# Building Confidence: An Ontological Approach to Assurance of Safety-Critical Systems

Odd Ivar Haugen
Group Research and Development department, DNV AS
Trondheim, NORWAY
e-mail: odd.ivar.haugen@dnv.com

***Abstract*—This paper presents an ontological framework for the assurance of safety-critical systems, focusing on the foundational relationship between knowledge, confidence, and risk. Society is growing increasingly intolerant of risk in high-tech systems; therefore, stakeholders must be provided with justified confidence in a system's safety. This confidence is not based on mere compliance with standards and guidelines, but on a robust assurance framework that demonstrates that the system behaves safely. The proposed ontology meets this need by defining assurance as a process that generates explicit knowledge to reduce uncertainty. At this framework's core, knowledge acts as the "hub" of assurance; when systematically represented in an assurance case, it directly influences stakeholder confidence. In this assurance, we generate knowledge whose objectivity provides a robust metric for justifying claims. This approach ensures that arguments supporting system safety are not only coherent but also demonstrably strong, linking the assurance effort directly to risk levels. This ontological model provides a comprehensive and systematic methodology for demonstrating safety by connecting the elicitation of system requirements to the justification of claims. This work, therefore, offers a structured path for building and communicating grounds for justified confidence in the responsible deployment of complex and novel systems.**

***Keywords-assurance; confidence; knowledge; risk; safety.***

## I. INTRODUCTION

High-tech systems, with their increasing complexity and societal integration, necessitate rigorous methods to ensure their safe and responsible operation. In safety-critical systems, where failures can have catastrophic consequences, stakeholders like operators, regulators, and the public must have justified confidence that the system will behave as intended. However, traditional approaches that focus on compliance with established standards often fail to address the novel risks and emergent behaviours of modern technologies. This failure creates a need for a more foundational approach to safety assurance.

The field lacks a clear, underlying framework that explicitly defines how assurance activities build this necessary stakeholder confidence. While common practices like developing safety cases exist, they can become procedural exercises if they lack a robust ontology that connects arguments and evidence to a tangible reduction in uncertainty and risk. To make assurance efforts both efficient and effective, this paper addresses the need for a systematic model that explains the core relationships between knowledge, risk, uncertainty, and confidence.

Existing assurance methodologies provide structures for arguing about safety, but they do not always articulate the epistemic principles that govern why these arguments should be considered trustworthy. The core limitation of current practices is the frequent disconnect between the assurance artefacts produced and the fundamental goal of cultivating a justified belief in the system's safety among diverse stakeholders. The need to bridge this gap motivates our work, which establishes a clear line of reasoning from stakeholder concerns about potential losses to the justified claims made about system behaviour.

This paper introduces a comprehensive ontological framework for the assurance of safety-critical systems, positing that assurance is fundamentally an epistemic activity. This framework generates explicit knowledge to reduce uncertainty about a system's properties. Our central thesis is that knowledge serves as the "hub" of assurance; when systematically gathered, analysed, and presented, this knowledge provides the robust and justifiable grounds for stakeholder confidence.

To develop this framework, we first model the intrinsic connections between risk, confidence, and uncertainty, demonstrating how generating knowledge directly reduces epistemic uncertainty. We advocate for a systems approach, utilising the CESM metamodel (Composition, Environment, Structure, Mechanism) to analyse emergent properties, such as safety. As a core contribution, this work establishes objectivity as a multi-dimensional metric for evaluating the strength of knowledge. Finally, we organise this knowledge within a structured assurance case. This assurance case systematically links claims about system safety to their substantiating arguments, thereby providing a scrutable and justified basis for confidence.

This paper is structured as follows. Section II introduces the main concepts of assurance. Section III provides an overview of assurance and confidence. Section IV discusses the relationship between assurance and risk. Section V presents the systems approach and the CESM metamodel. Section VI addresses epistemology and justification. Section VII introduces objectivity as a metric of knowledge strength. Section VIII discusses assurance cases. Section IX covers stakeholder objectives and system requirements. Section X concludes the paper and outlines future work.

## II. MAIN CONCEPTS OF ASSURANCE

Assurance is about becoming confident that the system behaves in a way that is acceptable to the stakeholders. Here, stakeholders are seen as any person, group of persons, governmental regulator, society, or even the natural environment. In short, it is an entity that is affected by the behaviour of the system.

A claim is a property of interest about the system. The claims can be thought of as system requirements; that is, ”this” is how the system should behave in order for the system to be accepted by the stakeholders. Analysing the previous statement reveals, as a first approach, the four principal criteria that must be in place to achieve acceptance:

1) the system requirements must reflect the interest of the stakeholders,
2) refining these requirements into technical specifications must maintain the essence of these requirements,
3) the system’s adherence to these requirements must be secured and adequately substantiated,
4) 1, 2, and 3 must be communicated to the stakeholders or their representatives in such a way that they can make intelligible decisions.

It is clear from the above items that the key to system acceptance is *knowledge*. Indeed, knowledge may be said to be the ”hub” of assurance. The stakeholders must know that the system behaves acceptably. Knowledge is a prerequisite for confidence, which reduces the uncertainty about the system.

Confidence is different from trust. Confidence is something that can be merited through demonstrating adequate capability; trust, however, has to be earned through time; that is, trust is closely connected to an agent’s intention. This means that confidence can be merited through demonstrating adequate capability (technical system and responsible agent); trust must be earned through time by a responsible agent adhering to sound and recognised ethical principles.

As assurance is about providing grounds for justified confidence, this paper will therefore focus on how to demonstrate adequate system capability so that the stakeholders can make intelligible decisions based on their knowledge and, thereby, their level of confidence in the system.

It should be noted that assurance is an epistemic activity, while risk management encapsulates both epistemology and intervention in the real world [1] [2].

The system capability, in this context, is equivalent to how the system behaves under normal operation and in abnormal situations.

The system risk is defined as the ”effect of uncertainty on objectives” [3] and reflects the consequences and uncertainties that the system causes losses for stakeholders. The uncertainty is here divided into two types: epistemic and aleatory [4].

Item three in the above list requires that the system behaviour adherence to the requirements is substantiated; that is, claims about the system must be substantiated through sound and relevant argumentation. For an argument to be sound, it must be generated in accordance with acknowledged methodologies using reliable tools and adequately skilled people.

To assess the soundness and the strength of arguments, an assessor not only needs to be a subject matter expert but also needs guidance about what can be regarded as acceptable methods and processes to develop arguments; that is, he needs guidance about the argument’s *objectiveness*. A higher degree of objectivity increases the strength of the argument, which is necessary when the risk is high, such as for safety-critical systems.

## III. Assurance and confidence - an overview

Confidence can be thought of, in statistical terms, as a quantitative measurement of uncertainty, e.g., an interval indicating the confidence that the value of a parameter is likely to fall within. However, confidence may also be thought of as a feeling that reflects the coherence of the information and the cognitive ease of processing it [5]. Assurance is defined as ”grounds for justified confidence that a claim has been or will be achieved” [6]. The definition does not limit assurance to either type of confidence; hence, assurance addresses both.

Both types of uncertainties pose challenges. The frequentist approach to quantifying uncertainty requires robust statistical data. Here lie a few major obstacles, some of which are: the inherent complexity of many safety-critical systems, the novelty of the technology, statistically significant data from rare events, and assigning probabilities to inherently social aspects.

The second type of confidence also poses challenges. We cannot base decisions concerning the safety and well-being of stakeholders and society on pure feelings but on strong knowledge based on facts and trustworthy evidence.

Therefore, assurance may provide grounds for justified confidence through uncertainty quantification only if based on robust statistics, that is, knowledge about properties of the statistical distribution of the parameter in question, and/or judgemental assessments only if based on sound argument substantiating the truthfulness of the claim.

Therefore, the immediate goal, or primary effect of assurance, is to generate knowledge, knowledge to decrease or establish the uncertainty about a claim, addressing both types of uncertainty when appropriate. A Functional Analysis System Technique (FAST) diagram illustrates the relation between assurance, knowledge and confidence (Figure 1).

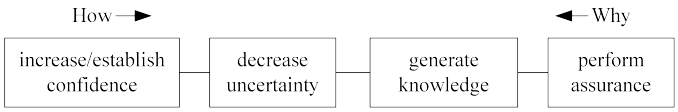


Figure 1. FAST diagram connecting knowledge to confidence.

A FAST diagram is read either way from left to right by asking, *How* is this function achieved? Or, from right to left by asking *Why* does this function need to be achieved.

As knowledge is the ”hub” of assurance, knowledge must be treated systematically and expressed explicitly to enable it to be rigorously scrutinised. This is to avoid that confidence being based on unsubstantiated feelings and pure guesswork. The assurance case is a systematic and explicit way of representing and treating knowledge.

As safety is an emergent property [7], the knowledge about the truthfulness of the claim must address all system aspects that affect emergence. Elements necessary in analysing emergent behaviour in engineered socio-technical systems are encapsulated in the systems approach.

Figure 2 depicts how the different items of assurance are related.

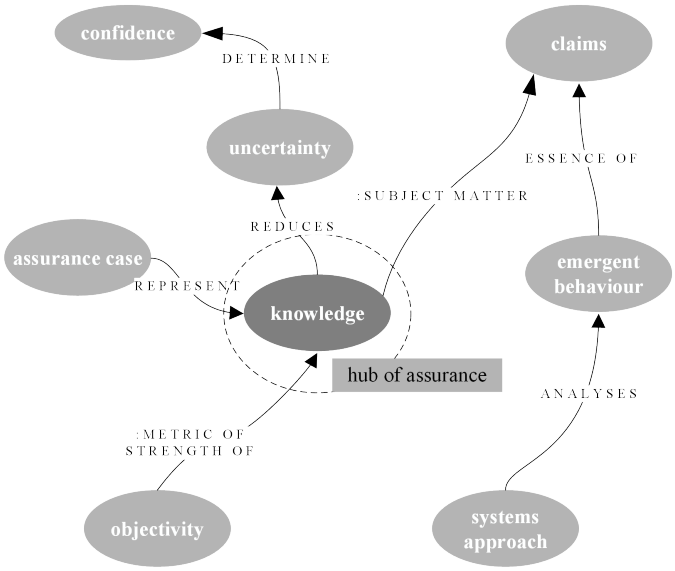


Figure 2. Overview of the ontology of assurance.

Intuitively, the higher the risk that the system poses to stakeholders, the higher confidence we need that it will indeed behave as expected. As knowledge reduces uncertainty and increases confidence, we need a way to assess its strength. Assessing the strength of knowledge is key to adjusting the assurance effort to risk level.

## IV. Assurance and System Risk

Figures 1 and 2 showed how knowledge generated in the assurance effort reduces uncertainty, and that uncertainty determines confidence. Moreover, as earlier established, uncertainty is one part of the risk concept. Hence, assurance and risk are connected through uncertainty (Figure 3).

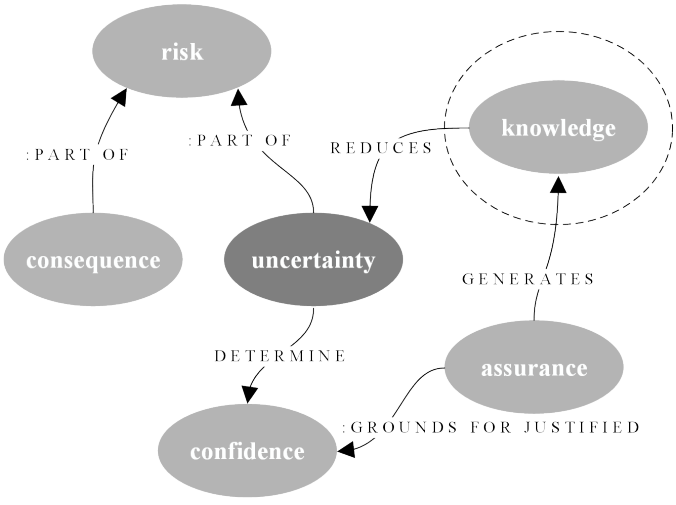


Figure 3. Assurance is connected to risk through uncertainty.

There is, however, another connection in addition to the one mentioned above. In the top right corner of Figure 2, it is indicated that the subject matter of the knowledge is the claim. Claims are statements about system properties that address the system requirements elicited by stakeholders and their concerns and objectives (Figure 4).

