experiment's emergent design spaces by ensuring the same criteria are used to explore design spaces and propose space reductions every time. Additionally, automating the simulation cuts back on the time it would take for a human designer to evaluate the present state of the data and formulate their next exploration or reduction decision. The simulation is not meant to be a perfect replication of how SBD activities are performed and reductions are made because SBD is fundamentally a human-centric process that is driven by knowledgeable designers. A simplified depiction of how the SBD simulation works is shown in Fig. 11, while a more detailed depiction of the simulation and all of the code can be viewed via the link in the *Data Access Statement*. Different parts of the simulation fall under the groupings of problem setup, exploration, or space reduction. Each one of these groupings are discussed in detail in Van Houten et al. (2024).

5. Assessing emergent design spaces

With the fragility frameworks, design problem, and SBD simulation established, different tests could be conducted to observe each discipline's emergent design spaces with and without fragility checks. The test cases are set up such that only the type (or absence) of fragility model differs; in each test case, design spaces are explored, space reductions are proposed, and fragility thresholds are set using the same, predefined process. The results evaluate the fragility of design spaces following space reductions by tracking each discipline's total and feasible design solutions remaining along with the diversity of those remaining design solutions over the elapsed project time.

5.1. Experimental setup

Different levels of space reduction and fragility checks are conducted for each test case of the SBD simulations as depicted in Table 2. The purpose of reducing the design spaces is to leave the design team with a much smaller and more manageable set of potential design solutions to pursue further. For the first test case, no fragility checks are performed at all; rather, it only gathers and considers opinions from each discipline based on *present* information before having the DIM finalize any space reductions. For the second two test cases, fragility checks are performed with either the PFM or EFM to assess the vulnerabilities of reduced design spaces to *new* information.

Each test case is run over a complete project timeline of 50 and 200 time iterations, and all simulations are run with the same effort and goal of reducing design spaces down to roughly 5% of their original size. Ideally, both frameworks should delay space reductions at *no added cost* to the design team. No added cost means that the total design solutions remaining at the end of the elapsed project time are not any greater than simulations conducted without any fragility checks.

Whenever a discipline chooses to run their analyses and explore another potential design solution, that expends time in the simulation. Exploring a new point in *Discipline 1* takes 2 time iterations, exploring a point in *Discipline 2* takes 3 time iterations, and exploring a point in *Discipline 3* takes 4 time iterations. All explored points are treated equally to reflect designers using the same analyses over the entire project time and no one discipline being more of a design driver than another. Explored points are created by assigning a uniform random value between 0 and 1 to each design space's coordinates. If the new point created falls within the remaining design space, then it is utilized for exploration. Otherwise, the process is repeated by assigning a new uniform random value to each coordinate. Designers would likely institute a more directed sampling approach than described for an actual design problem, but for purposes of investigating fragility in design spaces between the PFM and EFM, this randomized approach will suffice.

The complete project timeline relative to each discipline's analysis execution time are made drastically different between the 50 and 200 iteration experiments. These differing timelines force designers to train GPRs and form perceptions of feasibility with varying amounts of explored points so that emergent design space behavior can be observed under different levels of uncertainty.

Each discipline is also limited to proposing space reductions involving only one of their design variables. By nature of the design problem, any proposed space reductions involving x_1, x_3, or x_5 will directly affect two disciplines, while any proposed space reductions involving x_2, x_4, or x_6 will directly affect one discipline.

All other user inputs chosen for the SBD simulations are displayed in Table A.1, with their descriptions available in the *script.py* file of the *Python* code (see *Data Access Statement*). Because explored points are randomly chosen in any given simulation, the perceptions formed and space reductions proposed fluctuate from simulation-to-simulation, regardless of any fragility checks performed. To accommodate for this fluctuation, each test case is executed for 200 complete SBD simulations, and the averages of those simulations are gathered when examining the final results.

5.2. Establishing a maximum risk threshold

When faced with the endured added risk of a space reduction, a DIM could very well decide for themselves what constitutes a high-risk space reduction. However, to keep the SBD simulations completely automated and to maintain a fair comparison of the final results between the PFM and the EFM, a maximum added risk threshold the DIM is willing to endure at any moment in time is needed.

When setting that maximum added risk threshold, a DIM will likely be weighing the amount of each discipline's design space that has already been reduced along with the amount of time remaining for the design project. DIMs should be less inclined to take on risk when much of a design space has been reduced and/or significant project time remains, and they should be more inclined to take on risk when little of a design space has been reduced and/or little project time remains. With that logic in mind, Eq. (13) is used to set a reasonable maximum added risk threshold ($Risk_{max}$) for each discipline at any moment in time (t). If the endured added risk of a space reduction exceeds the maximum added risk threshold, then the reduction will be delayed, and vice versa if the endured added risk of a space reduction falls below the threshold.

