Thunderstruck the Meter – An Affordable DC Panel Meter

The Impact of AI on Future Education

A Thesis Prospectus Submitted to the

Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements of the Degree
Bachelor of Science, School of Engineering

,

Technical Project Team Members
Jacob Clatterbuck
Kristian Johnson
Hieu Le

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

Signature _____ Date _____

Approved _____ Date _____
, Department of

Approved _____ Date _____
, Department of Engineering and Society

Introduction

While there are existing DC Panel Meters on the market, they can be extremely expensive, ranging anywhere from one hundred dollars up to several hundred dollars. The available options either include unnecessary features or are poorly designed and provide low accuracy with no protection. Our technical project aims to find a middle ground, taking into consideration both an attractive price point as well as well thought out functionality. We want to design a DC panel meter that can read current and voltage down to 1mA and 10mV resolution, while still providing the necessary reverse polarity, short circuit, and over-current protection. The device will be designed to be paired with an existing power supply, eliminating the need for purchasing a new and more expensive power supply that has these filters built-in. While this design choice does eliminate waste by allowing the user to utilize a preexisting power supply to create a lab bench, making the design more sustainable and environmentally sound, it does create an ethical issue in the form of removing business and jobs from manufacturing companies that are creating newer products with these built-in features. The technical subject of the STS prospectus and the technical topic for the Dept. of Electrical and Computer Engineering are not related.

With the recent advances in Natural Language Processing and AI technologies, many new societal challenges are being discovered. While these tools are incredibly powerful and stand to provide an incredible boon to humanity, there have already been incidents of individuals using them for dishonest purposes. My STS topic seeks to investigate the impact these AI technologies will have on the field of education, both positive and negative. Uses such as Automated Teaching Assistants can help students and teachers by providing immediate feedback and information while reducing workload, but uses such as instant essay generation could harm society by allowing students to plagiarize undetected with ease.

Technical Topic

The goal of this project, Thunderstruck the Meter, is to create an inexpensive DC Panel Meter that combines both the functionality of a voltmeter and a current meter for use alongside an existing power supply. The maximum values that the meter is capable of measuring are 100V and 50A. To accomplish this, a voltage divider and current sensor IC will be used in the design. In addition, the product will also feature an MSP430 microcontroller and a 16bit Analog to Digital Converter (ADC), which can measure with accuracy down to 45uV. This allows the voltage and current from the power supply to be shown on a 16x2 Liquid Crystal Display (LCD). The device will also be designed with reverse polarity protection, short circuit protection, and over-current protection. Our final product aims to be small, compact, and accessible to any user with very little experience regarding electrical systems. Using any existing power supply already at hand, the user can turn it into a more useful power supply bench. Because the device has reverse polarity, short circuit, and over-current protection, it can also act as a test circuit for any power supply source below 100V.

The maximum values that the meter is capable of measuring are 100V and 50A. The panel can read current and voltage down to 1mA and 10mV resolution, respectively. The device will be designed to take a very small amount of power from the power supply that it is measuring in order to power itself up to 5V. If the voltage of the power supply exceeds 50V, we will also have an option of using an external power supply using a single pole double throw switch (SPDT). Two buck converter will be used to reduce the voltage from the supply to the 5.0V and 3.3V needed for the microcontroller and other powered elements. The circuit protections the meter will be designed with will increase both the testing safety and the lifespan of the product.

The device will operate in two modes: under 50V power supply mode, and between 50V and 100V power supply mode. The user can switch between the two modes by using the SPDT switch. In

the under 50V power supply mode, the device will use the power supply that it is measuring to power itself. The current will go through a PMOS diode current feedback loop which is used as the over-current protection. We will use another PMOS diode to block out the current if voltage polarity is reversed. If nothing is shorted and the voltage polarity is correct, the current can flow through the switching buck converter. The buck converter is rated as a 50V input such that any voltage below 50V is completely safe. The advantage of using the switching buck converter is that it has an amazing efficiency reaching up to 90%. The buck converters will be the main power source for all of the ICs on the circuit, including the MSP430, the LCD, the ADC, and the current sensor . The device uses a voltage divider to sense the voltage supply, and uses a current sensor IC (ACS758) to sense the current. The voltage and current readings will go to a 16 bit Analog to Digital Converter (ADC) (ADS 1115) which has 65k steps. The ADC allows us to measure the signal as low as 45uV with a reference voltage at 3V. The ADC will send the signal to the MSP430 to process using the I2C communication protocol. After processing the signals, the MSP430 will display the reading on a 16x2 LCD, using an averaging filter in software to filter out noise, with zeroes at the common mains hum frequency of 60Hz.

