

---

## FUNDING CALL FOR ONLINE LEARNING COURSES IN RESPONSIBLE AI (2022/23)

### CALL DOCUMENT AND DETAILED GUIDANCE

#### **Summary**

The Alan Turing Institute is pleased to announce a new funding call aimed at the development and delivery of online learning courses on the broad topic of Responsible AI. This work is part of our commitment to offering immersive learning opportunities in data science and AI with a focus on ethical and responsible approaches to designing, building and deploying AI for social good.

All researchers affiliated with a UK-based university, research institution or organisation are invited to submit proposals. Awards of up to £30,000 will be made for the successful applications. Each grant should primarily cover the costs associated with course development and delivery time.

The primary outcome of the project proposals will be self-paced learning activities and resources, and applicants will have the option to teach those courses synchronously.

*This work builds on the ‘Pedagogic innovation in (non-)cognate Data/AI Education through industry-academic co-creation’, which was supported by Towards Turing 2.0 under the EPSRC Grant EP/W037211/1 and The Alan Turing Institute. In addition, Accenture as a Turing strategic partner has provided significant resources towards this funding call.*

#### **Key dates**

- Call opens: 6 May 2022
- Drop-in Q&A: 11 May 2022 (16:00 – 17:00 BST)
- Collaborative sandpit: 16 May 2022 (14:30 – 16:00 BST)
- Call closes: 13 June 2022 (17:00 GMT)
- Outcome decisions: End of June 2022
- Project delivery: 1 July – 20 December 2022

## **Scope and Audience**

This funding call is aimed at developing learning resources focused on Responsible AI, or related areas in Data and AI Ethics.

The ethos of the call is rooted in academic-industry co-creation of learning activities and resources, and therefore the Turing encourages applicants to consider how they could work with academics and / or industry practitioners to design the learning materials iteratively and collaboratively. The proposals can also aim to bridge the gap across disciplines through training activities that support contextualisation of training methodologies.

There are two tracks that applicants should consider when applying:

### **a) Priority topics**

The Turing in collaboration with Accenture is seeking to fund training development in the following areas as a priority:

- Responsible AI – Fairness (intermediate and expert level)
- Responsible AI – Transparency (intermediate and expert level)

*A detailed breakdown of the training topics and expected outcomes can be found in [Appendix 1](#).*

### **b) General Responsible AI topics**

Proposals outside the Priority thematic areas will be considered if the applicant(s) are able to justify how the proposed course relates to Responsible AI, how the course is innovative and adds additional value to the existing Responsible AI online learning landscape.

## **Target Audience**

All courses must be suitable for researchers and industry professionals at an **intermediate to expert level**.

Activities will be advertised to our research community, including university network, subscribers to the Turing mailing lists, and the Turing network of collaborators. Anyone in the world will be able to sign up for the online learning courses, provided they have the prerequisite knowledge required for each course.

## **Format**

All materials must be designed for use as asynchronous, self-paced online courses. The expected length of the courses is minimum 5 hours up to 40 hours of learner engagement. This may include any time expected from them to go through the prerequisite materials.

*Edit: We will consider shorter or longer courses if a justification is put in place.*

All learning materials will reside on the Turing Online Learning Platform, and therefore applicants will be expected to use standard document formats and to design the courses to be effective without live instructor support. The platform supports Slack integration and live forum discussions, and applicants could make use of these to engage with the learner audience.

Optionally, proposals may also include live delivery elements, which can range from a single synchronous Q&A session to a live version of the full course. If including any optional live elements, please consider formats that can be delivered either in person or via Zoom.

## **Collaborative applications**

Applicants are welcome to apply individually if they can deliver the proposed activity without significant input from others. Alternatively, applicants can co-apply with existing collaborators from their network or attend the sandpit calls and potentially identify new collaborators that would help strengthen their application.

### **About the sandpits**

The Turing together with collaborators at the University of Newcastle are organising two sandpit calls to stimulate collaboration and identify new opportunities for joint bids to the funding call.

