Proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023)
Edited by Mário P. Brito, Terje Aven, Piero Baraldi, Marko Čepin and Enrico Zio
 Published by Research Publishing, Singapore.
doi: 10.3850/978-981-18-8071-1_P659-cd

Emerging Technology Certification Risk Assessment with ETHICIST

Shamal Faily, Rob Ashmore
Defence Science and Technology Laboratory, Portsdown West, UK. E-mail: {sfaily,rdashmore}@dstl.gov.uk

Risk stakeholders need help attending to certification challenges associated with emerging technologies on critical systems. Assessing emerging technology risk needs to account for its relationship with our technology and standards, and the processes and tools need to be accessible to different stakeholders. In this paper, we present ETHICIST: a systematic approach for assessing and managing emerging certification technology risk. Our approach uses multi-criteria decision analysis and concept mapping to account for different attributes of certification risk. It also visualises the cascading impact on other technology, regulations, and systems. We illustrate this approach by considering the certification risk of additive manufactured wearable computing elements for a military air system.

Keywords: Certification risk, Multi-Criteria Decision Analysis, Concept Maps

1. Introduction

Risk stakeholders need help identifying, understanding, and prioritising certification challenges associated with novel and new technologies. Certification typically represents an unavoidable barrier to technology exploitation, particularly in regulated domains like military aviation. Assessing these challenges is made difficult by an intrinsic lack of context. Imagining some pre-existing technology operating on an existing platform is comparatively easy. However, uncertainty about the form of some emerging technology, a lack of clear operating context, and sparse technology and certification subject matter expertise make the challenge difficult.

Using pre-existing risk assessment approaches to evaluate and make sense of novel technology risk in isolation is hard. Some technology might contribute to or depend on some other form of novel technology; this could change our opinion of the risk of operating both. Additionally, system-level architectural approaches that enable existing technologies may no longer be appropriate for emerging technology.To overcome these difficulties to meet the initial challenges, we contend that any approach must address four characteristics.

First, technology readiness is necessary but not sufficient for capturing the different facets of emerging technology risk. Therefore, risk assessment needs to account for different attributes associated with technology and certification risk. Moreover, although the risk ratings will rely on subject matter expertise, this expertise should be both transparent and amenable to exploration from different perspectives. This should make it possible to examine the impact of changing the factors feeding into any risk assessment algorithm, and how these factors might be weighted.

Second, emerging technology is not an island. Any approach should not just consider how technology is categorised, but how it relates to other technolgies, both existing and emerging. For example, a technology's risk may not be objectively rated as a significant, but its dependency or contribution to the operation of other technology may warrant a review of its importance.

Third, any approach should consider the impact of technology concepts not just to each other, but the standards they need to be certified against, and the platforms or systems they need to run on. As difficult as mapping some emerging technology to orthogonal concepts might be, insights into relevant regulations and systems could identify areas for further investigation into the implications of some risk. Moreover, it is also useful to capture where some technology has some indeterminate impact on regulations or systems, particularly if it operates in contexts with other technology with equally indeterminate implications.

Finally, the outputs of any approach should

be maintainable and accessible. No single person has the subject matter expertise to maintain the outputs of such an approach. As such, knowledge of technology concepts and risks should be maintained orthogonally, and configuration controlled. By doing so, others can easily maintain the models as our understanding of technology, regulations, and systems evolves.

To address these challenges, this paper presents ETHICIST (**E**merging **T**ec**H**nology cert**I**fi**C**ation r**I**sk regi**ST**er): a systematic approach for assessing and managing emerging certification techology risk. In Section 2, we describe the related work upon which our approach is based, before presenting the elements of ETHICIST in Section 3. We present a worked example of applying ETHICIST in Section 4, before concluding with some implications of our work in Section 5.

2. Related work

2.1. *Emerging Technology Assessment*

Technology Readiness Levels (TRLs) (Mankins, 1995) are the most commonly used metric for assessing the maturity of some technology, where the lower the readiness level, the more uncertain the route to integration and deployment might be. Although not designed as a risk metric, many people use TRLs as such. There are, however, known to be discrepancies between how technologists and customers perceive TRL (Frerking and Beauchamp, 2016), and providing assurance for a TRL assessment in the presence of uncertainty (Mankins, 2009b).