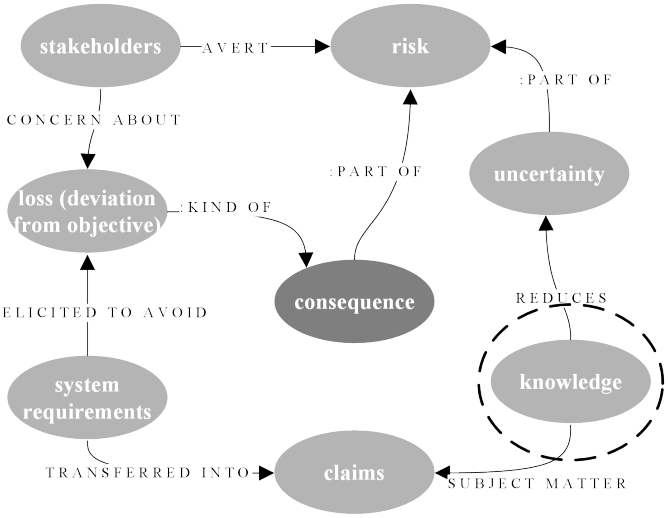


Figure 4. Assurance is connected to risk through claims.

Stakeholders are generally risk avert [5] and are concerned about the consequences of losses. They need adequate confidence that potential losses are acceptable. Assurance addresses these concerns by generating knowledge about the truthfulness of claims made about the system properties.

Risk can be reduced by altering the system design or operational conditions. These risk-reducing strategies affect the consequence and/or the aleatory uncertainty. However, as this paper is concerned with assurance, which is an epistemic endeavour, these two strategies are not further discussed. Their relationship to assurance is discussed in [1]; on the relationship between assurance and risk management.

## V. Assurance and the Systems Appoach

A way to understand and analyse complex systems and emergence, is to model the system behaviour in terms of its composition, structure, mechanisms and the environment in which it operates. These system aspects are termed the CESM metamodel [8]:

- **Composition** (C): Collection of all the parts or objects in the system.
- **Environment** (E): Systems outside (excluded from) the target system, but act upon, or are acted upon by, the target system.
- **Structure** (S): The relationships and bonds among the system agents and between the system agents and the environment.
- **Mechanisms** (M): The processes that make the system behave in the way that it does.

The emergent behaviour becomes a function of the above elements; that is, any system $s$ may be modelled, at any given instance, as the quadruple: $\mu(s) =< C(s), E(s), S(s), M(s) >$. As $\mu(s)$ is an emergent property, and emergent properties exist on different levels of abstraction (LoA) [9], the CESM must also be instantiated at these LoAs.

This can be visualised by the system triangle (Figure 5) where the corner of the triangle illustrates "CSM" encapsulated by "E". The "system" in the middle represents $\mu(s)$. $\mu(s)$ emerges, therefore, as a result of the conceptual interaction between the corners of the triangle, but also between the triangle and the environment (E). To move the analysis between the LoAs, a rule-based gradient is used, termed the gradient of abstraction (GoA).

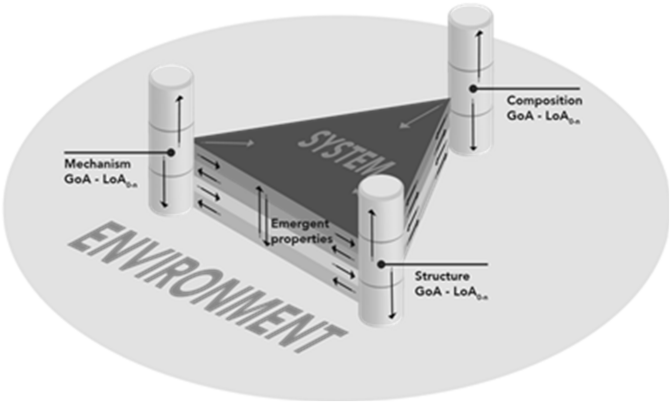


Figure 5. The CESM triangle showing.

For each element in the CESM metamodel, we can assign different system model categories [10]:

- Composition: **Object model** representing the system elements and components and their ontological relationship to each other.
- Environment: Also modelled as a system containing all aspects of the CESM metamodel, which means that the environment must be represented by models representing the composition, structure and mechanisms (our target system is part of the environment of its environment).
- Structure: **Agent model** includes entities, such as controllers, actuators, sensors, humans, and subsystems. The agent concept includes authority, responsibility, goals, concerns, motivation, and wishes (humans).
- Mechanisms: **Function model** represents the operations that must be performed (by the agents) to achieve goals.

Examples of system model instantiation of the agent model is the control structure known from Systems-Theoretic Process Analysis (STPA) [7]. Another agent model may focus more on the agent's goals, motivation, concerns and wishes, like a model used in a stakeholder analysis where social and business aspects are emphasised.

A function model may focus on the preconditions, resources, and timing for achieving it, like the model used in the Functional Resonance Analysis Method (FRAM) [11].

The functional dependencies between functions, like in FAST [12] may be used as GoA to move the analysis between abstraction levels, that is, to represent the system at different LoAs [13].

The systems approach described above, used in assurance, can be summarised by the following statements [2]:

- The conceptual interaction between the system composition (C), environment (E), structure (S), and mechanisms (M) models the system behaviour.
- The kind and number of levels of abstractions (LOAs) used in the modelling is determined by the knowledge sought through the assurance effort.
- The systems approach is used in every aspect of the assurance effort, such as system description, describing the system boundary, describing the environment in which the system is operating, system analysis, verification and validation, and elicitation of system requirements.

Figure 6 depicts the relationship between the systems approach and assurance.

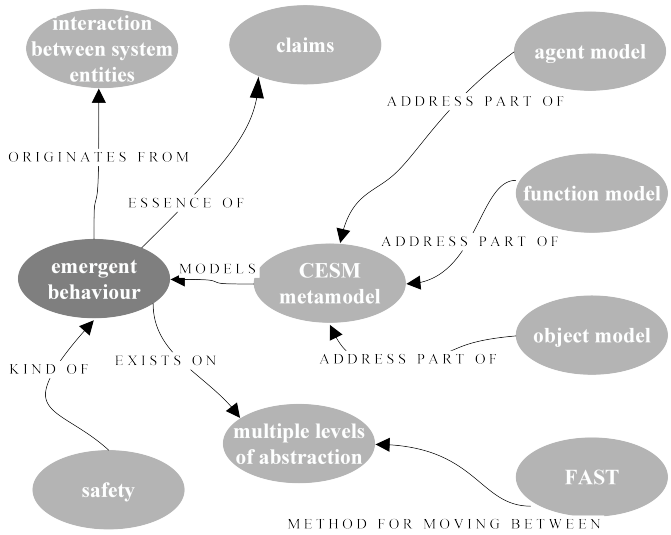


Figure 6. The systems approach is connected to assurance through the claims or requirements.

