RAIN
ROBOTICS AND AI IN NUCLEAR

MANCHESTER 1824
The University of Manchester
Autonomy and Verification Group

Safety Cases for the use of Autonomous Systems in Nuclear Environments

Chris Anderson

Research Associate

RAIN March 2022 Meeting 16th March 2022

Overview

Motivation

- Very little knowledge of how to construct safety cases for robots utilising autonomy and AI in civil nuclear applications
  - The safety case for some types of autonomy is well understood
- A large part of this is that few understand
  - the technology, and
  - how to construct a safety case

A high level overview of:

- actions that can be taken now
- a pointer to possibilities for the near to medium term

Consideration:

- For most nuclear tasks the environment is well constrained, but still there are challenges.

Basis for Safety Case strawman

ROBOTICS AND AI IN NUCLEAR

MANCHESTER 1824
The University of Manchester
Autonomy and Verification Group

A safety case framework

- based very loosely on A2I2 (Lilypad ASV + BlueROV)
- hypothetical surface vehicle:
  - utilises AI
  - carries out a survey of the spent fuel storage pond
- addresses an assumed hazard
  - collision
- recognises that there potentially are other hazards. e.g.:
  - propeller splash
  - unretrievable due to complete robot system failure
  - explosion due to H_2 evolution at the surface of the pond

- Define the task
- Formally identify and analyse hazards and place in the Hazard Log with a tolerable and ALARP mitigation strategy

ROBOTICS AND AI IN NUCLEAR

Other Robotic Systems

The University of Manchester
Autonomy and Verification Group

Other robotic systems were considered for this work. e.g. Vega

- Teleoperated
- No autonomy
- Deployed in a vent channel at Dounreay to survey contamination
- safety case (summary):
  - Hazards requiring mitigation. e.g. lanyard for recovery
  - Hazards requiring no mitigation. e.g. collision

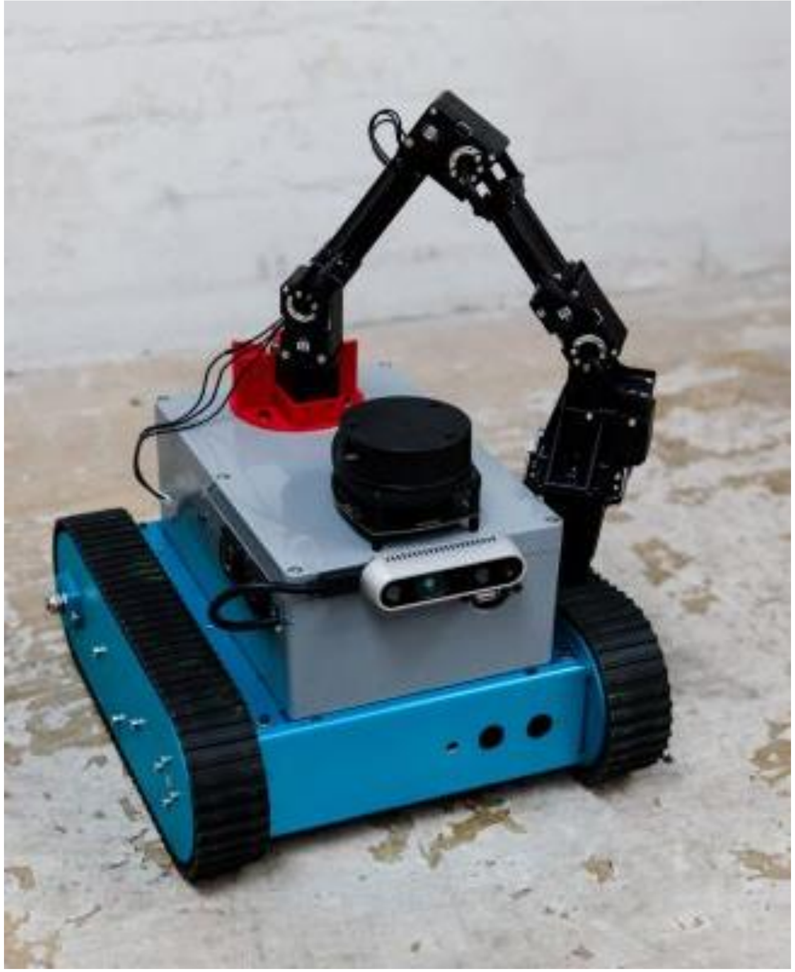

ROBOTICS AND AI IN NUCLEAR

Identification of Hazards

The University of Manchester
Autonomy and Verification Group

- It is important that the hazards are identified and analysed holistically for the robot within its application task and environment

