AI-Powered Learning: Personalizing Education for each Student

Flora Amato[1], Antonio Galli[1], Michela Gravina[1], Lidia Marassi[1], Stefano Marrone[1] and Carlo Sansone[1,*]

[1]Dipartimento di Ingegneria Elettrica e delle Tecnologie dell'Informazione (DIETI), University of Naples Federico II, Via Claudio,21, 80125, Naples, Italy

Abstract

Artificial Intelligence (AI) has the potential to enhance the traditional educational approach through the use of E-Learning and Massive Open Online Courses (MOOCs). Indeed, the application of AI-based techniques to MOOCs opens to providing a wide range of high-quality courses to a large global audience, increasing the accessibility of education and improving the learning process's effectiveness. At the same time, the growing popularity of MOOCs highlights the need to carefully consider AI use and potential negative consequences, prioritizing ethical considerations when developing and implementing this kind of technology. In this paper, we describe the use of AI techniques, specifically Deep Learning (DL), discussing the advantages, problems, and ethical concerns associated with using generative models in education. We present a case study implemented at the University of Naples Federico II that utilizes Deep Fakes to generate MOOC lessons. Furthermore, we report details about the Human-Centred AI Master's Programme (HCAIM), which supports the legal, regulatory-compliant, and ethical adoption of AI. We also introduce the use of ChatGPT, an AI-based teaching support tool, and discuss its benefits and potential risks. Overall, this paper emphasizes the importance of considering ethical concerns when implementing AI in education and highlights the potential benefits and challenges associated with classical as well as AI-based MOOCs.

Keywords

Massive Open Online Courses (MOOCs), Deep Fake, E-Learning, ChatGPT, Human-Centred AI

1. Introduction

Artificial Intelligence (AI) has the potential to revolutionize education, particularly through online learning platforms like Massive Open Online Courses (MOOCs), which provide unlimited access to educational materials. AI can enhance the quality, accessibility, and effectiveness of education by personalizing learning experiences, facilitating automatic grading, and enabling intelligent tutoring systems. In particular, one of the significant advantages of using AI in education is the ability to tailor learning experiences to different user typologies. A notable example is LessonAble, a new pipelined methodology introduced in a recent study that leverages the concept of deep fakes to generate MOOC visual content directly from a lesson script. The goal is to create a tool that relieves lecturers of the burden of writing and recording lectures while generating high-quality metadata to support impaired students.

Ital-IA 2023: 3rd National Conference on Artificial Intelligence, organized by CINI, May 29–31, 2023, Pisa, Italy

*Corresponding author.

flora.amato@unina.it (F. Amato); antonio.galli@unina.it (A. Galli); michela.gravina@unina.it (M. Gravina); lidia.marassi@unina.it (L. Marassi); stefano.marrone@unina.it (S. Marrone); carlosan@unina.it (C. Sansone)

0000-0002-5128-5558 (F. Amato); 0000-0001-9911-1517 (A. Galli); 0000-0001-5033-9617 (M. Gravina); 0000-0001-6852-0377 (S. Marrone); 0000-0002-8176-6950 (C. Sansone)

However, it is crucial to consider ethical concerns when developing and implementing generative technologies in the context of learning. Indeed, while tools such as ChatGPT may represent a new advance, they can also pose ethical difficulties, such as perpetuating bias and inequality, privacy concerns, and the risk of fake news or manipulated opinions. The cybersecurity industry should also consider the consequences of a cyber attack on these models. It is important for designers and society as a whole to reflect on the ethical implications of these tools and to develop ethically responsible and sustainable models that take into account issues of equity, privacy, and accountability. This requires a multidisciplinary approach and a thorough assessment of social and ethical impacts.

The Human-Centred AI Master's Programme (HCAIM) is a project that aims to cope with this wide range of needs and issues by supporting the legal, regulatory-compliant, and ethical adoption of AI in education and other fields. It started in 2021 and is being conducted by four universities, three research institutes, and three SMEs from five European countries, with the support of the European Platform for Digital Skills and Jobs. The project seeks to promote the development of AI technologies that can enhance the quality, accessibility, and effectiveness of education by personalizing learning experiences and facilitating automatic grading. However, the project also acknowledges the ethical concerns that must be considered when developing and implementing generative tech-

nologies in education, such as the risk of perpetuating bias and inequality, privacy concerns, and the risk of fake news or manipulated opinions. The HCAIM project emphasizes the importance of a multidisciplinary approach and a thorough assessment of social and ethical impacts to ensure the development of ethically responsible and sustainable AI models that take into account issues of equity, privacy, and accountability.