The system safety requirements are formulated as safety claims. As safety is an emergent property that emerges through the interaction between the system entities, it can be modelled through the CESM metamodel.

## VI. Assurance, epistemology and justification

Recall that the concept of risk incorporates, in addition to the consequence, two kinds of uncertainties: epistemic and aleatory. Strengthening the knowledge reduces epistemic uncertainty. If the risk is high, as in safety-critical systems, the argument supporting the claim must be strong. The strength of the argument and, thereby, the strength of the knowledge reduces the epistemic uncertainty and, thereby, the risk.

The classic definition of knowledge is: "Justified True Belief" (JTB). Although this definition has been under scrutiny for centuries and has been shown to have weaknesses [14], it must be linked to accessible facts about the subject matter. Moreover, building confidence through knowledge requires, not only apparently truthful propositions (claims), but also that the reasoning is sound, relevant and adequate; that the proposition is justified: *"Someone who is very confident but for the wrong reasons would also fail to have knowledge"* [15].

The reason for believing that a proposition represents the truth must be justified.

Justification may be thought of as an argument for why we hold certain beliefs or why we think those beliefs are reasonable and true. These justifications may be under the law or before God. However, in the context of assurance, justifying beliefs must be based on knowledge, or, in other words, on epistemic justification [16]. (A safety-critical system needs, of course, to conform to laws and regulations; however, the point is that the justification must be based on knowledge.)

Assurance seeks epistemic justification to establish whether a proposition can be turned into a belief, that is, belief through warranted propositions.

Belief revision is the process of changing beliefs based on new data [17]. It is important to emphasise that good reasoning is no guarantee of truth. Seeking the truth and believing to have found it using sound methods and reasoning is no guarantee of actually having found it.

Justifying a proposition may, in principle, entail an infinite chain of justifications (infinitism): The justification of the justification of the justification... This is, of course, unacceptable. The question, then, is when to stop this chain of justifications.

One strategy is to continue until the supporting justifications become self-evident, that is, propositions that do not need further justification (foundationalism). This kind of justification results in a hierarchy of propositions, and the "bottom" of this hierarchy consists of fundamental propositions, that is, self-justified propositions.

Alternatively, we may ensure that the propositions support each other, that is, the propositions are coherent (coherentism). With this strategy, there are no fundamental propositions. Critiques claim that this strategy can lead to circular argumentation [16].

A reasonable approach is to combine the two strategies, that is, ensuring coherence within the set of propositions and justification, and stopping the chain of justification when reaching a self-justified proposition.

In practice, one may not reach a self-evident fundamental level for several reasons. One reason may be that there is a dispute about whether such a level is actually reached; another reason may be that continuing the chain of justification requires disproportionate resources. Therefore, there may be residual uncertainty as to whether a proposition represents the truth.

Showing compliance towards an international industry standard is often regarded as such a self-justified belief. Providing evidence that a system complies with such a standard is often regarded as adequate for believing a proposition, e.g., that a system is reliable, fair, safe, and secure, as an international standard should reflect good industry practice. However, e.g., artificial intelligence (AI) is a novel technology that, even if there exists a relevant international standard, it may not be regarded as self-justified because the standard itself does not necessarily reflect any industry practice (because there do not exist any such practice), or at least the practice may be inadequate. This means that it might be necessary to continue the justification chain further when assuring novel complex systems, e.g., based on AI.

Other sources of uncertainty include evidence that weakens the proposition or a lack of available evidence. Moreover, other obstacles may hinder the generation of additional evidence, such as technical limitations, ethical concerns, lack of statistical data, or other practical causes.

There is no universal uncertainty threshold for when an agent will accept a proposition and when he rejects it. Moreover, given a justification of a proposition, there is no universal law governing the level of uncertainty an agent will feel about its truthfulness.

Belief revision depends not only on the properties of the justification of the proposition but also on the agent's epistemic state, that is, the agent's required rationality to turn a proposition into a belief, prior belief and any other properties important for the agent to represent facts about the world.

The uncertainty threshold for an agent's belief revision also depends on aspects such as the risk (perceived and/or actual) of accepting or rejecting a proposition (including being indifferent). Moreover, an agent's level of uncertainty, given a justification of a proposition, depends not only on the strength of the justification, but also on aspects such as the degree of being susceptible to cognitive biases [5] and rhetoric. Obviously, we should strive to minimise aspects of belief revision that are unrelated to the properties of the justification.

Perhaps the most commonly known is the so-called confirmation bias, that is, our tendency to seek evidence that confirms our prior beliefs. However, most other cognitive biases are at work, like the illusion of understanding and what you see is all there is (WYSIATI), that is, our tendency of believing that we understand complex topics by filling in the information gaps and the epistemic gaps so that the story becomes compelling and coherent, which leads to confidence in the truthfulness of the story (or proposition in this case).

An agent's prior beliefs cannot, and should not, be controlled and cannot be totally known. Nevertheless, prior belief is central to belief revision. Data-oriented Belief Revision (DBR) [18] (simplified illustration in Figure 7) is a model of belief revision that can illustrate the role of prior belief in belief revision.

After new data is available about a proposition (External data), the data is assessed to determine their relevance and strength, possibly forming a new or updated belief set, termed *belief selection* in Figure 7. This process regulates the interaction between data and beliefs, what to believe in, and with what strength.

As belief revision is tightly connected to the agent's prior beliefs and possible degrees of cognitive biases, we cannot assess the epistemic strength of the justification by appealing to the agent's prior beliefs, or what seems to be "very reasonable" and the like. What seems reasonable is an internal feeling in each agent and is largely based on their current epistemic state.