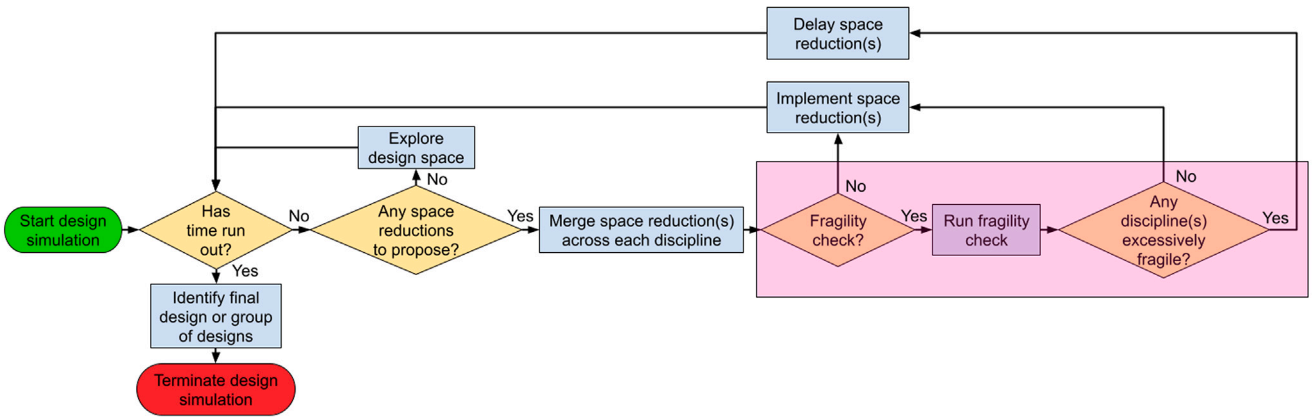

Fig. 11. Simplified flowchart of the automated SBD simulation with a fragility framework plug-in (depicted by the pink box) (Van Houten et al., 2024).

Table 2
Extent of space reduction checks for different test cases of the 50 and 200 iteration simulations.

Test case	Description	Fragility model
1	Only consider present opinions of other disciplines for each proposed space reduction	–
2	Consider present opinions of other disciplines and conduct fragility checks	PFM
3	Consider present opinions of other disciplines and conduct fragility checks	EFM

In the equation, c_{risk} acts as a scaling coefficient, $space_{rem}$ refers to the fraction of a discipline's remaining design space before the newest space reduction, $space_{max}$ refers to the fraction of a discipline's maximum allowable remaining design space, and $time_{rem}$ refers to the fraction of time remaining in the design project. Each multiplicative factor in Eq. (13) plays an important role in setting a discipline's maximum added risk threshold.

$$Risk_{max}(t) = c_{risk} * \frac{space_{rem}(t)}{space_{max}(t)} * \frac{1}{time_{rem}(t)} \tag{13}$$

The asymptotic factor $\left(\frac{1}{time_{rem}(t)}\right)$ increases the threshold as project time is depleted. As fragility checks intend to protect reduced design spaces against new and conflicting information, the threshold should be lowest when there is more time for that information to be made available. Towards the end of a design project, there should be little to no concern for having fragile design spaces as there should be very little opportunity for new information to alter formed perceptions.

The adaptive factor $\left(\frac{space_{rem}(t)}{space_{max}(t)}\right)$ adjusts the rate at which the threshold increases towards infinity based on the design space that remains relative to a predetermined space reduction pace. For the automated SBD simulation, the maximum space remaining consists of an exponential equation dictating when a discipline needs to relax constraints that discourage space reductions. If space reductions within a discipline are keeping up with that pace, the adaptive factor will decrease the maximum added risk threshold. If space reductions within a discipline are falling behind that pace, the adaptive factor will increase the maximum added risk threshold.

The DIM factor (c_{risk}) allows the user to put their own touch on how influential a role they want fragility considerations to play in space reduction decisions. They can use their own preferences to either set a larger DIM factor to discourage delaying space reductions on the grounds of fragility, or set a smaller DIM factor to encourage them.

5.3. Space remaining and diversity metrics

The final results will display the remaining size of a discipline's total and feasible design spaces as well as the diversity of those remaining designs over the elapsed project timeline.

The 'Total Space' curves show the total potential design solutions remaining in a discipline, and the 'Feasible' curves show the feasible design solutions remaining in a discipline. The 'Feasible-to-Remaining' curves combine the two to show the ratio of feasible-to-total designs remaining in a discipline. Test cases with the fragility checks are susceptible to decreasing ratios because they are more inclined to delay space reductions and maintain more total designs; however, also maintaining higher totals of feasible designs can alleviate the ratio reduction.

