

AI for Social Good

Artificial Intelligence (AI) personalised Education helping socially excluded people to prosper.

Philip Treleaven¹ and Daniel Brown^{1,2}
¹University College London and ²MegaNexus

Abstract

‘The AI genie is out of the bottle, for good’.

The imminent AI ‘tsunami’ (led by Generative AI) will be hugely disruptive for business and society. However, it also offers pioneering opportunities for social good, notably personalised education/training for socially excluded and disadvantaged groups (e.g., people in prison, people on probation, refugees, long term unemployed, low-aspiration and disadvantaged young people in education), thereby transforming their future and addressing major social problems. The UK has over 80,000 prisoners and education is the best deterrent to reoffending, which costs £18bn pa (UK Parliament, 2022).



Figure 1: the AI Genie is certainly out of the Bottle.

For context, we are now on the threshold of the AI (cf. Fourth Industrial) Revolution. Thirty years ago, we had the IT (cf. Third Industrial) Revolution: where it was essential for everyone to learn IT skills, such as Microsoft Office, to future-proof themselves. The AI revolution will be characterised by specialised AI *assistant* tools controlled by a human user, interacting through natural language processing (text, speech, image, video). The AI assistant will do the preliminary work, provide a review of the draft, and answer questions; with the human user verifying and completing the material. As an illustration of AI tools supporting professionals, a legal AI assistant (cf., Harvey.AI) might draft a contract, give comments on potential ambiguities and risks; with the expert human lawyer reviewing and correcting the contract. This is referred to as ‘human-in-the-loop’.

Generative AI can positively impact society in at least five ways:

- **Professional productivity tools** – AI assistants for professionals (e.g., lawyers, accountants, architects, teachers, HR, civil servants, social workers etc.). An example being the legal AI assistant described above.
- **Consumer ‘self-help’ tools** – AI assistants to support consumers in achieving things for themselves. An example is tax returns, pension advice, or consumer rights.
- **Education tools** – specialised Chatbot-driven education and training tools. For example, a generative AI tool trained on a specific textbook with a Chatbot user interface, whereby the student can hold a dialog with the book.
- **Lifestyle enhancement** – a ‘catch-all’ for AI assistants to improve individual’s wellbeing. Examples include avatar friends for children, digital companions for care home residents, tech advice for the elderly etc.

- **Social good** - personalised education/training for *socially excluded people* (e.g., prisoners, refugees, long-term unemployed, vulnerable adults, young people with limited opportunities in poorly performing schools) helping the poorest people in society prosper.

There are many threats ascribed to AI, but the intended contribution of this paper is to discuss *AI for Social Good*, introducing underlying AI technologies; reviewing role model STEM¹ education platforms supporting teachers and students; and presenting as a case study the **GenAIE** (Generative AI Education) platform supporting keyworkers to help prisoners, probationers, disadvantaged young people in education and refugees. Our mission is two-fold: firstly, to *AI-empower* keyworker professionals through ‘hands-on’ AI training using GenAIE; and secondly, to *AI-educate* disadvantaged users with personalised educational content. As a starting point, UCL and MegaNexus are working with educational professionals to produce personalised training content for the specific sector of Justice, including prisons and probation, described below, which evidence and demonstrate the positive power of generative AI for social benefit.

1. AI for social good

Thirty years-ago Society experienced the IT (cf. Third Industrial) Revolutions. Now as they say: ‘the AI genie is out of the bottle’. Society is on the threshold of the AI (cf. Fourth industrial) Revolution with its huge impact on business, professional employment, and students’ education. Textbooks will come with a Chatbot tutor. A Lawyer will no longer write a contract from scratch; a generative AI tool will produce a draft contract, identifying the key elements and risks for the legal professional to review (cf. human-in-the-loop). Clearly this constitutes a huge disrupter for the professions. Furthermore, and interestingly, aspiring programmers are already using generative AI (plus no-code and low-code tools) and ‘GitHub Copilot’ support to prototype programs. However, everyone will benefit from basic programming skills to read and customize their programs.

The AI revolution can help keyworker professionals supporting socially excluded and disadvantaged groups of people, empowering these keyworkers to produce personalised education/training content for their clients, contributing to solving some of Society’s major social problems.

The core AI education/training ingredients being brought into the mix include:

- **AI for social good** – focusing on providing AI productivity tools for professionals and education/training tools for socially excluded people; the lowest social stratum in a country or community, consisting of the poor, sick and unemployed.
- **Human-in-the-loop** – self-drive AI tools for professionals, a blend of supervised machine learning and active learning where humans are involved in both generating the training material and testing stages of building personalised educational content.
- **Professionals** – professionals in control of the AI assistants, providing educational/training material and monitoring its usage.
- **Socially excluded** - the lowest social stratum in a country or community, consisting of the very poor, and disadvantaged children, with very little power or chance to improve their lives.