Given its weaknesses, several replacements for and extensions to TRL have been proposed for assessing technology in context. For example, "Integration Readiness Levels" for considering different architectural views of a system (Jesus and Chagas Junior, 2022), "Human Readiness Levels" considering the readiness of technology to human operators (Salazar and Russi-Vigoya, 2021), and "Community Maturity Levels" that result from a cluster analysis of communities of system components (Goldschmid and Corns, 2021). However, integration and community maturity assume at least some basic level of system maturity, and human-readiness considers only operators directly using some technology not the broader community of stakeholders, e.g. engineers deploying it, and decision makers who indirectly use or benefit from it.

2.2. *Multi-Criteria Decision Analysis*

There have been proposals to formally integrate risk and TRL assessment (Mankins, 2009a) and, more recently, multi-attribute analysis has been suggested as a means of better incorporating uncertainty and stakeholder perspectives when making sense of technology and life cycle risk (Moni et al., 2020).

Multi-attribute analysis techniques help evaluate alternatives when different objectives are in tension (Bunn, 1984). Multi-attribute utility theory has long been used by policy stakeholders to evaluate different policy options in the presence of uncertainty, e.g. (Communities and Local Government, 2009; Fujiwara and Campbell, 2011), and do so in a manner transparent to different stakeholders. Previous work by Butler (Butler and Fischbeck, 2002) has also demonstrated how risk assessment can be framed as a multi-attribute decision problem, where risk is evaluated based on an additive value model (Edwards, 1977) following the formula :

$$v(x_1, x_2, \ldots, x_n) = \sum_{i=1,n} w_i v_i(x_i) \tag{1}$$

where $v_i(x_i)$ is a value function over levels of risk attribute x_i, and w_i is a weight applied to the value function.

2.3. *Concept Mapping*

Concept maps are sensemaking tools that connect ideas, objects, and events within a domain and, in doing so, help organise and visualise knowledge (Martin and Hanington, 2012). First proposed as a learning tool (Novak and Gowin, 1984), they help individuals make sense of and share discourse around concepts and how they relate to each other (Sutherland and Katz, 2005).

Concept mapping can be a cognitively demanding process but, with the aid of lightweight software tools for modelling and version control, task

complexity does not increase as maps become more elaborate (Faily et al., 2012). This is important for ensuring the maintainability of concept maps, given their potentially long lives.

3. Approach

To overcome the challenges highlighted in Section 1, a framework – ETHICIST (**E**merging **T**ec**H**nology cert**I**fi**C**ation r**I**sk regi**ST**er) – was devised to assess and manage emerging technology risk in context with other novel technology and cogent regulatory and system elements. The framework can be used to not only *assess* the risks to certification, but *manage* them as its outputs consistitute a risk register.

ETHICIST is applied to some collection of emerging technology concepts, and entails two parallel activities. Multi-Criteria Decision Analysis is used to carry out an independent risk assessment of each item of technology. Concept mapping is used to model the relationship between technologies and their impact on standards and systems they are deployed to. Further analysis is then carried out to colour concept map nodes based on the certification risk of the technology items, and the items contributing to them.

3.1. *Identify attributes and values*

Four attributes of each technology are proposed to assess technology certification risk: Technology Readiness & Disruption (t), Certification Readines (c), Scale of Complexity (x), Customer Demand (d). The first two attributes were proposed to assess the level of uncertainty associated with a technology's general maturity for certified operation and its potential for innovation disruption. The latter two attributes were proposed to consider uncertainty associated with its context, i.e. its scale and customer expectations. Each attribute is scored on an ordinal scale, where each ordinal attribute value is associated with a corresponding quantitative value v, where $0 \leq v \leq 1$, and the values are evenly distributed between 0 and 1. The attributes are intentionally agnostic of safety, security, and operational capability, but have direct or indirect implications on all three dimensions of system performance. For example, complex technology is more likely to have a larger cyber attack surface, and errors occurring with complex human or machine interaction are more likely to lead to faults.