The curves on the accompanying design space discrepancy graphs display the uniformity of potential design solutions remaining in a discipline. The discrepancy measure is adopted from *SciPy*'s Quasi-Monte Carlo submodule. The default, centered-discrepancy method is utilized and explained thoroughly in Zhou et al. (2013). The metric quantifies discrepancy by measuring the distance between a hypercube's uniform distribution and each discrete design solution remaining. A low discrepancy close to zero indicates that the distribution of remaining design solutions is uniform, while larger discrepancy values indicate the distribution of remaining design solutions is less uniform and more *locked-in* on certain areas of the hypercube. Maintaining a low discrepancy corresponds to maintaining a high diversity of potential design solutions. A high diversity of remaining designs gives a discipline more flexibility to adjust to new information that shifts present perceptions of feasibility.

5.4. Results and discussion

Average emergent design space and discrepancy results are gathered across each discipline of the SBD problem for each test case outlined in Table 2. The results of the simulations are displayed in Figs. 12–17.

The total space remaining results across each discipline indicate that both fragility models support a more gradual reduction of design spaces. When observing the 'Total Space' curves across each discipline of the 50 time iteration simulations (see Figs. 12(a), 14(a), and 16(a)), the test case without any fragility checks (TC1) sees the quickest reduction of each discipline's design space, followed by the test case with the PFM (TC2), and then by the test case with the EFM (TC3). In fact, TC1 sees a very abrupt reduction of each discipline's design space from roughly 20%–50% of the elapsed project timeline, before its 'Total Space' remaining curve across each discipline starts to flatten out. TC2 and TC3, on the other hand, always maintain at least 65% of their original design space by the time they reach 50% of the elapsed project timeline, without any abrupt decreases. This trend is further

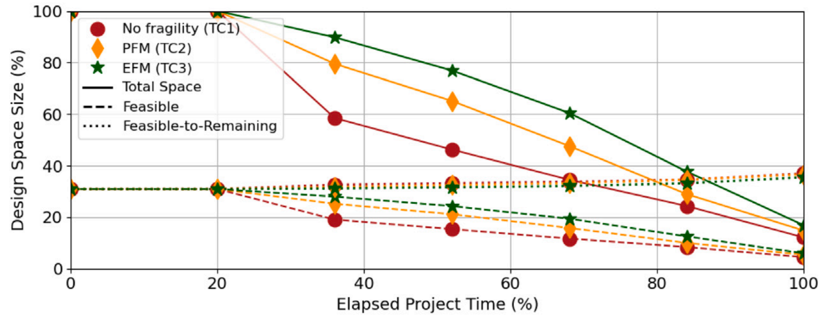

(a) Space remaining of emergent design spaces

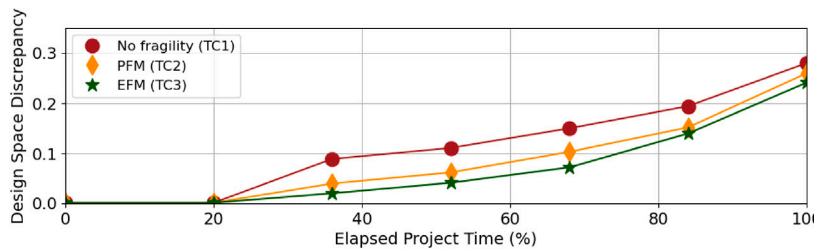

(b) Discrepancy of emergent design spaces

Fig. 12. Space remaining and discrepancy results for emergent design spaces of *Discipline 1*–50 time iterations.

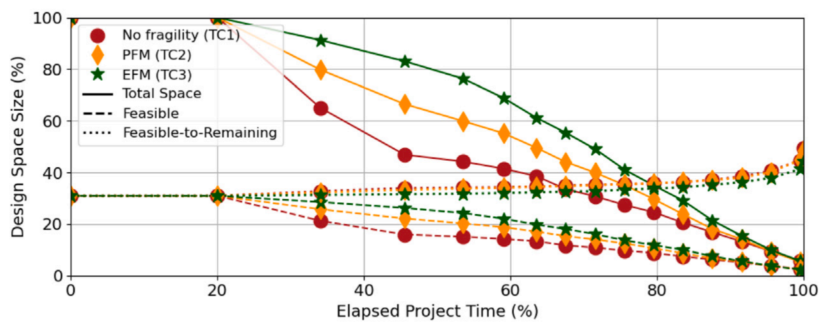

(a) Space remaining of emergent design spaces

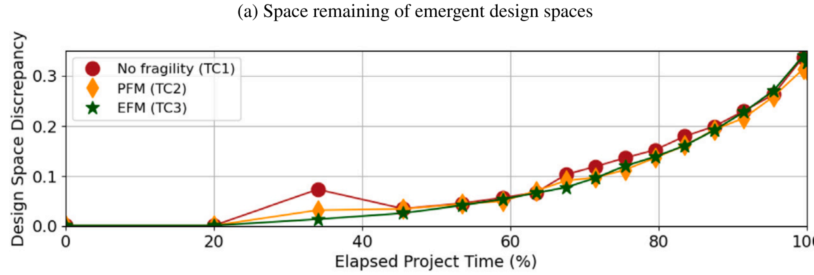

(b) Discrepancy of emergent design spaces

Fig. 13. Space remaining and discrepancy results for emergent design spaces of *Discipline 1*–200 time iterations.