¹ STEM - Science, Technology, Engineering, and Mathematics education

Ecosystem	Prisons	Probation service	Refugee counsellors	Education	Social Workers	Charities
Education/Training	Career training	Career training	Social training	inspiring disadvantaged children	Care workers	Support
Keyworker professionals	Ministry of Justice	Probation workers	Caseworker	Teachers	Social workers	Charity workers
Socially excluded groups	Prisoners	Probationers	Refugees Migrants	Low aspiration students	Carers in the Home	Drug addiction Alcohol abuse Betting addiction

Figure 2: AI Education/Training for Social Good

2. STEM Education Platforms

STEM education platforms for supporting teachers and students provide exemplars for what we are seeking to achieve with our Generative AI for the socially excluded and disadvantaged groups. STEM Education, at its core, simply means educating students in four specific disciplines, namely, Science, Technology, Engineering, and Mathematics (collectively shortened as STEM).

There are over 600 UK organisations involved in STEM education, from professional institutions to learned societies, universities, museums and science discovery centres, specialist education enrichment providers, teacher professional development organisations and subject-specific associations (RAEng, 2016). The STEM education support offered by these organisations covers the full range of primary, secondary, and further education, vocational education, apprenticeships, and university programmes. The provision ranges from teacher continual professional development (CPD); talks and lectures; classroom-based as well as extra-curricular activities to enhance and enrich the STEM curriculum; specialist and subject-specific educational resources; STEM degree and career pathway mappings and resources and much more.

The main aim of STEM education platforms and networks is to inspire, inform and encourage young people to consider STEM degrees and careers and to help strengthen and diversify the UK STEM workforce. To achieve this, they focus on the following key areas:

- Increased understanding and attitudes towards engineering among young people and their key influencers (for example, parents, teachers, carers, youth workers etc.).
- Increasing diversity and widening access to under-served communities and under-represented groups.
- Increasing and offering high quality support and further development for teachers of STEM subjects, through innovative teaching and employer engagement methods.
- Embedding STEM throughout primary and secondary education with a focus on early interventions.
- Supporting teaching and learning in further education and promoting technical and vocational routes into STEM.
- Development of specialist resources focusing on advice and guidance on STEM career and degree pathways across different sectors, including applying to university, work experience, research and industry placements.

Using generative AI ChatGPT for personalised education/training for low-aspiration and disadvantaged young people takes STEM education to a whole new level.

3. AI Algorithms

To demystify generative AI (e.g., ChatGPT), algorithms cover three broad domains: Computational Statistics (e.g., Monte Carlo methods), Complex Systems (e.g., Agent-Based systems), and Artificial Intelligence (e.g., Artificial Neural Networks). See Figure 3.

- **Computational Statistics** - computationally intensive statistical methods.
- **Complex Systems** - system featuring a large number of interacting components whose aggregate activity is nonlinear.
- **AI Algorithms** - mimicking a new form of human learning, reasoning, knowledge, and decision-making.
 - Knowledge-based systems - where knowledge is explicitly represented as IF-THEN rules.
 - Evolutionary algorithms - algorithms for global optimization inspired by biological evolution (e.g., Genetic Algorithms).
 - Machine learning – algorithms with the ability to learn without explicit programming and can change when exposed to new data.

Figure 3: Algorithm taxonomy (Koshiyama et al, 2020)

At a simple level, AI machine learning covers:

- **Traditional AI** - ML models to identify patterns within a training data set and make predictions.
- **Generative AI** - used to describe any type of AI used to dynamically create new ‘human-like’ texts, images, speech, video, programs, or synthetic data.

And for the future:

- **Algorithmic General Intelligence (AGI)** - solving problems as well as humans; faced with an unfamiliar task, the AGI system could find a solution.
- **Algorithmic Superintelligence (ASI)** – solving all problems better than people (ASI, 2023) across a comprehensive range of categories and fields of endeavour.

ChatGPT and Generative AI

ChatGPT and its growing number of rivals, including Google’s Bard, Anthropic, Character.ai, Cohere and others have proven capability of holding human-level conversations, but are also subject to generating inaccurate, unethical, and mis information. Fascinating is *AI hallucination*: a confident response that is biased, too specialised, even totally wrong (Ji et al, 2022). These AI hallucinations can be caused by various factors but the most common is when it has limited training data and hasn't been taught to say: “I don't know that answer.” Instead of saying that, it will make something up that seems like it could in fact be the answer. This illustrates the huge importance of the keyworker professional ‘human-in-the-loop’ driving the AI algorithm.

OpenAI’s ChatGPT is acknowledged as a leap in innovation (cf. iPhone); GPT 4 is trained on 170 trillion parameters. Plug-ins are now available using Microsoft 365 Copilot and GitHub, with numerous companies launching generative AI ad-ins to their products, such as Bloomberg GPT (Bloomberg, 2023). And on the horizon is Algorithmic Superintelligence (ASI) far surpassing most gifted human minds.

