Safety Science 172 (2024) 106424

Contents lists available at ScienceDirect

Safety Science

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

Understanding safety engineering practice: Comparing safety engineering practice as desired, as required, and as observed

Matt Osborne *, Richard Hawkins, Mark Nicholson, Rob Alexander

Assuring Autonomy International Programme, University of York, United Kingdom

ARTICLE INFO

Keywords:
Safety practice
Safety framework
Safety process
Safety lifecycle
Safety improvements

ABSTRACT

Good safety management means that continuous attempts are made to improve safety engineering practice. These improvements are often through creating interventions to perceived problems. Many of these interventions seem to have been largely ineffective, suggesting that they may not be addressing the real impediments to good safety engineering practice. We do not argue that existing tools for improving safety engineering practice (such as checklists) are necessarily deficient, rather we challenge whether they are being employed to correct the causes of impediments to better practice. Safety practice 'As Observed' (the actual safety engineering activities performed) is informed by defined processes (safety practice 'As Required'). These processes aim to ensure practice achieves the best safety outcomes (safety practice 'As Desired'). For many different and complex reasons 'As Observed' safety practice may not be equivalent to the safety practice 'As Required'. Similarly safety practice 'As Required' may not be equivalent to safety practice 'As Desired'. All of these discrepancies could play a significant role in poor safety engineering practice. By exploring these discrepancies it becomes possible to understand the causes of deficiencies in practice, and to start to propose effective interventions. In this paper we introduce and discuss a process for understanding safety engineering practice based around modelling safety practice 'As Desired', 'As Required', and 'As Observed', and the interactions between these elements. We describe how this process can be used to evaluate safety engineering practice and inform the design of effective improvements. We present an example of how the process may be applied to understand safety practice for software safety assurance in the military domain.

1. Introduction

1.1. The problem

A variety of stakeholders, such as companies, organisations, standards committees, and researchers, have tried to implement changes to safety engineering practice over the last few decades to 'fix' issues, yet these seem to have been largely ineffective and problems persist. Often, proposed 'fixes' offer nothing more than a new variant of existing analysis methods which are only evaluated in terms of how they could have prevented a famous accident/incident (decried as YAAPing by Rae et al. in Rae et al. (2020).

Surprisingly few empirical investigations have been undertaken into what constitutes good fixes, nor why improvement attempts have historically not been effective. Anecdotal evidence suggests that issues with safety engineering practice persist. This may be for instance, due to only a limited subset of the elements of safety engineering practice being considered. An example of a fix identified by focusing on a limited subset of the elements of safety engineering practice is the use of checklists to improve adherence to the steps of a procedure. However, adherence may not be the problem. It might be that the procedure was inappropriate or incorrect. In this case, the fix would need to focus on the identification and control of processes. Rasmussen (1997) focused on the control of work processes to avoid "accidental side effects causing harm to people, environment, or investment".

Any, all, or none of these examples could be appropriate explanations. We currently do not have the evidence. We need empirical data to establish whether this is true.

Rooksby et al. (2009) set out to ethnographically observe the relationships between the documented or expected procedures of software testing and the 'reality' of practice, and what the testing "actually involved". In doing so they were not concerned with having preconceived ideas on what "ought to be done", rather they set out to observe and characterise the socio-technical issues emanating from the practice of software testing.

* Correspondence to: Assuring Autonomy International Programme, Department of Computer Science, University of York, Deramore Lane, York, YO10 5GH, United Kingdom.
E-mail address: matthew.osborne@york.ac.uk (M. Osborne).

https://doi.org/10.1016/j.ssci.2024.106424
Received 16 August 2023; Received in revised form 13 December 2023; Accepted 4 January 2024
Available online 13 January 2024

Before embarking on a program of safety engineering practice improvement we believe we must be as confident as possible that we have both identified where any issues lie, and that we are addressing the complete set of problems and their causes. We start therefore from the position taken by Rae et al. (2020) in that we must describe current work before changes are prescribed. We aim to "capture, describe, and document, as well as conceptualise" (Von Krogh et al., 2012) safety engineering practice so that appropriate theorising can proceed. This paper introduces a process to improve our ability to understand safety engineering practice.

1.2. The contribution

In this paper, we present a process for studying the phenomenon (Von Krogh et al., 2012) of safety engineering practice. By 'safety engineering practice', we mean the activities carried out to assure and demonstrate the safety of a system. A 'system' is a combination of interacting elements organised to achieve one or more stated purpose (ISO/IEC, 2008). This safety engineering practice requires input by safety engineers, but also involves many other engineering disciplines and other specialisations. For brevity we now refer to the activities as 'safety practice', and for clarity refer to the multi-disciplinary professionals working on the activities as 'safety practitioners'.

We have created a process to understand safety practice because anecdotal evidence suggests that issues with safety practice persist, and such persistent issues COULD be due to poor practice, but we need empirical data to establish whether this is true.

Before researchers, safety practitioners, or managers embark on a program of safety practice improvement we believe we must be as confident as possible that we have both identified where any issues lie, and that we are addressing the complete set. Moreover, we must be confident that we are addressing appropriately-distal causes of issues rather than only proximal causes or indeed just symptoms. The process we present is therefore designed to help researchers and safety practitioners understand safety practice in its entirety.

Fixes to safety practice have to date been implemented despite surprisingly few empirical investigations (Rae et al., 2020), and this may lead us to question the effectiveness of these fixes in addressing the real impediments to good safety practice. We argue that many implemented changes to safety practice were ineffective because the analysis that revealed their necessity was limited to a consideration of only a limited subset of all the elements that constitute safety practice. We return to consider the elements which constitute safety practice later.

There have been a few empirical investigations of how safety practice is carried out (e.g. Rae et al. (2020)). These studies appear to have been centred in a 'generalised area' of practice, and there are limitations of these studies as applied to safety engineering practice. Safety researchers have to date considered both the safety of working practices, and the work of assuring the safety of working practices (e.g. Provan et al. (2019)). These investigations have been referred to using the form (Safety) <**Work As**> <**X**>, such as:

- Work as **Desired:** how people would like work to take place (Rae et al., 2021)
- Work as **Imagined:** what people imagine everyday work to be (Hollnagel, 2018)
- Work as **Done:** that carried out by the workforce (Provan et al., 2019)
- Work as **Documented, or Observed:** an observed and documented assessment of work carried out by the workforce (Hollnagel, 2012).

For example, Hollnagel has described the differences between the construct of "Work as Imagined" and "Work as Done", noting that this could be extended further through analyses of different "lenses" of 'Work as X' (such as between work as documented and work as observed) (Hollnagel, 2012). Such a theoretical construct allows the analyst to ethnographically identify differences between Work as Imagined and Work as Done (such as the work of Provan et al. (2019) in considering the safety of work and safety work), thereby making changes to either — in order to improve safety and resilience.

A potential limitation of these investigations is that they tend to be centred on investigating theoretical discrepancies between two elements of safety practice and tend to compare these two elements of <**Work As**> <**X**> as an omni- or bi-directional relationship, which makes an implicit assumption that any issues with practice emanate from only the elements under consideration. Examples of such investigations of safety practice include Hollnagel's description of the construct of "Work as Imagined" and "Work as Done" (Sujan et al., 2019), and the empirical investigations of this construct by researchers such as Provan, Rae, and Dekker (in Provan et al. (2019)).

Further, the terms used in these investigations (<**Work As**> <**X**>) are heavily entrenched in the discipline of Occupational Health and Safety (OHS), or the "safety of work" (Provan et al., 2019), and do not yet extend to consider the work of a safety practitioner working to assure the safety of a complex safety-critical system.

Each of these omni- or bi-directional approaches (Hollnagel's characterisations of Work as Imagined and Work as Done, Provan et al.'s considerations of the differences between the work of safety, and safety work, and Gawande's use of checklists) are focused on assessing specific elements of safety practice. Should changes be made as a result of these assessments, we cannot currently be sure whether such approaches are identifying and fixing real problems, as not all elements of safety practice have been considered concurrently. Nor do we know whether such 'fixes' are introducing new issues, or undermining other elements of safety practice. Notwithstanding, such approaches have become the accepted norm when issues with safety practice are suspected. Noting that academics have pre-conceived ideas on what the issues and fixes may be (Rooksby et al., 2009), and that "safety professionals are not confident operationally of how to create safety improvement" (Provan et al., 2019), the question therefore arises, how do we improve safety practice?

To understand safety practice fully, we need to have a model which can be used to discuss and evaluate safety practice in an ontologically-robust manner that allows the different elements of safety practice to be represented equally. We have therefore developed a model which is capable of linking the different elements together. This allows us to evaluate their relationships, and reveal any nuances and subtleties.

The safety work we are concerned with is the practice of the safety practitioner, who must ensure and demonstrate the safety of a system. To the authors' knowledge, no process exists for understanding safety (engineering) practice per se. This paper provides such a process.

First we introduce and discuss our model and process for understanding safety practice, and describe how the steps which underpin the process can be used to evaluate practice as a precursor to identifying effective improvements. Thus, the application of this process will enable mitigation research (Rae et al., 2020) into effective remedies for the identified deficiencies. This mitigation research is outside the scope of this paper, but we make recommendations for future work in Section 5.

The remainder of this paper is structured as follows. In Section 2 we present the model for understanding safety engineering practice, and discuss the elements and inter-relationships between them. In Section 3 we give an overview of the process for understanding safety engineering practice, before expanding on the process to demonstrate how models of practice are created and assessed in Sections 3.4 and 3.5 respectively. In Section 4 we provide an example of how we have applied this process, before considering the next steps in Section 5.

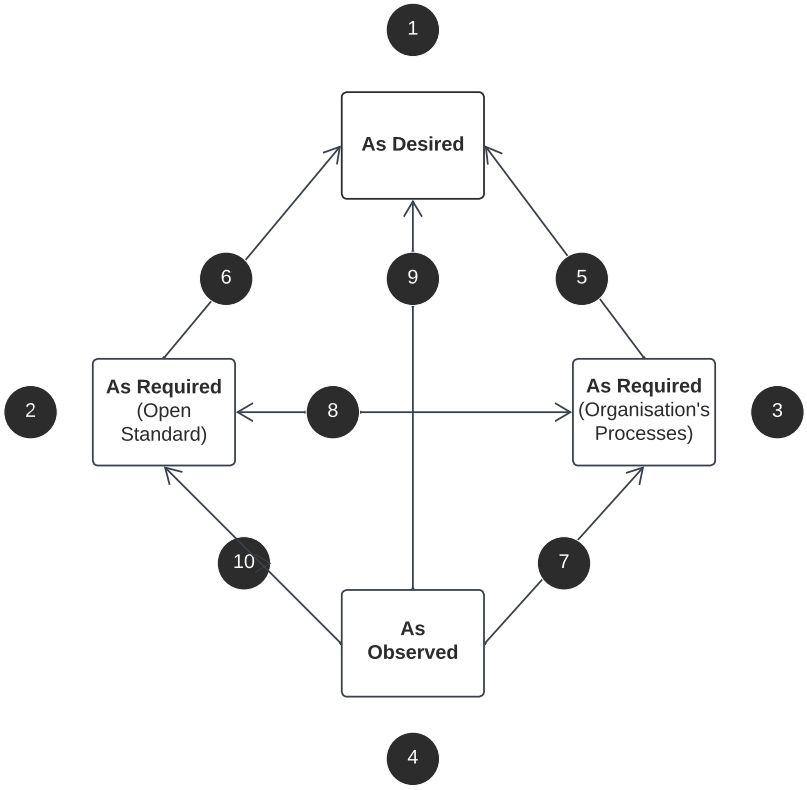

Fig. 1. The elements of safety engineering practice.

2. A model of safety engineering practice

In this section we present our process to model and understand safety practice. The aim is to enable an understanding of how safety practice is carried out, and why it is done in the manner that it is. Specifically, we seek an understanding of how safety practice is desired to be, how the desired practice is imparted to those required to enact it, and how safety practice as desired is interpreted and implemented by safety practitioners. The resulting model is useful because it creates a detailed representation of safety practice whilst remaining as simple as possible.

