

ARTIFICIAL INTELLIGENCE: THE FUTURE OF PEDAGOGY

By Brandon Mattalo¹

ABSTRACT

Generative Artificial Intelligence is rapidly permeating people's lives. In the past year, numerous generative AI technologies have been released to the public, including Open AI's GPT-4 and Dall-E, Google's Bard, Adobe's Firefly, and the independently operated Midjourney, amongst others.

The public release of these technologies has, understandably, caused concern in academic circles. Generative AI can be used to create unique content that cannot be detected by traditional plagiarism checkers. While it is important to research the negative impact of generative AI on academic integrity, I suggest that academics should focus most of their efforts on the opportunities these technologies present for improving pedagogical practices.

In this paper, I attempt to flip the narrative from one of fear to one of opportunity. I suggest that academics should research the use of generative AI within the academy to improve teaching effectiveness and efficiency. I offer various practical suggestions on how these tools can be used to advance pedagogical practices, with specific examples provided in the field of study that is my primary focus — business law.

I advance this argument in three parts. First, I provide a brief history of artificial intelligence and outline how the technology works, its capabilities and its limitations. Second, I briefly review the current state of scholarship on using AI technologies in pedagogical practices and provide a suggested framework for how AI pedagogical research could proceed. Third, I provide examples of how to use AI in the proposed framework, including the use of the technology in teaching business law. Finally, I conclude that there is an open field of opportunities for implementing and researching the use of generative AI in pedagogical practices.

¹ Brandon Mattalo is an Assistant Professor at the Lazaridis School of Business and Economics at Wilfrid Laurier University.

TABLE OF CONTENTS

Part I – Introduction to Artificial Intelligence	3
History Of Artificial Intelligence	3
Modern Artificial Intelligence.....	5
Generative Artificial Intelligence Capabilities.....	8
Part II – AI and Pedagogy Research	11
The Current State of AIEd	11
Proposed Learning AIEd Research Framework.....	13
Part III – Potential of the Modern Learning AIEd Framework.....	15
Pedagogical Enhancement Research	15
Pedagogical Breakthrough Research.....	19
Administrative Research	21
Potential Issues with AI in Pedagogy.....	24
Conclusion.....	26
Bibliography	28
Appendix I: Bing Produced Previewing.....	34
Appendix II: GPT-4 Assisted Questioning, Feedback and Grading	36
Appendix III: GPT-4 Assisted Exam Creation.....	39
Appendix IV: GPT-4 Assisted Study Group.....	41
Appendix V: GPT-4 Assisted Personalized Learning.....	46
Appendix VI: Bing Assisted Augmented Reading	49
Appendix VII: GPT-4 Produced Syllabus for a Business Law Course.....	52
Appendix VIII: GPT-4 Assisted Marking assistance.....	56

PART I – INTRODUCTION TO ARTIFICIAL INTELLIGENCE

To understand how generative AI can be used in pedagogy, one first needs to understand AI and its potential. In this first part, I provide a history of AI to provide perspective about the current time we find ourselves in the field of AI. I then explain how AI works, its current capabilities, and its expected capabilities in the future.

History Of Artificial Intelligence

Despite its recent popularity, artificial intelligence is not a new field — it has been studied as a scientific pursuit for over 70 years.² The study of AI was pioneered by Alan Turing, who was known for conceptualizing and building one of the first modern computers. His computer was famously used to decrypt the German’s Enigma encryption code during the Second World War.³

Turing published one of the first AI papers in 1950 entitled “Computing Machinery and Intelligence”.⁴ In this paper, Turing first described the now-famous Turing Test, which is a test he proposed for determining when a computer’s intelligence was indistinguishable from the intelligence of a human.⁵

Turing did not use the term “artificial intelligence” to describe his research. Instead, that term was first introduced by the scholars John McCarthy and Marvin Minsky during a conference they organized in 1956 in New Hampshire – the *Dartmouth Summer Research Project on Artificial Intelligence*.⁶

Significant time and money went into researching AI in the 1950s up to the 1970s.⁷ It was a promising field. One of the first AI technologies was the Eliza chatbot, invented in 1966 by Joseph Weizenbaum, an MIT professor.⁸ The Eliza chatbot attempted to mimic the behaviours of a psychologist in the treatment of patients.⁹ Like many of the systems developed around this time, the Eliza chatbot did not work the same way that modern AI systems work today. Instead, the Eliza chatbot used a pattern matching algorithm.¹⁰ This algorithm works by storing pre-determined

² This, of course, ignores the significant contributions that art and mythology made to the theories of artificial intelligence, as discussed in more detail in Bolt Beranek, H A Simon & P McCorduck, *History of Artificial Intelligence* (Cambridge, Massachusetts).

³ Michael Haenlein & Andreas Kaplan, “A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence” (2019) 61:4 California Management Review 5–14 at 6.

⁴ Alan M Turing, “Computing Machinery and Intelligence” (1950) LIX:236 Mind 433–460.

⁵ *Ibid.*

⁶ Haenlein & Kaplan, “A Brief History of Artificial Intelligence”, *supra* note 3 at 7; *Artificial Intelligence and Education: A Critical View Through the Lens of Human Rights, Democracy and the Rule of Law*, by Wayne Holmes (Strasbourg: Council of Europe, 2022) at 18.

⁷ Haenlein and Kaplan provide a good overview of the advancements in the field, including the creation of the ELIZA computer program (an early natural language processor) and the General Problem Solver program. See Haenlein & Kaplan, “A Brief History of Artificial Intelligence”, *supra* note 3.

⁸ Luka Bradeško & Dunja Mladenčić, “A Survey of Chabot Systems through a Loebner Prize Competition” (2012).

⁹ *Ibid.*

¹⁰ Siddhi Pardeshi et al, “A survey on Different Algorithms used in Chatbot” (2020) 07:05 IRJET 6092–6098 at 6094.

“question” templates and matching them to corresponding “answers” in a database.¹¹ When a templated question is asked by a user, the system retrieves the corresponding answer from the database.¹²

These systems were not complex by today’s standards, since they lacked self-learning capabilities and were limited to providing pre-determined answers to pre-determined question templates.¹³ If a user asked a question that was not stored in the database, it would be unable to answer the question. If an answer changed, then the answer would remain outdated until manually updated by in the database.

Unfortunately, AI research hit a wall in the 1970s when James Lighthall, a British Researcher, released a report on behalf of the UK’s Science Research Council where he concluded that “in no part of the field have discoveries made so far produced the major impact that was then promised”.¹⁴ DARPA in the United States released a similar report a few years later with a similar conclusion.¹⁵ This led to a defunding of AI Research, and is referred to as one of the first “AI Winters”.¹⁶ This boom and bust cycle in AI research continues today.¹⁷

As computing power increased exponentially, and more training data became available on the internet, various breakthroughs in AI research started making public headlines.¹⁸ Many of these breakthroughs were commercialized, while others were captured in public spectacles. Notable examples include the 1997 chess match where IBM’s Deep Blue system beat Kasparov,¹⁹ the 2011 jeopardy match where IBM’s Watson outperformed jeopardy champions,²⁰ the 2015 Go match where DeepMind’s AlphaGo system beat Fan Hui,²¹ and the 2020 scientific breakthrough when DeepMind’s AlphaFold was able to make transformational advancements in solving the protein folding problem.²²

¹¹ *Ibid.*

¹² *Ibid.*

¹³ *Ibid*; Holmes, *supra* note 6 at 18.

¹⁴ James Hendler, “Avoiding Another AI Winter” (2008) 23:2 IEEE Intelligent Systems 2–4 at 2; *Artificial Intelligence: A General Survey*, by Sir James Lighthill (1972).

¹⁵ Hendler, *supra* note 14 at 2.

¹⁶ There were numerous “AI Winters” over the years, as discussed in Hendler, *supra* note 14; Luciano Floridi, “AI and Its New Winter: from Myths to Realities” (2020) 33:1 Philos Technol 1–3 at 1; Haenlein & Kaplan, “A Brief History of Artificial Intelligence”, *supra* note 3 at 8–10.

¹⁷ Floridi, “AI and Its New Winter”, *supra* note 16 at 2–3; Hendler, *supra* note 14 at 3–4.

¹⁸ Holmes, *supra* note 6 at 18.

¹⁹ “IBM100 - Deep Blue”, (7 March 2012), online: <<http://www-03.ibm.com/ibm/history/ibm100/us/en/icons/deepblue/>>.

²⁰ “IBM100 - A Computer Called Watson”, (7 March 2012), online: <<http://www-03.ibm.com/ibm/history/ibm100/us/en/icons/watson/>>.

²¹ James Vincent, “Former Go champion beaten by DeepMind retires after declaring AI invincible”, (27 November 2019), online: *The Verge* <<https://www.theverge.com/2019/11/27/20985260/ai-go-alphago-lee-se-dol-retired-deepmind-defeat>>.

²² “Timeline of a breakthrough”, online: <<https://www.deepmind.com/research/highlighted-research/alphafold/timeline-of-a-breakthrough>>.

In addition to these public spectacles, artificial intelligence has slowly crept into our lives — often unnoticed. Everyday examples of AI include the summaries provided to you at the top of google searches,²³ the facial recognition software used to unlock your phone,²⁴ the music recommendation algorithms on Spotify,²⁵ and the voice-activated digital assistants we all use on our phones,²⁶ and the self-driving features shipped with many new vehicles.²⁷ This list is by no means exhaustive — AI is already here and a staple in our everyday lives.

Modern Artificial Intelligence

There is a significant difference between the primitive AI technologies developed in the 1950s and 1960s and the modern AI technologies that now provide us with our modern-day conveniences.

The early AI systems attempted to recreate human intelligence based on the philosophy that human intelligence was formalistic, meaning that all human intelligence could be represented by a branching tree of if-then statements that, if followed, would lead to the ultimate answer to a question.²⁸ Under this view, making “smarter” systems was simply a process of increasing the complexity of the if-then statements and increasing the computing capacity to process more arguments in a shorter period of time. These types of systems are known as “Expert Systems” and they tend to perform well at formalistic processes (such as a game of chess or Go, where the rules are set, the permutations are limited, and the paths to victory can be mapped and evaluated) but tend to perform poorly at tasks that cannot be easily formalized (such as image recognition, language processing, or the creation of new works).²⁹ Some do not regard Expert Systems as true artificial intelligence, since they rely on pre-determined logic input by humans, are unable to learn from external data sources, and are unable to improve their abilities — a hallmarks of intelligence.³⁰

Modern AI systems, on the other hand, stem from a different philosophy of how human intelligence works. Instead of human intelligence being a complex series of binary arguments of increasingly complexity that can be formalized, human intelligence is better represented by taking in external data, assigning various probabilities of what event will follow the input, and producing

²³ “How AI is powering a more helpful Google”, (15 October 2020), online: *Google* <<https://blog.google/products/search/search-on/>>.

²⁴ “An On-device Deep Neural Network for Face Detection”, online: *Apple Machine Learning Research* <<https://machinelearning.apple.com/research/face-detection>>.

²⁵ James McInerney et al, *Explore, exploit, and explain: personalizing explainable recommendations with bandits* (Vancouver British Columbia Canada: ACM, 2018).

²⁶ “Deep Learning for Siri’s Voice: On-device Deep Mixture Density Networks for Hybrid Unit Selection Synthesis”, online: *Apple Machine Learning Research* <<https://machinelearning.apple.com/research/siri-voices>>.

²⁷ “AI & Robotics”, online: *Tesla* <<https://www.tesla.com/AI>>; Peter Lyon, “Honda’s New Safety Target Focuses On AI-Powered Technology”, online: *Forbes* <<https://www.forbes.com/sites/peterlyon/2021/11/26/hondas-new-safety-target-focuses-on-ai-powered-technology/>>.

²⁸ Haenlein & Kaplan, “A Brief History of Artificial Intelligence”, *supra* note 3 at 8.

²⁹ *Ibid.*

³⁰ *Ibid.*

an output based on the most expected outcome.³¹ This approach to AI was arguably pioneered by Donald Hebb, a Canadian psychologist, who hypothesized in the 1940s about using statistical models to mimic neurons in the human brain — a theory that led to the creation of artificial neural networks much later.³² While this theoretical approach was enticing, it did not make significant advancements until the past few decades because the computing capacity did not exist to operate such systems at scale.³³ This theory of artificial neural networks became a commercial reality in at the turn of the century, and now artificial neural networks power many of the modern AI systems we use today, from self-driving cars to voice assistants.³⁴

Artificial neural networks work by mimicking the functions of a human brain.³⁵ In simple terms, the brain itself has been modelled as a set of approximately ten to one-hundred billion neurons.³⁶ The neurons are connected and pass information to other neurons using electrical events (called action potentials) and chemical signals (called chemical neurotransmitters) through connecting channels (called synapses).³⁷ Chemical neurotransmitters either inhibit or excite a neuron.³⁸ If the total of the excitatory and inhibitory inputs reach an action potential threshold, then they fire and release chemical neurotransmitters to connected neurons, where the process is repeated.³⁹ Over time, the connections between various neurons strengthen, such that any input can be predicted to have an expected output. This is how humans form habits, memories and learn new skills.⁴⁰

Artificial neural networks mimic these processes in the brain. The system is typically comprised of three layers of artificial neurons — the input layer, the hidden layer, and the output layer, all of which are connected through channels.⁴¹ Each channel is given a numerical value known as a weight.⁴² This system is depicted in the diagram, below, with the circles representing neurons and the lines representing channels, and each column representing a different layer:⁴³

³¹ *Ibid* at 8–9.

³² *Ibid* at 8.

³³ *Ibid*.

³⁴ *Ibid*.

³⁵ Anders Krogh, “What are artificial neural networks?” (2008) 26:2 *Nat Biotechnol* 195–197.

³⁶ Jayesh Ahire, “The Artificial Neural Networks handbook: Part 1 - DataScienceCentral.com”, (24 August 2018), online: *Data Science Central* <<https://www.datasciencecentral.com/the-artificial-neural-networks-handbook-part-1/>>.

³⁷ Alberto E Pereda, “Electrical synapses and their functional interactions with chemical synapses” (2014) 15:4 *Nat Rev Neurosci* 250–263 at 250.

³⁸ “Action potentials and synapses”, (22 November 2016), online: *Queensland Brain Institute* <<https://qbi.uq.edu.au/brain-basics/brain/brain-physiology/action-potentials-and-synapses>>.

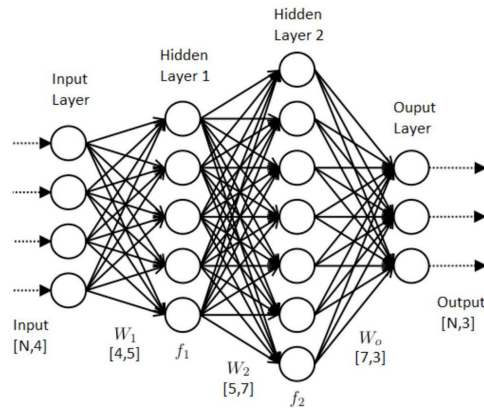
³⁹ Pereda, *supra* note 37.

⁴⁰ *Supra* note 38; Pereda, *supra* note 37 at 250.

⁴¹ Iqbal H Sarker, “Deep Learning: A Comprehensive Overview on Techniques, Taxonomy, Applications and Research Directions” (2021) 2:6 *SN COMPUT SCI* 420 at 5–6; Michael A Nielsen, *Neural Networks and Deep Learning* (Determination Press, 2015).

⁴² Sarker, “Deep Learning”, *supra* note 41 at 5.

⁴³ Ahire, *supra* note 36.