Instead, the agent needs to be nudged towards sound rationality of assessing uncertainty using a more comprehensive framework of thinking about the level of uncertainty (epistemic strength of the justification), without being forced into an

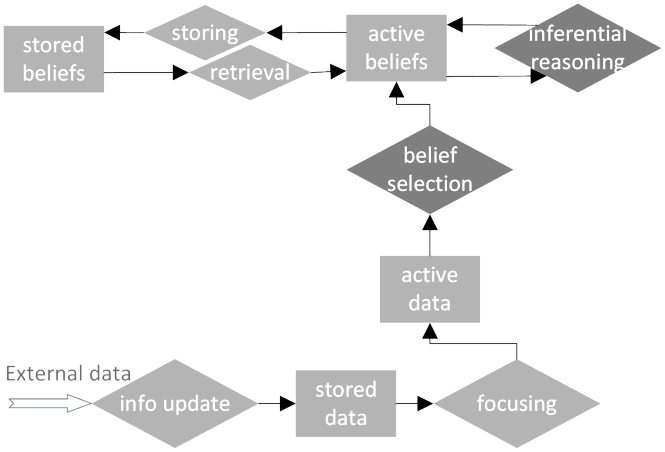


Figure 7. Simplified epistemic processing in DBR [18].

epistemic straitjacket of predefined categories of epistemic levels.

We want to decrease uncertainty as the risk of accepting a false proposition increases. The opposite may not be so obvious, that we also want to decrease uncertainty when risk increases by rejecting a true proposition. Accepting a false proposition, on the one hand, or rejecting a true proposition, on the other, represents assurance risk.

Assurance risk, that is, the risk of making wrong decisions due to weak or inaccurate knowledge, is one kind of risk in the context of assurance. The other kind of risk is the system risk, that is, the undesired consequences with associated uncertainty of operating a system in the real world.

Decreasing uncertainty to the point of accepting a proposition, or in other words, revising one's belief, can be achieved by both strengthening the justification that the proposition is true, and/or by increasing effort in seeking justification that the proposition is false without finding such justification. Sometimes, the only way to justify a proposition $p$ is to find a strong justification that $\neg p$ is not the case.

A famous statement from software testing illustrates this: *Software testing cannot prove the absence of bugs, only their presence*. A proposition that some software code is bug-free $p$ cannot be proven through testing alone. Software testing tries to find bugs, and when no bugs are found, one may start to believe $p$ because one hasn't found evidence that $\neg p$ is the case. However, as most testing is non-exhaustive, not finding bugs does not mean the absence of bugs.

A way to accommodate proper assessment of knowledge built on epistemic justification is through argumentation. While belief revision describes how we should update our beliefs, argumentation is a way to make belief revision occur. "The two concepts are two sides of the same epistemic coin" [19] [18].

## VII. Objectivity - a metric of strength of knowledge

By generating knowledge about the system, the epistemic uncertainty about deviation from objective changes, that is, knowledge about how an accident may occur or the potential consequence should it occur. High risk means severe potential consequences combined with a large degree of uncertainty (epistemic and/or aleatory). As knowledge decreases uncertainty, high risk requires strong knowledge, that is, knowledge substantiated with strong grounds for justification.

Justification, and thereby knowledge, is, among other things, based on artefacts representing the system and its properties, together with how these artefacts are interpreted, that is, the reasoning used to conclude based on these artefacts. Artefacts, such as training data, algorithms, source code, and system descriptions, may represent the system directly. Other kinds of artefacts may indirectly represent it, e.g., artefacts generated through verification, such as test cases, test results and results from inspections and reviews. The strength of knowledge is directly linked to these artefacts and the process of generating and collecting them.

Distinguishing weak from strong knowledge requires a metric by which the strength of knowledge can be assessed. By comparing the definitions of knowledge and assurance, we recognise the similarities. Both definitions contain the term "justified": The degree of justification for a true belief (knowledge) - the grounds for justified confidence (assurance). Degree of justification is central in assessing both strength of knowledge and degree of confidence (via uncertainty as shown in Figure 2). A high degree of confidence requires strong ground for justification.

Objectivity encapsulates the aspects important for assessing the degree of justification, that is, the strength of knowledge. Hence, the strength of knowledge is measured through the degree of objectivity. The likelihood that the result of an enquiry represents the truth increases if it is conducted in an objective manner, including the artefacts produced and used in that enquiry.

Ensuring consistency and repeatability in our enquiries requires that the concept of objectivity be described. Objectivity in this context is a multi-dimensional, non-orthogonal and non-binary concept [20]. Hence, objectivity cannot be treated in a reductionist manner.

There are three categories (i.e., dimensions) that lay out the space of objectivity [20] [14] (Figure 8):

1) properties and processes by which the artefacts are generated
2) reasoning, or the thinking about those artefacts
3) social processes concerning items 1 and 2.

**Item 1** is about interacting with the system and its stakeholders during its entire lifecycle. It is about the choice of methods, how they are applied, and how those decisions influence the properties of the outcomes, that is, the artefacts. This category also includes procedures, methods, techniques, first principles in physics, standardised equations, algorithms, etc.

**Item 2**, this category is about how people and organisations think and the reasons and positions they take based on their interests and roles. This includes the involved assurance agent's values and independence from the developer.

**Item 3** is about the social processes that advocate different viewpoints, such as agreement among subject-matter experts

on the suitability and correct use of methods for generating artefacts and on how to think about those artefacts. This kind of objectivity can be thought of as a form of inter-subjectivity and is strengthened if the group consists of individuals with different but relevant competencies. The content of standards is a result of such agreements.

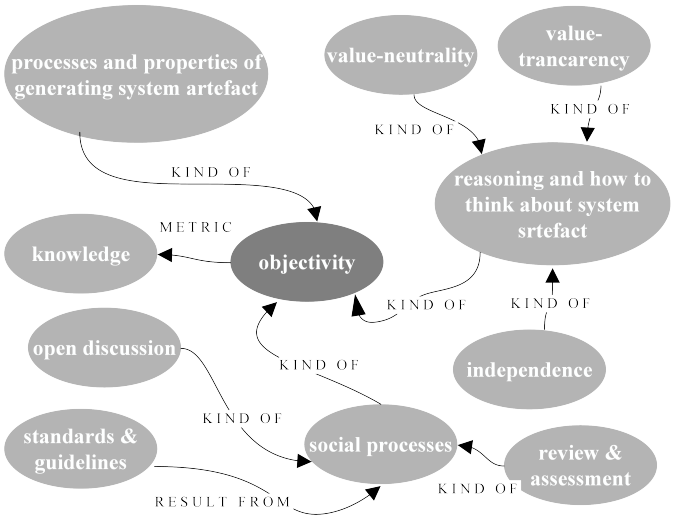


Figure 8. Categories of objectivity.