We deliberately avoid theorising as to whether and why poor safety practice exists within an organisation, rather we adopt the phenomenon-based research suggested by von Krogh, Rossi-Lamastra and Haefliger, and create a process which will allow an organisation to identify and gather relevant data using our innovative research design (Von Krogh et al., 2012).

An organisation such as a developer of safety-critical systems, or a committee of experts aiming to standardise safety practice must establish a notion of what idealised safety practice must constitute. Such organisations must then develop a means of imparting this idealised practice to those safety practitioners required to carry it out. Finally, the safety practitioners working for an organisation must carry out the practice in the manner intended.

Logically, this constitutes three elements of safety practice, which improves on the simpler model of 'work as imagined versus work as done' (Rae et al., 2021).

The three elements of safety practice, spanning the idealised concept of what best practice *should* be, the manner in which it is imparted to practitioners, and the reality of practice as carried out by practitioners are shown in Fig. 1. Each of the elements (the numbers relating to Steps that are described as the process is explained), along with the relationships between them is shown:

- Safety Practice as Desired
- Safety Practice as Required
- Safety Practice as Observed.

All existing (safety) (<**Practice As**> <**X**>) can be mapped onto these three elements, and whilst we argue these elements are necessary, we cannot yet argue whether this is sufficient — although further instantiations of the model, and implementations of the associated process will reveal the levels of confidence in sufficiency.

In the following subsections we take each element in turn:

2.1. Safety practice as desired

Safety Practice as-desired is characterised by a set of safety objectives that are held by stakeholders within a Project.[1] Together these safety objectives embody the safety philosophy and risk appetite of the Project. Complete and correct compliance with safety practice as-desired should manifest in a product which is acceptably safe to operate in a given operating environment. Of course, poorly identified, or articulated, or incompatible safety objectives may be a source of safety process failures.

For reasons we next discuss, there is no 'one-size-fits-all' approach to defining safety practice as desired across all sectors and applications, and we do not prescribe what as-desired practice is constituted by in all cases. However, we do provide a process through which a project can establish and define as-desired safety practice, and we return to the specifics of this process later.

[1] A Project can emanate from any industry and application, and its level of design abstraction can range from a full product, through systems, items, or down to the hardware and software levels. A project may involve one or more organisations.

2.1.1. A discussion on safety practice as-desired

For a project that wishes to understand safety practice, defining what constitutes as-desired practice may be the most complex and challenging element. Even if a project leader believes that their as-desired practice is clearly defined and described, it may not actually be explicitly documented anywhere — or at least it may not be described in a manner which allows its philosophical attributes to influence an engineered design. There are many potential reasons for this.

From the experience of the authors, Open Standards (such as BSi (2010) are often held to represent safety engineering practice as-desired in the form of codified expertise (Asplund et al., 2020). We argue that this cannot in fact, be considered to be sufficient for as-desired practice. Such standards contain a mixture of normative requirements, informative guidance, and safety philosophy. Whilst normative requirements can be measured and even audited against, it is not immediately clear how the required intent and safety philosophy of a Standard can be validated and verified in a design. Nor can the committee responsible for an Open Standard presuppose the safety philosophy and risk appetite of an organisation which aims to comply with it.

Whilst compliance with requirements can be measured through qualitative and quantitative means, and processes can be assessed as to their correctness and completeness, it is not possible to measure the alignment of intent and philosophy of a standard in an auditable manner. This presents an open question of whether a safe design is achieved by following a standard, or whether a safe design is achieved because the safety practitioners cared enough to deliver it.

Standards committees therefore face the challenge of ensuring that their standard will represent the overarching intent and safety philosophy for any organisation that elects to comply with it as a means of assuring the safety of their design. Consider BS EN 61508 (BSi, 2010), which was established as a unified, generic standard that aims to achieve functional safety by minimising risk through the application of Safety Integrity Levels (SIL) to safety functions. A challenge that arises through offering such a 'pan-industry' approach is whether the standard can appeal to the different safety philosophies of disparate projects across industry, many of whom will also have their own disparate supply chain.

Different Open Standards have been designed with variations in the means of their intended application, and predicated on different principles of risk management. For example, whilst BS EN 61508 (BSi, 2010) may be formally certified against for compliance, the standard is not designed to be applied in tandem with a regulatory or certifying body specifically. This is in contrast with the ARP 4754A (ARP) suite of standards (SAE Aerospace, 2010), which provide "safety recommended practice" as an acceptable means of compliance for certification by regulatory bodies (such as the FAA and EUROCAE).

The ARP's safety philosophy is to moderate the severity of outcome through the application of Design Assurance Levels (DAL) to components and systems. A challenge for the Standard's body here is ensuring that its safety philosophy is applicable to civil and military manufacturers of aerospace systems — both rotary and fixed wing (and indeed to organisations outside of the aerospace sector which have adopted it). Projects and regulatory bodies in these disparate sectors may have differing approaches to risk which are predicated on factors other than the moderation of the severity of outcome (alone). They may also have differing safety philosophies.

Standards such as BS EN 61508 and ARP 4754A are predicated on philosophies that safety is achieved through the moderation of risk or severity of outcome, yet an project's philosophy may not be founded in risk or severity reduction per se. The as-desired safety practice of a project may be founded on principles (e.g. Hawkins et al. (2013a)), systems theory (e.g. Leveson and Thomas (2018)), Normal Accident Theory (as discussed in Haavik (2021)) or predicated on specific attributes such as resilience (e.g. Hollnagel et al. (2006)), or high levels of reliability (e.g. Roberts (1990)). Further, a project may also operate a safety management system which is based on either centralised control (i.e. 'Safety I'), or guided adaptability (i.e. 'Safety II') (Provan et al., 2020).

Whilst we accept that a project can impart some aspects of idealised practice through normative requirements, a final challenge concerns the efficacy of the processes, procedures, techniques, and methods that manage these requirements. Consider again BS EN 61508 (BSi, 2010), for which compliance is met by achieving its objectives; which are held to have been met if the applicant meets the requirements (clauses); which in turn can be instantiated by following recommended techniques and measures (with accompanying levels of importance predicated on the SIL). What is not known is whether this process (objectives met by requirements, which are instantiated by techniques and measures) manifests in a safe design directly. These objectives, requirements, techniques and measures are not derived from an evaluation of empirical data, but predicated on expert judgement and opinion — based on the experience of the standard's committee members.

Similarly, UK Defence Standards require both clauses and objectives to be met. Defence Standard 00-055 (MoD, 2016) requires not only that its specific clauses be met, but demands compliance with its objectives (expressed as 5 Principles). One clause requires the Contractor to select an Open Standard to comply with. Once selected, the proposed Standard is agreed with the MoD and complied with. Additional work is then required to conform with the appropriate 'Military Delta'. The 'Military Delta' stipulates additional requirements to recover perceived shortfalls in the selected standard when used in a military context. Whilst the 5 Principles are laudable (predicated on Hawkins et al. (2013a)), it is not immediately clear how the Principles are met by compliance with the clauses of the Standard itself, the selected Open Standard, nor the Military Delta required by the Defence Standard.

Because of the complexities and challenges of defining as-desired safety practice we have discussed, and the acceptance that as-desired safety practice must be sector- and application-specific, we do not assert in this paper what safety practice as-desired *should* be, but explain how the element of as-desired practice can be modelled and evaluated.

2.2. Safety practice as required

Safety Practice as-required is constituted by a set of processes which are designed to instantiate safety practice as-desired when followed by a safety practitioner. Safety practice as-required is a representation of how a project explicitly requires its personnel to carry out safety practice. Used by standards bodies and organisations alike, safety practice as-required describes the processes required to be followed by safety practitioners.

The activity to convert safety practice as-desired into safety practice as-required is not considered as part of this process as there exist many projects emanating from disparate industries and applications, and there can be no 'one-size-fits-all' approach to this. The process to understand safety practice *does* provide a multi-directional mechanism for assessing the elements of safety practice however, as we will discuss.

2.2.1. A discussion on safety practice as required

There are two types of as-required safety practice, the first being an Open Standard such as the ARP 4754A suite of publications, which requires the adoption of a 'V-Model' lifecycle that aims to show the interaction between safety processes and design and development processes, and is used in an iterative and concurrent manner from 'Platform' level down to 'Item' and 'Software' levels. Another example of as-required practice expressed by a standard is that specified by BS EN 61508, and whilst this standard does not require a specific development lifecycle model, it requires that its objectives are met. Each objective is achieved through compliance with the standard's requirements (clauses) — which may in turn be instantiated by the use of selected techniques and methods.

As shown by Step number 8 in Fig. 1, and discussed in Section 2.1.1, Standards Committees design and compile their Open Standards in a

manner that expresses a set of lifecycle activities which – when adopted by a project, and followed by the project's safety practitioners – aims to comply with the standard's version of as-required safety practice.

We must also look to the practice required by organisations who employ safety practitioners within a manufacturing, design, or procurement setting. Such practice is normally documented by the developing/acquiring organisation, and expressed as a lifecycle of processes that are undertaken throughout the product lifecycle. Such an organisation may develop its own practice predicated on the requirements and informative guidance expressed by an Open Standard, as the organisation seeks to demonstrate due diligence based on conformance with an Open Standard (Habli, 2017) by incorporating the Standard's knowledge within its processes (Antonino et al., 2014).

As this practice represents the Intellectual Property Rights (IPR) of the organisation, we refer to such processes as 'Closed Standards'. These Closed Standards may offer different representations of safety practice to Open Standards, but they are protected by IPR which prevents, for example, the widespread sharing of practice and issues arising. When considering the development of complex 'Systems of Systems' (SoS), the constituent systems/components may also be developed by differing organisations who adhere to different Open/Closed safety Standards. Should these differences between Open and/or Closed standards exist, we need to understand the reasons for this, and what impact they may have on safety practice as-observed.

Akin to Morley et al.'s argument that research into policies for safely implementing Artificial Intelligence (AI) into the health care domain requires a prospective approach that is cognisant of the complexities of individuality, and social and organisational structures; any subsequent mitigation research into improving the effectiveness of Open/Closed standards should take into account the complex inter-relationships of all elements of safety practice (Morley et al., 2022).

Of course, organisations do not directly follow an Open Standard — they create their own set of internal processes and procedures which may be derived with the intent of meeting a specific and selected Open Standard (as shown in Fig. 1). An organisation *may* select an Open Standard and create their processes and procedures such that they will comply with the as-required practice required in the Open Standard. Open Standards have their own challenges, and these challenges may be masked or exacerbated as an organisation develops their processes in a manner that may lose the intent of the Open Standard.

The challenges associated with the creation and maintenance of Standards are well-documented (see Habli (2017) for example), but we offer no a priori hypotheses regarding any contribution made by the use of as-required processes (Von Krogh et al., 2012), but present a process to determine whether and how a standard can contribute to the problem.

A challenge for any project is to define as-required processes and procedures which meet their as-desired safety practice AND to express this in a way that clearly and explicitly imparts the requirements, intent, and philosophy to those individuals charged with implementing it.

Individuals control a lot of detailed safety practice as we shall discuss in the next section. Further, individuals form both organised and unorganised groupings within an organisation and will learn, modify behaviours, and create their own norms and rules for self-governance (Von Krogh et al., 2012). Our process seeks to specifically elicit such nuanced practices and behaviours, and how they relate to a project's processes (Von Krogh et al., 2012).

2.3. Safety practice as observed

Safety practice as-observed is a representation of how a project's safety practitioners carry out their work. This safety practice as-observed is nominally controlled and directed through defined processes, and there will often be multiple sources of these processes which "can be represented, with imperfect fidelity, through standardised models and procedure" (or safety practice as-required) (Rae et al., 2020).

For example, we may observe in an organisation that work on eliciting software safety requirements starts before component-level safety requirements are fully established. This may deviate from the lifecycle model that portrays a chronological and sequential hierarchical decomposition of safety requirements.