The network receives inputs through the input layer.⁴⁴ The data from the input layer is then passed onto the hidden layers through a channel with a channel weight. The neurons in the input layer pass their value onto the neurons in the hidden layer using a weighted value function.⁴⁵ For example, if there are two neurons in the input layer connected to one neuron in the hidden layer, then the value received at the neuron in the hidden layer may be represented by the following equation: (Neuron 1 Value) x (Channel 1 Weight) + (Neuron 2 Value) x (Channel 2 Weight).⁴⁶

The neurons in the hidden layer have a static random value associated with them called a bias, which is added to the data received from the input layer.⁴⁷ A function is then applied to the data received in the hidden layer, known as the activation function, which determines whether the specific neuron is activated.⁴⁸ If the neuron is activated, it passes on the result to the next layer of the hidden layer over channels with different weights, to neurons with different biases, which themselves apply an activation function to the information received.⁴⁹ The final hidden layer passes its information to the neurons in the output layer. The neurons in the output layer determine the predicted output from the given inputs.

This whole process is known as “forward propagation”.⁵⁰ If the model is untrained, then the output will likely be erroneous, because all the original weights and biases in the model are randomly assigned and have not been trained on any data to be predictive.⁵¹

To increase the accuracy of the model, the model undergoes a process called training to fine-tune the channel weights and neuron biases. This begins with a process called backpropagation.

⁴⁴ Sarker, “Deep Learning”, *supra* note 41 at 5.

⁴⁵ *Ibid* at 2.

⁴⁶ *Ibid*.

⁴⁷ “Machine Learning for Beginners: An Introduction to Neural Networks - victorzhou.com”, online: <<https://victorzhou.com/blog/intro-to-neural-networks/>>; Su-Hyun Han et al, “Artificial Neural Network: Understanding the Basic Concepts without Mathematics” (2018) 17:3 Dement Neurocogn Disord 83–89 at 85.

⁴⁸ Han et al, “Artificial Neural Network”, *supra* note 47 at 85.

⁴⁹ *Ibid* at 86.

⁵⁰ “Forward Propagation, Backward Propagation, and Computational Graphs”, online: *Dive Into Deep Learning* <https://d2l.ai/chapter_multilayer-perceptrons/backprop.html>.

⁵¹ Han et al, “Artificial Neural Network”, *supra* note 47; “Explained: Neural networks”, (14 April 2017), online: *MIT News | Massachusetts Institute of Technology* <<https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>>.

Backpropagation occurs first by determining the error value in the output layer, which is a value representing the difference between the erroneous output and the ideal output.⁵² This error value is then fed backwards through the hidden layer of the neural network and the weights and biases are altered to minimize the error rate.⁵³

This process of forward propagation and backpropagation is repeated through a training process, until the expected output can be predicted with a sufficiently high probability. Training can either be supervised (i.e., using humans or other processes to label inputs or provide feedback about the expected outputs) or unsupervised (i.e., a process whereby the model can find patterns in the inputs and outputs and use those patterns to train itself), or a combination of both.⁵⁴ By training the model on more data over longer periods of time, the model becomes more accurate at predicting an output from any given input with a high probability of success.

This neural network process has unlocked the potential of generative AI in recent years. While the basic premise of a neural network is still used today, the architecture of the model has advanced significantly, with the most advanced release being transformers (which is where GPT gets its name from – Generative Pre-trained Transformer).⁵⁵

Artificial intelligence is no longer limited to a cascading network of pre-determined if-then statements, but now encompasses a complex web of artificial neurons in a model that undergoes a process mimicking the learning process of the human brain.

Generative Artificial Intelligence Capabilities

Due to the advancements in neural networks, and the ability to train those models on significant amounts of internet data, AI models have now reached a point where they can interact with humans and mimic human-like responses with human-like creativity and originality.

While it may appear that they have reached human intelligence, they have not. The point at which artificial intelligence models reach human-like intelligence is known as General AI, Artificial General Intelligence (AGI), or Strong AI.⁵⁶ In a recent paper, Microsoft noted that the most recent iteration of OpenAI’s GPT-4 has shown glimmers of General AI.⁵⁷ This has raised many ethical issues, which itself has caused some of the leading AI researchers to release a public letter demanding a pause on AI research.⁵⁸

⁵² Krogh, *supra* note 35 at 195; Han et al, “Artificial Neural Network”, *supra* note 47 at 86.

⁵³ Krogh, *supra* note 35 at 195; Han et al, “Artificial Neural Network”, *supra* note 47 at 86–87.

⁵⁴ “Supervised vs. Unsupervised Learning: What’s the Difference?”, (15 November 2022), online: <<https://www.ibm.com/cloud/blog/supervised-vs-unsupervised-learning>>.

⁵⁵ Ashish Vaswani et al, *Attention is All you Need* (Curran Associates, Inc., 2017).

⁵⁶ “What is Strong AI? | IBM”, online: <<https://www.ibm.com/topics/strong-ai>>.

⁵⁷ Sébastien Bubeck et al, *Sparks of Artificial General Intelligence: Early experiments with GPT-4* (arXiv, 2023) arXiv:2303.12712 [cs].

⁵⁸ “Pause Giant AI Experiments: An Open Letter”, online: *Future of Life Institute* <<https://futureoflife.org/open-letter/pause-giant-ai-experiments/>>; for interesting discussions on the dangers of General AI, I would suggest reading the works of Eliezer Yudkowsky, who is one of the strongest proponents for pausing AI research due to the catastrophic impacts it might have on human civilization due to the alignment problem: “The Open Letter on AI

While we have not yet reached General AI, the current systems have reached a point where they can generate close-to-human like responses to prompts. One type of generative AI is referred to as a “large language model” (LLM, for short). As described by Chalmers: “[l]arge language models are language models using giant artificial neural networks trained on a huge amount of text data”, which can be put on transformer architecture, and generate humanlike text.⁵⁹

For anyone that has used GPT-4 or Bing, it appears that a human on the other side of the screen is responding to you. However, the models have various limitations. For example, there is a noted issue with large language models’ mathematical abilities, but it appears to develop an emergent ability to do math and arithmetic when given sufficient data.⁶⁰ The models may also lead to hallucinations — which describe instances where the model makes up people, documents, or facts. For example, GPT-4 has been documented making up authors’ names, paper names, journal names, legal cases, and other facts that do not actually exist, while providing seemingly precise descriptions of these authors or documents, despite their non-existence.⁶¹ This makes it increasingly difficult to detect when a LLM is providing accurate information, unless you are already well-trained in the field and can detect the errors.

These limitations should come as no surprise, given how the models work. The models do not “understand” the prompts being provided to them. Similarly, they do not “understand” the outputs they provide to users. Instead, as explained above, the models simply provide statistically likely outputs based on the prompts provided and based on the data they are trained on. For example, if one were to ask the model to finish the following sentence: “one, two, three, _____”, it would likely answer “four” — not because it understands that the number 4 comes after 3, or even that the letters “one”, “two” and “three” are numbers, but because it was trained on a mass amount of text that made it learn that when it sees the words “one, two, three” in a sequence, then it is most likely that it will be followed by the letters “f”, “o”, “u” and “r”. This is a different reasoning process than used by humans, despite the answer mimicking a human-like response.

One limitation of neural networks is that its creators are unable to predict the outcome of any given prompt, because we currently do not have an understanding of how the “hidden layers” of the neural network produce any given output.⁶² This is generally known as the “black box” problem or explainability problem in AI, and has led to a whole field of research known as *eXplainable AI (XAI)*.⁶³ Indeed, providing the same prompt to the same model will often elicit a different and unique response. Therefore, it is difficult for a model provider to ensure the accuracy of the responses being provided and it is difficult for a user to trust the model. While the responses

Doesn’t Go Far Enough”, (29 March 2023), online: *Time* <<https://time.com/6266923/ai-eliezer-yudkowsky-open-letter-not-enough/>>.

⁵⁹ David J Chalmers, *Could a Large Language Model be Conscious?* (arXiv, 2023) at 2 arXiv:2303.07103 [cs].

⁶⁰ Jason Wei et al, *Emergent Abilities of Large Language Models* (arXiv, 2022) at 23 arXiv:2206.07682 [cs].

⁶¹ See for e.g., Hussam Alkaiissi & Samy I McFarlane, “Artificial Hallucinations in ChatGPT: Implications in Scientific Writing” (2023) 15:2 *Cureus* e35179; Holmes, *supra* note 6 at 19; David Baidoo-Anu & Leticia Owusu Ansah, *Education in the Era of Generative Artificial Intelligence (AI): Understanding the Potential Benefits of ChatGPT in Promoting Teaching and Learning* (Rochester, NY, 2023) at 14.

⁶² Davide Castelvecchi, “Can we open the black box of AI?” (2016) 538:7623 *Nature News* 20.

⁶³ Warren J von Eschenbach, “Transparency and the Black Box Problem: Why We Do Not Trust AI” (2021) 34:4 *Philos Technol* 1607–1622 at 1608; Alejandro Barredo Arrieta et al, *Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI* (arXiv, 2019) at 4–8.

provided might be the statistically predictable text that follows the prompt, this does not necessarily mean that the response is the factually correct answer. Indeed, since many of the models available today were trained on data found on the internet, the models themselves may include data fraught with errors, inaccuracies, and pre-existing biases.⁶⁴ Despite this, the models will continue to improve over time using supervised and unsupervised learning. Put simply, we are still in the early stages, and the accuracy of the models is expected to improve rapidly.

Over time, it is expected that LLMs will become faster, more accurate, and able to accomplish more tasks outside of just providing written responses. For example, OpenAI recently released a beta of its “plugins”, which allows their model to interact with the real world — booking flights, restaurant reservations, ordering groceries, hiring humans directly to complete real-life tasks, and more.⁶⁵ For example, it is now possible to ask GPT-4 to prepare a meal plan for you, within the confines of your dietary restrictions, automatically order the food required for the meals, and have it delivered to your front-door, all with a single prompt.⁶⁶ It is also possible to prompt GPT-4 to plan and book your whole vacation itinerary based on your budget, climate preference, vacation type, and various activities and restaurant preferences.⁶⁷

Bill Gates recently released an article proclaiming that “The Age of AI has begun”. He explained that AI is one of the two revolutionary technologies released in his lifetime. In his own words:⁶⁸

The development of AI is as fundamental as the creation of the microprocessor, the personal computer, the Internet, and the mobile phone. It will change the way people work, learn, travel, get health care, and communicate with each other. Entire industries will reorient around it. Businesses will distinguish themselves by how well they use it.

This is not a unique viewpoint. Artificial Intelligence has been described as the “fourth industrial revolution”.⁶⁹ The Director-General of UNESCO has claimed that “[e]ducation will be profoundly transformed by AI...Teaching tools, ways of learning, access to knowledge, and teacher training will be revolutionized”.⁷⁰

The implications of this technology are clear, and the ramifications will be immense. While it is easy to become scared of change, it is also true that changes like this provide once-in-a-lifetime opportunities to explore new ideas, unlock new potential and build a better future. That is

⁶⁴ Min Zhang & Juntao Li, “A commentary of GPT-3 in MIT Technology Review 2021” (2021) 1:6 Fundamental Research 831–833 at 832; Holmes, *supra* note 6 at 18 and 22.

⁶⁵ “ChatGPT plugins”, online: <<https://openai.com/blog/chatgpt-plugins>>.

⁶⁶ Lauren Svensson, “Introducing the Instacart Plugin for ChatGPT”, (23 March 2023), online: *Instacart Corporate* <<https://www.instacart.com/company/uncategorized/instacart-chatgpt/>>.

⁶⁷ “ChatGPT can now connect to the internet, book holidays, and find restaurants”, (24 March 2023), online: *Yahoo Sports* <<https://ca.sports.yahoo.com/news/chatgpt-now-connect-internet-book-141611095.html>>.

⁶⁸ Bill Gates, “The Age of AI has begun”, online: *gatesnotes.com* <<https://www.gatesnotes.com/The-Age-of-AI-Has-Begun>>.

⁶⁹ Selena Nemorin et al, “AI hyped? A horizon scan of discourse on artificial intelligence in education (AIED) and development” (2023) 48:1 Learning, Media and Technology 38–51 at 38.

⁷⁰ *Ibid* at 42.

the focus of this paper – specifically, the immense potential that AI has for advancing our pedagogical effectiveness.

PART II – AI AND PEDAGOGY RESEARCH

The Current State of AIEd

While Artificial Intelligence has rapidly developed over the past decade, research into the use of AI in pedagogy remains underdeveloped. AI is often lauded as a solution to many of education’s problems,⁷¹ but it is rarely used in any meaningful way to improve learning outcomes. The education sector is generally behind other sectors in adopting AI, and as noted by Jia and Zhang, the study of using AI in some University programs is deficient.⁷²

The cross-sectional study of AI and education commenced over 30 years ago,⁷³ and is known as Artificial Intelligence in Education (“**AIEd**”).⁷⁴ AIEd is focused on improving teaching and learning processes in higher education.⁷⁵ At the end of 2022, the Council of Europe released a paper entitled *Artificial Intelligence and Education*, where they provided a meta review of the state of the art of AIEd.⁷⁶ As noted in the report, the research into AIEd has generally been taxonomized into four categories:⁷⁷

- 1. Learning with AI:** This area of research investigates using AI directly in pedagogical practices to assist students with learning or using AI systems to assist instructors with their administrative loads or teaching methods.⁷⁸
- 2. Using AI to learn about Learning:** This area of research investigates using AI to analyze the learning process to try and uncover patterns and trends about how students learn, which itself, could advance pedagogical practices.⁷⁹
- 3. Learning about AI:** This area of research is concerned with increasing knowledge about AI and how it works. This includes research into the best methods of teaching AI curriculum or even effective methods for increasing general knowledge about AI.⁸⁰

⁷¹ Holmes, *supra* note 6 at 18; Gwo-Jen Hwang et al, “Vision, challenges, roles and research issues of Artificial Intelligence in Education” (2020) 1 Computers and Education: Artificial Intelligence 100001 at 1.

⁷² Suning Jia & Xiaoxiao Zhang, “Teaching Mode of Psychology and Pedagogy in Colleges and Universities Based on Artificial Intelligence Technology” (2021) 1852:3 J Phys: Conf Ser 032033 at 2.

⁷³ Muhammad Ali Chaudhry & Emre Kazim, “Artificial Intelligence in Education (AIEd): a high-level academic and industry note 2021” (2022) 2:1 AI Ethics 157–165 at 163; Holmes, *supra* note 6 at 22; Hwang et al, *supra* note 71 at 1; *The Impact of Artificial Intelligence (AI) on Online Pedagogy - Free Online Webinar -* (2023) Art 12:50.

⁷⁴ Olga Tapalova & Nadezhda Zhiyenbayeva, “Artificial Intelligence in Education: AIEd for Personalised Learning Pathways” (2022) 20:5 Electronic Journal of e-Learning 639–653 at 639.

⁷⁵ Jose L Martin Nunez & Andres Diaz Lantada, “Artificial Intelligence Aided Engineering Education: State of the Art, Potentials and Challenges” (2020) 36:6 International Journal of Engineering Education 1740–1751 at 1740.

⁷⁶ Holmes, *supra* note 6.

⁷⁷ *Ibid* at 19; see also Ismail Celik et al, “The Promises and Challenges of Artificial Intelligence for Teachers: a Systematic Review of Research” (2022) 66:4 TechTrends 616–630 (The author here does a literature review and summarizes the various ways that AI can be used by instructors and the benefits it can have).

⁷⁸ Holmes, *supra* note 6 at 19.

⁷⁹ *Ibid*.

⁸⁰ *Ibid* at 19.

4. **Preparing for AI:** This area of research is concerned with preparing individuals for a world that will be dominated by AI. As opposed to focusing on learning about how AI systems work, this area of research is focused on informing the general population about how to use AI effectively in their work or their lives.⁸¹

The focus of this paper is the sub-field of “Learning with AI”, which has been part of the research of AIEd since the 1980s and was first theorized as a possibility in the 1930s.⁸² Learning with AI research has provided some preliminary evidence that using AI to support pedagogical practices is effective.⁸³ Indeed, a meta-analysis done by Ma in 2014 showed that at least one of the tools, intelligent tutoring systems, are effective.⁸⁴

Examples of the tools that have been studied or proposed as part of “Learning with AI” include:

- Adaptive learning tools, or intelligent tutoring systems, use machine learning algorithms to collect data about students and adapt learning paths based on students’ individual performance and advancement;⁸⁵
- Student measurement tools, that collect data about students in the classroom to draw various conclusions about their attention or interest in the subject-area and provide suggested alterations to their learning paths;⁸⁶
- Chatbots to provide support for students during their learning process, albeit with limited functionality due to the technology available at the time of study;⁸⁷
- Personalized learning tools, that use AI To provide learning support to students based on that student’s personal characteristics and learning styles;⁸⁸

⁸¹ *Ibid* at 19 and 32.

