Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 153 (2019) 27–36

Procedia Computer Science

www.elsevier.com/locate/procedia

17th Annual Conference on Systems Engineering Research (CSER)

Artificial intelligence analytics with Multi-Attribute Tradespace Exploration and Set-Based Design

Matthew E. Fitzgerald[a] and Adam M. Ross[a]

[a] The Perduco Group, 3610 Pentagon Blvd Suite 110, Beavercreek, OH 45431, USA

Abstract

Data-driven design approaches such as Multi-Attribute Tradespace Exploration and Set-Based Design are increasing in popularity due to their ability to capture broader decision spaces than traditional point-based design. These methods share many of the same features and have complementary goals. Artificial intelligence offers a way to process the large amounts of data created by these methods in a fast and objective manner, supporting the insights of subject matter experts. This paper discusses the intersection of these three research areas and demonstrates an approach for combining these techniques to rapidly identify the most value-driving decisions available to the design team.

Keywords: multi-attribute; tradespace exploration; set-based design; artificial intelligence; acquisition; analytics; clustering

1. Introduction

As computational power has increased over time, system architects and designers have sought to further improve their analytics capability by collecting/generating and leveraging ever-larger amounts of data. Where "design" was once a task conducted by siloed domain experts, it is now characterized by increasing collaboration and joint analysis. The design process remains rooted in the knowledge and expertise of its practitioners, but the rising wealth of data available for analysis provides a new avenue for the application of artificial intelligence (AI) and the possibility of capturing emergent insights that may run counter to expectations or prevailing wisdom. Multi-Attribute Tradespace Exploration (MATE) and Set-Based Design (SBD) are well suited for establishing a foundational source of data that can be leveraged in this way. This paper provides an introduction to these two analysis frameworks/techniques, discusses their relationship to one another specifically with their ability to support the same analysis goals, and demonstrates the application of basic AI to a ground vehicle design problem formulated with MATE in a way that can augment the definition of "sets" in SBD.

10.1016/j.procs.2019.05.052

2. Multi-Attribute Tradespace Exploration

MATE is an analysis framework within the larger tradespace exploration paradigm [1], which emphasizes the evaluation of many different alternative solutions – including some that are expected to perform poorly – in contrast to traditional point-based design, which concentrates on a more detailed assessment of handful of viable designs. The "multi-attribute" component refers to MATE's emphasis on value modeling: the capturing of stakeholder preferences on many performance attributes and their aggregation (to the appropriate degree) in order to facilitate smoother exploration of the tradespace and increase the ease with which "good" choices can be found. Generally, MATE is applied with multi-attribute utility (MAU) value models, but any value modeling approach can be used, such as the analytic hierarchy process (AHP) or cost-benefit analysis.

MATE is cast as a three-phase framework, with multiple tradespace-related processes within each. The three phases are Define, Generate, and Explore. The Define phase is devoted to problem formulation. As a model-centric, data-driven approach, it is paramount that the scoping and purpose of the application is given due consideration in order to both ensure that all critical aspects of the decision are modeled and able to be incorporated into the tradespace, as well as to minimize the potential for wasted effort spent modeling extraneous or superfluous aspects. MATE encourages that these framing tasks are performed in conjunction with stakeholders, in accordance with the principles of value focused thinking and value-driven design: making sure that the pertinent questions are answered and the stakeholders will ultimately be satisfied by the decision. The Generate phase involves the creation of evaluative and value models that match the problem formulation and are able to calculate the performance/value criteria for all of the alternatives and scenarios of interest. Finally, the resulting tradespace data is Explored in the last phase, using a variety of tradespace visualization and analysis techniques.

