AI as yet another tool in undergraduate student projects: preliminary results

Iván J. Pérez-Colado
Departamento de Ingeniería del Software e Inteligencia Artificial
U. Complutense de Madrid
Madrid, Spain
ivanjper@ucm.es

Manuel Freire-Morán
Departamento de Ingeniería del Software e Inteligencia Artificial
U. Complutense de Madrid
Madrid, Spain
manuel.freire@ucm.es

Antonio Calvo-Morata
Departamento de Ingeniería del Software e Inteligencia Artificial
U. Complutense de Madrid
Madrid, Spain
toni@ucm.es

Víctor M. Pérez-Colado
Nord University
Bodø, Norway
victor.m.colado@nord.no

Baltasar Fernández-Manjón
Departamento de Ingeniería del Software e Inteligencia Artificial
U. Complutense de Madrid
Madrid, Spain
balta@fdi.ucm.es

***Abstract*—How do students use artificial intelligence tools in coursework projects when given the liberty to do so, with the only requirement of documenting how, where and why? We describe experiences with two groups of undergraduates in courses related to serious game authoring and human-computer interaction, both carried out in the second semester of 2023. In the serious games course, students were given the option of following a teacher-developed methodology for generating graphical assets for their serious games using a set of generative AI tools. This methodology was explained in the class but not hands on lab was carried out. In the interaction course, students were free to choose which AI tools to use when designing their system or in the development of their project documentation. Despite the limited number of participants (41 in total) we can see very different views and degrees of involvement: while some tried to use AI for as many tasks as possible, others considered that the learning curve for those tools was too steep to be worthwhile. Both experiences included a free-text survey at the end, and taken together, provide insights into how both supervised and unsupervised generative AI use could impact undergraduate projects in similar subjects. In addition to describing how students chose to use the tools, and the main takeaways from their survey response, we also discuss some of the ethical aspects about the access to the tools and what should be the minimal conditions to be met to allow the equitable use of AI in the classroom.**

Keywords—AI in education, generative artificial intelligence, game development, serious games authoring, goal-driven design

I. Introduction

Since the late 2022 launch of OpenAI's ChatGPT, generative AI has been a source of strong emotions in academia. On one hand, some believe that use of AI should be strictly forbidden, for example due to the ease with which freely available tools can write plausible essays given minimal prompting; while others advocate for greater adoption, given both the potential advantages of generative AI and the difficulties in detecting and restricting its use. Both groups are very concerned about its impact on the evaluation of students and the possible adjustments that would have to be made to long-established practices.

In this paper, we report on the results of envisioning AI as just another tool that students can choose to use in creative undergraduate projects. Students from two courses were given the option of using generative AI in their projects, with the sole requirement of documenting that use and answering a set of questions on the degree to which they used it, together with the purpose of such use and its on their work.

The next section explores related work and introduces the types of generative AI that participating students reported using, and, in the case of the course on serious games, the AI-assisted asset-creation methodology suggested to participants. Section III describes the two experiences and Section IV their results, while Section V contains a discussion and future work.

II. Background and Related Work

Generative AI uses trained models to generate new synthetic content, typically in response to a textual prompt. OpenAI's ChatGPT relies on a generative pre-trained transformer (GPT) architecture, trained with huge datasets from different sources such as the Wikipedia, collections of digitized books, or web crawls [1]. Since the appearance of ChatGPT as an online service, other systems based on the same principles have emerged [2], such as Google's Bard, Microsoft's Copilot AI (which is designed to interoperate with Microsoft's desktop productivity applications [3]), X's Grok, Facebook's Llama, or Amazon's Q. Collectively, we will refer to text-generating AI as large language models, or LLMs.