In this paper, we report and deepen all these topics, highlighting for each the central role that the Universities can, should and must have.

2. LessonAble: Leveraging Deep Fakes in MOOC Content Creation

In recent years, Massive Open Online Courses (MOOCs) have spread exponentially, with their global market size expected to grow from USD 3.9 billion in 2018 to USD 20.8 billion by 2023[1], at a Compound Annual Growth Rate (CAGR) of 40.1% during the forecast period. This success is mostly due to the wide range of benefits they offer, such as the prospect to rethink course content based on analytics and the opportunity to provide different course experiences through A/B testing to know which educational experience is more effective. MOOCs [1] are modern online courses for many participants at the same time ("massive"), without access restrictions ("open"), and in a course format (with video lectures and integrated tests). As a consequence, the students have the comfort of studying from home, self-paced learning, and much more.

An educational institution creates the learning content of a MOOC to teach and train students and experts. However, creating the MOOC content, i.e. a video lesson, often implies following a script (a text) already defined, with the author required to interpret it following every line of the defined text instead of recording a lesson on a wimp (as during a classical frontal lecture) to generate high-quality metadata (e.g. dubs) to support impaired students. Thus, despite this need tends to be extremely time-consuming, it is often a mandatory fair requirement. To take the best from this need, in [2] we introduced LessonAble, a pipelined methodology leveraging the concept of Deep Fakes for generating MOOC (Massive Online Open Course) visual contents directly from that lesson script. The idea is to realise a tool that supports content generation by relieving the lecturer of the duty of both writing the lesson script and recording its lecture, automatically generating the latter from the former. To achieve this, the proposed pipeline consists of three main modules: the audio generation module, the video generation module and a lip-syncing step.

The proposed methodology is a pipelined architecture summarized in Figure 1 and designed to create, by sequential steps, the video lesson generating the voice, the video, and lip-syncing the latter on the former. In particular, the implemented work is mainly focused on making the modules independent from each other, giving the possibility to replace or add new modules without the need for rewriting the whole architecture.

The *Voice Generation* module aims to generate an audio file (e.g., in .wav format) reproducing the user's voice from a written text. To create a module able to generate audio in a target subject's voice, it is needed to collect a dataset consisting of audio samples belonging to the target subject associated with their transcription. Besides the particular technology used to generate the voice, during the design of LessonAble we observed that the dataset collection stage can strongly affect the final quality. Indeed, although the formatting step often depends on the used text-to-speech technique, the dataset should always respect some characteristics, including both short and long voice clips having tone and pitch differences, avoiding wrong or broken files or background noise. We leveraged the NVIDIA Tacotron 2 [3] as Text-To-Speech (TTS) model to generate audios of the target subject, as in our preliminary analysis it outperformed all the other TTS systems, such as concatenate and parametric baseline ones. NVIDIA recommends using 16-bit audio with a sampling rate of 22050 Hz. Indeed, this bit-depth provides a good signal-to-noise ratio and the sampling rate provides a good inference-speed-to-audio-quality ratio because it covers most of the human voice frequency range (80Hz to 16kHz).

In the *Video Generation* module we used the First Order Motion Model (FOMM) [4], a deep learning-based approach able to generate a deep fake video by animating a target subject image by using a driving video sequence. Despite other approaches being available, such as X2Face [5] and Monkey-Net [6], we decided to use FOMM as it proved to obtain more realistic and versatile videos. To generate videos with a desired emphasis, LessonAble is able to choose the most suited target image and driving sequence, by extracting the metadata from the audio file. Thanks to this configuration, the algorithm is able to figure out which expression should be used in every specific part of the lesson, allowing the generation of more natural videos, as the expression, and movements change in different frames.

