

GenAI Ethics

A task sheet for students to work through several times and hopefully then internalise.

Evaluation criteria	Notes
<i>How can we identify and mitigate bias in GenAI outputs?</i>	Fairness, Bias, and Inclusion:
<i>Who is responsible when GenAI perpetuates or amplifies stereotypes?</i>	Fairness, Bias, and Inclusion:
<i>Does GenAI treat all dialects, cultures, and languages fairly?</i>	Fairness, Bias, and Inclusion:
<i>Are GenAI systems being designed for the needs and values of the global South and marginalized communities?</i>	Fairness, Bias, and Inclusion:
<i>Who gets to shape the future of GenAI — and who is being left out of that conversation?</i>	Fairness, Bias, and Inclusion:
<i>Should users always be informed when they're interacting with GenAI rather than a human?</i>	Transparency, Consent, and Privacy:
<i>How can GenAI systems be made more explainable to non-experts?</i>	Transparency, Consent, and Privacy:
<i>What level of transparency should be required about how a GenAI system was trained and what data it used?</i>	Transparency, Consent, and Privacy:
<i>Is it ethical to train GenAI on publicly available data without explicit consent?</i>	Transparency, Consent, and Privacy:
<i>Should individuals have a right to opt out of having their data used in AI training?</i>	Transparency, Consent, and Privacy:
<i>Who is accountable when GenAI causes harm — the developer, deployer, or user?</i>	Responsibility, Regulation, and Control:
<i>Should GenAI systems have ethical "safeguard triggers" or constraints to prevent misuse?</i>	Responsibility, Regulation, and Control:
<i>How should the law evolve to handle damages caused by autonomous GenAI systems?</i>	Responsibility, Regulation, and Control:
<i>How should society respond to the use of GenAI to generate deepfakes or propaganda?</i>	Responsibility, Regulation, and Control:
<i>Can GenAI-generated content be used to manipulate public opinion — and who should regulate that?</i>	Responsibility, Regulation, and Control:
<i>Can GenAI-generated content infringe on copyrights or trademarks?</i>	Ownership, Labour, and Creativity:
<i>Who owns the output of a GenAI system — the user, the company, or the model itself?</i>	Ownership, Labour, and Creativity:
<i>Should GenAI be required to cite sources or inspirations for its outputs?</i>	Ownership, Labour, and Creativity:
<i>What is the ethical obligation of companies using GenAI to workers whose jobs it may replace?</i>	Ownership, Labour, and Creativity:
<i>Is it ethical to use GenAI to replace human labour in creative fields like art, journalism, or music?</i>	Ownership, Labour, and Creativity:
<i>How do we balance the environmental cost of training and running GenAI with its benefits?</i>	Wellbeing, Autonomy, and Sustainability:
<i>Should there be sustainability standards for GenAI models and infrastructure?</i>	Wellbeing, Autonomy, and Sustainability:
<i>How can we prevent overreliance on GenAI in decision-making processes (e.g., healthcare, law)?</i>	Wellbeing, Autonomy, and Sustainability:
<i>Should GenAI be allowed to generate mental health advice, diagnoses, or medical guidance?</i>	Wellbeing, Autonomy, and Sustainability:
<i>How can we protect personal or sensitive data from being leaked by GenAI?</i>	Wellbeing, Autonomy, and Sustainability: