# Towards a Clearer Understanding of Context and Its Role in Assurance Argument Confidence

Patrick John Graydon

Mälardalen Real-Time Research Centre, Mälardalen University, Västerås, Sweden
patrick.graydon@mdh.se

**Abstract.** The Goal Structuring Notation (GSN) is a popular graphical notation for recording safety arguments. One of GSN's key innovations is a context element that links short phrases used in the argument to detail available elsewhere. However, definitions of the context element admit multiple interpretations and conflict with guidance for building assured safety arguments. If readers do not share an understanding of the meaning of context that makes context's impact on the main safety claim clear, confidence in safety might be misplaced. In this paper, we analyse the definitions and usage of GSN context elements, identify contradictions and vagueness, propose a more precise definition, and make updated recommendations for assured safety argument structure.

**Keywords:** Assurance argument, safety case, safety argument, goal structuring notation, context, confidence, assured safety argument.

## 1 Introduction

Developers of some safety-critical systems develop a *safety case* that contains both safety evidence and an argument linking that evidence to safety claims [1,2]. The *Goal Structuring Notation* (GSN) is a popular graphical notation for recording these *safety arguments* [2,3,4]. One of GSN's key innovations is a *context* element for linking to contextual information (which is not necessarily about the system's operating context). However, definitions of context in GSN admit multiple interpretations. Moreover, a recent proposal for a clearer argument structure, namely *assured safety arguments*, demonstrates an understanding of context elements that is at odds with existing definitions [5]. If argument readers and writers do not share an understanding of the meaning of context that makes context's impact on the truth of the safety claim clear, confidence in the safety claim might be misplaced with disastrous consequences. This paper makes four contributions toward a clearer understanding of context in GSN arguments:

- A review of the definitions and uses of context elements in GSN
- Identification of contradictions and vagueness in existing notions of context
- A precise definition in terms of normative models of inductive argument
- Recommendations for applying the proposed definition, including new guidance for structuring assured safety arguments

A. Bondavalli and F. Di Giandomenico (Eds.): SAFECOMP 2014, LNCS 8666, pp. 139–154, 2014.

In Sect. 2, we analyse the definitions of GSN context elements given in authoritative sources and show that these admit multiple interpretations. In Sect. 3, we examine GSN context elements as used in *assured safety arguments*. In Sect. 4, we show that these definitions are inherently contradictory and explore the consequences of that contradiction. In Sect. 5, we propose and defend a definition of context in GSN given in terms of normative models of inductive argument. Finally, we discuss related work in Sect. 6 and conclude in Sect. 7.

## 2 Context in the Goal Structuring Notation

In some domains, developers of critical systems construct an *assurance case*. When the critical property is safety, assurance cases are known specifically as safety cases. A safety case is a 'structured argument, supported by a body of evidence, that provides a compelling, comprehensible and valid case that a system is safe for a given application in a given environment' [1, Sect. 9.1]. The argument explains how the evidence relates to safety objectives [2, Sect. 1.2.1].

GSN is one of two popular graphical notations for recording assurance arguments [3,4]. Figure 1 presents an example that the *GSN Community Standard* gives to illustrate the notation [3]. *Goal* element G1 presents the argument's main claim. Arrows with filled heads indicate that G1 is *SupportedBy* goals G2 and G3: the control system is deemed acceptably safe because all identified hazards have been eliminated or sufficiently mitigated and the software has been developed to an appropriate safety integrity level. *Strategy* S1 explains *how* goals G4–G6 support goal G2. *Solution* Sn1 provides evidence supporting the claim in G4. Context elements C1 and C2 are asserted at goal G1 using the open-ended *InContextOf* arrow. This paper considers the function of such context elements.

Consider three potential interpretations of the meaning of asserting C2 at G1: (1) the arguer asserts that the system as operated matches the referenced definition, (2) the arguer is identifying the system the argument is about, and (3) the arguer is identifying a document that the reader can refer to for details about the system. These alternatives have different impacts on the argument's soundness: (1) presents a claim that must be checked because false premises undermine conclusions; (2) is clarification that cannot be said to be true or false; and (3) has indeterminate impact because a reader could look up anything. We now turn to normative sources for help choosing the correct interpretation.

### 2.1 Kelly's *Arguing Safety*

One of the first specifications of GSN appeared in Kelly's DPhil Thesis [2]. Kelly introduces context elements into GSN 'in order to be able to represent the context in which a safety argument is stated and, thus, how the argument relates to, and depends upon, information from other viewpoints' [2, Sect. 3.3]. Context elements have 'two possible forms: as a reference *to* contextual information [and] as a statement *of* contextual information'. Providing context elements 'allows reference to where [the concepts used in a goal] are fully defined' [2, Sect. 3.5.2].

