



Experience AI

Educator guide: Foundations of AI



This resource is licensed by the [Raspberry Pi Foundation](https://www.raspberrypi.org/) under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International Public License (CC BY-NC-ND 4.0). For more information on this licence, see creativecommons.org/licenses/by-nc-nd/4.0/.

Contents

Introduction	2
Curriculum design	3
The approach	3
Coherence and flexibility.....	3
Knowledge organisation.....	3
Core principles	4
Inclusive and ambitious.....	4
Research-informed.....	4
Time-saving for educators.....	5
UNESCO AI competency framework.....	5
Structure of the unit	6
Lesson outlines	6
Lesson 1: What is AI?.....	6
Lesson 2: How computers learn from data.....	7
Lesson 3: Bias in, bias out.....	7
Lesson 4: Decision trees.....	8
Lesson 5: How to solve problems with machine learning models.....	8
Lesson 6: Model cards and careers.....	9
Guidance for adapting lesson content.....	10
Principles for adaptation.....	10
Examples of lesson adaptations.....	11
Progression	15
Progression within the unit.....	15
Pedagogy	16
Assessment	19
Formative assessment.....	19
Summative assessment.....	19
Adapting to your setting.....	19
Raspberry Pi Foundation	20

Introduction

Experience AI is an educational programme that offers cutting-edge resources on artificial intelligence (AI) and machine learning (ML) to students. Amid young people's increasingly independent use of these technologies, Experience AI fosters skills such as problem solving, innovation, and critical thinking, empowering students to use AI and ML in informed, responsible ways.

Developed by the Raspberry Pi Foundation in collaboration with Google DeepMind, the Experience AI 'Foundations of AI' unit is a comprehensive set of six lessons that introduces 11- to 14-year-old students to a range of AI and ML concepts. Students are given an introduction to what AI is and then explore, among many other things, how ML models are built, the impact of bias in data, decision trees, and the AI project lifecycle. They will also see how AI and ML technologies may affect them – both now and in the future – as well as learning about the ever-increasing variety of AI-related careers.

Designed with non-specialist educators in mind, the Foundations of AI unit has everything educators need to deliver the lessons. The materials provided include lesson plans, slide decks, videos, and guides for hands-on activities that can be used in the classroom. In addition, students will have the opportunity to create their own ML models and gain practical experience with ML learning tools. Importantly, the unit does not require educators or students to have experience in programming or computer science.

Curriculum design

The approach

Coherence and flexibility

The Foundations of AI Lessons are designed to be taught in order, with each of the six lessons building logically on the previous one. This ensures that students develop a strong foundational knowledge in key AI concepts and competencies before progressing to more advanced topics such as bias, decision trees, and ML model creation.

Each lesson comes with a detailed plan that provides suggested timings for the activities, allowing educators to adjust their pacing based on their classroom needs while maintaining the coherence of the unit's learning goals. For example, educators can choose to spend longer on topics that spark more interest amongst their students, or shorten activities to fit with time constraints. For more guidance on how to adapt lessons, please refer to the 'Guidance for adapting lesson content' section later in this guide.

Knowledge organisation

Experience AI has been designed around the ['SEAME'](https://rpf.io/blog-seame-framework) (rpf.io/blog-seame-framework) model, which structures the students' learning journey based on social and ethical aspects (SE), applications (A), models (M), and engines (E). These different layers of AI systems are explored in age-appropriate ways, guiding students through the ethical and social implications of AI as well as more technical aspects, such as model development and practical applications.

With the incorporation of the SEAME model in Experience AI, we have ensured that students can engage with AI technologies in a well-rounded manner, equipping them not only with the technical knowledge related to AI technologies, but also with an awareness of ethical issues like data bias and misinformation.

Core principles

Inclusive and ambitious

We recognise the critical importance of preparing the next generation of students to become informed citizens who understand AI technologies and their complexities. With such understanding, students can contribute their unique perspectives to the ongoing development of these technologies, helping to bring about a digital future that is more inclusive and ethical.

With this in mind, we've developed Experience AI to be as accessible as possible. We aim to empower all students, regardless of the social and cultural contexts in which they live – Experience AI is designed for an international audience of diverse individuals. The resources give students choices in projects and allow them to pursue their own particular interests, ensuring that no matter who they are, they'll find the lessons engaging and relatable.

We've also made things accessible from the educator side. We've ensured that the lessons can be delivered by non-specialist educators who are not familiar with AI topics. We provide comprehensive guidance in all the provided materials so that any educator will feel confident and comfortable even when they don't have a background in computer science, AI, ML, or another technical subject.