Technology Readiness & Disruption considers how ready the emerging technology concept is for operational use. A concept with a low score will have been subject to only a preliminary level of scientific evaluation, whereas a concept with a high score may already be in operational use in limited or constrained circumstances. This attribute is similar, but not identical to, TRL. However, Technology Readiness in ETHICIST also considers the potential of disruptive innovation. We do not consider whether the disruption is positive or negative, only the level of uncertainty it precipitates. A low level of disruption might occur at the component level as new use cases are identified. A high level of disruption might occur if a higher performing, but less robust component is introduced, e.g. through machine learning. Unforeseen external factors might also disrupt comparatively mature technology, e.g. security vulnerabilities like Spectre (Kocher et al., 2019) which shed light on long-standing design decisions around speculative execution on modern microprocessors. A high-level of disruption could lead to uncertainty around how the technology might be used or deployed, which could cause problems if idiomatic practice evolves in a manner not conducive to military air domain utility. The attribute labels for d range from 1 to 9, and lower values corresponding to less certainty, with values distributed equally across this range, i.e. $9 = 0.11$, $8 = 0.22$, ... $1 = 1$.

Certification Readiness considers how ready the technology is for certification or qualification. The certification readiness of some technology can catch up with technology readiness if the latter becomes mature and stable. For example, as technology readiness of object orientation heightened and it became a dominant paradigm for software design and development at the start of the millennium, extensions were added to key standards (e.g. DO-178C), which subsequently raised its certification readiness. The values for this attribute are Immediate ($I = 0$); High ($H = 0.33$);

Medium ($M = 0.66$); Low ($L = 1$).

Scale of Complexity considers the scope of some emerging technology; the broader its scale or coverage, the more complex it is likely to be. This results in a greater risk to certification because a larger value of information has to be provided. For example, a concept which is entirely software bound might, with sufficient time and effort, be amenable to validation and verification. Emergent or non-deterministic behaviour is more likely to be yielded from systems composed of hardware and people, or systems composed of other systems.

The values for this attribute are Software ($S = 0.2$); Software & Hardware ($SH = 0.4$); Software, Hardware, and People ($SHP = 0.8$); Systems of Systems ($SoS = 1$).

Customer Demand captures the level of customer "pull" for some technology. It might appear odd that a customer is unduly concerned about some technology running on a system, given the level of indirection between the former and the latter. However, this attribute could be based not only the perceived utility in the technology, but also on social factors. Customer demand could be influenced by a range of factors, e.g. individual perception, personal experience, cost, or knowledge of privileged information that cannot be disclosed. The values for this attribute are None ($N = 0$); Low ($L = 0.33$); Medium ($M = 0.66$); High ($H = 1$).

3.2. *Weight attributes*

As some attributes will be more important to decision makers than others, the Swing Weight method (Keeney and Raiffa, 1993) is used to assign attribute weights. The method requires envisaging a scenario with the worst of all possible attribute values, prioritising the attribute with the greatest potential for impact by assigning the maximal value (100), and assigning values between 0 and 100 to the remaining attributes, with values of relative importance to the first attribute.

3.3. *Score attributes*

Based on Formula 1, the risk of a technology concept x is calculated using Formula 2:

$$Risk_x = w_r v_r(x_r) + w_c v_c(x_c) + w_d v_d(x_d) + w_o v_o(x_o) \quad (2)$$

For example, consider the emerging technology *Quantum Clocks* (q), where all attributes are equally weighted as 0.5, and attribute values are $t = 4$, $c = L$, $x = SH$, and $d = N$.

$$\begin{aligned} Risk_q &= w_t v_t(x_q) + w_c v_c(x_q) + w_x v_x(x_q) + w_d v_d(x_q) \\ &= (0.5 \times v_t(4)) + (0.5 \times v_c(L)) + (0.5 \times v_x(SH)) \\ &\quad +(0.25 \times v_d(N)) \\ &= 0.33 + 0.5 + 0.2 + 0 \\ &= 1.03 \\ &\approx 1 \end{aligned}$$

The final step in this activity is to assign a risk rating based on Table 1. The rating represents the recommended posture that should be taken to manage the emerging technology risk. The ranges are not evenly distributed because, given the importance of certification, there is a bias towards greater leadership of risk.