The *Lip-syncing* module represents the final phase of the LessonAble pipeline, where we used an unconstrained method that is independent of training data and thus can be applied to generic videos and audios. In particular, we exploit the Wav2Lip library [7], a tool addressing the problem of lip-syncing by talking face video of an

[1] https://www.marketsandmarkets.com/Market-Reports/massive-open-online-course-market-237288995.html

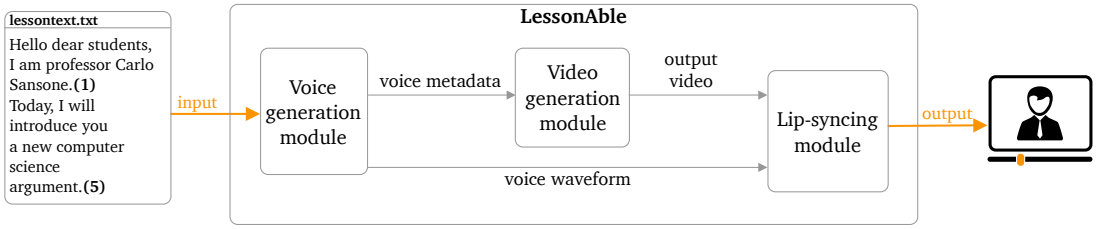

Figure 1: The LessonAble pipeline consisting of three main modules: the *Voice generation*, the *Video generation* and the *Lip-syncing*. On the left, a lesson script is used as input to the voice generation module. The voice module generates both a voice waveform and a voice metadata file containing the duration of each synthesized sentence and the markdown associated to it, to serve as input to the video generation module. Finally, the lip-syncing module synchronises the lips in the generated video, based on the generated voice, producing the final lesson.

arbitrary identity and matching a target speech segment in a dynamic manner. Designed by the same authors of LipGan [8], it makes a step forward toward high-accuracy lip-syncing by using a more accurate discriminator and other smart tricks to make the synchronized video more natural and realistic.

As a case study, we generated a deep fake MOOC lesson for professor Carlo Sansone, from the University of Naples, Federico II, testing the generation process in two languages, English and Italian. An example of Lessonable output is available at this link. We believe our efforts and our ideas in this approach can lead to new directions, such as the possibility of creating real fake MOOC content hubs. Furthermore, updating courses will be much easier as authors will only need to edit the text, while the video will be generated automatically. Future research should focus on the application of expressions to the generated audio, providing the possibility to add slides during the lesson. The result will be an extremely intuitive tool to support MOOC content generation leveraging for a good purpose Deep Fakes, a technology often associated with AI misuses, such as fake news, hoaxes and scams. Thus, a side effect of this work is also to further highlight that AI is only a (very powerful) tool, which misuse is to be associated with the user and not with the tool itself.

3. How can Chatgpt help students?

Nowadays, several companies have released free or open-source demos of their language models, which are attracting a great deal of interest. ChatGPT is one of the first examples of these chat-based generative models, whose original responses to each user input and ability to adapt to a wide range of topics make it suitable for performing different tasks. As a result, ChatGPT has quickly become popular among students as a valuable tool to improve their academic performance and save time in studying, and many of them use ChatGPT to generate topic ideas, outlines and even complete drafts of essays. Although questions are being raised about the possibility that its use may become uncontrolled, becoming a ploy to make an algorithm do the work of students, it should not be forgotten that any technology in itself is never harmful. However, the recognised capabilities of this tool make it necessary to reflect on the possible negative consequences that ChatGPT might entail. From an ethical point of view, while misuse of this language model should certainly be discouraged, it might be good to understand how instead this tool might help if it were included in a controlled way within educational institutions. ChatGPT may in fact prove to be a valuable resource for both students and faculty. In Naples (Italy) a collaboration with the Department of Electrical Engineering and Computer Science (DIETI) at the Department of Social Sciences at the Federico II University has been initiated, with the main objective of implementing digital solutions to support teaching in the humanities area. Specifically, the University of Naples Federico II intends to use ChatGPT in a controlled manner within its degree programs. In particular, the University is using this model contextually in the course Social Epidemiology, Algorithms and Big Data (included in the curriculum of the bachelor's degree in public, Social and Political Communication). ChatGPT is currently being used as a teaching support for faculty members, with a special interest in the activities of:

- Learning support, to help students understand concepts covered in lectures or clarify any doubts they may have about specific topics;
- Teacher training, to train teachers on best teaching practices, pedagogy, and how to develop effective teaching materials;
- Student feedback, to collect anonymous and rapid feedback from students.