Research-informed

All Experience AI lessons were developed using a research-informed approach built on insights from a series of [research seminars](https://rpf.io/ai-blogs) (rpf.io/ai-blogs) on AI and data science education hosted by the Raspberry Pi Foundation in 2021 and 2022, as well as ongoing research at the [Raspberry Pi Computing Education Research Centre](https://rpf.io/cerc) (rpf.io/cerc) at the University of Cambridge. Working collaboratively alongside industry experts at Google DeepMind, we've also ensured that the lessons are grounded in cutting-edge research and pedagogy while addressing gaps in existing AI education resources.

Time-saving for educators

The Foundations of AI unit is designed to save educators time by providing detailed lesson plans, slide decks, worksheets, and more, all of which you can easily adapt to suit the needs of your students. These resources are accessible to non-specialist educators and are versatile enough to be used in various settings, including assemblies and youth clubs.

UNESCO AI competency framework

The [AI competency framework for students by UNESCO](https://rpf.io/unesco-ai) (rpf.io/unesco-ai) calls for a human-centred and technical approach to AI in education to ensure that future generations are not only competent users and creators of AI technologies, but also critical consumers and ethical decision-makers. These competencies are closely mirrored in the Foundations of AI unit, reflecting a shared commitment to equipping younger generations with both the knowledge and ethical understanding that they will need to effectively navigate and contribute to the ever-evolving field of AI.

The heat map below represents how the learning objectives for the Foundations of AI unit map onto the curricula aspects in UNESCO's AI competency framework for students.

Competency aspects	Progression levels		
	Understand	Apply	Create
Human-centred mindset	Human agency	Human accountability	Citizenship in the era of AI
Ethics of AI	Embodied ethics	Safe and responsible use	Ethics by design
AI techniques and applications	AI foundations	Application skills	Creating AI tools
AI system design	Problem scoping	Architecture design	Iteration and feedback loops

Structure of the unit

There are six lessons in the Foundations of AI unit. They are designed to be taught in order – each lesson builds conceptually on the previous one. The earlier lessons aim to foster a strong foundational understanding of AI systems, with students exploring different ML models and engines, and learning how ML models process data. Social and ethical issues will be explored throughout the unit, while towards the unit’s end students will have the opportunity to solve real-world problems using their AI skills.

Each lesson comes with videos that introduce new AI or ML concepts, so that students can hear from industry experts and see real-world examples of how these technologies are used. This will not only help students connect their learning to potential career paths in the AI field, but also highlight careers in other industries that will be influenced by AI and ML.

Lesson outlines

Lesson 1: What is AI?

The first lesson is designed to be an introduction to AI and the potential benefits and drawbacks of AI systems. Through interactive activities, such as playing noughts and crosses (tic-tac-toe) against an algorithm, students will reflect on the concept of ‘intelligence’, compare rule-based and data-driven approaches, and evaluate two AI applications along with their societal impacts.

Objectives

Students will be able to:

- Describe the difference between data-driven and rule-based approaches to application development
- Name examples of AI applications
- Outline some benefits and issues of using AI applications

Key vocabulary

Artificial intelligence (AI), algorithm, data, rule-based, data-driven, model, generative AI, computer vision

Lesson 2: How computers learn from data

Lesson 2 builds on students' understanding of AI by focusing on the role of data and ML in AI systems. Differentiating between rule-based and data-driven applications, students will investigate how ML models are created, with a focus on supervised learning and classification.

Objectives

Students will be able to:

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

Key vocabulary

Machine learning (ML), training data, supervised learning, unsupervised learning, reinforcement learning, classification, class, label

Lesson 3: Bias in, bias out

In lesson 3, students will apply what they learned about classification and supervised learning from lesson 2, and create their own ML model using the tool [Machine Learning for Kids](https://rpf.io/ml4k) (rpf.io/ml4k). Through this activity, students will explore how bias in data can influence ML outcomes. The importance of using training data from diverse sources will also be emphasised.

Objectives

Students will be able to:

- Describe the impact of data on the accuracy of an ML model
- Explain the need for both training and test data
- Explain how bias can influence the predictions generated by an ML model

Key vocabulary

Artificial intelligence (AI), machine learning (ML), supervised learning, classification, training data, test data, accuracy, bias, data bias, societal bias

Lesson 4: Decision trees

Building on students' understanding from previous lessons, this lesson will introduce students to decision trees: a type of ML model that uses classification. Students will learn about the structure of decision trees – in particular, how decision trees process data and predict labels. They will use training data related to astronomy to develop their own decision tree using [Machine Learning for Kids](https://rpf.io/ml4k) (rpf.io/ml4k), simulating how scientists might classify new stars discovered by the James Webb telescope.