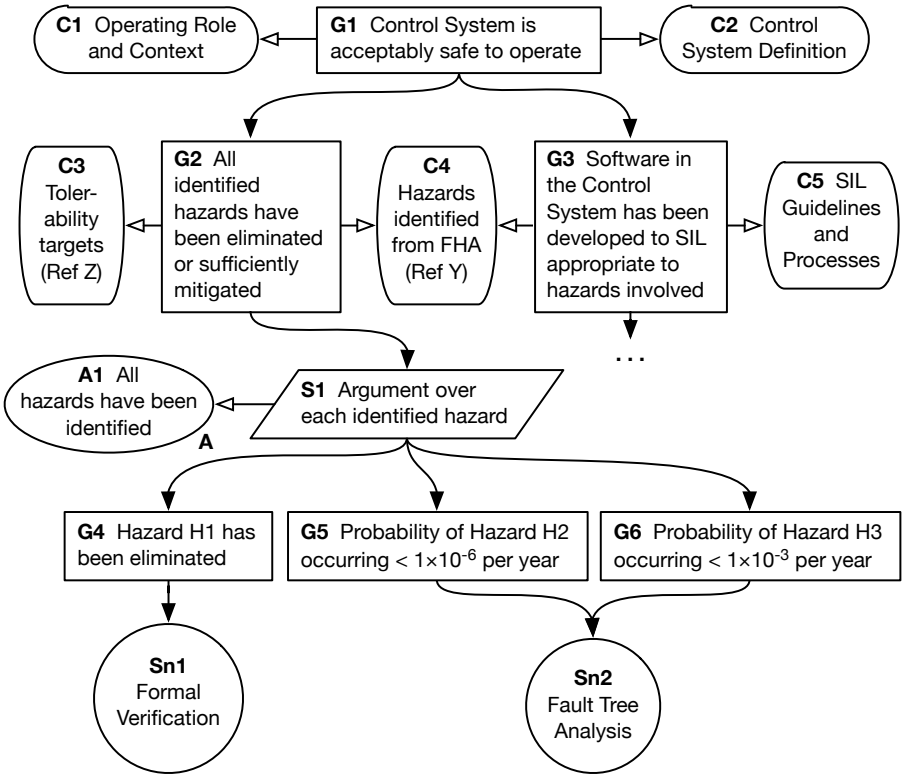


**Fig. 1.** Extract from 'An Example Goal Structure' (Fig. 6 from [3])

But the thesis does clearly describe how context affects the meaning and soundness of arguments. Kelly introduces Toulmin's normative model of informal, inductive argument [6] as background [2, Sect. 2.6.3], but does not describe the function of context elements in terms of any normative argument model. Instead, he gives the examples depicted in Fig. 2 [2, Sect. 3.3 (emphasis mine)]:

> The claim that all applicable hazards have been complied with [sic] is set in the context of whatever is determined as an applicable standard. C1 ... refers to the set of standards identified as applicable (e.g. pointing to the document or file location / section where applicability is discussed and defined). The second example shows an argument ... (S1) ... over ... all hazards. ... S1 is only truly defined when *the basis over which it is stated is made clear*. C2 refers to where the identified hazards are discussed and defined within the supporting safety case documentation. The [third] example ... shows context being used to *communicate the basis* on which a piece of evidence (solution) is being put forward. ... C3 makes clear that the fault tree evidence referred to by Sn1 depends upon

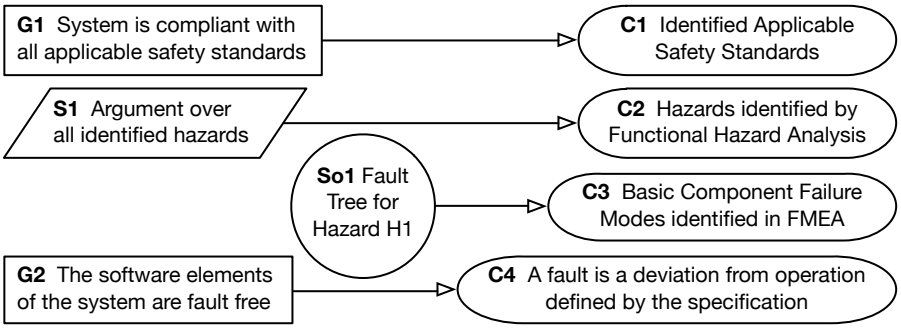


**Fig. 2.** Combination of 'Example Uses of GSN Context' (Fig. 26 from [2]) and 'Example Use of Context Statement' (Fig. 27 from [2])

> the failure rates provided by the more primitive FMEA (Failure Modes and Effects Analysis) evidence. ... [The fourth example illustrates] an 'immediate' contextual statement used to *clarify the basis* of [a] goal .... C4 is phrased as a statement that helps *define ... the basis* of G2. Without C4, ... a reader of G2 may adopt an alternative meaning.

We will return to what might be meant by 'define the basis' in Sect. 2.3.

In GSN, context asserted at a goal is *inherited* by all goal, strategy, and solution elements supporting that goal. Considering the example of an 'argument over all identified hazards' strategy expressed in the context of a hazard log, Kelly writes that 'all the goals and solutions underneath are *also* expressed in the context of the hazard log' [2, Sect. 4.4.3.2].

### 2.2 The GSN Community Standard

More than a decade after GSN's introduction, a consortium of GSN users wrote the *GSN Community Standard* to 'provide a comprehensive, authoritative definition of the Goal Structuring Notation' [3]. The standard introduces context elements by noting that 'when documenting a GSN *goal* or *strategy* it can also be important to capture the context in which the claim or reasoning step should be interpreted. This is done in GSN by documenting context'. Like *Arguing Safety*, the standard explains that a context element may contain 'a reference to contextual information or a statement' [3, Fig. 7]. Part 1 (which defines GSN) offers no normative model of how context affects the meaning and validity of arguments [3]. However, Part 2 clarifies what context *isn't* [3]:

- 'In GSN, *context* elements should not be used to refer to information which is intended to support the validity of a claim. Such information ... should be represented using a GSN *solution* element' [3, Sect. 2.6.2.1].
- '*Context* elements are sometimes used where a GSN *assumption* or *justification* may be more appropriate' [3, Sect. 2.6.2.3].

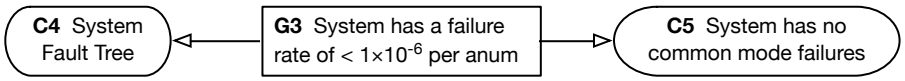


**Fig. 3.** Combination of 'Incorrect Use of Context (as a Solution)' (Fig. 52 from [3]) and 'Incorrect Use of Context (as an Assumption)' (Fig. 53 from [3])

Figure 3 depicts the example with which the standard illustrates these points. 'Context C4 is incorrectly associated with Goal G3 as evidence [supporting] the failure rate claim .... The correct way to represent this relationship is to associate the System Fault Tree with Goal G3 as a GSN *solution*. ... [Context C5] would be more appropriately rendered as an *assumption*' [3, Secs. 2.6.2.2–3].

Like *Arguing Safety*, the standard states that context is inherited: 'contextual information associated with a claim made in a particular goal is understood to be in scope for all sub-goals of that goal' [3, Sect. 2.3.3.4]. Discussion with Kelly suggests a fourth possible meaning of context: what is inherited is the understanding created by asserting contextual information at a claim, not the contextual information itself [7]. Returning to C2 in Fig. 1, that interpretation would be that it is the clarification of 'Control System' created by asserting C2 at G1, not the control system definition itself, that is inherited by G2, G3, etc.

Unlike *Arguing Safety*, the *GSN Community Standard* explicitly addresses conflicting context: 'nothing in the supporting argument for the goal to which the context is applied should contradict or undermine the relationship between the goal and the context' [3, Sect. 1.3.7, emphasis removed].

### 2.3 Interpreting GSN's Definition of Context

*Arguing Safety* and the GSN Community Standard are the two most authoritative definitions of GSN. Given what they say about context elements, we return to the assertion of C2 as context for G1 in example argument given in Fig. 1 and discussed in Sect. 2. Interpretation (1) of this context assertion as a claim that the system as operated matches the referenced definition cannot be correct because it contradicts the prohibition on context introducing information on which the validity of a claim depends [3, Sect. 2.6.2.1]. Interpretation (2) of the context assertion as identifying the system the argument is about seems plausible because it allows reference to where the concepts in G1 are defined [2, Sect. 2.6.3]. But interpretation (3) of context C2 *also* seems plausible. Kelly's use of the phrase 'define the basis of' [2, Sect. 3.3] and the standard's use of the phrase 'capture the context in which the claim ... should be interpreted' [3, Sect. 2] seem to suggest that the reader should keep the entire contents of the control system definition in mind when interpreting G1 and the entirety of the argument supporting it. Those contents cannot be used as a premise, but might presumably clarify the meaning of any part of the argument.

Some examples in *Arguing Safety* seem to be very clearly intended to be read using interpretation (2). For example, a context element in one example reads,

‘“Sufficient” = platform meets target failure rate of $1 \times 10^{-6}$ per flight hour’ [2, Fig. 44]. Context C1 in Fig. 1 seems to be more consistent with interpretation (3): there is no mention of the system’s operational role or operating context in goal G1. Other examples don’t seem to clearly fit either of those interpretations. For example, context C3 in Fig. 2 seems to be better explained as documenting the *provenance* of an evidence item than as explaining the meaning of the text in solution So1 or offering information that would help interpret the evidence. One might regard C3 as explaining what is meant by ‘fault tree’, but simply knowing the failure modes would not help to interpret the strength and meaning of that evidence. It is knowing the provenance of the fault tree – which would not be documented in the referenced FMEA results – that would aid this interpretation.

## 3 Assured Safety Arguments

Hawkins et al. have proposed *assured safety arguments* as a means of more clearly communicating both (1) how evidence supports system safety claims and (2) why that argument establishes sufficient confidence in the main safety claim [5]. An assured safety argument contains two distinct sub-arguments:

1. ‘A *safety argument* that documents the arguments and evidence used to establish direct claims of system safety’
2. ‘A *confidence argument* that justifies the sufficiency of confidence in this safety argument’ [5, emphasis mine]

Later discussion [7] resulted in adding a *conformance argument* to document how developers interpreted and conformed with relevant standards [8].

### 3.1 Structure of an Assured Safety Argument

Assured safety arguments simplify and clarify the safety rationale by relocating information that does not explain how evidence supports the safety claim. Information that increases confidence – by, for example, testifying to the quality or relevance of the evidence – is presented in a separate confidence argument. *Assurance Claim Points* (ACPs) link inferences, evidence assertions, and context assertions in the safety argument to relevant parts of the confidence argument.

Figure 4 reproduces an example used to illustrate ACPs. The square decorations ACP.A4, ACP.A1, and ACP.A3 identify the assertion of context elements DIP.A4, DIP.A1, and DIP.A3, respectively, at goal DIP.G1. ACP.S1 identifies the inference of DIP.G1 from premises DIP.G2–DIP.G6 using the argument strategy DIP.S1. ACP.A2 identifies the assertion of context DIP.A2 at that inference step. (The diamond decorations on goals GIP.G2–GIP.G6 are from GSN’s pattern extension and indicate that these goals require support that is not shown here [3].)