Controlled use of this digital tool can help teachers improve their approach and create a better learning envi-

ronment for students. ChatGPT should also be used for providing information and guidance to students on study programs, course requirements, available services, and other useful information as if it were an extension of the teaching secretariat. And if these tools can be useful to students, they in turn can also help improve and analyse the performance of AI models. Currently, the University of Naples is also considering including Federico II students in a social experiment to assess ChatGPT's learning behaviour/ability with respect to constant interaction with a specific group of users. The general idea would be to proceed, after a constant period of interaction with the algorithm, to take in targeted information to evaluate its consequences. The dialogue between the areas of computer science and the humanities will thus be handled both from a practical point of view, through the implementation of digital solutions to support the Department of Social Sciences, and by providing DIETI with a concrete opportunity to analyse the possible ethical consequences that the use of these solutions entails. As the technology continues to develop, it is likely that we will see more universities adopting ChatGPT for a variety of applications. The role of academic institutions, in the development and dissemination of models, should not be underestimated. Indeed, the main reason for the rise in popularity of chat-themed generative models is their ease of use and accessibility, which is a concern with respect to the risk of people adapting to this AI without thinking through the possible negative consequences. What seems necessary, then, is for society to be pushed toward the prudent use of AI through increased awareness of new technologies.

3.1. Ethical issues related to generative models' use

While this tool may represent a new advance, the worst risk is that it is used without thinking through the possible ethical difficulties that this AI, like all generative models, might pose. Generative models are trained on data, which means they can easily perpetuate bias and inequality if present in the training set. The use of generative models can also pose concerns about intellectual property, privacy, and the risk of fake news or manipulated opinions. The cybersecurity industry should also consider the possible consequences of a cyber attack on these models. Indeed, ChatGPT itself could quickly generate targeted phishing emails or malicious code for malware attacks. The potential impact of generative models on society also raises several questions of accountability, as careless use of these models can have unforeseen effects. As is often the case with technologies, ethical responsibility seems to be shared. It is therefore important that not only designers of language models work responsibly and consider ethical issues in the development and use of models, but also that society itself reflects on the changes that these tools might produce. It cannot be overlooked that these systems deal with human language processing, and thus directly affect the ethical dimension as well. It is important to develop ethically responsible and sustainable generative models that take into account issues of equity, privacy, and accountability. It is equally essential that there be more transparency about the use of these models and their implications. This requires a multidisciplinary approach to the design and implementation of generative models, involving a variety of stakeholders and a thorough assessment of social and ethical impacts. Technological development also implies social responsibility, so that the potential of this technology does not become a disadvantage instead.

4. HCAIM

The Human-Centered AI Master's Programme (HCAIM) is an interdisciplinary program that provides students with a comprehensive understanding of designing, developing, and evaluating human-centred AI systems [9]. The programme's curriculum is designed to equip students with the necessary skills and knowledge to work in various domains, including healthcare, education, and finance, where AI has a significant impact. The program is a joint initiative of four prestigious European universities, namely Budapest University of Technology and Economics, HU University of Applied Sciences Utrecht, Technological University Dublin, and University of Naples Federico II. The consortium also comprises prominent experts in computer science, social sciences, and humanities who are dedicated to providing students with a comprehensive and interdisciplinary education in human-centred AI. In particular, HCAIM is backed by a network of industry partners and stakeholders who provide real-world case studies and research opportunities to enrich students' learning experiences. The program partners include companies (Real AI, Nathean Technologies Ltd, Fiven) and research institutions (National Research Council of Italy - CNR, Ireland's Centre for Applied AI - CeADAR, European Software Institute Center Eastern Europe - ESI CEE), all sharing a common goal of developing AI systems that are centred on human needs. Additionally, HCAIM benefits from the expertise of a diverse group of international scholars and researchers who contribute to the teaching and research activities of the program. The HCAIM programme is supported by the European Platform for Digital Skills and Jobs (INEA/CEF/ICT/A2020/2267304), which is focused on promoting digital skills and expertise in Europe.

4.1. Mission and Aim

HCAIM's objective is to cultivate a new generation of AI experts who prioritize the human perspective when designing, developing, and implementing AI systems. The program aims to impart students with the expertise and knowledge necessary to create AI solutions that are socially responsible, technically sound, and ethically acceptable. HCAIM is also committed to producing innovative educational materials that will be used consistently at partner institutions, fostering the exchange of project ideas and students among participating countries. Furthermore, the program offers short-term Erasmus actions that facilitate the exchange of knowledge and ideas among students.

The program is designed to promote interdisciplinary collaboration among computer science, social sciences, and humanities, equipping students with a comprehensive skill set that can be applied across various domains and contexts. It is this collaborative and interdisciplinary approach that makes HCAIM stand out in the field of AI education. Although the precise programs vary to some extent across participating universities based on their respective traditions, resources, and legal and cultural environments, the diverse array of elective courses available enables students to build their own personalized HCAIM portfolio, based on their interests in different domains. Nonetheless, the program maintains a consistent and unified curriculum across all partner institutes. This is achieved through the collaboration of the HCAIM consortium, which has developed all courses and lectures in a collective effort to provide students with a comprehensive and well-rounded education. Through this collective approach, HCAIM students will gain a shared understanding of the key concepts, principles, and techniques related to the development and implementation of human-centred AI systems. This ensures that students will acquire a consistent and cohesive foundation of knowledge, regardless of the university they attend.