Objectives

Students will be able to:

- Describe how decision trees can be used to build an ML classification model
- Describe how training data changes an ML model
- Explain why ML is used to create decision trees

Key vocabulary

Decision tree, feature, node, root node, decision node, leaf node, classification, explainability

Lesson 5: How to solve problems with machine learning models

Lesson 5 introduces students to the AI project lifecycle. Students will follow the stages of the AI project lifecycle and then pick a real-world problem to solve. They will create their own ML model to solve the problem, training the model and testing it to determine its accuracy.

Objectives

Students will be able to:

- Describe the stages of the AI project lifecycle
- Use an ML tool to import data and train an ML model
- Test and examine the accuracy of an ML model

Key vocabulary

AI project lifecycle, data cleaning, machine learning (ML) model, class, label, training, testing, accuracy, confidence score, confidence threshold

Lesson 6: Model cards and careers

Building on students' previous work from lesson 5, this final lesson focuses on the final stages of the AI project lifecycle: evaluating and explaining a model. Students will be introduced to model cards, which they will use to share information about how to use their models, the results of testing their models, and their models' limitations – with the aim of providing transparency. Students will also explore AI-related careers and applications, and gain insights into the various opportunities in AI and related fields.

Objectives

Students will be able to:

- Evaluate an ML model
- Produce a model card to explain an ML model
- Recognise the range of opportunities that exist in AI-related careers

Key vocabulary

AI project lifecycle, machine learning (ML) model, model cards, class, label, training, testing, accuracy, confidence score, confidence threshold

Guidance for adapting lesson content

Depending on your context, the timing of your lessons, or the confidence of your students, you may wish to adapt the lesson content. The materials are designed to allow you to extend or shrink the activities to suit your needs and give your students more space to explore new concepts and hands-on activities.

Principles for adaptation

Sequence is important

The concepts and skills have been ordered purposefully to build on from one another, so when adapting the lessons make sure to keep the order of the concepts intact. Avoid re-ordering the introduction of concepts, but feel free to re-order individual activities to start or end sessions at a suitable time.

Introductions and recaps

Throughout the unit there are semantic waves, which are designed to carefully introduce, unpack, and repack concepts. When adapting lessons you might break these waves, which is fine, but be sure to recap learnings at the end of each session and reintroduce important concepts at the start of a session. This will ensure that students are brought back to the correct point on a wave.

Use the learning graph

The unit learning graph will support you in keeping the concepts in order and providing guidance on what concepts to recap and introduce in your lesson adaptations. It's worth noting that the decision trees activity is out on a limb in the learning graph, so if any activity needs to be significantly reduced, this one is a good choice.

Examples of lesson adaptations

Below are examples of ways to split the six default lessons into nine shorter lessons, with options for extending lessons to last for an hour and for approaching the content at a slower pace.

New lesson 1 – What is AI?

Total activity time: 45 mins

Activities:

- **What is intelligence?** Starter activity from default version of lesson 1
- **The ‘intelligent’ piece of paper.** Introduction from default version of lesson 1
- **What is artificial intelligence (AI)?** Activity 1 from default version of lesson 1
- **AI applications – generative AI.** Activity 2 from default version of lesson 1
- **AI or not AI?** Plenary activity from default version of lesson 1

Extend: AI applications – generative AI – give students more time to play around with image generators

Shrink: AI applications – generative AI and AI or not AI? – reduce each activity overall so together they fit into 30 mins

New lesson 2 – Driven by data

Total activity time: 35 mins

Activities:

- **AI applications – computer vision.** Activity 3 from default version of lesson 1
- **Is a ‘smart’ speaker an AI application? Why?** Starter activity from default version of lesson 2
- **What is machine learning?** Activity 1 from default version of lesson 2.

Extend: AI applications – computer vision – allow students to look at more images

Shrink: Is a ‘smart’ speaker an AI application? Why? – reduce the activity to 30 mins

New lesson 3 – Classification

Total activity time: 30 mins

Activities:

- **The types of machine learning.** Activity 2 from default version of lesson 2
- **Classification.** Activity 3 from default version of lesson 2
- **Classification – your turn.** Plenary activity from default version of lesson 2

Extend: Classification – your turn – allow students to come up with their own classes to solve a problem close to them

New lesson 4 – Supermarket AI application

Total activity time: 30 mins

Activities:

- **The three different types of machine learning.** Starter activity from default version of lesson 3
- **Supermarket AI application.** Introduction from default version of lesson 3
- **Training a model.** Activity 1 from default version of lesson 3