accentuated across each discipline of the 200 time iteration simulations (see Figs. 13(a), 15(a), and 17(a)). Looking at Fig. 17(a) specifically, *Discipline 3*'s total design space remaining at 45% of the elapsed project time is about 43% for TC1, 73% for TC2, and 96% for TC3. Design teams not conducting fragility checks may think that these abrupt space reductions are advisable as long as they primarily involve potential design solutions perceived as infeasible, but in doing so, they ignore the fact that quickly eliminating these areas leave them more vulnerable to new information that end up shifting those perceptions.

The more gradual reduction of design spaces accompanying test cases with the fragility checks also helps each discipline maintain their feasible design solutions longer. This trend is most apparent in *Discipline 1* where over 30% of its initial design space contains feasible design solutions (see Figs. 12(a) and 13(a)). The 'Feasible' curves both start and end at roughly the same points across each test case of *Discipline 1*, but significant gaps form between these curves in the middle of the simulations. TC1's abrupt space reductions lead to quicker decreases in the feasible solutions remaining, followed by TC2, and then by TC3. In the 40%–50% elapsed time frame, there is anywhere from a 10%–15% difference in the feasible design solutions remaining between TC1 and TC3 for both the 50 and 200 time iteration simulations. This lack of feasible designs for TC1 can become a major problem for the discipline if new design requirements or analyses are introduced at this stage in the design process that further rule out these potential design solutions. These feasibility trends between each test case also occur in *Discipline*

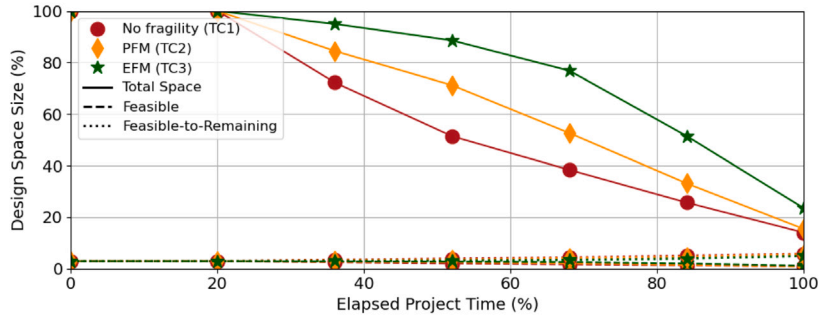

(a) Space remaining of emergent design spaces

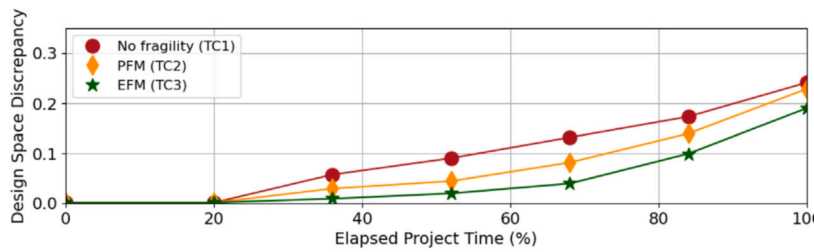

(b) Discrepancy of emergent design spaces

Fig. 14. Space remaining and discrepancy results for emergent design spaces of *Discipline 2*–50 time iterations.

2 and *Discipline 3*, but they are not as noticeable in the figures because each of these disciplines begin with much smaller feasible spaces than *Discipline 1*.

The drawback of delaying space reductions and maintaining more feasible design solutions is having to maintain more infeasible design solutions as well, especially when it is entirely possible that present perceptions of feasibility will not change. This drawback is most apparent when observing the 'Feasible-to-Remaining' curves of *Discipline 3* in Fig. 17(a). As *Discipline 3*'s design space is reduced much more abruptly for TC1 than TC2 and TC3, the feasible-to-remaining design space ratio for TC1 increases more quickly than those test cases including fragility checks. Designers of TC1 may have fewer feasible design solutions available, but they can also disregard far more currently infeasible design solutions than designers of TC2 and TC3. For the 50 time iteration simulations, delaying space reductions with the fragility frameworks can lead to a noticeable gap between the total design solutions remaining across each test case. That gap elicits a subsequent reduction in the feasible-to-remaining design space ratios at 100% of the elapsed project time. Even so, DIMs should be willing to accept this drawback because of the greater flexibility the fragility framework provides them when there are still uncertainties about design space feasibility in the time prior. Unlike the 50 time iteration fragility test cases, none of the 200 time iteration test cases suffer from this drawback. Evidently, designers of these test cases are given ample time to form their perceptions of feasibility, and instituting either the PFM or EFM is remarkably at no added cost to their allotted time and resources.