Generative AI has massive implications for business, education, employment, but also law enforcement and crime. As commented by the UK Sunday Times newspaper (Fortson, 2023) “the ability to entwine a chatty, randy, bot with the image or video of a [real] person made to order to meet a user’s preferences is [already] here”.

Key terms are:

- **Generative Pre-trained Transformers (GPT)** - a family of AI models that uses the transformer architecture and is a key advancement in artificial intelligence (AI) powering generative AI applications such as ChatGPT (Amatriain et al., 2023).
- **Large Language Models (LLM)** - LLMs are a subset of AI trained on a vast quantity of (online) data to produce human-like responses to dialogue or other natural language inputs (Riedl, 2023).

- **Generative AI** - a broad label for any type of AI used to dynamically create new 'human-like' content (Lawson, 2023).

4. Review of Generative AI Tools

Generative AI can positively impact society in a number of ways:

- **Professional productivity tools** – AI assistants for professionals (e.g., lawyers, accountants, architects, creatives, HR, civil servants etc.).
- **Consumer 'self-help' tools** – AI assistants to support consumers in achieving things for themselves. An example is personalised investment, tax returns or pension advice.
- **Education tools** – specialised AI-driven education and training tools.

There are literally thousands of generative AI tools entering the marketplace (Qian, 2023). We highlight some professional and consumer applications, and AI tools as exemplars.

Professional GPT applications

In the business area, generative AI is producing a new generation of AI professional productivity tools. Examples include:

- **GitHub copilot** - brings the power of generative AI to engineering teams, accelerating the speed of software development and innovation.
- **Harvey.AI** - assists with legal contract analysis, due diligence, litigation, and regulatory compliance and can help generate insights, recommendations, and predictions.
- **Legalfly** - provides contract review & drafting experience to save time and bring AI into legal teams.
- **Casetext** - is an AI-powered legal research platform that uses natural language processing (NLP) to analyse legal documents and identify relevant case law.
- **RadiologyGPT** - a Large Language Model for Radiology.
- **Skillate** - streamlines the recruiting process by providing enhanced applicant experiences, intelligent recruiting services backed by AI, and people analytics.

Consumer GPT applications

As expected, there are a growing number of consumer AI tools especially in the creative industries, such as art, music, video and even beauty.

- **Dall-E** - is an AI system that can create realistic images and art from a description in natural language.
- **Hemingway** - highlights aspects of "poor" style including overlong sentences, passive voice, and excessive adverb use.
- **Voicify** - an AI-based music creation platform that allows users to generate custom songs in with the voices of their favourite artists.
- **Bing and Bard** - are AI enhanced search engines using interfaces to Large Language Models as a complement to a more standard search experience

Generative AI Tools

Generative AI tools are either general-purpose (e.g., ChatGPT, LaMDA) or domain specific (e.g., GitHub Copilot):

- **ChatGPT** – an AI chatbot that uses natural language processing to create humanlike conversational dialogue enabling users to refine and steer a conversation.
- **LaMDA** - a natural language processing tool developed by Google with 137 billion parameters.

- **GitHub Copilot** - a cloud-based AI tool developed by GitHub and OpenAI to assist programmers by auto-completing code.