The sandpits will bring together academics and industry practitioners from a variety of disciplines. It will briefly introduce principles of effective instructional design and provide applicants with networking opportunities.

Following the initial sandpits, applicants will be offered individual sessions to help refine their joint proposals prior to submission. During these sessions the Turing will provide feedback to help strengthen their applications. Applicants will also have access to a drop-in clinic for further advice. All will take place via Zoom.

### **How to sign up for the sandpits**

Interested applicants can sign up for one or both the sandpit calls by submitting an entry to [this form](#). Further details will be shared closer to date to those who register.

Dates for the sandpits:

- 12 May 2022 (10:00 – 12:00 BST)
- 16 May 2022 (14:30 – 16:00 BST)

### **How to submit a collaborative application**

Applicants who are submitting a collaborative application will need to decide who is the lead applicant, in the event that two or more applicants have already started an application. The lead applicant will then be asked to insert the details of the other collaborators in the application form and provide a matrix of responsibilities shared between all. The lead applicant can also invite

collaborators to work together on the application by inviting them through the application system (more information provided in the form). It is the responsibility of the lead applicant to liaise with the collaborators in a timely manner and ultimately submit the application before the deadline.

### **Available funding**

The Turing can award the following grants to successful applicants:

#### **a) Endorsement / Refinement (up to £10,000)**

Suitable funding track for already-existing materials that require updates, refinements, or the creation of additional supporting material to make it suitable for self-paced learning on the Turing's Online Learning Platform.

An additional amount up to £3000 is available if the proposal includes the provision of a live training element.

#### **b) Adaptation / Development (up to £30,000)**

Appropriate if the course requires significant development work (i.e., the training materials do not exist at the time of the call / need to be adapted for new target groups), alongside implementation of the online learning course for self-paced learning.

An additional amount up to £3000 is available if the proposal includes the provision of a live training element.

### **Applicable to all**

Applicants are expected to justify the costs in the proposal. Funding will be proportional to the scope and development time required for the activity to be successful and suited for this call. The Turing reserves the right to make conditional offers depending on what the reviewers deem appropriate funding for each eligible, high-quality application.

Where applicants are submitting collaborative applications, the lead applicant will need to ensure the funds reach the collaborators and to pay for any other costs approved for the proposed learning activity.

The Turing will provide organisational and promotional support and will work alongside the trainers to ensure the learning activity will be successful and suited for the purpose. We do not expect funded initiatives to generate profit.

*For details of eligible costs for this funding call, please see [Appendix 2](#).*

### **The application process**

Please follow [this link](#) to start or resume an application.

Applicants are required to complete all mandatory questions and submit the application before the deadline in order to be considered for funding.

The application form includes questions addressing the following:

- Personal and professional details (for the main applicants and other potential collaborators)
- Project details: title and summary of the proposal, alignment with Responsible AI, personal motivation, optional inclusion of a live delivery element, project timeline
- Learning activity details: learning objectives, course structure outline, approaches to enable sustainability and evaluation, prerequisite knowledge / skills, approaches to enable effective self-paced learning

### **Application guidelines**

Before completing the application form, please ensure you have considered the following:

<b>Who can apply</b>	The call is currently open to anyone with expertise in topics related to Responsible AI, as well as those with a holistic understanding of the theme. A minimum of 3 years of research experience is required either in an industry, research, or public sector role.  Interested applicants must be affiliated with a UK institution.
<b>Alignment to the theme</b>	Proposals must demonstrate relevance to the topic of Responsible AI, and the added value they will bring to the research and industry learner community.
<b>Learning objectives</b>	Proposals should include a clear description of what participants will achieve by completing the course, and how the materials will facilitate these outcomes.
<b>Prerequisites</b>	Resources should clearly state any prior knowledge or skills that learners will need in order to benefit from the course.
<b>Online Learning Platform</b>	All courses will be hosted on the Turing's Online Learning Platform. Any activity or file formats not compatible with the platform may be linked to external platforms such as GitHub, Google Collab, AWS, YouTube. Successful applicants will be provided detailed guidance and on-hand support throughout the implementation process.
<b>Development and delivery timeline</b>	All materials must be created and delivered between July and December 2022. Exact schedule and project milestones will be agreed with the Skills Team upon a successful application outcome.
<b>Training at scale</b>	Applicants are encouraged to consider approaches to enable training at scale i.e., teach others how to deliver the course as educators, where there is a live delivery element. Successful applicants will be provided