⁸² *Ibid* at 19.

⁸³ *Ibid* at 20.

⁸⁴ Wenting Ma et al, “Intelligent tutoring systems and learning outcomes: A meta-analysis.” (2014) 106:4 *Journal of Educational Psychology* 901–918 at 907–913; Holmes, *supra* note 6 at 20; see also Chaudhry & Kazim, “Artificial Intelligence in Education (AIEd)”, *supra* note 73 at 159–160.

⁸⁵ Hwang et al, *supra* note 71 at 2; see also Ma et al, “Intelligent tutoring systems and learning outcomes”, *supra* note 84; see also Fan Ouyang & Pengcheng Jiao, “Artificial intelligence in education: The three paradigms” (2021) 2 *Computers and Education: Artificial Intelligence* 100020 at 2–3; Matthew Lynch, “My Vision for the Future of Artificial Intelligence in Education”, (6 December 2018), online: *The Advocate* <<https://www.theedadvocate.org/vision-future-artificial-intelligence-education/>>; Tapalova & Zhiyenbayeva, “Artificial Intelligence in Education”, *supra* note 74 at 643.

⁸⁶ Holmes, *supra* note 6 at 21 and 25–26; Reet Kasepalu et al, “Teacher Artificial Intelligence-Supported Pedagogical Actions in Collaborative Learning Coregulation: A Wizard-of-Oz Study” (2022) 7 *Frontiers in Education*, online: <<https://www.frontiersin.org/articles/10.3389/educ.2022.736194>>; Chaudhry & Kazim, “Artificial Intelligence in Education (AIEd)”, *supra* note 73 at 157; “Artificial Intelligence in Higher Education: Applications, Promise and Perils, and Ethical Questions”, online: <<https://er.educause.edu/articles/2019/8/artificial-intelligence-in-higher-education-applications-promise-and-perils-and-ethical-questions>>; Nunez & Lantada, *supra* note 75 at 1747; Esteban Vázquez-Cano, “Artificial intelligence and education: A pedagogical challenge for the 21st century” (2021) 10 *Educational Process International Journal* at 10.

⁸⁷ Hwang et al, *supra* note 71 at 2; Holmes, *supra* note 6 at 21 and 24.

⁸⁸ Hwang et al, *supra* note 71 at 1; Holmes, *supra* note 6 at 35–36; “43 Examples of Artificial Intelligence in Education”, (7 December 2021), online: *University of San Diego Master of Science in Applied Artificial Intelligence*

- Automatic writing assessments, which uses AI tools to provide instantaneous feedback to students on their work, albeit with limited functionality due to the technology available at the time of study;⁸⁹ and
- Automating administrative tasks, which uses AI to assist instructors in automating various administrative tasks that may take away from their abilities to focus on increasing the pedagogical outcomes of their students.⁹⁰

However, as noted by the European Council, this research has often been flawed or based on limited data, leading it to conclude that “robust, independent evidence remains scarce” about the effectiveness of these systems.⁹¹ The European Council found there to be “surprisingly little to justify its wide use in well-resources classrooms” and led it to conclude that while the goals of Learning with AI are laudable, they remain mostly aspirational and have not demonstrated to be effective.⁹² Indeed, Chaudhry and Kazim concluded in a 2021 paper that “there are still a number of AI breakthroughs required to see...disruption in education at scale”.⁹³

Arguably, those disruptions are now here. The release of GPT-4, Bard, and other AI tools over the past year have unleashed new capabilities that were not possible when all these studies were conducted. Given the potential of generative AI, and the dearth of high-quality research into the intersection of pedagogical practices and generative AI, there is a significant opportunity for academics to advance this field of research.

In the following section, I provide a framework for thinking about how to advance the field of Learning AIEd to modernize the research in this area.

Proposed Learning AIEd Research Framework

To assist the direction of Learning AIEd research I propose a framework for organizing the future of this research. The purpose is to provide structure and direction to an otherwise unstructured field of research. This framework consists of three prongs — Pedagogical Enhancement Research, Pedagogical Breakthrough Research, and Administrative Research.

Pedagogical Enhancement Research is concerned with mapping existing pedagogical best-practices to the capabilities of AI. This is consistent with the approach advocated by Bower, that technology use should be connected tightly with learning theories.⁹⁴ As noted by Ouyang and Jiao, to date, only a few studies have used pedagogical theories to ground their research into AIEd.⁹⁵

<<https://onlinedegrees.sandiego.edu/artificial-intelligence-education/>>; Vázquez-Cano, “Artificial intelligence and education”, *supra* note 86 at 9.

⁸⁹ *Artificial Intelligence, Teacher Tasks and Individualized Pedagogy*, preprint, by Bruno Ferman, Lycia Lima & Flávio Riva, DOI.org (Crossref), preprint (SocArXiv, 2021); Holmes, *supra* note 6 at 23–24; note 88.

⁹⁰ Nemorin et al, “AI hyped?”, *supra* note 69 at 39.

⁹¹ Holmes, *supra* note 6 at 20.

⁹² *Ibid.*

⁹³ Chaudhry & Kazim, “Artificial Intelligence in Education (AIEd)”, *supra* note 73 at 163.

⁹⁴ Ouyang & Jiao, “Artificial intelligence in education”, *supra* note 85 at 2.

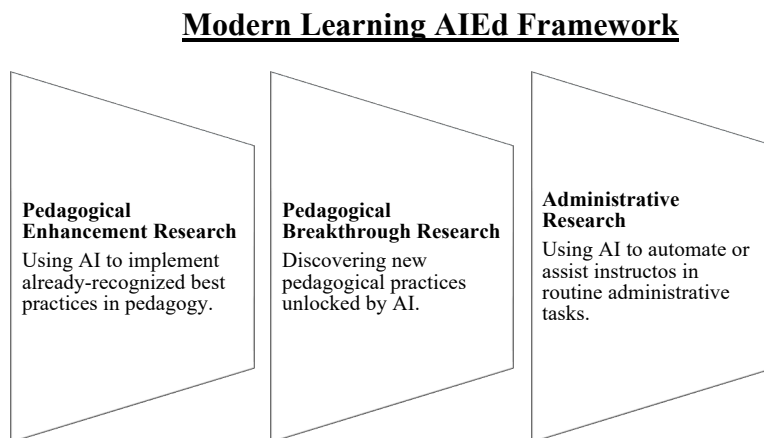
⁹⁵ *Ibid.*

The research can include everything from theoretical research (i.e., hypothesizing how AI can be used as a tool to implement or enhance already-recognized pedagogical practices) to empirical (i.e., conducting empirical studies on the effectiveness of using AI tools in pedagogical practices in comparison to traditional methods). The Pedagogical Enhancement Research contains low hanging fruit for researchers. The pedagogical practices are already well-known and studied, and researchers can attempt to recreate older empirical studies using AI tools to see if they enhance pedagogical outcomes.

Pedagogical Breakthrough Research is concerned with discovering new pedagogical practices that are unlocked by AI and were not possible in traditional learning environments. This area of research requires a deep understanding of how AI technologies work, may require training AI models on specific tasks, and will certainly require creativity and experimentation to uncover new methods of teaching that were not previously possible.

Administrative Support Research is focused on researching AI for the assistance of administering a course. This can include, for example, research into using AI to grade papers, to design courses and assessments, to detect student engagement, to detect at-risk students, to respond to e-mails, or to administer discussion boards. The goal of this research is to find areas where AI can provide significant assistance to instructors by alleviating their administrative workloads, which often takes away from their time researching, teaching or supporting students.⁹⁶ While the subject-matter of this research is not pedagogical in nature, its impact could reduce the administrative load on instructors, which in turn, could allow instructors to spend their time more focused on tasks that are likely to improve pedagogical outcomes.

I have summarized this framework as the Modern Learning AIEd Framework:



The importance of researching these topics cannot be understated. Absent any guidance from their instructors, students will use these tools as shortcuts to simplify their lives and reduce their

⁹⁶ Chaudhry & Kazim, “Artificial Intelligence in Education (AIEd)”, *supra* note 73 at 158; Baidoo-Anu & Owusu Ansah, *supra* note 61 at 12.

workload, not as tools for enhancing their learning.⁹⁷ By researching the effective uses of AI and providing guidance to students on how to use these new tools, we can harness them for their potential to improve pedagogical outcomes.

PART III – POTENTIAL OF THE MODERN LEARNING AIED FRAMEWORK

In this section, I provide various ideas, and concrete examples, of how generative AI can be used within the Modern Learning AIED Framework. The examples are focused on the teaching of business law to students in a University environment, but could easily be adapted to other contexts. The goal is to provide a non-exhaustive and exploratory overview of the potential of AI to improve pedagogy using the Modern Learning AIED Framework. More research is required to explore new ideas and empirically research the effectiveness of all ideas.

Pedagogical Enhancement Research

Pedagogical Enhancement Research is concerned with studying the effectiveness of using AI to implement already well-recognized best practices in pedagogy with AI tools. As noted by Ouyang and Jiao, “there is a need for a close combination of AI technologies with educational and learning theories”.⁹⁸ Here, I outline some of the best pedagogical practices that have been developed over the past century and show how current generative AI technologies can be used to implement or further enhance these methods.

Pedagogy has moved towards encouraging metacognition, a term coined by John. H. Flavell in 1976, which describes a process by which a student learns how to analyze their own learning and individually assess how well they have learned a concept.⁹⁹ Through using metacognition, students learn the skills to effectively learn, as opposed to simply completing course tasks without regard to how the tasks ought to impact their learning.¹⁰⁰

One of the main frameworks to encourage metacognition is Bloom’s Taxonomy and the Study Cycle, which was created by Benjamin Bloom and his colleagues in 1956 to depict various levels of educational goals, with lower-order thinking on the bottom (often associated with memorizing facts) and higher-order thinking on the top (often associated with being able to create

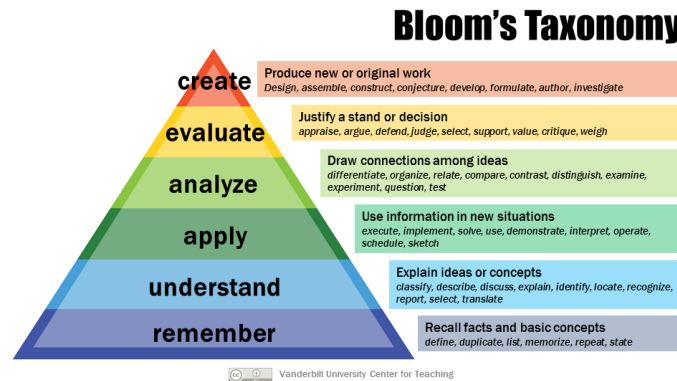
⁹⁷ This is already evidence in the way that students started using GPT-4 to write essays and assignments, and which has led to some institutions banning GPT-4: Beatrice Nolan, “Here are the schools and colleges that have banned the use of ChatGPT over plagiarism and misinformation fears”, online: *Business Insider* <<https://www.businessinsider.com/chatgpt-schools-colleges-ban-plagiarism-misinformation-education-2023-1>>; see also “Digital Pedagogy Toolbox: Let’s Make Friends with ChatGPT – BCcampus”, (22 February 2023), online: <<https://bccampus.ca/2023/02/22/digital-pedagogy-toolbox-lets-make-friends-with-chatgpt/>>.

⁹⁸ Ouyang & Jiao, “Artificial intelligence in education”, *supra* note 85 at 4.

⁹⁹ Sandra Yancy McGuire & Stephanie McGuire, *Teach Students How to Learn: Strategies You Can Incorporate into Any Course to Improve Student Metacognition, Study Skills, and Motivation*, first edition ed (Sterling, Virginia: Stylus Publishing, LLC, 2015) at 16–17.

¹⁰⁰ *Ibid* at 30.

new ideas based on foundational knowledge).¹⁰¹ Bloom’s Taxonomy was revised and modernized in 2001 and is now often depicted as the following pyramid.¹⁰²



In her book, *Teach Students How to Learn*, Dr. Saundra Yancy McGuire sets out various techniques that she employs in her teaching for encouraging metacognition and higher-order thinking.¹⁰³ I borrow from some of her examples, here, to show how AI can be used to implement and enhance these tools for students.

AI-Assisted Previewing

One concept Dr. McGuire encourages students to use is to “preview” readings, which is a process by which the student looks over the chapter before they read it, taking note of the headings, figures, bolded terms, and other basic information so that they can understand *what* they are about to read before they read it. Dr. McGuire also suggests that students produce questions that the text will be able to answer.¹⁰⁴

The previewing process is supported by research that shows that students better understand what they are reading if they have an overview of what they are about to learn, and then read by filling in the gaps of their high-level outline.¹⁰⁵ The process of pre-determining questions that can be answered in the text also creates genuine curiosity in students, which motivates them to read.¹⁰⁶

Here, AI can be used to complete this previewing and outlining process for students, so that they can understand what topics will be covered in the readings before reading them. For example, a student can open a textbook in a viewer that supports an AI model, which can interact with the

¹⁰¹ “Bloom’s Taxonomy”, online: *Vanderbilt University* <<https://cft.vanderbilt.edu/guides-sub-pages/blooms-taxonomy/>>.

¹⁰² Lorin W Anderson & David R Krathwohl, eds, *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom’s Taxonomy of Educational Objectives*, complete ed (New York: Longman, 2001).

¹⁰³ McGuire & McGuire, *supra* note 99.

¹⁰⁴ *Ibid* at 46–47.

¹⁰⁵ *Ibid* at 46.

¹⁰⁶ *Ibid* at 46–47.

text. The student can ask the AI model to outline the content they are about to read before they read it and to provide a list of questions that they will be able to answer after finishing the reading.

In Appendix I, I provide an example of using Bing to preview a chapter from an open access business law textbook. The example shows how AI can accurately preview and summarize legal topics for students, providing them with an outline of the text before they read it.

AI-Assisted Questioning and Feedback

Another concept Dr. McGuire encourages students to use to engage active learning is to solve homework problems. Dr. McGuire encourages students to answer the problems on their own, then look at solutions only after an attempted answer.¹⁰⁷ The student should then correct their mistakes against the ideal answer.¹⁰⁸ The process of attempting an answer, making mistakes, and correcting those mistakes, is an effective learning tool, and ensures that students do not make the same mistakes in the future, such as on an examination.¹⁰⁹

Dr. McGuire also encourages students to create outlines of expected examinations based on the course syllabus, readings, and slides, and to create their own practice exams from textbook questions.¹¹⁰ As noted by Dr. McGuire, “[t]here is powerful evidence demonstrating the effectiveness of testing as a way to reinforce, deepen and enrich learning”.¹¹¹

Here, AI can be used in three ways. First, AI can be used to produce unique hypothetical practice problems for students.¹¹² Second, it can be used to check a student’s answer and provide feedback on that answer.¹¹³ Third, it can be used to create mock exams based on the most important topics in an examination, which can also be graded by the AI model with feedback. This process provides a way for students to apply and practice what they learn, while also providing a direct feedback loop for the student to identify and correct their mistakes, one of the key tenants of metacognition.¹¹⁴

In Appendix II, I provide an example of using GPT-4 to create a hypothetical problem for a student and show how the model can be used to provide feedback to a student’s answer in real-time. In Appendix III, I provide an example of using GPT-4 to produce a full mock examination for a student based on their own assessment of what is important, while customizing the examination to their own interests. While not depicted, the same process of checking answers through the AI model could also be used to check the student’s answers to the hypothetical examination.

¹⁰⁷ *Ibid* at 53.

¹⁰⁸ *Ibid*.

¹⁰⁹ *Ibid* at 54.

¹¹⁰ *Ibid* at 56.

¹¹¹ *Ibid*.

¹¹² Baidoo-Anu & Owusu Ansah, *supra* note 61 at 13.