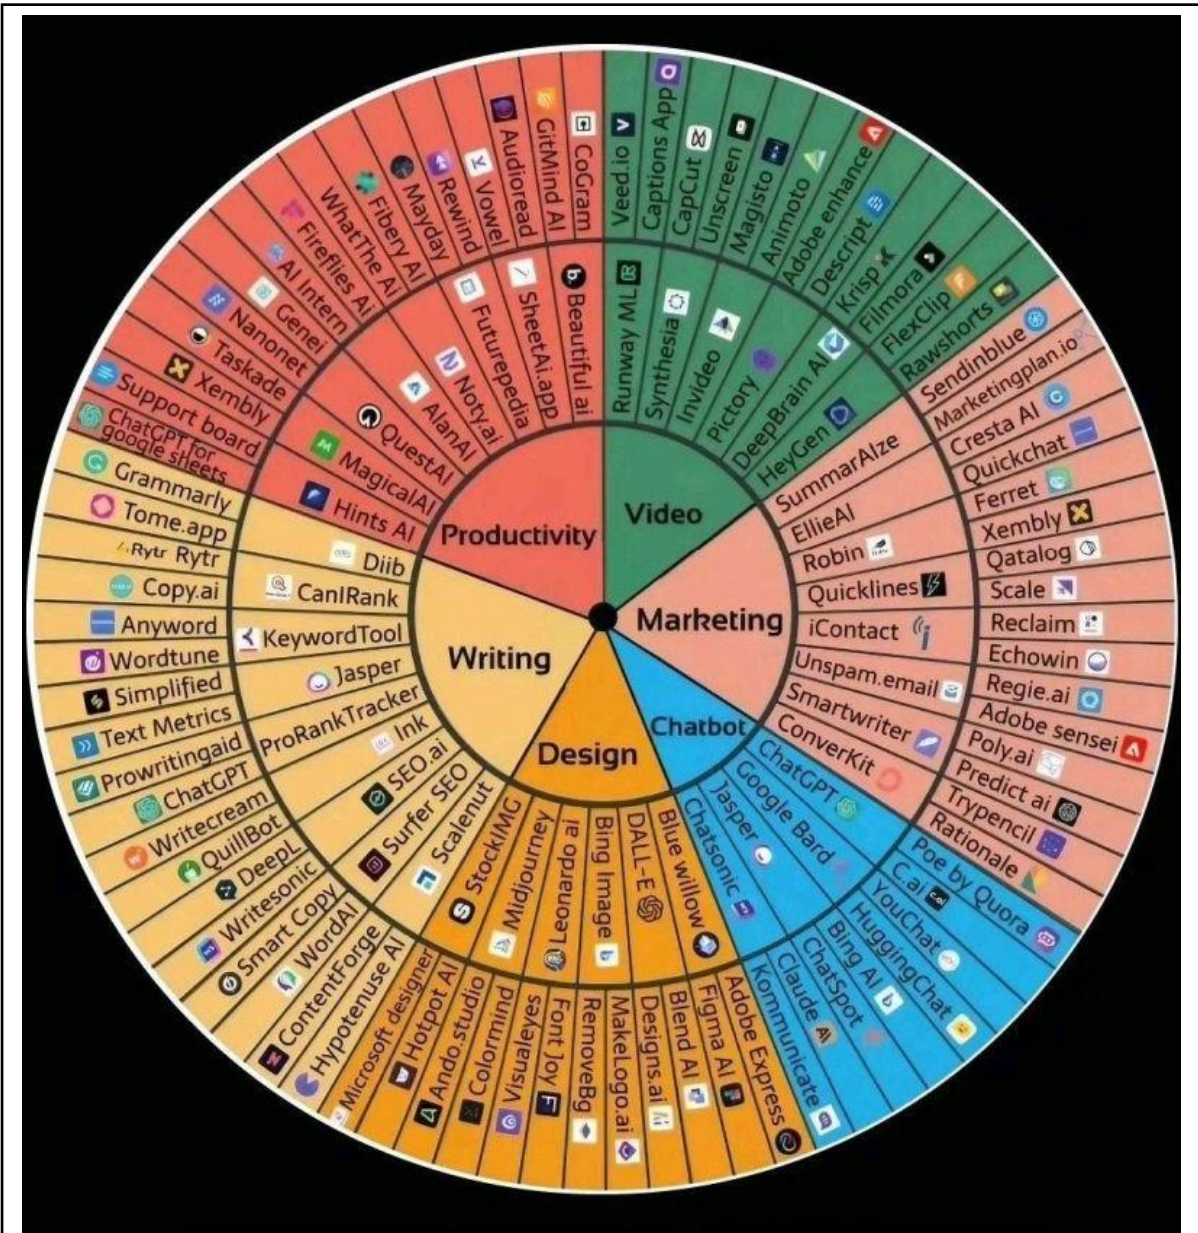


Figure 4: 120 AI Tools (LinkedIn, 2023)

The plethora of AI tools and Apps are multiplying and already supporting increasing volumes of professionals and consumers alike. However, whilst they have on the one hand been accompanied by many prophecies of machines overrunning human beings, on the other hand they have not been at all focused on supporting those who could benefit the most: the socially excluded and disadvantaged groups at the fringes of society who do not have the resources to seek them out.

5. Generative AI training and risks

The AI tools we have developed are to support keyworker professionals to produce and monitor personalised educational/training content, rather than giving end-users free rein. As an illustration, if prisoners are given uncontrolled access to a chatbot they can ask “what skills do I need to become a chef?” or equally “what do I need to escape from prison.” In computing, this *walled garden* provides

an environment that moderates the user's access to network-based content and services. In effect, the walled garden directs the user's navigation within particular areas to enable access to a selection of material and prevent access to other material.

AI Training Data

Training data refers to the massive datasets used to teach AI models to detect patterns, make predictions and generate new information. The performance and propensity for incorrect responses, misinformation, and *hallucinations* of an AI system hinge critically on the quality and diversity of its training data (Ji et al, 2022). If data used lacks diversity or contains biases, the AI can develop blind spots and make unreliable judgments, especially in unfamiliar cases. For example, an AI system trained primarily on images of young, light-skinned individuals may have limited ability to recognize darker-skinned or older faces. In practice, limited data availability, inconsistent labelling, and lack of diverse perspectives can lead to gaps in training data. Continually assessing and strengthening AI training data are crucial to developing systems that produce fair, accurate and trusted results for all. Careful curation must remain a collaborative effort between data scientists, experts, and communities affected by AI systems.