Each ACP is a pointer to a separate portion of the confidence argument. Figure 5 reproduces an example Hawkins et al. give to illustrate assurance arguments [5, Fig. 17]. Goals CC1.3 and CC2.3 are associated with ACP.A1; together,

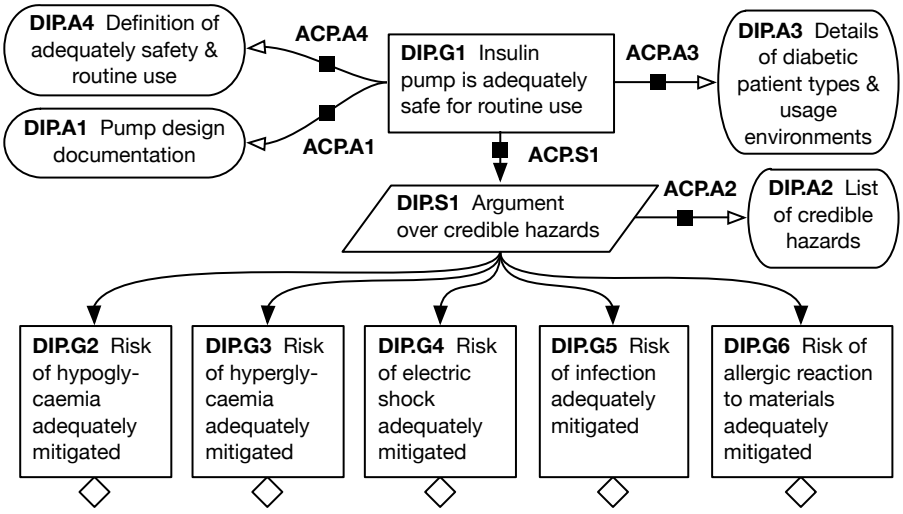


**Fig. 4.** 'High-level safety argument for an insulin pump' (Fig. 16 from [5])

the arguments supporting them show why we can have confidence in the assertion of 'pump design documentation' as context for the claim that the 'insulin pump is adequately safe for routine use'. (Presumably, given other patterns in the paper, ACP.A1 is attached to a different goal not shown in the original figure. That goal would read 'sufficient confidence exists in the assertion of DIP.A1 as context at goal DIP.G1' and be supported by CC1.3 and CC2.3.)

### 3.2 Confidence Argument Structure

Figure 6 reproduces Hawkins et al.'s illustration of a confidence argument's top-level structure. The argument claims that confidence in the safety argument's main safety claim is justified because each of the safety argument's components (inferences, solutions, and context) is fit for the purpose it serves. Instantiations of confidence patterns of the kind shown in Fig. 5 demonstrate that fitness.

There are several ways to describe confidence in assurance claims, each with its own benefits and drawbacks [9]. Hawkins et al.'s confidence argument patterns use a form of Baconian probability [9,10,11]. That is, they enumerate plausible *defeaters* of the argument – things that might directly rebut a claim or undermine the reasoning supporting it – and describe why those defeaters are thought to be implausible and/or the residual likelihood of them acceptable. (Some small degree of doubt is inevitable: even a machine checked, deductive proof might be wrong if the proof checker is faulty or was used improperly.)