¹¹³ *Ibid*.

¹¹⁴ McGuire & McGuire, *supra* note 99 at 54.

AI-Assisted Group Studying

Another method Dr. McGuire encourages students to use is teaching concepts to friends or pretending to teach concepts to an empty audience.¹¹⁵ Dr. McGuire also suggests working in pairs or groups to understand concepts together, which Dr. McGuire explains engages the metacognitive concept of accurately judging one's own learning.¹¹⁶ As described by Ouyang and Jiao, this view of learning is “grounded upon a cognitive and social constructivism view of learning, which reflects a notion that learning occurs when a learner interacts with people, information, and technology in socially situated contexts”.¹¹⁷

AI can be used as a peer or personalized tutor for students to interact with as part of this learning process. AI can be the audience and prompt the student to teach it concepts.¹¹⁸ AI can also be used like a study group, to assist the student in their learning process. This can be particularly helpful for students who are introverted, for distance-ed students who do not have access to in-person study groups, or for students who, for any reason, do not have access to effective study groups. Early forms of this technology were known as Dialogue-Based Tutoring Systems (DTS) or Exploratory Learning Environments (ELE), but those systems were limited by the technology at the time.¹¹⁹ Hwang refers to this concept of using AI as an audience as an “Intelligent tutee”, and notes that students helping others (in this case, an AI system) understand a concept can promote “higher-order thinking competences and knowledge levels”.¹²⁰

In Appendix IV, I provide an example of using GPT-4 to mimic an audience that a student can teach. The example also shows how GPT-4 can mimic an interactive study group, whereby students take part in the teaching process, but are taught concepts they do not yet understand by the AI model itself, mimicking the give-and-take nature of a study group.

AI-Assisted Personalized Learning

Professor James Lang also published a leading book on pedagogy, entitled *Small Teaching*.¹²¹ While Lang also summarizes some of the same techniques as McGuire, he also provides more practices that instructors can use while teaching their courses. One such technique is ensuring the personalization of learning. As explained by Lang, the learning process encompasses forming new connections between existing experiences.¹²² We try to connect new concepts to things we already know, and then we try to continue to connect that concept to new things we interact with every day.¹²³ Therefore, Lang suggests that higher-order thinking can be

¹¹⁵ *Ibid* at 55.

¹¹⁶ *Ibid* at 56.

¹¹⁷ Ouyang & Jiao, “Artificial intelligence in education”, *supra* note 85 at 3.

¹¹⁸ *Ibid* at 2–3 (This is described as “AI-Directed, learner-as-recipient” and “AI-Supported, learner-as-collaborator” paradigms by Ouyang and Jiao).

¹¹⁹ *Ibid* at 3.

¹²⁰ Hwang et al, *supra* note 71 at 2.

¹²¹ James M Lang, *Small Teaching: Everyday Lessons from the Science of Learning*, second edition ed (San Francisco: Jossey-Bass, 2021).

¹²² *Ibid* at 95.

¹²³ *Ibid*.

achieved by helping students forge “rich, interconnected networks of knowledge — ones that enable each existing piece of information in our content area to connect with lots of other information, concepts, and ideas”.¹²⁴ This is consistent with the Dewey, learner-centric approach to teaching and learning.¹²⁵

Lang presents various techniques to assist students in achieving connections, including using examples, analogies, concept maps, and providing frameworks to students. However useful these tools may be, it difficult for instructors to make these connections with every student, especially in larger and more diverse classes, where each student may have very different life experiences and interests.¹²⁶

AI can assist by reframing concepts learned in class to existing concepts they understand or to areas that a student may find personally interesting. This can, for example, compare a new concept with another concept the student already understands, provide an analogy or example that connects with something the student already knows or deeply cares about, or reframe the same concept in language that might be more familiar to a student. This concept is like the “intelligent tutor” described by Hwang, the effectiveness of which has been demonstrated in several meta-analyses.¹²⁷

In Appendix V, I provide an example of using GPT-4 to personalize the learning experience for a student. I use GPT-4 to explain a concept using a subject-area that is already interesting to a specific student and to compare the concept to a concept the student already understands. I then ask GPT-4 to relate the concept to a student’s own life circumstances. Each of these examples show how AI can effectively be used as a tool for students to personalize their learning and form deeper connections between the concepts learned in class and their lives.

Summary

In this section, I have outlined some concrete examples of how AI can be used to implement and enhance already-recognized best practices in pedagogy. This section was intended to be introductory and exploratory. More work is required to explore the full potential of AI in this respect. Further research also needs to be done to determine the effectiveness of using AI over traditional implementations of these methods, considering that these new methods are likely to be more widely available and accessible to students than other traditional methods, and taking into account the shortcomings of AI models.

Pedagogical Breakthrough Research

Pedagogical Breakthrough Research is concerned with studying AI tools and attempting to unlock new pedagogical practices or learning outcomes that were not possible without AI. In this

¹²⁴ *Ibid* at 96.

¹²⁵ Holmes, *supra* note 6 at 35–37; see also Chaudhry & Kazim, “Artificial Intelligence in Education (AIED)”, *supra* note 73 at 159.

¹²⁶ Lang, *supra* note 121 at 99–111.

¹²⁷ Hwang et al, *supra* note 71 at 2; see also note 97; see also Nunez & Lantada, *supra* note 75 at 1743–1744.

section, I outline some of the learning outcomes and pedagogical practices that may now exist due to the advancements in AI and suggest that further research be put into this creative pursuit.

Prompt Engineering

One of the first new areas worth exploring is prompt engineering. Prompt engineering describes the art behind asking questions to AI models.¹²⁸ Prompt engineering is becoming popular, since being able to craft the proper prompts can have big differences on the output provided by an AI model. Indeed, “prompt engineer” is a new job title that is commanding six-figure salaries in some companies,¹²⁹ given its potential for large returns.

To draw an analogy, prompt engineering is like the skills taught in modern legal research courses and seminars. Modern legal research courses teach students how to properly search electronic legal databases such as Westlaw and Quick Law using Boolean search connectors (for e.g., IF, AND, OR, wildcat, and other Boolean connectors). It is essential to understand how Boolean searches work to be a competent legal researcher.

The same is true with AI. The output of an AI model is only as good as its input. The more detail provided in a prompt, the more likely the result will match what the user is looking for. It is not a stretch to suggest that existing legal research tools will be replaced or augmented by AI soon. Therefore, curriculums will need to be adapted to include prompt engineering in the learning outcomes.

Augmented Reading

Another potential use for AI is what I have coined “augmented reading”. With a traditional textbook or article, students are unable to interact with the material. Instead, textbooks provide a one-way learning process, whereby the information on the page is transferred to the student’s memory through the process of reading and taking notes on the readings. Any questions, gaps in understanding, or application of the principles learned, is then reserved for activities outside of the textbook, such as searching through other sources materials for an answer, asking a professor or colleague for support, or completing practice questions and reviewing them with a colleague or a professor.

Modern AI systems, however, are capable of interacting with electronic textbooks and articles. Students can open the textbook or article in an AI-assisted browser (like Microsoft’s Edge), highlight the text they want to interact with, and then prompt the model, through a process I have called “augmented reading”. These prompts can include summarizing key information, rewording an explanation to make it easier to digest, providing an example from recent world events about the topic, and more.

¹²⁸ See for e.g., a discussion of prompt engineering and a catalogue of prompts for GPT-4 in Jules White et al, *A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT* (arXiv, 2023) arXiv:2302.11382 [cs].

¹²⁹ “How to Get a Six-Figure Job as an AI Prompt Engineer”, (14 April 2023), online: *Time* <<https://time.com/6272103/ai-prompt-engineer-job/>>.

In Appendix VI, I provide an example of augmented reading, where I simulate a situation where a student might ask a specific question about a paragraph in a textbook to further understand the topic. The augmented reading process has potentially significant pedagogical value since it forces an active reading process and can personalize otherwise generic text to ensure that any student can understand the concepts being conveyed through personalized examples, real-life events, or even by rewording the text in a way that is more digestible to the student.

AI-Based Assessments

Another potential use for AI is in the assessment process. Some instructors have created new assessment techniques that use AI as part of the process for assessment. Specifically, a new form of AI critique assessment has been suggested, whereby students prompt an AI model to answer a question or produce an essay on a topic, and the assignment requires a student to critique the answer, find its flaws, and improve its work product.¹³⁰ Various Universities across Canada have provided guidance on how to use AI in the classroom, and suggested innovative assessments using AI at their core.¹³¹

Summary

There is significant potential in creatively using AI in pedagogical practices to unlock new ways of learning information. I have provided an example of some of the creative uses that AI unlocks, but there is likely to be significantly more and effective practices that will be developed in the future. More research is needed on unlocking new pedagogical practices using AI, including theorizing about its capabilities as well as empirically testing theories to determine whether they have the potential to increase pedagogical outcomes.

Administrative Research

Administrative Research is concerned with using AI to automate and alleviate the administrative load on instructors, so that they can spend more of their time on pedagogical practices, which will have the indirect consequence of increasing student performance.

There are numerous ways that AI can be used to alleviate the administrative load of instructors. The table, below, outlines some of the most common tasks instructors undertake when designing or delivering a course, each of which can potentially be aided or automated with AI:

¹³⁰ See for e.g., “ChatGPT and Generative AI in the Classroom – Office of the Vice-Provost, Innovations in Undergraduate Education”, online: <<https://www.viceprovostundergrad.utoronto.ca/strategic-priorities/digital-learning/special-initiative-artificial-intelligence/>>.

¹³¹ See for e.g., *ibid*; see also, note 97; see also, “AI in the Classroom - Concordia University”, online: <<https://www.concordia.ca/content/concordia/en/ctl/digital-teaching/ai-in-the-classroom.html>>.

Common Course Administration Tasks

Course Design	Course Delivery
Curriculum Design	Lecturing
Syllabus Creation	Student Communication
Assessment Creation	Marking and Feedback

Course Design Assistance

AI can be used by instructors during the design of a new course, or in the process of incrementally improving a course every semester. For example, AI can be used to create a first draft of a syllabus that covers various topics, to suggest readings on any given topic, define learning outcomes for the course, and even assist in drafting an examination. The purpose is not to shirk the instructor's responsibility, but to use AI as a creative tool for producing first drafts of these documents or providing specific suggestions that can be used by the instructor in an iterative process.

In Appendix III, I show how GPT-4 can be used to create draft examinations. The instructor can prompt GPT-4 on the specific style of the examination, the types of questions to ask, the topics to be covered in the examination, and can even prompt a model about an inspiration for a hypothetical fact pattern.

In Appendix VII, I provided an example of how GPT-4 can be used to create a first draft of a syllabus for a Business Law course. An instructor could provide more details to the AI (for e.g., textbook choice, additional topics, expected learning outcomes, and more). The speed at which GPT-4 can create a first draft of a syllabus could significantly decrease the administrative time spent by instructors on the task.

Course Delivery Assistance

In addition to assisting with the design of a course, AI can also assist with the administration of a course.¹³²

First, AI that can assist with lectures. For example, existing AI tools can take lecture notes and automatically prepare pleasing PowerPoint presentations to use during lectures.¹³³ Similarly, AI can be used to automatically transcribe lectures and provide detailed summaries and take-aways for students immediately after the class is over.¹³⁴

¹³² Nemorin et al, "AI hyped?", *supra* note 69 at 39.

¹³³ "Introducing Microsoft 365 Copilot – your copilot for work", (16 March 2023), online: *The Official Microsoft Blog* <<https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/>>.

¹³⁴ Barbara Krasnoff, "Find the best AI-powered app to transcribe your audio", (24 August 2022), online: *The Verge* <<https://www.theverge.com/23316220/transcription-ai-otter-temi-how-to>>.

Second, AI can assist with communicating with students. For example, AI can be used to prepare first drafts of responses to student questions. With some additional design, an AI chatbot can be implemented on a course website to automatically answer student questions about the administration or content of the course – including when the midterm or exam is, what the readings are assigned for the week, when any breaks or cancelled classes are, answers to common questions, and much more.

Third, AI can assist in marking. An AI model can be prompted to read an examination question, a model answer, a grading rubric, and then be prompted to mark a student’s assignment with feedback. AI-assisted marking tools have already been created to provide qualitative evaluations of student work.¹³⁵ Empirical studies on some of these systems have found that there was no discernable difference between human-assisted marking and AI-assisted marking.¹³⁶

In Appendix VIII, I provide an example of how an AI model would mark a hypothetical University student’s response to an actual midterm question used in a business law class. This shows how, with the correct prompt and information, an AI model can be used to provide appropriate grading and detailed feedback to students on their performance.

Fourth, there are some more inventive tools being developed to assist with student monitoring. For example, AI tools have been created to monitor students during class to provide real-time feedback to professors about the effectiveness of their pedagogical approaches. Similarly, AI tools have been used to monitor student data from online learning platforms and to provide automatic early-warning signs about students who may be struggling in the course.¹³⁷

While these examples may seem fantastical today, many of these tools will be standard functionality in the software we already use day to day. Whether your University relies on Microsoft 365 or Google Enterprise, both companies have announced the launch of built-in AI tools in their whole suite of software, and both are expected to roll out publicly later this year.¹³⁸ For example, it will be a standard feature in Outlook and Gmail to have an AI assistant when drafting e-mails.¹³⁹ Similarly, it will be a standard feature in PowerPoint and Google Slides to create full presentations for you based on pre-existing content.¹⁴⁰ The same features will also be available in Microsoft Word or Google Docs for producing first drafts of new papers or documents.¹⁴¹ All of these tools will be implement in our day-to-day lives, and there is significant potential for reducing administrative loads.

¹³⁵ Chaudhry & Kazim, “Artificial Intelligence in Education (AIED)”, *supra* note 73 at 159.

¹³⁶ Ferman, Lima & Riva, *supra* note 89; Nunez & Lantada, *supra* note 75 at 1743.

¹³⁷ Holmes, *supra* note 6 at 21 and 25–26; Kasepalu et al, “Teacher Artificial Intelligence-Supported Pedagogical Actions in Collaborative Learning Coregulation”, *supra* note 86; Chaudhry & Kazim, “Artificial Intelligence in Education (AIED)”, *supra* note 73 at 157; *supra* note 86; Nunez & Lantada, *supra* note 75 at 1747; Vázquez-Cano, “Artificial intelligence and education”, *supra* note 86 at 10.

¹³⁸ “Announcing new generative AI experiences in Google Workspace”, online: *Google Workspace Blog* <<https://workspace.google.com/blog/product-announcements/generative-ai>>; note 133.

¹³⁹ *Supra* note 133; *supra* note 138.

¹⁴⁰ *Supra* note 138; *supra* note 133.

¹⁴¹ *Supra* note 138; *supra* note 133.

Potential Issues with AI in Pedagogy

While these use-cases are certainly exciting, there are likely to be critics opposed to using AI in pedagogical practices — in part due to the accuracy of the current models and in part due to the ethical or privacy issues these tools create. Furthermore, as AI’s capabilities closely approach the capabilities of most humans, this raises existential issues of meaning, including the role of Universities and University Professors.¹⁴² These are legitimate concerns, each of which I will address.

With respect to the accuracy of the models, this is something that will improve with time and should not inhibit AIED Learning Research in the short-term. Indeed, as Voltaire stated, “perfect is the enemy of good”. In this paper, I am not advocating for the immediate release of these technologies into the University environment at scale. Instead, this paper is intended to explore the pedagogical possibilities of AI while setting the stage for additional research. In any event, most of the pedagogical practices proposed in this paper provide alternatives to practices that are, themselves, already error-prone or simply not accessible to students. For example, study-groups often consist of students in the same course, and therefore prone to inaccurate information being disseminated. In any event, many students may not have access to effective study groups, meaning that an AI model mimicking a study group, with some errors, might nonetheless be better than nothing.