AI Generated Information

An interesting new role is *Prompt Engineering*, the art of crafting prompts that effectively instruct GPT tools to generate the desired output. Information challenges include:

- **Opaque information** - hard to understand; not clear or lucid; and obscure generated information.
- **Inaccurate information** – data that is inaccurate, incomplete, or inconsistent, but not deliberately created.
- **Misinformation** – false or inaccurate information, especially that which is deliberately intended to deceive, including false, disinformation which is intended to mislead, especially propaganda issued by a government organization to a rival power or the media.
- **Inappropriate information** - inappropriate content consists of information or images that that should not be generated or provided, that might lead to unlawful or dangerous behaviour. For example, a prisoner asking for advice on escaping, or a user asking how to make an illegal drug.
- **AI Hallucination** – generation of outputs that may sound plausible but are either factually incorrect, or unrelated to the given context; with the algorithms generating additional nonsensical content when challenged by the user (Ji et al, 2022).

Professional in the Loop

The above is the overwhelming reason for the 'Professional in the loop', to produce educational/training content and ensure that the generated information is moderated. Without proper safeguards, such technology could enable, for example, misinformation and 'deepfakes' to deceive audiences.

AI *hallucinations* (Ji et al, 2022) pose a significant challenge as generative AI systems become more advanced and widely adopted. An AI hallucination refers to when an AI model generates a confident yet inaccurate response. The causes of AI hallucinations are multifaceted. Limited or biased data used to train the AI can skew its knowledge and judgments. Lack of data on certain populations or topics can lead to overconfidence in generating erroneous information about those areas. Flaws in the AI system itself, including lack of safeguards to detect and mitigate unreliable responses, enable hallucinations to arise and spread. Addressing AI hallucinations remains an open challenge requiring persistent human governance and guidance.

6. Case study: GenAIE GPT platform

GenAIE (Generative AI Education) was developed as an automated content generation platform for professionals. Firstly, it provides public sector keyworkers with an introduction to AI and an exemplar of AI productivity tools. Secondly, it provides eLearning content across the educational and vocational spectrum. GenAIE builds upon the MegaNexus Community Campus software used in the UK Probation Service to provide a personalised learning engine that builds learning pathways for each user.

The Community Campus educational platform provides a secure means of access to selected learning materials at different levels and in different languages that align to the needs of each individual learner. This in turn requires content that is appropriate to the requirements of each individual in terms of vocational area, skills, learning abilities, job aspirations and language. GenAIE is a natural development of Community Campus. GenAIE supports keyworker professionals to build courses, lessons and resources using generative AI. This accelerates development and supports moderated subject matter delivery. In addition, all content materials can be provided in multiple languages important, especially for disadvantaged users with limited English.

The requirement was for a system to support keyworker professionals in course planning and in automatically building lesson content in different formats and media and at different levels in different languages. Personalised to individuals' learning abilities and tailored to the key stage level of learning needed. Educational AI automatically provides lesson plans, resource hyperlinks and extra-curricular activities such as flashcards, quizzes and wordsearches, with the ability for human teachers to edit, amend and insert additional information to the lesson content.

Content also needs to be published in a variety of formats. For example, with Moodle a course between 2-10 modules provides teachers with a backbone of subject knowledge, alleviating lesson preparation time to allow focus on targeted individuals whilst bringing creativity and life-support to the subject. In addition, AI systems need to provide consistent assessment points to enhance understanding of students' progress giving better guidance towards the grades needed to pass and to excel.

GenAIE (Generative AI for Education) is a fully automated content generation platform for professionals. It supports keyworkers to build, within minutes, courses, lessons and resources inclusive of flexible modules and lessons, extra-curricular resources and audio assistance using generative artificial intelligence.

This accelerates course material development, supporting moderated subject matter delivery across the educational and vocational spectrum, to enable teachers to best focus on communication with and support for students.

GenAIE can operate either independently or with an existing LMS platform to provide personalised education with learning pathways for each user and automated content provision configured for the needs of each individual. This content is designed according to the requirements of each individual in terms of vocational area, skills, learning abilities, job aspirations and language with all content materials available in multiple languages, especially significant for disadvantaged users with limited English.

The GenAIE platform utilises a microservices architecture and generative AI to wholly automate educational and vocational content creation; it is cloud based for scalability. This architecture deploys the application as a collection of services, providing the framework to develop, deploy, and customise services independently.

GenAIE leverages a set of different fine-tuned Transformer-based natural language models. Natural language generation workflows chain together prompt engineering, large language models (LLMs) and helper agents into pipelines. These are coordinated by different agents or Genies (Curriculum, Student, Teacher and Moderator) with different motivations performing different tasks. They ensure that the content produced is appropriate for a varied set of stakeholders with different needs and wants.

As a result, GenAIE supports professionals to generate curriculum resources on any subject matter, with support for multi-language content and learning levels spanning from early primary to advanced.

