2018 Conference on Systems Engineering Research

Foundational and precautionary considerations for value-driven tradespace exploration and analysis

Adam M. Ross[a]*, Matthew E. Fitzgerald[a], Randy K. Buchanan[b]

[a]The Perduco Group, 3610 Pentagon Blvd Suite 110, Beavercreek, OH 45431, USA
[b]Institute for Systems Engineering Research, US Army Engineer Research and Development Center, Vicksburg, MS 39180, USA

Abstract

As the use and adoption of tradespace exploration-related techniques for acquisition decision-making grows, it is important to revisit the foundations to ensure a common basis for understanding. This paper describes a synthesized set of core mechanics that make up common value-driven tradespace exploration and analysis (TSE&A) techniques, with a set of associated processes. Key strengths and misconceptions are described to foster continued dialogue about appropriate application of TSE&A and avenues for further research and development. Lastly, precautions in several areas related to practical application of TSE&A at scale are described, along with descriptions of how overcoming these challenges can advance the state of the art in TSE&A.

Keywords: tradespace exploration; value; foundations; misconceptions; TSE&A; processes;

1. Introduction

A key challenge in modern acquisition is balancing the appropriate commitment of significant resources against the uncertain development of complex systems with promised capabilities. On top of this challenge is the uncertain value of these systems in the face of uncertain and changing threats and opportunities. Once made, the consequences of an acquisition decision will echo throughout the system's lifecycle, both enabling and limiting the ultimate value that it can deliver. Resilient decisions and systems would increase the ability of the DoD to deal with both anticipated and emergent threats and opportunities. A first step is considering a broad array of possible decisions and understanding the consequences of choosing one over another. Cost-capability analysis (CCA) is one such approach, broadening the acquisition question into a generalized tradeoff between costs and benefits of alternatives, allowing for situational awareness of what is possible and what is affordable.

Value-driven tradespace exploration and analysis (TSE&A) is a generalization of CCA, encoded as a set of techniques that enable the data-driven discovery of superior and resilient systems that are more affordable and more capable. Harnessing modern computing technologies, and leveraging the state of the art in systems engineering, decision analysis, and visual analytics, formalized TSE&A is poised to become a game-changing capability for guiding and supporting complex and impactful acquisition decisions, both in the DoD and beyond. Providing a solid foundation for TSE&A is essential for ensuring not only its appropriate role in supporting data-driven decision-making, but also its effective application across domains and decision types.

* Corresponding author. *E-mail address:* adam.ross@theperducogroup.com

2. Foundations

Tradespace exploration and analysis is an approach whereby users delay making selection decisions about what alternative "best" satisfies a particular problem (or need) at hand. Instead, a large set of possible alternatives are simultaneously considered, with evidence gathered as to each of their abilities to "solve the problem." Exploration involves attempting to better characterize some aspect of this flow (e.g. what alternatives should I consider? What data best characterizes compelling evidence for each alternative? What are my most important criteria for selecting better alternatives? How do I know when I have an answer?). Analysis involves an algorithmic and usually quantitative generation and reduction of data associated with the alternatives in order to provide further evidence for choosing particular alternatives as a solution. Value-driven TSE&A explicitly uses the ends (what makes a solution good?) to orient and drive the creation of the alternatives, thereby increasingly the likelihood of finding "better" solutions. In this way "values" are judgments that are layered on top of the evaluations of the alternatives, which represents a key activity in decision-making that always occurs. Explicitly considering values in TSE&A enables their impact to be explored (e.g. how do changes in my range "requirement" impact the attractiveness of alternatives?). Core mechanics, key strengths, and misconceptions of value-driven TSE&A will now be discussed.

2.1. Core mechanics

Based on a synthesis of the literature, interviews, and the authors' research experience based on fifteen years of conducting TSE&A studies, Fig. 1 describes a general three-layer approach for conducting a TSE&A study[1,2,3,4,5,6,7]. The first layer –*define*– encompasses the identification and definition of the decision to be made (addressing either a problem or opportunity). This would entail determining the key stakeholders, project scope, constraints on possible solutions, as well as concepts for potential alternatives and the uncertainties that might impact the success of a solution. The second layer –*generate*– encompasses the data-generation layer of TSE&A, where alternatives and scenarios are specified, evaluated, and valued to provide the basis for decision making in the next layer. The third layer –*explore*– encompasses the primary data exploration and analysis-supported decision-making activities, usually involving some combination of visualizations, analyses, sensitivity explorations, and consideration of uncertainty response strategies. Each of the layers both feed-forward and feed-back, as users seek to gain insights and confidence in the results. The processes associated with each layer shown in the figure are described in Table 1.

