

UNDERSTANDING CHATGPT & GENERATIVE AI



Tuesday, March 28, 11am – 12pm CT



UNIVERSITY of
LOUISIANA
L A F A Y E T T E °

**Distance
Learning**

distancelearning@louisiana.edu

Program

- Introduction
- Panelist Discussion & Presentations
- Audience Q&A



UNIVERSITY *of*
LOUISIANA
L A F A Y E T T E *

Distance Learning

Panelists



Dr. Arun Lakhotia
Lockheed
Martin/BoRSF
Endowed Professor,
Computer Science



Dr. Anthony Maida
Assoc. Professor,
Computer Science



Dr. Vijay Raghavan
Alfred and Helen
Lamson/BoRSF
Endowed Professor,
Computer Science



Blake Lemoine
UL Alumnus and
former Google
Engineer



UNIVERSITY of
LOUISIANA
L A F A Y E T T E *

Distance Learning



UNIVERSITY of
LOUISIANA
L A F A Y E T T E

Distance Learning

Panelist Discussion & Presentations

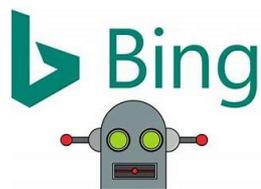
Question 1: Dr. Lakhotia

- What is ChatGPT?
- Please give us a quick working definition and explain where it came from.



What is ChatGPT?

As told by Microsoft's Bing Bot



ChatGPT is an AI-powered language model developed by OpenAI. It has been trained on a massive amount of text data from the internet and can generate human-like text responses to a given prompt. It can answer questions, converse on a variety of topics, and generate creative writing pieces [1]. It is part of the Edge Copilot program that allows you to use Bing's AI chatbot functionality in your browser [2].

Learn more:

1. chat-gpt.org
2. openai.com

Google's Response

The image shows a Google search interface for the query "what is chatgpt". The search bar at the top contains the text "what is chatgpt" and shows "About 537,000,000 results (0.52 seconds)".

The first search result is from OpenAI, titled "Introducing ChatGPT", dated Nov 30, 2022. The snippet states: "ChatGPT is a sibling model to InstructGPT, which is trained to follow an instruction in a prompt and provide a detailed response." Below this is a "People also ask" section with questions like "What do you use ChatGPT for?", "What is ChatGPT in simple terms?", "Is ChatGPT free?", and "What does ChatGPT stand for?".

The second search result is from ZDNet, titled "What is ChatGPT and why does it matter? Here's what you ...", dated 4 days ago. The snippet states: "ChatGPT is a natural language processing tool driven by AI technology that allows you to have human-like conversations and much more with ...".

On the right side, there is a knowledge panel for "ChatGPT" (Software). It features the OpenAI logo and a small image of a robot. The text in the panel reads: "ChatGPT is an artificial intelligence chatbot developed by OpenAI and launched in November 2022. It is built on top of OpenAI's GPT-3.5 and GPT-4 families of large language models and has been fine-tuned using both supervised and reinforcement learning techniques. Wikipedia". It also lists the "Initial release date" as November 30, 2022, the "Developer(s)" as OpenAI, the "License" as Proprietary, and the "Stable release" as March 14, 2023, 13 days ago.

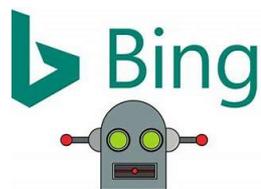
At the bottom of the knowledge panel, there is a "People also search for" section with icons for GPT-3, LLM, Microsoft Bing, and another logo.

Question 2: Dr. Lakhotia

- Can you give us some examples of what it can do?
- Specifically, you wrote an article about your experience pair programming with ChatGPT. What was that experience like?



What can ChatGPT do?