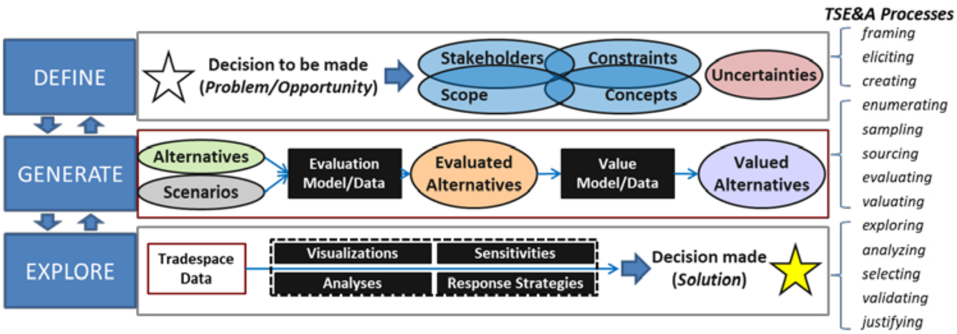

Fig 1. Overview of the MATE framework

Following the MATE framework from start to finish results in a large quantity of data on many design alternatives, allowing analysts to discover patterns in the tradespace that are difficult or impossible to see when viewing only a few point designs. Patterns and other insights are often obscured when comparing only "good" choices. By learning more about the underlying relationships between controllable design parameters and resulting system performance, analysts who use MATE are better able to build intuition for how the individual choices they make affect the system as a whole. This can lead both to the direct identification of high-value designs as well as an improved ability to creatively expand the decision space, leveraging the insights and intuition gained from the exploration to propose solutions that may have been excluded from the initial scope. Importantly, MATE allows the comparison of alternatives that are vastly different in concept on the same value-oriented dimensions, supporting this spiral of broad decision making and preventing "lock-in" on a single concept.

3. Set-Based Design

SBD is a design paradigm that, like tradespace exploration, is primarily positioned in contrast to traditional point-based design [2, 3]. Its particular emphasis is on delaying the selection of a specific design and instead supporting more abstract analysis on sets of alternatives, making it well-suited for conceptual design by leaving the commitment to specific design decisions for later detailed design efforts. In order to enable this type of thinking, SBD explicitly divides the designer-controlled variables/choices into two categories: design set drivers and design set modifiers. This conceptual partitioning highlights the decisions that define the platform (drivers) as separate from those that are more fungible details (modifiers). Sets are defined by fixed drivers, with allowed variation in the modifiers. The drivers/modifiers framework echoes the more general concepts of **architecture** (a set with fixed drivers but variable modifiers) and **design** (an instance of a set with fixed modifiers). Typically, the identification of design set drivers is performed by domain experts who possess knowledge of which variables will both drive the most substantial tradeoffs in the decision space and be the most difficult to change during detailed design.

After defining the drivers and modifiers, a typical SBD process will involve the separate analysis of different domain teams, each responsible for advocating for their own needs. In a classic point-design study this approach is risky, as the individual domain teams will usually propose vastly different designs with no clear path to reconciling those differences. By focusing on the sets, SBD forces the domain teams to leave open a wider range of possibilities, increasing the likelihood that there is overlap between them. If the sets proposed by the domain teams do not overlap, the gradual introduction of requirements on the set drivers forces the teams to iteratively adjust their proposed set until they do (as shown in Fig 2). This gradual reduction in scope also potentially enables an increase in model fidelity as designers zero-in on a region of interest. The expected output of an SBD conceptual study is the selection of one set with which to move forward into detailed design. A specific design within that set may be used as a baseline for the following analysis, but with the understanding that the chosen design set modifiers are subject to further change.

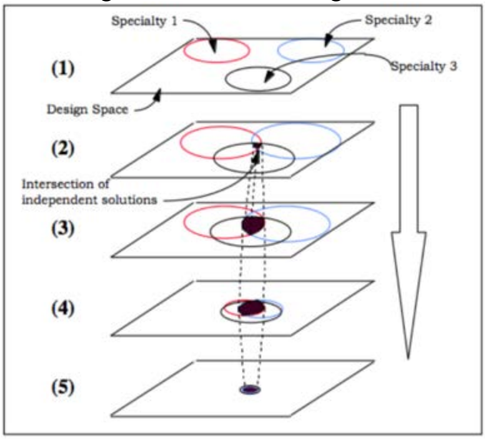

Fig 2. Set-Based Design process [2]

A key challenge for SBD is how to evaluate and compare the sets. Because sets are by definition a collection of similar designs with a range of performance, there is not a clear one-to-one assignment of "goodness" to a set which contains more available tradeoffs within itself (assuming that the level of fidelity in the model is detailed enough to incorporate the impact of set modifiers in addition to the set drivers). Generally, this will involve taking sample designs from within the set, evaluating them, and assigning the set some sort of score based on their performance. The taking of sample designs, when executed in large quantities, draws a very obvious parallel with tradespace exploration and it should be noted that, as with many popular developing techniques incorporating the work of multiple authors, the terminology for SBD has some differences between sources that occasionally overlap with similar fields. For example, Small et al. [4] and Parnell [5] use the phrase "tradespace exploration" in conjunction with SBD, likening SBD to a tradespace exploration of many alternatives (versus a "tradespace exploration of traditional point-based designs," which is using "tradespace exploration" to mean something closer to classic tradeoff analysis than the tradespace exploration paradigm that inspired MATE).