Figure 5 summarises course generation flow, built around a collection of ‘Genies’, each with a specific function:

- **Curriculum Genie** - ensures compliance against known curriculums, legislation, and regulations.
- **Teacher Genie** - generates lesson plans, learning material, activities, signposted resources.
- **Student Genie** - ensures language and content are appropriate for the level expected.
- **Moderator Genie** - guardrails around profanity, triggering language and inappropriate materials.

During the production of content each ‘Genie’ operates with the others to ensure that moderated course content aligns to a set curriculum, covers appropriate signposted resources, and accords with students’ expected levels, (vocational) interests and capabilities. To support a keyworker in producing educational course material, GenAIE utilises a multi-step collaboration approach:

- **Curriculum Bot** - performs focused LLM retrieval to source related curricula.
- **Teacher Genie** - reviews and customizes the curriculum.
- **Outline Agent** - uses prompt templating to generate a course outline based on a requested number of lessons.
- **Teacher Genie**- refines the outline.
- **Lesson Agent** - leverages fine-tuned Text to Text Transfer Transformer (T5) to generate individual lessons.
- **Teacher Genie** - edits lessons prior to finalization.
- **Resource Agent** - uses AI visual generation tools and LLMs to generate supplemental assets including audio, images and quizzes and multiple languages.
- **Export** - the curriculum packages are exported to an appropriate format (e.g., SCORM) for Learning Management System integration (e.g., pptx for editing, or PDF for print).

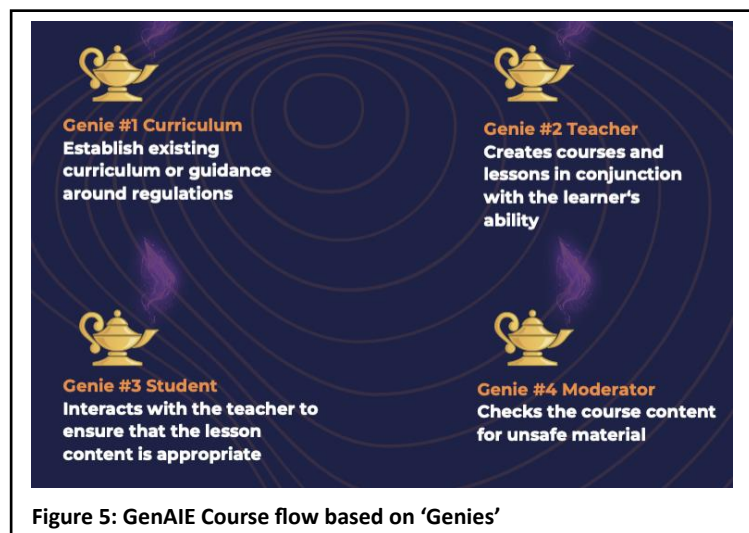


Figure 5: GenAIE Course flow based on ‘Genies’

Knowledge Transfer

In summary, our ‘AI for Social Good’ mission is to empower keyworkers through knowledge transfer using a three-stage process:

- **GenAIE deployment** – UCL and MegaNexus work with educational professionals to produce personalised training content for a specific sector, such as people in prisons and people on probation described below, to demonstrate the constructive power of generative AI.
- **Keyworker AI training** – a training and awareness course for keyworker professionals on AI and emerging technologies and their impact on social services. As part of the training, professionals prepare sample educational material using GenAIE.
- **Keyworker empowerment** – professionals independently use GenAIE to produce educational material and also monitor its deployment and usage.

7. Case Study: GenAIE Education for people in the justice system

As a case study we look specifically at education for people on probation, and the preparation of personalised educational material by keyworker specialists.

As background, reoffending in the UK has a £18bn (\$23bn) pa cost (UK Parliament, 2022) with one of the greatest recidivism influences being lack of education and training. A major challenge facing keyworker teaching professionals supporting learning for probationers is targeting relevant course materials that support each individual on their personal pathway at the right learning level, in an appropriate format with courses that meet their interests and (vocational) needs. The volume of permutations required to supply these materials has historically been very challenging with the cost of bespoke configuration of content a major obstacle.

The UK Probation Service required a large-scale provision of educational materials to People on Probation (PoPs) to support them on a variety of vocational pathways. The key requirement was to provide moderated content generation that keeps expert human educators in control of the course generation process but saves time in developing appropriate effective content at different levels, in different languages calibrated to the capabilities of different learners. The aim is to allow teachers to focus on curriculum and learning pathway design and to gauge the enhancements and improvements needed for content and curricula without having to work on the more onerous and time-consuming course materials implementations that are less effective uses of their human skills.

The GenAIE platform is used by keyworkers to provide a variety of tailored educational/training materials calibrated to individual learners. Course materials are produced through generative AI using a system developed in a partnership between a team at UCL and MegaNexus, to support development of online learning materials for socially excluded people on probation.