### 3.3 Context as Used in Assured Safety Arguments

Hawkins et al. describe the meaning of a context element linked to a goal as an

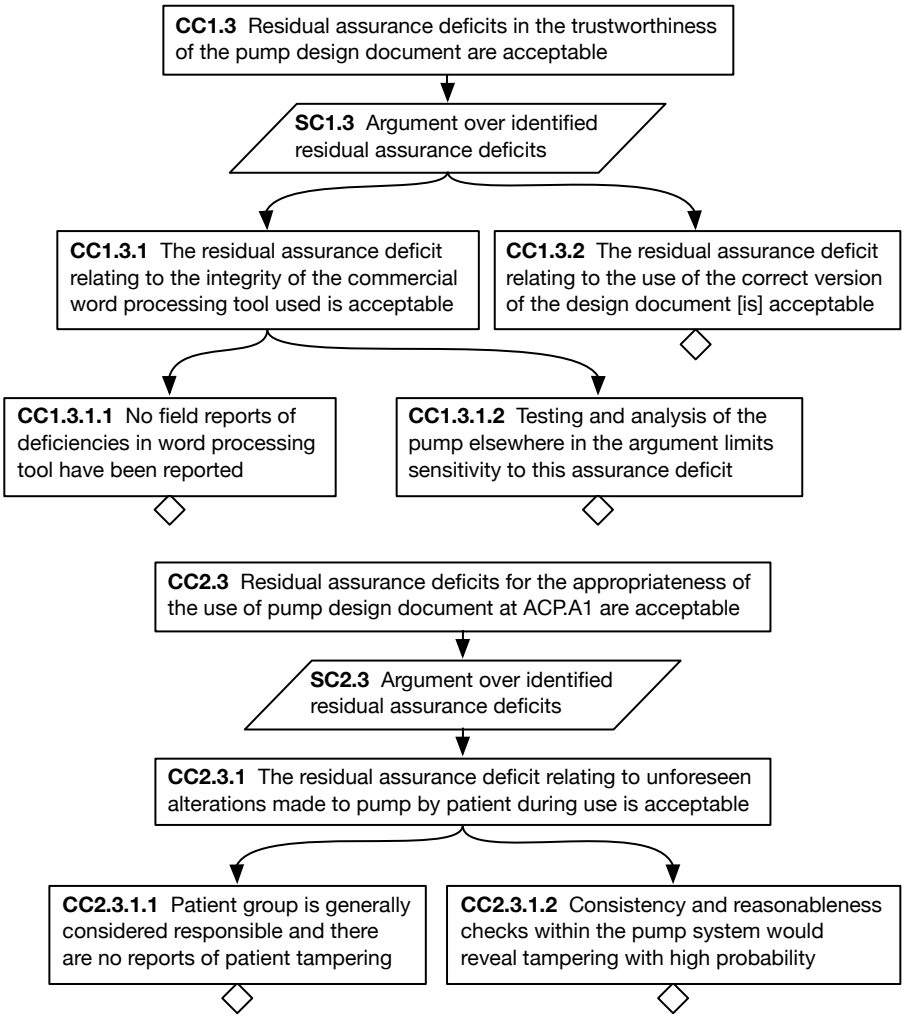


**Fig. 5.** 'Part of the confidence argument for ACP.A1' (Fig. 17 from [5])

> assert[ion] that the context is appropriate for the elements to which it applies. For example, consider a context reference to a list of failure modes for a particular piece of equipment. The introduction of this context element when arguing about the safety of that piece of equipment implicitly asserts that the list of failure modes referred to is appropriate to the application and operating context in question.

The appropriateness of context must be considered throughout the part of the argument that inherits the context: 'the assurance of the strategy depends upon the confidence that the context ... stated is appropriate for that strategy and its

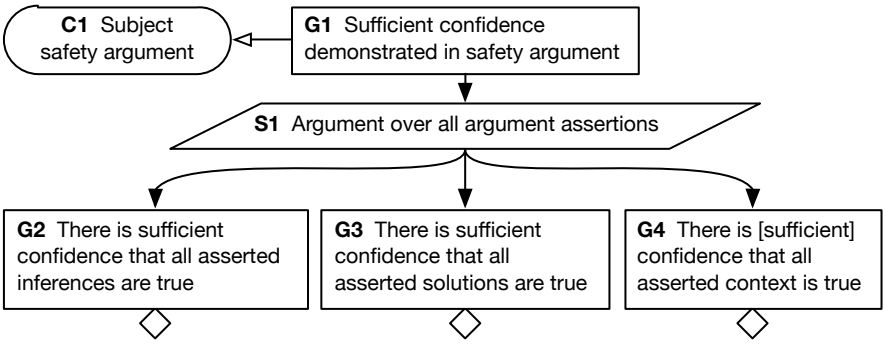


**Fig. 6.** 'Representing an overall confidence argument' (Fig. 15 from [5])

subgoals' [5]. But considering the appropriateness of context alone is insufficient: 'in addition to the appropriateness of the context, it is also necessary to provide an argument as to the trustworthiness of the context in question'.

Hawkins et al. do not provide a testable definitions of what it means for context to be 'appropriate' and 'trustworthy', although in the latter case they say that 'the concept of trustworthiness relates to freedom from flaw' [5]. However, they do provide examples. Referring to a generic argument over all hazards, they write that for a hazard list to be appropriate context 'there must be confidence that the hazard list is appropriate with respect to the system, application, and context'. In the case of the example shown in Fig. 4, they give the confidence argument fragment depicted in Fig. 5. They also write of the meaning of ACP.A2 that 'there is sufficient confidence that the list of credible hazards is complete and correct. Inadequate definition of a hazard or omission might invalidate the safety claim'. The context must be 'true', as goal G4 in Fig. 6 puts it.

The assertion of a context element in an assured safety argument seems to be mostly clearly defined as the making of two claims:

1. *Acceptable instantiation.* The identified thing *is* the kind of thing implied by the ordinary meaning of the term used to represent it.
2. *Fitness for role.* The identified thing has all of the properties that the entire applicable portion of argument needs it to have.

These claims then serve as implicit premises throughout the inheritance area. In Toulmin's terms [6], to assert a context element in an assured safety argument is to assert that acceptable instantiation and fitness for role are warrants that can be implicitly used in any of the affected reasoning steps.

To illustrate this definition of context, consider the example in Fig. 4. A hazard list must be *the* hazard list created for the system in question because the ordinary meaning of the words 'all credible hazards' is that they are hazards of the system in question. The hazard list must also be 'complete and correct' because inference DIP.S1 would be invalid if supporting goals DIP.G2–DIP.G6 did not accurately portray all relevant hazards. The hazard list can be assumed to have these properties throughout the argument supporting goals DIP.G2–DIP.G6.

Acceptable instantiation and fitness for role are not simply properties that must be true as prerequisites for judging whether an argument is sufficiently compelling. There are such properties, for example argument clarity and comprehensibility. But the example argument depends upon these as premises. DIP.G1 in Fig. 4 lacks a sub-goal or justification claiming that the list of hazards identified by DIP.A2 is complete and correct. Such support is used in similar reasoning steps in plain safety arguments, including in Hawkins' *High Level Software Safety Argument Pattern* [12]. The absence of such support here can only mean that the assertion of DIP.A2 is meant to demonstrate that there are no credible hazards that are not covered by one of the goals DIP.G2–DIP.G6.

## 4 The Problem of Conflicting Definitions of Context

Section 2.3 discussed how context is (somewhat vaguely) defined in authoritative guides to GSN. Section 3.3 showed that context elements in assured safety arguments function as claims of acceptable instantiation and fitness for role. These definitions are mutually exclusive. Clarification and identification of reference material cannot introduce new claims. Introduction of claims contradicts the prohibition on using context elements 'to refer to information which is intended to support the validity of a claim' [3, Sect. 2.6.2.1].