  This applies to:
  - existing robots (although substantiation can be very difficult/impossible)
  - proposed robots
- Hazards should be identified for all phases/tasks of the robots lifecycle
  - *(design, build and test)*
  - commissioning
  - operation (for now this is on-site testing)
  - recovery
  - maintenance
  - disposal

  Addressing normal and abnormal operations
- Apply high level principles (which the site licensee will have) to identify hazards and determine completeness

Bow Tie Diagram

Most of our robotic autonomy efforts!

Adapted from: Niklas Möller et al. (2018). Handbook of Safety Principles, https://doi.org/10.1002/9781119443070

RAIN
ROBOTICS AND AI IN NUCLEAR

Tolerability and ALARP

MANCHESTER 1824
The University of Manchester
Autonomy and Verification Group

Off-site (General public)

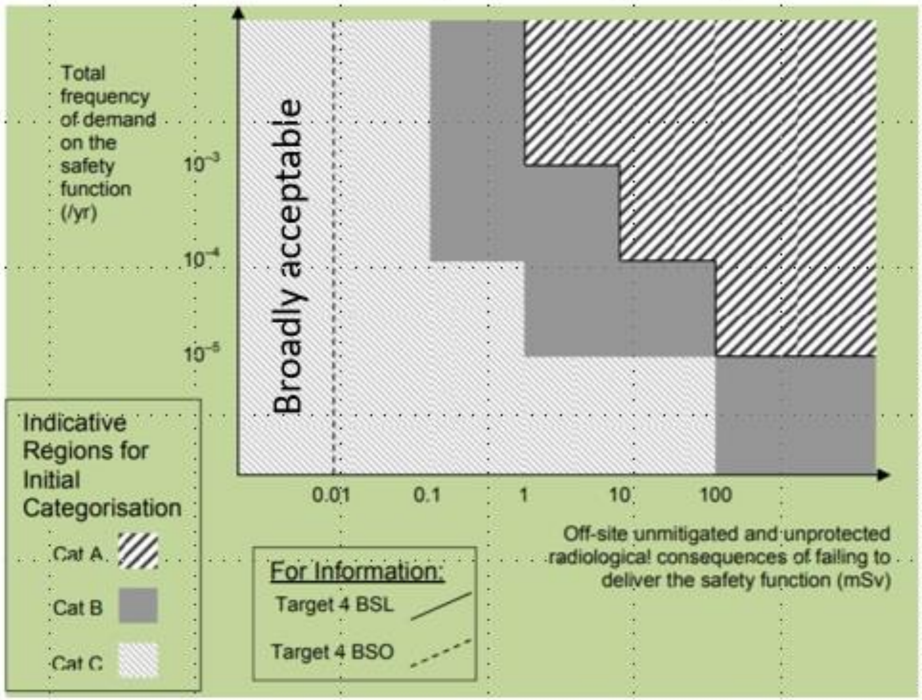

On-site (Worker)