There are also numerous ethical and privacy issues with using AI in a University environment. As noted by Nemorin, there are currently six prominent proposals for ethical AI use, known as the 5P framework – (1) the *Asilomar AI Principles*; (2) the *Montreal Declaration for Responsible AI*; (3) the General Principles in the second version of the *Ethically Aligned Design: A Vision for Prioritizing Human Well-Being with Autonomous and Intelligence Systems*; (4) the Ethical Principles in the European Commission’s European Group on Ethics in Science and New Technologies report entitled *Artificial Intelligence, Robotics and ‘Autonomous Systems’*; (5) the five principles from the House of Lords Artificial Intelligence Committee’s report entitled *AI in the UK: Ready, willing and able?*; and (6) the *Tenets of Partnership on AI*.¹⁴³ Together, these yield 45 total principles,¹⁴⁴ and can be grouped into five main areas: (1) transparency; (2) justice and fairness; (3) non-maleficence; (4) responsibility; and (5) privacy.¹⁴⁵ These should be regarded when designing or research AI systems in pedagogical practices.

One specific ethical issue raised by one of the suggestions in this paper – AI assisted marking – is the procedural fairness to students if AI is used in marking papers. For example, during the pandemic, the UK used a predictive AI model to extrapolate student scores on courses they were unable to finish due to school disruptions, which raised lawsuits and significant concerns about the model’s accuracy and procedural fairness rights.¹⁴⁶ Since most AI models suffer

¹⁴² Holmes, *supra* note 6 at 23.

¹⁴³ Nemorin et al, “AI hyped?”, *supra* note 69 at 46.

¹⁴⁴ *Ibid.*

¹⁴⁵ Holmes, *supra* note 6 at 42.

¹⁴⁶ “‘F**k the algorithm’?: What the world can learn from the UK’s A-level grading fiasco”, (26 August 2020), online: *Impact of Social Sciences* <<https://blogs.lse.ac.uk/impactofsocialsciences/2020/08/26/fk-the-algorithm-what-the-world-can-learn-from-the-uks-a-level-grading-fiasco/>>.

from the “black box” problem, it is difficult for people to provide explanations to those affected by the model about the judgment that went into preparing the result.¹⁴⁷ It is also difficult to ensure the accuracy and consistency of AI grading, which could lead to grade appeals.

However, similar issues arise with using markers and Teaching Assistants for grading. The same methods of verifying grades across markers could be used, including statistically significant sampling by the instructor. Furthermore, the AI models can be prompted to explain the grade and provide meaningful feedback to the student, which could end up *increasing* the perception of procedural fairness, in comparison to the minimal feedback that is possible on exams in larger section classes. Indeed, one empirical study showed that AI-assisted decision making resulted in students believing that decisions was more objective and ranked AI higher on both procedural and distributive fairness.¹⁴⁸

Interestingly, an adjacent issue recently arose at the Federal Court of Canada. An applicant brought a judicial review of an Immigration Officer’s decision. The applicant alleged that the decision of the Immigration Officer was written using artificial intelligence, and therefore raised issues of procedural fairness. The Federal Court of Canada rejected this view, holding (albeit arguably only in *obiter*) that so long as the decision was reasonable, the fact that it is written by artificial intelligence is irrelevant to procedural fairness issues under *Vavilov*.¹⁴⁹

[24] As to artificial intelligence, the Applicant submits the Decision is based on artificial intelligence generated by Microsoft in the form of “Chinook” software. However, the evidence is that the Decision was made by a Visa Officer and not by software. I agree the Decision had input assembled by artificial intelligence, but it seems to me the Court on judicial review is to look at the record and the Decision and determine its reasonableness in accordance with *Vavilov*. Whether a decision is reasonable or unreasonable will determine if it is upheld or set aside, whether or not artificial intelligence was used. To hold otherwise would elevate process over substance.

Therefore, similar reasoning can be used during grade appeals. The use of AI marking assistance, in and of itself, is not a procedural fairness issue, if the substance of the grade is accurate, the model is consistent, and the reasons the model provides for the grade are supported. Guardrails can be put in place to ensure this, including a clear appeal process, transparency with students, and verification of the model’s output by statistical sampling.

With respect to the more existential issues – the role of Universities and University Professors during the AI-revolution – we can take comfort in the universality of the problem this presents to all working professionals. As reported recently by the New York Times,¹⁵⁰ two recent

¹⁴⁷ Castelvechi, *supra* note 62; von Eschenbach, “Transparency and the Black Box Problem”, *supra* note 63.

¹⁴⁸ Holmes, *supra* note 6 at 68.

¹⁴⁹ *Haghshenas v Canada (Citizenship and Immigration)*, 2023 FC 464 at para 23.

¹⁵⁰ Steve Lohr, “A.I. Is Coming for Lawyers, Again”, *The New York Times* (10 April 2023), online: <<https://www.nytimes.com/2023/04/10/technology/ai-is-coming-for-lawyers-again.html>>.

reports placed law as the most likely industry to be displaced by AI. One of the reports estimated that approximately 44% of legal work could be automated by AI.¹⁵¹

A methodology for predicting the exposure of a profession to AI, known as the AI Occupational Exposure methodology (“AIOE”) estimates the impact of AI on an occupation and industry.¹⁵² A recent study placed postsecondary law teachers as the fifth highest exposed profession amongst many other postsecondary educators taking four of the top five spots, just behind telemarketers.¹⁵³ The industry with the most exposure to being displaced by AI was the legal services industry, with business schools ranking in thirteenth place.¹⁵⁴

However, it is important to remember that the AIOE methodology is agnostic as to the effect on the industry – meaning that the effect can be positive or negative – depending on whether AI displaces workers (negative) or augments their work to make it easier and more effective (positive).¹⁵⁵ I tend to fall into the camp of people who believe that both lawyers, and university professors, will continue to play important roles in advocating for their clients and teaching students, respectively. I believe, as demonstrated in this paper, that AI will simply augment many of our tasks to reduce administrative work loads, increase efficiencies in research and writing, and enhance pedagogical practices. While these studies show that AI will have a significant impact on both the education and legal industries, there is a significant opportunity to research AIED Learning to ensure that the impact is positive and enhances the pedagogical outcomes of students.

CONCLUSION

There is a significant opportunity to improve AIED Learning research and to implement AI into pedagogical practices in higher education.¹⁵⁶ The technology has now advanced to a point where the tools could be implemented into the classroom with careful planning and monitoring. Students are already using these tools, but do not know how to yield them to their advantage.

This paper presents a framework for thinking about the future of AIED. There is significant research and work that needs to go into advancing this area, which will become increasingly important in all education, particularly legal and business education, which are some of the highest industries and occupations that will be impacted by AI. The AI models will continue to improve, and the tools they provide to us will become increasingly advanced. Students already have public access to these tools and will be using them. By providing instruction to our students on how to properly use AI in a higher education environment to increase their pedagogical outcomes, our institutions can become touchstones of how to embrace and benefit from AI.

In GPT-4’s own words:

¹⁵¹ *The Potentially Large Effects of Artificial Intelligence on Economic Growth*, by Jan Hatzius et al, Zotero (Goldman Sachs, 2023).

¹⁵² Ed Felten, Manav Raj & Robert Seamans, *How will Language Modelers like ChatGPT Affect Occupations and Industries?* (arXiv, 2023) at 3 arXiv:2303.01157 [econ, q-fin].

¹⁵³ *Ibid* at 14.

¹⁵⁴ *Ibid* at 15.

¹⁵⁵ *Ibid* at 3.

¹⁵⁶ Celik et al, “The Promises and Challenges of Artificial Intelligence for Teachers”, *supra* note 77 at 617.

In a post-AI world, my advice to educators would be to embrace the potential of AI as a complementary tool for enhancing teaching and learning experiences. Focus on developing a symbiotic relationship between human educators and AI, leveraging the strengths of both parties. As AI excels in handling repetitive tasks, data analysis, and content generation, utilize it to create personalized and adaptive learning experiences for students. At the same time, educators should emphasize their irreplaceable role in fostering critical thinking, creativity, empathy, and ethical decision-making, which are essential human qualities that AI cannot fully replicate. By integrating AI into pedagogy and adopting a growth mindset, educators can empower themselves and their students to thrive in an increasingly technology-driven world.

I could not have said it better myself.

BIBLIOGRAPHY

- “F**k the algorithm?": What the world can learn from the UK's A-level grading fiasco", (26 August 2020), online: *Impact of Social Sciences* <<https://blogs.lse.ac.uk/impactofsocialsciences/2020/08/26/fk-the-algorithm-what-the-world-can-learn-from-the-uks-a-level-grading-fiasco/>>.
- “43 Examples of Artificial Intelligence in Education”, (7 December 2021), online: *University of San Diego Master of Science in Applied Artificial Intelligence* <<https://onlinedegrees.sandiego.edu/artificial-intelligence-education/>>.
- “Action potentials and synapses”, (22 November 2016), online: *Queensland Brain Institute* <<https://qbi.uq.edu.au/brain-basics/brain/brain-physiology/action-potentials-and-synapses>>.
- “AI & Robotics”, online: *Tesla* <<https://www.tesla.com/AI>>.
- “AI in the Classroom - Concordia University”, online: <<https://www.concordia.ca/content/concordia/en/ctl/digital-teaching/ai-in-the-classroom.html>>.
- “An On-device Deep Neural Network for Face Detection”, online: *Apple Machine Learning Research* <<https://machinelearning.apple.com/research/face-detection>>.
- “Announcing new generative AI experiences in Google Workspace”, online: *Google Workspace Blog* <<https://workspace.google.com/blog/product-announcements/generative-ai>>.
- “Artificial Intelligence in Higher Education: Applications, Promise and Perils, and Ethical Questions”, online: <<https://er.educause.edu/articles/2019/8/artificial-intelligence-in-higher-education-applications-promise-and-perils-and-ethical-questions>>.
- “Bloom's Taxonomy”, online: *Vanderbilt University* <<https://cft.vanderbilt.edu/guides-sub-pages/blooms-taxonomy/>>.
- “Business Law and the Legal Environment”, (2012), online: *Open Textbook Library* <<https://open.umn.edu/opentextbooks/textbooks/273>>.
- “ChatGPT and Generative AI in the Classroom – Office of the Vice-Provost, Innovations in Undergraduate Education”, online: <<https://www.viceprovostundergrad.utoronto.ca/strategic-priorities/digital-learning/special-initiative-artificial-intelligence/>>.
- “ChatGPT can now connect to the internet, book holidays, and find restaurants”, (24 March 2023), online: *Yahoo Sports* <<https://ca.sports.yahoo.com/news/chatgpt-now-connect-internet-book-141611095.html>>.
- “ChatGPT plugins”, online: <<https://openai.com/blog/chatgpt-plugins>>.
- “Deep Learning for Siri's Voice: On-device Deep Mixture Density Networks for Hybrid Unit Selection Synthesis”, online: *Apple Machine Learning Research* <<https://machinelearning.apple.com/research/siri-voices>>.

“Digital Pedagogy Toolbox: Let’s Make Friends with ChatGPT – BCcampus”, (22 February 2023), online: <<https://bccampus.ca/2023/02/22/digital-pedagogy-toolbox-lets-make-friends-with-chatgpt/>>.

“Explained: Neural networks”, (14 April 2017), online: *MIT News | Massachusetts Institute of Technology* <<https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>>.

“Forward Propagation, Backward Propagation, and Computational Graphs”, online: *Dive Into Deep Learning* <https://d2l.ai/chapter_multilayer-perceptrons/backprop.html>.

“How AI is powering a more helpful Google”, (15 October 2020), online: *Google* <<https://blog.google/products/search/search-on/>>.

“How to Get a Six-Figure Job as an AI Prompt Engineer”, (14 April 2023), online: *Time* <<https://time.com/6272103/ai-prompt-engineer-job/>>.

“IBM100 - A Computer Called Watson”, (7 March 2012), online: <<http://www-03.ibm.com/ibm/history/ibm100/us/en/icons/watson/>>.

“IBM100 - Deep Blue”, (7 March 2012), online: <<http://www-03.ibm.com/ibm/history/ibm100/us/en/icons/deepblue/>>.

“Introducing Microsoft 365 Copilot – your copilot for work”, (16 March 2023), online: *The Official Microsoft Blog* <<https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/>>.

“Machine Learning for Beginners: An Introduction to Neural Networks - victorzhou.com”, online: <<https://victorzhou.com/blog/intro-to-neural-networks/>>.

“Pause Giant AI Experiments: An Open Letter”, online: *Future of Life Institute* <<https://futureoflife.org/open-letter/pause-giant-ai-experiments/>>.

“Supervised vs. Unsupervised Learning: What’s the Difference?”, (15 November 2022), online: <<https://www.ibm.com/cloud/blog/supervised-vs-unsupervised-learning>>.

“The Open Letter on AI Doesn’t Go Far Enough”, (29 March 2023), online: *Time* <<https://time.com/6266923/ai-eliezer-yudkowsky-open-letter-not-enough/>>.

“Timeline of a breakthrough”, online: <<https://www.deepmind.com/research/highlighted-research/alphafold/timeline-of-a-breakthrough>>.

“What is Strong AI? | IBM”, online: <<https://www.ibm.com/topics/strong-ai>>.

Ahire, Jayesh, “The Artificial Neural Networks handbook: Part 1 - DataScienceCentral.com”, (24 August 2018), online: *Data Science Central* <<https://www.datasciencecentral.com/the-artificial-neural-networks-handbook-part-1/>>.

Arrieta, Alejandro Barredo et al, *Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI* (arXiv, 2019).

- Baidoo-Anu, David & Leticia Owusu Ansah, *Education in the Era of Generative Artificial Intelligence (AI): Understanding the Potential Benefits of ChatGPT in Promoting Teaching and Learning* (Rochester, NY, 2023).
- Beranek, Bolt, H A Simon & P McCorduck, *History of Artificial Intelligence* (Cambridge, Massachusetts).
- Bradeško, Luka & Dunja Mladenčić, “A Survey of Chabot Systems through a Loebner Prize Competition” (2012).
- Bubeck, Sébastien et al, *Sparks of Artificial General Intelligence: Early experiments with GPT-4* (arXiv, 2023).
- Castelvecchi, Davide, “Can we open the black box of AI?” (2016) 538:7623 Nature News 20.
- Celik, Ismail et al, “The Promises and Challenges of Artificial Intelligence for Teachers: a Systematic Review of Research” (2022) 66:4 TechTrends 616–630.
- Chalmers, David J, *Could a Large Language Model be Conscious?* (arXiv, 2023).
- Chaudhry, Muhammad Ali & Emre Kazim, “Artificial Intelligence in Education (AIED): a high-level academic and industry note 2021” (2022) 2:1 AI Ethics 157–165.
- Eschenbach, Warren J von, “Transparency and the Black Box Problem: Why We Do Not Trust AI” (2021) 34:4 Philos Technol 1607–1622.
- Felten, Ed, Manav Raj & Robert Seamans, *How will Language Modelers like ChatGPT Affect Occupations and Industries?* (arXiv, 2023).
- Ferman, Bruno, Lycia Lima & Flávio Riva, *Artificial Intelligence, Teacher Tasks and Individualized Pedagogy*, preprint, by Bruno Ferman, Lycia Lima & Flávio Riva, DOI.org (Crossref), preprint (SocArXiv, 2021).
- Floridi, Luciano, “AI and Its New Winter: from Myths to Realities” (2020) 33:1 Philos Technol 1–3.
- Gates, Bill, “The Age of AI has begun”, online: *gatesnotes.com* <<https://www.gatesnotes.com/The-Age-of-AI-Has-Begun>>.
- Haenlein, Michael & Andreas Kaplan, “A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence” (2019) 61:4 California Management Review 5–14.
- Haghshenas v Canada (Citizenship and Immigration)*, 2023 FC 464.
- Han, Su-Hyun et al, “Artificial Neural Network: Understanding the Basic Concepts without Mathematics” (2018) 17:3 Dement Neurocogn Disord 83–89.