Extend: Training a model – allow students to keep training the existing model or make their own models

New lesson 5 – Bias

Total activity time: 35 mins

Activities:

- **Bias.** Activity 2 from default version of lesson 3
- **Student timetable model.** Activity 3 from default version of lesson 3
- **Reducing bias.** Plenary activity from default version of lesson 3
- **What is classification?** Starter activity from default version of lesson 4

Extend: Bias – allow students to go back to their models and correct for bias

New lesson 6 – Decision trees

Total activity time: 50 mins

Activities:

- **What does a model look like?** Introduction from default version of lesson 4
- **Decision trees.** Activity 1 from default version of lesson 4
- **How decision trees are made.** Activity 2 from default version of lesson 4
- **Using machine learning to create a decision tree.** Activity 3 from default version of lesson 4

Shrink: How decisions trees are made – reduce the activity so it has only one example

New lesson 7 – Solving problems with models

Total activity time: 40 mins

Activities:

- **Decision trees in medicine.** Plenary activity from default version of lesson 4
- **Order the stages of the AI project lifecycle.** Starter activity from default version of lesson 5
- **User-focused approach.** Introduction from default version of lesson 5
- **Stage 1: Defining the problem & Stage 2: Preparing the data.** Activity 1 from default version of lesson 5
- **Stage 3: Training the model.** Activity 2 from default version of lesson 5

Extend: Stage 1: Defining the problem & Stage 2: Preparing the data – perhaps allow students to gather their own data

Shrink: Order the stages of the AI project lifecycle – consider removing this activity

New lesson 8 – Evaluation and explainability

Activities:

- **Stage 4: Testing the model.** Activity 3 from default version of lesson 5
- **Reporting on the accuracy of a model.** Plenary activity from default version of lesson 5
- **Predicting future crime.** Starter activity from default version of lesson 6
- **Evaluation and explainability.** Introduction from default version of lesson 6
- **Using a model card.** Activity 1 from default version of lesson 6

Total activity time: 43 mins

Extend: Reporting on the accuracy of a model – allow students to test their models more fully

Shrink: Using a model card – remove the practical aspect of the activity and verbally explain the concept of model cards

New lesson 9 – Careers in AI

Activities:

- **Create your own model card.** Activity 2 from default version of lesson 6
- **Careers in AI and machine learning.** Activity 3 from default version of lesson 6
- **The use of AI applications in other fields.** Plenary activity from default version of lesson 6

Total activity time: 30 mins

Extend: The use of AI applications in other fields – turn this into a group activity where students research and propose a data-driven solution for something they're interested in


Progression

Progression within the unit

We've carefully considered how students might progress through this unit, particularly in relation to the AI concepts included in the lessons. [Learning graphs](https://rpf.io/xailg) (rpf.io/xailg) have been produced so educators can visually illustrate this progression. Students will need to understand some concepts and skills before learning others – the learning graphs show how the concepts and skills are related.



Concepts and skills



The learning graphs are provided in three formats to demonstrate how learning progresses against three measures:

- Concepts and skills
- The SEAME framework (rpf.io/seame)
- Bloom's Taxonomy (rpf.io/blooms)

Pedagogy

AI is a broad and relatively new topic, but teaching it doesn't require new strategies: educators can build on a range of familiar pedagogical practices to deliver effective lessons to their students.

We've embedded these practices within the slides and activities for each lesson, but we recommend that educators use their professional judgement to review them, selecting and applying the practices relevant for their students.

Lead with concepts

Support students in the exploration of the subject area by introducing key AI concepts before engaging in hands-on activities. This approach ensures that students develop an understanding of the concepts before applying their knowledge. Videos can be used to support educators in delivering these concepts, easing the pressure on educators who may not have the subject expertise.

Work together

Encourage collaboration, specifically structured group tasks. Working together stimulates classroom dialogue, articulation of concepts, and shared understanding.

Get hands-on

Use hands-on activities so students can apply their learning and understand abstract concepts. These activities will not only support students as they engage more critically and deeply with AI technologies, but also foster their critical thinking and problem-solving skills.

Unpack, unplug, repack

Teach new concepts by first unpacking complex terms and ideas, exploring these ideas in unplugged and familiar contexts, then repacking this new understanding into the original concept. This approach, called the 'semantic wave' (the-cc.io/gr06), can help students develop a solid understanding of complex concepts.

Model everything

Model processes or practices – everything from reading decision trees to training ML models – using techniques such as worked examples. Modelling is particularly beneficial to novices, providing scaffolding that can be gradually taken away.