	detailed guidance and on-hand support throughout the implementation process.
<b>Intellectual Property</b>	All training initiatives will be licensed with a <a href="#">CC-BY 4.0 licence</a> which allows the Turing to make the learning activities open source, while crediting the instructors who develop the lessons. Please familiarise yourself with what the licence entails.
<b>Letter of approval</b>	You will be asked to upload a signed letter of approval from your Head of Department/ Unit, or line manager (as applicable). We will provide the template and the letter only requires sign-off. This is required from each collaborator, where an application involves multiple co-applicants.
<b>Equality, diversity, and inclusion</b>	Please think of concrete examples of how equality, diversity and inclusion can will be considered and accounted for within your resources. Please also consult the <a href="#">Web Content Accessibility Guide</a> .

## **Evaluation Criteria**

Applications will be assessed according to the following criteria:

- How well does the proposal align with the topic of Responsible AI?
- Is the proposal suitable and engaging for a self-paced online course?
- Is the proposed course innovative and does it target a missing gap in the online educational space?
- Is the project timeline clear and feasible before December 2022, and where applicable, have the co-applicants worked out a matrix of responsibilities?
- Is the proposed material appropriate for the target audiences?
- Is there evidence that the applicant(s) have considered Equality, Diversity and Inclusion?
- Is there a good plan to ensure sustainability of the open-source materials?

Proposals submitted in this call will be evaluated by members of the Turing Skills Group (TSG), with the support of the Skills Team. The TSG is responsible for the development and quality of the Turing's training and development strategy and activities.

## **Any other questions**

If you have any questions or anything is unclear from this document / application form, please get in touch with the team at [skills@turing.ac.uk](mailto:skills@turing.ac.uk)

## **Appendix 1: Priority topics and audiences**

### **Responsible AI – Fairness and Transparency**

Responsible AI is the practice of designing, building, and deploying AI in a manner that contributes to a well-functioning society and economy, with fairness, reliability, security, and appropriate transparency and privacy at the core. Responsible AI helps researchers, businesses and other organisations to build trust in their AI models and scale these models with confidence.

Responsible AI is, in essence, all about doing AI well. There has been a great deal of previous research to develop theoretical techniques that can be used to power AI applications, but questions relating to whether AI is being deployed responsibly have received far less attention, at least until relatively recently.

Responsible AI is a priority area for the Turing, Accenture and for many academic and commercial organisations and is why this area has been chosen as the focus for the work outlined here.

While courses on any aspect of Responsible AI will be considered, the topics of particular priority are:

1. Fairness
2. Transparency

### **Target learner audience**

For the purposes of this funding call, all proposed courses must be aimed at an Intermediate or Expert audience, as described below. This applies regardless of whether a course covers a Priority topic, or General Responsible AI topics.

<b>Level</b>	<b>Learner description</b>
<b>Intermediate</b>	I have some familiarity in this area, but I need to know more – either because I will be specialising in this field, or because my research or role requires a more detailed appreciation of the subject than covered at an introductory level. I need to develop relevant skills and an understanding of key theoretical concepts, but I do not need to dive in to all the details.
<b>Expert</b>	I specialise in this field and learning to this advanced level is a key part of my research or professional development. I will use this information routinely in my study or work and need to have a thorough understanding of the detailed technical picture and underpinning theory. I will look to pass my skills in this field on to others in the future.