An important activity in assurance is the generation and collection of evidence through verification and validation (V&V). V&V is described through two properties: 1) The level of intensity in the V&V effort, and 2) the level of rigour in the V&V effort [21]. V&V intensity is connected to the size of the scope, the number of system artefacts investigated, and the level of V&V involvement in each phase of the system lifecycle. V&V rigour is connected to comprehensiveness and thoroughness, leaving less room for logical inconsistencies and contradictions in the results, that is, performed with different levels of formality concerning techniques and documentation. One useful metaphor describing the relationship and difference between the two properties may be that increased V&V intensity makes the mesh width smaller and smaller, while increasing the V&V rigour means that each mesh is investigated closer and closer.

The output from the V&V effort is the evidence representing the system properties of interest, such as safety, reliability, robustness and security. V&V intensity and rigour affect the evidence properties [21] such as quality, capability, and coverage.

Confidence is a result of the assessment of the strength of justification and knowledge through the degree of objectivity. Furthermore, through the V&V intensity and rigour, and the resulting evidence properties. The assessment cannot be a simple checklist, which results in a numerical score aggregated as a simple sum or a single-dimensional category. The strength of knowledge must be assessed in each particular project in the context of a totality. That is, the strength (of knowledge) is not a resultant property of the degree of objectivity (and V&V), but emergent. Assessing the truthfulness (strength of justification and knowledge) of claims made about emergent properties in novel, complex safety-critical systems depends on the judgement of experts in the relevant disciplines. It is guided by the objectivity criteria described here.

This position does not preclude the use of quantitative or probabilistic measures where they are epistemically justified and appropriate; rather, it asserts that no single quantitative score can substitute for expert judgement when assessing the strength of knowledge concerning emergent properties in complex safety-critical systems.

## VIII. Assurance case - a systematic way to represent knowledge

The assurance case is a way to represent knowledge (Figure 9 and Figure 2). At its core, an assurance case consists of a hierarchy of claims and arguments, including evidence that substantiates those claims. The claims are equivalent to the before-mentioned propositions, and the argument is equivalent to the before-mentioned justifications. Moreover, claims can be understood as a reformulation of system requirements. A question may be how to lay out and organise arguments, which is the topic of this section.

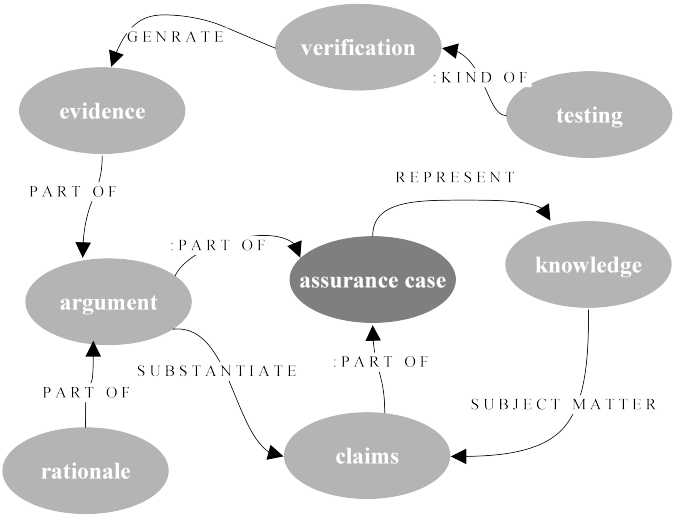


Figure 9. The assurance case represents the knowledge in an assurance effort.

One of the most recognised and influential argument schemas is the one described by Stephen Toulmin in his 1958 book ”The Uses of Arguments” [22]. Toulmin’s motivation was to create a richer format that better reflected how people argued in reality, rather than the more formal and traditional format consisting of premise and conclusion.

The argument layout consists of six elements [23]: Claim (or Conclusion) (C), Data (D) (or Datum, Toulmin uses both terms), Warrant (W), Qualifier (Q), Backing (B), Rebuttal (R) (Figure 10).

In the simplest form, (D) may be some evidence that proves that (C) is the case. The transition between (D) and (C) may not be trivial, so a warrant needs to act as an inference licence between (D) and (C); that is, (W) acts as a bridge between (D) and (C). (W) may also be challenged, so a backing (B) may be

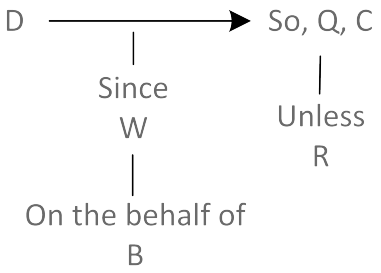


Figure 10. General layout of an argument [23, p. 97].

needed to support (W), that is, why (W) holds. (Q) indicates the strength of the step (i.e., strength of the ”bridge”) from (D) to (C). (R) indicates circumstances in which (W) may not hold.

Although the elements of an argument described by Toulmin are necessary aspects of an epistemic justification substantiating a proposition or assertion, the schema, in its simplest form, is insufficient for assurance of complex systems. The schema needs to be expanded.

Firstly, in the assurance of complex systems, there are many claims. System claims represent statements about the system properties and its use. These requirements address many system properties, including safety. Moreover, the claims must be refined at several levels of abstraction (LoAs). The LoAs link back to the LoAs connected to the systems approach and foundationalism.

Secondly, although one of Toulmin’s key motivations was to enable ”practical assessment of arguments” [23], he did not discuss aspects of argument assessment in detail. Clearly, when, e.g., a (top) claim is refined into two or more subclaims with accompanying justification, assessing the strength of each argument needs to be aggregated in some way to reflect the confidence in the top claim. Moreover, each element in the argumentation schema should be assessed, resulting in a network of assessments across different elements of an argument at different LoAs.

Several expanded argument schemas based on Toulmin have been developed, such as Goal Structuring notation (GSN) [24] and Trust-IT [25].

An assurance case organises these arguments systematically and in a structured manner, and represents the knowledge generated in the assurance (Figure 2). Different ways are possible based on the various argument schemas, such as [24] or [26]; both are compatible with [6]. A metamodel of an assurance case may also be found in [27].

## IX. Stakeholder’s objectives and system requirements

Stakeholders hold objectives and pursue goals through utilising the system; that is, they use the system for a reason. A system’s mission is expressed as system requirements, which are derived from these objectives.

The stakeholders may be users, developers, and bystanders who have nothing to gain from the system but may be affected by it. Through its legislation and standards, the government represents stakeholders that cannot be consulted directly, such as the natural environment, future generations, the general public, children, etc. In such cases, conformance to standards means meeting stakeholders’ objectives and interests.

Stakeholders need confidence that their objectives are fulfilled or will be, or that a deviation from those objectives is acceptable. Implicitly, stakeholders also hold the objectives of being safe, secure, and treated fairly. These objectives may not be directly linked to the reason for developing and using the system in the first place (i.e., the mission). The system requirements must incorporate such implicit objectives. These kinds of system requirements can be termed mission-supporting requirements, or non-functional requirements [28], or even system constraints (Prof. Nancy Leveson terms this ”safety constraints”; however, when expanding the scope of such requirements to other system quality characteristics, they can be termed as ”system constraints”.) [7] (Figure 11).

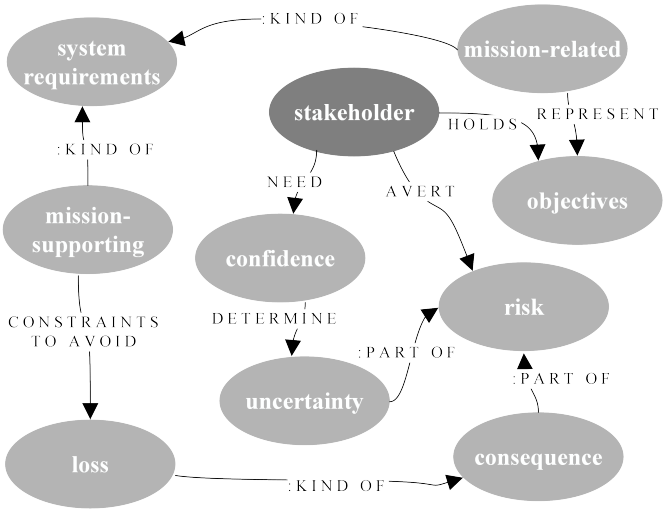


Figure 11. Stakeholders hold objectives that determine the system requirements.

In the context of assuring AI systems, mission-supporting system requirements should be based on a set of ethical principles, such as: [29].

Conflicts often arise between requirements directly related to the mission of the system and the mission-supporting requirements. Moreover, similar conflicts may also arise between the objectives and goals of different stakeholders, and even between different objectives of the same stakeholder (e.g., long-term vs. short-term goals). One understanding of ethics is: ”the identification, study, and resolution or mitigation of conflicts among competing values or goals” [30]. The assurance effort should document the trade-offs made between competing goals.

## X. Conclusion and future work

This paper has presented a comprehensive ontology for the assurance of safety-critical systems, positing that assurance is fundamentally an epistemic activity. The core thesis establishes

knowledge as the central ”hub” of assurance; its systematic generation and explicit representation are the primary means of reducing epistemic uncertainty, which in turn builds justified stakeholder confidence. We have demonstrated how stakeholder concerns about potential loss are translated into system requirements and safety claims. These claims are substantiated by the assurance process, which generates knowledge structured and presented within an assurance case. The framework employs the CESM metamodel as a foundational systems approach to analyse the system behaviour and emergent properties, such as safety, that these claims address. Furthermore, we have introduced a multi-dimensional concept of objectivity as a critical metric for evaluating the strength of this knowledge. This metric ensures that the assurance effort is commensurate with the level of system risk.

This work shifts assurance from a traditional, compliance-focused approach to a more foundational and systematic methodology. This ontological model provides a scrutable and reasoned pathway for demonstrating a system's safety by clearly articulating the relationships between risk, knowledge, and confidence. This pathway is particularly significant for novel and complex systems, where established standards may be inadequate, and a deeper justification of safety is required to gain stakeholder acceptance. The framework offers practitioners a structured methodology to connect high-level stakeholder objectives directly to the evidence and arguments that form the basis of a safety case.

The scope is intentionally focused on assurance as an epistemic endeavour—the generation of knowledge to reduce epistemic uncertainty. Consequently, we did not discuss other vital risk management strategies in detail, such as altering system design to mitigate consequences or reduce aleatory uncertainty. Additionally, while the concept of objectivity provides guidance, experts must ultimately assess the strength of knowledge through a nuanced process that relies on judgement rather than a simplistic quantitative measure.

Building upon this foundation, further work could focus on operationalising the multi-dimensional objectivity metric into practical assessment tools. Practitioners could then apply the complete ontological framework to specific, challenging domains such as autonomous systems or artificial intelligence. A final valuable avenue for inquiry would be to investigate methods for aggregating argument strength across multiple levels of abstraction within a complex assurance case.

Ultimately, this paper provides a robust and coherent ontology for building justified confidence in the safety of complex systems. This approach offers a rigorous and defensible foundation for the responsible design, deployment, and operation of safety-critical complex systems by grounding assurance in the systematic generation of knowledge.

## REFERENCES


[1] O. I. Haugen, ‘Integrating Assurance and Risk Management of Complex Systems’, in *2025 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, Vienna, Austria: IEEE, 2025-10, pp. 6739–6746, ISBN: 979-8-3315-3358-8. DOI: 10.1109/SMC58881.2025.11342869. Accessed: 2026-04-13.

[2] DNV, *DNV-RP-0671 Assurance of AI-enabled systems*, Recommended Practice, 2023-09.

[3] International Organization for Standardization, *ISO/IEC/IEEE 31000 - Risk management*, International Standard, 2018-02. Accessed: 2026-04-13.

[4] C. R. Fox and G. Ülkümen, ‘Distinguishing Two Dimensions of Uncertainty’, in *Perspectives on Thinking, Judging, and Decision Making: A Tribute to Karl Halvor Teigen*, Universitetsforlaget, 2011, pp. 21–36, ISBN: 978-82-15-01878-2.

