


How to investigate when a robot causes an accident, and why it's important that we do

March 25 2022, by Keri Grieman



Credit: AI-generated image ([disclaimer](#))

Robots are featuring more and more in our daily lives. They can be incredibly useful (bionic limbs, robotic lawnmowers, or robots which [deliver meals](#) to people in quarantine), or merely entertaining (robotic dogs, dancing toys, and acrobatic drones). Imagination is perhaps the only limit to what robots will be able to do in the future.

What happens, though, when robots don't do what we want them to—or do it in a way that causes harm? For example, what happens if a bionic arm is involved [in a driving accident](#)?

Robot accidents are becoming a concern for two reasons. First, the increase in the number of robots will naturally see a rise in the number of accidents they're involved in. Second, we're getting better at building more complex robots. When a robot is more complex, it's more difficult to understand why something went wrong.

Most robots run on various forms of [artificial intelligence](#) (AI). AIs are capable of making human-like decisions (though they may make objectively good or bad ones). These decisions can be any number of things, from identifying an object to interpreting speech.

AIs are trained to make these decisions for the robot based on information from vast datasets. The AIs are then tested for accuracy (how well they do what we want them to) before they're set the task.

AIs can be designed in different ways. As an example, consider the robot vacuum. It could be designed so that whenever it bumps off a surface it redirects in a random direction. Conversely, it could be designed to map out its surroundings to find obstacles, cover all surface areas, and return to its charging base. While the first vacuum is taking in input from its sensors, the second is tracking that input into an internal mapping system. In both cases, the AI is taking in information and making a decision around it.

The more complex things a robot is capable of, the more types of information it has to interpret. It also may be assessing multiple sources of one type of data, such as, in the case of aural data, a live voice, a radio, and the wind.

As robots become more complex and are able to act on a variety of information, it becomes even more important to determine which information the robot acted on, particularly when harm is caused.

Accidents happen

As with any product, things can and do go wrong with robots. Sometimes this is an internal issue, such as the robot not recognizing a voice command. Sometimes it's external—the robot's sensor was damaged. And sometimes it can be both, such as the robot not being designed to work on carpets and "tripping." [Robot accident investigations](#) must look at all potential causes.

While it may be inconvenient if the robot is damaged when something goes wrong, we are far more concerned when the robot causes harm to, or fails to mitigate harm to, a person. For example, if a bionic arm fails to grasp a hot beverage, knocking it onto the owner; or if a care robot fails to register a distress call when the frail user has fallen.

Why is robot accident investigation different to that of human accidents? Notably, robots don't have motives. We want to know why a robot made the decision it did based on the particular set of inputs that it had.

In the example of the bionic arm, was it a miscommunication between the user and the hand? Did the robot confuse multiple signals? Lock unexpectedly? In the example of the person falling over, could the robot not "hear" the call for help over a loud fan? Or did it have trouble interpreting the user's speech?

The black box

Robot accident investigation has a key benefit over human accident

investigation: there's potential for a built-in witness. Commercial airplanes have a similar witness: [the black box](#), built to withstand [plane crashes](#) and provide information as to why the crash happened. This information is incredibly valuable not only in understanding incidents, but in preventing them from happening again.

As part of [RoboTIPS](#), a project which focuses on responsible innovation for social robots (robots that interact with people), we have created what we call the [ethical black box](#): an internal record of the robot's inputs and corresponding actions. The ethical [black box](#) is designed for each type of robot it inhabits and is built to record all information that the [robot](#) acts on. This can be voice, visual, or even [brainwave activity](#).

We are testing the ethical black box on a variety of robots in both laboratory and simulated accident conditions. The aim is that the ethical black box will become standard in robots of all makes and applications.

While data recorded by the ethical black box still needs to be interpreted in the case of an accident, having this data in the first instance is crucial in allowing us to investigate.

The investigation process offers the chance to ensure that the same errors don't happen twice. The ethical black box is a way not only to build better robots, but to innovate responsibly in an exciting and dynamic field.

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

Citation: How to investigate when a robot causes an accident, and why it's important that we do

(2022, March 25) retrieved 8 April 2024 from <https://techxplore.com/news/2022-03-robot-accident-important.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.