For many different and complex reasons, safety practice as-observed may not be equivalent to safety practice as-required, and this could contribute to achieving poor safety engineering outcomes. We must therefore strive to ensure safety practice as-observed aligns with safety practice as-required — and that both elements of practice are fit for purpose (i.e. that safety practice as-required is aligned with safety practice as-desired as well).

Of course, it may be the case that those charged with undertaking as-observed practice are aware of shortcomings in safety practice as-required and that they have made subtle (perhaps hidden) improvements on safety practice as-required in an attempt to align with safety practice as-desired. For example, in the example above of software safety requirements being partly derived before component-level requirements are finalised, this may be a positive deviation, in that it allows the design activities to progress within the required timescales in the specific context of the organisation that does this. Because deviations may be positive, any evaluation should be capable of identifying any subtle improvements made by those carrying out safety practice as-observed over that stated in the as-required processes — including any changes resulting in meaningful engagement with internal or external communities of practice (Asplund et al., 2020).

It is of course possible that any such 'improvements' may not actually improve safety practice, and may instead undermine other elements of practice, and may also introduce new issues. By considering all elements of safety practice (and the interrelationships between them) in a holistic manner, it should be possible to identify and mitigate the risk of unintentionally undermining safety practice.

Beyond deficiencies in the as-required process, safety practice as-observed may contain elements of what Dekker refers to as "malicious compliance" — where those charged with implementing safety practice as-required carry out processes that they know are inadequate (Dekker, 2017); or instances of work that do not contribute to achieving or demonstrating safety (or "safety clutter" (Provan et al., 2019). As-observed practice may also have instances of what Provan refers to as "role retreat" (where workers just perform their role only as defined (i.e. work to role)), or covert work systems (where work as-observed is hidden from 'outsiders' due to the fear that it will be stopped or changed, thereby making work more difficult for front-line teams) (Provan et al., 2020).

Safety practice as-observed is normally measured by auditing against work as-required. This misses the nuances and intricacies of what actually happens 'As Done'. Indeed, Provan notes that safety work needs to adapt and deviate from plans, rules, roles, and procedures because of the dynamic and emergent nature of complex systems (Provan et al., 2020).

3. A process to understand safety engineering practice

To carry out an effective evaluation of safety practice, our approach is to first create a clear and equal model of the three main elements of safety practice (Steps 1, 2, 3, and 4 in Fig. 1). Having this clear and equal model of the elements which constitute safety practice facilitates a like-for-like comparison. The comparisons are shown at Steps 5, 6, 7, 8, 9, and 10 of Fig. 1. Each of the Steps in Fig. 1 are listed below and described in more detail in the Appendix.

1. **Desired:** the as-desired representation
2. **Required (Open):** Representation of an Open Standard
3. **Required (Closed):** Representation of a Closed Standard

4. **Observed** Representation of practice as-observed
5. **Required (Closed) v Desired:** Comparison of an organisation's safety process with safety practice as-desired
6. **Required (Open) v Desired:** Comparison of an Open Standard with safety practice as-desired
7. **Observed v Required (Closed):** Comparison of observed practice with the organisation's safety engineering process
8. **Required (Closed) v Required (Open):** Comparison between the organisation's safety engineering lifecycle and the Open Standard which may have informed its development
9. **Observed v Desired:** Comparison of observed practice with safety practice as-desired
10. **Observed v Required (Open):** Comparison of observed practice with an Open Standard.

Take particular note of Step 9 which indicates that the evaluation should also consider safety practice as-observed directly with respect to the as-desired model. By performing this step, researchers can identify whether safety practitioners are overcoming deficiencies in the as-required practice in order to comply with safety practice as-desired (intentionally or otherwise). This ensures the evaluation identifies any subtle improvements made by those instantiating organisational safety engineering processes over the as-desired model. By considering all elements of safety practice (and the interrelationships between them) our process is able to identify, and mitigate the risk of unintentionally undermining safety practice.

Step 10 of Fig. 1 is in place as it is entirely possible that safety practitioners charged with implementing as-required practice may indirectly/directly appeal to the normative requirements and/or informative guidance from an Open Standard they are familiar with. This could be to recover perceived shortfalls in the as-required processes, or could simply be a default to a standard they know well.

Any identified differences, deviations, and impediments will enable an assessment of whether, and how, safety engineering guidance and/or practices need to change. The proposed process will identify and help mitigate these impediments in a way that is compatible with the work as-observed profile identified.

It is reasonable to argue that all of these elements of safety practice, once represented in text or graphically, are simply a form of "Work as Imagined" (Hollnagel, 2012). Whilst this is a potential weakness of the process, we assert that we need to transform each element of safety practice into comparable models that describe each element as accurately as possible, and in a manner that facilitates analysis.

We now consider each element of safety practice in turn, the white numbers in the black boxes refer to the numbered process Steps in Fig. 1.

3.1. Safety Engineering Practice As Desired 1

The first element of safety practice to model is that of safety practice as-desired — pertinent to the industry/application/technology specifics of the project in question. Because as-desired practice must be sector- and application-specific, we do not assert what as-desired safety practice *should* be, but demonstrate how it is to be modelled, then evaluated — enabling an assessment of widely-held safety practice as-desired itself. We provide an example of modelling an as-desired safety practice in Section 4.

It is important to note that any deficiencies in an as-desired model *may* be an oversight, but it may also be deliberate — reasonably relying on process at the as-required or as-observed level to remove them. As well as identifying any elements that must be specifically targeted to ensure gaps are closed in the as-desired level, modelling safety practice as-desired also facilitates an evaluation of compliance/conformance of safety practice as-required.

3.2. Safety practice as required 2 3

There are two elements of as-required safety practice which must be modelled. The first element is that represented in Open Standards such as SAE Aerospace (2010) or BSi (2010). Standards such as these prescribe a set of lifecycle activities that are argued to represent good practice. The second element is those practices described by organisational practice (Closed Standards), and these may, or may not have been designed as a means to implement the described lifecycle of a specific Open Standard. Our process is designed to allow both ways to be modelled, and any relationships between them to be evaluated.

3.3. Safety practice as observed 4

Safety practice as-observed is the model of safety work carried out by employees of an organisation. Safety practice as-observed *could* be interpreted as many different things, including:

- The assessment of intrinsic risks associated with everyday work practices (**the safety of work**)
- The analyses and methods undertaken to evaluate the safety of a design (**safety work**)
- The analyses, methods and monitoring required to assure the safety of a product in service (**safety work**).

Whilst our process is not restricted to a particular phase in a project's lifetime (from concept through to disposal), in this paper we define safety practice as-observed as being limited to 'safety work', and not the 'safety of work' (Provan et al., 2019). Our process enables safety practice as-observed to be modelled, and facilitates an evaluation of its relationships with safety practice as-required, and safety practice as-desired.

The first step of the process is to model the disparate elements that constitute safety practice.

3.4. A process for modelling safety practice

Our process requires the creation of models that are a faithful representation of the key elements of safety practice, but which are also as simple as possible. One of the weaknesses of many safety process representations is that although they portray the activities to be undertaken, they do not consider the attributes of activities such as timing constraints (Hawkins et al., 2013b), timing requirements (Vilela et al., 2017) commercial/contractual complexities (Squair, 2006; Reinhardt and McDermid, 2012), the resources required to undertake the activities (nor the attributes thereof), nor the intricate inter-relationships and inter-dependencies between activities. Our models of safety practice must therefore be capable of representing:

- Safety lifecycle activities
- Inputs to and outputs from activities
- Process interactions
- Relationships, dependencies, and constraints
- Methods or techniques that control an activity (such as an international standard that guides the conduct of a safety engineering activity)
- Resources expended (both people and materiel).

Further, to facilitate evaluation of safety practice, the selected representation must also be capable of:

- Representing the strength of links between activities
- Representing levels of agreement (e.g. between work as-required and work as-desired)
- Representing levels of compliance/conformance
- Representing identified deficiencies

- Describing optionality and multiplicity (such as that used in Goal Structuring Notation (SCSC, 2021)
- Quality attributes, including time (expressed as either a calendar date or phase in the lifecycle); personnel attributes including qualifications, experience, independence, authority, and role; and data format and contents.

The complex inter-relationships between activities, the required quality attributes of activities and the produced and consumed artefacts, coupled with a need to determine levels of compliance, agreements, and deficiencies, suggest that a textual representation would not be suitable as a representation, as it would require multiple cross-references, and could need to be re-read numerous times to decipher meaning, as Weaver discovered when considering text-based safety arguments (Weaver, 2003).

Owing to these complexities, a graphical representation is a critical enabler for understanding safety practice. Whilst we have to date used an adapted version of FRAM (Hollnagel, 2012) (see examples of using our process in Section 4), any suitable representation can be used, so long as it is capable of modelling the requisite attributes of safety practice as discussed above.

3.4.1. Modelling safety practice as desired 1

A project can determine and assert its safety practice as-desired by establishing, and expressing the objectives by which it will deliver a product which is safe to operate within a stated operating environment. Such objectives will likely be informed by the project's risk appetite, and legislative and regulatory obligations; and the need for a project to simultaneously minimise expenditure and maximise profitability.

This process to understand safety practice does not guarantee that a project's as-desired safety practice will be complete, nor correct; rather it provides a mechanism to assess its 'goodness' — as we will demonstrate.

3.4.2. Modelling safety practice as required 2 3

We suspect that safety practice as-observed differs from safety practice as-desired, and as-required, so we illustrate a process that can be used to create as-desired and as-required models (Fig. 2). Here, the lifecycle activities are modelled with reference to the as-required representation, or as-desired criteria, along with the six aspects of each activity that must be asserted. For activities we expect to model the following aspects which are taken from the ever-increasing existing body of knowledge on what aspects are required for processes to successfully complete (Hollnagel, 2012). We cannot yet argue these aspects are a complete list of those required, but take confidence from this ever-maturing body of knowledge:

- Resources consumed by the activity (which includes both human resource and materiel)
- Inputs to the activity
- Output from the activity
- Methods/techniques that can be used to carry out an activity
- Controls that constrain or define the activity
- Time by which the activity must take place.

As the lifecycle model progresses through activities, a level of abstraction is reached at which the inputs to, and/or outputs from an activity will not require further consideration; merely an article that is consumed or produced by an activity. Examples of this would be an International Standard that controls how an analysis such as a HAZOP is to be undertaken, or a Project Management Plan that is not under the control of the safety team.

The lowest level of abstraction that our representation will model is referred to as an artefact — a deliverable/item that supports/constrains an activity or is produced as the result of an activity. For artefacts we expect to model the following aspects:

- Human Resources
- Methods/techniques
- Data inputs (to activities)
- Data outputs (from activities).

Each of these require specific quality criteria and time constraints to be defined for them if we are to assert/assess safety practice as-required. The required quality criteria for aspects will differ – as illustrated by the examples given in Table 1 (noting that Time constraints are required for all – and therefore not included in this Table).

As can be seen in Table 1, the kinds of quality criteria that need to be defined for artefacts are wide-ranging, will differ across the types of artefact, and are not always required or relevant (for all artefacts). Generating an aspect for each would be cumbersome in modelling terms — with many not required for specific instantiations (i.e. 'Format' would not be required for a 'Human Resource' artefact type).

Whilst activities are represented with the inputs and outputs thereof, having both an input AND an output to an artefact may be superfluous in certain circumstances, as we may only need to represent the existence of the artefact. In other words, there is no modelling benefit in separately modelling the input (activity which created the artefact) from the output (activity to which the artefact contributes). To ensure we can adequately model the required aspects of all instantiations of an artefact, we therefore model the following attributes:

- Time or phase required by ('T')
- Quality Criteria ('Q')
- Existence (positive/negative) ('E').

It is possible that assessments of safety practice may reveal instances where the existence of artefacts are not explicitly stated by the lifecycle under analysis (e.g. in the example at Section 4). It is also possible that instances may be encountered where an artefact would have to be produced or updated for the lifecycle to achieve a satisfactory outcome, but where no producing activity is explicitly stated. To enable a modelling of internal consistency, our recommendation therefore is to categorise these as Inferred Activities. We have established three categories of modelled artefacts:

- **Explicit:** artefacts that are explicitly described, and have a consuming or producing activity that is clearly stated
- **Inferred:** artefacts that are discussed without any consideration of their creation or management, and with no consuming or producing activity that is explicitly stated. (e.g. if a standard says that "assumptions must be managed", we can infer the existence of an 'Assumptions' artefact)
- **Orphan:** artefacts that are explicitly described, but have no stated activity that produces them.