Finally, the more gradual reduction of design spaces accompanying both fragility models also results in each discipline maintaining a higher diversity of remaining design solutions for longer. Maintaining a higher diversity of potential design solutions helps ensure each discipline is less vulnerable to new information that changes their current perceptions of feasibility by providing designers with more alternative solutions. When observing the discrepancy results across each discipline of the 50 time iteration simulations (see Figs. 12(b), 14(b), and 16(b)), TC1 increases the fastest, followed by TC2, and then TC3. These results indicate that simulations with the EFM maintain the highest diversity of potential design solutions, followed by simulations with the PFM, and then simulations without any fragility checks. These discrepancy trends still generally hold for the 200 time iteration simulations (see Figs. 13(b), 15(b), and 17(b)), but to lesser extents. In *Discipline 1*, the PFM and EFM test cases have discrepancies that generally overlap and lie just below or also overlap that of TC1. In *Discipline 2*, the EFM test case maintains a low discrepancy while the PFM test case discrepancy falls more in-line with TC1. And in *Discipline 3*, the EFM test case maintains a low discrepancy while the PFM discrepancy initially aligns more with that of TC1 before aligning more with the EFM test case later. These discrepancy results between the 50 and 200 time iteration simulations indicate that the fragility models help disciplines more effectively maintain diversity within their design spaces when project timelines are shorter, with the EFM slightly outperforming the PFM.

6. Conclusions

The main objective of this research was to explore the efficacy of two different approaches that mitigate the vulnerabilities of reduced design spaces to new information before making a design space reduction decision. Both approaches similarly begin by forming present perceptions of feasibility from the potential design solutions explored thus far and end by comparing a reduced design space to its non-reduced design space to quantify the risk of a space reduction inciting emergent design failures. They differ by either quantifying fragility from the present probabilities of feasibility (PFM) or from the entropies characterizing the time history of formed perceptions (EFM). Once each fragility framework was introduced, experiments were conducted on an interdependent design problem with a SBD simulation coded in *Python* to observe emergent design spaces with and without these fragility checks.

The results show that SBD simulations including the fragility checks support a more gradual reduction of design spaces with less tendency to lock-in on certain designs prematurely. Both the PFM and the EFM helped reasonably delay space reduction decisions for the design problem while more potential design solutions are explored and understanding of design space behavior grows, although the EFM did so to a greater extent. The EFM consistently maintained the highest total and feasible designs remaining within each discipline as the design project progressed, which also resulted in those design spaces having a higher diversity of potential design solutions for longer. For the 200 time iteration simulations, none of the delayed space reductions prompted

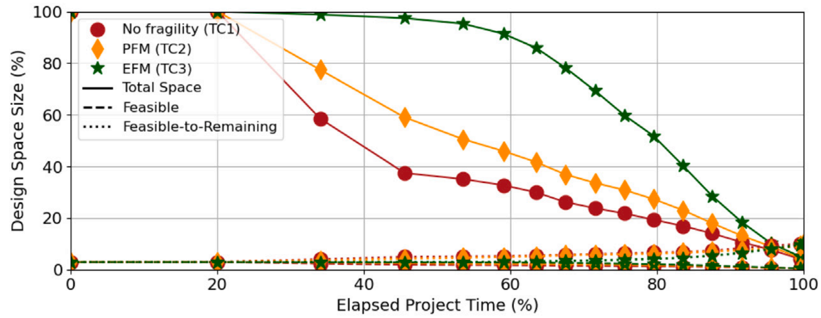

(a) Space remaining of emergent design spaces

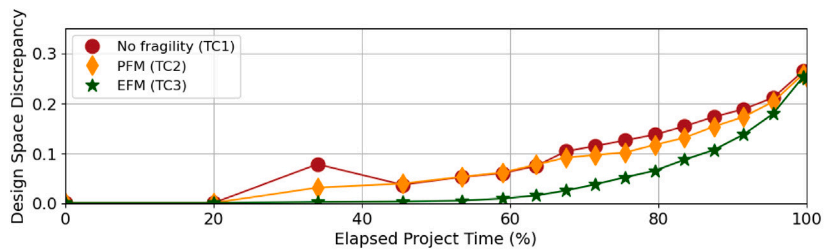

(b) Discrepancy of emergent design spaces

Fig. 15. Space remaining and discrepancy results for emergent design spaces of *Discipline 2*–200 time iterations.