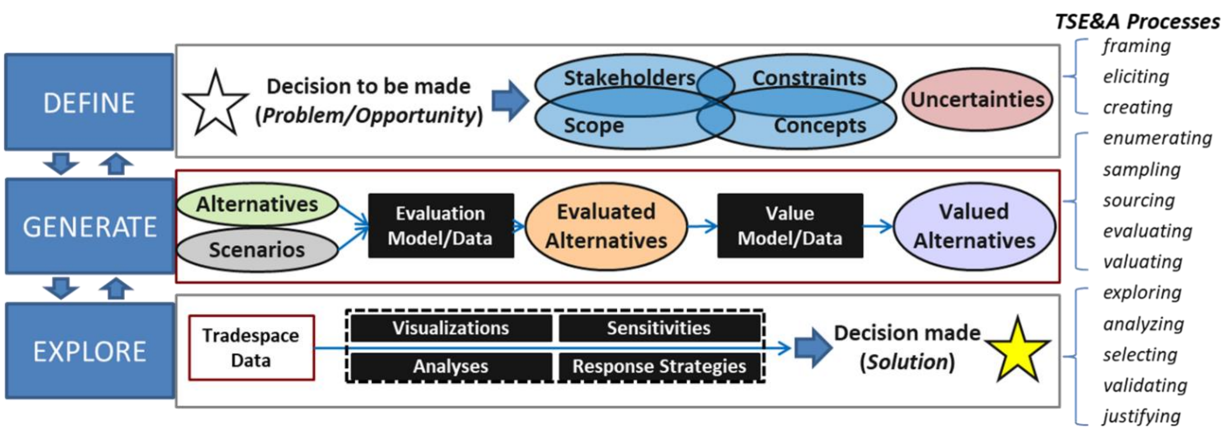

Fig. 1. The three key layers in general tradespace exploration and analysis (TSE&A) along with core concepts and associated processes

Table 1. Tradespace exploration and analysis (TSE&A) processes descriptions

Process	Description
Framing	The definition of the driving problem to be solved, including relevant stakeholders who specify how alternatives should be judged, scope extent of what to consider in the study, as well as study constraints
Eliciting	The definition of specific judgment criteria and limitations on possible solutions as specified by stakeholders, policy, and other data (e.g. MOEs, MOPs, TPMs, inherited technologies and systems, budget, schedule, etc.), as well as key uncertainties
Creating	The definition of the alternatives (e.g. design) space from which particular alternatives can be specified (e.g. spanned by a design variable set), as well as the potential scenario space (e.g. contexts and needs, both short run and long run)
Enumerating	The definition of particular alternatives from the alternatives space and particular scenarios from the scenario space
Sampling	The definition of a particular subset of alternatives from the enumerated set intended for evaluation (e.g. through design of experiments or expert rules), including definition of particular subset of scenarios from the enumerated set
Sourcing	The determination of how alternatives will be evaluated within and across scenarios, usually consisting of the identification and procurement of appropriate model(s) and/or dataset(s)
Evaluating	The evaluation of sampled alternatives via evaluative model(s)/data in terms of desired metrics (e.g. performance and resource-related metrics) often done through modeling and simulation
Valuating	The valuation of evaluated alternatives via value model(s) in terms of aggregate perceived benefits and costs as specified by stakeholders (e.g. utility and lifecycle cost)
Exploring & Analyzing	The intentional investigation of relationships and patterns between the input space and outputs space, resulting in knowledge and insights for the analyst/decision maker
Selecting	The decision on one or more alternatives as "answering the question" posed by the study, for example providing the "best" benefit at a given cost across considered use contexts
Validating	The act of confirming the selected alternative(s) actually address(es) the driving problem (i.e. satisfy stakeholder needs)
Justifying	The intentional act of compiling the evidence needed to support the solution selection recommendation(s)