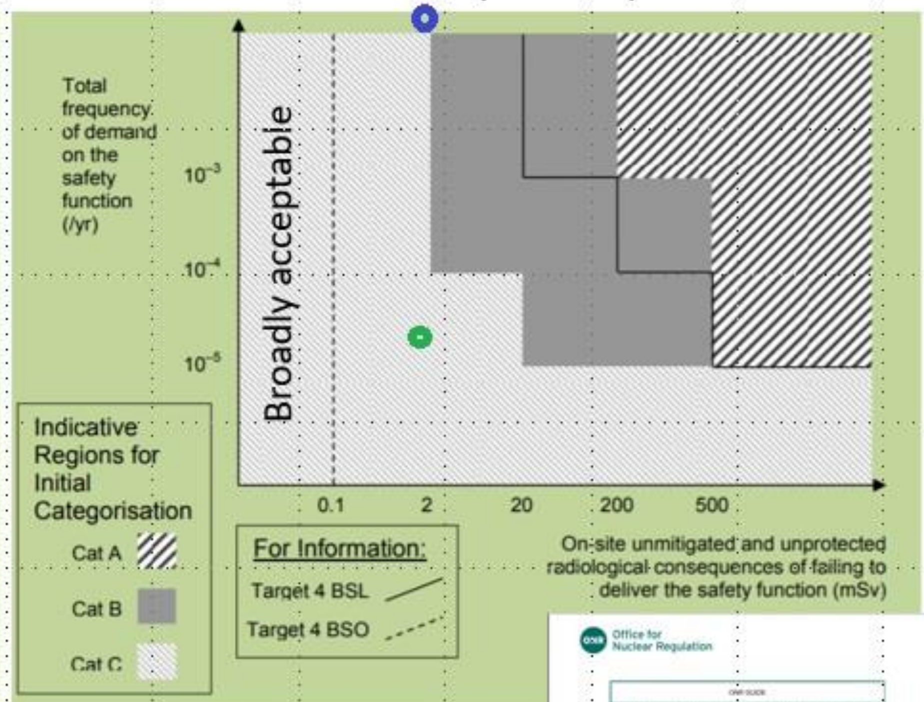

0.01mSv 100g bag of Brazil nuts

e.g. For a 10mSv on-site dose, residual risk should be <1E-4 /y and ALARP

TAG 94 Rev 2

Realisation of Safety Functions

ROBOTICS AND AI IN NUCLEAR

The University of Manchester
Autonomy and Verification Group

A Safety Function (SF) can be realised as either:

- a function which is diverse, independent and segregated from the control system, inc. sensors, control and actuators (guards)
- the control function itself within the control system
- a combination of guard and control system

The guard and/or control system must:

- lend itself to design, implementation, verification & validation to the degree required by the hazard analysis and the safety requirements (functional and non-functional) imposed on it
- Meet all deterministic requirements. e.g. the severity of the hazard may be such that a diverse SIF is required, therefore negating the use of the high integrity control system
- meet the probabilistic claim required by the hazard analysis and the safety requirements (functional and non-functional) imposed on it

Realisation of Safety Functions

ROBOTICS AND AI IN NUCLEAR

The University of Manchester

Autonomy and Verification Group

Safety Functions (SF) are realised:

- by Structures, Systems and Components (SSC) (also known as Safety Instrumented Functions (SIF))
- using appropriate standards and Recognised Good Practice (RGP)
  - e.g. IEC 61508
- Demonstrating Production Excellence (PE)
  - showing good control of the robot's development and verification lifecycle
- Independent checking of the final validated software (in its target hardware deployment) and of the testing programme (ICBM).

