ROCKET SCIENCE

Artificial Intelligence Policy

Developed by Rocket Science Final

August 2025

Contents

1. Introduction

This document sets out Rocket Science's approach to the use of Artificial Intelligence (AI) across our business. Fundamentally our work relies on human connection and relationships. The underlying principles for all our work will be human-led, human-informed and human-first.

As a social purpose business, we are committed to using AI in ways that promote justice, equity, diversity, and inclusion (JEDI). This includes scrutiny of how AI tools may reproduce or exacerbate bias, exclusion, and/or harm, and ensuring that their use aligns with our organisational values and advances our commitment to enabling everyone to live happy, healthy, and fulfilling lives.

We also want to ensure any use of AI in the business is ethical, considered, and appropriate. We will:

- Critically assess, on a case-by-case basis, all AI-supported work regarding its potential impact on marginalised individuals and communities.
- Avoid using AI in contexts where there exists high potential to reproduce or exacerbate inequities.
- Involve a diversity of stakeholders (e.g., clients, communities) in evaluating the appropriateness and potential implications (positive and negative) of AI use.
- Maintain full transparency with clients, participants, and wider stakeholders regarding any use of AI tools.
- Regularly review the social and environmental justice impacts of our AI-supported work and adjust practices accordingly – to be led by our Justice, Equity, Diversity and Inclusion Working Group.

This policy sets out the overall principles which underpin our decision-making on the use of AI in the business and provides specific guidance about how you can and cannot use it in your job at Rocket Science.

This will be a living document. We will formally review this every quarter to take account of

- the pace at which AI is developing.
- what we are learning from its use in the business and
- what we learn more about the positive and negative impacts of AI on the people, communities, and industries we work in.

The person responsible for the implementation and oversight of the policy is the Managing Director.

2. What is AI and for what can it be used?

Artificial Intelligence (AI) is used in computer programmes and digital tools, using substantial amounts of data to perform tasks. AI can recognise patterns in data, process it according to instructions, and use it to make predictions. AI is not 'intelligent' in the same way as humans, for example it cannot apply reasoning or critical thinking, and it does not understand what words or phrases mean, but it can sort, and process information quickly and then perform a range of tasks.

It is used in many digital platforms and apps that we use regularly, from smart speakers like Alexa, to enhancing medical screening, TV streaming recommendations, and making predictions about whether someone should be offered a bank loan.

AI has been in use for many years. More recently, Generative AI (GenAI) has become popular. GenAI creates new content based on the data it processes and the prompts that users give; this might be text in the case of something like writing assistants in Word or Google Docs, ChatGPT, or images, for example produced by Midjourney.

This explainer from the BBC provides a good explanation of how AI works, and some of the controversies surrounding it.

3. AI Policy General Principles

The use of AI within Rocket Science will be guided by the ways we run the business, through our values and principles.

- Equitable: we will proactively assess whether the use of AI risks reinforcing structural inequalities, particularly those affecting under-represented groups. We will strive to ensure AI-supported work reflects a diversity of perspectives, particularly those of people with lived experience of injustice, exclusion, and/or marginalisation.
- Accountable: we recognise that decisions informed by AI can have real-world impacts, particularly for marginalised groups. We will remain accountable for the outputs of AI-supported work and invite scrutiny, feedback and challenge from peers, clients, and communities.
- Proportionality: we recognise that the environmental costs of AI disproportionately impact marginalised communities, including Global Majority individuals, women, and disabled people. We will be limited to specific uses and where more suitable alternatives are unavailable. We will also use AI tools and models that compare favourably to others in terms of environmental impact (e.g., ChatGPT alternatives, local models). Such proportionality ensures that our AI use aligns with our Sustainability Policy and associated Carbon Reduction Plan.

It will also comply with our research ethics policy.

- Our use of AI will **respect** and **preserve** the rights and privacy of participants.
- Our use of AI will be **informed** by our duty to protect research participants from harm.
- We will continue to **design** our research and processes to ensure integrity and high quality.
- We will continue to require **informed consent** from all research participants, including explaining clearly if any AI tools will be used with participant data.
- We will **not compromise** on confidentiality, data security, or compliance with GDPR when using AI tools.

For more details, see the following section on guidance and your responsibilities. You must adhere to the guidance set out in this policy when you are considering using AI at work. You should also make sure you are adhering to our ethical, data protection and safeguarding policies when you use AI at work. Speak to your supervisor if in any doubt. Do not try to guess what is acceptable if you are not sure.

4. AI Usage: Guidance and Responsibilities

Risks and Opportunities of using AI at work

While AI is not a new technology, its availability for use by non-experts and the range of tasks it promises to deliver has expanded enormously in the past few years. The introduction of Generative AI has particularly captured the public imagination. AI is not without its downsides however, and anyone considering using it needs to be aware of both the risks and opportunities.

Risks

Impacts on workers

- Many jobs will change with the continued rollout of AI. The number that completely disappear may not be as many as commonly predicted, but many will require new skills, and the tasks that make up individual job roles may change significantly. Some jobs might become less appealing or push out people who do not have the necessary skills.
- There is some evidence that workers who use AI a lot at work have lower job satisfaction.
- One reason for this is that jobs can become focused on giving tasks to AI and checking the outputs, limiting the role of human creativity and critical thinking.
- If a lot of tasks are given to AI instead of other team members, team interaction may be reduced, affecting collaboration and staff wellbeing.
- There is evidence that intensive use of AI reduces our critical thinking skills, and makes us reliant on AI tools, affecting our overall abilities and quality of work.

Limitations of AI-produced outputs

- Generative AI regularly creates false information, sometimes called 'hallucinations.' It may get facts wrong, create entirely false references and sources, or mix up facts to create misleading statements.
- Reinforcing bias and inequities: AI tools may reproduce and even amplify existing societal biases related to race, gender, class, disability, and other protected characteristics. These biases can shape outcomes in subtle ways – for example, by marginalising certain voices in analysis or reinforcing 'deficit' narratives about communities. Without explicit mitigation, AI use has the potential to undermine inclusive practice.
- Generative AI can also create harmful and illegal content.
- Generative AI appears to 'think' and 'understand' but it is putting content together based on the statistical likelihood of, for example, one word following another. It does not understand the words or images that are generated, so it can generate odd, poor quality, generic and misleading outputs.
- It also does not know or understand the context in which you are living or working, so the outputs it generates may not be relevant to your social, political, or cultural context.
- Because AI does not understand the content it analyses or creates, tasks like creating an article summary do not necessarily capture the things that you need it to capture.

Problems with the AI business model

- Data privacy and security - companies use the data you upload to train their AI tools, and confidential or personal details may be given by AI as an answer to a query. This may breach data protection legislation including GDPR or DPAs, as well as being ethically unsound.
- Generative AI requires enormous amounts of training data, and in many cases, this is copyrighted and has been used without creators' permission.
- AI tools are often a 'black box' in which the way they work is hidden from users, resulting in a lack of transparency and explainability.
- Many AI tools begin by offering a free product but over time either introduce charges, or launch new higher-spec, paid-for products while cutting back the quality of the free ones. Reliance on free products risks being cut off if fees are introduced and may degrade human skills.
- Although exact figures are hard to come by, there are significant environmental impacts from AI. Enormous amounts of power are needed to run the data centres which AI needs to function and large volumes of water and extensive air conditioning are needed to cool them. People living close to data centres report disturbing levels of noise pollution.
- AI is pitched by some companies as the answer to everything. Without proper engagement and training, AI may be over-used for tasks it is not suited to, causing harm to their businesses, people's livelihoods, and the information environment.
- AI can benefit processing for grant makers and automate safeguarding checks. However, the challenge is that charities/grant applicants are also using AI to build their policies and applications to funders. There is a need to read the application to spot usage of AI and focus on facts and consider discussions/interviews where appropriate.

Opportunities

In the right circumstances, and with appropriate safeguards, AI can be useful in the following ways:

- Automating or streamline repetitive or time-intensive tasks to enable workers to focus on more complex or creative aspects of their jobs.
- Analysing large volumes of non-sensitive data, and identifying trends and patterns
- Improving the efficiency of logistics, coordination, task management and back-office functions, supporting team working and collaboration
- Interacting with clients, providing customer service or other outward-facing services
- Generating first-draft ideas, prompts and brainstorming.
- In certain circumstances, AI has the potential to generate insights which are too time-consuming for humans to generate.
- Checking or improving the quality of grammar and spelling, particularly if someone has challenges such as dyslexia or is writing in a second language.

A range of AI products are on the market which are designed to support different business functions (not all of these will be compatible with the guiding principles of this policy):

- **HR** - recruitment, internal policy management, wellbeing, learning and development, performance management.

- **Project management** - notetaking, scheduling, managing, and analysing workloads, document processing, creating presentations.
- **Research and writing** - finding sources, summarising text, writing assistants, knowledge management, transcription, translation, document formatting.
- **Creative and communications** - image generation, video generation, social media management, image manipulation, graphic design, data visualisation
- **Finance** - automated calculations, forecasting, processing payments and invoices, data analysis.
- **Grant making** – automated eligibility builder based on criteria, due diligence, report building.