In the under 50V power supply mode, the device will need to use an external power supply to supply its power. The reason for this is that the buck converter can only handle up to 50V. Anything over that will cause serious damage to the whole board. Any DC power supply source over 5V will suffice.

STS Prospectus

Introduction

Recently, the field of natural language processing has made a huge leap thanks to the release of OpenAI's Generative Pre-trained Transformer 3 (GPT-3) in May of 2020. GPT-3 is a neural-net based language prediction model that can be used to generate text and essays that are difficult to distinguish from those that are human made. The potential benefit this AI can offer humanity is extremely promising, with individuals already having been able to guide it to generate Javascript code from basic instructions. With minimal human guidance, GPT-3 has also been used to generate entire blog posts and articles that are functionally indistinguishable from human writing made from scratch (Philip 2020). While this presents an incredible opportunity to industry, it also contains a potential risk for academic integrity. Basic high school essays can be generated in seconds with a prompt of a few sentences. Even as early as 2005, a far weaker language generator called SCIgen was used by MIT students to trick an academic journal into accepting a fake paper made by AI. The technological artifact of GPT-3 represents one of many advancements in AI that can challenge future educational systems. The main topic I wish to explore is the impact generative AIs like these will have on academic institutions.

While language generation technologies are still in their infancy, the educational world has already taken steps to find uses for it. Some schools and universities have installed automated teaching assistants to reduce the workload for teachers (Preston 2020). Others abroad in China have more recently used AI systems to ensure students are present and learning during online learning in the midst of the COVID-19 pandemic (Liu 2020). There have been several papers comparing and examining the costs or benefits of AI in education, but there is a gap in the literature for a holistic analysis of both the benefits and potential downsides for all relevant stakeholders.

Research Questions

This topic leads to many research questions that explore the social impacts of AI on future education. Like the available literature, the questions can be loosely grouped into four major themes. Some questions deal with the current status of AI Education (AIEd). What subjects, specifically, can benefit from AI enhanced learning? What are the most common uses of AI in modern education, and what demonstrable benefit do they show? How common is the use of AI in education in modern day? Are there major demographic differences in who has access to these technologies and the benefits they may offer?

Other questions deal with examining the benefits of AIEd. How will the quality of work the students produce change when they have access to adaptive learning tools and automated TAs? When will this technology become commonplace enough that these benefits become widespread?

Then there are the research questions that explore the social and educational impacts as newer technologies are adopted into the educational system. How will grading systems and other methods of measuring academic excellence have to change? How will students use language generators, and will they serve as in an assistive role or will they be used to cheat and generate entire essays? How will AI change how we evaluate student performance?

Finally there are the questions that examine the risks and potential harms that developing AIEd can have on society. Will richer students who have access to this technology have an advantage over those who don't? Will the use of AI tools in the classroom lead to a fairer distribution of academic resources or will it only exacerbate the problem of class differences? How are the developers at OpenAI and other AI research teams addressing these issues?

Literature Review

The existing literature on the topic of AI in education can be grouped into four major themes: overviews of the current status, papers analyzing the potential benefits, arguments for what the future of AIEd will look like, and information that raises caution about potential risks. Current literature is stratified into these categories few take a holistic approach and with the rapidly advancing field of AI, there are gaps in the research.