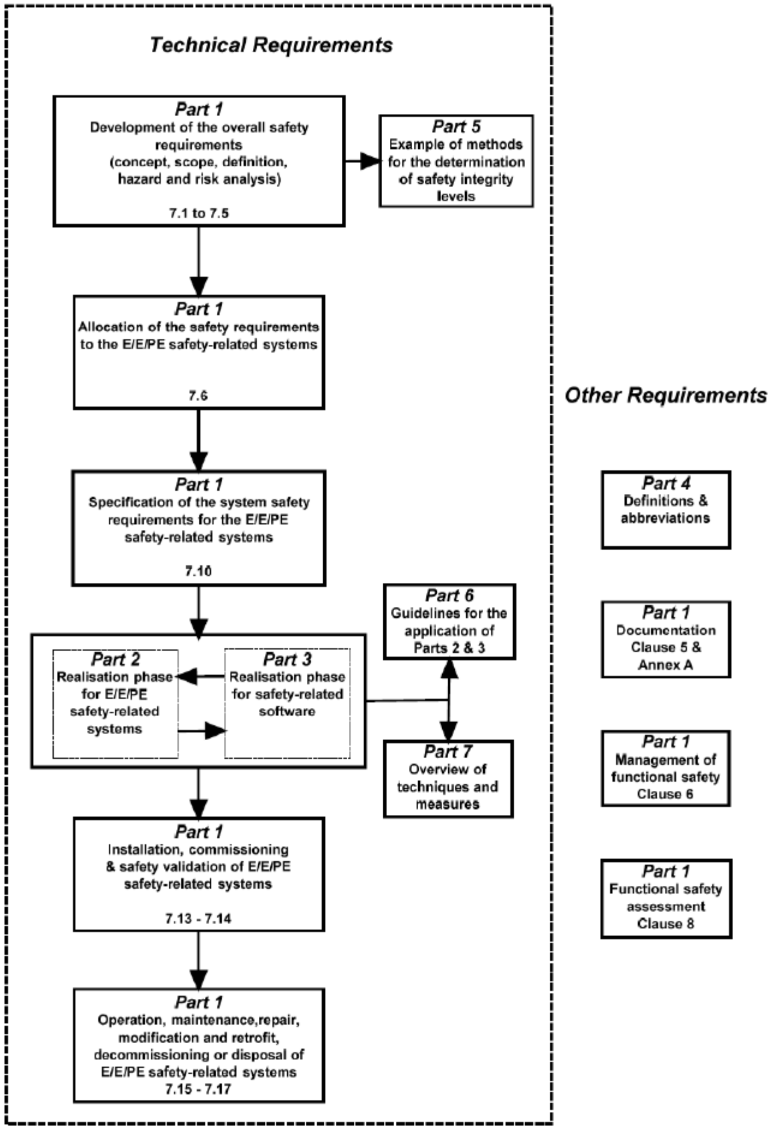

COTS Robots

Possibly identify the generic failures of a COTS or 'research development' robot (e.g. freezing, uncommanded movement),

- analyse how these relate to the identified hazards
- bound the robot accordingly

Difficulties adopting a COTS or 'research development' solution

- Very difficult to show PE and ICBM
- Proven-in-use in general never provides enough confidence that the equipment deployed in the application is tolerable and ALARP

Better then to use a 'simple' guard around the whole or part of the control system than try to substantiate COTS or 'research development' solution

Hazard Identification for the Strawman: Hazard 1

RAIN
ROBOTICS AND AI IN NUCLEAR

MANCHESTER 1824
The University of Manchester

Autonomy and Verification Group

Collision

– Consequence: damage to the pond lining, resulting in a leak of liquor from up to 5 cm below the surface of the liquor

- Occurrence Barrier
  - The likelihood of this consequence has been reduced to tolerable and ALARP by the presence of a safety instrumented function SIF.
- Protection Barrier
  - As the occurrence has been reduced to tolerable and ALARP no protection barrier is necessary, however, as recognised good practice (RGP) and for defence in depth the following provides a protection barrier
    - The pond is bunded and can easily contain the maximum volume of liquor that could leak
    - Radiologically and waterproof PPE for all workers within 10m of the edge of the pond.
- Mitigation Barrier
  - As the occurrence has been reduced to tolerable and ALARP no mitigation barrier is necessary

Hazard Identification for the Strawman: Hazard 1

The University of Manchester

Autonomy and Verification Group

However!

It may be possible to argue:

by analysis that the maximum collision energy ($\frac{1}{2}mv^2$) could not possibly damage the structural integrity of the pond.