To illustrate the ease with which this risk assessment can be supported, we managed the emerging technology risk using a version control Excel workbook. Excel was chosen due to its ubiquity and interoperability, but a web-based front end would also be suitable. Formulas calculate risk values based on on drop-down values for t, c, x, and d, and anonymised participant response data is used to automatically calculate the respective weights for each attribute.

Table 1. Emerging Technology Risk Ratings

Rating	Lower Bound	Upper Bound
Ignore	0	0.2
Watch	0.2	0.4
Influence	0.4	0.6
Lead	0.6	1

3.4. *Model concept maps*

ETHICIST relies on three orthogonal concept maps, with each map modelled in Dot (AT&T, 2012) and consequentially can be version controlled. The *technology* concept map describes the contribution relationships between different forms of emerging technology. The *standards* concept

map is a regulatory hierarchy specific to the type of system that requires certification. The *system* concept map captures relationships between functional elements of the system that the emerging technology might deploy to or heavily influence the design of. We acknowledge there is no certainty to where technology will be deployed, but placing the technology in context of a system further helps make sense its risk.

The technology concepts are enumerated to identify relationships between technology and standard (and guidelines, etc.) concepts indicating regulations governing the technology, and technology and system concepts indicating where technology partially or fully operationalises the system concept. Where a technology concept has some potential but indeterminate impact on systems or standards, this is reflected by mapping the technology to the respective map in general. All three concept maps are then merged into a single, consolidated map.

3.5. *Generate risk concept map*

Risk concept maps are generated based on a selection of one, some, or all emerging technology concepts. Based on the technology concepts selected, two steps are performed.

The first step entails calculating cascaded technology risk resulting from concept map relationships. We consider a relationship from technology A to B as one where the risk associated with A contributes to an increase in the risk associated with B, where the concept map is a direct acyclic graph.

The approach for calculating this risk is specified in Algorithm 1. The $cascadedRiskScore$ takes as input the technology name, a risk model dictionary mapping technology to risks, a dictionary mapping technology to the set of technology contributing to it, and a set of technology with cascaded risk scores. If the technology's cascaded score has already been calculated, or there is no contributing technology then the pre-existing risk score is returned (Lines 2–9). The $cascadedRiskScore$ algorithm is called recursively for each technology contributing to the technology item of interest (Line 10); each of these items is output from the $edge$ function parametrised by the originally inputted technology name. The cascaded score for each attribute is based on maximal attribute value for each contributing risk's attribute value and the attribute value for the non-cascaded risk (Lines 11–14). Before the cascaded risk value is returned, the risk model is updated to reflect the cascaded risk score for the technology (Line 15).

Algorithm 1: Cascaded risk score

Input : $tech$ - technology name, $tech\_risk$ - dictionary of node risks, $tech\_conts$ - dictionary of technology contributing to each node, $calculated$ - set of node names with calculated cascaded risks
Data: c_tech - name of contributing node, $contRisks$ - set of risks contributing to the node
Output : $cascadedRisk$ - cascaded technology risk

