Raspberry Pi Foundation

Lesson 2: How computers learn from data

Enjoyed teaching these lessons? Found a mistake? Share your feedback at rpf.io/experienceai-feedback.

Introduction

In this lesson, students will build on the new view of artificial intelligence from Lesson 1, with a particular focus on the use of data in AI systems. The activities will help students think critically about which parts of a system use AI principles, and the role machine learning plays in creating the models introduced in Lesson 1.

First, students will consider the functionality of a 'smart' speaker, with the aim of identifying which uses involve data-driven techniques and which do not. Next, they will be introduced to the definition and description of 'machine learning' and its role in the landscape of AI. Students will hear from experts about the different types of machine learning and the problems its use can help solve.

Finally, students will learn about a specific example of machine learning — classification. This is where algorithms are used to classify (group) data into categories (called 'classes'), and example data that has already been labelled must be used to train the algorithms.

Learning objectives

- Define machine learning's relationship to artificial intelligence
- Name the three common approaches to machine learning
- Describe how classification can be solved using supervised learning

Key vocabulary

Machine learning, training data, supervised learning, unsupervised learning, reinforcement learning, classification, class, label

Preparation

You will need:

- Slides
- Activity 1 worksheet and answer sheet
- Activity 2 teacher question sheet

- Activity 3 worksheet and answer sheet

Subject knowledge:

The language in this lesson has been chosen very carefully to avoid anthropomorphising (giving human qualities to) AI applications, focusing more on AI as a field of research and an idea of making models that complete complex tasks. This can be very difficult to do, but is worth keeping in mind during discussions in this lesson.

Activity 1: This activity will introduce the term 'machine learning', so it is recommended that you familiarise yourself with the definition below:

> Machine learning (ML) is a way of building artificial intelligence systems using statistics, instead of by writing out the rules in a program. Machine learning systems are called 'data-driven', because they use a lot of examples — data — to work. Examples of machine learning applications include processing speech, recognising faces, recommending songs, and diagnosing illnesses. Unlike traditional programming where step-by-step instructions are written to solve a problem, machine learning systems are designed to make predictions based on lots of data. For this reason, machine learning systems can produce useful results even for complicated problems that are not well defined. However, for the same reason, machine learning systems can also be less consistent and harder to debug than rule-based programming.

It would also be beneficial for you to watch the video on slide 11 before the lesson, so you can anticipate any questions your class may have.

Activity 2: It is recommended that you watch the video on slide 20 before the lesson. The three different types of machine learning are described in this video. Some technical terms are introduced, but the aim of this activity is to introduce the students to these types of machine learning, and not for them to completely understand the intricacies of each one. Most of the rest of the unit will focus on **supervised learning**, so it is particularly important to familiarise yourself with that part of the video.

Activity 3: This activity introduces a specific use of machine learning — classification. This is where a machine learning model is used to classify (categorise) data into groups (called 'classes'). The model will produce a 'label' (a prediction of which class the data belongs to). It is recommended that you review these key terms before the lesson.

There are two examples of classification in this activity — sentiment analysis and computer vision. Familiarising yourself with these will help you answer any questions the students may have. Example questions and answers have been provided in the lesson plan below.

Assessment opportunities

Activity 1 has differentiation built in. Students will likely find it easier to hypothesise what data the model might need to 'learn' from if they choose the "Identifying animals in a video" option.

Activity 2 presents an opportunity to assess students' comprehension of the key information provided in the video with an informal, low-stakes quiz. Take this opportunity to expand on any of the characteristics that many students in the class get wrong.

Outline plan

Please note that the slide deck labels the activities in the top right-hand corner to help you navigate the lesson.

**Timings are rough guides*

Starter activity (Slides 2–3) 3 mins	**Is a 'smart' speaker an AI application? Why?** As the students enter the classroom, display slide 2 with the question "Is a 'smart' speaker an AI application? Why?" As the students are arriving, ask them to discuss their answers and their reasoning with the person sitting next to them. Engage students who have been in the classroom the longest and ask probing questions on any parts of their reasoning that are not clear. The aim of this activity is not to arrive at a definitive answer, but to encourage students to recall the ideas presented in Lesson 1 about rule-based and data-driven applications. It is also a chance for students to think in this way about a common item they might already be aware of. A short description of some example uses of a 'smart' speaker is provided on the slide to support students who may not be familiar with them to engage in the discussion. When you are ready to begin the lesson, show slide 3 and run through the objectives with the class.
Introduction (Slides 4–9) 7 mins	**Breaking down a 'smart' speaker** Display slide 4, which shows two examples from the end of Lesson 1 that the students categorised as 'Not AI' and 'AI'. Move on to slide 5, and use these two examples to remind the students of the

