RECOMMENDATIONS FOR UVA COMPUTER SCIENCE CURRICULUM TO CREATE INDUSTRY READY GRADUATES

ATTITUDES TOWARDS AI GENERATED ART

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Andrea Jerausek

November 1, 2022

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Rider Foley, Department of Engineering and Society

Briana Morrison, Computer Science

Increasing the Supply of Competent Developers

Software engineering is a growing field with a high demand for skilled labor. This means that a high supply of competent new graduates is crucial for the maintenance of existing technology and development of new technology (Spichkova, 2019). The need for competent software developers makes it essential to understand what kind of curriculums produce the most industry ready software engineers (Spichkova, 2019).

This begs the question, how is UVA's computer science (CS) program preparing students for industry in comparison to other university curricula that utilize open source (studied based on systematic literature review methodologies) or multi-semester Capstone projects (from Arizona State University Polytechnic) (Alasbali 2015; see also Gary 2008). Further, how do bootcamps programs, such as Makers Academy or CodeClan, get individuals technically proficient quickly, in comparison to four-year degree programs (Wilson 2017)? While exploring advantages of various curriculums, it is important to consider whether industry demands too much from young developers. Expertise in any field requires years of experience, and asking for high levels of expertise from new graduates may be more counterproductive than helpful. Specifically, setting excessively high expectations for entry level roles may needlessly discourage people from entering the field. It is essential to investigate which computer science curriculums best get students industry ready so that CS graduates from the University of Virginia will leave academia as competent developers.

Bettering Computer Science Education

The computer science curriculum at the University of Virginia could be improved to help students build more practical skills to ensure they are industry ready. Meta analysis of research relating to the topic of creating industry ready developers during undergraduate studies is examined. In this evaluation of the current curriculum, success is defined by the results of third party studies. More specifically, I compare non-traditional curriculums, such as programs that include open source projects or that were created for experimental purposes, to UVA's traditional computer science program. Due to the rapidly changing and expansive nature of technology, it is impossible to create a curriculum that will fully satisfy all levels of industry; however, there are improvements that traditional curricula can make to create better software engineers, including increasing focus on soft skills such as collaboration and communication.

Attitudes towards AI Generated Art

Although bettering computer science education is an important topic, the next portion of this paper focuses on art generated by artificial intelligence (AI). Everytime a new AI model is created, controversy is sparked and concerns about the over reliance of technology arise. One source of controversy emerges due to people viewing AI as a "threat or rival" to human labor (Mazzone & Elgammal, 2019). It is feared that since robots have replaced humans in completing menial tasks in factories, they will eventually replace humans for creating art. Artists' fear of replacement by AI is derived from the high quality of recent AI produced art (Roose, 2022). There is also a concern about whether AI generated art is stealing work from human artists as AI often uses previous artwork for training or stylizing purposes (Salkowitz, 2022). Negative bias toward AI generated art is also derived from the belief that AI is unable to be creative because of its inability to communicate ideas from art to the audience (Yamshchikov & Tikhonov, 2018). Not only that but according to one study, researchers found that opinions toward AI generated artwork were negatively affected when participants were reminded that the AI does not have a sense of self (Lima et al., 2021). This indicates that people's ontological analysis of machines may prevent them from even considering art created by AI as true art (Lima et al., 2021).

Although there are a lot of negative opinions directed towards AI generated artwork, one study found that people tend to unknowingly rate AI generated artwork about the same as human generated art when considering aspects such as originality, degree of improvement or growth, development of personal style, experimentation or risk taking, and communication of ideas (Hong & Curran, 2019). Since a disparity between opinions towards AI generated art is created when the artist's identity is transparent, it makes one question whether these negative biases are rational. Because negative perceptions of AI generated artwork might be irrational, it is important to verify whether people's understanding of AI models and model training will influence their perception of the technology.

Before defining particular models this paper will investigate, we should instead establish a basic understanding for the terms AI, models, and training. Firstly, let's define what differentiates general programs from artificial intelligence. When general programs execute, they either take in an input or execute a unit of code when asked and produce an output that is deterministic. However, with AI, the decision making of this software should mimic that of a human. Therefore, while one can have a set of expectations for the output an AI will produce, it may do something random. This aspect of randomness is similar to the decision making process of a human. For instance, if Bill goes to get coffee from Starbucks everyday before work, we can hypothesize that he will go get coffee from Starbucks tomorrow. However, Bill may choose to be adventurous one day and go to Dunkin Donuts instead for no clear reason. AI may make similarly random decisions when run.

AI models are concerned with *how* a particular AI will make decisions. On a lower level, the how, relates to specific techniques employed that will influence the output the AI will produce. The techniques are heavily dependent on various mathematical concepts. Asking various AI models to complete a specific task is comparable to asking two different people to drive to Starbucks. They'll both probably get in a car and attempt to reach Starbucks (or at least a coffee shop) but the exact route they take to reach their destination may be different.

Training AI models directly relates to the input you feed the algorithm. In order to produce answers similar to humans, the algorithm needs to experience information the same way humans do. The training will influence the AI models' outputs greatly. This training is related to how humans are influenced by the environments they grow up in. Someone growing up in rural Texas versus New York City are going to have different values, personalities, and biases. The initial environment they grow up in will directly influence how each person behaves later on in life. The same is true for AI models. If you train the model on one specific set of inputs, it will later produce outputs that are directly biased towards what they were initially exposed to.

Since models and their training are mimicking human development and decision making, evaluating all models equally is unfair. Analyzing a variety of models is important for gaining an understanding about people's exact perception on AI generated art. The specific models this research will focus on are DALL・E, CAN, and GLIDE. The DALL・E model does text to image generation, and it generates art with an autoregressive approach (predicting future outputs based on previous inputs) (Ramesh et al., 2021). DALL・E was trained with 12 billion inputs, where text and images were fed into the model as one stream of data, that were unlabeled (no human verified image description) images obtained from the internet (Ramesh et al., 2021).

The CAN (creative adversarial network) model generates art after following a two step process which includes learning and creative steps. In the learning step, this network called a discriminator learns about different art and styles based on previous inputs (Elgammal et al., 2017). Then it initiates a creative process where a network called a generator attempts to create art when given a random input that was not given to the discriminator (Elgammal et al., 2017). The generator attempts to make art that the discriminator deems as art but will also confuse the discriminator in determining the style of the art produced (Elgammal et al., 2017).

The GLIDE (Guided Language to Image Diffusion for Generation and Editing) model works by utilizing concepts of a diffusion model (Sapunov, 2022). The diffusion model involves taking in input, adding and deleting information from it (adding noise), then asking the model to reconstruct the input by eliminating the noise (Nichol, 2022). The model utilizes two different guidance strategies (CLIP and classifier-free) (Nichol, 2022). The guidance strategies basically lead the noisy image back to some average form. The final image will be different from the original and without noise.

For this paper, it is relevant to understand the current technology for AI generative art because we are interested in investigating how the techniques of each model will influence the perception of the art that an AI generates. It would also be interesting to understand if there would be a difference of opinion on AI generated art between various social groups if the models used to generate art were not a black box technology.

Since the analysis of particular AI models is important to this research, defining the model in a particular way is important with this assessment. Models will be described using Latour's definitions of prescription, circumscription, and description in an effort to better analyze what issues each social group has with an AI model. Prescription is defined as the "moral and ethical dimension of a mechanism" (Latour, 1992, pp. 157). The paper will use prescriptions to define *what* values AI generated art has and who the technology discriminates against. Circumscription explains the limits the technology has due to external factors (Latour, 1992, pp. 162-163). Here, it will be used to identify *when* art generated by AI affects users. Description of

AI generated art defines the exact mechanism that the technology is built on and it will be used to define *how* AI generates art (Latour, 1992).

Research Question and Methods

This research endeavors to answer the question of: How do people's perception of AI generated art vary depending on the specific technology (model) used to generate the art? This question is important because it could help identify specific issues people have with AI generated art. After specific issues are identified, software engineers can then approach a solution to develop better models.

This question will be assessed by creating focus groups and conducting interviews. Focus groups will be conducted with the general public and artists. Interviews will be held with software developers.

For both groups, I will ask them to analyze a series of art generated by the previously listed models before and after explaining how each model works. The criteria for analysis will be a rating system (1-5) where each work can be rated in categories described by Hong and Curran (2019): originality, degree of improvement or growth, composition, development of personal style, degree of expression, experimentation or risk taking, aesthetic value, and successful communication of ideas. Interviews with software developers will follow a similar process just at an individual level. If opinions change, before and after explaining how models work, will ask questions of why opinions change.

The results of the experiment will be analyzed using a variety of statistical t-tests to reveal the differences between unpaired group.

Conclusion

Overall, the problem is that if developers do not understand the issues people have from AI generative art, then they cannot make improvements to their models. Improvements to AI generative art would have the potential to completely alter how the medium art is viewed and valued. I think the paper will unveil biases people have towards AI generative art and we will be able to link if ignorance plays a factor on its negative perception.

References

Alasbali, N. & Benatallah, B. (2015). Open source as an innovative approach in computer science education a systematic review of advantages and challenges. *2015 IEEE 3rd International Conference on MOOCs, Innovation and Technology in Education*. http://dx.doi.org/10.1109/mite.2015.7375330

Elgammal, A., Liu, B., Elhoseiny, M., & Mazzone, M. (2017). CAN: creative adversarial networks generating “art” by learning about styles and deviating from style norms. *ArXiv: Artificial Intelligence*. https://doi.org/10.48550/arXiv.1706.07068

Gary, K. (2008). The Software Enterprise: Practicing Best Practices in Software Engineering Education.. Retrieved from https://www.researchgate.net/publication/263602679_The_Software_Enterprise_Practicing_Best_Practices_in_Software_Engineering_Education

Hong, J. W., & Curran, N. M. (2019). Artificial Intelligence, Artists, and Art: Attitudes Toward Artwork Produced by Humans vs. Artificial Intelligence. *ACM Transactions on Multimedia Computing, Communications, and Applications*, *15*(2s), 1-16. https://doi.org/10.1145/3326337

Latour, B. (1992). Where are the Missing Masses? The Sociology of a Few Mundane Artifacts. In W. E. Bijker, & J. Law (Eds.), *Shaping Technology/Building Society: Studies in Sociotechnical Change* (pp. 225-258). Cambridge, Massachusetts: MIT Press.

Lima, G., Zhunis, A., Manovich, L., & Cha, M. (2021). On the social-relational moral standing of ai: an empirical study using ai-generated art. *Frontiers in Robotics and AI, 8*. doi:10.3389/frobt.2021.719944

Mazzone, M., & Elgammal, A. (2019). Art, creativity, and the potential of artificial intelligence. *Arts, 8*(1), 26. doi:10.3390/arts8010026

Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., & Chen, M. (2022). GLIDE: towards photorealistic image generation and editing with text-guided diffusion models. *ArXiv: Artificial Intelligence*. https://doi.org/10.48550/arXiv.2112.10741

Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., . . . Sutskever, I. (2021). Zero-shot text-to-image generation. *ArXiv: Artificial Intelligence*. https://doi.org/10.48550/arXiv.2102.12092

Roose, K. (2022). AI-Generated Art Won a Prize. Artists Aren't Happy. [The New York Times]. Retrieved from https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html

Salkowitz, R. (2022). AI Is Coming For Commercial Art Jobs. Can It Be Stopped? [Forbes]. Retrieved from https://www.forbes.com/sites/robsalkowitz/2022/09/16/ai-is-coming-for-commercial-art-jobs-can-it-be-stopped/?sh=11c9baf454b0

Sapunov, G. (2022). OpenAI and the road to text-guided image generation: DALL·E, CLIP, GLIDE, DALL·E 2 (unCLIP) [Medium article]. Retrieved from https://blog.inten.to/openai-and-the-road-to-text-guided-image-generation-dall-e-clip-glide-dall-e-2-unclip-c6e28f7194ea.

Spichkova, M. (2019). Industry-Oriented Project-based Learning of Software Engineering. *International Conference on Engineering of Complex Computer Systems*. doi: 10.1109/ICECCS.2019.00013.

Wilson, G. (2017). Building a new mythology: The Coding Boot-camp Phenomenon. *ACM Inroads* 8(4), 66–71. http://dx.doi.org/10.1145/3132706

Yamshchikov, I. P., & Tikhonov, A. (2018). I feel you: What makes algorithmic experience personal? *Electronic Workshops in Computing*. doi:10.14236/ewic/evac18.32