It is vital that all readers of an argument understand the same meaning of its context elements. If they do not, confidence in safety claims might be misplaced. For example, consider multiple reviewers collaborating to review of a large argument in parts [13]. Suppose that reviewer $A$ examines the assertion of a hazard list document identifier as context and interprets it as explaining the term 'hazard list'. Suppose that reviewer $B$ reviews a supporting portion of argument and interprets the context assertion as claims of acceptable instantiation and fitness for role. Because $A$ sees no need to check either property and $B$ assumes that they have been checked, neither will check it. The system might be put into service despite not addressing a significant hazard.

Returning to the example in Fig. 1, suppose that reviewer $C$ interprets this 'basis' of goal G1 as simply scoping the situations to which the argument applies. Suppose that stakeholder $D$ reads in the referenced documentation a claim about what the operating context *is*. $D$ might assume that review had confirmed that it was acceptable to assume that the system would be used in this way while $C$ might not see the need to check that assumption.

## 5 Proposed Treatment of Context and Confidence

Given the harm that misinterpretation of context might bring, GSN users should adopt a single, normative definition. This section gives and justifies our proposal.

A useful definition of GSN context elements must satisfy two requirements:

1. *Means to perform the functions that people have been using GSN context elements to perform must be preserved.* If the definition precludes using context elements to meet a need that GSN users have used them to meet, we must also propose an alternative means of meeting that need.

2. *The effect of context elements on confidence in the argument's main claim must be well defined.* Understanding this effect is a precondition for defining an effective argument review process and, ultimately, for using an argument to make certification or acceptance decisions.

### 5.1 Our Proposal: GSN Context Elements as Explications

We propose defining context elements as explicating terms used in the argument.

*Form.* Context element text must be of the form '$X$: $Y$' where $X$ is a phrase and $Y$ is its explication. $Y$ should identify relevant documentation where appropriate.

*Scope.* The explication applies to (i) the element $e$ at which the context $c$ is asserted, (ii) any goal, strategy, or solution in the same argument module that directly or indirectly supports $e$ through IsSupportedBy relationships, and (iii) any justification, assertion, or confidence element in the same module asserted as context to an element to which $c$ applies as per rules (i) and (ii).

*Effect.* Arguments should be understood as if explicated terms were replaced by their explications.

*Uniqueness.* Arguers may not assert two explications for the same term that apply to the same element.

*Non-circularity.* Arguers may not assert explications such that any term is directly or indirectly explicated in terms of itself.

*Presentation.* Explicated terms appearing in GSN elements should be visually distinguished from non-explicated text. For example, explicated terms might be presented in a different font, in italics, in a different colour, underlined, or some combination of these. Hyperlinks should be used where practicable.

*Loaded language.* Arguers should not use context to phrase arguments in terms whose plain-language meaning might cause misunderstanding of the argument.

### 5.2 An Illustrative Example

To illustrate the proposal given in Sect. 5.1, consider the example given in Fig. 4. Figure 7 presents a version of that argument revised to reflect our definition of context. Context elements C1–C3 now clearly explicate terms used in goal G1. C1 clarifies which insulin pump we mean and that we mean it as delivered, not just as designed. C2 refers to documentation giving the relevant definition of 'adequately safe'. C3 clarifies what we mean by 'routine use', thus limiting the scope of the argument to that use. For clarity, we introduce goal G2 to separate (a) the inductive leap related to the relationship between safety and hazard management from (b) the argument-by-cases over the set of identified hazards.

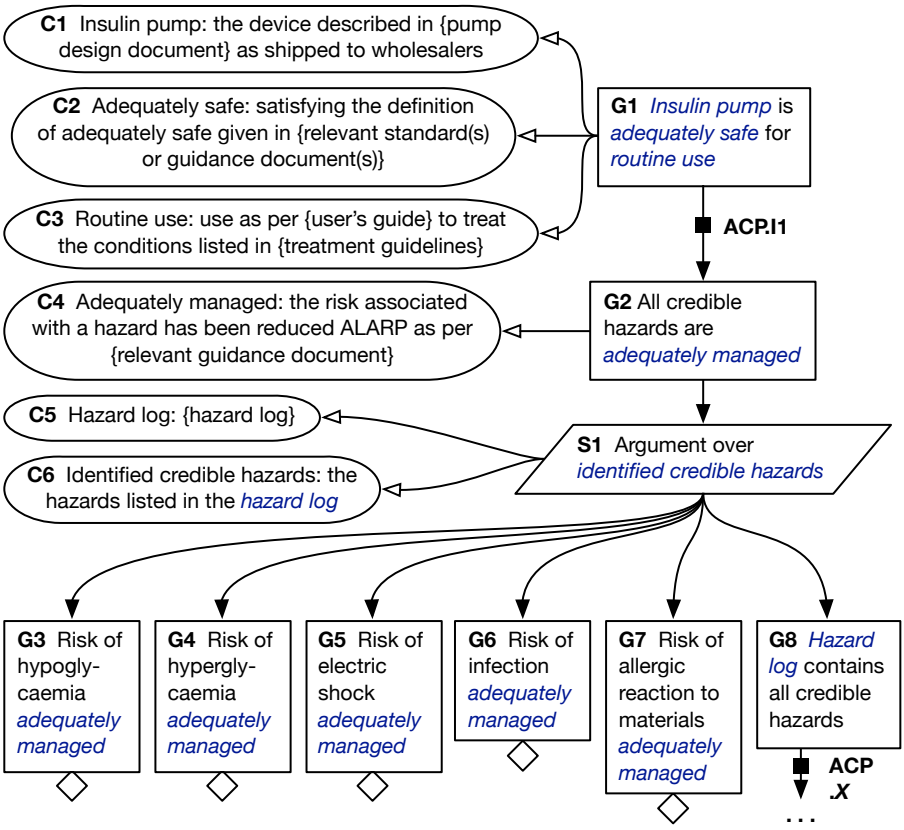


**Fig. 7.** Revised version of argument given in Fig. 4. The text in braces stands for details that identify the documents in question (e.g., document and version numbers).