Our approach to the risks and opportunities of AI will be continually reviewed in our governance and management of the business.

5. How do we approach using AI at work?

General approach

We prioritise quality, data protection and transparency and if AI will compromise any of these, you should not use it.

Understand and use the tools appropriately.

- Make sure that you have a general understanding of any AI tools you are using and do some due diligence. Who makes them? Where are they based? What are their policies on safety, privacy, data security, copyright, environmental impacts? What data have they been trained on? What is their position on using data which might be under copyright?
- Some non-AI tools and software are more than adequate for the things we need them to do; do not jump straight to AI tools when there might be something else more suitable.
- Build a team repository of approved tools and talk to each other about what you have tried and whether it adds value. Talk to your line manager about what you want to try and why and have a regular slot at whole team meetings where you share learning, questions and make joint decisions on any difficult calls.
- When you have the option, use AI tools that are built into software you already use and trust e.g. Microsoft Teams transcription rather than a standalone AI transcription product, and don't use any free AI tools as these will be less secure and give poor quality outputs
- You can use an AI tool to get information about how to do a more complex task, for example a data analysis method, but you should not use AI to perform the task.
- Do not use GenAI as a search engine, which is not what it is designed to do. You could use it to help you find trustworthy sources of information, but do not use the outputs of GenAI without thorough checking first. Learning how to give GenAI good prompts will improve the outputs it gives you.

AI is a tool, not a replacement for human knowledge and skills.

- Do not lose sight of where humans are best, and the value of hiring Rocket Science for a project - asking the right questions, drawing on your knowledge and experience as a team, understanding your client and the ecosystem they operate in, seeing connections and knowledge of current practice which generic AI tools will not have. GenAI will not help you create new knowledge and understanding or help you develop your institutional memory i.e. the shared repository of knowledge that you build together over the course of many projects.
- Think of AI, particularly GenAI, as an assistant, not a highly experienced professional. It is well suited to certain tasks, but it needs close supervision, and its outputs need to be checked thoroughly.
- Bring your critical thinking to your use of AI: engage with how it works and reflect on the results it generates, do not just accept everything at face value. Apply the same rigour you would in the rest of your work to understanding AI and making informed judgements on where it is appropriate in the business.
- Be transparent with your team and your clients about what AI you are using and how, and why you believe it will improve the quality of your work. Be prepared to take on board other people's views, and particularly with research participants make sure you have alternatives available if someone is not comfortable with what you are proposing.

- Only use it for initial ideas for writing, not final versions. Anything produced by GenAI needs to be subject to proper critical analysis and thought, and any final outputs that have used GenAI at any point should be the work of the human creator.

Your responsibilities when using AI.

- Use AI tools in areas where you are already competent, and do not rely on it to make important decisions. Do not expect it to be able to do everything but focus on using it on tasks for which it is designed.
- Check **EVERYTHING.** The 'hallucinations' i.e. invented information and sources that GenAI can create sound very convincing. The point of GenAI is to generate content that appears credible, but it may not be factually accurate. This also applies to the AI generated results provided by Google search. If we use AI generated content that is factually inaccurate, our work may suffer, and our reputation may be damaged.
    - Do not cite or use AI content that does not reference a source. Where it does provide a source, always go to the original to check the information provided to you. Even when a chatbot references a genuine source it may not have picked out the right data, or it may have misinterpreted the data or what you have asked it to do.
    - Is the data or content provided by AI up to date? Check this when you check sources - when were they written or produced? Have they been superseded?
    - For quantitative data, always take the data from the original source, not an AI-produced summary. Do not pick up a context-free statistic from a chatbot and reproduce it in your work: go to the source to check if it is correct, and make sure you understand the wider context.
    - Because GenAI does not understand the text or other content it creates, make sure that you thoroughly check the logic in any text it produces there may be contradictions or inconsistencies.
    - Is the information provided by AI relevant to the context that you are working in? Beware of responses that are based on, for example, information from the US rather than UK.
    - Some chatbots are not connected to the internet. These may not provide the most up-to-date information or sources.
    - Some sources are not available to any chatbots, for example academic journals which are beyond a paywall. There are some exceptions e.g. where academic publishers are collaborating with AI firms to allow them to train models on their papers but never assume that an AI tool has access to all relevant sources.
- Be aware of biases. AI image creation or editing will recreate existing biases around gender or race for example. Text generators will certainly have been trained on data which reflects societal biases and so will reproduce these in the text they generate. Think through the biases that might exist in the work you are doing:
    - are you working with a community that is routinely excluded?
    - Are there pre-existing concerns around gender, race or any other biases which might be amplified through using AI? Ask yourself: Do I really need to use this tool? Who benefits from the use of this tool? Who might be excluded or harmed?
    - Consider whether the underlying data and training sets represent diverse communities and geographies (many AI tools are trained on US- and Euro-centric data).