1 **Function** *cascadedRiskScore(tech, tech_risk, tech_conts, calculated)* **is**
2 $cascadedRisk \leftarrow$ tech_risk $tech$;
3 **if** $tech \in calculated$ **then**
4 return $cascadedRisk$;
5 **end**
6 **if** $tech \notin tech\_conts$ **then**
7 $calculated \leftarrow \{tech\} \cup calculated$;
8 return $cascadedRisk$;
9 **end**
10 $contRisks \leftarrow$ map (λ c_tech : String • cascadedRiskScore(c_tech $tech\_risk$ $tech\_conts$ $calculated$) , edges $tech$) ;
11 $cascadedRisk.t \leftarrow$ max(map (λ r : Risk • $r.t$, $contRisks$) $\frown$ $cascadedRisk.t$);
12 $cascadedRisk.c \leftarrow$ max(map (λ r : Risk • $r.c$, $contRisks$) $\frown$ $cascadedRisk.c$);
13 $cascadedRisk.x \leftarrow$ max(map (λ r : Risk • $r.x$, $contRisks$) $\frown$ $cascadedRisk.x$);
14 $cascadedRisk.d \leftarrow$ max(map (λ r : Risk • $r.d$, $contRisks$) $\frown$ $cascadedRisk.d$);
15 $tech\_risk \leftarrow tech\_risk \oplus \{tech \mapsto cascadedRisk\}$;
16 $calculated \leftarrow \{tech\} \cup calculated$;
17 **return** $cascadedRisk$;
18 **end**

In the second step, the impact of emerging technology risk is applied to related standard or system concepts. The impact on a standards concept is based on the median of the certification readiness values of contributing technologies. The impact on a system concept is based on the median of

the Technology Readiness & Disruption values of contributing technologies. In both cases, thes attribute values are considered in isolation, i.e. not based on cascaded values.

To reflect the calculated risk and impact ratings, concept map nodes were filled based on the colour scheme shown in Figure 1. To facilitate rapid re-generation of concept maps based on changes to technology concept risk values and concept maps, a collection of Python scripts were used to process the Excel workbook containing the risk data, and automatically re-generate the concept maps.

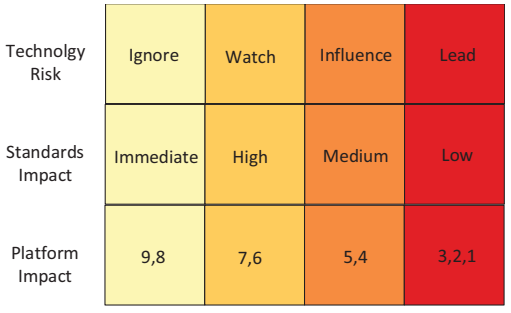

Fig. 1. Concept map node colours

4. Worked Example

We evaluated the impact of emerging technology risk with ETHICIST for 93 selected emerging techology on military air platforms. The emerging technology concepts were cogent to autonomy, computational hardware, infrastructure, novel forms of regulation, software, and tools. The concept map for standards was a hierarchy of regulations for military airworthiness. The concept map for systems was based on the relationship between avionics system components.

Using a combination of subject matter expertise and desk research, attribute values were completed for each emerging technology item. An initial set of weights was elicited from 13 software & systems dependability experts based on the responses to a hypothetical scenario. The scenario (below) was based on a novel air platform (MAKKA) and a novel piece of technology (Pandora).

A sudden conflict in Europe has forced the immediate deployment of the MAKKA platform. The central component of MAKKA is Pandora. Pandora has been subject only to a very preliminary scientific evaluation, yet could disrupt most elements of MAKKA, e.g. change working practices of air and ground crews, programmable elements, and paradigms for developing and maintaining software running on it. Pandora also requires substantial interaction across different domains, and interaction with hardware, software, and human elements within them. However, for reasons that have not been elaborated on, Pandora's inclusion in MAKKA is essential.

Participants emailed responses to two questions. First, if they had the budget to completely remove risks associated with one of technology readiness, certification readiness, complexity, and customer demand, which would they choose. Second, with a more limited budget, prioritise the remaining risks for removal, indicating the level of importance of each compared to the attribute selected for the first question.

For brevity, we consider only a small subset of these concepts in this paper. Specifically, we examine emerging technology associated with additive manufactured components for wearable computing devices. The individual, non-weighted attribute values for different technology concepts is specified in Table 2, together with the risk rating. The attribute values were set based on subject matter expertise, and the risk rating was assigned based on the weighted risk score calculated by Formula 2. Figure 2 shows the impact of the cascaded scores when projected onto the concept maps. For example, *Model-Based Engineering* has a comparatively low risk rating when considered independently, but when rated with respect to emerging technology influencing it, e.g. domain specific languages specifying the models, and data safety guidance constraining them, this rating notably increases.