Figure 6 illustrates GenAIE course development. For the ministry of Justice, a range of courses were developed to support learners in prison and on probation using generative AI. These incorporated the adversarial genie approach highlighted above to ensure that content is made appropriate for the audience, targets relevant levels, satisfying curricula requirements where needed and allows humans in the loop to provide governance and encouragement. All materials produced are subject to human teacher approval before being passed into appropriate formats on the Community Campus. This inclusion of humans in the loop is a key form of governance and an effective use of people resources providing a critical guardrail and which is far more cost effective and efficient than generating course materials from scratch.

Firstly, course subject matter is selected by the Head of Education, Training & Employment Improvement (HETEI) at the HM Prison & Probation Service (HMPPS) National Community Payback Team. GenAIE then develops the course materials in text format. These are reviewed by the HETEI and added to or amended where appropriate. Following this, GenAIE develops the full course content materials including audio and video materials in multiple languages. These are then automatically

exported in a structured format to the secure Community Campus Learning Management System (LMS).

For the Probation Service, GenAIE provided a series of nine complete introductory courses for People on Probation through an internet accessible platform portal available online through the Community Campus. The primary objective of these courses is to improve the employability of individuals engaged in Unpaid Work.

As discussed, GenAIE created customised courses for Community Campus supporting people on probation. These are specifically tailored to meet the audience’s requirements. These courses have adopted an animated style and feature a diverse range of characters representing varied demographics, reflecting the diversity of today’s society. Moreover, all courses are available in 10 languages, aligning with HMPPS’s cohort. Community Campus is accessible across England and Wales. As of 31st October 2023, the application had gained over 21,000 users and accumulated over 125,000 hours of Education, Training and Employment (ETE) learning.

In addition, significant positive qualitative feedback was received from both keyworkers and people on probation, as highlighted examples:

“Community Campus has been a real benefit to Probation, it is something that’s been required for a long time, it has taken us away from the fluffy courses that we didn’t have any control of”.

Community Payback Operational Manager.

“I have always wanted to do a plumbing course, and this has given me the start”.

Person on Probation regarding Introduction to Plumbing course.

Besides the training material, each course video integrates quizzes to encourage better learning outcomes. GenAIE has significantly reduced the time and effort required for resource creation, leading to significant cost savings while effectively addressing the personalised needs of UK HMPPS. Following the success of GenAIE in supporting people on probation, HMPPS has adopted the GenAIE platform to support people in prison.

The graphs below show the level of engagement received since the launch in December 2022. UK HMPPS has seen exponentially increasing results in terms of active users and total learning hours. As of October 2023, GenAIE reached a milestone of 125,454 hours across all courses.