2.2. Key strengths

The strengths of TSE&A as an approach for conceptual design and acquisition decision support are generally well-known, as they have been communicated extensively throughout the history of the field. Some recent examples follow. A recent paper addressing the use of TSE&A broadly identified four key positive outcomes of using the approach[8]:

- Knowledge development – learning underlying cause-and-effect relationships with value
- Refining preferences – "design by shopping" paradigm allows the stakeholder to gradually reveal their own preferences rather than requiring an exact elicitation[9]
- Choosing a concept – the central act of making a decision
- Advancing the state of the art – incorporates modern concepts such as value-focused thinking and an expanded scope far beyond traditional point-design studies

Another perspective, following a survey of TSE&A tools, identified six main functions enabled by TSE&A that support the decision-making process[7]:

- Capturing value – using value-focused methods to customize for user perception
- Multi-Disciplinary Optimization – allowing for maximization of chosen objective functions
- Statistical data analysis – identifying patterns and developing forecasts
- Visualization – displaying information with intuitive graphics
- Decision analysis – evaluating sets of alternatives against elicited preferences
- Project/Portfolio management – scoping to include resource allocation across a portfolio of assets

These strengths only scratch the surface, but most stem from delaying premature decision making, revisiting and questioning assumptions, enabling knowledge gain through developing intuition about potentially complex relationships between what stakeholders ask for and what is technically possible, fostering open dialogue around problem framing and solution approaches, and identifying and making explicit the characteristics of a good solution.

2.3. Misconceptions

TSE&A, as a growing field, has experienced substantial change over the past 10-15 years. The sophistication of the TSE&A research community (and practitioners) has increased dramatically, yet some misconceptions have persisted from the early days of TSE&A. These misconceptions are mostly perceived weaknesses of "classic" or "one-shot" TSE&A as performed by inexperienced designers and engineers, when current best practices have addressed these by various means. The following subsections will highlight a handful of these concerns and how experienced TSE&A practitioners can address them.

2.3.1. "TSE will stifle innovation"

When viewing TSE&A as a tool for design, some designers feel that a rigid modeling and analysis structure such as TSE&A has a negative effect on innovation – limiting designer creativity and expertise by forcing the solution to come from the set of enumerated and evaluated designs. From the perspective of the TSE&A community, this is a strawman argument: the exploration of TSE&A has always been first and foremost a technique to be used as a <u>decision aid</u> between potential identified solutions, not a tool for assisting or replacing the innovation/identification of new solution concepts. That said, the stages of TSE&A prior to the evaluation of alternatives still rely heavily on designer creativity through the innovation and enumeration of system concepts and their associated controllable design variables and feasible levels. If new or innovative solutions to a problem are desired, then innovation must be applied at this point in the design process – not during the exploration and analysis stage. However, lessons learned during exploration *can* be utilized to drive innovation when defining the next set of concepts, if TSE&A is iterated.

It is also possible to utilize TSE&A as a means of guiding targeted innovation. Because the exploration and analysis of TSE&A is (often) performed on simulated data, it is possible to compare not only currently actionable/feasible solutions but also <u>hypothetical</u> solutions. For example, when designing a new aircraft, currently-impossible levels of engine fuel efficiency and/or thrust could be included in the tradespace evaluations. These solutions will nominally be superior to known-feasible solutions and can be compared against them on value-oriented dimensions to capture the possible value of technological innovation. The value-added of different innovations can also be compared to each other in order to direct innovation research resources into the most promising areas of design.

2.3.2. "A tradespace covers only a fraction of possible solutions"