that damage to the contents of the pond does not create any safety concern

Avoidance of Collision SIF Method 1 Diverse Guard

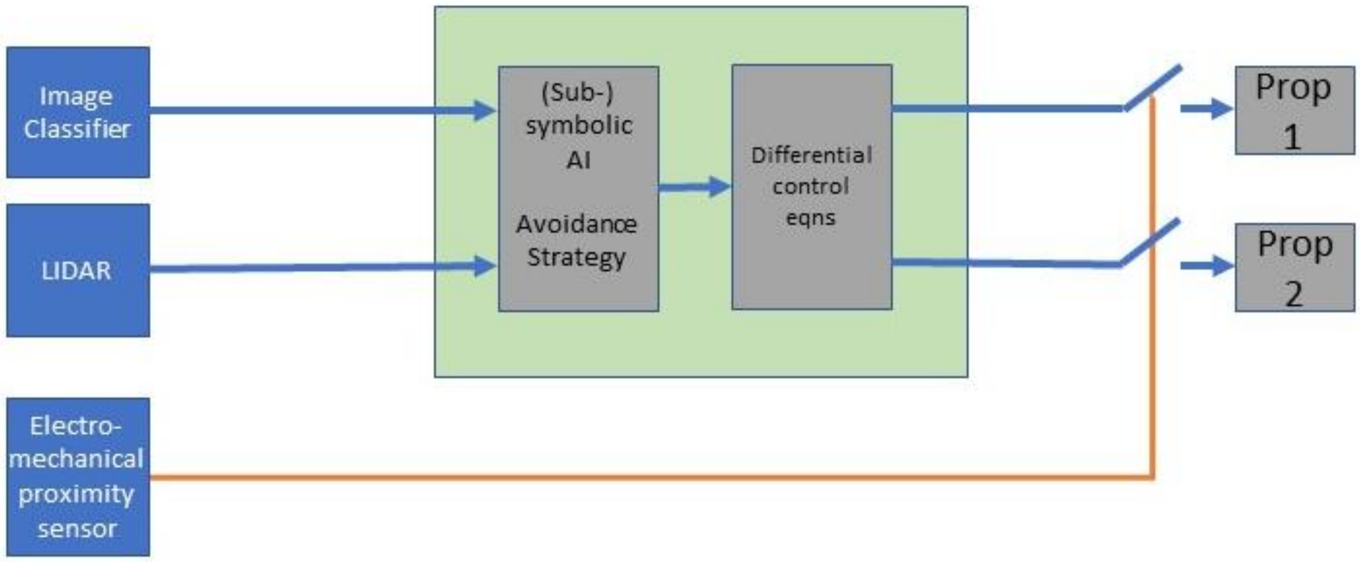

ROBOTICS AND AI IN NUCLEAR

Avoidance of Collision SIF Method 2 Rules Based Reasoning

The University of Manchester