Inferred Activities and Artefacts, and Orphan Artefacts are denoted in our graphical representation using red colour-coding, and examples of this can be seen in our application example in Section 4.

The steps of the process for modelling safety practice as-desired and as-required are shown in Figs. 2 and 3 with the numbers in the accompanying text corresponding to the numbered steps in the Figures. The entire process is described in full at Appendix.

1. **Establish As Desired or As Required Practice**: This involves identifying standards/organisational practice and as-desired practice relevant to the safety engineering practice being evaluated.
2. **Model Activities**: By scrutinising the publication or criteria it is possible to identify all activities required by the lifecycle. Once these are modelled we can represent the aspects of each process activity:
    - Techniques/Methods: means by which an activity is fulfilled (e.g. carrying out a HAZID as a technique for carrying out 'Hazard Identification')

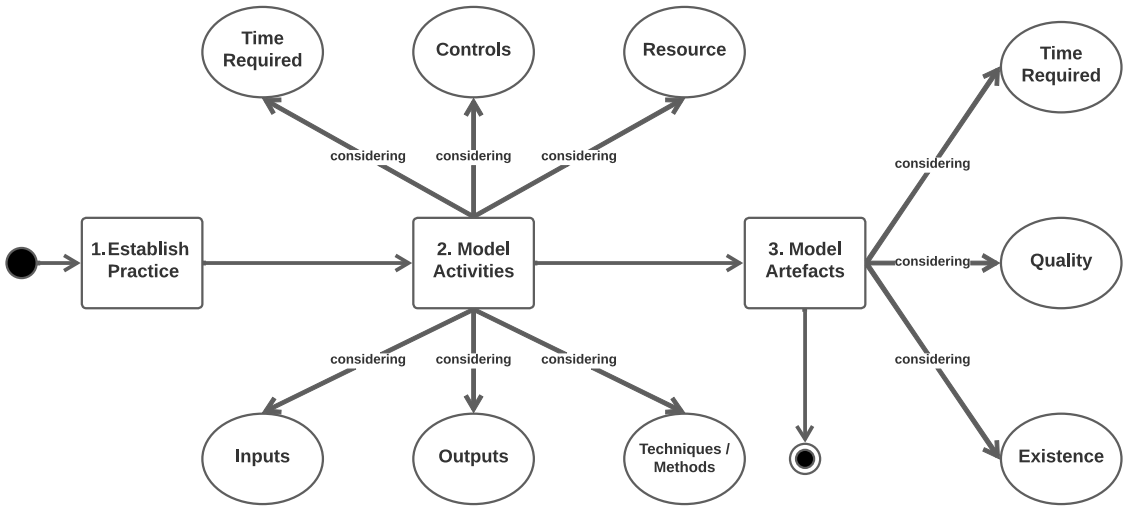

Fig. 2. Modelling safety practice.

Table 1
Aspects and quality criteria.

Aspect	Quality criteria	Examples
Human Resource	Qualifications	Safety Engineer with MSc
	Experience	5 Years experience
	Independence	Not involved in the design
	Authority	Authorised signatory
	Role	Software Safety Consultant
Method/technique	Relevance (to an activity)	ISO Standard for HAZOP
	Approved/Recommended status	Formally Issued at Rev C
Data (inputs/outputs)	Format	DOORS export in Excel
	Specific Contents	Requirements Specification Measures of Performance Maturity Level
	Generating Resource	Requirements Engineer
	Receiving Resource	Owner

- Inputs/Outputs: stipulated inputs to and outputs from each required activity
- Time: expressed as the point by which the activity should start and/or be complete by (or perhaps not start until). This may be expressed as a calendar date, a dependent activity, or phase in the program
- Controls: aspects that control how an activity is undertaken (e.g. a recognised Open Standard that controls how functional safety analysis is to be undertaken)
- Resources: person(nel) required to undertake the task, and any materiel required to complete it.

3. **Model Artefacts**: Artefacts represent the lowest level of abstraction, and are modelled as inputs to/outputs from an activity. Artefacts are deliverables or items that support/constrain an activity, or are produced as a result of an activity (e.g. an artefact describing the resource required in support of an activity, or a report produced as a result of an activity). To ensure all required aspects of an artefact are considered, artefacts have the following represented:

    - Time: expressed as the point by which the activity should start and/or complete by. This may be expressed as a calendar date or phase in the program
    - Quality Criteria: quality attributes required of an artefact, such as the skills and experience required of the person charged with carrying out an activity, or the format and contents of a report
    - Existence: does the artefact (yet) exist? This attribute is used to consider whether the artefact needs to be produced ahead of the supported activity (and therefore whether another activity should be modelled, or a dependency placed on a department other than Safety Engineering); or whether a person exists within the organisation who has the requisite skills or independence, for example.

The accuracy of the represented model determines the robustness of the resultant evaluation, and so it is vital that any model be as accurate and complete as possible as a representation of safety practice. Organisations may have competitive advantage, or security concerns that leads them to withhold parts of their safety lifecycle processes from the analyst or researcher. Such instances must be explicitly disclosed by the respondent, along with an assessment from the analyst as to whether the strength of any inferences made on the model(s) are impacted by the absence of such data.

3.4.3. Modelling safety practice as observed 4

Safety practice as-observed represents the observed practice of people performing safety activities. Rather than relying on a suite of documents, eliciting safety practice as-observed requires a form of ethnographic study (O'Reilly, 2012) of actual safety practice. To our knowledge however, such ethnographic studies on the work of safety practitioners (specifically safety engineers) are not carried out due to

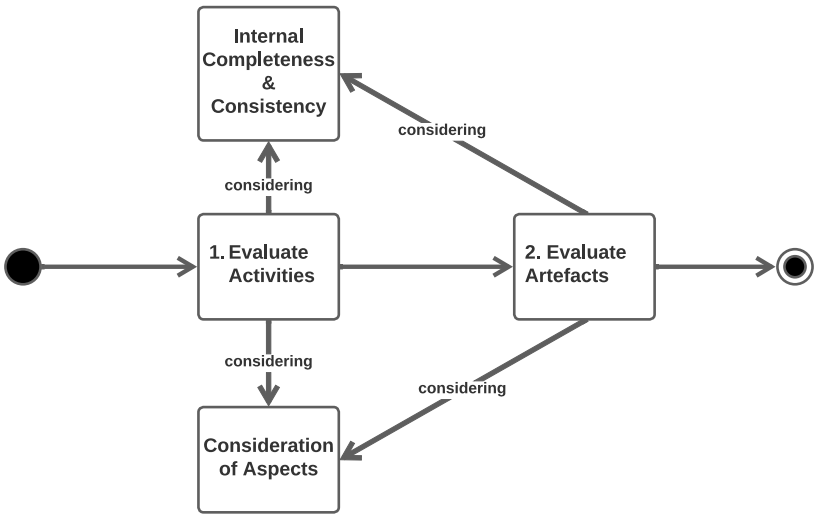

Fig. 3. Assessing safety practice.

the substantial time taken (dependent on the type of practice, the technology involved, and the length of the project/programme); the need for impartial and independent observers; and/or the cost of such an undertaking.

Should a full ethnographic study be infeasible, then an alternative method is to employ a series of interviews. Care must be taken if this approach is taken however, as this element of safety practice can no longer be considered as safety practice as-observed, and instead morphs into safety practice "as-disclosed" (Shorrock, 2020), and presents many opportunities for bias to skew the data (Somekh and Lewin, 2005; Yin, 2014; Marschan-Piekkari and Welch, 2004).

3.5. Assessing safety practice

Having modelled the relevant aspects of safety practice, an assessment is performed to identify deficiencies in practice, and impediments to the adoption of safety practice as-desired. The assessment process is illustrated in Fig. 3. During the assessment stage, we also make use of simple colour-coding of the as-desired, as-required, and as-observed models (with appropriate definitions for red/amber/green as highlighted below).

3.5.1. Assessing the internal completeness and internal consistency of safety engineering practice as required [2] [3]

The assessment has the following two steps:

1. **Evaluate Activities**. Each activity in the lifecycle of required practice is assessed for:
    - Completeness and Consistency: are there enough activities commensurate with achieving the required outcome; and does each activity have sufficient supporting sub-activities to ensure it can be completed to a sufficient level?

    For example, to produce the artefact 'Preliminary System Safety Assessment', does the modelled safety practice identify all activities that are reasonably required to produce it? Examples here would be 'Review Preliminary Platform Safety Assessment', 'Review System Safety Requirements' and perhaps 'Carry out Fault Tree Analysis'. An analyst with an understanding of the activities required to produce such a safety assessment will be able to assess whether it is reasonable that the safety assessment will be produced. Although this is perhaps a subjective assessment, it is an early identifier as to the sufficiency of the safety practice under review.

    - Consideration of Aspects: is there sufficient detail in the description of the aspects required of/produced by the identified activities to have confidence that sufficient consideration is given to Inputs, Outputs, Time, Techniques and Methods, Controls, and Resources?

    For example, the activity 'Review System Safety Requirements' should have aspects that denote the time it should be completed by; a defined set of inputs and outputs; the techniques and methods by which the Safety Requirements are to be reviewed; any controlling procedures (such as a Process Instruction); and the resources expended by the activity.
2. **Evaluate Artefacts**: Each artefact is assessed for:
    - Completeness and Consistency: are there enough artefacts to enable successful completion of all activities; does every activity produce an artefact; and does each activity have an adequate set of required inputs and outputs?

    Many Open Standards do not explicitly document the specific artefacts that are consumed by an activity, but merely infer them (such as stating that the user should "check the validity of any assumptions", without ever referencing where such assumptions are derived from, nor what should be the arbiter of 'validity'. Open Standards also often fail to consider who needs an artefact, and in what format. This may not be an inadequacy on the part of the standard, however — which may reasonably rely on those implementing the standard's processes to consider such details. Whether this is a deliberate and reasonable assumption will be revealed through the modelling of organisational processes, and observations of practice.

    - Consideration of Aspects: is there sufficient detail of the aspects to denote when they need to be produced or used; are sufficient quality attributes considered; and does the artefact yet exist, or does it need to be created (in which case further producing/consuming activities may need to be modelled)?

    Whilst standards and organisational process denote the production of activities, they do not always specify when the artefact is needed. For aspects that are required inputs to activities, there is not always consideration of when the artefact is needed by, who produces it, nor to what quality (such as format). On the occasion that resources are mentioned, the quality attributes of the resource are not stipulated (such as qualifications, training, experience, and independence). Future analysis of organisational processes and Open Standards

(subsequent steps in the process) will uncover whether this is an acceptable omission from an as-required perspective, and follow-up interviews with practitioners/management will reveal whether such potential shortcomings are overcome — intentionally or otherwise.

3.5.2. Evaluating safety engineering practice as required 5 6 8

We now need to compare the relationships between safety practice as-required with safety practice as-desired, so an evaluation activity is undertaken to consider the levels of compliance between the as-required model and safety practice as-desired (Steps 5 and 6 of Fig. 1), as well as (where applicable) between the organisation's internal as-required model and the model of the applicable Open Standard (Step 8 of Fig. 1). When evaluating the levels of compliance, we recommend the adoption of the following colour-coding scheme:

- **GREEN**: The link between the activities/sub-activities, or the considered aspects of the activity meets the claims required of/agrees with the comparison model in full
- **AMBER**: The link between the activities or sub-activities, or the considered aspects of the activity only partially meets claims required of/partially agrees with the comparison model
- **RED**: The link between the activities or sub-activities, or the considered aspects of the activity meet no aspect of the claims required of/does not agree with any of the comparison model.