This argument suggests that the "best" answer from a tradespace is not truly the "best." It is primarily targeted at two aspects of the tradespace: (1) the large number of controllable design variables that must be omitted for the sake of managing computational effort, and (2) the discretization of continuous variables. The small nugget of truth in this argument (likely based on observations of inexperienced TSE&A practitioners) is that some people do ascribe too much confidence in the "bestness" of the answer that is chosen via TSE&A. However, proper TSE&A practice addresses these concerns in two important ways. First, there are techniques available for TSE&A practitioners to quickly and accurately identify the most important value-driving design variables, such as the Design Value Mapping (DVM)[10]. It is relatively uncontroversial to suggest that a small handful of decisions often determines the vast majority of value (the Pareto principle), and therefore the most impactful value tradeoffs can be captured as long as this subset of drivers can be incorporated into the tradespace (e.g. wingspan is more important to model than the head type of the screws on an aircraft). Second, most modern TSE&A is designed to be iterative – exploring and locking in "big" decisions before iterating with successively "smaller" decisions before finalizing a concept and entering detailed design. Any deviations from "bestness" can be removed in these iterations as scope tightens and detail improves, including those deviations derived from discretizing continuous variables – either by increasing the fineness of the discretization or eventually moving into a low-level optimization.

In general, the constraints of procedural generation/enumeration and evaluation of system alternatives – a finite set of design variables, limited to the domain for which the model is valid, and used to assess only tractably modeled attributes – does limit the scope of TSE&A from considering the infinite space of conceptual design. However, these limitations are not unique to TSE&A but are rather a fundamental aspect of all model-based engineering efforts, including multi-variable optimization. Such a limitation in scope is the price that must be paid to allow computers to evaluate potential system designs faster and cheaper than via human expertise or prototyping.

2.3.3. "Tradeoffs can mislead or be actively manipulated"

This argument is a challenge to the fundamental value of performing tradeoff analysis, suggesting that such analysis is unreliable either by chance or by deliberate manipulation. For many years, practitioners were taught that the "knee" of the Pareto front in the standard benefit-cost tradespace was the first place to look for the "best" solution, and to this day the "knee" has a particular perceptual attraction to many TSE&A practitioners – it appears visually to be the best "bang for your buck" solution. However, when either dimension of the benefit-cost space uses a non-ratio scale (including most non-monetized value models), the "knee" is not meaningfully superior to other Pareto efficient alternatives. Accordingly, recent TSE&A research has deemphasized the "knee" in favor of more holistic analysis of the Pareto front in order to avoid the potential for misleading analysis based on purely visual circumstances[4,5]. The penetration of this idea into practice has been slowly following.

Because the “knee” is generally identified visually, it is also subject to change depending on the chosen axis limits and scale. This does open the door for potential deliberate manipulation, to place a favored alternative on the “knee” to give it additional attention. However, such behavior should not be considered a fault of TSE&A but rather the result of unscrupulous, advocacy-based analysis. In fact, such advocacy bias is substantially more difficult to implement in the many-alternative framework of TSE&A than in either point-design studies (e.g. the common “not good enough”, “too expensive”, and “just right” three-solution study) or multi-variable optimizations (where the parameters can be tweaked until the preferred solution is returned).

2.3.4. “Utility theory measures total preference, and should not separate cost and benefit”

This final argument is a mathematical challenge to the principle of using utility theory to model subcategories of value, but is unfounded. The original von Neumann-Morgenstern utility theory[11] very specifically has only two postulates: (1) the subject can always determine which of two packages they prefer or if they are indifferent between them (complete ordering), and (2) that this evaluation can include probabilistic events (lotteries). They go on to describe the creation of a vector-basis of utility based on the weakening of the complete ordering postulate. This would be the case if the subject is unable to say which of two benefit-cost packages they prefer (e.g. low-cost, low-benefit or high-cost, high-benefit) or displays intransitive preferences in those dimensions. In such a case, separate utilities for each dimension – with complete ordering within each – can form a utility vector with no adverse effects beyond the accompanying increase in complexity. This type of incomplete ordering between benefit and cost has been observed in the field of complex system design and is justification for the conceptual positioning of TSE&A as a means of trading benefit and cost, either or both of which can be modeled with utility theory if they are comprised of multiple attributes. Incomplete ordering is also the main barrier to the reduction of value to a meaningful one-dimensional function, with the goal of layering “isovalue” curves on top of the benefit-cost tradespace. In practice, such isovalue curves can be estimated by asking stakeholders to identify alternatives to which they are indifferent.

3. Precautions

Despite the strengths of TSE&A, practitioners must continue to take appropriate precautions. Certain problem features present challenges that require careful management of both human and computing effort – failure to do so can result in faulty insights and/or inferior decision making that fails to align with stakeholder preferences. The following subsection will discuss a few of these challenges and how they can be managed with effective practice.

