SACE

Develop compelling safety cases for autonomous systems

What is SACE?

SACE is the first methodology for creating a safety case for an autonomous system (AS) in a complex environment. It's a downloadable framework being used in industries from energy to transport enabling safety engineers and assessors to take a holistic view of the AS and its operating environment.

Who is SACE for?

Our SACE guidance is aimed at three primary audiences:

Safety Engineers

SACE enables safety engineers to understand what must be done to provide the required assurance of the safety of an AS operating in a complex environment and create a safety case for that system.

System Developers

SACE enables system developers to understand the safety assurance considerations and activities when developing an AS and generate the artefacts required to support the safety case.

System Assessors and Regulators

SACE enables system assessors and regulators to understand what should be the focus of review, and the criteria by which the sufficiency of the assurance activities and artefacts should be judged.

Why should you use SACE?

SACE is the only freely-available guidance which shows how a detailed and compelling safety case can be created for an autonomous system, taking account of the system and its interaction with the environment.

Our guidance:

- Builds on established and trusted system safety approaches.
- Enables engineers and assessors to take a holistic view of the safety of an autonomous system within its environment.
- Is written in a style and format that is easy to pick up and use by a range of stakeholders.
- Works in conjunction with our AMLAS framework to create a comprehensive and compelling safety case for the autonomous system as a whole.

SACE creates a compelling safety case which positively impacts market readiness and regulatory compliance.

The expertise behind SACE

SACE has been developed by the Centre for Assuring Autonomy, a £10m partnership between Lloyd's Register Foundation and the University of York. Our team of globally recognised experts has been advancing the safety of complex systems through pioneering research for over 30 years.

We worked with an international community of developers, regulators, and researchers to develop and release SACE as freely accessible guidance, with proven real-world impact and benefit. Our SACE guidance is based on:

1. Sound research – peer-reviewed and published in leading academic venues.
2. Empirical evaluation – evaluated in real-world, credible contexts.
3. Accessibility – practical guidance disseminated online and through CPD training.

Through SACE, we, at the Centre for Assuring Autonomy, are leading the way in ensuring society can safely benefit from artificial intelligence and autonomous systems.

How does SACE work?

The assurance activities of SACE are complementary to, and integrate with, existing system safety processes.

SACE consists of eight stages, which are:

Stage 1: Operating context assurance

The aim of this stage is to define and validate the operating context for an AS. This defines the scope under which we are able to demonstrate that the AS is safe to operate. Demonstrating assurance that the operating context is defined correctly at an appropriate level of detail is particularly challenging for an AS that operates in complex open environments.

Stage 2: AS hazardous scenarios identification

This stage is concerned with identifying and defining potentially hazardous scenarios for the AS. Hazardous scenarios are those that the AS may encounter during its operation that could, under certain conditions, lead to an unsafe outcome. For AS we focus in particular on the interactions between the AS and elements of the operating environment, and on the decisions that are made by the AS as part of its autonomous operation.

Stage 3: Safe operating concept assurance

At this stage the safe operating concept (SOC) for the AS is defined and validated. The SOC specifies the requirements for sufficiently safe operation of the AS within the operating context previously defined, taking account of the identified hazardous scenarios. As part of the SOC it may also be necessary to define additional constraints on the operation of the AS to ensure safety is maintained under identified system or environment conditions.

Stage 4: AS safety requirements assurance

This stage considers how the safety requirements are defined and validated at each level of decomposition in the AS development process such that traceability can be demonstrated and maintained.

Stage 5: AS design assurance

At this stage, the sufficiency of the design of the AS is justified with respect to the defined safety requirements. This step is iterative as it considers the assurance of the design of the AS across multiple levels of design decomposition. This stage involves creating design proposals to meet the safety requirements, carrying out analysis of those proposals, updating the safety requirements or changing the design in response, doing additional analysis, and so on for all levels of the AS design.

Stage 6: Hazardous failures management

This stage addresses the identification and mitigation of hazardous failures of the AS. This considers the design of the AS at each level to determine how hazardous failures could arise as a result of that design. This is a crucial activity since, even where the design has completely implemented all of the identified safety requirements, it still may be the case that the AS is capable of doing something else, under certain conditions, that could be hazardous.

Stage 7: Out of context operation assurance

Although undesired, an AS may spend some time operating outside the defined operating context whilst still operating autonomously. This could be unsafe, since autonomous operation is only assured for safety within the defined operating context. This stage seeks to provide assurance that the AS will remain safe even when operating out of context (for example by safely handing over control to a human).

Stage 8: AS verification assurance

This stage generates evidence to demonstrate that the safety requirements specified for the AS are satisfied and justifies the sufficiency of the verification activities.

Centre *for* Assuring Autonomy

Work with us

There are several ways the Centre for Assuring Autonomy can support your organisation in developing safety assurance cases.

Training

Our bespoke training and education sessions enable industry, regulators, and policy makers to develop the expertise necessary to ensure that autonomous systems are brought safely to market and into operation.

Research

We can help you solve emerging and critical research questions around safe AI through our multidisciplinary team approach.

Partnerships and collaborations

Our team of experts works with organisations to help develop safety cases, address particular challenges, or develop compliance procedures.

Consultancy

Our bespoke consultancy service enables you to build a package of support which meets organisational needs and access the right experts for your AI and AS challenges.

Our safety assurance guidance and frameworks

SACE is part of a suite of freely-available frameworks and guidance which organisations can use to develop safety cases for autonomous systems and AI, including Frontier AI.

Read our other factsheets in this series

- BIG Argument
- AMLAS
- PRAISE

University of York

LR Foundation

+44 (0)1904 325345
assuring-autonomy@york.ac.uk

cfaa.bsky.social
assuring-autonomy