[5] D. Kahneman, *Thinking, Fast and Slow*, 1st ed. New York: Farrar, Straus and Giroux, 2011, ISBN: 978-0-374-27563-1 978-0-374-53355-7 978-0-606-27564-4.

[6] International Organization for Standardization, *ISO/IEC/IEEE 15026 Systems and software engineering–Systems and software assurance –Part 1: Concepts and vocabulary*, International Standard, 2019-03. DOI: 10.1109/IEEESTD.2019.8657410. Accessed: 2026-04-13.

[7] N. G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*. Cambridge, Massachusetts: MIT Press, 2012-01. DOI: 10.7551/mitpress/8179.001.0001. Accessed: 2022-10-21.

[8] M. Bunge, *Emergence and Convergence: Qualitative Novelty and the Unity of Knowledge* (Toronto Studies in Philosophy). Toronto ; Buffalo: University of Toronto Press, 2003, ISBN: 978-0-8020-8860-4.

[9] O. I. Haugen, ‘Safety assurance of complex systems Part 2: Assurance and analysis’, DNV AS, Høvik, Norway, Whitepaper, 2019. Accessed: 2026-04-13.

[10] O. I. Haugen, ‘A Systems Approach to Modelling Emergent Behaviour in Maritime Control Systems Using the Composition, Environment, Structure, and Mechanisms (CESM) Metamodel’, in *The Fifteenth International Conference on Performance, Safety and Robustness in Complex Systems and Applications*, vol. ISSN: 2308-3700, Nice, France: Think Mind, 2025-06, pp. 1–8, ISBN: 978-1-68558-280-7. Accessed: 2026-04-13.

[11] E. Hollnagel, *FRAM: The Functional Resonance Analysis Method, Modelling Complex Socio-Technical Systems*. Ashgate Publishing Limited, 2012.

[12] C. W. Bytheway, *FAST Creativity & Innovation: Rapidly Improving Processes, Product Development and Solving Complex Problems*. Fort Lauderdale, Fla: J. Ross Pub, 2007, ISBN: 978-1-932159-66-0.

[13] O. I. Haugen, ‘The Systems Approach’, in *Demonstrating Safety of Software-Dependent Systems; With Examples from Subsea Electric Technology*, T. Myhrvold and M. van der Meulen, Eds., DNV AS, 2022, pp. 145–163, ISBN: 978-82-515-0324-2. Accessed: 2026-04-13.

[14] O. I. Haugen, *An epistemic approach to confidence through objectivity in assurance of safety-critical complex systems*, 2024-12. Accessed: 2026-04-13.

[15] J. Nagel, *Knowledge: A Very Short Introduction* (Very Short Introductions 400), First edition. Oxford: Oxford University Press, 2014, ISBN: 978-0-19-966126-8.

[16] J. C. Watson, *Epistemic justification*, https://iep.utm.edu/epi-just/. Accessed: 2026-04-13.

[17] M. A. Falappa, G. Kern-Isberner and G. R. Simari, ‘Belief Revision and Argumentation Theory’, in *Argumentation in Artificial Intelligence*, I. Rahwan and G. R. Simari, Eds., 1st ed., Boston, MA: Springer Dordrecht Heidelberg, 2009-07, ISBN: 978-0-387-98196-3 978-0-387-98197-0. DOI: 10.1007/978-0-387-98197-0.

[18] F. Paglieri and C. Castelfranchi, ‘The Toulmin Test: Framing Argumentation within Belief Revision Theories’, in *Arguing on the Toulmin Model: New Essays in Argument Analysis and Evaluation*, D. Hitchcock and B. Verheij, Eds., Dordrecht: Springer Netherlands, 2006, pp. 359–377, ISBN: 978-1-4020-

4938-5. DOI: 10.1007/978-1-4020-4938-5_24. Accessed: 2023-11-06.

[19] F. Paglieri and C. Castelfranchi, *Argumentation and Data-oriented Belief Revision: On the Two-Sided Nature of Epistemic Change*, https://cmna.csc.liv.ac.uk//CMNA4/B.pdf, 2004-01. Accessed: 2026-04-13.

[20] H. E. Douglas, *Science, Policy, and the Value-Free Ideal*. University of Pittsburgh Press, 2009, ISBN: 978-0-8229-6026-3. DOI: 10.2307/j.ctt6wrc78. JSTOR: j.ctt6wrc78. Accessed: 2023-10-10.

[21] O. I. Haugen, 'Safety assurance of complex systems Part 3: Verification and evidence', DNV, Høvik, Norway, Whitepaper, 2019. Accessed: 2026-04-13.

[22] B. Verheij, 'The Toulmin Argument Model in Artificial Intelligence – Or: How semi-formal, defeasible argumentation schemes creep into logic', in *Argumentation in Artificial Intelligence*, 1st ed., Springer New York, NY, 2009-01, pp. 219–238.

[23] S. E. Toulmin, *The Uses of Argument*, 2nd ed. Cambridge University Press, 2002, ISBN: 978-0-521-53483-3.

[24] *Goal Structuring Notation Community Standard*, https://scsc.uk/gsn-standard, 2021-05. Accessed: 2026-04-13.

[25] J. Górski, Ł. Cyra, A. Jarzębowicz and J. Miler, *Argument Strategies and Patterns of the Trust-IT Framework*, 2008-01. Accessed: 2026-04-13. [Online]. Available: https://www.researchgate.net/publication/229034967%5C_Argument%5C_Strategies%5C_and%5C_Patterns%5C_of%5C_the%5C_Trust-IT%5C_Framework.

[26] *Argevide - System assurance management tools - Assurance cases*, https://www.argevide.com/home/, 2023-10. Accessed: 2024-01-15.

[27] *Structured Assurance Case Metamodel (SACM)*, https://www.omg.org/spec/SACM, 2023-10. Accessed: 2026-04-13.

[28] A. van Lamsweerde, *Requirements Engineering: From System Goals to UML Models to Software Specifications*. Chichester, England ; Hoboken, NJ: John Wiley, 2009, ISBN: 978-0-470-01270-3.

[29] High-Level Expert Group on AI, 'Ethics Guidelines for Trustworthy AI', European Commission, B-1049 Brussels, Tech. Rep., 2019-04. Accessed: 2022-06-10.

[30] L. McDaniel, *What Is Bioethics?*, https://bioethics.msu.edu/about/what-is-bioethics. Accessed: 2026-04-13.