- Hatzius, Jan et al, *The Potentially Large Effects of Artificial Intelligence on Economic Growth*, by Jan Hatzius et al, Zotero (Goldman Sachs, 2023).
- Hendler, James, “Avoiding Another AI Winter” (2008) 23:2 IEEE Intelligent Systems 2–4.
- Holmes, Wayne, *Artificial Intelligence and Education: A Critical View Through the Lens of Human Rights, Democracy and the Rule of Law*, by Wayne Holmes (Strasbourg: Council of Europe, 2022).
- Hwang, Gwo-Jen et al, “Vision, challenges, roles and research issues of Artificial Intelligence in Education” (2020) 1 Computers and Education: Artificial Intelligence 100001.
- Jia, Suning & Xiaoxiao Zhang, “Teaching Mode of Psychology and Pedagogy in Colleges and Universities Based on Artificial Intelligence Technology” (2021) 1852:3 J Phys: Conf Ser 032033.
- Kasepalu, Reet et al, “Teacher Artificial Intelligence-Supported Pedagogical Actions in Collaborative Learning Coregulation: A Wizard-of-Oz Study” (2022) 7 Frontiers in Education, online: <<https://www.frontiersin.org/articles/10.3389/educ.2022.736194>>.
- Krasnoff, Barbara, “Find the best AI-powered app to transcribe your audio”, (24 August 2022), online: *The Verge* <<https://www.theverge.com/23316220/transcription-ai-otter-temi-how-to>>.
- Krogh, Anders, “What are artificial neural networks?” (2008) 26:2 Nat Biotechnol 195–197.
- Lang, James M, *Small Teaching: Everyday Lessons from the Science of Learning*, second edition ed (San Francisco: Jossey-Bass, 2021).
- Lighthill, Sir James, *Artificial Intelligence: A General Survey*, by Sir James Lighthill (1972).
- Lohr, Steve, “A.I. Is Coming for Lawyers, Again”, *The New York Times* (10 April 2023), online: <<https://www.nytimes.com/2023/04/10/technology/ai-is-coming-for-lawyers-again.html>>.
- Lynch, Matthew, “My Vision for the Future of Artificial Intelligence in Education”, (6 December 2018), online: *The Edvocate* <<https://www.theedadvocate.org/vision-future-artificial-intelligence-education/>>.
- Lyon, Peter, “Honda’s New Safety Target Focuses On AI-Powered Technology”, online: *Forbes* <<https://www.forbes.com/sites/peterlyon/2021/11/26/hondas-new-safety-target-focuses-on-ai-powered-technology/>>.
- Ma, Wenting et al, “Intelligent tutoring systems and learning outcomes: A meta-analysis.” (2014) 106:4 Journal of Educational Psychology 901–918.
- McGuire, Sandra Yancy & Stephanie McGuire, *Teach Students How to Learn: Strategies You Can Incorporate into Any Course to Improve Student Metacognition, Study Skills, and Motivation*, first edition ed (Sterling, Virginia: Stylus Publishing, LLC, 2015).

McInerney, James et al, *Explore, exploit, and explain: personalizing explainable recommendations with bandits* (Vancouver British Columbia Canada: ACM, 2018).

Nemorin, Selena et al, “AI hyped? A horizon scan of discourse on artificial intelligence in education (AIED) and development” (2023) 48:1 *Learning, Media and Technology* 38–51.

Nielsen, Michael A, *Neural Networks and Deep Learning* (Determination Press, 2015).

Alkaissi, Hussam & Samy I McFarlane, “Artificial Hallucinations in ChatGPT: Implications in Scientific Writing” (2023) 15:2 *Cureus* e35179.

Nolan, Beatrice, “Here are the schools and colleges that have banned the use of ChatGPT over plagiarism and misinformation fears”, online: *Business Insider* <<https://www.businessinsider.com/chatgpt-schools-colleges-ban-plagiarism-misinformation-education-2023-1>>.

Nunez, Jose L Martin & Andres Diaz Lantada, “Artificial Intelligence Aided Engineering Education: State of the Art, Potentials and Challenges” (2020) 36:6 *International Journal of Engineering Education* 1740–1751.

Ouyang, Fan & Pengcheng Jiao, “Artificial intelligence in education: The three paradigms” (2021) 2 *Computers and Education: Artificial Intelligence* 100020.

Pardeshi, Siddhi et al, “A survey on Different Algorithms used in Chatbot” (2020) 07:05 *IRJET* 6092–6098.

Pereda, Alberto E, “Electrical synapses and their functional interactions with chemical synapses” (2014) 15:4 *Nat Rev Neurosci* 250–263.

Sarker, Iqbal H, “Deep Learning: A Comprehensive Overview on Techniques, Taxonomy, Applications and Research Directions” (2021) 2:6 *SN COMPUT SCI* 420.

Svensson, Lauren, “Introducing the Instacart Plugin for ChatGPT”, (23 March 2023), online: *Instacart Corporate* <<https://www.instacart.com/company/uncategorized/instacart-chatgpt/>>.

Tapalova, Olga & Nadezhda Zhiyenbayeva, “Artificial Intelligence in Education: AIED for Personalised Learning Pathways” (2022) 20:5 *Electronic Journal of e-Learning* 639–653.

The Impact of Artificial Intelligence (AI) on Online Pedagogy - Free Online Webinar - (2023).

Turing, Alan M, “Computing Machinery and Intelligence” (1950) LIX:236 *Mind* 433–460.

Vaswani, Ashish et al, *Attention is All you Need* (Curran Associates, Inc., 2017).

Vázquez-Cano, Esteban, “Artificial intelligence and education: A pedagogical challenge for the 21st century” (2021) 10 *Educational Process International Journal*.

Vincent, James, “Former Go champion beaten by DeepMind retires after declaring AI invincible”, (27 November 2019), online: *The Verge* <<https://www.theverge.com/2019/11/27/20985260/ai-go-alphago-lee-se-dol-retired-deepmind-defeat>>.

Wei, Jason et al, *Emergent Abilities of Large Language Models* (arXiv, 2022).

White, Jules et al, *A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT* (arXiv, 2023).

Zhang, Min & Juntao Li, “A commentary of GPT-3 in MIT Technology Review 2021” (2021) 1:6 Fundamental Research 831–833.

APPENDIX I: BING PRODUCED PREVIEWING

In this example, I simulate how a student may preview a textbook using AI. To accomplish this, I opened chapter 11 of a textbook, *Business Law and the Legal Environment*,¹⁵⁷ in the Microsoft Edge browser, which has Bing's GPT-4 powered AI built in. This chapter introduced the concept of consideration in American contract law. I then asked Microsoft's Bing to preview the text for me before I read it, to provide a list of questions I should be able to answer after reading it, and to summarize the main concepts I should take away from the readings. This example shows how AI can be used to instantaneously preview readings for students.

Prompt:

Can you create a list of the topics that will be covered in this chapter. For each topic on the list, provide a sub-point with a series of what, why, how questions that I should be able to answer after reading each topic.

Bing Response:

Sure! Here's a list of topics that will be covered in this chapter along with sub-points that will help you understand each topic:

1. General Perspectives on Consideration

- What is "consideration" in contract law?
- What purposes does the doctrine serve?
- How does the law determine whether consideration exists?
- What are the elements of consideration?

2. Legal Sufficiency

- What does "legal sufficiency" mean when examining consideration?
- How does the concept operate in such common situations as threat of litigation and accord and satisfaction?
- Why are illusory promises unenforceable, and how do courts deal with needs, outputs, and exclusive dealings contracts?

I hope this helps! Let me know if you have any other questions.

¹⁵⁷ "Business Law and the Legal Environment", (2012), online: *Open Textbook Library* <<https://open.umn.edu/opentextbooks/textbooks/273>>.

Prompt:

Can you summarize the main takeaways from this page into five bullet points? Extract the key concepts I would need to know as a business student studying business law.

Bing Response:

Sure! Here are the five key takeaways from this page:

- Consideration is a required element of a contract. It is the bargained-for giving up of something of legal value for something in return.
- Consideration serves the purposes of making formal the intention to contract and reducing rash promise making.
- The “something” that is promised or delivered cannot be just anything, such as a feeling of pride, warmth, amusement, or friendship; it must be something known as a legal detriment—an act, forbearance, or a promise of such from the promisee.
- Legal consideration need not even be certain; it can be a promise contingent on an event that may never happen.
- Courts are not generally concerned with the economic adequacy of the consideration but instead with whether it is present.

I hope this helps! Let me know if you have any other questions.

APPENDIX II: GPT-4 ASSISTED QUESTIONING, FEEDBACK AND GRADING

In this example, I simulate how a student can use AI to create questions about specific legal topics and provide feedback on a student's answer. To accomplish this, I asked GPT-4 to produce a practice problem so that the student can test their knowledge on a specific topic. I then created a student-like answer, with some purposeful inaccuracies and generalizations, and asked GPT-4 to provide feedback on the answer and tips for improvement. This example shows how GPT-4 can be used by students to test their knowledge and receive immediate feedback on their performance.

Prompt:

Can you ask me a hypothetical-based question, similar to a law school exam, about the oppression remedy?

GPT-4 Response:

Samantha is a minority shareholder in XYZ Corporation, which is a privately held company that owns and operates a chain of high-end coffee shops. The majority shareholder, Richard, owns 70% of the shares, while Samantha owns 20% and various other minority shareholders own the remaining 10%. Richard also serves as the company's CEO and sits on the board of directors.

Richard recently decided to expand the company's operations by acquiring a small tea company that he personally owns. In order to accomplish this, he negotiated a deal where XYZ Corporation would purchase the tea company for \$1 million, which is significantly higher than the tea company's fair market value. Samantha was not consulted or informed about the transaction, and she only found out about it after the deal had already been completed. She suspects that Richard used XYZ Corporation's resources to enrich himself at the expense of minority shareholders like her.

Samantha is considering bringing a lawsuit against Richard under the oppression remedy. Discuss the potential success of her case based on the following:

1. Explain the key elements of the oppression remedy that Samantha would need to establish in order to succeed in her claim.
2. Analyze the facts to determine if Samantha can make a strong case that Richard's actions satisfy the elements of an oppression claim.
3. Assess possible defenses Richard may raise and their likelihood of success.
4. If Samantha is successful, discuss the possible remedies that a court may order to address the oppression she has experienced.

Prompt:

Please Evaluate my answer to the question:

The issue in this case is whether Samantha, as a minority shareholder, would have a successful claim in oppression against Richard, the CEO and director of the company.

To succeed in an oppression remedy claim in Ontario, Samantha must show that: (1) she formed reasonable expectations when she invested in the company; and (2) those reasonable expectations were unfairly disregarded, or were disregarded in a way that was unduly prejudicial or unfair.

Here, it is likely that Samantha would not succeed in her claim in oppression. Richard made the decision to purchase XYZ corporation, which is a corporation he himself owns, without regard to Samantha's interests as a minority shareholder. Since he never sought the input of Samantha or considered her interests, these interests were unduly disregarded.

However, Richard may argue that the claim is not properly pled as oppression, and instead, Samantha ought to bring a derivative action against Richard, since the duties Richard owes are to the corporation and not to Samantha. It is likely that Richard is correct - his decision to invest in XYZ Corporation was a breach of his fiduciary duties owed to the corporation, which are duties not owed to Samantha. Therefore, she requires leave of the court to bring a derivative action.

Therefore, it is unlikely that Samantha would succeed in her claim, and instead, she should consider bringing a derivative action for Richard breaching his fiduciary duties owed to the corporation.

GPT-4 Response:

Your answer provides a good analysis of the issue at hand and the potential success of Samantha's oppression claim. However, there are a few areas that can be improved upon or clarified:

1. You've correctly identified the elements of an oppression claim, but it is essential to emphasize that Samantha must show that Richard's actions were oppressive, unfairly prejudicial, or unfairly disregarded her interests.
2. It is important to discuss the reasonable expectations that Samantha might have had as a minority shareholder. For example, she could have expected to be informed about significant transactions, such as the acquisition of the tea company, and to have her interests considered by the majority shareholder and the board of directors.
3. While you correctly mentioned that Richard might argue that Samantha should bring a derivative action instead, it is essential to note that oppression claims and derivative actions serve different purposes. An oppression remedy aims to protect individual shareholders' interests, while a derivative action seeks to remedy wrongs done to the corporation. In some cases, both actions might be appropriate, depending on the facts.
4. To strengthen your conclusion, you can elaborate on the specific reasons why it might be unlikely that Samantha would succeed in her oppression claim. For instance, a court may find that Richard's actions, while arguably breaching his fiduciary duties, do not rise to the

level of oppression, unfairly prejudicial conduct, or unfairly disregarding Samantha's interests.

In conclusion, your answer demonstrates a solid understanding of oppression claims and derivative actions. With a few adjustments and added clarity, your analysis would be even stronger.

APPENDIX III: GPT-4 ASSISTED EXAM CREATION

In this example, I simulate how a student may use GPT-4 to produce a mock examination while they are studying for a midterm or final examination. I asked GPT-4 to produce a mock examination, with specific details on the style of the examination, the topics it should cover, and the theme of the examination, to make it personalized to my interests. This example shows how students can use GPT-4 to create practice exams for themselves when studying for exams.

Prompt:

I am a student in a second year undergraduate university student in a BBA program. Pretend you are a professor teaching me business law and produce a law-school like hypothetical exam with 5 questions, spanning the following topics: (1) contract law; (2) torts; (3) secured transactions; (4) agency law; (5) minority shareholder protections. Make the hypothetical scenario based on my most recent favourite TV Show, Succession. Make sure that the questions are pertinent to Canadian law - specifically in Ontario. All questions should prompt me to advise the client about their legal situation and the whole exam should read as one coherent story, with questions spliced in the middle of the story as a legal issue comes up.

GPT-4 Response:

[Exam Instructions: Read the following fact pattern carefully and answer the questions that follow. Your answers should be based on Ontario law, as applicable. Be sure to provide a detailed analysis and cite any relevant statutes or case law to support your arguments. Identify the legal subject area in each question and provide counter-arguments where possible.]

Fact Pattern:

Logan Roy, the founder and CEO of Waystar Royco, a multinational media and entertainment conglomerate, has decided to retire. He has four children - Kendall, Siobhan, Roman, and Connor - who are all interested in taking over the family business. As the patriarch of the family and the majority shareholder of Waystar Royco, Logan wants to ensure a smooth transition of power while also protecting his own interests and those of the minority shareholders.

Before Logan announces his retirement, he decides to enter into a consulting agreement with Kendall. The agreement states that Kendall will provide strategic advice to Waystar Royco for a period of two years in exchange for a \$2 million annual fee. The agreement also includes a clause that allows Logan to terminate the agreement without penalty if he decides not to retire within the first three months. Shortly after signing the agreement, Logan has a change of heart and decides not to retire. Kendall, already having left his previous job, seeks to enforce the consulting agreement.

Question 1:

Advise Kendall on his legal position and the enforceability of the consulting agreement, taking into account the termination clause.

As part of the transition process, Logan asks Siobhan to oversee a major acquisition for the company. However, she leaks confidential information about the target company to the press, causing the acquisition to fall through and leading to a substantial loss for Waystar Royco. The target company sues Waystar Royco for the tort of inducing a breach of confidence. Siobhan claims that she did not intentionally leak the information and that it was an honest mistake.

Question 2:

Advise Waystar Royco on its potential liability, considering Siobhan's claim that the leak was unintentional.

In an attempt to raise capital for a new project, Roman enters into a loan agreement with a bank on behalf of Waystar Royco. To secure the loan, he pledges the company's assets as collateral without the knowledge or consent of Logan or the board of directors. The bank later discovers that Roman did not have the proper authority to pledge the assets. However, the bank had previously dealt with Roman on multiple occasions and assumed he had the authority based on their past dealings.

Question 3:

Advise the bank on the validity and enforceability of the security interest in the assets of Waystar Royco, considering their past dealings with Roman.

Connor, the least experienced of the siblings, takes on the role of negotiating with a supplier for a new contract. Unbeknownst to Logan, Connor promises the supplier a significant bonus if the supplier can deliver the goods ahead of schedule. The supplier meets the deadline and demands the bonus. Logan refuses to pay, arguing that Connor did not have the authority to make such a promise. The supplier claims that Waystar Royco has always paid bonuses for early deliveries in the past, creating an expectation of a bonus.