3.5.3. Evaluating safety practice as observed 7 9 10

Once the analyst has evaluated as-required safety practice, they can then evaluate safety practice as-observed. Primarily, they should do this in terms of its relationships with safety practice as-required and as-desired. They can, again, use colour-coding to represent the level of consistency between as-required (Closed *and* Open) and as-desired practice.

- **GREEN**: The link between the activities/sub-activities, or the considered aspects of the activity agrees with that of safety engineering practice as-required/as-desired in full
- **AMBER**: The link between the activities or sub-activities, or the considered aspects of the activity partially agrees with that of safety engineering practice as-required/as-desired
- **RED**: The link between the activities or sub-activities, or the considered aspects of the activity does not agree with any part of safety engineering practice as-required/as-desired.

Once the analyst has evaluated the as-observed practice in this way, they can then hold follow-up interviews with representatives of the organisation whose processes are under analysis in the context of the respondent's industry sector, organisational hierarchy, and the product being created by the respondent's organisation. These follow-up interviews should aim to identify:

- The reasons for limited areas of agreement
- The reasons why there are areas of no agreement
- Evidence regarding the validity of the as-desired model (is the as-desired model complete and correct?)
- Any difficulties or complexities behind the areas of limited or no agreement
- Whether and why any areas of limited/no agreement contribute to meeting any shortfalls with an organisation's processes with respect to an Open Standard (Step 8 of Fig. 1)
- Whether and why any areas of limited/no agreement contribute to meeting any shortfalls with an organisation's processes and that of the as-desired model (Step 9 of Fig. 1)

As a result of the follow-up interviews, the analyst must determine whether and how the models and subsequent evaluations are challenged, and whether they need to be modified as a result of new information.

3.6. Potential outcomes

Evaluation of safety practice may reveal elements of practice which are in agreement or compliant with each other, or some areas of discrepancies between the different elements of safety practice. There are many potential reasons for any equivalence or discrepancy, and we now discuss some of these potential outcomes and reasons.

Organisational practice improves on practice required by an Open Standard, or as-desired practice. Should the as-required safety practice of an organisation's lifecycle improve on safety practice expressed in Open Standards, or the representation of as-desired safety practice, this may suggest that the practices required of Open Standards or the as-desired model are a cause of impediments to good safety practice. Alternatively, perhaps the practitioner is implicitly or explicitly aware of the shortcomings of such standards and has evolved local processes in isolation of the standards. In such cases one may consider research into improving the as-required (Closed) practice and/or use this as a mechanism to research potential amendments to the practices extolled in as-required (Open) practice.

Organisational practice is deficient when compared to the as-desired model. This may indicate that issues with safety practice manifest in the interpretation of safety practice as-desired into organisation-described processes. Targeted interviews with the organisation may indicate where the issues lie.

Internal Inconsistency. It may be revealed that safety practice as-required (by Open and/or Closed Standards) is inconsistent, preventing the safety practice as-required from ever being adopted as safety practice as-observed. Should this be revealed, it must be highlighted to the organisation as part of follow-up interviews, and/or the relevant standard's committee should be notified.

Safety practice as-observed is equivalent to that stipulated in organisational processes. As the evaluation process moves from the as-required safety practice to safety practice as-observed, safety practice as-observed may be equivalent to as-required (Closed) practice. Assuming that there is at least some equivalency of the organisation's safety practice with the as-desired model of safety practice, equivalency between organisational practice and safety practice as-observed will suggest that organisational processes are being fully implemented. However, follow-up interviews as part of evaluating safety practice as-observed may identify difficulties in implementing the organisation's processes, and the existence of such difficulties may suggest issues with the organisational processes exist.

Safety practice as-observed improves on organisational practice. Should safety practice as-observed improve on the as-required representation of organisational safety practice, this may suggest those charged with implementing the organisation's processes are aware of the limitations, inefficiencies, inaccuracies, or unrealistic expectations of their organisation's processes and have adopted approaches to compensate. There may even be elements of core and discretionary work (perhaps owing to any deliberate vagaries of processes, and hence by recourse to what a collective of engineers engaged with them believe is required (Asplund et al., 2020), assumptions made by the practitioner, or tensions arising through power relationships (Rae et al., 2020) that require investigation. Through targeted follow-up questions as part of Steps 7 and 10 of Fig. 1, it may be possible to identify any impediments or difficulties that have led to a circumvention of process; and ultimately characterise and suggest mitigation research accordingly.

4. Application of part of our process

Here we use an example of applying our process which highlights the importance of the inter-relationships between the elements which constitute safety practice. The example in question concerns the development of software in a safety-critical system of an aircraft that is being developed for the UK Military.

4.1. How we applied the process to this example

We have carried out a complete application of the process to understand safety practice (in support of our ongoing research into improving software safety assurance practice) and here we describe part of that application in detail.

The project against which we applied the process involved the supply of software which is used in a safety-critical application for the UK Military Air Domain. The project selected the 4+1 Principles as the basis of establishing the as-desired model of software safety practice predicated on the assertion in Hawkins et al. (2013a) that they are *"constant across domains and across projects, and can be regarded as the immutable core of any software safety justification"*. The use of these principles was further justified by the project as they continue to be widely adopted for use in functional safety — including their incorporation as the overarching principles and objectives of UK Defence Standards (MoD, 2016, 2017).

Whilst the principles were argued to be appropriate, they needed to be stated in a manner which facilitated a measurement of compliance against them. As such a set of criteria was created. These criteria represent Step 1 of the process to understand safety practice.

The first principle of the 4+1 Principles requires that *"software safety requirements shall be defined to address the software contribution to system hazards"*. The following measurable criteria was established and, if met, would ensure the 4+1 Principles of Software Safety Assurance (Hawkins and Kelly, 2012) are complied with:

- A clear description of the software in the system will be provided
- The operating context of the system in which the software resides will be described
- A clear description of the system in which the software resides will be provided
- The system hazards to which software may contribute will be identified
- The specific failure modes by which software contributes to the identified system hazards will be described
- The software contribution to the identified system hazards will be acceptably managed through the elicitation of software safety requirements that specify the required behaviour(s); for each identified software contribution, to each system hazard
- All software safety requirements will be atomic, unambiguous, defined in sufficient detail, and verifiable.

Software safety practice as-required (Open) was held by the project to be the ARP 4754A suite of publications (SAE Aerospace, 2010) (Step 2), and this was compared to the as-desired criteria (Step 6).

Applying our process in Fig. 1 to the UK Military example required us to carry out the following steps:

1. **Model As-Desired Practice:** For safety critical software this is modelled as a set of criteria which if met, will meet the 4 + 1 software safety assurance principles.
2. **Model As-Required (Open) Practice:** In this case the suite of standards described in ARP 4754A (SAE Aerospace, 2010).
3. **Model As-Required (Closed) Practice:** The safety engineering processes defined by the organisation who is developing the safety-critical software.
4. **Model As-Observed Practice:** How the safety engineers and software engineers developing the safety-critical software are observed to perform their safety activities.
5. **Compare As-Required (Closed) Practice with As-Desired Practice:** Assess the extent to which the organisation's processes comply with the 4 + 1 Principles.
6. **Compare As-Required (Open) Practice with As-Desired Practice:** Evaluate the extent by which the open standard meets the 4 + 1 Principles.
7. **Compare As-Observed Practice with As-Required (Closed) Practice:** This step is an evaluation of how well safety practice as-observed compares with that required by the organisation's defined safety process.
8. **Compare As-Required (Closed) Practice with As-Required (Open) Practice:** A comparison between the organisation's defined safety engineering process, and that required by the Open Standard (SAE Aerospace, 2010).
9. **Compare As-Observed Practice with As-Desired Practice:** This step makes it possible to identify whether the safety engineers and software engineers are overcoming perceived deficiencies in the as-required practice in order to meet the as-desired model (intentionally or otherwise). Should there be any deviation between as-observed and as-desired practice, the reasons for the deviations must logically manifest in the relationship between safety practice as-observed and that required by the Open Standard, and/or organisational process.
10. **Compare As-Observed Practice with As-Required (Open) Practice:** This step is in place as it is entirely possible that software safety practitioners charged with implementing as-required practice may indirectly/directly appeal to the normative requirements and/or informative guidance from an Open Standard they are familiar with.

Because we have created realistic models of the elements which constitute safety practice, and have undertaken an analysis of the relationships which exist between each model, we can highlight any impediments or deficiencies to implementing as-desired safety practice — as we now discuss.

In Section 3.4 we noted that our analysis needs to use a graphical notation — here, we have used an adapted version of FRAM (Hollnagel, 2012). We chose FRAM for the graphical notation because it is capable of easily representing the activities that should be carried out and the considerations for each activity to succeed. FRAM also enables the identification of agreements and disagreements between models of safety practice in a simple manner, and the representation of complex empirical data in a manner that is easy to digest and comprehend.

4.2. Results

Here we provide details of a single criterion we evaluated as part of the research, namely the provision of a **clear definition of software within the system** (required by Principle 1 of the 4 + 1 Principles' criteria). The accompanying Figs. 4 and 5 represent only a snapshot of the full model of Principle 1 (which can be found at Osborne (2021), yet provide clear examples of potential deficiencies of the as-required safety practice asserted by Open Standards. Through the evaluation of both organisational processes designed to meet ARP 4754A, and the evaluation of safety practice as-observed (both subsequent steps in the process to understand safety practice) we can confirm whether these highlighted deficiencies of ARP 4754A are reasonably left to organisational processes and/or safety practice as-observed to recover them.

When we modelled ARP 4754A, it very quickly became apparent that no individual artefact is created by the ARP 4754A lifecycle which contains a clear definition of the software within a system. But, rather than dismiss the standard as being wholly deficient against this criterion, we identified a number of artefacts that are produced by the ARP 4754A lifecycle that *could* reasonably be expected to contain a clear definition of some, or all of the software in the system (either in a single artefact or collectively across them). The identified ARP 4754A artefacts were:

- 'System Architectures'
- 'System Description and Environment'
- 'Software Design Details'

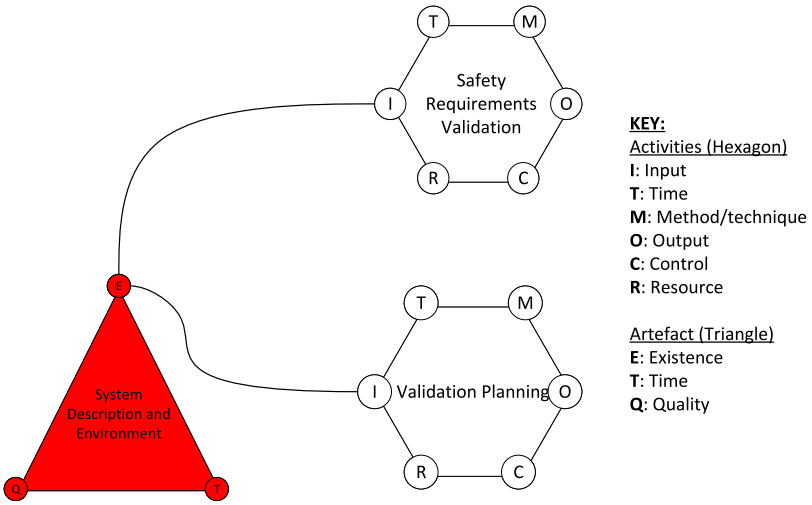

Fig. 4. Inferred artefacts.

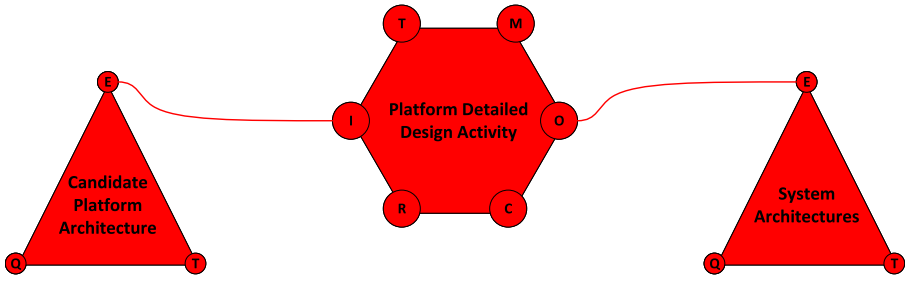

Fig. 5. Inferred activity.