Autonomy and Verification Group

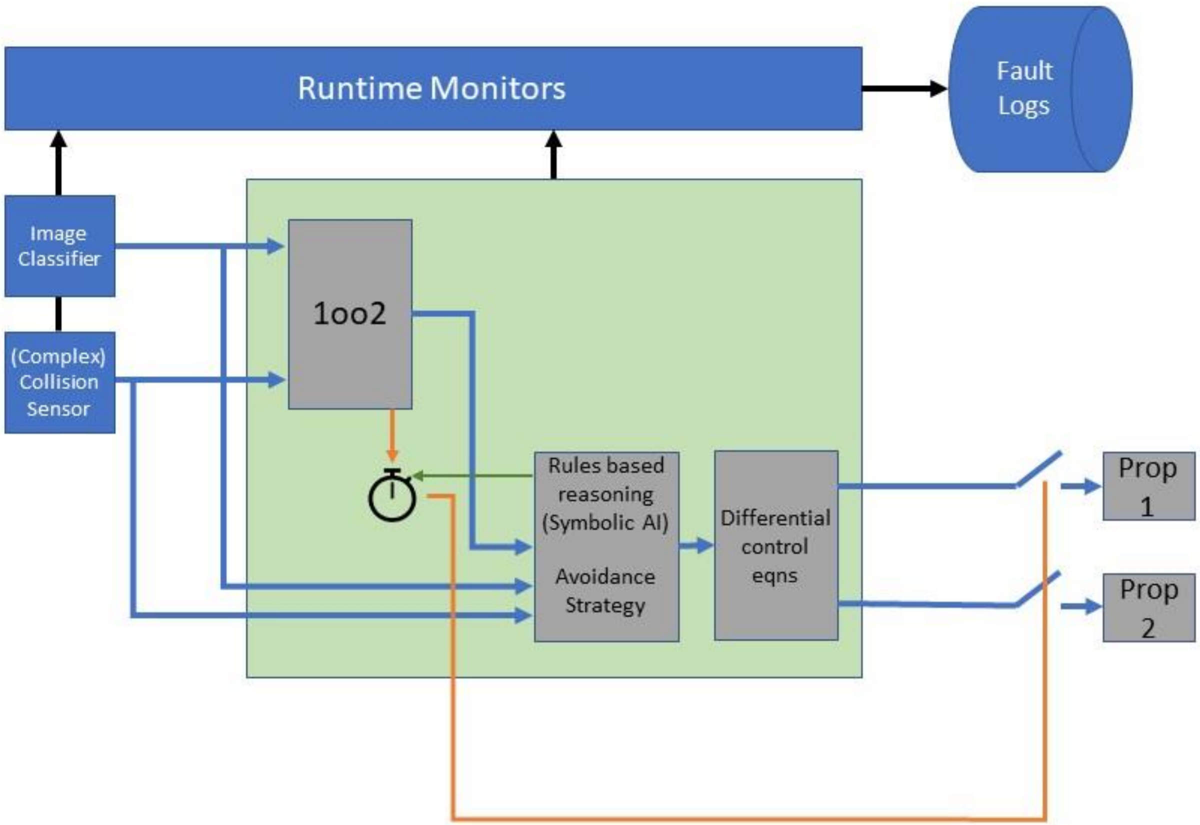

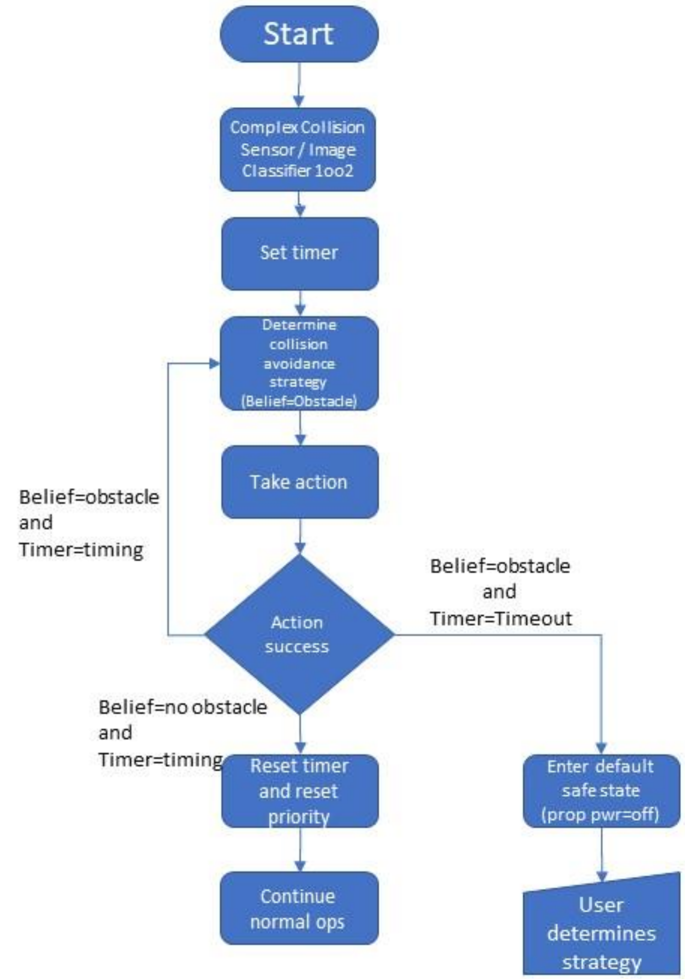

ROBOTICS AND AI IN NUCLEAR

Safety Case Formats
CAE, GSN and text

MANCHESTER 1824
The University of Manchester
Autonomy and Verification Group

How do we document a safety case?