4. MATE and SBD alignment

MATE and SBD are fundamentally aligned in terms of their objectives, with the same analysis goals and similar approaches for achieving them. Consider this quote on the desired outcomes of SBD, from Singer et al. [6], emphasis ours:

> First, one would expect to have identified a **manageable set of design parameters that have been determined to be**

principal factors in achieving maximum design value. Next, one would expect to have determined which of the set is more important than the others. One would expect to have **identified which design attributes and measures are most important in differentiating among the most promising design combinations**. One would also expect to be able [to] **comparatively evaluate the most promising designs** in an analysis framework that capitalizes on the current best knowledge of design parameters and system attributes to assess total value. One would also expect to be able to **examine the impact of changes in attribute preferences** on the best design recommendation.

The bolded items are also common stated objectives in most MATE applications, particularly the attention paid to stakeholder preferences which is indicative of a value-focused methodology. Because of their similarities, the insight gathered from MATE and SBD could easily flow in *either* direction depending on the perspective of a given study: MATE could be used as a computational/evaluative framework for an SBD study or SBD could be a particular perspective applied to the analysis of a MATE study, focusing on sets of similar designs rather than individual alternatives. And the core idea of viewing many alternative designs as members of a smaller number of cohesive alternative sets has manifested as "bubble" tradespaces in the work of both fields, as in Fig 3 – a take on the classic scatterplot that emphasizes regions of the tradespace with similar alternatives. For SBD the bubbles correspond to the sets and for MATE they correspond to shared design variables that are considered important: "drivers" of the architecture in the same sense that SBD uses to define the sets up front.

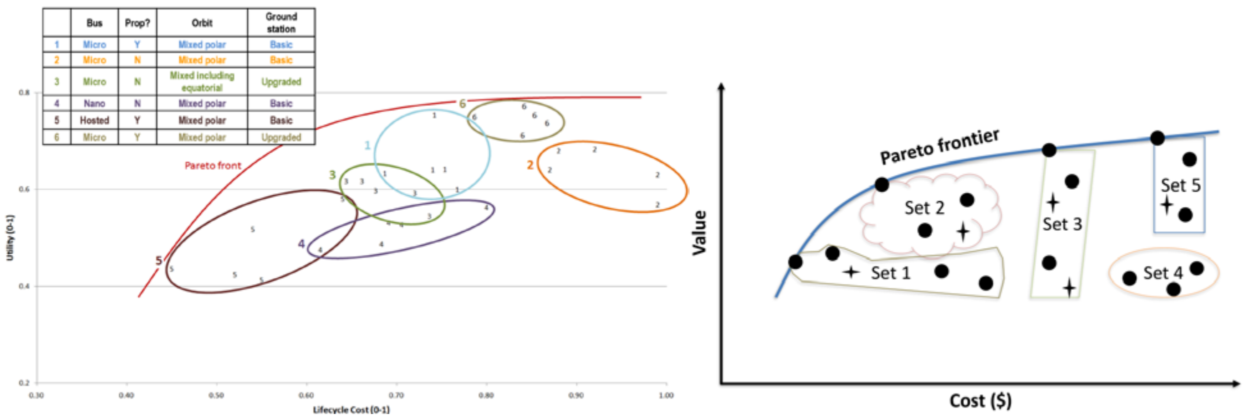

Fig 3. Bubble tradespaces outlining the regions of different sets in both MATE research [7] and SBD research [5]

5. Leveraging artificial intelligence with MATE to augment SBD

The data-driven approaches of both MATE and SBD demand similar required levels of scoping and modeling effort (an increase over point-design methods) but in turn enable the wealth of insight that is available to analysts who use them to structure their problems. However, most applications of these techniques have continued to rely heavily on human expertise to perform the analysis and exploration of the tradespace. In the case of MATE, this has been implicitly inherited from the expertise and workflows familiar to its practitioners. For SBD, the individual exploration of domain teams is considered a key part of the process. Though great success has been met with both approaches, the application of AI to these large datasets has the potential to powerfully supplement human analysis by revealing unexpected or counterintuitive tradeoffs.