5. Discussion and Conclusion

In this paper, we presented ETHICIST: a systematic approach for assessing and managing emerging certification techology risk. In doing so, the approach has two benefits.

First, ETHICIST provides transparency without de-anonymising contributing participants. Be-

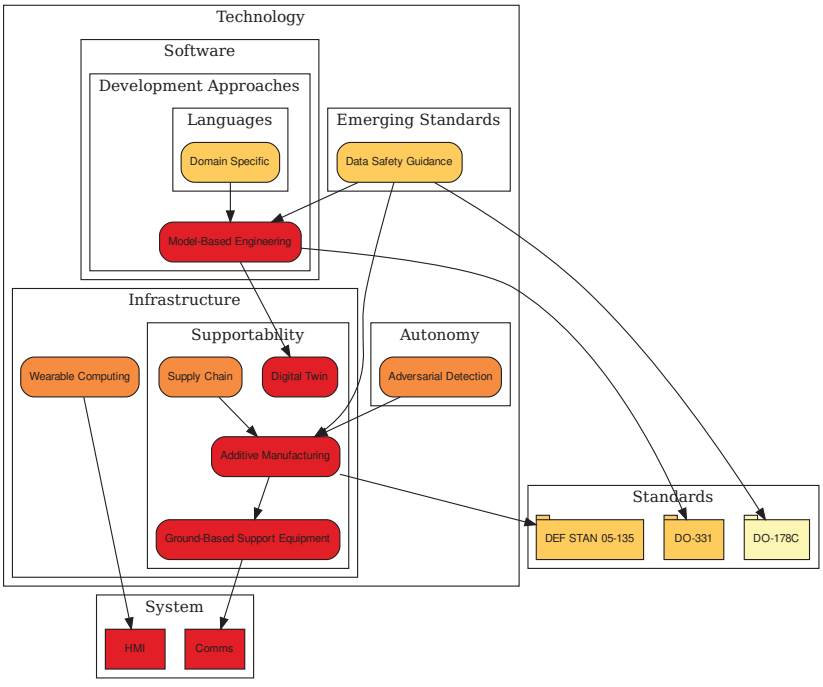

Fig. 2. Emerging technology risks for printable wearable technology

Table 2. Additive Manufacturing and Wearable Computing Certification risks

Technology	t	c	x	d	Rating
Additive Manufacturing	9	H	SH	H	Influence
Adversarial Detection	6	M	S	M	Influence
Data Safety Guidance	3	I	S	L	Watch
Digital Twin	6	M	S	H	Influence
Domain Specific	9	M	S	N	Watch
Ground-Based Support Equipment	9	M	SH	L	Watch
Model-Based Engineering	9	H	S	H	Watch
Supply Chain	9	M	SHP	H	Influence
Wearable Computing	7	M	SHP	L	Influence

cause the participants have been anonymised in the risk model, it is possible to inspect the participant data, weights, and explore the impact of modifying attributes for some technology, e.g. based on feedback for different participant cohorts, changes to customer perceptions, or advances in technology. Because the concept map data is also transparent, it is easy to explore the implications of emerging technology on individual systems in light of different regulatory combinations. The use of multi-attribute decision analysis also makes both the approach and the outputs familiar to policy makers, which may choose to take further actions based on the results.

Second, both the approach and tools used by ETHICIST are scalable. It is comparatively easy to obtain stakeholder input to revise the risks and concept maps. For example, survey research could be used to obtain participant data from a large range of stakeholders or subject matter experts, and the use of configuration control allows changes to be easily tracked. The tools selected are robust enough to scale to a greater range of participants, and significantly larger technology, system, and regulatory models.