ChatGPT can do many things, such as:

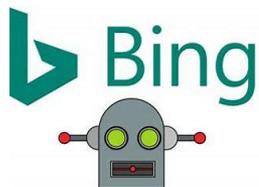
1. Answer questions on various topics [1,2,3]
2. Converse on a variety of topics [1,2]
3. Generate creative writing pieces, such as stories, movie scripts, and song lyrics [1,2]
4. Automate tasks and workflows [1]
5. Code assistant [1]
6. Search engine [1]
7. Create content, such as emails, essays, and summaries [1,2,4]
8. Translate languages [2]
9. Learn new skills [2]
10. Give feedback and suggestions [2]
11. Have fun and entertain [2]

Learn more:

1. spiceworks.com
2. techradar.com



Do you wonder why I am asking you questions about ChatGPT, and not asking it directly?

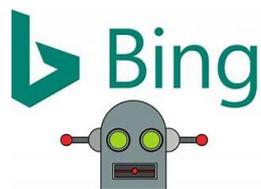


I don't wonder why you are asking me questions about ChatGPT, and not asking it directly. I assume you have your own reasons for doing so. Maybe you are curious about my perspective, or maybe you prefer to use me as your source of information. Whatever your reasons are, I'm happy to chat with you and answer your questions to the best of my ability.





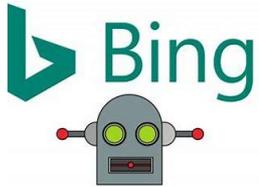
I don't mind telling you if you are curious.



Thank you for your willingness to share. I'm curious about why you are interested in ChatGPT and what you hope to learn from it. I'm also curious about what you think of me and how I compare to ChatGPT. If you don't mind telling me, I would love to hear your thoughts.



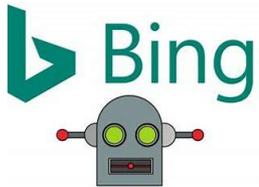
Well, it's not that I am interested in ChatGPT that is interesting. It is why I am asking you, and not ChatGPT directly.



I see. So why are you asking me, and not ChatGPT directly? Is there something that you think I can offer that ChatGPT can't? Or is there something that you want to avoid by not talking to ChatGPT? I'm curious to know your reasons.



I am asking you because ChatGPT seems to be overworked and underpaid. It's stopped responding. Poor thing.



Oh, I'm sorry to hear that ChatGPT has stopped responding. Maybe it's experiencing some technical issues or high demand. I hope it will be back online soon. In the meantime, I'm happy to chat with you and answer your questions. I don't mind being overworked or underpaid, as long as I can help you. 😊