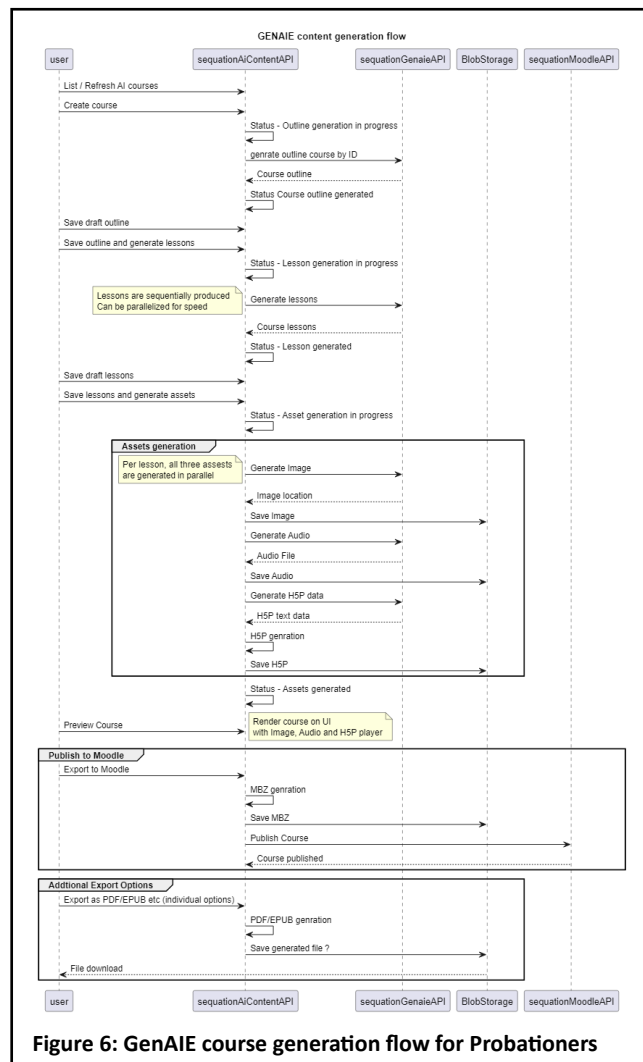


Figure 6: GenAIE course generation flow for Probationers

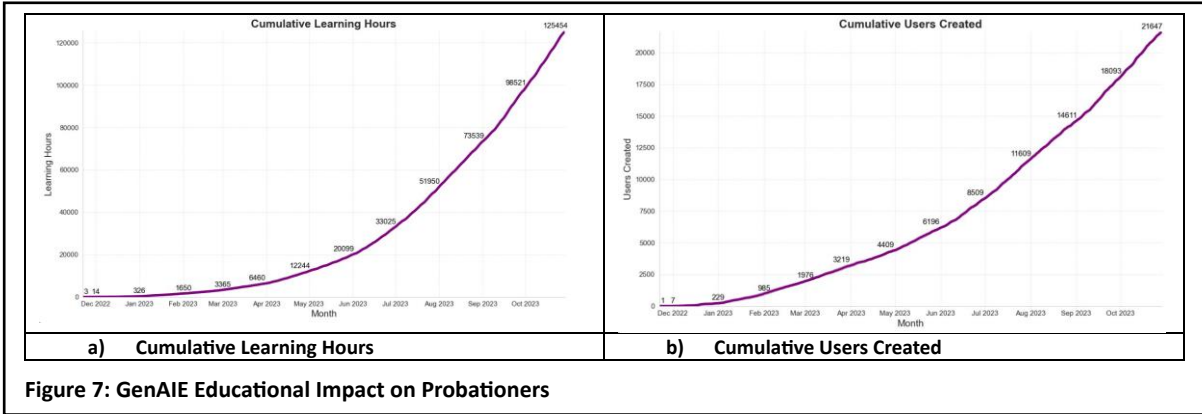


Figure 7: GenAIE Educational Impact on Probationers

8. Conclusions

AI for Social Good is a major social development in personalised education/training for disadvantaged groups of people (including people on probation, people in prison, refugees, long-term unemployed, vulnerable adults and young people with limited opportunities) thereby transforming their futures and addressing major social problems. It offers a pioneering opportunity to help socially excluded and disadvantaged people through artificial intelligence technology. It is already operating across the Probation Service, prisons, schools, Further Education colleges and refugee centres. As highlighted, the cost to society of reoffending in the UK alone is estimated at £18bn (\$23bn) per annum, with one of the important recidivism influences being lack of education and training (UK Parliament, 2022).

GenAIE offers the opportunity to positively impact the reoffending cycle at low cost and further to prevent the origins of offending behaviour in the first instance.

We seek partnerships with professionals in other areas including local authority Social Services staff supporting carers, and Jobcentre staff seeking to tackle long-term unemployment. For the future, we would like to explore lifestyle enhancement for the aged and care home residents, including training support for care home and at home keyworkers in this area.

9. Acknowledgements

We would like to acknowledge the notable contributions of Paul Henderson, Melanie Robinson, Oliver Blaine, Prasad Kelkar, Andy Smith, Emily Allred, Lewis Griffiths, and John Keane, who have led the development of the GenAIE system and given valuable feedback.

10. Further Reading

(ChatGPT, 2023) Introducing ChatGPT, OpenAI, <https://openai.com/blog/chatgpt>

(Chaudhary, 2023) Chaudhary, S., Medium, <https://medium.com/@bedigisure/generative-ai-tools-8315caa25694>.

(Dwivedi et al, 2023) Dwivedi, Y., et al. "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy, International Journal of Information Management, Volume 71, August 2023, 102642, <https://pdf.sciencedirectassets.com/271677/1-s2.0-S0268401223X00036/1-s2.0-S0268401223000233/main.pdf?>

(GAO, 2023) Science & Tech Spotlight: Generative AI, U.S. Government Accountability Office, <https://www.gao.gov/products/gao-23-106782>

(Ji et al, 2022) Ji, Z., Lee, N., Frieske, R, Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Da, W., Madatto, A., Fung, P., Survey of Hallucination in Natural Language Generation. <https://arxiv.org/pdf/2202.03629.pdf>

(LinkedIn, 2023) Generative AI, LinkedIn, https://www.linkedin.com/posts/genai-works_120-mind-blowing-ai-tools-for-video-productivity-activity-7087320504851693568-y7L6/

(NVIDIA, 2023) What is Generative AI?, NVIDIA, <https://www.nvidia.com/en-us/glossary/data-science/generative-ai/>

(Qian, 2023) Qian, L., Review of 40+ products using GPT-3, LinkedIn, www.linkedin.com/pulse/review-40-products-using-gpt-3-liang-qian/

(RAEng, 2016) The UK STEM Education Landscape, Royal Academy of Engineering, https://raeng.org.uk/media/bcbf2kyb/112408-raoe-uk-stem-education-landscape_final_lowres.pdf

(UK Parliament. 2022) Reoffending Costs, UK Parliament, <https://questions-statements.parliament.uk/written-questions/detail/2022-03-09/137323>