Before proceeding further, it may be helpful to clarify what is meant by *artificial intelligence* in this context. The concept of AI – the ability for machines to imitate and/or replicate behaviors of humans – has been an umbrella for a wide range of mathematical and technical research since at least the 1950s [8]. The boundary of the field has become blurry as research has progressed over that time, with humans becoming more and more comfortable interacting with (and relying upon) machines displaying some level of intelligence, which has made for a "moving goalpost" to the question "what is AI?". At this time, the colloquial meaning of AI is essentially synonymous with *machine learning*: the idea that a machine can teach itself without the direct oversight/programming of its human creator. For the level of discussion in this paper, we will divide machine learning into three rough areas of research, with the caveat that this is not an exhaustive list and that many applications must rely on the blending of these techniques:

- **Supervised Learning** – The training of a machine with matched pairs of inputs and outputs, with the end goal of being able to accurately predict outputs from a set of inputs
- **Reinforcement Learning** – The training of a machine by allowing it to make decisions, assess the "goodness" of the result, and iteratively refine itself to make decisions that maximize that "goodness"
- **Unsupervised Learning** – The training of a machine to recognize patterns in otherwise unstructured input data, typically to help users better understand the data and find explanatory relationships

There is currently considerable research on difficult, modern problems within all of these categories, applying innovative mathematical techniques such as Bayesian belief networks and neural networks to domains such as image recognition, game strategy, and fraud detection. These applications generally have massive search spaces, necessitating the use of high-performance machine learning techniques to exceed human performance [9]. However, even relatively simple applications of these types of AI tasks have the potential to significantly augment human-alone performance. Indeed, each of these three categories of machine learning has corresponding "simpler" tasks that can be tackled with mathematical techniques that most practicing analysts will be familiar with. For example, *regression* is a form of supervised learning, one in which we teach a machine to execute the same tasks that a student might perform when calculating a linear regression with pencil and paper in Algebra II (with some additional complexity to capture more complex functions). Similarly, the *optimization* of a black-box model via any algorithm that is allowed to iteratively sample new "guesses" – particularly heuristic optimization techniques such as genetic algorithms – is a rudimentary form of reinforcement learning. These basic techniques are used to great effect in many engineering applications, whether or not they are actively considered to be "AI". The ability to apply these analysis concepts to a tradespace-oriented methodology demonstrates the potential for a future in which more advanced mathematical techniques (e.g. neural networks) can also be adopted, with the potential to both speed up and improve the quality of the resulting AI insights on these problems.

Particularly relevant to the discussion of MATE and SBD is the application of unsupervised learning to the definition of the sets. The sets of SBD are defined *a priori* by the expert-driven categorization of design variables as drivers or modifiers. The main purpose of this choice in problem formulation is to leverage human expertise to structure the way in which the full space of alternatives will be trimmed down across iterations of the process. Similar approaches have been taken in MATE studies, by creating bubble tradespaces for different architectures and discarding those architectures/bubbles that do not ever approach the Pareto front of the tradespace. One advantage that MATE has in this regard is that the delineation of the sets can optionally occur *ex post facto* rather than *a priori*, allowing the analyst to determine which design variables are <u>actually</u> drivers by exploring the data rather than relying solely on expertise. However, the central task of both these approaches – the division of the full tradespace into sets of somehow similar alternatives – is a *clustering* task that can fall under the AI purview of unsupervised learning. What if an AI was able to define these sets better than a human expert? Or, if not better, at least present some compelling insight or a different way of looking at the problem than originally expected? The following section will present an example of data-driven clustering to a ground vehicle design problem, allowing for sets to be clustered fully automatically and have the results presented to the analyst for rapid assessment.

6. Ground vehicle design example

This section will be a simple demonstration of the potential for AI to leverage a MATE dataset to recommend meaningful definitions of sets that could be used to supplement SME judgement when framing an SBD effort. The case in question is a notional ground vehicle, based on a model created for a scalable-aggregation MATE study [10]. The vehicle alternatives are defined by 9 design variables, shown in Table 1 with their valid ranges.