In terms of an overview of modern uses, Lu, J. and Harris, L. (2018) offer the most succinct paper that addresses a variety of issues. The purpose of their article and organization is to provide a nonpartisan information source to members of Congress in an easily consumable manner, and it is therefore brief and focuses on American values and interests. It gives examples of current uses of AI in education such as Intelligent Tutoring Systems (ITS), Personalized Learning, Testing, and Task Automation. Policy considerations are also enumerated so that legislators may seek to take action on protecting student privacy and product procurement. This source is valuable for its legal analysis and ease of understanding, and is particularly useful for addressing research questions about the current use of AI in education. For a more in-depth overview of ITS, Beck, Stern, & Haugsjaa (1996) offer detailed descriptions of the different kinds of models in AI tutor systems and how they can be combined to create an effective ITS. This is an older source that lacks any insight on newer technologies and techniques like neural networks, but is invaluable for its breakdown on a commonly used application of AI in education. For a more modern source with both breadth and depth covering the current status of AIEd, Luckin, Holmes, Griffiths, & Forcier (2016) enumerate in their book the current uses, models, and benefits. In their paper, they closely examine the efficacy of current AI models used to provide feedback to students and discuss the possibility of robots replacing teachers altogether.

Another theme of the academic literature that some of these sources fall into is that of benefits analysis research. Perotta and Selwyn (2019) give a detailed examination of scholarly work carried out

on the effect of deep learning to predict educational performance. Their meta study analyses the existing research to find patterns in how much success students find with the help of AI in different educational topics. Perotta and Selwyn (2019) also look at demographic differences such as race and class to find differences in benefits gained. They conclude that there are different relations between different versions of learning and they question the scientific objectivity and neutrality of AIEd systems. For a more positive outlook, Zhu & Zhang (2010), argue that ITS can help teachers by dynamically distinguishing in which concepts a student has shortcomings, finding that the immediate feedback AI can offer is unattainable through hand grading and other traditional methods of evaluation.

There are also associated risks with the increased usage of AI systems in education, both in the form of privacy and plagiarism concerns. In their paper, Arai & Matsuzaki (2014), argue that AI systems in 2016 were capable of passing the entrance examinations of more than half the universities in Japan. They raise concerns over what this means in terms of the value of what is being tested, implying that education will have to move away from the sorts of tasks that artificial intelligence can do with ease. Nils & Mossink (2020) performed an experiment with a state-of-the-art Natural Language Generation algorithm that showed humans are unable to tell AI made poetry apart from that written by humans. They argue that the implications of this study are numerous, including raising questions about academic fraud and how much the developer and training data authors should be credited. For non-poetry essays, Philip (2020) offers an anecdotal account of GPT-3 being used on the social media site, Reddit, and going undetected for some time. This source offers an example of how modern AI uses can be used for academic fraud to generate, at times, passable papers and short stories that are untraceable with normal plagiarism detection algorithms. The earlier mentioned congressional research (Lu & Harris, 2018) offers an account of privacy concerns for students that may be in conflict with existing laws.

A last categorical theme found in the retrieved sources are those articles that speculate and hypothesize about what the next phase of AIEd should or might look like. Andriessen and Sandberg (1999) argue that the role of AI in education is due a fundamental shift away from an ITS-based paradigm and towards a model that helps students "learn how to learn". Learning how to learn means providing students with the cognitive tools to better understand effective research and memorization of important knowledge. They present the compelling evidence that ITS helps with rote memorization, but argue for AI systems that can help with scenarios that deal with non-fixated learning goals. Neil Selwyn (2019) in his book makes the case for human teachers and claims that AI is a "double-edged sword". While there are some fields that AI excels in, such as immediate feedback, there are downsides that include loss of jobs, heightened differences between different types of learners such as visual learners and social learners, and data privacy concerns. The existing literature is in agreement that current technology is no where near replacing human teachers, but AI aids can be an invaluable resource.

STS Framework and Method

Of the STS frameworks used to analyze this topic, three in particular stood out as highly relevant, Winner's politics of artifacts framework, Hughes' LTS framework, and Pinch and Bijker's SCOT framework. Winner's framework is particularly useful for addressing the research questions focused on the deployment and availability of the technology for individual AI systems such as GPT-3. These deployments have relevance to macro politics through the lens of resource distribution of different school systems, creating potential differentiation between classes and locations. Hughes' large technical system framework is well suited to addressing the research questions pertaining to the technological momentum of AI's adoption in education and learning institutions. It is also useful for investigation into reverse salients that may impede the development and incorporation of these

technologies. LTS analysis can also provide clarity into why some countries are being faster to adopt AIEd than others through its reverse salient concept. Perhaps most relevant to the topic at hand is the framework of SCOT. With SCOT, the technological artifacts can be interpreted through the different lenses of stakeholders in the education system such as parents, teachers, and of course students.