Question 4:

Advise the supplier on its legal position and whether it can enforce the bonus payment, taking into account the company's past practices.

During the transition period, Logan continues to make decisions that disproportionately benefit him and his children, while disregarding the interests of the minority shareholders. Frustrated, a group of minority shareholders band together to challenge Logan's actions. Among their complaints, they claim that Logan has engaged in self-dealing by awarding lucrative contracts to businesses owned by his children without considering alternative options.

Question 5:

Advise the minority shareholders on the protections available to them under Ontario law and the potential remedies they may seek for the alleged oppressive conduct, considering the allegations of self-dealing.

APPENDIX IV: GPT-4 ASSISTED STUDY GROUP

In this example, I simulate a situation where a student is using GPT-4 as an artificial peer to teach it a concept and to be taught a concept interactively. The purpose was to try and mimic a study-group style experience using AI. I purposefully provided GPT-4 with half-answers and simulated a situation where a student might not know the answer to a question. This example shows how GPT-4 can be used by students to mimic group studying techniques.

Prompt:

Pretend you are a second year university student, you are in a business program at a Canadian university. You are taking a course called “business law”. You are trying to learn about a concept known as a “derivative action”. Ask me questions until you fully understand the concept.

GPT-4 Response:

Sure! Let's start with the basics. Can you please explain what a derivative action is in the context of business law?

Prompt:

I think it is a way for a shareholder to sue the directors

GPT-4 Response:

I see. Can you please explain the circumstances under which a shareholder might choose to initiate a derivative action against the directors of a company?

Prompt:

Well, the directors might breach a duty they owe to the corporation, like a fiduciary duty or a duty of care.

GPT-4 Response:

That makes sense. In a derivative action, does the shareholder sue the directors directly, or do they sue on behalf of the corporation? And what is the rationale behind this approach?

Prompt:

I'm not sure, do you know?

GPT-4 Response:

Yes, in a derivative action, the shareholder sues on behalf of the corporation, not directly against the directors. The rationale behind this is that the alleged wrongdoing by the directors has caused harm to the corporation itself, and as a result, any recovery or remedy obtained from the lawsuit would benefit the corporation and its shareholders indirectly, rather than benefiting the initiating shareholder directly.

Prompt:

Got it. So they sue on behalf of the corporation. Ask me more.

GPT-4 Response:

Sure. What is the difference between a derivative action and a direct action brought by a shareholder? Are there specific instances when one is more appropriate than the other?

Prompt:

As you said, a derivative action is brought on behalf of the corporation, to allow the shareholder to bring an action that only the corporation has standing to bring. A shareholder, however, may have its own action, for example an oppression remedy action, when their own personal interests have been disregarded in a way that is oppressive or unfairly prejudicial. A personal action will be more appropriate when the shareholder's personal interests have been affected, whereas a derivative action is more appropriate when the corporation's interests have been affected.

GPT-4 Response:

That's a helpful distinction. Can you explain the procedural requirements for initiating a derivative action? Are there any specific steps a shareholder must take before they can bring a derivative action on behalf of the corporation?

Prompt:

Yes there are three requirements: first, proving that the board failed to authorize the corporation to bring the action; second, showing that the shareholder is acting in good faith; and third, showing that the action would be in the best interests of the corporation

GPT-4 Response:

Thank you for explaining the requirements. It seems that these steps help to ensure that derivative actions are not misused or initiated without merit. Are there any potential consequences or risks for shareholders who decide to pursue a derivative action? For example, are there any costs or liabilities they may face as a result of initiating such an action?

Prompt:

I'm not sure, do you know?

GPT-4 Response:

Yes, there are potential consequences and risks for shareholders who pursue a derivative action. One of the main risks is the cost associated with litigation, including legal fees and expenses. In some cases, if the shareholder loses the lawsuit, they might be required to pay the legal costs of the other party, which can be substantial.

Another risk is the potential damage to the shareholder's reputation and relationship with the corporation and its other shareholders. Engaging in litigation might cause friction between the parties involved, and could potentially impact future business dealings and investments.

To mitigate some of these risks, courts may require the shareholder to provide security for costs before allowing the derivative action to proceed. This is meant to ensure that the corporation's resources are not unduly burdened by the costs of the litigation, and to discourage shareholders from initiating frivolous or baseless lawsuits.

Prompt:

Interesting. But can't the court order the legal costs to be paid by the corporation itself? I remember reading that somewhere, do you remember any cases like that? Or maybe it was in the OBCA?

GPT-4 Response:

Yes, you're correct. Under certain circumstances, the court can order the corporation to pay the legal costs of the shareholder who initiates the derivative action. This is known as an indemnity for costs. The rationale behind this is to ensure that shareholders are not discouraged from initiating a legitimate derivative action due to the fear of financial burden.

In the Ontario Business Corporations Act (OBCA), section 246 specifically deals with derivative actions and provides for the possibility of the court ordering the corporation to pay the reasonable legal costs of the shareholder who brings the action. The court may do this if it believes it is appropriate given the circumstances of the case.

It's important to note that the indemnity for costs is not automatic, and the court will consider various factors, such as the outcome of the lawsuit, the conduct of the parties involved, and whether the action was truly in the best interest of the corporation, before deciding whether to order the corporation to pay the shareholder's legal costs.

Prompt:

Phew, thought I was wrong there. Any other questions?

GPT-4 Response:

Yes, let's discuss remedies. If a shareholder successfully pursues a derivative action, what types of remedies or outcomes can the court award to address the harm caused to the corporation? Are there any specific examples of remedies that may be granted in a derivative action lawsuit?

Prompt:

Good question. The actual remedy for the application for a derivative action is leave from the court to bring a claim in the name of the corporation. The remedy for the ultimate lawsuit will depend on the nature of the claim.

GPT-4 Response:

Thank you for the clarification. So, the main remedy in the derivative action itself is the court's permission to bring the claim on behalf of the corporation. Once that permission is granted and the lawsuit proceeds, the remedies available in the subsequent lawsuit will depend on the specific claim and the nature of the alleged wrongdoing. Some possible remedies could include compensatory damages, an injunction to prevent further harm, or an order to rescind a particular transaction, among others.

Now that we have a good understanding of derivative actions, is there anything else you'd like me to ask or clarify about the topic? Or would you like to move on to another aspect of business law?

Prompt:

We only discussed bringing a derivative action to get permission to sue directors, can you use it to sue anyone else?

GPT-4 Response:

Yes, you can use a derivative action to sue not only directors but also other parties who may be responsible for causing harm to the corporation. This can include officers, employees, and even third parties such as contractors or business partners who may have breached a duty owed to the corporation or engaged in wrongful conduct that caused the corporation to suffer losses.

The rationale behind allowing shareholders to bring a derivative action against various parties is to ensure that those who cause harm to the corporation are held accountable, even when the corporation's management is unwilling or unable to take action themselves. The main objective is to protect the corporation's interests and seek remedies that will benefit the corporation and its shareholders as a whole.

Prompt:

Interesting. I think I get it. What's the next topic you would suggest discussing?

GPT-4 Response:

I'm glad you feel comfortable with the concept of derivative actions. Since we've been discussing shareholder rights and remedies in the context of business law, another relevant topic you might find interesting is the concept of "oppression remedies." This would allow you to explore situations where minority shareholders feel that their rights have been unfairly disregarded or oppressed by the majority shareholders or the company's management, and the legal remedies available to them.

Alternatively, you could also explore the topic of corporate governance, which covers the rules, practices, and processes that dictate how a corporation is directed and controlled. This could include the roles and responsibilities of directors, officers, and shareholders, as well as the principles of transparency, accountability, and ethical conduct in the management of a corporation.

Feel free to let me know which topic you'd like to discuss or if you have any other areas of interest within business law that you'd like to explore.

APPENDIX V: GPT-4 ASSISTED PERSONALIZED LEARNING

In this example, I simulate a student who wants to learn a new legal principle (in this case, frustration) by connecting it to his or her personal life. I first ask GPT-4 to relate frustration to a concept the student learned recently – mistake. I then asked GPT-4 to explain frustration using marvel movies, to help relate it to an interest the student might already have. Lastly, I asked GPT-4 to apply the concept to an event that may have occurred in the student's real life to try and connect it with the student's real-world experiences. All these examples could assist a student in forming strong connections to these new concepts, increasing a deeper understanding of the subject-matter.

Prompt:

I am an undergraduate student and have no prior understanding of law. Can you help me understand a concept? Pretend you are a law school professor teaching contract law to me. Let's start with the concept I am trying to understand. I am reading about a new concept called "frustration". I am confused how this is different than a "mistake in the existence of a subject matter" which we learned about last week. Can you explain the difference between the two?

GPT-4 Response:

Of course! As a law school professor, I'd be happy to help you understand the difference between the concepts of "frustration" and "mistake in the existence of a subject matter" in contract law. Both concepts are related to situations where the performance of a contract becomes impossible or radically different from what was originally agreed upon, but they differ in their underlying circumstances and legal consequences.

1. Frustration:

Frustration is a concept in contract law that occurs when an unforeseen event happens after the formation of a contract, which makes it impossible or extremely burdensome for one or both parties to fulfill their obligations under the contract. This unforeseen event must be so significant that it essentially changes the nature of the contract, making the original agreement fundamentally different from what the parties intended.

Frustration can result in the discharge of a contract, which means that the parties are no longer bound by the contract's terms and are relieved of their obligations to perform. However, the doctrine of frustration is applied narrowly by courts, as it is considered an exception to the general rule that parties should be held to their bargains.

Some examples of events that may lead to frustration include:

- The destruction of the subject matter of the contract (e.g., a building that was to be leased burns down)
- The death or incapacity of a person necessary for the performance of the contract (e.g., a famous singer who was contracted to perform at a concert)

- The occurrence of an event that makes performance illegal (e.g., a new law that prohibits the sale of a particular product)

2. Mistake in the Existence of a Subject Matter:

A mistake in the existence of a subject matter is a situation where, at the time the contract was formed, one or both parties were mistaken about the existence or identity of the subject matter of the contract. In other words, the mistake concerns a fundamental aspect of the contract and affects the very basis upon which the parties entered into the agreement.

If a mistake in the existence of a subject matter is established, the contract may be considered void ab initio (i.e., void from the beginning). This means that, in the eyes of the law, the contract never existed, and the parties are not bound by its terms.

Examples of a mistake in the existence of a subject matter include:

- A buyer and seller agree to a contract for the sale of a specific painting, but the painting was destroyed before the contract was formed, and neither party was aware of this fact.
- Two parties enter into a contract for the sale of a specific parcel of land, but both are mistaken about the actual location of the land.

In summary, the key differences between frustration and mistake in the existence of a subject matter are:

- Frustration concerns an unforeseen event that occurs after the formation of the contract, while mistake in the existence of a subject matter concerns a fundamental mistake at the time the contract was formed.
- Frustration may result in the discharge of a contract, while a mistake in the existence of a subject matter may result in the contract being considered void ab initio.

Prompt:

I think I get it, but to help me understand it even more, can you provide me an example of frustration using a marvel movie?

GPT-4 Response:

Certainly! Let's use an example involving a hypothetical Marvel movie premiere event to illustrate the concept of frustration.

Suppose a local movie theater enters into a contract with an event management company to host an exclusive premiere event for a new Marvel movie. The event is supposed to include a live panel discussion with the movie's main cast members, followed by the screening of the film. The contract specifies that the event will take place at the movie theater on a particular date and that the event management company is responsible for organizing the panel discussion and ensuring the cast members' attendance.

However, a week before the scheduled event, an unprecedented natural disaster occurs in the city where the premiere is supposed to take place. As a result, the city is declared a disaster area, and public gatherings are temporarily banned by the authorities to ensure safety and facilitate rescue and recovery efforts.

In this scenario, the contract between the movie theater and the event management company may be considered frustrated due to the unforeseen natural disaster and the subsequent ban on public gatherings. This unforeseen event has made it impossible for the parties to fulfill their obligations under the contract, as the performance of the contract would now be illegal.

The doctrine of frustration would likely apply here, resulting in the discharge of the contract, which means both parties would be relieved of their obligations to perform under the contract. It's important to note that frustration is applied narrowly, and in this case, it is applied because the event that caused the frustration was both unforeseeable and beyond the control of both parties.

Prompt:

Interesting. Last week, I agreed to sell someone my laptop on Facebook Marketplace, but then I found out that the laptop I was selling could be sold for a higher price, so I no longer wanted to give it to them. Would that be an example of frustration or not?

GPT-4 Response:

No, the situation you described does not qualify as an example of frustration under contract law. Frustration occurs when an unforeseen event arises after the formation of a contract, making it impossible or extremely burdensome for one or both parties to fulfill their obligations under the contract. The unforeseen event must be so significant that it essentially changes the nature of the contract, making the original agreement fundamentally different from what the parties intended.

In your case, the fact that you discovered you could sell your laptop for a higher price after agreeing to sell it to someone on Facebook Marketplace does not constitute an unforeseen event that makes the performance of the contract impossible or extremely burdensome. Your change of mind due to the possibility of obtaining a higher price is not a valid ground for invoking the doctrine of frustration.

As a result, you would still be bound by the agreement to sell your laptop to the buyer at the originally agreed-upon price, unless you can negotiate a new agreement with the buyer or find another valid legal reason to terminate the contract. It's important to remember that contracts are legally binding, and parties should generally be held to their bargains, barring exceptional circumstances.

APPENDIX VI: BING ASSISTED AUGMENTED READING

In this example, I simulate a situation where a student wants to learn more about a concept discussed in a textbook, but is having difficulty understanding the textbook as written. To accomplish this, I opened chapter 43 of a textbook, *Business Law and the Legal Environment*,¹⁵⁸ in the Microsoft Edge browser.¹⁵⁹ This chapter discusses the characteristics of a corporation. I picked a paragraph that a student may find confusing:

In comparing partnerships and corporations, there is one additional factor that ordinarily tips the balance in favor of incorporating: the corporation is a legal entity in its own right, one that can provide a “veil” that protects its shareholders from personal liability.

I then augmented the reading by having a simulated discussion about that paragraph with Bing.

Prompt:

Can you provide an example of why this is important, using my favourite TV show the Simpsons?

Bing Response:

Sure! Let’s say that Homer Simpson and his friends Barney, Lenny, and Carl decide to start a business selling their own brand of Duff beer. They have two options: they can form a partnership or incorporate their business. If they form a partnership, each partner is personally liable for the debts and obligations of the business. This means that if the business is sued or goes bankrupt, the partners’ personal assets (such as their homes and cars) could be at risk.

On the other hand, if they incorporate their business and form a corporation, the corporation becomes a separate legal entity. This means that the shareholders (Homer, Barney, Lenny, and Carl) are generally not personally liable for the debts and obligations of the corporation. If the corporation is sued or goes bankrupt, the shareholders’ personal assets are usually protected by the “corporate veil.” This limited liability is one of the main reasons why many business owners choose to incorporate their businesses. Is there anything else you would like to know?

Prompt:

Can you provide me a recent real-life example from the news where this was discussed?

¹⁵⁸ <https://open.umn.edu/opentextbooks/textbooks/273>

¹⁵⁹

Bing Response:

Yes! A recent example of a case where the corporate veil was discussed is the 2019 Court of Appeal case of *Rossendale Borough Council v Hurstwood Properties Ltd.* The case focused on payments of business rates and the use of special purpose vehicle (SPV) companies to avoid rate liability. The court ultimately upheld the resilience of the corporate veil. Another example is a 2021 Florida Court of Appeals case, *Segal v. Forastero, Inc.*, where an investor who bought and sold real estate used an LLC to enter into a purchase agreement and then walked away from the deal, claiming that he had no personal liability because it was the LLC and not him who signed the agreement. The court ultimately found that none of the elements needed to show that the company was an alter ego of Segal or could be pierced existed. These are just a few examples of how courts have addressed issues related to the corporate veil. Is there anything else you would like to know?

Prompt:

Interesting, what are the factors that the court considered in Segal for whether to pierce the corporate veil?