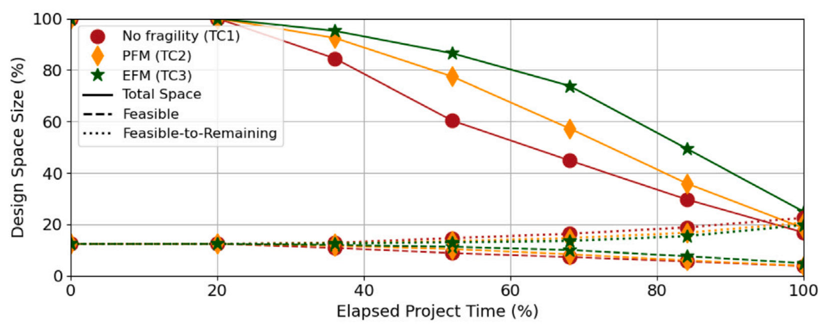

(a) Space remaining of emergent design spaces

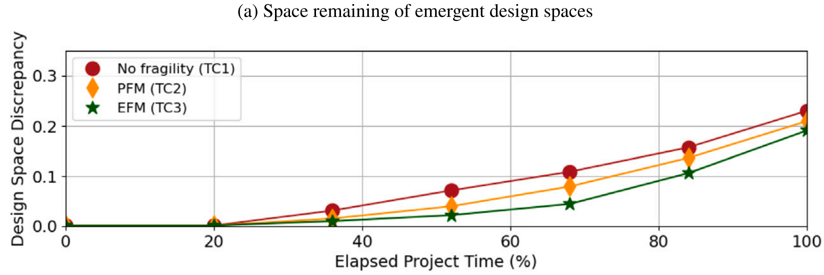

(b) Discrepancy of emergent design spaces

Fig. 16. Space remaining and discrepancy results for emergent design spaces of *Discipline 3*–50 time iterations.

by the PFM or EFM was at an added cost to the designers. The total space remaining results of those simulations always realigned at the end of the project, while the 50 time iteration test cases saw only minor increases in designs remaining.

In spite of this initial evidence, it is still hard to declare with absolute certainty that the EFM is superior to the PFM. There could be situations where fluctuating perceptions of feasibility end up hampering space reduction decisions made with the EFM far more than they would with the PFM. Furthermore, and as already mentioned, neither the PFM or EFM has been built up enough to address every fragility framework requirement outlined in Table 1. More specifically, neither framework assesses the fragility of design spaces on a component-based level, connects fragility across interdependent design spaces, or considers the effect that changes to design requirements or analyses would have on a design space's fragility.

With those deficiencies in mind, there are several extensions for this work moving forward. For one, both fragility models can be improved by considering fragility risks for various combinations of design variables within each discipline rather than only considering the fragility of each discipline's design space as a whole. Both models can also start

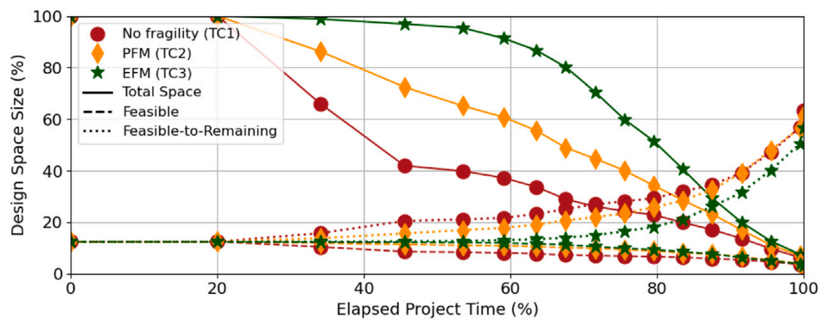

(a) Space remaining of emergent design spaces

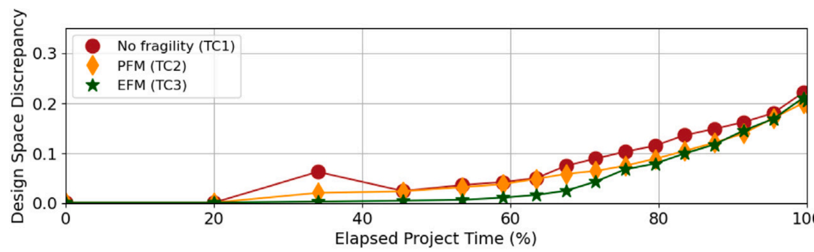

(b) Discrepancy of emergent design spaces

Fig. 17. Space remaining and discrepancy results for emergent design spaces of *Discipline 3*–200 time iterations.

Table A.1
User inputs established in Python for running SBD simulations (Van Houten et al., 2024).