**Topic areas**

**Fairness**

This work will ideally interface with and build on material from the below:

Reference 1: Assessing and Mitigating Bias and Discrimination in AI

[alan-turing-institute/bias-in-AI-course \(github.com\)](https://alan-turing-institute.github.io/bias-in-AI-course/)

Level	Course objectives & Learning outcomes
<b>Intermediate</b>	<p>Learners will become familiar with</p> <ul style="list-style-type: none"> <li>• the landscape of algorithmic bias assessment and mitigation methods,</li> <li>• open-source fairness tools and frameworks,</li> <li>• trade-offs to be made</li> </ul> <p>Materials should go beyond binary classification (as presented in Reference 1) and explore, at a high level, fairness in more complex model types, such as multi-class classification, regression, unsupervised learning (such as clustering), reinforcement learning and recommender systems.</p> <p>The course should also allow the learner to gain hands-on experience with Jupyter notebook-based tutorials based on supplier-determined reference cases and datasets.</p> <p>Upon completion of this learning, the student should have a solid understanding of:</p> <ul style="list-style-type: none"> <li>• the mathematical representations and statistical interpretations of fairness assessment methods, covering the domain of “group fairness” such as disparate impact, demographic parity, equalized odds, equal opportunity, and the domain of “individual fairness”, as well as the respective trade-offs to be made and incompatibility of fairness metrics, particularly considering the non-discrimination criteria “independence,” “separation” and “sufficiency”;</li> <li>• the applicability of algorithmic fairness methods, for example when comparing AI models for discrete classification problems (e.g. consumer gets loan or not) vs. continuous regression problems (e.g. consumer gets credit limit X);</li> <li>• the spectrum of algorithmic bias mitigation techniques, from data pre-processing, AI model training/in-processing to post-processing techniques, along with first basic practical examples using open-source fairness tools and frameworks.</li> </ul>
<b>Expert</b>	<p>This course will extend the material in Reference 1 to consider model types beyond binary classification and will explore, in some detail, fairness in</p>

	<p>models such as multi-class classification, regression, unsupervised learning (such as clustering), reinforcement learning and recommender systems.</p> <p>Upon completion of this training, the learner should be comfortable with the primary fairness metrics and mitigation techniques, their implementations in Python and their applicability to relevant AI use case domains. This includes an advanced understanding of:</p> <ul style="list-style-type: none"> <li>• Extending group fairness to consider multiple protected groups at once and the interaction between them</li> <li>• Individual level fairness e.g. counterfactuals</li> <li>• How and why fairness of a live system can change once live, e.g. drift etc. and how to mitigate against this</li> <li>• Implementations and limitations of common fairness metrics such as equalized odds, predictive parity, counterfactual fairness and demographic parity;</li> <li>• Common (debiasing) data-preprocessing methods on which basis AI models are trained, incl. “Fair Representation Learning,” “Disparate Impact Removal,” “Optimized Preprocessing,” and “Reweighting”, followed by hands-on demonstrated applicability for known reference cases;</li> <li>• Common in-processing methods constraining or altering the training of AI models, incl. “Adversarial Debiasing,” “Fair Linear Regression,” “Fair Bayesian Networks,” and “Meta Fair Classifiers”, followed by hands-on demonstrated applicability for known reference cases;</li> <li>• Common post-processing methods mitigating biased outputs of trained AI models, incl. “(Calibrated) Equality of Odds,” “Reject Option Classification,” “Graph Laplacian Regularization for Individual Fairness,” “Post-Processing for both Individual &amp; Group Fairness,” and “Multiaccuracy Boost” , followed by hands-on demonstrated applicability for known reference cases.</li> </ul>
--	--

It is highly desirable that the above material includes a final assessment of understanding upon completion, with a clear recommendation for the pass/fail boundary.