A dataset of 6480 alternatives was sampled and evaluated; for each combination of the full-factorial sample of discrete design variables, 30 random samples of the continuous variables (wheel base, engine power, and fuel tank size) were taken. This is a reasonable analogue to how SBD sets are often sampled, with some combination of "forced" variation to populate multiple sets, but also with random sampling to fill out the space. However, we will not be defining the sets *a priori*, and instead apply a clustering technique to automatically identify the most impactful decisions: those that most strongly differentiate the sets in terms of the benefit and cost that they provide. The evaluative model calculates the unit acquisition cost for a vehicle as well as a set of performance attributes (including weight, range, acceleration, weapon resistance, and others) that are used to calculate benefit via a multi-attribute utility (MAU) function. For the purposes of this clustering task, the details of these models are irrelevant and are thus omitted; however, it is important to note that the following analysis could be performed on any metrics of interest and is not specific to this dataset nor the use of a utility model.

Table 1. Ground vehicle design variables and sampling ranges

Design Variable	Valid enumerated range
Wheel base	8 – 14 ft
Engine power	200 – 500 hp
Number of powered axles	{1, 2}
Fuel tank size	4 – 10 ft^3
Tire type	Street, weather, bulletproof
Suspension type	Spring, air
Body type	Open, closed, armored
Underbody	Flat, V-shape
Fire suppression	None, water, foam

To quantify the value impact of choosing between sets, we will consider *convex hulls* of subsets of the sampled alternatives. This is similar to the "bubble" tradespace in that each hull encloses the set, and can provide a low-clutter substitute visualization for the typically crowded scatterplot while still allowing the analyst to identify the benefit/cost tradeoff patterns that result from varying design variable choices (a task traditionally performed by coloring each point in the scatterplot according to the corresponding design variable level). Fig 4 shows an example of convex hulls calculated and drawn around the three levels of the fire suppression design variable, fully encircling the black (0 = no suppression), brown (1 = water-based), and tan (2 = foam-based) sets.

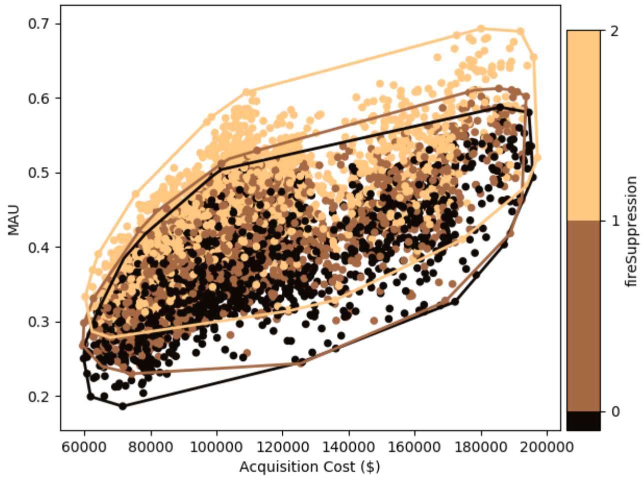

Fig 4. Ground vehicle tradespace colored by Fire Suppression system, with convex hulls defining three sets

A group of hulls defined this way is a potential definition of sets. To assess them, we will define their *differentiation* of the space as a function of how they overlap, containing members of other sets within their bounds:

$$differentiation = 1 - \frac{avgMembership - 1}{\#\, hulls - 1}$$

In this equation, *avgMembership* is the average number of convex hulls a randomly selected alternative in the tradespace is contained by. The resulting function ranges from 0 to 1 and can be represented as a percentage of full differentiation. Differentiation is 0 when *avgMembership* is equal to the number of hulls: implying that each point is contained in every hull and thus they perfectly overlap. Differentiation is 1 (or 100%) when *avgMembership* is 1: implying that the hulls are completely disjoint and have no overlap. Proposed set definitions with high differentiation outline highly important decisions, because they necessarily must drive significant trades in the value space. Note that this measure of overlap is based on the alternatives themselves, which is mathematically rigorous even when one or more dimensions is a non-ratio number type (including MAU in this example), in contrast to methods that rely in part on the area of overlap or distance between points (e.g. force-directed graphs).