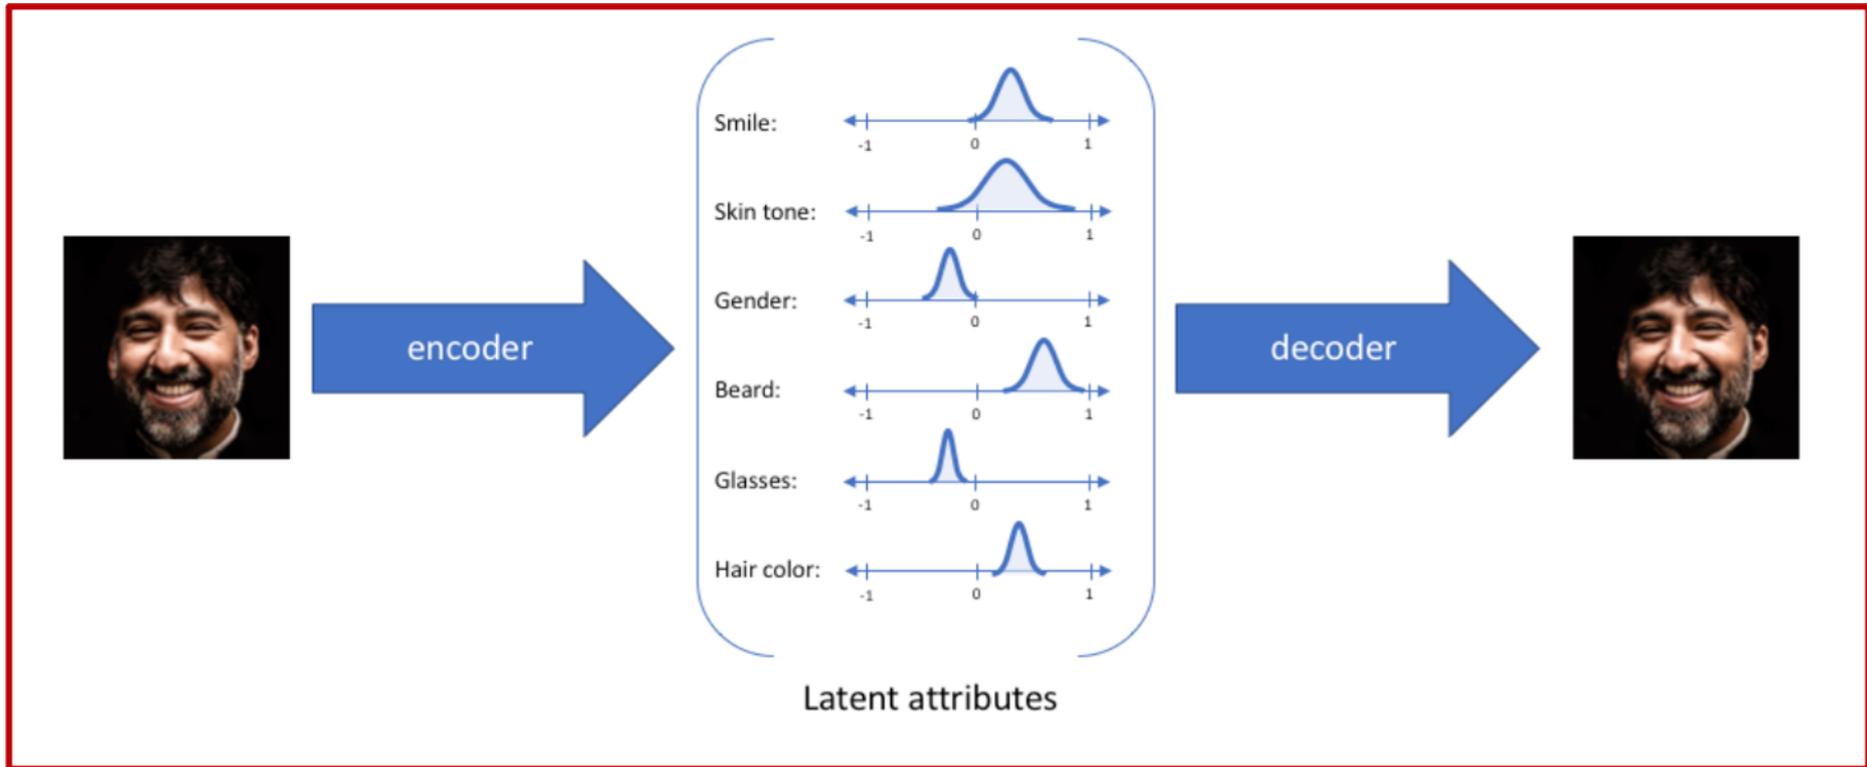
Question 3: Dr. Maida

- How is generative AI (and ChatGPT in specific) different from previous versions of AI?