Risk assessment in ETHICIST remains no less subjective than for risk assessment in general. Subjectivity exists around the expertise contributing to, and the assumptions underpinning the risk data and concept maps. This is particularly the case for customer demand and associations between emerging technology concepts where subject matter expertise was limited. The feedback from subject matter experts was also open to interpretation, particularly given the difficulties

some participants had responding to the scenario in Section 4. Fortunately, projecting the outputs of the risk assessment to the concept maps and analysing the results provided some level of sensitivity analysis for the multi-criteria decision analysis. The validity of the results obtained could still be further validated with additional subject matter expert reviews of the technology risk attributes, and concept maps - particularly on the impact on the standard and system maps. This would also reinforce the transparency of the outputs of this risk assessment approach.

Acknowledgement

This document is an overview of UK MOD sponsored research. The contents of this document should not be interpreted as representing the views of the UK MOD, nor should it be assumed that they reflect any current or future UK MOD policy.

References

AT&T (2012, June). Graphviz web site. `https://www.graphviz.org`.

Bunn, D. W. (1984). *Applied Decision Analytics*. McGraw Hill.

Butler, S. and P. Fischbeck (2002). Multi-Attribute Risk Assessment. *Proceedings of the Symposium on Requirements Engineering for Information Security (SREIS) 2002*.

Communities and Local Government (2009). *Multi-criteria analysis: a manual*. Department for Communities and Local Government.

Edwards, W. (1977). How to Use Multiattribute Utility Measurement for Social Decisionmaking. *IEEE Transactions on Systems, Man, and Cybernetics 7*(5), 326–340.

Faily, S., J. Lyle, A. Paul, A. Atzeni, D. Blomme, H. Desruelle, and K. Bangalore (2012). Requirements Sensemaking using Concept Maps. In *Proceedings of the 4th International Conference on Human-Centered Software Engineering*, pp. 217–232. Springer.

Frerking, M. A. and P. M. Beauchamp (2016). Jpl technology readiness assessment guideline. In *2016 IEEE Aerospace Conference*, pp. 1–10.

Fujiwara, D. and R. Campbell (2011). *Valuation Techniques for Social Cost-Benefit Analysis: Stated Preference, Revealed Preference and Subjective Well-Being Approaches - A Discussion of the Current Issues*. HM Treasury and Department for Work and Pensions.

Goldschmid, J. and S. Corns (2021). A cluster-based framework for interface analysis in large-scale aerospace systems. *Systems Engineering 24*(5), 339–351.

Jesus, G. T. and M. F. Chagas Junior (2022). Using systems architecture views to assess integration readiness levels. *IEEE Transactions on Engineering Management 69*(6), 3902–3912.

Keeney, R. L. and H. Raiffa (1993). *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. Cambridge University Press.

Kocher, P., J. Horn, A. Fogh, , D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom (2019). Spectre attacks: Exploiting speculative execution. In *40th IEEE Symposium on Security and Privacy (S&P'19)*.

Mankins, J. C. (1995). Technology Readiness Levels: A White Paper. Technical report, NASA.

Mankins, J. C. (2009a). Technology readiness and risk assessments: A new approach. *Acta Astronautica 65*(9), 1208–1215.

Mankins, J. C. (2009b). Technology readiness assessments: A retrospective. *Acta Astronautica 65*(9), 1216–1223.

Martin, B. and B. Hanington (2012). *Universal Methods of Design: 100 Ways to Research Complex Problems, Develop Innovative Ideas, and Design Effective Solutions*. Rockport.

Moni, S. M., R. Mahmud, K. High, and M. Carbajales-Dale (2020). Life cycle assessment of emerging technologies: A review. *Journal of Industrial Ecology 24*(1), 52–63.

Novak, J. D. and D. B. Gowin (1984). *Learning How To Learn*. Cambridge University Press.

Salazar, G. and M. N. Russi-Vigoya (2021). Technology readiness level as the foundation of human readiness level. *Ergonomics in Design 29*(4), 25–29.

Sutherland, S. and S. Katz (2005). Concept mapping methodology: a catalyst for organizational learning. *Evaluation and Program Planning 28*(5), 257–269.