    - When using AI to support research, evaluation and learning, actively seek to counterbalance the limitations of the tool by drawing on diverse voices, co-design methods and lived experience.
- Never submit AI-generated content as a final draft. If you use it to get started on things e.g. producing ideas, suggestions of sources to look at, or summaries of articles, you must build on these starting points with your own work. No AI-generated content should be directly reproduced in a final output: do not just copy and paste AI content into your work. All outputs will be reviewed by a Director/Head for quality assurance.
- Disclosure is essential. If you have used AI at any point in a piece of work, you must disclose this within the final version of the work. You should note what tool(s) you used, for what purposes, and how you have ensured the accuracy and quality of the work.
- You remain responsible. The person who creates an output, be it an article, blogpost, report, briefing, literature review, marketing content, financial reports or anything else, is responsible for that content, whether they have used any AI tools.

If you are found to have used AI in your work without disclosing it to your manager, and/ or you have used a tool which has not been discussed and approved within the organisation, this will be grounds for disciplinary review. If a project lead/ manager suspects that AI has been used without disclosure or authorisation they may ask you about how you produced the work.

How to make decisions about data and AI

Recognise that any data you feed into an AI tool is certainly not secure, and act accordingly. Do not input any of the following types of data into an AI tool:

- Personally identifiable information e.g. names of research subjects, grantees and or beneficiaries
- Primary research data such as interview or focus group transcripts, survey responses or any other original material that you have received from a research participant or subject or otherwise generated. This is particularly important for any material which contains any details on that individual's personal or family life, medical conditions, sensitive personal circumstances, or anything else that they have a reasonable expectation will remain private.
- Any documentation or content from a client which is not in the public domain.
- Any confidential documentation relating to a bid or request for proposal, including documents from potential clients that are not in the public domain, and internal bidding documents.
- Research findings or recommendations for a project which will not be published / released by the client.
- Any financial information relating to Rocket Science or its clients, including funders and grantees.
- Any content which contains someone else's Intellectual Property, which is not in the public domain, including from grantees.

If in doubt, talk to your project director. Do not use GenAI tools to produce content that is potentially harmful, misleading, offensive, or illegal.

6. When to use AI

Flow chart for AI decision-making.