	two different approaches that are the focus of this unit: 1. Rule-based 2. Data-driven Show slide 6. Ask the students, **"**What can you use a 'smart' speaker to do?" Take answers from the class. Students might use some language that humanises the 'smart' speaker, such as "listens" or "understands". You should try to challenge this language with responses such as: • "Does a microphone 'listen'? Or is that something humans do?" • "What does 'understand' mean? Can a computer 'understand' anything? Computers process inputs and produce outputs based on the programs running at the time; there is no understanding involved." Use the animation on the slide to reveal four example uses. This is not an exhaustive list, and students might have thought of some others (such as using a 'smart' speaker to interact with other internet-enabled devices like lights). Click to activate the next animation and remind the students that these are all things that **they use** a 'smart' speaker to do, not things the 'smart' speaker does itself. Devices like these are tools and do not have the ability to understand or answer questions themselves. Click again to show a key question: "Which uses would benefit from a data-driven approach?" Ask students for opinions on each of the uses on the slide, and remind them to consider the complexity of each of the uses. Show slide 7, which explains that interpreting voice commands can be really complex, as everyone speaks in a slightly different way; for example, people speak in different languages and have different accents. A data-driven approach, with a model, would be beneficial. Move on to slide 8, which explains that answering questions requires a search engine and lists some ways a model might help provide high-quality answers to the user. Display slide 9, which encourages the students to be more nuanced in their discussions of AI applications, acknowledging that not every use of these applications uses AI techniques. This should further cement the idea that AI applications are not 'an AI', but a collection of processes that can be both data-driven and rule-based.

Activity 1 (Slides 10–19) 15 mins	**What is machine learning?** Display slide 10. Ask the students, **"**If I wanted the model used in a 'smart' speaker to accurately interpret commands from as many people as possible… What would I need to train the model?" Students might remember the link to data from Lesson 1, so subtly encourage them to do so. Use the animation on the slide to show them that "data" is the right answer. Play the video on slide 11, which will introduce the students to machine learning. Next, use slide 12 to show the students the first part of the definition of 'machine learning'. Highlight that the models that the students have been learning about are **created** using machine learning techniques. Although we have been calling them 'AI models' so far, a more accurate name is '**machine learning models**'. This is the term that is used widely in the field and will be used from this point on in the lessons. On slide 13, there is a Venn diagram to show how machine learning fits into the field of AI. The first animation tells students that machine learning is a part of AI and not the whole of the field. Students may ask what type of AI does not use machine learning, and you can explain that some AI applications rely on extreme amounts of processing power to run complex rule-based systems that do not use machine learning techniques, but that this is an older view of AI. The next animation will help with this explanation. It reiterates the concept learned in Lesson 1 that some AI applications will still use rule-based systems, but the current state of AI (referring to the field of study) is focused on the use of machine learning models, so that is why it is the focus of this unit. Show slide 14, which features a later part of the definition of 'machine learning' focused on **predictions**. Machine learning is beneficial when a problem is very complex and the exact steps to solve it are not easily defined; otherwise, you would use a rule-based approach (like the noughts and crosses example from Lesson 1). Slide 15 focuses on predictions. This is a key learning for the students, as the outputs of machine learning models include some uncertainty. You can ask the students where they have seen signs of this uncertainty already, reminding

	the class of the **confidence scores** they saw in the computer vision example in Lesson 1. Move on to slide 16 and hand out the worksheet. Explain that the students need to be critical when deciding whether to use machine learning models when solving a problem. This activity will encourage them to think critically about three example uses of an application, evaluating whether a data-driven approach would be beneficial. Give them 5 minutes to consider the uses and discuss with a partner whether they think it would be best to use instructions or a model for each of them. Once the class has finished, move on to slide 17 and gather feedback. Slide 18 reveals some answers. You will notice that there is some speculation in them, and this is purposeful. It is very difficult to say for certain that any problem **definitely** requires a particular approach. Students should see that calculating the distance to a star could probably be achieved with a rule-based instructions approach and the other two example uses would probably require a model. Move on to slide 19 and ask the students to complete the second part of the activity. They should choose one of the uses requiring a model and speculate on the data they would need to create a machine learning model to complete the task. Students will likely find it easier to think about the data for identifying animals in a video, rather than for forecasting air pollution. The aim of this activity is not for the students to get the answers exactly right, but for them to start thinking about models in more concrete terms, as things that are made and not 'magic boxes' that produce outputs on their own. You can find example answers in the answer sheet to help guide the discussion of the two models and the data required.
Activity 2 (Slides 20–32) 10 mins	**The types of machine learning** Play the video on slide 20, which will make students aware of the different approaches to machine learning. **Note:** It is recommended that you watch the video before the lesson so that you know what is addressed. Students are not expected to understand the intricacies of each of the different types of machine learning, so the rest of this activity is a quiz to help

	them cement the key knowledge from the video. You can use the slide deck to show the questions and answers (slides 21–31), or if you prefer, you can skip the slides and read the questions out using the teacher question sheet. Students can indicate which type of machine learning each statement refers to in a number of ways: • Holding up coloured paper • Writing “1”, “2”, or “3” on a whiteboard • Holding up a number of fingers • Moving to one of three designated areas in the classroom The mode of delivery is entirely up to you. The important part is that you can assess the students’ comprehension of the video through this quiz. Use slide 32 to recap the key characteristics of the different types of machine learning for the students to remember.
Activity 3 (Slides 33–45) 15 mins	**Classification** Display slide 33 and explain that you are going to focus on one particular type of problem that machine learning can help solve — **classification**. Classification uses **supervised learning** to categorise data. There are some key terms to introduce here: classification, class, and label. Explain that during classification, data is grouped into categories — or classes. The image on the slide shows examples of three classes: Apple, Banana, and Orange. Move on to slide 34 and explain that a classification model is trained with example data (training data). This training data has already been labelled by a human. Ask the students, “What labels would you apply to these images?” Show slide 35 to display the labels for each example image. Display slide 36 to expand on the training process. Explain that the more example data used when training a classification model, the more accurate the model will be. Display slide 37 to show what happens once training is complete: new data can be fed into the model and it will predict a class from the options it has

	been trained with. Note the inclusion of a confidence score with the prediction, to further cement the idea of uncertainty. Ask the students, “Would you be happy to use a model that was 92% confident when identifying an orange?” Then ask, “Would you be happy driving in a car that was 92% confident when identifying a human in the road?” This suggests that different levels of uncertainty are acceptable in different situations. Display slide 38, which shows the definitions of ‘class’ and ‘label’ to help cement them for the students. To introduce the students to some real-world contexts in which classification is used, the next few slides (39–43) summarise two examples. Display slide 39 and introduce **sentiment analysis**, which is when a model is trained to label comments and reviews as generally positive or negative by examining the text for keywords and, of course, sentiments. As you mention the text, use the animation on the slide to highlight which part of the review would be used to train the model. Move on to slide 40 and explain that the classes in this case would be Positive and Negative. Explain that sentiment analysis models are beneficial because reading a large volume of reviews is very time-consuming for a human to do, and these models can generalise from large sets of data whether the product or service is well received or needs improvement. Show slide 41 and introduce **computer vision**, which students experienced in the examples of AI applications in Lesson 1. An important point to note is that the model can recognise many objects in a single data point (an image, in this example). The model will be trained with images, both of individual objects and of multiple objects, like the image shown on the slide. Slide 42 shows all the different classes that the algorithm has found in the featured image, to give students an idea of the breadth of labels that can be applied. Move on to slide 43, which shows a summary of some of the classes, and explain that computer vision models are useful because multiple objects can be recognised in a photo or video, which is useful for technologies such as driverless cars. You can ask students what classes of objects a car might need to recognise,

	for example: • Pedestrians • Traffic lights • Curbs • Other vehicles • Signs Use slide 44 to explain that now the students have an idea of what classification is, they are going to consider three example classification problems, what type of data they would use, and what some of the classes for each of them could be. Highlight that almost any type of data can be used for a classifier, not just images or text. Show slide 45 and hand out the worksheet. Give the students 5 minutes to complete the task and then run through their answers with the whole class. You can find some example answers in the answer sheet and listed below. Assigning a genre to music: • Data types: Audio, such as MP3, WAV, etc. • Classes: Any commonly used genre title, such as Pop, Rock, Hip-hop, Drum and bass, etc. Identifying spam emails: • Data types: Text, numbers, and URLs • Classes: Spam, Not spam, and For review Language detection: • Data types: Text, audio, and video • Classes: Any language, such as English, French, Arabic, British Sign Language, etc.
Plenary (Slides 46–51) 5 mins	**Classification — your turn** Throughout this lesson, the students have been exploring the role of both machine learning algorithms and data in AI applications. This final activity will allow the students to consider the learning in the real-world context of a project that is making a huge difference in animal conservation. Show slide 46 and explain the context to the students: DeepMind are using

	classification in the Serengeti (a national park in East Africa) to help protect endangered species. Move on to slide 47. Tell the students that they are going to "help" by pre-labelling some data to use when training a classification model for this project. The next four slides feature images captured by remote cameras in the Serengeti. These are real images that were actually used to train the classification model used by DeepMind and the conservation team. For each image, the students need to pick a label from the available classes to apply to the image, getting them involved in this amazing use of machine learning models for good. The answers are as follows: • Slide 48: elephant • Slide 49: gazelle • Slide 50: lion • Slide 51: zebra