## Transparency

This work will interface with and build on material covered in the paper below:

Reference 2: “AI in Financial Services” (section 5, “AI Transparency”, Pages 45 – 63 inclusive): [https://www.turing.ac.uk/sites/default/files/2021-06/ati\\_ai\\_in\\_financial\\_services\\_lores.pdf](https://www.turing.ac.uk/sites/default/files/2021-06/ati_ai_in_financial_services_lores.pdf)

Level	Course objectives & Learning outcomes
-------	---------------------------------------

<p><b>Intermediate</b></p>	<p>In Chapter 5 of Reference 2 (Pages 45 – 63 inclusive) a structured discussion is presented that provides a good intermediate-level overview of Transparency as applied to Responsible AI in the Financial Services sector.</p> <p>Work within this project will look to build on this content by:</p> <ul style="list-style-type: none"> <li>• Adding <b>additional</b> discussions of, and use cases for, industry sectors other than Financial Services</li> <li>• Converting the content into a format suitable for self-study</li> </ul>
<p><b>Expert</b></p>	<p>In this course, we are interested in learning the state of the art in AI Transparency. This material will be suitable for specialist practitioners in the field of Responsible AI and is likely to take a “Masterclass” or “Workshop” format, led by recognised SMEs in the field being covered.</p> <p>We are looking to cover two themes (topics) in this training, with the choice of themes being agreed with the successful applicants prior to development work commencing. In agreeing this choice, we will look to be guided by the supplier’s current areas of research and areas in which they are demonstrating thought leadership and/or developing novel approaches, techniques and models.</p> <p><i>Please note: output for this 4<sup>th</sup> Priority Learning area is different from the rest of the call, as it involves a live delivery element only – 2 x masterclass / workshop sessions with themes to be agreed covering the state of the art in AI Transparency. There will be a maximum of 20 attendees at each session.</i></p>

It is highly desirable that the above material includes a final assessment of understanding upon completion, with a clear recommendation for the pass/fail boundary.

**Appendix 2: Eligible costs**

Please use your discretion when working out the costing breakdown, and please use the below as a guide. Applications that include other type of costs, or where the rates / type of expenses are different from this guide, will be considered provided there is justification in place.

Funding will be proportional to the scope and development time required for the activity to be successful and suited for this call. The Turing reserves the right to make conditional offers depending on what the reviewers deem appropriate funding for each eligible, high-quality application.

Type of cost	Notes
<p><b>Development and delivery time</b></p>	<p>Applicants can budget £40 per hour for course development time for the lead applicant and any co-applicant / collaborator on the application, and commensurate to the scope of the proposal (see <i>Available Funding</i>).</p> <p>This would cover ‘lesson delivery’ time in the form of recording videos, and live teaching where the proposal includes a live delivery element.</p> <p>Teaching assistants and other support educators are eligible for funding and we expect a lower hourly rate for these, in the range of £20 per hour and in accordance to the standard UKRI pay.</p>
<p><b>External service providers</b></p>	<p>Applicants can consider procuring services from external providers / suppliers, in areas such as graphic design and e-learning, among others.</p> <p>The Turing has internal expertise and resources in certain areas (e.g., e-learning design, accessibility and EDI assessments, video editing) and while we will try to make use of these where feasible, applicants are also encouraged to consider working with other suppliers to enhance the quality of their proposed projects.</p>
<p><b>Software and online materials</b></p>	<p>If an activity requires specialist software, the Turing can either seek to provide the software licence from the core IT resources or reimburse the organiser(s) for any expenses incurred. The same applies for accessing paid materials.</p> <p><i>Please note the Turing promotes open science and we encourage the use of freely accessible datasets and conventional open-source software (e.g. Python, R).</i></p>

<b>Recording costs</b>	<p>While videos / recordings are not a requirement, we anticipate that video presentations will be key for presenting content in many courses, and as such, we want to ensure educators are able to create high quality recordings.</p> <ul style="list-style-type: none"><li>• Local travel costs to the Turing's office where recording can take place</li><li>• Costs to rent an alternative space for recording</li><li>• Costs to hire recording equipment</li></ul>
------------------------	---

Costs not mentioned above could be considered on a case-by-case basis. Please contact the Skills team at [skills@turing.ac.uk](mailto:skills@turing.ac.uk) to discuss before submitting your proposal.