New goal G8 is the claim of hazard log completeness that justifies strategy S1. G8 could be replaced by a justification or an away goal asserted as justification using the InContextOf relationship [12]. Those alternatives have the stylistic benefit of distinguishing the claim about hazard log completeness from the claims to have managed each hazard. However, those alternatives also have an area effect that complicates both change management and argument review.

The *scope* rule dictates that explications of 'insulin pump', 'adequately safe', and 'routine use' are applicable in all goals, strategy S1, and context elements C4–C6. The *uniqueness* rule would preclude asserting a competing explication of 'adequately managed' at goals G2–G8 or strategy S1. The *loaded language* recommendation suggested the change from 'credible hazards' in the original to 'identified credible hazards' in S1.

### 5.3 Assessing the Proposal: Performing All Context Functions

The examples in *Arguing Safety* [2], the GSN community standard [3], and the assured safety argument paper [5], Hawkin's software safety argument patterns [12], and our own experience suggest at least six functions that arguers *might* want context elements to perform: (1) to explain a term's meaning, (2) to link arguments to other documents, (3) to assert an implicit premise, (4) to identify background information, (5) to document the circumstances in which the argument was made, and (6) to make GSN elements less verbose.

*Explaining terms.* Our proposed definition clearly serves this purpose.

*Linking the argument to documents.* Our definition facilitates linking to explain terms more clearly. Solution elements link arguments to evidence. Linking without a clear purpose is disallowed to prevent confusion. We would replace C3 in Fig. 2 with an ACP on the solution linking to an argument that explains the fault tree's provenance and that provenance's effect on confidence.

*Asserting an implicit premise.* GSN offers two ways to assert a premise throughout an argument (i.e., as an implicit warrant in Toulmin's model [6]): justification elements and away goals asserted as context. In any case, this function might be overused. Local scope simplifies change and review and we might only need the goal G8 in Fig. 7 as a premise in this particular reasoning step.

*Identifying background information.* Background information can help to understand and validate an argument (e.g., show that it is not oversimplified [14]). But premises should be introduced using a goal, away goal, or justification element, links to details can be made as described above, and it is not clear that 'background information' serves any other useful purpose. Categorising information as either what-the-arguer-means (which can be accepted) or as evidence claims (which must be checked) facilitates argument validation.

*Documenting the argument writer's circumstances.* Documenting the circumstances under which an argument was made might aid interpretation. But much of this (e.g., the colour of the author's clothing) is irrelevant. Moreover, different people might interpret the remainder differently. Restricting context elements to explication forces arguers to identify which meanings are influenced by circumstance and (more importantly) what those meanings are meant to be.

*Making GSN elements less verbose.* A single artefact might serve multiple roles in an argument. For example, a hazard log might serve as a list of hazards, information about hazard severity, or an indication of project status [7]. Explication-only context might not reduce verbosity as well as unrestricted context because authors might have to reword element text to include the explicated term or assert context multiple times to fill multiple roles. We consider this an acceptable price for increased clarity.

**Table 1.** How our definition of context solves other definition's problems

| Problem | Solution |
|---|---|
| The effect on argument confidence is unclear. See Sect. 4. | The assurance argument should be judged as if explicated terms had been replaced by their explications. See Sect. 5.4. |
| A document asserted as the basis for a goal (and possibly all of its supporting reasoning, depending on how inheritance is interpreted) could be understood by different people as explaining different things. See Sect. 2. | The context explicitly identifies the term being explicated. The *presentation* recommendation reminds readers that a term is explicated. See Sect. 5.1. |
| Readers might interpret different parts of a linked document as the explanation of a term. See Sect. 2. | Arguers understanding context as explication will craft explications to resist misinterpretation. Reviewers will help by pointing out vagueness. Guidance created to clarify terms used in requirements might help further, e.g. by eliminating hedge words such as 'usually' or 'generally' [15]. |
| Readers might or might not interpret context assertion as a claim that referenced material is what its title suggests and is fit for purpose. See Sect. 3.3 and Sect. 4. | Our definition precludes this interpretation. The *loaded language* recommendation, the *presentation* recommendation, and appropriate review reduce the risk that the explicated term's plain-language meaning will colour understanding of the argument. See Sect. 5.1. |

### 5.4 Assessing the Proposal: The Effect on Confidence

The main virtue of the definition of context in Sect. 5.1 is a well-defined impact on argument confidence: the argument should be assessed as if all explicated terms had been replaced by their explications. In Toulmin's terms, context as we define it is simply a mechanism for replacing shorthand text used in data, warrants, claims, reservations, and qualifications [6]. Table 1 shows how our proposal addresses the confidence-related problems that other definitions have.

### 5.5 Impact on Assured Safety Arguments

The organisation of the confidence argument (depicted at the top level in Fig. 6) must change: context as defined in Sect. 5.1 cannot be said to be 'true' or false. There is no need to argue over context elements because we argue instead over the elements and relationships whose meanings they clarify. The burden of demonstrating hazard log completeness, carried at ACP.A2 in the original formulation in Fig. 5, is carried by the evidence and inferences supporting goal G8 in Fig. 7.