Generative AI are widening their capacities, from text generation to other areas such as code, graphics, animations and

Partially funded by the Ministry of Education (PID2020-119620RB-I00, SNOLA Network RED2022-134284-T), by the European Commission (Gaming4Coding Erasmus+ project 2021-1-SE01-KA220-SCH-000023932), and by the Telefónica-Complutense Chair on Digital Education and Serious Games

video generation. This is expected to have a deep impact in many industries, such as videogames, that traditionally relied on experts to create such materials. For image generation, OpenAI released DALL-E in 2021, which could generate images from textual descriptions after training on a large set of image-text pairs (400M, in the initial release of DALL-E [4]). Improved systems have followed, including significant updates to DALL-E, StabilityAI's Stable Diffusion (SD), Google's Imagen, or Midjourney – with the market for generative AI art estimated at 48B USD in 2023, to the alarm of many artists [5]. A key contention is the problem of training data: the above-listed generative AIs were all trained with image-text pairs where most images had a questionable copyright status. In comparison, attempts to create copyright-free generators are only in their infancy [6], [7].

There is intense competition to fill in gaps in the functionality of generative AI. Initial versions of LLMs were bad at math and coding, but both aspects have received significant attention. Since LLMs tend to appear very confident in their output, despite often making it up (sometimes referred to as "hallucinations" [8]), efforts have been made to introduce a measure of external checking in the tools themselves. In the case of image-generation, newer versions are much better at preserving facial symmetry and hand anatomy, to cite two problems that were often found in earlier versions. There is also a vibrant community developing add-on functionality, for example to cater to specific art-styles, specific needs such as posing or adding facial expressions, or even the generation of artificially generated video. Beyond text and images, there are also several systems that can recognize and/or generate audio.

According to industry associations, most adults in affluent regions such as the EU or the USA play videogames, which have significantly surpassed movies in terms of economic weight. Given the importance of this industry, popular videogames tend to be visually stunning, with elaborate and consistent graphics and music. Creating these artistic assets is a major driver of cost in videogame creation, both in terms of time and expertise, and top-tier results are completely out of reach for developers without the necessary knowledge or budgets. High expectations are a barrier in the development of serious games by small teams that are still learning how to use the necessary tools. Simplifying and automating art asset creation tasks with generative AI can allow teams to use their resources more effectively and explore a greater range of ideas, especially in the initial, prototype stages of game design.

A. *Generative AI to create art assets for serious games*

In a previous study from the authors [6], we tested a generative AI methodology to create graphical assets for serious games, and specifically to generate both background and character assets. Our goal was to simplify the asset creation process while maximizing the author's control of the output, guaranteeing both quality and consistency when producing assets in the context of a project. Consistency between generations is important to avoid breaking immersion, and to allow incremental changes as projects mature. Our toolset can be used in a controlled and iterative way to produce backgrounds and characters using the Txt2Img feature of Stable Diffusion[1] (SD). During the afore-mentioned study, local, as opposed to internet-hosted versions of SD were used. We used an automated installer[2], which results in a self-hosted, web-accessible instance of SD with advantages such as free and open access with unlimited usage (online versions generally require subscriptions for heavy use), privacy for participants, and advanced fine-tuning options. Some of these fine-tuning tools were key parts of the methodology, including LoRA (Low-Rank Adaptation) models, used to focus on a specific art style; ControlNet[3][9] to provide additional constraints, such as poses; or OpenPose[4], to author pose information.

In addition to asset creation, the methodology had three subsections for modifications: adding details to already-created backgrounds, adding new character poses, and adding new character facial expressions. Modifications relied on combining the Img2Img feature with additional restrictions, including references to the character sheet and pose, with additional parameters when authoring new character variations. Taking together asset-creation and modifications, our methodology greatly reduced the time required to produce assets during the creation of a serious game as a Master of Design thesis project: a total of 20 scenarios, 5 characters and 8 cutscenes were created in a total of 72h of work. According to self-estimations, the production time was reduced by 63%, and time savings could potentially increase to 86% once the process is mastered.