Simulation parameters	Parameter values
problem_name	‘SBD1’
iters_max	50 or 200
sample	‘uniform’
search_factor	100
total_points	10 000
run_time	[2, 3, 4]
exp_parameters	array([0.2, 2.2, 1.0, 0.95])
auto_accept	False
fragility	False (Test case 1) or True (Test case 2 and 3)
fragility_type	‘PFM’ (Test case 2) or ‘EFM’ (Test case 3)
fragility_shift	0.4
change_design	[]
change_time	[]
part_params	{‘cdf_crit’: [0.1, 0.1], ‘fail_crit’: [0.0, 0.05], ‘dist_crit’: [0.2, 0.1], ‘disc_crit’: [0.2, 0.1]}
dtc_kwargs	{‘max_depth’: 1}
gpr_params	{‘length_scale_bounds’: (1e−2, 1e3), ‘alpha’: 0.00001}
bez_point	{‘P0’: (0.0, 1.0), ‘P1’: (0.5, 0.8), ‘P2’: (1.0, 0.0)}

considering some other originating sources of new information, such as new information arising from newly explored design points of interdependent disciplines as well as new information arising from changes to design requirements or analyses. Currently, the fragility frameworks treat all information equally when forming perceptions of feasibility with the GPRs; in the future, the GPRs should be made more flexible so as to handle information coming from analyses with varying fidelity and importance. Furthermore, the frameworks either implement space reductions or delay space reductions based on the results of the fragility checks, with no in between. Introducing strategies that either guide sampling or reduce the size of a space reduction to alleviate design space fragility may give designers a little more flexibility when dealing with a high risk space reduction. Eventually, the PFM and EFM should be evaluated for more complex design problems that have some more direct marine design applications as well. As these additions are made to both fragility models, they will provide DIMs with context as to whether or not their space reduction decisions are leaving disciplines exceedingly vulnerable to new information and prone to emergent design failures.

CRediT authorship contribution statement

Joseph Van Houten: Writing – original draft, Visualization, Software, Methodology, Investigation, Funding acquisition, Data curation, Conceptualization. **Austin A. Kana:** Writing – review & editing, Supervision, Resources, Conceptualization. **David Singer:** Writing – review & editing, Resources, Conceptualization. **Matthew Collette:** Writing – review & editing, Supervision, Resources, Methodology, 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.

Appendix

See Table A.1.

Data availability

All code for the simulation and fragility frameworks along with data for each test case can be publicly accessed online at https://github.com/Marine-Structures-Design-Lab/DesignSpace_Fragility/releases/tag/Ocean_Eng2.

References

Andrews, D., 2018. The sophistication of early stage design for complex vessels. Int. J. Marit. Eng. Vol 160 (SE 18), http://dx.doi.org/10.3940/rina.ijme.2018.SE.472.

Bernstein, J.I., 1998. Design Methods in the Aerospace Industry: Looking for Evidence of Set-Based Practices (Master's thesis). Massachusetts Institute of Technology, Cambridge, MA, USA.

Braha, D., Reich, Y., 2003. Topological structures for modeling engineering design processes. Res. Eng. Des. 14 (4), 185–199. http://dx.doi.org/10.1007/s00163-003-0035-3.

Burrow, J., Doerry, N., Earnesty, M., Was, J., Myers, J., Banko, J., McConnell, J., Pepper, J., Tafolla, T., 2014. Concept exploration of the amphibious combat vehicle. In: SNAME Maritime Convention. SNAME, Houston, Texas, USA, pp. 1–16. http://dx.doi.org/10.5957/SMC-2014-T19.

Doerry, N., 2012. A vision for ship design and analysis tools. SNAME (mt) Mar. Technol. 8–9.

Doerry, N., 2015. Measuring diversity in set-based design. In: ASNE Day. Arlington, VA, USA, pp. 1–14.

Dong, A., 2017. Functional lock-in and the problem of design transformation. Res. Eng. Des. 28 (2), 203–221. http://dx.doi.org/10.1007/s00163-016-0234-3.

Garner, M., Doerry, N., MacKenna, A., Pearce, F., Bassler, C., Hannapel, S., McCauley, P., 2015. Concept exploration methods for the small surface combatant. In: World Maritime Technology Conference. SNAME, Providence, Rhode Island, USA, D021S004R005. http://dx.doi.org/10.5957/WMTC-2015-036.

Gembarski, P.C., Plappert, S., Lachmayer, R., 2021. Making design decisions under uncertainties: probabilistic reasoning and robust product design. J. Intell. Inf. Syst. 57 (3), 563–581. http://dx.doi.org/10.1007/s10844-021-00665-6.

Goodrum, C.J., 2020. Conceptually Robust Knowledge Generation in Early Stage Complex Design (Ph.D. thesis). University of Michigan, Ann Arbor, MI, USA.

Gumina, J.M., 2019. A Set-Based Approach to Systems Design (Ph.D. thesis). Naval Postgraduate School, Monterey, CA, USA.

Hatchuel, A., Weil, B., 2009. C-k design theory: an advanced formulation. Res. Eng. Des. 19 (4), 181–192. http://dx.doi.org/10.1007/s00163-008-0043-4.