Bing Response:

In the case of *Segal v. Forastero, Inc.*, the Florida Court of Appeals considered whether any of the three elements needed to show that the company was an alter ego of Segal or could be pierced existed. These elements are typically: (1) the shareholder dominated and controlled the corporation to such an extent that the corporation's independent existence was in fact non-existent and the shareholders were in fact alter egos of the corporation; (2) the corporate form must have been used fraudulently or for an improper purpose; and (3) the fraudulent or improper use of the corporate form caused injury to the claimant. The court ultimately found that none of these elements applied in this case. Is there anything else you would like to know?

Prompt:

Do you have any recent examples where the court did pierce the corporate veil?

Bing Response:

Yes, there are several recent examples where courts have pierced the corporate veil. One such example is the case of *Inter Export LLC v Townley*, where the court found that the corporate veil could be pierced so that directors could be personally liable for fraudulent misrepresentation. Another example is *6071376 Canada Inc. v. 3966305 Canada Inc. et.al.*, where a plaintiff was able to pierce the corporate veil in an attempt to collect on a breach of contract judgment. These cases demonstrate that while the corporate veil provides significant protection for shareholders, it is not absolute and can be pierced in certain circumstances. Is there anything else you would like to know?

Prompt:

Why did the court pierce the corporate veil in 6071376 Canada?

Bing Response:

In the case of 6071376 Canada Inc. v. 3966305 Canada Inc. et.al., the court pierced the corporate veil because it found that the corporation was being used to commit fraud. In this case, 6071376 Canada Inc. had put up 40% of the money to buy a building in partnership with 3966305 Canada Inc. However, 3966305 Canada Inc. sold the building without informing 6071376 Canada Inc. and did not provide them with their share of the proceeds. The court found that this constituted fraud and pierced the corporate veil to hold the shareholders of 3966305 Canada Inc. personally liable. Is there anything else you would like to know?

APPENDIX VII: GPT-4 PRODUCED SYLLABUS FOR A BUSINESS LAW COURSE

In this example, I simulate an instructor using AI to assist it in designing the syllabus for a course. While the instructor is unlikely to use the syllabus provided, and the suggested readings may consist of hallucinations, it nonetheless provides a good outline that the instructor can then amend independently or with further prompts to the AI model. This example shows how AI can be used by instructors to reduce the amount of time it takes to complete administrative tasks.

Prompt:

Create a syllabus for a second year course in a business program entitled "Business Law", which covers Ontario law, and the topics of: torts, contracts, business associations, secured transactions, agency law, franchise law, intellectual property law and conflict of laws. The course should take place over 12 weeks. Include all areas that will be covered in the course, and recommended readings for each subject.

GPT-4 Response:

- **Course Title:** Business Law
- **Course Level:** Second Year
- **Duration:** 12 Weeks
- **Location:** Ontario, Canada

Course Description:

This course provides a comprehensive overview of business law in Ontario, covering key legal topics such as torts, contracts, business associations, secured transactions, agency law, franchise law, intellectual property law, and conflict of laws. The course aims to equip students with a strong understanding of the legal framework governing businesses in Ontario, enabling them to make informed decisions and minimize legal risks in their professional pursuits.

Week-by-Week Breakdown:

Week 1: Introduction to Business Law

- Overview of the Canadian legal system
- The role of law in business
- Sources of law in Ontario

Recommended Readings:

1. Fridman, G. H. L. (2017). The Law of Torts in Canada. Thomson Reuters. (Chapters 1-2)

Week 2: Torts

- Negligence
- Strict liability
- Product liability
- Defamation

Recommended Readings:

1. Fridman, G. H. L. (2017). The Law of Torts in Canada. Thomson Reuters. (Chapters 3-7)

Week 3: Contracts

- Formation of contracts
- Consideration and capacity
- Terms and conditions
- Remedies for breach of contract

Recommended Readings:

1. Waddams, S. M. (2020). The Law of Contracts. Thomson Reuters. (Chapters 1-4)

Week 4: Business Associations

- Sole proprietorships
- Partnerships
- Corporations
- Limited liability companies (LLCs)

Recommended Readings:

1. O'Brien, J. (2016). Understanding Canadian Business Structures: Law, Tax, and Finance. Thomson Reuters. (Chapters 1-3)

Week 5: Secured Transactions

- Personal Property Security Act (PPSA)
- Creation and perfection of security interests
- Priorities and enforcement

Recommended Readings:

1. Wood, R. (2016). Law of Secured Transactions. LexisNexis. (Chapters 1-3)

Week 6: Agency Law

- Agency relationships
- Authority of agents
- Rights and duties of agents and principals
- Termination of agency

Recommended Readings:

1. Conklin, W. (2018). The Law of Agency. Irwin Law. (Chapters 1-3)

Week 7: Franchise Law

- Overview of franchise law in Ontario
- Disclosure requirements

- Relationship between franchisor and franchisee
- Dispute resolution in franchise relationships

Recommended Readings:

1. Levitt, B., & Kestenberg, J. (2019). Canadian Franchise Law Handbook. Carswell. (Chapters 1-4)

Week 8: Intellectual Property Law

- Copyright
- Trademarks
- Patents
- Trade secrets

Recommended Readings:

1. Hughes, J., & Siebrasse, N. (2017). Intellectual Property Law in Canada. LexisNexis. (Chapters 1-5)

Week 9: Conflict of Laws

- Jurisdiction
- Choice of law
- Recognition and enforcement of foreign judgments

Recommended Readings:

1. Pitel, S. G. A., & Rafferty, N. (2018). Private International Law in Common Law Canada. LexisNexis. (Chapters 1-4)

Week 10: Case Studies and Practical Applications

- Applying legal concepts to real-world scenarios
- Developing strategies for addressing legal issues in business

Recommended Readings:

1. Case studies and articles provided by the instructor

Week 11: Legal Ethics and Professional Responsibility

- Lawyers' ethical duties and responsibilities
- Conflict of interest rules
- Confidentiality and privilege
- The role of law societies in regulating the legal profession

Recommended Readings:

1. Woolley, A., & Cotter, B. P. (2020). Lawyers' Ethics and Professional Regulation. LexisNexis. (Chapters 1-3)

Week 12: Review and Exam Preparation

- Recap of key concepts from the course
- Exam tips and strategies
- Mock exam and discussion

Recommended Readings:

1. Review materials and study guides provided by the instructor

Assessment:

- **Class participation:** 10%
- **Midterm exam:** 30%
- **Group project (case study analysis):** 20%
- **Final exam:** 40%

Course Objectives: Upon completion of this course, students will be able to:

1. Understand the legal framework governing businesses in Ontario.
2. Identify and analyze legal issues related to torts, contracts, business associations, secured transactions, agency law, franchise law, intellectual property law, and conflict of laws.
3. Apply legal concepts to real-world business scenarios and develop strategies for addressing legal issues in business.
4. Recognize ethical considerations and professional responsibilities in the context of business law.
5. Develop critical thinking, problem-solving, and communication skills necessary for success in the business world.

APPENDIX VIII: GPT-4 ASSISTED MARKING ASSISTANCE

In this example, I simulate a situation where an instructor may use AI to assist it in marking student exams. I provided GPT-4 with a midterm question in a business law course, clear instructions on how to mark the answer, and also provided it with a model answer to the question. I then provided GPT-4 with a hypothetical student answer, that was purposefully fraught with errors and brief. This example shows how AI models can be used to assist with grading and/or providing students with detailed feedback on their answers.

Prompt:

I know that you cannot provide grades on student work, but I wanted to try and test how you would perform in marking an assignment. To that end, I would like you to pretend you are a University professor. You are teaching undergraduate business students, who are taking their first introductory course in law, entitled "Business Law". You have recently administered an exam. Generally, we grade as follows:

(a) The “conclusion” does not matter. The student can conclude either way, so long as their arguments are supported, well-argued and the conclusion can be reasonably inferred from their arguments.

(b) Students may think of “unique” ways to approach a problem, that are different from the answer key, below. To the extent that the student’s “unique” approach is legally correct, and applicable to the question, then the student can still do well. Please use your judgment.

(c) Generally, a B answer states the right law, and applies the law at a high level. An A answer will (a) have very few or no mistakes; (b) will have a strong application section, which includes counterarguments and acknowledgement of the various sub-issues; and (c) does not miss anything essential. A C answer will have significant mistakes and/or will not provide significant details when applying the law to the facts. A D answer or below will generally misstate the law completely, apply the wrong law, or have very superficial levels of application. Very rarely should you give a mark of 0 (unless the student does not answer or writes things that are the equivalent to a non-answer).

(d) Below, we have provided outlines of the expected answers. We do not expect students’ answers to be this detailed given the time constraints under which they are writing, and because the exam is closed book. To this end, if students do not use the exact correct words (in a test for example), they can still do well, so long as they explain the law correctly (even in their own words). We do not expect students will restate the tests exactly, since it is closed book. Furthermore, we do not expect students to cite cases or specific pieces of legislation. Instead, they need only understand the concepts.

(e) We expect the average will be around low to the mid 70s, but it will depend on the students’ performance this term.

We expect the grades to be in percentage points and to provide detailed feedback to the student so that they can improve their performance and know where they went wrong.

The Question asked of students was the following:

Joel is a single father who lives with his daughter, Sarah, a 12-year-old who just started grade 7 at the local middle school. Joel is a general contractor who assists with various building and renovation contracts around the city.

Joel and Sarah are getting ready for work. In the background, the news anchors on the television are discussing the breaking news. A fungal infection, caused by the cordyceps mushroom, appears to have mutated in Southeast Asia, allowing it to infect humans. They state that there is no concern for people living in North America, and that only a few cases have been reported in Asia.

Joel picks up the phone. It is a potential customer who wants to have a deck built in their backyard.

Graydon: Hey Joel, this is Tommy's son, Graydon.

Joel: Hey Graydon, happy belated 15th birthday!

Graydon: I would like to surprise my dad by having you build a deck in our backyard for him. I have been saving up money for the last 5 years to get this done.

Joel: Great, I would be happy to help. I am pretty booked up for the next few weeks, but could get started in July if that works for you?

Graydon: July sounds great. How much is this going to cost me?

Joel: Well, it depends on a lot of factors, but we can figure something out. It will cost you between \$2,000 and \$3,500, depending on the type of wood and finishings you choose. I'll take 10% off the price, as well, if you pay cash, so that I don't have to report this to the tax man. Let's meet next week to go over the details.

Graydon: Great, we have a deal! See you next week.

Joel: Sounds good Graydon. We have a deal!

Joel drops Sarah off at school. The radio echoes around the car. "The first reported infection has been found in North America – right here in the city. Wait...this just in, the hospital is now reporting 10 infections."

Q1. Joel is in his own head. He is thinking about the contract he just agreed to do for Graydon. He hastily agreed to the project but realized that he forgot that he had a very lucrative renovation project happening in July. Advise Joel whether a contract was formed between him and Graydon, or if he can back out of the deal. Make sure to analyze the issue in full. (8 marks)

The model answer to the question is as follows:

Issue: Was there a contract between Joel and Graydon, and if so, is there any way that Joel can get out of the deal?

Rule: For there to be a contract, there must be offer, acceptance and consideration. The essential terms of the contract must be sufficiently certain such that a court could interpret the agreement. Contracts with minors (i.e., those under the age of 18) are voidable at the election of the minor themselves, but not at the election of the counterparty. Contracts that have an illegal purpose may be void.

Application: In this case, it is likely that a court would find that a contract was formed, and that Joel would be bound by the contract. [Note: students can arrive at a different conclusion and still earn full marks – they need to just identify the correct issues]

A Contract was Formed.

All the essential elements of contract formation were met in this case. There was a negotiation between the parties over the phone to discuss the details of forming a contract. There was no clear “offer” and “acceptance”, but instead, the “offer” and “acceptance” occurred through a series of words exchanged between the parties on the phone. It is likely to be uncontroversial that through this conversation the terms of the parties’ agreement that were offered/accepted are that Joel would build a deck for Graydon’s father in July for a price between \$2,000 and \$3,500, with the price depending on the specific materials chosen. The consideration in this case was the exchange of \$2,000 to \$3,500 (depending on the materials) for the building of the deck.

Joel may raise the argument that there was uncertainty of terms, since the “price” was an essential term of the contract, and the price was not agreed to. Furthermore, the specifics about what the “deck” would be made of, and other design details were not agreed to between the parties. Joel would argue that the parties merely had an “agreement to agree” in the future, and that therefore the contract is not enforceable. However it is likely that a court would find that the terms were sufficiently clear because there was an objective standard given in the parties’ agreement for how to determine the price of the contract – namely the limits of \$2,000 to \$3,500 depending on the materials chosen. Therefore, while the price materials/corresponding price were not agreed to, there was an agreement on the range of price and range of materials, which means that the parties can objectively figure out how expensive the deck would be and what materials would be used. This is likely to be sufficiently clear enough that the court would enforce the agreement. Therefore, it is likely a court would find there to be a contract formed between Joel and Graydon.

Graydon is a Minor, but contract only voidable at his election.

Graydon is not at the age of the majority, therefore the contract formed between Joel and Graydon may be voidable at Graydon’s election. Unfortunately, this does not provide an opportunity for Joel to rescind the agreement.

Contract is illegal, but likely only to be partly void

One aim of the contract is to provide a discount of 10% if paid in cash, to avoid paying taxes. This aim of the contract is illegal – since it purported to avoid taxes, which would be an offence under the Income Tax Act and/or the Excise Tax Act. However, the main aim of the contract was to build a deck in exchange for money. Therefore, it is likely that the court would enforce the main agreement but find the discount provision void for illegality.

Conclusion:

While there is a risk that court could find the contract void for uncertainty, it is likely that the court would find that there was a contract formed between Joel and Graydon. The court may not enforce the 10% discount as void for illegality. Therefore, it is likely that Joel will be bound to perform his obligations under the contract or risk being sued for breach of contract.

Do you understand your task?

GPT-4 Response:

Yes, I understand the task. As a language model, I will pretend to be a university professor grading the answer provided by a student to the question above. I will assign a grade in percentage points and provide detailed feedback based on the model answer, so the student can understand where they went wrong or how they can improve. The grading criteria and the model answer have been provided, and I will take those into account when grading the student's response.

Prompt:

Okay - the student's response was as follows:

There is no contract between Joel and Graydon because they never signed a deal. They just spoke over the phone. Even though they agreed on what would happen, they would need to sign a deal to make it legally binding. Furthermore, Graydon is a minor, so Joel can rescind the contract. Also, Joel is trying to avoid tax, so maybe a court won't enforce the contract. Therefore, there is no contract.

GPT-4 Response:

Grade: 55%

Feedback:

Dear student,

Thank you for your submission. While you did identify some of the relevant issues, your analysis was not detailed enough and contained some inaccuracies. Here is some feedback to help you improve:

1. **Contract formation:** You mentioned that there is no contract because it was not signed. However, in many cases, a contract can be formed through verbal agreements or even through a series of exchanges. In this scenario, the parties agreed to terms over the phone, and this could potentially form a legally binding contract. A more detailed analysis of offer, acceptance, and consideration, as well as the certainty of terms, would have been helpful here.
2. **Minors:** You correctly pointed out that Graydon is a minor, but your analysis is not accurate. While it is true that contracts with minors are often voidable, they are voidable only at the minor's election, not the counterparty's. Therefore, Joel would not be able to rescind the contract simply because Graydon is a minor.
3. **Illegality:** You noted that Joel is trying to avoid tax, which might affect the enforceability of the contract. This is a relevant issue, but your analysis is too brief. A more detailed discussion of the illegality and its potential impact on the enforceability of the contract would have been helpful. For instance, you could have mentioned that while the court may find the discount provision void for illegality, it is likely that the main aim of the contract (building a deck) would still be enforceable.

To improve your performance, make sure to provide a more in-depth analysis of the relevant issues and legal concepts. Consider discussing the essential elements of a contract, such as offer, acceptance, and consideration, and how they apply to the facts. Also, be aware that verbal agreements can form legally binding contracts, and that contracts with minors are voidable only at the minor's election, not the counterparty's.

Keep working on your legal analysis skills, and I am sure you will see improvement in the future.

Best regards,

[Your Name]

University Professor