- 'Software Load Control Records'
- (Software) 'Loading Data'
- 'Software Configuration Management Records'
- 'Software Configuration Index'
- 'Candidate Platform Architecture'.

The modelling and analysis undertaken revealed issues with these artefacts. First, we discovered that these artefacts, and the activities that produce/consume them, are not explicitly defined by the ARP 4754A lifecycle. For example, the artefacts 'System Description and Environment', and 'Software Design Details', which may reasonably be considered the most obvious artefacts to contain a clear description of the software, are in fact 'orphan' artefacts (i.e. there is no identified activity that produces them).

Second, we discovered a number of the listed artefacts are inferred (artefacts which are not explicitly consumed or produced by stated activities in the ARP 4754A lifecycle):

- 'System Description and Environment' is an inferred artefact discussed at the 'Platform' level as an input to the activities of 'Validation Planning' and 'Safety Requirements Validation' (Fig. 4). When discussing the 'Validation Process Model' at Section 5.4.2, ARP 4754A notes that "*inputs to the validation process may include a description of the system (including the operating environment)*". The ARP does not define when this data is created, nor by what activity, however.
- 'Software Design Details' is also an inferred artefact — discussed at the 'Platform' level as an input to the activity 'Safety Requirements Verification'. Section 2.5.6 of DO-178C (RTCA, 2011) notes that "*Software design details that relate to the system functionality need to be made available to aid system verification*", yet the nature of these details is not expanded on, nor is the activity that produces them stated.
- 'System Architectures' is an artefact discussed by the (also inferred) activities 'Platform Detailed Design Activity' and 'System Detailed Design Activity' (Fig. 5). Section 4.6 of ARP 4754A notes that "*candidate system architectures are derived from the activity 'System Requirements Identification' which are iteratively and recursively evaluated using the PSA* [Platform Safety Assessment], *PSSA* [Preliminary System Safety Assessment], *and CCA* [Common Cause Analysis] *processes in order to establish their feasibility in meeting the requirements*". At some stage, each candidate architecture must be formally endorsed through design decision(s) (if they are not we can never be confident that we have instantiated the optimal design solution); yet such an activity is not considered by the ARP. Taking this example of 'Candidate System Architecture' it is not acceptable for a completed product to be accepted into service with a design that was predicated on candidate architectures. The ARP does not consider how the design maturity of such artefacts is finalised, however. Nor does the ARP consider when this should be done, by whom, nor to what quality criteria.

The ARP discusses these inferred activities in a manner that suggests a significant contribution to safety. This is a concern, as they are not described by the ARP at a level of detail that ensures they will be created and delivered on time with the required level of quality.

4.3. Summary

Here we summarise and discuss the ability of the ARP to meet Principle 1 of the 4+1 Principles:

- Principle 1 is not met by the ARP. Only one criterion can be argued to be met (the description of the intended operating environment), the remainder being non-compliant, or partially compliant

- There are artefacts that one would expect (through both experience and by recourse to the wider safety science literature) to be created or contributed to in support of functional safety management which are not created by this lifecycle. Examples include:
  - A Software Design Description
  - A Functional Requirements Specification
  - A list of System Hazards
  - An Operational Requirements Specification
  - A Hazard Log.
- Artefacts and activities are inferred to take place (and are therefore modelled to achieve internal consistency), but are not considered explicitly by the ARP. Examples include:
  - System Description and Environment
  - System Architecture (only candidate system architectures are mentioned by the ARP)
  - Item Architecture (only candidate item architectures are mentioned by the ARP)
  - System Detailed Design Activity.
- There exists no clear mapping between hazards and safety requirements at any level of design abstraction
- Some criteria required for the start of the design process are only (partially) met at the assurance or certification phase of the lifecycle.

Whilst we discovered principled omissions from our evaluation of whether ARP 4754A meets the 4+1 Principles, we also discovered instances of *potentially problematic* omissions. Such problematic omissions may or may not be reasonably be left to other elements of safety practice to recover. Examples include:

- No consideration is given to the resources (human-power) that are consumed by activities. This may be a deliberate decision to rely on individual organisations, and our process is designed to investigate this as part of Step 5
- No consideration is given to when activities should be commenced nor completed, nor to when artefacts are required to be produced/provided. This is not considered in terms of calendar, phase, nor stage. This may be a deliberate decision to rely on individual organisations to define this, and this is investigated as part of Step 5
- The methodology by which a safety activity should be undertaken is not considered — either as a mandatory method, or a suite of options (presented with selection considerations). Although the ARP may rely on sister publications such as ARP 4761, there exists a risk that without explicit appeal to artefacts which constrain safety-related activities, that organisations will use processes which may be sub-optimal. Whether an organisation specifies the methods or techniques to be used is investigated as part of Steps 3 and 8.

5. Conclusions

In this paper, we have argued the importance of understanding safety practice, and of establishing the true causes of poor safety practice before applying interventions which are not based on theory or evidence. We have also argued that establishing the true causes of poor practice will allow effective interventions to be made. Based on this, we have therefore proposed a process for modelling and assessing safety practice based around three distinct elements of safety practice as-desired, as-required, and as-observed, and the evaluation of the inter-relationships between these elements.

The adoption of our process will not preclude the use of frameworks such as ‘Work as Imagined’ versus ‘Work as Done’ (to rectify identified issues with ‘Work as Described’, or ‘Work as Done’), nor the use of complimentary activities such as checklists (to assure alignment between ‘Work As Described’ and ‘Work As Done’). However, these and other existing approaches may be at the detriment other potential interventions (such as a change to a process rather than relying on a compliance check against the existing process), and form only part of a wider process for improving safety practice. We provide such a process, which can also be employed to explore whether any interventions could manifest in negative, and unanticipated side-effects.

Modelling any element of safety practice requires the analyst to compile an accurate and complete representation of the data, and a reliable, impartial interpretation of the practice being represented. Threats to the completeness of data may be caused by a lack of access to it (such as commercially sensitive information contained in an organisation’s process descriptions). In such cases, the organisation needs to declare that information has been withheld, and provide an assessment as to whether the strength of any inferences made on the model(s) are impacted by the absence of such data.

5.1. Future work

We invite safety engineers and managers in any project involved in the design, manufacture, and analysis of safety-related products to use our process to understand and assess their own safety practice. The step-by-step process instructions have not been included in the main body of the paper for brevity, but are included in Appendix.

Having applied the process, we seek feedback on how the process could be improved or tailored to specific industrial or technological application.

For our own part, we are currently applying the process presented in this paper to undertake a full evaluation of the practice of software safety assurance in the military domain. In this paper we have shown a very small snapshot of the evaluation of an as-required model with respect to an as-desired model (Steps 1, 2 and 6 of the process). Our ongoing work has applied the remaining steps of the process to investigate the same application domain through interaction with practitioners from a number of representatives of the organisation developing safety critical software for military aircraft. We will make the results of this evaluation available in due course. There are of course many different safety disciplines to which this process can be applied to evaluate safety practices, and we too plan to apply the process in other domains to assess its applicability. We encourage safety researchers to do the same, and encourage safety practitioners and safety managers to use it as a tool for measuring the effectiveness of their own safety practice.

CRediT authorship contribution statement

Matt Osborne: Writing – original draft, Visualization, Methodology, Conceptualization. **Richard Hawkins:** Writing – review & editing, Supervision. **Mark Nicholson:** Writing – review & editing, Supervision. **Rob Alexander:** Writing – review & editing.

Declaration of competing interest

Declarations of Interest: None.

Appendix. A framework and process for understanding safety practice

A.1. Safety engineering practice

The main elements of safety practice, and their relationships are shown in Fig. 6 — with each number representing an activity within the process to understand safety practice.

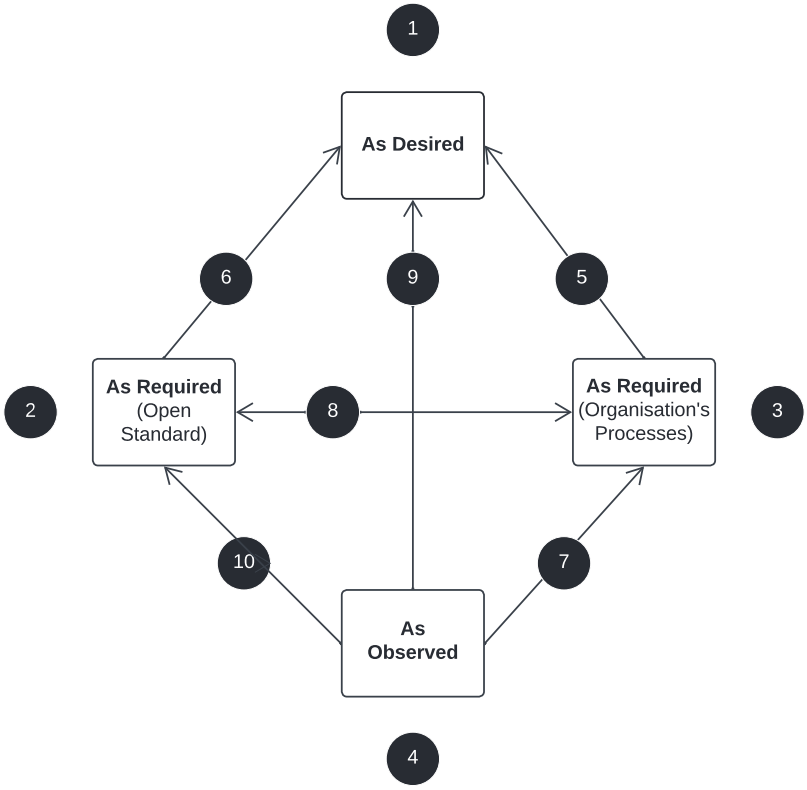

Fig. 6. The elements of safety engineering practice.

A.1.1. The framework:

1. As-desired safety practice model
2. Safety practice as represented by an Open Standard
3. An organisation's safety engineering processes (as-required)
4. Safety practice as carried out (as-observed)
5. An organisation's safety engineering processes (as-required)
6. Degree of conformance between the Open Standard with the safety practice as-desired
7. Degree of conformance between the safety practice (as-observed) and the safety engineering practice (as-required)
8. Degree of conformance between the organisation's safety engineering lifecycle (as-required) and the Open Standard which informed/influenced its development
9. Degree of conformance between the safety practice (as-observed) with the as-desired model
10. Degree of conformance between the safety practice (as-observed) with safety practice (as-required).

To instantiate this framework the following activities must be undertaken:

1. Representation of the as-desired model
2. Representation of an Open Standard using the selected notation
3. Representation of an organisations' safety engineering processes (as-required) using the selected notation
4. Representation of safety practice as carried out (as-observed) using the selected notation
5. Comparison of an organisation's safety engineering processes (as-required) with the as-desired model
6. Comparison of an Open Standard with the as-desired model
7. Comparison of safety practice (as-observed) with the safety engineering processes formulated by an organisations' lifecycle (as-required)
8. Comparison between the organisations' safety engineering lifecycle (as-required) and the Open Standard which informed/influenced its development
9. Comparison of safety practice (as-observed) with the as-desired model
10. Comparison of safety practice (as-observed) with an Open Standard.

A.1.1.1. Identify the activities and articles employed. To instantiate the framework effectively, the process models:

- The activities required of a safety engineering lifecycle
- The required inputs to, and outputs from each activity in a safety engineering lifecycle
- The interactions required throughout a safety engineering lifecycle
- The relationships, dependencies, and constraints
- Methods or techniques that control an activity (such as an international standard that guides the conduct of a safety engineering activity)
- The resources required for each activity (both people and materiel).

To ensure widespread use, the selected representation is ready for use with minimal adaptation, capable of use without the need for proprietary software, saveable in a portable format, capable of construction and analysis in the absence of formal modelling knowledge, understandable and interpretable in the absence of prior ontological knowledge/experience, and capable of construction in the absence of complex background databases.