As part of a post-analysis of the results prior to the present study, we identified at least two parts of the process that could be simplified by the introduction of Stability Matrix[5] and ComfyUI[6]. StabilityMatrix is a local instance manager for SD with a UI that allows the user to easily download and install extensions such as fine-tuning models, add-ons, and new tools. This tool replaces the SD WebUI installer while providing new UIs that simplify the management of extensions and fine-tuning tools such as checkpoints, or LoRA, and making them easy to access across the different SD-related tools. ComfyUI is a node-based editor for SD that allows users to create complex SD workflows without programming. It provides improved visual understanding of the creation flows, and it even includes a system to export and import templates, which can later be used to greatly reduce asset creation times, and lower configuration errors while retaining a large degree of control. However, while ComfyUI has several advantages in the background creation process, the de-facto Web UI is simple enough for new users to work with for new asset generation. Because it exposes users to the full parameter flow, usage of ComfyUI can be overwhelming to novices, and we thus relegate it to character creation and later fine-tuning.

[1] Stable Diffusion: github.com/CompVis/stable-diffusion
[2] SD WebUI: github.com/AUTOMATIC1111/stable-diffusion-webui
[3] ControlNet Plugin for SD: github.com/Mikubill/sd-webui-controlnet
[4] OpenPose Editor: github.com/fkunn1326/openpose-editor
[5] StabilityMatrix: github.com/LykosAI/StabilityMatrix
[6] ComfyUI: github.com/comfyanonymous/ComfyUI

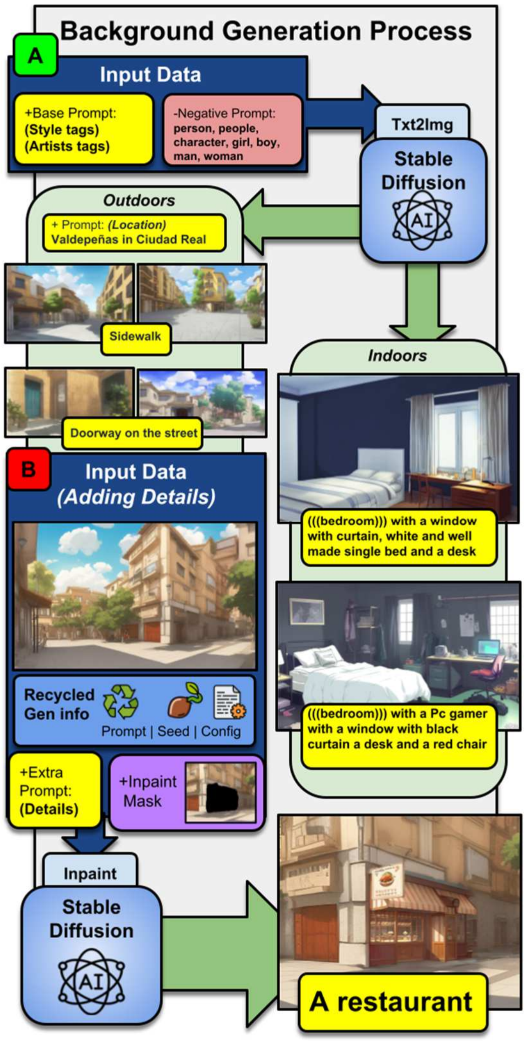

Fig. 1. Description of the scenario generation process, showing: A) background generation and B) fine details manipulation. From [10].