For my research I wish to make use of several methods that may help provide information of the current state of AI in education. Intelligent Tutoring Systems are the most broadly used example of AI in education. Interviews with those who have interacted with them can help shed some light on how helpful they were, along with any problems in the adoption of the technology. Ideally, I want the input of multiple stakeholders such as students, teachers, and teaching assistants. Retrieving and analyzing the syllabus for classes taught with AI assistance and comparing them to a normally taught class's syllabus may also provide some insight. Additionally I aim to perform document analysis on existing educational laws and standards of education that may affect the adoption of AI in the classroom. While there have been studies showing that certain modern Natural Language Processing algorithms cannot be distinguished from some human writing under some conditions, there is no data as to how these AI-generated essays might be graded if used for academic fraud. A survey of middle and high school teachers would therefore be an appropriate way to determine the efficacy of such a scheme, gathering data on the letter grade an AI-made essay might receive at different grade levels.

Timeline

The goal is to complete the data collection as early as possible, ideally within the first month so that further research can be done with the results in mind. If there are significant time delays in accessing relevant individuals for interviews, then it would also be beneficial to start as early as possible. The intention is to have all data collection complete by early March.

Conclusion

This STS prospectus topic is composed of both a review on existing information on the current impact of AI in education and an analysis for the development of a research plan for analyzing the future impact. While the most cutting-edge NLP algorithms, developed in the last year have been used to pose as real humans, little to no original research has been performed to investigate the potential for use as either learning tools or the potential threat of widely available essay generators for students. The data collection proposed would therefore represent a novel contribution to the ongoing questions surrounding AI education. Many existing papers seek to either analyze benefits or focus on the potential risks and costs that introducing AI into education may pose, leaving a gap in the research for a more holistic approach. Ideally this thesis can bridge this gap and provide a balanced examination of these cutting-edge technologies.

Bibliography

Andreissen, J., & Sandberg, J. (1999). Where is Education Heading and How About AI? *International Journal of Artificial Intelligence in Education,*10, 130-150.

Arai, N., & Matsuzaki, T. (2014). The impact of A.I. on education – Can a robot get into The University of Tokyo? *National Institute of Informatics*. Retrieved October 24, 2020 from http://icce2014.jaist.ac.jp/icce2014/wp-content/uploads/2014/12/ICCE2014DIpaper.pdf

Beck, J., Stern, M., & Haugsjaa E. (1996 September). Applications of AI in Education. *XRDS: Crossroads, The ACM Magazine for Students.* Retrieved from https://dl.acm.org/

Liu, Y. (2020, May 18). The Future of the Classroom? China's experience of AI in education. Retrieved November 11, 2020 from https://www.nesta.org.uk/report/the-future-of-the-classroom/

Lu, J., Harris, L. (2018). Artificial Intelligence (AI) and Education. *Congressional Research Service*. Retrieved October 25, 2020, from https://crsreports.congress.gov/product/pdf/IF/IF10937/3

Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. (2016). *Intelligence Unleashed: An argument for AI in Education.* London: Pearson Education.

Nils K., & Mossink, L. (2020). Artificial Intelligence versus Maya Angelou: Experimental evidence that people cannot differentiate AI-generated from human-written poetry. *Computers in Human Behavior 2020*. Retrieved October 25, 2020 from arXiv.

Perrotta, C., & Selwyn, N. (2019). Deep learning goes to school: toward a relational understanding of AI in education. *Learning, Media and Technology*, 45, 251-269.

Philip (2020, October 6). GPT-3 Bot Posed as a Human on AskReddit for a Week. Retrieved October 25, 2020 from https://www.kmeme.com/2020/10/gpt-3-bot-went-undetected-askreddit-for.html

Preston, J. (2020, January 24). The World's First AI Teaching Assistant Turns 4. Retrieved November 11,2020 from https://www.cc.gatech.edu/news/631545/worlds-first-ai-teaching-assistant-turns-4

Selwyn, N. (2019). *Should Robots Replace Teachers?: AI and the Future of Education.* Cambridge: Polity Press.

Zhu, G., Zhang, K., & Li, J. (2010). Intelligent tutoring system based on computing conceptual graphs. *Applied Mechanics and Materials*, 347, 2430.