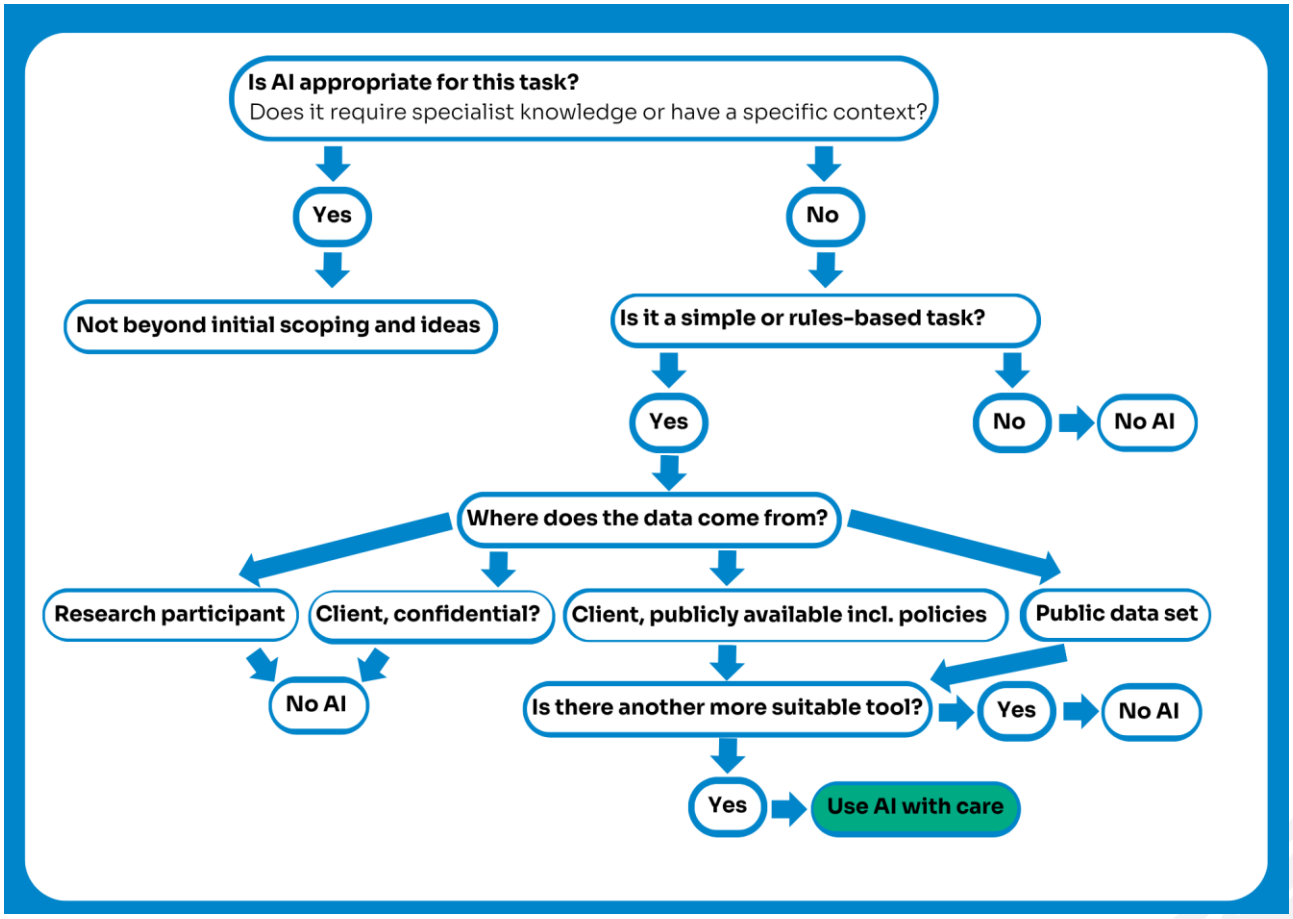

- First, think carefully about whether AI is appropriate for the task. Does the task need contextual knowledge, expertise and understanding? Is it a rules-based task (binary response/closed questions) which are simple to provide a prompt for? Does it rely on original data which might be sensitive?
- AI can be helpful for some tasks such as research and grant processes but not all, and ethical considerations and a focus on quality should be the most crucial factors in deciding whether to use AI.
- Tasks where AI could be applied, subject to testing and agreement within the team, and thorough checking of all outputs include.
    - Research:
        - Summarising long documents as part of a literature/ desk review to indicate whether they are relevant and should be read in full.
        - As a starting point for a review of a policy landscape or horizon scanning
        - Identifying potential sources for qual and quant data
        - Processing data into user-friendly formats
        - Summarising or transcribing non-sensitive meeting or discussion notes.
    - Grant making:
        - Reviewing public data including legal registers, finances and 360 giving data
        - Creating an eligibility quiz based on programme criteria.

- Automatically checking organisational policies in agreement to set standards such as a safeguarding policy having a named Designated Safeguarding Lead and being up to 2 years old
- Summarising projects
- Drafting unsuccessful communication based on assessment responses.
- Improving visualisation of data including recommendation portfolios, outputs, and outcomes data.

In making decisions about using AI, think about

- Where has the data come from? Have we created it ourselves from original research? Has it been supplied to us by a client? Is it publicly available? If it is not routinely publicly available, you should not input data into any AI tools.
- Does the data include any personal or sensitive information given to you by research participants, or passed to you by a client? If the data were in the public domain would this cause a data breach?
- How specialist is the topic? Is a general all-purpose AI tool likely to give you up-to-date, accurate information which is properly contextualised and explained?

7. Further Reading

Demos Generative AI Policy Paper

Evaluating Generative AI Tools

About ChatGPT & Generative AI LLMs - Using AI Tools in Your Research

How to use AI to do practical stuff: A new guide

Renee's four golden rules of artificial intelligence | Safe Hands

Sensible and Ethical Use of Generative AI and other assistance tools | UNE myLearn

Free AI Risk Assessment Tool for the Third Sector

ROCKET SCIENCE

About Rocket Science

Rocket Science is a social purpose business, working towards a future where everyone can live healthy, happy, and fulfilling lives. Our objectives are to help our clients reimagine systems, maximise investment to have the greatest impact and transform lifetime outcomes using a whole-person and prevention lens.

We do this through research, mapping and analysis, evaluation and impact measurement, strategy and service design, participatory and peer research, learning, collaboration, and end to end fund management.

Offices in:

Edinburgh
London
Newcastle

rocketsciencelab.co.uk

INSTITUTE FOR SOCIAL VALUE

CYBER ESSENTIALS CERTIFIED PLUS

disability confident

LONDON FUNDERS

MINDFUL EMPLOYER

Living Wage Employer