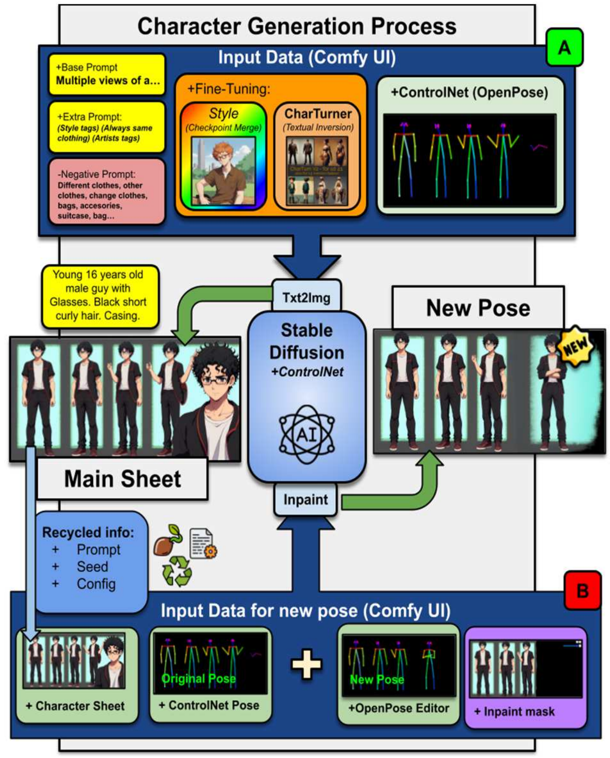

Fig. 2. The AI-assisted character generation process used in the Serious Game experience. Several tools are combined to generate initial character sheets (A) and later fine-tune them to add new poses required by the game script (B) and/or facial emotions.

III. Materials and Methods

We carried out two experiences with university students. Both experiences were carried out by the authors in subjects from two different degrees at Computer Science School at the Complutense University in Madrid, Spain, in the second half of the 2023 academic year. Students were free to use (or avoid using) any AI tool when working on their subject final projects, which were mandatory and worth most of the grade for their respective subjects. The main difference between the design of both experiences is that, in the serious games course, students were provided with guidance on how to use a specific set of tools, while those in HCI course were not provided any specific guidance. As to the participants, it must be noted that those in the serious game experience are much more art-conscious than their HCI peers, because they have several fine arts and even sound-design courses as part of their game design curriculums. This had important effects in their perception of AI.

A. *Experience 1: Serious Games*

This is an elective subject within a four-year videogame development degree, teaching students the basics of designing and creation of serious games, that is, games where the main goal is not entertainment, but for example education or increasing awareness on an issue. In this serious game course, students were introduced in class to a generative AI art-asset creation methodology previously developed by the authors (see [10]). This allowed students to see the possibilities offered by a set of popular AI tools, which can help significantly in the development of art assets for serious games. However, the students were given complete freedom to decide whether they wanted to use it, and what they would use it for. Students were also offered help on demand if they decided to apply the methodology in their projects.

Students (N=26) worked on their projects in teams of 3, with the goal of creating a playable prototype of a serious game in 2 months and a half, after being introduced to the basic theory of how to create such games. A total of 9 groups worked on the course, including one group of 2 students. They were first introduced to the technology in a 2-hour presentation, with live examples, of using the methodology described in the following subsection. They then had two weekly 2-hour sessions and were provided feedback on their effort at several game development milestones: choice of topic, educational design, general game design, and initial prototype. Finally, they had an additional month of unsupervised time to refine their serious games prototypes.

The methodology includes two main steps, which allow generation of the main assets required to build a point-and-click serious game: scenario and character generation.

1. *Scenario generation*

Point-and-click adventure games typically consist of several scenarios, essentially equivalent to scenes in a theater piece. In terms of art assets, each scene has a corresponding background, such as a classroom or a bus stop. Occasionally, items in a scene can be interactive; in this case, they need to be generated separately, so that they can be independently handled by the game (for example, highlighted, placed into the player's inventory, or added and removed from scenes). The output of this step is a set of artistically compatible backgrounds for each scene, and another set of items that can be placed in parts of specific scene backgrounds as the story requires. Fig. 1 illustrates the process of generating backgrounds with AI. Note the use of "inpaint" to specify regions to be re-drawn, for example to add or modify elements.

2. *Character generation*

Inside scenes, the game displays the actors or characters – depending on the specific game, the author may choose to either draw the player-controlled character on screen (3rd person view) or only display what this character would see (1st person view). Regardless, since the game genre relies heavily on interacting with non-player characters (NPCs), it is important to display them adequately inside each scene, and closeups of their heads when displaying dialogue, so that players can see who "says" each line and the responses of other characters, including emotional cues, to what is being said. Character generation, illustrated in Fig. 2, has 3 sub-steps: first, for each character, a concept sheet must be developed, displaying the character in a multitude of poses, such as standing up, facing in different directions, or sitting down; then, fine-tuning of poses may be required, to fit with the script; and finally, facial expressions may need to be added to some of the poses, allowing characters in the game to express emotion.

Note that generation of scenario and graphic resources is not the only field of application of generative AIs to serious games. Students were also allowed to use AI to generate supporting narrative, code and conversations for the game, if they documented such use appropriately. Additionally, a course requirement for student projects was that they were to be published in GitHub with a free license, and that any student use of AI to generate code, text or images had to be clearly described in project documentation. Finally, at the end of the semester, students had to make a presentation of the projects, which included a 2-minute video showcasing their results.

B. *Experience 2: Interactive Systems Design*

This is another elective subject within a 4-year computer science degree which teaches students the basics of designing usable systems following Alan Cooper's Goal-Directed Design (GDD) methodology [11]

This subject is taught in English and usually has over 50% international students. Conversely, this means that roughly half of the students are not international, and their English may not be as good as their Spanish. The course is mainly assessed on a final project basis with students organized into groups of 3 to 5 students each. Students were asked to conceptualize and design the entire interaction of a complex application in a domain of their interest. It is designed for mobile first and it was expected to be able to interact with other devices, such as smart watches, or voice/screen operated devices. Students had to interview potential users to elicit system requirements, model them (using "personas" as a generic user representation tool) and document the entire application of the GDD methodology in their corresponding project. It is therefore a module that requires a lot of coordination and documentation work.

In this case, students (N=22) were not given any lectures on the use of AI and were left free to use it for any task where they felt that it could either add value or avoid repetitive work.

C. *Questionnaries*

Students from both experiences had to report which tools they had used and for what purpose and how they used them in the final project documentation. To this end, they were requested to fill out a questionnaire via an online form with the following 6 questions:

1. Did you use AI tools in the project? If not, why not?
2. Do you think AI tools should be allowed to be used in other modules? Why? Give positive and negative examples of its use.
3. What have you used these AIs tools for? (*for each experience, several possible examples were provided here*) Specify what you have generated and how useful you have found the use of AI tools.
4. Do you think their use of AI has saved you work time? How much (make an estimate)?

Use of this questionnaire was intended to better understand the students' perception of these technologies and try to understand the main reasons that lead to the use or non-use of AI, such as its perceived benefits and drawbacks, which could also have implications in extending their use to future editions of the course, or different courses in the same School.

IV. RESULTS

A. *Experience 1: Serious Games*

At the end of the course, 9 games of very diverse genres and themes were presented. These results are described below:

A orillas del Duero is a single player resource management game, where players take the role of a farmer living in a village on the verge of abandonment. The game aims to raise awareness about the concept of a hollowed-out Spain and the abandonment of villages due to lack of resources. This group did not use any AI in their game.

Mind in Conflict, a virtual reality game in which the player puts himself in the shoes of a person with phonophobia and agoraphobia. The game aims to raise awareness of the anxiety that people with these disorders can suffer. This project, besides being in 3D, has no text and is focused on the experience through sound and immersion. The group generated the game's logo and title using AI.

The Morality of AI is a 2D decision-making game aimed at philosophy and ethics teachers, an interactive tool for their students to reflect on the morality and ethics of some situations. The group used ChatGPT as a consulting tool.

Mined Stories is a set of mini-games and puzzles that aims to immerse the player in a journey through the various regions of Andalusia to discover, preserve and celebrate the cultures and traditions of this region. The group used AI as a consulting tool.

Amazon no name is a 3d game where the player is a bird that must find food to survive in a vast forest. The educational objective is to make players aware of the devastating effects of deforestation and its consequences on both the environment and the local fauna, killing millions of animals. This group did not use any AI.

Eco Swipe is a 2D card and decision-making game. The game is intended to help players understand the intricate relationship between ecology and economics, encourage more informed and responsible decision making in real life, and to and promote a stronger commitment to sustainability and environmental conservation. sustainability and environmental conservation. This group did not use any AI.

GeoGuide is a 3D simulation game, focused on teaching geography. The player pilots a plane transporting customers around the world to the destination they request and learns all kinds of interesting facts about the countries visited. The group did not use AI.

Change wheels is a 2D point&click game, focused on raising awareness of the challenges faced by people with disabilities in everyday environments, and highlighting situations that are made worse by indifferent attitudes that often go unnoticed. Its authors made the heaviest use of AI among all groups, focusing mainly on generating 2D images for characters and scenery (illustrated in Fig. 3). They also used it to generate certain conversations.

Fig. 3. Two screenshots from *Change wheels*, a serious game developed by students to increase awareness of wheelchair-accessibility. The game's graphical assets were generated with AI.

Vector Golf is a puzzle and mini-golf game. The main goal is to help high school students to visualize vectors, so that they can understand them both as concept and in terms of vector operations such as addition or scalar multiplication. These concepts are useful, for example, at the high school physics level. The group did not use any AI.

Regarding the students' own responses on the use of AI, 9 of the participants used it in their games, while 12 did not. The most popular uses of AI were text generation (75%), brainstorming (33.3%), coding (22.2%), English proofreading (22.2%), logo creation (22%), and character and background creation (11%). The asset creating methodology was only used by one of the participants, reportedly saving over 10 hours of time in asset creation. In general, respondents showed positive reception towards the use of AI in the class, as only 3 students did not recommend its usage.

The reasons reported by the students for not using AI (12) include personal preference (mentioned by 41.6% of respondents), that the existence of freely available assets made it unnecessary to create them (33.3%), and ethical objections regarding how those image generation AIs were trained: 25% of participants noted that there are serious copyright questions surrounding image-generation AI, and one of them went as far as to provide several references of recent copyright-infringement AI headlines.

Even though the methodology for the creation of graphic resources that was explained to the students had been improved since its first version, and perhaps because the explanation was did not include hands-on installation and configuration of the tools, the students still considered it too complex. This was noticed not only on the reasons reported in the questionnaire, but also on the problems observed when students tried to generate art using AI. In students' computers, the main process-breaking issue was the insufficient amount of video card memory in many systems, while SD required at least 16GB of video memory to work. Other minor issues included lack of local drive space, CPU performance, and platform incompatibilities.

Finally, despite the advantages in freedom and privacy for participants when using a self-hosted generative setup, setting up the necessary programs correctly is still a process-breaking barrier for many students, due to the novelty, frequent undocumented changes and low maturity of several of the packages used.

B. Experience 2: Interactive Systems Design

Students had to design systems based on a description proposed by the teacher or propose other systems of similar complexity in an area of their interest, using the Goal-Directed Design (GDD) methodology. One of the teacher-proposed systems was a management system for a shared apartment, which would include both the organization of cleaning and maintenance tasks and the economic aspects of shared expenses, including both recurrent expenses such as rent and purchases and extraordinary ones such as improvements or parties. Another system was aimed at improving health habits, considering diet, physical activity and sports (acquired through a smartwatch), rest patterns, and so on in order to make suggestions for healthier living to the user. In addition, this system allowed users to set personal challenges for improvement and provided personalized alerts and monitoring of those alarms both from the mobile and from the smartwatch. Yet another system proposed by students managed travel expenses with friends, incorporating more sophisticated aspects such as support for different currencies, interest for late payment, compensation of expenses between members or inclusion of photos of the invoices of the expenses incurred.