AI can utilize this metric to quantify the salience of different clusters, scanning the space of potential clusterings in order to identify the highest-differentiation combinations, gaining insight into what drives value for the system and potentially leveraging that insight into the definition of sets for SBD. For this example, we will solve the problem using a brute force algorithm, testing various binnings for each design variable in search of the most-differentiating combination. The discrete variables for this case are all ordinal (having an implied greater/less than relationship), so they are tested with bin edges at each possible level. For example, a three-level variable such as body type would test the three-set clustering of {open, closed, armored} and also the two-set clusterings of {open or closed, armored} and {open, closed or armored}. The continuous variables were tested with all possible clusterings of up to four sets, with bin edges allowed at the 10% quantiles of the dataset (10%, 20%, 30%, etc.). Future applications of this technique could use advanced AI algorithms to more intelligently search for the binnings that maximize differentiation without wasting computational effort on likely-poor clusterings, and increase the fineness with which continuous variables may be partitioned. Additionally, the technique is further generalizable to combinations of design variables, as long as each alternative can be assigned to a single cluster (e.g. vehicles with small engines vs. vehicles with large engines and large cargo vs. vehicles with large engines and small cargo): this quickly leads to a combinatory "explosion" of the search space and is thus infeasible for brute force, but is achievable for more sophisticated search methods.

Table 2 shows the results of this AI search, displaying the set clustering with highest possible differentiation for each design variable, sorted by differentiation. For this tradespace, it appears that body type – the shape and armor applied to the vehicle – is the primary driver of value. Interestingly, the most impactful decision here is not between all three body types (the three-set solution), but rather the decision solely between the armored variant and the two not-armored variants (a two-set solution). An analyst interested in comparing these possible set definitions can bring up the results for the suboptimal clustering as well, as shown in Fig 5 (this time rendered without the individual scatterplot points to reduce clutter): the three-set solution, when plotted, shows a relatively large overlap between the open (0) and closed (1) sets that reduces the differentiation of the clustering from 85% to 65%.

Table 2. Optimal clustering and differentiation scores for each design variable (base context)

Design Variable	Best Clustering	# Sets	Differentiation (%)	Rank
Body type	Open/Closed, Armored	2	84.7	1
Wheel base	8 – 8.9 ft, 8.9 – 14 ft	2	34.4	2
Underbody	Flat, V-shape	2	25.3	3
Fire suppression	None/Water, Foam	2	11.2	4
Engine power	200 – 471 hp, 471 – 500 hp	2	9.8	5
Fuel tank size	4 – 4.7 ft^3, 4.7 – 10 ft^3	2	6.9	6
Suspension type	Air, Spring	2	2.6	7
Number of powered axles	1, 2	2	1.7	8
Tire type	Street, Weather, Bulletproof	3	1.5	9

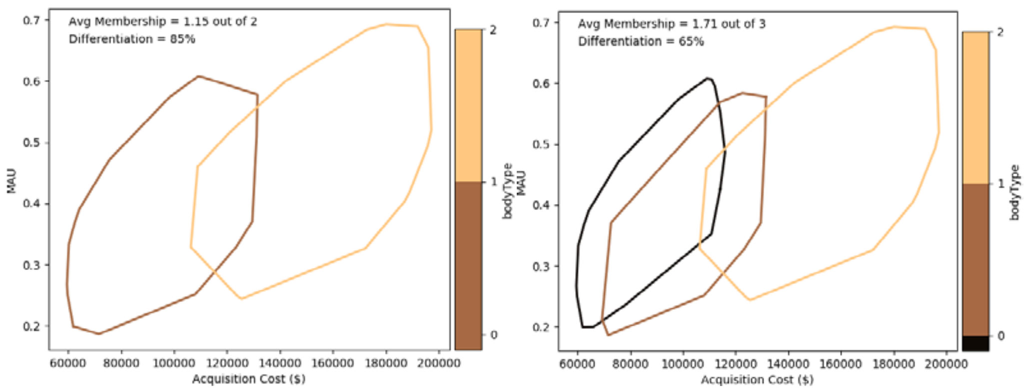

Fig 5. Optimal body type clustering (left) vs sub-optimal clustering (right) with levels 0 and 1 separated

It is important to remember that the power of this technique comes from the automation of allowing AI to recommend the optimal clustering configurations. Rather than requiring an analyst to test various partitions or rely solely on prior knowledge or expertise, the computer is able to automatically find the best set configuration for each variable and can immediately direct analyst attention to the most impactful decisions available to the design team: in this case the decision to add armor would drive a substantial increase in cost, with possible utility gains once outside the cost range of the non-armored designs. The clustering can be used to objectively define the design set drivers for an SBD process for maximal tradespace coverage and minimal overlap and rework of the different teams, each of which would take charge of a single set.