CAE

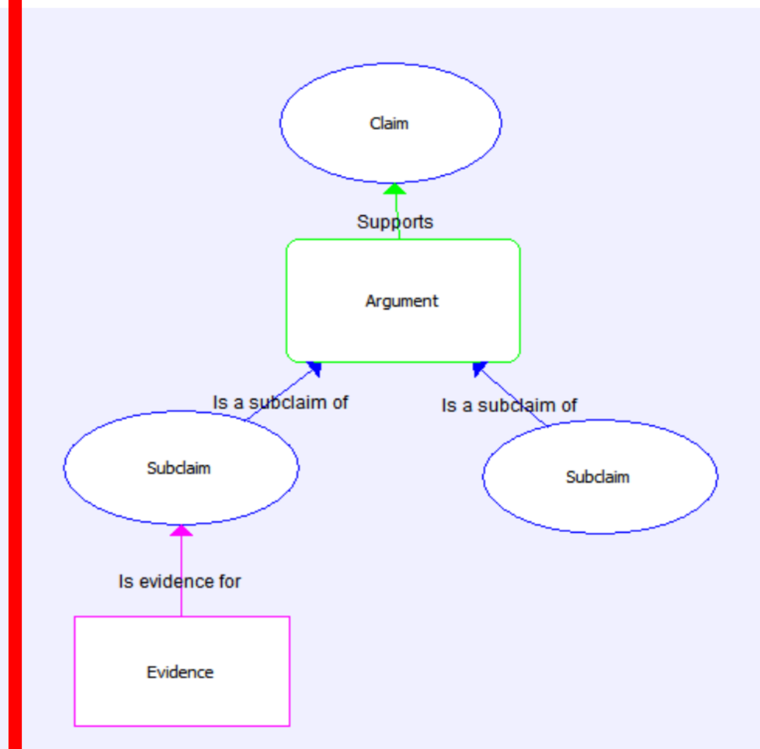

Tense: Past

ASCE (https://www.adelard.com/asce/choosing-asce/index/) | CAE

GSN

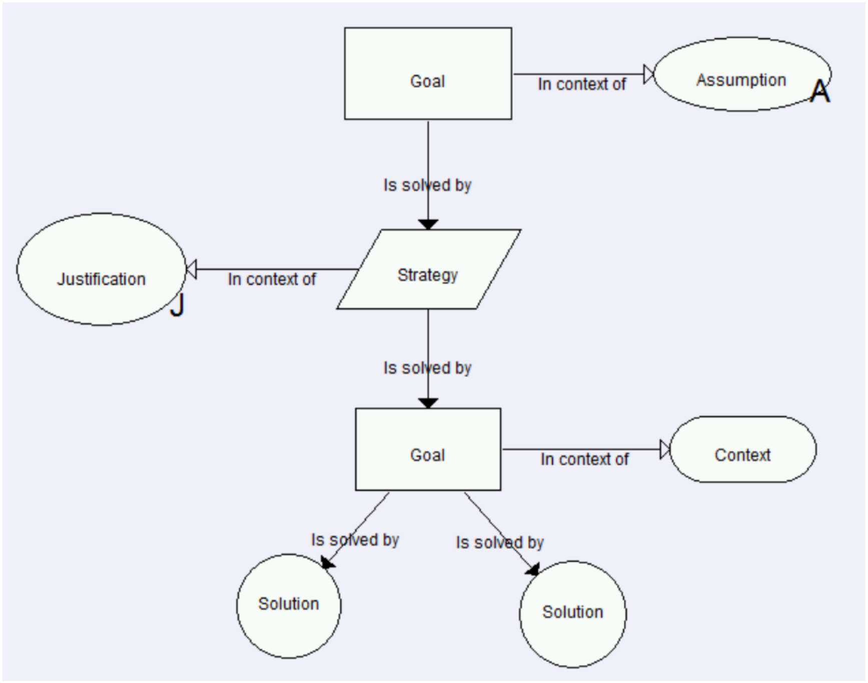

Tense: Future

ASCE (https://www.adelard.com/asce/choosing-asce/index/) | GSN

ROBOTICS AND AI IN NUCLEAR

Take Home Message

The University of Manchester
Autonomy and Verification Group

Safety Case

Alternative link

Take Home Message

Don't panic

But don't leave it to the end

Make the Safety Case a fundamental part of the robotics project

OR

Embed elements of the Safety Case into the project

- Define the task
- Identify the hazards
- Avoid making decisions which could make it difficult to retrospectively correct
- Be prepared to have industry take your autonomous robotics and develop it through an appropriate safety lifecycle

ROBOTICS AND AI IN NUCLEAR

Acknowledgements

The University of Manchester
Autonomy and Verification Group

Particular thanks to:

- Richard Cooper (SL)
- Andrew White (ONR)

ROBOTICS AND AI IN NUCLEAR

The University of Manchester

Autonomy and Verification Group

Thank you for listening