Designing these complex systems using the GDD methodology requires complex documentation, since it is necessary to make a general planning of the project, identify potential users to be interviewed and perform an analysis of other applications that may compete with the proposal. Transcripts of user interviews are important to identify the key design elements from the user's point of view. From there, as previously mentioned, users are modeled by means of one or more generic personas, and the conceptualization phase of the user interface itself begins. In this case, students first created individual paper-based interfaces, which were later examined together to create an initial group design. Subsequently, this initial design was refined and evolved by means of a digital prototype, which students could choose to create in either low fidelity (using Balsamiq) or in high fidelity (using Figma). Both prototyping tools, Balsamiq and Figma, had been presented and practiced with in the lab at the beginning of the course. Finally, the designed systems were evaluated heuristically and with users to obtain feedback to improve them and create the final versions. Having to document the whole process, the students are faced with managing a project with a complex final report. It is in these more repetitive aspects that we believe AI could be used to help students focus on the more creative and conceptual tasks.

Only 1 student avoided the use of AI, asserting that it was unsuitable for creative work. The remaining 21 did use it, with the most popular use being text summarization (66.6%), followed by user interview audio-to-text transcription (57.1%) and finally persona-generation and English proofreading (33.3% in both). Several students note the high quality of the speech-to-text transcriptions using freely available online tools, and also how painless it was to summarize those transcripts for inclusion into the project documentation. It is also interesting to see how many students report on improving their English by having a generative AI proofread drafts of the document, to search for better wording.

Given the copyright issues that surround generative AI, it is also interesting to see that some students report using generative AI precisely to avoid copyright issues – instead of using material found on the Internet in their interface mockups, they used AI to generate a set of icons from textual prompts. In their words, "online images have copyrights, those generated by stable diffusion don't". While many jurisdictions do not accept generated images as copyrightable, it is a misconception that they cannot infringe the copyrights of others, as copyright infringement does not depend on how the image was generated, but only on its contents; and it has been repeatedly shown that some prompts can, indeed, generate copyrighted material which was part of the training image-text pairs; several examples can be found in [6].

An additional controversial aspect is that, with the inclusion of new capabilities in generative AI systems (e.g., code generation), it seemed that they could also be used to prototype applications based on their textual description. When asked in the labs if they had tried to use it, the students stated that, at that time, it was still a limited capability and required a paid license, so they ruled out its use.

V. Discussion and Future Work

We have reported on two experiences showcasing the use of generative AI as a tool in academic projects, both in courses related to computer science. There is a notable difference between the enthusiastic use of text generation by participants in the second experience (project documentation in an HCI-related course) and the much lower adoption of participants in the first experience (a serious game for a course in videogame design). It is somewhat surprising that students in the interaction course are beginning to perceive that generative AI is an all-purpose tool that can address almost any type of problem. Not only did they use it to create drafts of text, but also to improve writing, summarizing, or transcribing audio to text. And all this instead of using existing and already proven tools (which may have already used AI techniques as well).

Analyzing student responses to the questionnaires, it becomes clear that tool usability and affordability is an important factor: while configuring image generation requires significant technical expertise and time-investment, it is by comparison very easy to generate text by asking a short question or pasting a chunk of text into a box and then adding requests such as "please summarize this". On reflection, we consider that hands-on use of generative technology would have been important to drive adoption of our AI-assisted art asset creation methodology for games for the first experience. It also raises the possibility of improving its usability by streamlining the whole methodology and wrapping it into a single standalone tool for the sole purpose of the creation of serious game assets.

A second observation is the strength of the ethical concerns expressed by game design students, which echoed the ongoing debate concerning copyright infringement in AI training – while no such concerns were expressed by the HCI group. Training in fine arts is probably an important factor for the greater awareness of such problems by the game design group. Despite the fair-use debate surrounding generative AI, such as those described in [7], this highlights the need for a trustworthy platform which students can trust as a legitimate learning tool. In this sense, efforts to create copyright-respecting generative AIs, such as [6], are promising.

A third observation from the creative perspective, is that many of the students decided to create their own art despite the opportunity to use AI instead. This hints that the low control and slow response times afforded by AI tools were, to these students, less tempting than tried and true digital art creation. These technological hurdles may be only temporary, as advancements in SD tools, such as SDXL-Turbo and LCM LoRA [12], [13], have been shown capable of producing images in nearly real time. This could allow artists to design, compose and manipulate images in a real-time iterative flow, positively impacting the artists' adoption of AI tools once they allow a more natural way of designing and creating.