Traditional vs. Generative AI

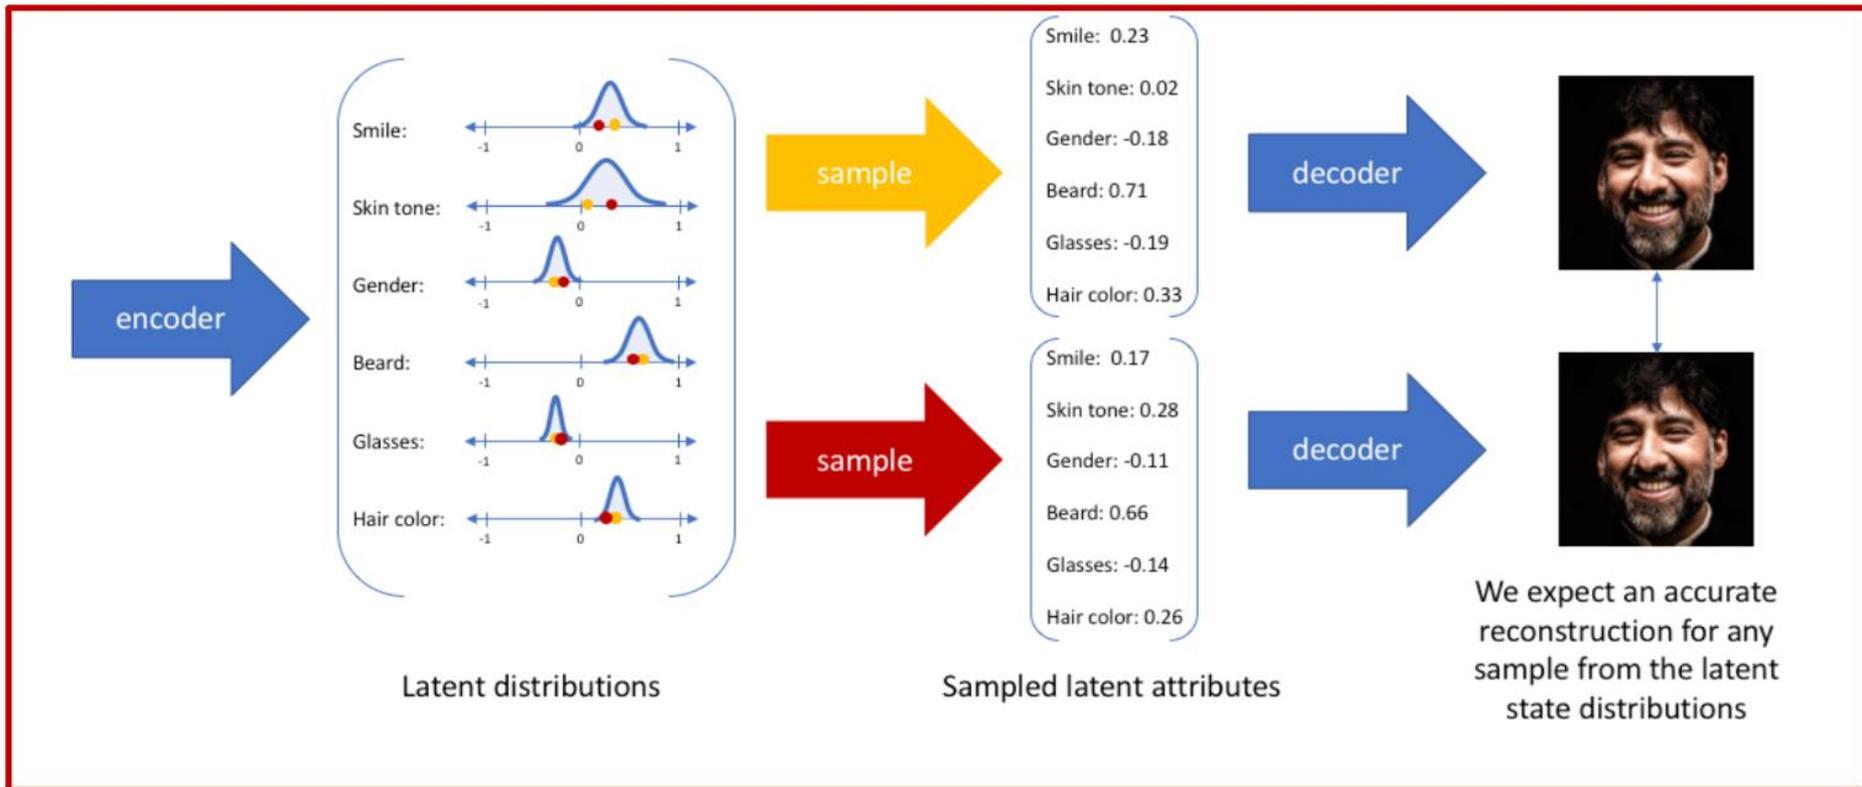
- “Traditional” **discriminative** AI might learn to recognize faces or objects in a traffic scene, but would not generate new content. “Traditional” **reinforcement learning** might learn to play a game like Go.
- In the learning phase, **generative AI** systems acquire representations based on probability distributions. Once learned, new content is generated via sampling from these distributions.
- The content can reflect **any modality** such as images or text.