Appendix B. The process

Each Step of the process has an Objective, Inputs, a Task, and its Outputs. Each step requires a Suitably Qualified and Experienced Person (SQEP) to undertake the process. However, what constitutes a SQEP individual is outside of the scope of this process, however.

B.1. STEP 1: Model safety practice as desired

OBJECTIVE: Define safety practice as desired for the organisation wishing to understand and assess their safety practice.

INPUTS: An organisation wishing to model Safety Practice as-desired requires the following inputs:

- A Safety Philosophy
- A Risk (acceptance and tolerance) Policy
- A Safety Management Philosophy (which may be instantiated as a Management System and Plan(s))
- A suitably qualified and experienced Safety Manager to determine, represent and agree with the product owners the model of as-desired practice.

TASKS:

1. Define Safety Practice As Desired
2. Create a tangible and measurable representation of Safety Practice As Desired.

Notwithstanding the complexities and challenges of this task, the organisation must create a representation of as-desired practice which is both tangible and measurable.

OUTPUTS: The output of this step is either a set of objectives, or a set of measurable criteria which can then be used to assess the other elements of safety practice for compliance.

B.2. STEP 2: Model safety practice as required (open)

OBJECTIVE: There are two ways in which safety practice 'as-required' is currently represented in industrial practice. The first way as-required practice is that represented as an Open Standard such as ARP 4754A or BS EN 61508. Standards such as these prescribe a set of lifecycle activities that are argued by their developing committees to represent good practice. The objective here therefore, is to employ a process to model the set of lifecycle activities described by the relevant Open Standard which may have influenced the development of organisational practice.

INPUTS: The single input for this task is an Open Standard which may have influenced the creation of organisational practice.

TASKS:

1. Identify all the activities and produced/consumed documents and other articles required by the standard
2. Identify the sequences of linked activities and articles
3. Represent graphically the lifecycle required by the standard as a sequence of linked activities and articles
4. Compile a report which defines the:
   Modelling process used
   Modelling symbology used
   Location of the model, and any proprietary software required to access it.

OUTPUTS: Two Outputs are created by this Step:

1. The appropriately represented As-Required (Open) Model
2. The Report accompanying the As-Required Model.

B.3. STEP 3: Modelling safety practice as required (closed)

OBJECTIVE: The second way in which safety practice 'as-required' is currently defined is by one generated by a specific organisation — or its 'Closed' Standard.

The objective here therefore, is to model the set of lifecycle activities described by the organisational processes and procedures.

INPUTS: The single input for this task is the Closed Standard which constitutes organisational practice.

TASKS:

1. Identify all the activities and produced/consumed documents and other articles required by organisational practice
2. Identify the sequences of linked activities and articles
3. Represent graphically the lifecycle required by the standard as a sequence of linked activities and articles
4. Compile a report which defines the:
   Modelling process used
   Modelling symbology used
   Location of the model, and any proprietary software required to access it.

OUTPUTS: Two Outputs are created by this Step:

1. The appropriately represented As-Required (Closed) Model
2. The Report accompanying the As-Required (Closed) Model.

B.4. STEP 4: Model safety practice as observed

OBJECTIVE: Safety practice 'as observed' represents the actual activities of those practitioners within an organisation. Instead of relying on a suite of documentary articles, safety engineering practice as-observed necessitates a form of independent ethnographic study. The objective here therefore, is to model the safety activities carried out by safety practitioners in a given organisation.

INPUTS: The single input to this task is an empirical report of as-observed safety practice.

TASKS:

1. Identify all the activities and produced/consumed documents and other articles carried out by the safety practitioner
2. Identify the sequence of activities carried out by the safety practitioner
3. Represent graphically the sequence of linked activities carried out and articles produced/consumed
4. Compile a report which defines the:
   Modelling process used
   Modelling symbology used
   Location of the model, and any proprietary software required to access it.

OUTPUTS: Two Outputs are created by this Step:

1. The appropriately represented As-Observed Model
2. The Report accompanying the As-Observed Model.

Having identified, modelled and represented the elements of safety practice, attention now turns to the process to assess safety practice.

B.5. STEP 5: Compare organisational practice with safety practice as-desired

OBJECTIVE: Organisational Practice must be capable demonstrably of complying with Safety Practice as-desired. The objective of this step is therefore to assess the levels of compliance between organisational practice and safety practice as-desired.

INPUTS: Two completed modelling elements of the framework instantiation process:

1. The As-Required (Closed) Model of Safety Practice
2. The As-Desired Model of Safety Practice.

TASKS:

1. Create a copy of the model of As-Required (Closed) practice created at Step Three

2. Using the newly-created model, create a representation of how as-required practice conforms with each subset/criteria of as-desired safety practice in turn
3. Taking each subset/criteria in turn, evaluate each contributing activity:

Internal Completeness and Consistency: are the activities correct and pertinent commensurate with achieving as-desired practice? Do the right amount of activities exist; and does each activity have the correct amount of supporting contributing activities to ensure it can be completed to the required level of compliance?

Consideration of Attributes: is the information stated for the attributes the correct information (i.e. Inputs, Outputs, Time, Techniques and Methods, Controls, and Resources); and is the correct amount of information given for the attributes for the as-desired practice to be met?

4. For each subset/criteria, evaluate each article which is produced/consumed by an activity:

Sufficiency: are there the correct amount of articles to enable successful completion of all activities, and are the articles the correct ones? Does every activity produce an article; and does each activity have the correct amount and type of articles (as inputs) to comply with the model of as-desired practice?

Consideration of Attributes: is the information stated for the attributes the correct information (i.e. Time, Quality Criteria and Existence) to denote when they need to be produced or used? Are the correct amount of quality attributes considered for each article, and are they the correct quality attributes for as-desired practice to be complied with?

5. Annotate the newly-created model denoting the levels of compliance with as-desired safety practice
6. Should potential deficiencies be evident, then follow-up research with the organisation should be undertaken to establish the reasons why
7. Compile a report which defines the:
    Modelling process used
    Modelling symbology used
    Location of the model, and any proprietary software required to access it.

OUTPUTS: Two Outputs are created by this Step:

1. The appropriately represented Model of As-Required (Closed) Practice Compliance
2. The Report accompanying the Model of As-Required (Closed) Practice Compliance.

B.6. STEP 6: Compare the open standard with safety practice as-desired

OBJECTIVE: A published Open Standard must be capable demonstrably of complying with Safety Practice as-desired. The objective of this step is therefore to assess the levels of compliance between an Open Standard and safety practice as-desired.

INPUTS: Two completed modelling elements of the framework instantiation process:

1. The As-Required (Open) Model of Safety Practice
2. The As-Desired Model of Safety Practice.

TASKS:

1. Create a copy of the model of As-Required (Open) practice created at Step Two
2. Using the newly-created model, create a representation of how as-required practice conforms with each subset/criteria of as-desired safety practice in turn
3. Taking each subset/criteria in turn, evaluate each contributing activity:

Internal Completeness and Consistency: are the activities correct and pertinent commensurate with achieving as-desired practice? Do the right amount of activities exist; and does each activity have the correct amount of supporting contributing activities to ensure it can be completed to the required level of compliance?

Consideration of Attributes: is the information stated for the attributes the correct information (i.e. Inputs, Outputs, Time, Techniques and Methods, Controls, and Resources); and is the correct amount of information given for the attributes for the as-desired practice to be met?

4. For each subset/criteria, evaluate each article which is produced/consumed by an activity:

Sufficiency: is there the correct amount of articles to enable successful completion of all activities, and are the articles the correct ones? Does every activity produce an article; and does each activity have the correct amount and type of articles (as inputs) to comply with the model of as-desired practice?

Consideration of Attributes: is the information stated for the attributes the correct information (i.e. Time, Quality Criteria and Existence) to denote when they need to be produced or used? Are the correct amount of quality attributes considered for each article, and are they the correct quality attributes for as-desired practice to be complied with?

5. Annotate the newly-created model denoting the levels of compliance with as-desired safety practice
6. Should potential deficiencies be evident, then follow-up research with the organisation should be undertaken to establish the reasons why
7. Compile a report which defines the:
    Modelling process used
    Modelling symbology used
    Location of the model, and any proprietary software required to access it.

OUTPUTS: Two Outputs are created by this Step:

1. The appropriately represented Model of As-Required (Open) Practice Compliance
2. The Report accompanying the Model of As-Required (Open) Practice Compliance.

B.7. STEP 7: Compare as observed practice with as required (closed) practice

OBJECTIVE: Safety Practice As-Observed may be different to, or the same as Safety Practice As-Required (Closed). The objective of this step is therefore to compare as-observed practice with the lifecycle of organisational practice.

INPUTS: Two completed modelling elements of the framework instantiation process:

1. The As-Required (Closed) Model of Safety Practice
2. The As-Observed Model of Safety Practice.

TASKS:

1. Create a copy of the model of As-Observed Practice created at Step Four
2. Using the newly created model, compare the levels of agreement between safety practice as-observed, and safety practice as-required (closed)

3. Annotate the newly created model denoting the levels of agreement between the two models of practice
4. Should differences be evident, then follow-up research with the organisation should be undertaken to establish the reasons why
5. Compile a report which defines the:
   Modelling process used
   Modelling symbology used
   Location of the model, and any proprietary software required to access it.

OUTPUTS: Two Outputs are created by this Step:

1. The appropriately represented Model of how As-Observed Practice compares with As-Required (Closed) Practice
2. The Report accompanying the comparison of how As-Observed Practice compares with As-Required (Closed) Practice.

B.8. STEP 8: Compare as required (closed) practice with as required (open) practice

OBJECTIVE: Safety Practice As-Required (Closed) may be different to, or the same as the Open Standard which may have informed its development (Safety Practice As-Required (Closed)). The objective of this step is therefore to compare both models of safety practice as-required.

INPUTS: For this Step to proceed, two modelling elements of the framework instantiation process must have already been completed:

1. The As-Required (Closed) Model of Safety Practice
2. The As-Required (Open) Model of Safety Practice

TASKS:

1. Create a copy of the model of As-Required (Closed) Practice created at Step Three
2. Using the newly-created model, compare the levels of agreement between safety practice as-required (Closed), and safety practice as-required (Open)
3. Annotate the newly-created model denoting the levels of agreement between the two models of as-required practice
4. Should differences be evident, then follow-up research with the organisation should be undertaken to establish the reasons why
5. Compile a report which defines the:
   Modelling process used
   Modelling symbology used
   Location of the model, and any proprietary software required to access it.

OUTPUTS: Two Outputs are created by this Step:

1. The appropriately represented Model of how As-Required (Closed) Practice compares with As-Required (Open) Practice
2. The Report accompanying the comparison of the two models of As-Required Practice.

B.9. STEP 9: Compare as observed practice with as desired practice

Along with Step 10, this is a conditional step which may not necessarily have an output. The Task for Steps 9 and 10 is identical, only the rationale behind any identified differences will differ.

OBJECTIVE: Having completed the model of as-observed safety practice, and completed the comparison with organisational practice, differences between the two models may have been identified. The objective here, therefore is to determine whether any differences in the as-observed model exist because those charged with implementing an organisation's safety lifecycle are aware of deficiencies in organisational practice with regards to achieving as-desired practice. It aims to determine whether any differences which may exist are additional activities to those required by organisational processes, or whether activities are carried out in a manner other than those required by organisational processes.

INPUTS: Two completed modelling elements of the framework instantiation process:

1. The Model of how As-Observed Practice compares with As-Required (Closed) Practice
2. The As-Desired Model of Safety Practice.

TASKS:

1. Determine whether any differences in Model of how As-Observed Practice compares with As-Required (Closed) Practice exist
2. Conduct further enquiries with the participant(s) in the observation of safety practice to determine the reasons behind the differences
3. Create a report that documents the reasons for the differences
4. Carry out further investigations with the organisation whose processes are under analysis.

OUTPUTS: The single output from this Step is a report outlining the differences between as-observed practice and as-observed (closed) practice — including the posited reasons for the differences.