Finally, further research is needed not only on the possible legal implications, but also on other ethical issues involved in the use of these AI technologies. The power and behavior of these tools is very different depending on whether you access free or paid versions, and so is the degree of privacy of the prompts. And in the case of trying to use open-source versions, they usually require much more effort on the part of the students or the results are not of such high quality. Therefore, in the medium and long term, if the use of these technologies is allowed, it is necessary to address what access to the tools and what should be the minimum conditions to be met to allow the equitable use of AI in the classroom.

References

[1] C. Raffel *et al.*, "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer," 2020. [Online]. Available: http://jmlr.org/papers/v21/20-074.html.

[2] A. Bozkurt, "Unleashing the Potential of Generative AI, Conversational Agents and Chatbots in Educational Praxis: A Systematic Review and Bibliometric Analysis of GenAI in Education," *Open Praxis*, vol. 15, no. 4, pp. 261–270, Nov. 2023, doi: 10.55982/openpraxis.15.4.609.

[3] Jared Spataro, "Introducing Microsoft 365 Copilot – your copilot for work - The Official Microsoft Blog," Official Microsoft Blog. Accessed: Jan. 10, 2024. [Online]. Available:https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/

[4] A. Radford *et al.*, "Learning Transferable Visual Models From Natural Language Supervision," *Proceedings of the 38th International Conference on Machine Learning, PMLR,* 139, 8748–8763.

[5] H. H. Jiang *et al.*, "AI Art and its Impact on Artists," in *AIES 2023 - Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, Association for Computing Machinery, Inc, Aug. 2023, pp. 363–374. doi: 10.1145/3600211.3604681.

[6] A. Gokaslan *et al.*, "CommonCanvas: An Open Diffusion Model Trained with Creative-Commons Images," *arXiv preprint* , vol. arXiv:2310.16825, Oct. 2023, [Online]. Available: http://arxiv.org/abs/2310.16825

[7] P. Samuelson, "Generative AI meets copyright: Ongoing lawsuits could affect everyone who uses generative AI," *Science (1979)*, vol. 381, no. 6654, pp. 158–161, Jul. 2023, doi: 10.1126/science.adi0656.

[8] T. R. McIntosh, T. Liu, T. Susnjak, P. Watters, A. Ng, and M. N. Halgamuge, "A Culturally Sensitive Test to Evaluate Nuanced GPT Hallucination," *IEEE Transactions on Artificial Intelligence*, 2023, doi: 10.1109/TAI.2023.3332837.

[9] L. Zhang, A. Rao, and M. Agrawala, "Adding Conditional Control to Text-to-Image Diffusion Models," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 3836–3847.

[10] I. J. P. Colado, V. M. P. Colado, A. C. Morata, R. S. C. Píriz, and B. F. Manjón, "Using New AI-Driven Techniques to Ease Serious Games Authoring," in *2023 IEEE Frontiers in Education Conference (FIE)*, IEEE, Oct. 2023, pp. 1–9. doi: 10.1109/FIE58773.2023.10343021.

[11] Alan Cooper, Robert Reimann, and David Cronin, *About Face 3:The Essentials of Interaction Design*. Wiley, 2007.

[12] A. Sauer, D. Lorenz, A. Blattmann, and R. Rombach, "Adversarial Diffusion Distillation," Nov. 2023, Accessed: Jan. 12, 2024. [Online]. Available: https://arxiv.org/abs/2311.17042v1

[13] S. Luo, Y. Tan, L. Huang, J. Li, and H. Zhao, "Latent Consistency Models: Synthesizing High-Resolution Images with Few-Step Inference," Oct. 2023, Accessed: Jan. 10, 2024. [Online]. Available: https://arxiv.org/abs/2310.04378v1