Other insights that can be gleaned from the data returned by the clustering analysis in Table 2 include things such as the relative value-driving impact of different design variables, ranging from very important (body type) to somewhat important (wheel base, underbody, fire suppression), to not important. We can also infer the areas of strong vs. weak impact of continuous variables: e.g. trading wheel base length drives more differences in value when wheel base is small than large, hence the clustering that isolates the smallest alternatives (8 to 8.9 feet) from the remainder of the tradespace

This analysis can also be used to explore decision sensitivity to uncertainty by considering changes in differentiation caused by changes to the parameterization of the evaluative model, such as would be done for a MATE study using Epoch-Era Analysis [11] to capture potential variability in the operational context of the system. This clustering analysis was repeated on the same set of alternatives but evaluated in a different epoch corresponding to bad weather, represented by changes in the underlying model parameters and therefore the shape and structure of the tradespace. These results are shown in Table 3, preserving the order of design variables from the first analysis and graying out clusters that are the same (or effectively the same). Prominent differences between contexts are shown in Fig 6:

- The new context is more challenging than the original context, resulting in fewer than 10% of the valid alternatives in the tradespace and therefore generally higher differentiation scores (less density = less overlap between sets).
- Body type increases to 100% differentiation, showing no overlap in this context.
- Tire type increases substantially in importance: despite being the least differentiating design variable in the base context we can see that it is now the fourth highest (46%). While the three types are essentially equivalent on dry roads, all-weather tires provide substantial additional utility in bad weather.

Table 3. Optimal clustering and differentiation scores for each design variable (bad weather context)

Design Variable	Best Clustering	# Sets	Differentiation (%)	Rank
Body type	Open/Closed, Armored	2	100	T1
Wheel base	8 – 8.8 ft, 8.8 – 14 ft	2	39.9	5
Underbody	Flat, V-shape	2	53.5	3
Fire suppression	None/Water, Foam	2	34.8	7
Engine power	200 – 218 hp, 218 – 500 hp	2	36.3	6
Fuel tank size	4 – 7, 7 – 7.6, 7.6 – 8, 8 – 10 ft^3	4	29.5	8
Suspension type	Air, Spring	2	8.2	9
Number of powered axles	2	1	Undefined	T1
Tire type	Street, Weather, Bulletproof	3	46.0	4

Another striking insight is that the number of powered axles has undefined differentiation, because 2-axle designs are the ONLY valid solution for this context (and thus only 1 set defined by this variable can exist). This type of edge case is still easily identifiable as a high-importance decision, which is interesting since, like tire type, this design variable was extremely unimportant in the original context. Other differences between contexts include engine power, a continuous variable, changing from being most differentiating at *high* power (set border at 471 hp) to most differentiating at *low* power (218 hp), implying that bad weather increases the downside risk of low-power vehicles. Fuel tank size is also now maximally differentiated with four sets instead of two, however it remains of relatively low importance.