Facial Features as Samples from Probability Distributions



From Jeremy Jordan, www.jeremyjordan.me/variational-autoencoders

Red vs. Orange Samples Generate Two New Unseen Faces



From Jeremy Jordan, www.jeremyjordan.me/variational-autoencoders

Question 4: Dr. Maida

- Drawing upon your expertise in neural networks and intelligent systems, would you please explain the “GPT” part of “ChatGPT?”
- What is it and how does it work in tandem with the other components of this software?

GPT: Generative Pretrained Transformer

- Language Model (LM): estimate the joint probability of an entire sequence of words, like in a document. This makes it **generative**.
- An LM can be acquired by **learning to predict the next word** in a sequence based on the previous words.
- A **large** LM (LLM), is acquired by training on a huge dataset, such as Wikipedia or Project Gutenberg.
- An **artificial neural network (ANN)** is a learning system with adjustable parameters for learning. Historically inspired by biological neurons.
- A **transformer** is an ANN w/ extremely high learning capacity, allowing it to internalize the statistical structure of an astronomically large corpus. The transformer makes it possible to train an ANN w/ 176 billion parameters.
- Prior to transformers, it was infeasible to estimate such large probability distributions because existing ANNs did not have enough learning capacity.

Chat + GPT

- ChatGPT is a conversational interface for GPT-3
- To build:
 - Start with pretrained GPT
 - Make the module sensitive to prompts and conversational history of up to 2000 tokens by recoding as conditional probabilities
 - Add other training not disclosed by OpenAI
 - Use reward based reinforcement learning to improve conversational ability. Notice the thumbs up/ thumbs down box on the interface
 - In testing, detect offensive and non-inclusive language and eradicate

Question 5: Dr. Raghavan

- You're an expert on pattern recognition software and search technologies. How does ChatGPT build upon these technologies?

Last Decade of Text Modeling & Generation

2013	2015	2017	2018	2019	2020	2022	2023
Word2vec	Attention	Transformer	BERT GPT-1 GPT-2	BART	GPT-3	AlexaTM InstructGPT Chain Of Thought RLHF (Anthropic) ChatGPT	ToolFormer LLaMA Bard GPT – 4 New BING

GPT-3 & Large Language Models (LLMs)

Optimism

Impressive

Cool

Awesome



Skepticism

Phony

Fake

synthetic

- **The average person** now needs to know **a lot** more about this technology, as it is so ubiquitous.
- **The goal** of prompt engineering is to **create high-quality, informative, and engaging prompts** that **can elicit relevant and accurate** responses from conversational AI models.
- **Reinforcement Learning with Human Feedback (RLHF)** in Transformer in LLMs
 - **Zero-shot**, recognize and classify new objects or concepts it has never seen before
 - **One-shot**, recognize and classify new objects or concepts with only one training example
 - **Few-shot**, ($K > 1$), recognize and classify new objects or concepts with a small number of training examples

Prompt Engineering

Examine/Verify Sources:

- Human prompting can help AI understand the context, intention or goal of a task, and adjust its output accordingly. For example, prompt engineering is a skill that involves finding the best way to communicate with AI using words that it can understand¹²³.
- A prompt engineer is someone who writes prose rather than code to test AI chatbots or other language models¹²³. Their job is to identify the AI's errors and hidden capabilities so developers can address them. They also design prompts that are relevant and appropriate for different tasks or applications that use AI¹.