3.1. Visualization scalability

One challenge faced when applying TSE&A with many alternatives is the effective visualization of each alternative. In particular, as the number of alternatives increases, many visualizations that rely on one-to-one visual representation demonstrate significant *occlusion*, in which some data is covered by overlapping data. This is a common problem in TSE&A when looking at large design spaces. Consider an example with two different samplings of the design space for the same study (Set A = Design of Experiment (DOE)-sampled and Set B = Genetic Algorthm (GA)-sampled). The occlusion becomes noticeably worse when moving from the $O(10^4)$ designs of Set A to the $O(10^5)$ of Set B, as fewer points are individually identifiable. In this case, occlusion limits our ability to see how densely populated different areas of the tradespace are, because the entire tradespace appears as a solid block. Occlusion can also create misleading visualizations, commonly when looking for patterns in features such as color, as in Fig. 2.

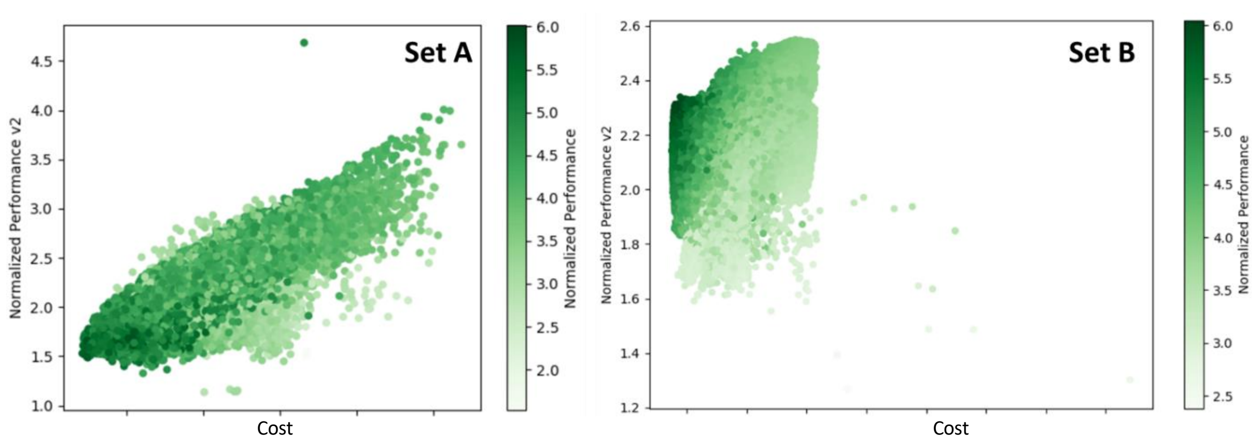

Fig. 2. Example Normalized Performance v2 vs Cost tradespace scatterplots (colored by Normalized Performance v1)

In these plots, Normalized Performance v1 (NP) is indicated by color, and Normalized Performance v2 (NP2) has taken the y-axis. First, we can see that the tradespace takes a common shape, with more NP2 benefit becoming achievable as more cost is spent. However, based on the color, we might infer very different insights from these two plots about the relationship between NP and NP2. In the Set A plot, it appears that NP and NP2 are roughly inversely proportional, with the highest NP scores (darkest colors) appearing in the lower left corner. In the Set B plot, it appears that the same designs are best in both NP and NP2, at least in the low-cost regime. However, neither of these conclusions is reliable given the large amount of occlusion: the colors that can be seen are merely the points that were plotted "on top," and may not reflect the points underneath. Importantly, this is not a problem that can be addressed simply by increasing computational power. One approach for addressing occlusion in large datasets is the use of a binned tradespace as in Fig. 3.