4.2. Structure

The HCAIM program is a flexible and adaptable master's degree offered by four European universities, each with their own unique approach to delivering the program. The duration of the program and the number of credits required to complete it depend on the delivering university, ranging from one year (60 ECTS) to two years (120 ECTS). Some universities are offering HCAIM as a stand-alone program, while others are integrating it into their existing programs. The HCAIM program is delivered entirely in English and comprises a combination of lectures, seminars, hands-on projects, and research activities. The program is designed to equip students with the necessary skills and knowledge to design, develop, and evaluate human-centred AI systems. It is an interdisciplinary program that fosters collaboration between computer science, social sciences, and humanities.

Students are required to complete a master's thesis as part of the program, providing them with an opportunity to apply the skills and knowledge acquired during the program to a real-world problem. The thesis is an important part of the program, allowing students to demonstrate their ability to apply critical thinking, problem-solving, and research skills. The HCAIM program is supported by a network of industry partners and stakeholders who provide real-world case studies and research opportunities for students. This collaboration between academia and industry ensures that students are exposed to the latest developments and trends in the field of AI.

5. Conclusion

In conclusion, the use of AI in education has the potential to revolutionize traditional educational approaches and improve accessibility, effectiveness, and personalization of learning experiences. In this paper we highlighted the importance of developing ethically responsible and sustainable AI models that take into account issues of equity, privacy, and accountability, and the central role that universities can, should, and must have in this process.

References

[1] J. Reich, Rebooting mooc research, Science 347 (2015) 34–35. URL: https://www.science.org/doi/abs/10.1126/science.1261627. doi:10.1126/science.1261627.

[2] C. Sannino, M. Gravina, S. Marrone, G. Fiameni, C. Sansone, Lessonable: Leveraging deep fakes in mooc content creation, in: Image Analysis and Processing–ICIAP 2022: 21st International Conference, Lecce, Italy, May 23–27, 2022, Proceedings, Part I, Springer, 2022, pp. 27–37.

[3] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. J. Skerry-Ryan, R. A. Saurous, Y. Agiomyrgiannakis, Y. Wu, Natural TTS synthesis by conditioning wavenet on mel spectrogram predictions, CoRR abs/1712.05884 (2017). URL: http://arxiv.org/abs/1712.05884. arXiv:1712.05884.

[4] A. Siarohin, S. Lathuilière, S. Tulyakov, E. Ricci, N. Sebe, First order motion model for image animation, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 32, Curran Associates, Inc., 2019.

URL: https://proceedings.neurips.cc/paper/2019/file/31c0b36aef265d9221af80872ceb62f9-Paper.pdf.

[5] O. Wiles, A. S. Koepke, A. Zisserman, X2face: A network for controlling face generation by using images, audio, and pose codes, CoRR abs/1807.10550 (2018). URL: http://arxiv.org/abs/1807.10550. arXiv:1807.10550.

[6] A. Siarohin, S. Lathuilière, S. Tulyakov, E. Ricci, N. Sebe, Animating arbitrary objects via deep motion transfer, CoRR abs/1812.08861 (2018). URL: http://arxiv.org/abs/1812.08861. arXiv:1812.08861.

[7] K. R. Prajwal, R. Mukhopadhyay, V. Namboodiri, C. V. Jawahar, A lip sync expert is all you need for speech to lip generation in the wild, CoRR abs/2008.10010 (2020). URL: https://arxiv.org/abs/2008.10010. arXiv:2008.10010.

[8] K. R. Prajwal, R. Mukhopadhyay, J. Philip, A. Jha, V. Namboodiri, C. V. Jawahar, Towards automatic face-to-face translation, CoRR abs/2003.00418 (2020). URL: https://arxiv.org/abs/2003.00418. arXiv:2003.00418.

[9] B. Feeney, M. Zuccarini, T. Singh, H. Aldewereld, S. Marrone, K. Quille, Developing a human centred ai masters: the good, the bad and the ugly, in: Proceedings of the 27th ACM Conference on on Innovation and Technology in Computer Science Education Vol. 2, 2022, pp. 660–661.