Create projects

Use project-based learning activities to provide students with the opportunity to apply and consolidate their knowledge of AI and ML technologies. Students can focus on developing ML models to help them better understand the decisions and trade-offs that AI developers make in real-world applications.

Add variety

Provide activities with different levels of direction, scaffolding, and support that promote learning. The activities can range from highly structured to more exploratory tasks. Adapting your instructions to suit different objectives will help keep all your students engaged and encourage greater independence.

Challenge misconceptions

Use formative questioning to uncover misconceptions and adapt teaching to address them as they occur.

Make concrete

Bring abstract concepts to life with real-world, contextual examples and a focus on interdependencies with other curriculum subjects. This can be achieved by embedding unplugged activities, analogies, storytelling around key concepts, and carefully crafted real-world examples into the lessons.

De-anthropomorphism

To support students in forming accurate mental models of AI and ML technologies, avoid [anthropomorphism](https://rpf.io/blog-anthropomorphism) (rpf.io/blog-anthropomorphism) and using words that may lead students to misunderstand machines as being human-like in their capabilities. For example, rather than saying “AI understands” or “it listens”, it is more accurate to describe AI tools as receiving inputs, processing data, and producing outputs. Adopting this sort of language will allow students to become discerning users and creators of AI technologies.

Data-driven versus rule-based

Incorporate [Computational Thinking \(CT\) 1.0 and 2.0 frameworks](https://rpf.io/ct2) (rpf.io/ct2) when teaching students about AI and ML technologies. While classical programming (CT 1.0) can be described as being rule-based and characterised by strict syntax and step-wise ordered code, ML introduces a data-driven approach (CT 2.0) that leverages vast amounts of data to identify patterns and make predictions. By regularly distinguishing between these two frameworks, students can gain a deeper understanding of how AI systems operate and the impact of different problem-solving approaches.

SEAME framework

Structure your students' learning journey around the ['SEAME'](https://rpf.io/blog-seame-framework) framework (rpf.io/blog-seame-framework), guiding them through the social and ethical aspects (SE) of AI and the applications (A), models (M), and engines (E) that drive AI systems. This structured approach empowers students to navigate and contribute to the AI field with both technical knowledge and ethical insight.

Assessment

Formative assessment

Each lesson in Experience AI includes formative assessment opportunities that are outlined in the lesson plans. These include informal quizzes, written activities, and self-guided technical tasks, giving you various opportunities to recognise and address any misconceptions your students may have.

The formative assessments are designed to be flexible and adaptable to your classroom to suit the needs of your students. Learning objectives are introduced at the start of each lesson in the slides, along with starter and plenary activities that support formative assessments.

Summative assessment

The Foundations of AI unit includes a summative assessment consisting of 19 questions. The document includes the assessment as well as the answers, and has been designed to support you to be able to quickly assess the progress made by your students and help identify any gaps in their learning. We have also designed the questions to make them suitable for uploading to online self-marking platforms such as Google Forms.

Adapting to your setting

As there are no universally agreed levels of assessment for school-aged students learning about AI, the assessment materials provided are designed to be used and adapted by educators in whichever way best suits their needs. Each lesson includes one assessment opportunity per learning objective (to inform either a formative or summative assessment) to help you assess your students' understanding. This could feed into your school's assessment process, aligning with its approach to assessment in other subjects.

Raspberry Pi Foundation

The Raspberry Pi Foundation is a UK-based charity with the mission to enable young people to realise their full potential through the power of computing and digital technologies.

Our vision is that every young person develops:

- The knowledge, skills, and confidence to use computers and digital technologies effectively in their work, community, and personal life; to solve problems and to express themselves creatively
- Sufficient understanding of societal and ethical issues to be able to critically evaluate digital technologies and their application, and to design and use technology for good
- The mindsets that enable them to confidently engage with technological change and to continue learning about new and emerging technologies

Our long-term goals

- Education: To enable any school to teach students about computing and how to create with digital technologies, through providing the best possible curriculum, resources, and training for teachers
- Non-formal learning: To engage millions of young people in learning about computing and how to create with digital technologies outside of school, through online resources and apps, clubs, competitions, and partnerships with youth organisations
- Research: To deepen our understanding of how young people learn about computing and how to create with digital technologies, and to use that knowledge to increase the impact of our work and advance the field of computing education

For more free support for teachers, including online courses to enhance your understanding of computing content and pedagogy, visit: raspberrypi.org/teach.



This resource is licensed by the [Raspberry Pi Foundation](https://raspberrypi.org/) under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International Public License (CC BY-NC-ND 4.0). For more information on this licence, see creativecommons.org/licenses/by-nc-nd/4.0.