But GSN also allows justification elements, assumption elements, and away goals to be the object of InContextOf relationships. Goal G4 in Fig. 6 would cover

those assertions if it read 'there is sufficient confidence that all assumptions, justifications, and away goals asserted as context are true'.

One might argue the need for an assurance claim point on the inference of goal G2 from goals G3–G8 through strategy S1. The associated confidence argument fragment would cite review evidence showing that goals G3–G7 cover all of the credible hazards listed in the hazard log. It is obvious that such a review will be performed. Since the inference admits no other assurance deficit, we see no reason to burden the arguer with writing such a confidence argument fragment.

### 5.6 Further Recommendations: Update Review Processes

An explication cannot be either true or false, but a poor explication might admit multiple interpretations, compromising the efficacy of argument review. Existing argument review processes include steps aimed an ensuring clarity [3,13,16]. GSN argument review processes should require reviewers to consider whether terms used in the argument have multiple meanings in general use and in the relevant technical domain(s). Terms with multiple meanings should be explicated, and explications should rule out unintended meanings of the explicated terms.

## 6 Related Work

*Arguing Safety* [2] and the *GSN Community Standard* [3] define context in GSN. Examples of context in the former clearly explicate terms. What is novel about our proposed definition of context is that we limit context to this function, thus making its impact on argument confidence clear.

Matsuno and Taguchi's proposed formalisation of GSN patterns [17] defines context elements as declarations of types and variables. The definition of context proposed in Sect. 5.1 is for arguments that have not been formalised.

Because the other popular graphical argument notation, CAE, has no context element, our proposed definition of context does not apply to it [4]. For similar reasons, our proposal does not apply to plain text or tabular arguments. Any informal argument might be vague, but other causes will apply in other notations.

## 7 Conclusion

In this paper, we reviewed how both Kelly's thesis [2] (the original normative definition of GSN) and the *GSN community standard* [3] define context elements. Neither defines context in terms of normative models of argument and both permit multiple interpretations. But both sources are clearly at odds with how context elements are treated in assured safety arguments: the former say that context elements cannot support the validity of claims, while the latter says that they do. To resolve this contradiction and bring clarity to the meaning of GSN, we proposed a more precise definition of the semantics of GSN context elements. We illustrated this definition and its impact on the structure of assured safety arguments by reworking a published example of an assured safety argument.

Any proposed change to language semantics – whether for a natural language, a programming language, or an argument notation – will fail if people choose not to adopt it. A key factor in the adoption of this change is whether the proposal addresses all of the functions for which arguers have been using GSN context elements. We have examined examples for evidence of such functions and found none, but very few published examples exist. The only practical way forward is to make this proposal public so that arguers can judge for themselves.

**Acknowledgments.** This research was funded by the Swedish Foundation for Strategic Research as part of the SYNOPSIS project and by the EU/Artemis as part of the nSafeCer project (grant 295373). We thank Pierre Loisy for inspiring this work and Tim Kelly and Iain Bate for helpful discussions of this paper.

## References

1. Defence Standard 00-56: Safety Management Requirements for Defence Systems, Issue 4, Part 1: Requirements (U.K.) Ministry of Defence (June 2007)
2. Kelly, T.P.: Arguing Safety — A Systematic Approach to Managing Safety Cases. DPhil thesis, University of York (September 1998)
3. Attwood, K., et al.: GSN Community Standard Version 1. Origin Consulting Limited, York (November 2011)
4. Bishop, P., Bloomfield, R.: A methodology for safety case development. In: Proc. Safety-Critical Systems Symposium (SSS) (1998)
5. Hawkins, R., Kelly, T., Knight, J., Graydon, P.: A new approach to creating clear safety arguments. In: Proc. Safety-Critical Systems Symposium (SSS), pp. 3–23 (2011)
6. Toulmin, S.E.: The Uses of Argument, Updated edn. Cambridge University Press, New York (2003)
7. Kelly, T.: Personal communication
8. Graydon, P., Habli, I., Hawkins, R., Kelly, T., Knight, J.: Arguing conformance. IEEE Software 29, 50–57 (2012)
9. Graydon, P.J.: Uncertainty and confidence in safety logic. In: Proc. Int'l System Safety Conference (ISSC) (2013)
10. McDermid, J.A.: Risk, uncertainty and software safety. In: Proc. Int'l Systems Safety Conference (ISSC) (2008)
11. Weinstock, C.B., Goodenough, J.B., Klein, A.Z.: Measuring assurance case confidence using Baconian probabilities. In: Proc. Int'l Wkshp. on Assurance Cases for Software-Intensive Systems (ASSURE) (2013)
12. Hawkins, R., Kelly, T.: A software safety argument pattern catalogue. Technical Report YCS-2013-482, University of York (2013)
13. Graydon, P., Knight, J., Green, M.: Certification and safety cases. In: Proc. Int'l Systems Safety Conference (ISSC) (2010)
14. Greenwell, W.S., Knight, J.C., Holloway, C.M., Pease, J.J.: A taxonomy of fallacies in system safety arguments. In: Proc. Int'l System Safety Conference (ISSC) (2006)
15. Wasson, K.S.: CLEAR Requirements: Improving Validity Using Cognitive Linguistic Elicitation and Representation. PhD thesis, University of Virginia (2006)
16. Kelly, T.: Reviewing assurance arguments — a step-by-step approach. In: Proc. Wkshp. on Assurance Cases for Security — The Metrics Challenge (July 2007)
17. Matsuno, Y., Taguchi, K.: Parameterised argument structure in GSN patterns. In: Proc. Int'l Conf. on Quality Software (2011)