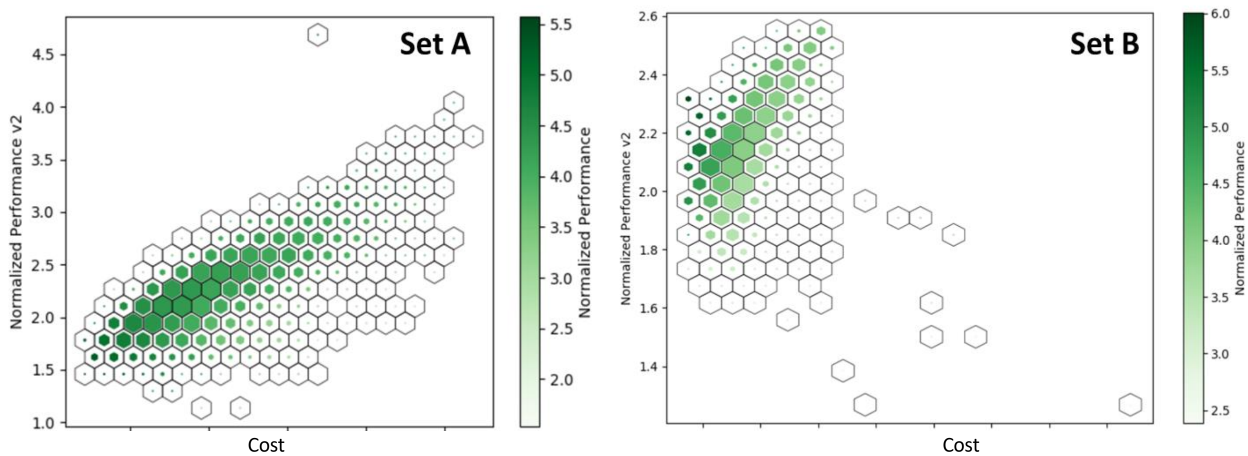

Fig. 3. Example Normalized Performance v2 vs. Cost binned tradespaces (colored by mean Normalized Performance v1)

These two plots show the same data, but presented in a hexagonal grid. Each hexagon summarizes all of the points contained in its bounds. The size of the colored area (how "filled" the hex is) corresponds to the *number* of contained points, and enables a more accurate reading of the density of points in different regions. The *color* represents a summary statistic for the color-value of all the contained points: in this case, the average NP. Now we can see with confidence that the surface-level insight we drew from the colored scatterplot of Set A was not very strong: the lower-cost designs do have a higher average NP than the high-cost designs but not by very much. On the other hand, the Set B insight appears to have been correct, with the highest-NP hexes located in the upper-left corner where NP2 is also high. What could explain the difference between the two? The upper-left portion of the Set B tradespace is outside the domain of the Set A tradespace; the GA found strictly superior designs to the DOE in that region, and they are very good in both benefit metrics. The portion of the Set B tradespace that does overlap with Set A also has roughly uniform NP.

We can also see that the Pareto front is more densely populated in Set B than Set A (larger filled hex sizes on the edge of the tradespace), as expected from an intelligently-searching GA compared to a DOE, which is likely to hit many more medium-attractiveness designs. However, the bulk of the Set B points are still located in the center of the tradespace. It is not reasonable to expect the Pareto front to be the *densest* part of the tradespace – it is, after all, the region that pushes the boundaries of design, and the random mutations occurring in a genetic algorithm are individually unlikely to be optimal – but in this case, it appears that none of the hexes on the Pareto front have more than ~10% of the designs that the 8 or so "full" hexes in the center do. This may merit further investigation to check if the GA is enumerating new alternatives effectively: for example, this pattern could be caused by an excessively high mutation rate or a selection method that insufficiently weights individual fitness.

3.2. Organizational buy-in and effort for value modeling

A significant part of TSE&A's value as an approach comes from its alignment with the principles of value-focused thinking and value-driven design. If the true value proposition of the system as defined by system stakeholders and decision makers is not captured, the insights derived from TSE&A will be strictly mechanical – how design variables impact attributes – and lack the connection to what is desired from the system. As such, maximizing the benefit of TSE&A requires participation from stakeholders. Unlike other analysis techniques that may be able to be completed with full effectiveness by engineers and analysts alone, practitioners seeking to fully realize the benefits of TSE&A (and justify the effort needed) must gather support from beyond the analyst, to include key stakeholders.

Unfortunately, capturing value propositions for complex systems can be effort intensive, particularly when considering large stakeholder networks or “big” problems such as systems-of-systems or portfolios. Because stakeholders and decision makers are often busy, they frequently balk at the prospect of eliciting and revising value statements for incorporation into a value model. As such, perhaps the most useful leverage to practitioners of TSE&A is organizational buy-in: a culture of value-focused thinking that supports the effort to make decisions that align with perceived value instead of just technical parameters alone.