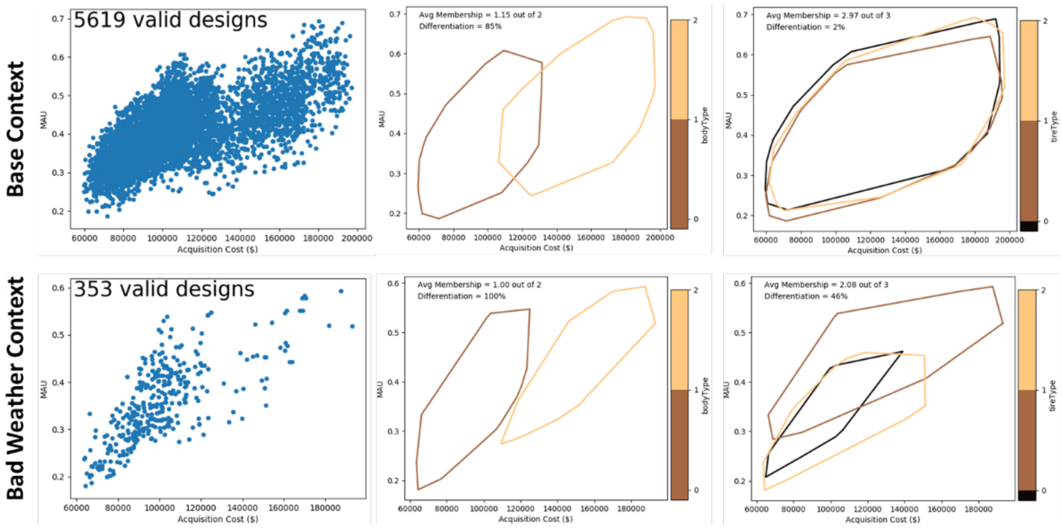

Fig 6. Comparison of scatterplot (L), body type sets (M), and tire type sets (R) for the two contexts

7. Conclusion

The use of AI to perform automated analysis of the large datasets generated during the application of early phase design techniques such as MATE and SBD has the potential to significantly change the way that engineers and analysts design systems. While brute force methods are able to solve the relatively simple problems shown in this paper – such as the partitioning of design variables into sets, and the ranking of their value-driving impact – their more important role is that of a proof-of-concept for the application of advanced AI methods such as deep learning, which can operate on the same types of unsupervised learning problems (or more specifically, clustering or pattern-finding). AI offers an opportunity to dramatically expand the scope and speed with which we can extract insights from complex problems: clustering multiple variables at once, intelligently searching to avoid time wasted on predictably overlapping clusters, and perhaps even combining the insights of multiple metrics (e.g. differentiation, number of valid designs in each cluster to avoid degenerate solutions) or across multiple operational contexts into a more holistic ranking of tradeoff importance. These objective insights offer a valuable perspective on the design problem that can supplement the knowledge and expertise of SMEs – whether providing additional support for their initial expectations or highlighting counterintuitive and unanticipated tradeoffs that impact system value. Future research will seek to expand on these concepts and incorporate new AI research with larger datasets, tackling "big data" problems for which manual analysis and exploration is insufficient for the identification of these patterns.

References

[1] Ross, A.M. and D.E. Hastings. (2005) "The Tradespace Exploration Paradigm," INCOSE International Symposium 2005, Rochester, NY.

[2] Bernstein, J.I. (1998) "Design Methods in the Aerospace Industry: Looking for Evidence of Set-Based Practices," MS Thesis, MIT.

[3] Doerry, N. (2012) "Modeling and Simulation Tools for Set-Based Design," NAFEMS 2012, Herndon, VA.

[4] Small, C., G.S. Parnell, R. Buchanan, M. Cilli, E. Pohl, S. Goerger, and Z. Wade. (2018) "A UAV Case Study with Set-based Design," INCOSE Int'l Symposium 2018, Washington, DC.

[5] Parnell, G. (2018) "Affordability and Risk Analysis Linkages in Analysis of Alternatives," 86th MORS Symposium, Monterey, CA.

[6] Singer, D.J., N. Doerry, and M.E. Buckley. (2009) "What is Set-Based Design?" ASNE Naval Engineers Journal, Vol 121, No 4, pp 31-43.

[7] Rader, A.A., F.T. Newland, and A.M. Ross. (2011) "An Iterative Subsystem-Generated Approach to Populating a Satellite Constellation Tradespace," AIAA Space 2011, Long Beach, CA.

[8] Brooks, R.A. (1991), "Intelligence without representation," Artificial intelligence, 47(1-3), pg. 139-159.

[9] Brown, N. and T. Sandholm. (2017) "Superhuman AI for heads-up no-limit poker: Libratus beats top professionals," Science, 10.1126/science.aao1733.

[10] Fitzgerald, M.E. and A.M. Ross. (2018) "Tailorable Tradespace Analysis of Mobility and Survivability for Ground Vehicles," SBIR I Report.

[11] Ross, A.M. and D.H. Rhodes. (2008) "Using Natural Value-centric Time Scales for Conceptualizing System Timelines through Epoch-Era Analysis," INCOSE International Symposium 2008, Utrecht, the Netherlands.