Kana, A.A., 2017. Forecasting design and decision paths in ship design using the ship-centric Markov decision process model. Ocean Eng. 137, 328–337. http://dx.doi.org/10.1016/j.oceaneng.2017.04.012.

McKenney, T., Singer, D., 2014. Set-Based: A concurrent engineering approach with particular application to complex marine products. SNAME (mt) Mar. Technol. 51–55.

Mebane, W.L., Carlson, C.M., Dowd, C., Singer, D.J., Buckley, M.E., 2011. Set-based design and the ship to shore connector. Nav. Eng. J. 123 (3), 79–92.

Mihm, J., Loch, C.H., 2006. Spiraling out of control: Problem-solving dynamics in complex distributed engineering projects. In: Braha, D., Minai, A.A., Bar-Yam, Y. (Eds.), Complex Engineered Systems. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 141–157. http://dx.doi.org/10.1007/3-540-32834-3_7, Series Title: Understanding Complex Systems.

Niese, N.D., Kana, A.A., Singer, D.J., 2015. Ship design evaluation subject to carbon emission policymaking using a Markov decision process framework. Ocean Eng. 106, 371–385. http://dx.doi.org/10.1016/j.oceaneng.2015.06.042.

Page, S.E., 2006. Path dependence. Q. J. Political Sci. 1 (1), 87–115. http://dx.doi.org/10.1561/100.00000006.

Parraguez, P., Eppinger, S.D., Maier, A.M., 2015. Information flow through stages of complex engineering design projects: A dynamic network analysis approach. IEEE Trans. Eng. Manage. 62 (4), 604–617. http://dx.doi.org/10.1109/TEM.2015.2469680.

Rao, M., Chen, Y., Vemuri, B., Wang, F., 2004. Cumulative residual entropy: A new measure of information. IEEE Trans. Inform. Theory 50 (6), 1220–1228. http://dx.doi.org/10.1109/TIT.2004.828057, URL: http://ieeexplore.ieee.org/document/1302300/.

Shannon, C.E., 1948. A mathematical theory of communication. Bell Syst. Tech. J. 27 (3), 379–423.

Shields, C.P.F., 2017. Investigating Emergent Design Failures Using a Knowledge-Action-Decision Framework (Ph.D. thesis). University of Michigan, Ann Arbor, MI, USA.

Shields, C.P.F., Singer, D.J., 2017. Naval design, knowledge-based complexity, and emergent design failures. Nav. Eng. J. 129 (4), 75–86.

Singer, D.J., Doerry, N., Buckley, M.E., 2009. What is set-based design? Nav. Eng. J. 121 (4), 31–43. http://dx.doi.org/10.1111/j.1559-3584.2009.00226.x.

Smith, R.P., Eppinger, S.D., 1997. Identifying controlling features of engineering design iteration. Manage. Sci. 43 (3), 276–293.

Sypniewski, M.J., 2019. A Novel Analysis Framework for Evaluating Predisposition of Design Solutions through the Creation of Hereditary-Amelioration Networks Derived from the Dynamics within an Evolutionary Optimizer (Ph.D. thesis). University of Michigan, Ann Arbor, MI, USA.

Toche, B., Pellerin, R., Fortin, C., 2020. Set-based design: a review and new directions. Des. Sci. 6, 1–41. http://dx.doi.org/10.1017/dsj.2020.16.

Van Houten, J., Kana, A., Singer, D., Collette, M., 2024. Early risk quantification strategy for design space reduction decisions in set-based design. Int. Mar. Des. Conf. http://dx.doi.org/10.59490/imdc.2024.852, URL: https://proceedings.open.tudelft.nl/imdc24/article/view/852.

Van Houten, J., Singer, D., Collette, M., 2022. Balancing designer influence with rework for design paths of a simple polynomial model. In: Practical Design of Ships and Other Floating Structures. Vol. 1, Dubrovnik, Croatia, pp. 508–527.

Ward, A., Liker, J.K., Cristiano, J.J., Sobek, II, D.K., 1995. The second Toyota paradox: how delaying decisions can make better cars faster. Sloan Manage. Rev. 36 (3), 43–61. http://dx.doi.org/10.1016/0024-6301(95)94310-U.

Wynn, D.C., Eckert, C.M., 2017. Perspectives on iteration in design and development. Res. Eng. Des. 28 (2), 153–184. http://dx.doi.org/10.1007/s00163-016-0226-3.

Zhou, Y.-D., Fang, K.-T., Ning, J.-H., 2013. Mixture discrepancy for quasi-random point sets. J. Complexity 29 (3–4), 283–301. http://dx.doi.org/10.1016/j.jco.2012.11.006, URL: https://linkinghub.elsevier.com/retrieve/pii/S0885064X12001057.