3.3. Stakeholder turnover for long-lived programs

Building upon the prior challenge, even if stakeholders do participate in the value modeling aspects of TSE&A an additional challenge remains: development timelines for complex systems are frequently longer than the timescale of leadership roles (particularly in the government). Though some programs are completed with a single program manager from start to finish, rarely do the key stakeholders that those managers must satisfy remain the same. Stakeholder turnover can result in dramatic and potentially unpredictable changes in the value proposition for a system. Organizational buy-in can help here again, if it ensures that the new stakeholder is at least willing to put in the effort to define and model their own preferences (needed for TSE&A to continue). However, because some design decisions are locked in gradually, the threat of stakeholder turnover also requires an increased focus on resilience concepts such as changeability or versatility. If TSE&A is used only to optimize performance for a single scenario, the resulting system will likely be brittle to any changes in stakeholder preference. Practitioners should incorporate uncertainty frameworks and analyses (e.g. Epoch-Era Analysis, sensitivity analysis, etc.) into their TSE&A studies or they risk being unable to remain relevant and insightful.

Alternatively, this problem could be addressed by changing the way such programs are conducted: either through longer-term leadership commitments or fixed program directions and/or requirements that new leadership is not allowed to alter. Of course, the barriers to this type of organizational change are incredibly high – especially for the government – so TSE&A practitioners should concentrate on managing stakeholder uncertainty with resilience.

3.4. Incorporating preexisting modeling/simulation/analysis

Because modeling and simulation (or data-sourcing more generally) is typically the most time-consuming part of a TSE&A study, practitioners often must incorporate preexisting models into the evaluation of the tradespace. This type of model reuse can save considerable time and effort but, if not treated with care, can be performed incorrectly and lead to faulty analysis. The central task is to connect the outputs of these evaluative models with the corresponding attributes of the value models defined by the stakeholders. Often, the evaluative models were not created with value-focused thinking in mind, causing their outputs to be more technical (or less applied) than the value-oriented terms a stakeholder might request. For example, a technical model of a car might return “handling” as a measure of side-force that can be applied while making a turn without slipping, while a stakeholder might think of “handling” as a shorthand for smoothness of ride. A TSE&A practitioner must ensure that these types of mismatches are addressed via conversion models (i.e. a translation from engineering-speak to value-speak). When time permits models to be developed specifically for a given TSE&A study, such a conversion model may not be required. However, this can cause practitioners to possibly forget the importance of such a conversion when performing future studies with repurposed models.

Similarly, using preexisting models will typically involve creating a composite model of multiple different models – for example, recycling both an engine model and a weight/sizing model to evaluate a car. However, combining preexisting models can be a trap for unwary TSE&A practitioners, as they can differ in ways that prevent a straightforward combination. Some models may rely on different assumptions and may be unreconcilable. Others may have different time scales of simulation, necessitating either interpolation or forecasting (based on their dependencies in the composite model). Others may have varying degrees of uncertainty, which then requires careful statistical accounting of uncertainty as it propagates through the model. TSE&A practitioners must be aware of these threats, any of which could compromise the effectiveness of the study. This is a contributing factor to why a TSE&A practitioner should have a working understanding of modeling and simulation: analyzing the data properly requires understanding (and vetting) the source of the data and cannot be fully separated from the evaluation task[12].

3.5. Training – moving beyond tool development

The expertise necessary to run a believable and insightful TSE&A study should not be underestimated. The TSE&A practitioner needs to be familiar with evaluative modeling/simulation, value modeling, visualization, and statistics to be able to integrate their contributions to TSE&A. It is easy to think of TSE&A as a mechanical in/out task akin to solving a math problem for the “best” alternative, but this ignores the difficulty of setting up a tradespace. There are a minimum of 12 distinct tasks necessary to create and analyze a tradespace listed in Table 1, many of which require or benefit from human intuition and expertise. The previous subsection also identified some of the practical challenges TSE&A practitioners face with respect to collecting data through multiple models, such that the data remains valid for tradespace analysis. In fact, “soft factors” such as “trust” and “confidence” have been recognized as increasingly important in order to leverage model-centric techniques to support consequential decision-making[13].