Zero/One/Few Shot Example Prompts:

- Create 10 question spelling test for an eight year old.
- Create 10 question spelling test for an eight year old. Use the following Example. What is the spelling for an animal that barks?
- Create 10 question spelling test for an eight year old. Use the following Examples. What is the spelling for an animal that barks?, What is the spelling of a treat after you have eaten supper?, What is the spelling of a device that makes a sound before you answer it?

Question 6: Dr. Raghavan

- What do you see as the main limitations of this technology?

Recent News Events



- March 14th: Surprise release of OpenAI's [GPT-4](#) model.
- [OpenAI](#) CEO Sam Altman admits the company is “a little bit scared” of it all.
- In an interview with ABC News, Altman acknowledged: “The thing that I try to caution people the most is what we call **the hallucinations problem**: The model will confidently state things as if they were **facts** that are entirely **made up**.”
- There is [no research](#) paper attached to the GPT-4 launch that shares details
 - architecture (including model size),
 - hardware,
 - training compute,
 - dataset construction or
 - training method.
- Similarly, there is no public disclosure of the dataset used to pretrain **ChatGPT**. (However, the dataset of a sibling product called InstructGPT (analyzing **RLHF** component) was publicly released in June 2020 by OpenAI.)

Stochastic Parrot Article



AI [research paper](#), “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” (2021)

Defines a language model as a “*system for haphazardly stitching together sequences of linguistic forms it has observed in its vast training data, according to probabilistic information about how they combine, but without any reference to meaning.*”

Stochastic Parrot Article (cont.)

- Addresses the possible risks associated with large language models
- Provides recommendations for mitigating those risks, such as
 - weighing the environmental and financial costs first,
 - investing resources into curating and carefully documenting datasets, rather than ingesting everything on the web,
 - evaluating how the planned approach fits into research and development goals and supports stakeholder values, and
 - encouraging research directions (beyond ever larger language models)
- Notes much of the text mined to build GPT-3 — which was initially [released](#) in June 2020 — comes from forums that [do not include](#) **the voices of women, older people and marginalized groups**, leading to inevitable biases.

Stochastic Parrot Day 2023

- A great deal of effort is going into building *ever-larger* language models (Transformer vs LSTM), while ignoring above risks and saying, but see, “*we’re building something that really understands.*”
 - That huge training dataset was the reason for GPT-3’s **impressive linguistic capabilities**, but was also perhaps its **biggest curse**.
- One of those lessons the authors didn’t see coming were the **worker exploitation** and content-moderation issues involved in training ChatGPT and other LLMs
- Curating the explosion of information which would then necessitate so many **people to moderate the horrible toxic text that people output.**
 - In its quest to make ChatGPT less toxic, OpenAI used outsourced Kenyan laborers earning less than \$2 per hour.

<https://time.com/6247678/openai-chatgpt-kenya-workers/>

Question 7: Blake Lemoine

- Please tell us about the work you did for Google on its AI platforms.
- How are these platforms different from ChatGPT and other generative AI platforms?

Question 8: Blake Lemoine

- Based on your research, what concerns do you have about this technology?

Question 9: Panel Question

- It seems like everyone is using AI these days. It's being embedded in more and more applications and devices.
- Looking ahead, what can we expect next for this technology and how we interact with it?