B.10. STEP 10: Comparing as observed practice with as required (open) practice

Along with Step 9, this is a conditional step which may not necessarily have an output. The Tasks for Steps 9 and 10 are identical, only the rationale behind any identified differences will differ.

OBJECTIVE: Having completed the model of as-observed safety practice, and completed the comparison with as-required (Open) practice, differences between the two models may have been identified.

The objective here, therefore is to determine whether any differences in the as-observed model exist because those charged with implementing an organisation's safety lifecycle are aware of aspects of organisational practice which are not in full agreement with an Open Standard.

INPUTS: Completed Model of how As-Observed Practice compares with As-Required (Closed) Practice.

TASKS:

1. Determine whether any differences in the as-observed model exist
2. Conduct further enquiries with the participant(s) in the observation of safety practice to determine the reasons behind the differences
3. Create a report that documents the reasons for the differences
4. Carry out further investigations with the organisation whose processes are under analysis.

OUTPUTS: The single output from this Step is a report outlining the differences between as-observed practice and as-observed (closed) practice — including the posited reasons for the differences.

Appendix C. Data management

The empirical data produced in this process provides an organisation with valuable insights into the current state of their safety practice. Some of the generated data may reveal the need for further, immediate recovery action, and some data may necessitate further research before any action is taken. Further, potential impediments and their proposed next steps are considered in Table 2. This is not presented as an exhaustive list, and the implementation of these next steps is outside the scope of this process.

Table 2
Potential impediments and their mitigation(s).

Identified potential impediment	Next steps
Non-compliance between As-Required and As-Desired Practice	1. Clear deficiency (i.e. Lack of 'activity x' requires creation of 'activity x') 2. Repeat Step 2/3 to ensure sufficiency 3. Repeat Step 5/6 to ensure deficiency has been cleared
Internal consistency deficiencies (insufficient information for activity to successfully conclude/for artefact to be produced)	1. Recover internal inconsistency 2. Repeat Step 2/3 to ensure deficiency is removed
Levels of disagreement between your organisation's process and the Open Standard which influenced its development (Not applicable if your organisation is a Standard's Committee)	1. Determine the reason for each disagreement considering: a. Is there evidence of safety clutter? b. Are the disagreements due to contractual/commercial complexities? c. Are the disagreements reasonable (i.e. is there an option asserted?) 2. Assess the impact of the disagreement in terms of whether this represents: a. An unsafe act b. A necessary deviation c. Surplus work activities 3. Identify potential mitigation options 4. Assess each mitigation option on the ability of the organisation to meet the as-desired criteria 5. Assess each mitigation option for whether an unintended consequence could manifest 6. Select mitigation and implement 7. Repeat Process Steps 8 and 5
Non-compliance between As-Observed and As-Required (Closed) Practice	1. Determine the reason for each non-compliance considering the following (not exhaustive): a. Whether non-compliance is an intentional deviation in order to meet As-Desired practice b. Whether there is a lack of clarity in the As-Required (Closed) Practice c. Whether the non-compliance has an impact on safety 2. Identify potential mitigation options 3. Assess each mitigation option on the ability of the organisation to meet the as-desired criteria 4. Assess each mitigation option for whether an unintended consequence could manifest 5. Select mitigation and implement 6. Repeat Step 7
Activities emanating from As-Observed Practice which comply with As-Required (Open) Practice — but which are not mandated by As-Required (Closed) Practice	1. Determine the reason for each activity considering: (not exhaustive): a. Whether the activity is an intentional act in order to meet As-Desired practice (i.e. recovering a perceived shortfall in As-Required (Closed) Practice) b. Whether the activity has a positive/negative impact on safety 2. Identify potential mitigation options 3. Assess each mitigation option on the ability of the organisation to meet the as-desired criteria 4. Assess each mitigation option for whether an unintended consequence could manifest 5. Select mitigation and implement 6. Repeat Steps 4 (partial) and 10 (as applicable)
Activities emanating from As-Observed Practice which comply with As-Desired Practice — but which are not mandated by As-Required (Open and/or Closed) Practice	1. Determine the rationale for the additional activities of the as-observed practice 2. Determine whether other activities required by as-required practice comply with the same requirement(s) of as-desired practice (using different activities) 3. Assess the data from Step 1 and Step 2 and establish which of the activities would be the most prudent to adopt or cease. 4. Repeat Steps 2, 3, and/or 4 as appropriate

References

Antonino, P.O., Trapp, M., Barbosa, P., Gurjão, E.C., Rosário, J., 2014. The safety requirements decomposition pattern. In: Int. Conference on Computer Saf., Reliab., and Security. Springer, pp. 269–282. http://dx.doi.org/10.1007/978-3-319-24255-2_20.

Asplund, F., Holland, G., Odeh, S., 2020. Conflict as software levels diversify: Tactical elimination or strategic transformation of practice? Saf. Sci. 126, 104682. http://dx.doi.org/10.1016/j.ssci.2020.104682.

BSi, 2010. Functional Safety of Electrical / Electronic / Programmable Electronic Safety Related Systems Parts 1-7. Standard BS EN 61508.

Dekker, S., 2017. Malicious compliance. Hindsight 25. Eurocontrol. URL: https://www.eurocontrol.int/sites/default/files/publication/files/hindsight25.pdf. Last (Accessed 08 Dec 2023).

Haavik, T.K., 2021. Debates and politics in safety science. Reliab. Eng. Syst. Saf. 210, 107547. http://dx.doi.org/10.1016/j.ress.2021.107547.

Habli, I., 2017. Safety standards: Chronic challenges and emerging principles. Handb. Saf. Princ. 732–746. http://dx.doi.org/10.1002/9781119443070.ch31.

Hawkins, R., Habli, I., Kelly, T., 2013a. The principles of software safety assurance. In: International Syst. Saf. Conference.

Hawkins, R., Habli, I., Kelly, T., McDermid, J., 2013b. Assurance cases and prescriptive software safety certification: A comparative study. Saf. Sci. 59, 55–71. http://dx.doi.org/10.1016/j.ssci.2013.04.007.

Hawkins, R.D., Kelly, T.P., 2012. A framework for determining the sufficiency of software safety assurance. In: 7th IET International Conference on Syst Saf.,Incorporating the Cyber Security Conference 2012. IET, pp. 1–6. http://dx.doi.org/10.1049/cp.2012.1529.

Hollnagel, E., 2012. FRAM: The Functional Resonance Analysis Method. Ashgate, ISBN: 978-1-4094-4551-7.

Hollnagel, E., 2018. Safety–I and Safety–II: The Past and Future of Safety Management. CRC Press, ISBN: 978-1-4724-2305-4.

Hollnagel, E., Woods, D.D., Leveson, N., 2006. Resilience Engineering: Concepts and Precepts. Ashgate Publishing, Ltd., ISBN: 978-0-7546-4641-9.

ISO/IEC, 2008. Systems and Software Engineering - System Lifecycle Processes. Standard ISO/IEC 15288:2008.

Leveson, N.G., Thomas, J.P., 2018. STPA Handbook. Cambridge, MA, USA.

Marschan-Piekkari, R., Welch, C., 2004. Handbook of Qualitative Research Methods for International Business. Edward Elgar Cheltenham, ISBN: 1845424344.

MoD, 2016. Requirements for Safety of Programmable Elements (PE) in Defence Systems Part: 01 : Requirements and Guidance. Standard Defence Standard 00-055 Part 1, Issue 5, Ministry of Defence.

MoD, 2017. Safety Management Requirements for Defence Systems Part 1: Requirements. Standard Defence Standard 00-56 Part 1, Ministry of Defence.

Morley, J., Murphy, L., Mishra, A., Joshi, I., Karpathakis, K., et al., 2022. Governing data and artificial intelligence for health care: Developing an international understanding. JMIR Format. res. 6 (1), http://dx.doi.org/10.2196/31623.

O'Reilly, K., 2012. Ethnographic Methods. Routledge, ISBN: 978-0-415-516181-5.

Osborne, M., 2021. ARP 4754A - as described principle 1. URL: https://www-users.york.ac.uk/~mo705/rc_images/arp4754a_asdesiredprinciple1_v01.pdf.

Provan, D.J., Rae, A.J., Dekker, S.W., 2019. An ethnography of the safety professional's dilemma: Safety work or the safety of work? Saf. Sci. 117, 276–289. http://dx.doi.org/10.1016/j.ssci.2019.04.024.

Provan, D.J., Woods, D.D., Dekker, S.W., Rae, A.J., 2020. Safety II professionals: How resilience engineering can transform safety practice. Reliab. Eng. Syst. Saf. 195, 106740. http://dx.doi.org/10.1016/j.ress.2019.106740.

Rae, A., Provan, D., Aboelssaad, H., Alexander, R., 2020. A manifesto for reality-based safety science. Saf. Sci. 126, 104654. http://dx.doi.org/10.1016/j.ssci.2020.104654.

Rae, A.J., Weber, D.E., Dekker, S.W., 2021. Work as planned, as done and as desired: A framework for exploring everyday safety-critical practice. In: Inside Hazard. Technological Syst.. CRC Press, pp. 115–132. http://dx.doi.org/10.1201/9780429281587.

Rasmussen, J., 1997. Risk management in a dynamic society: a modelling problem. Saf. Sci. 27 (2–3), 183–213. http://dx.doi.org/10.1016/S0925-7535(97)00052-0.

Reinhardt, D., McDermid, J., 2012. Contracting for assurance of military aviation software systems. In: Proc. of the Australian Syst. Saf. Conference-Volume 145. pp. 91–105.

Roberts, K.H., 1990. Managing high reliability organizations. Calif. Manag. Rev. 32 (4), 101–113. http://dx.doi.org/10.2307/41166631.

Rooksby, J., Rouncefield, M., Sommerville, I., 2009. Testing in the wild: The social and organisational dimensions of real world practice. Comput. Support. Cooperat. Work (CSCW) 18 (5–6), 559. http://dx.doi.org/10.1007/s10606-009-9098-7.

RTCA, 2011. Software Considerations in Airborne Systems and Equipment Certification. Standard RTCA DO-178C, RTCA.

SAE Aerospace, 2010. Aerospace Recommended Practice (R) Guidelines for Development of Civil Aircraft and Systems. Standard ARP 4654A, SAE Aerospace.

SCSC, 2021. Goal Structuring Notation Community Standard. Standard, The Assurance Case Working Group, SCSC-141C.

Shorrock, S., 2020. Proxies for work-as-done: 3. Work-as-disclosed. URL: https://humanisticsystems.com/2020/11/01/proxies-for-work-as-done-3-work-as-disclosed/.

Somekh, B., Lewin, C., 2005. Research Methods in the Social Sciences. Sage, ISBN: 0-7619-4401-X.

Squair, M.J., 2006. Issues in the application of software safety standards. In: ACM Int. Conference Proc. Ser., Vol. 162. pp. 13–26, ISBN: 1920682376.

Sujan, M.A., Furniss, D., Anderson, J., Braithwaite, J., Hollnagel, E., 2019. Resilient Health Care as the basis for teaching patient safety–A Safety-II critique of the World Health Organisation patient safety curriculum. Saf. Sci. 118, 15–21. http://dx.doi.org/10.1016/j.ssci.2019.04.046.

Vilela, J., Castro, J., Martins, L.E.G., Gorschek, T., 2017. Integration between requirements engineering and safety analysis: A systematic literature review. J. Sys. Softw. 125, 68–92. http://dx.doi.org/10.1016/j.jss.2016.11.031.

Von Krogh, G., Rossi-Lamastra, C., Haefliger, S., 2012. Phenomenon-based research in management and organisation science: When is it rigorous and does it matter? Long Range Plan. 45 (4), 277–298. http://dx.doi.org/10.1016/j.lrp.2012.05.001.

Weaver, R.A., 2003. The Safety of Software: Constructing and Assuring Arguments (Ph.D. thesis). University of York York, UK.

Yin, R.K., 2014. Case Study Research Design and Methods. Sage, ISBN: 978-1-4522-4256-9.