As such, there is a definite risk with assigning TSE&A tasks to new, untrained analysts. Though the value of sophisticated and intuitive TSE&A tools and software is appreciated by many people[7,14], TSE&A training is equally important given the significant impact that the human-in-the-loop has on the result of a study. Organizations that do not acknowledge the importance of training staff to both rigorously assemble models/datasets for TSE&A and deeply understand/perform the exploration and analysis tasks, will struggle to realize the full potential of TSE&A.

4. Conclusion

Sitting at the nexus of advancements in computation, decision analysis, visual analytics, model-centric engineering, and even behavioral economics, TSE&A is poised to offer compelling and powerful insights to the design and acquisition communities, helping to curtail and manage the complexity of real decision-making in dynamic and uncertain high dimensional problem spaces. But the real power of TSE&A stems not from its ability to manipulate large quantities of data in a structured manner, but rather from its ability to encourage explicit consideration and communication of how we both frame and seek solutions to our problems, and how it broadens our perspective around what a “good” solution might look like over time. Fostering open dialogue in the community around foundations, best practices, challenges, and research opportunities can help move TSE&A toward becoming an essential part of any critical decision-making endeavor.

Acknowledgements

This material is based upon work supported, in whole or in part, by the U.S. Department of Defense through the Engineered Resilient Systems (ERS) Program. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense.

References

1. Ross AM and Hastings DE. The tradespace exploration paradigm. INCOSE Int’l Symp 2005. Rochester, NY, July 2005.
2. Ross AM, McManus HL, Rhodes DH, and Hastings DE. Revisting the tradespace exploration paradigm: structuring the exploration process. AIAA Space 2010. Anaheim, CA, September 2010.
3. Ricci N, Schaffner MA, Ross AM, Rhodes DH, Fitzgerald ME. Exploring stakeholder value models via interactive visualization. 12th Conf on Sys Eng Research. (CSER14). Redondo Beach, CA, March 2014.
4. Ross AM, Rhodes DH, and Fitzgerald ME. Interactive value model trading for resilient systems decisions. 13th Conf on Sys Eng Research. (CSER15). Hoboken, NJ, March 2015.
5. Ross AM, Fitzgerald, ME, and Rhodes DH. Interactive evaluative model trading for resilient systems decisions., 14th Conf on Sys Eng Research. (CSER16). Huntsville, AL, March 2016.
6. Ender T, Goerger S, McGroarty C, Spero E. Engineered resilient systems: tradespace enabled decision making. 17th Annual NDIA Systems Engineering Conference, Springfield, VA, October 2014.
7. Spero E, Avera MP, Valdez PE, Goerger SR. Tradespace exploration for the engineering of resilient systems. 12th Conf on Sys Eng Research. (CSER14). Redondo Beach, CA, March 2014.
8. Collopy PD. Tradespace exploration: promise and limits. 15th Conf on Sys Eng Research. (CSER17). Redondo Beach, CA, March 2017.
9. Balling R. Design by shopping: a new paradigm? Proc of 3rd World Cong of Struct and Multidisc Opt (WCSMO-3). Buffalo, NY, May 1999.
10. Ross AM. *Managing unarticulated value: changeability in multi-attribute tradespace exploration.* PhD diss. MIT: Cambridge; 2006.
11. Von Neumann J, Morgenstern O. *Theory of games and economic behavior*. 3rd ed. Princeton University Press: Princeton, NJ; 1953.
12. Reymondet L, Rhodes DH, and Ross AM. Considerations for model curation in model-centric systems engineering. IEEE Syscon 2016. Orlando, FL, April 2016.
13. Rhodes DH, and Ross AM. A vision for human-model interaction in model-centric systems engineering. INCOSE Int’l Symp 2016. Edinburgh, UK, July 2016.
14. Spero E, Bloebaum CL, German BJ, Pyster A, Ross AM. A research agenda for tradespace exploration and analysis of engineered resilient systems. 12th Conf on Sys Eng Research. (CSER14). Redondo Beach, CA, March 2014.