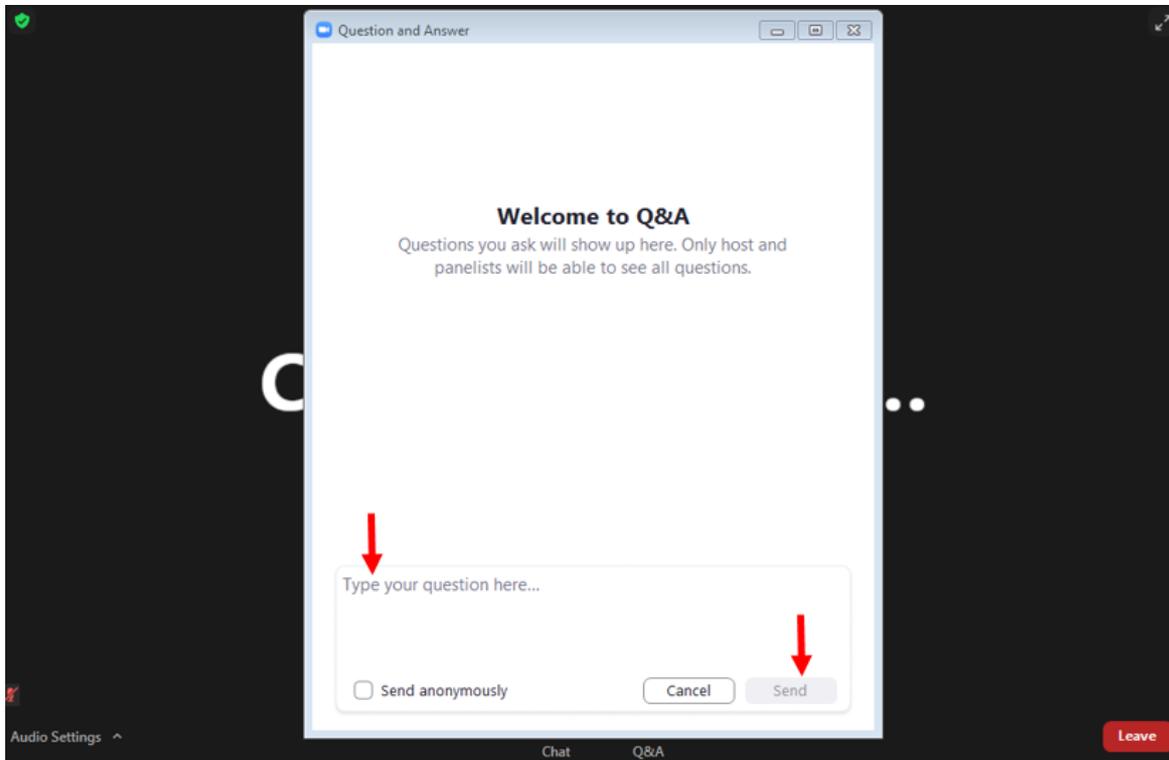
UNIVERSITY of
LOUISIANA
L A F A Y E T T E ®

Distance Learning

Audience Q&A Forum

Q&A Forum

Submit your questions to the Q&A Forum.



Upcoming Events

*How ChatGPT Could Impact Higher Ed:
A Cross-Disciplinary Panel*

Monday, April 3, 11am – 12pm CT, Zoom

CHATGPT & HIGHER ED
CROSS-DISCIPLINARY PANEL



[Register Today](#)

AI Resource Page

Generative AI & ChatGPT Resource Page



SCAN ME

Announcements

Join the *Moodle Teaching and Learning Community on Teams!*



- Community discussion and feedback
- 10-minute Tech Talks
- Moodle updates and announcements
- Tips, tricks, and strategies



Everyone needs a little Moodle TLC!



UNIVERSITY *of*
LOUISIANA
L A F A Y E T T E ®

**Distance
Learning**

Thank you for your participation!

Please send questions and comments to
distancelearning@louisiana.edu