AI Overview/Update

DCSS - Nov 16, 2023

What is Generative AI?

Generative AI (GenAI) is a type of artificial intelligence that can create new content, such as text, images, videos, or code, from the data it has been trained on. GenAI models learn the patterns and structure of their input data and then generate new data that has similar characteristics. *Wikipedia*

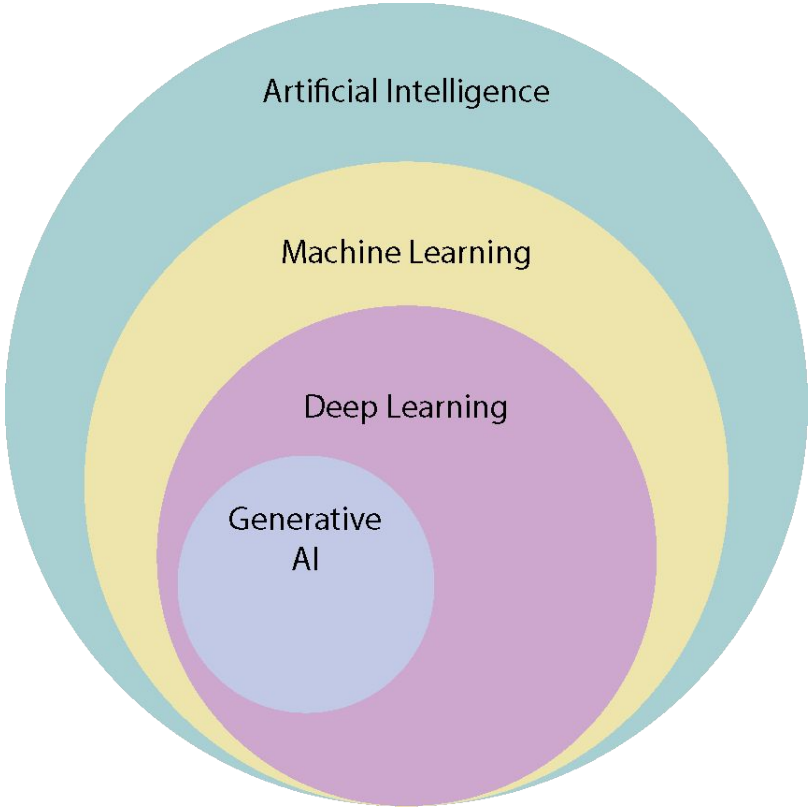

What are Large Language Models?

A large language model (LLM) is a specialized type of artificial intelligence (AI) that has been trained on vast amounts of text to understand existing content and generate original content.

Gartner

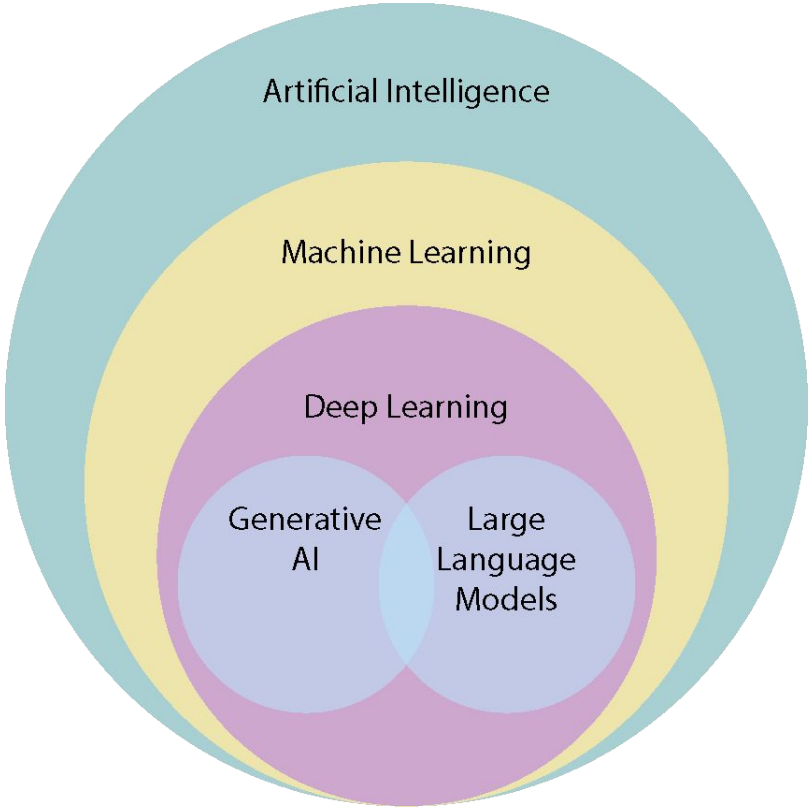

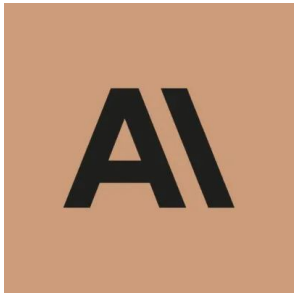


Bard

EU AI Act

'Artificial intelligence system' (AI system) means a system that is designed to operate with a certain level of autonomy and that, based on machine and/or human-provided data and inputs, infers how to achieve a given set of human-defined objectives using machine learning and/or logic- and knowledge-based approaches.

https://artificialintelligenceact.eu/the-act/

National Artificial Intelligence Initiative Act of 2020 (NAIIA)

The term "artificial intelligence" means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to—

(A) perceive real and virtual environments;

(B) abstract such perceptions into models through analysis in an automated manner; and

(C) use model inference to formulate options for information or action.

https://www.uspto.gov/sites/default/files/documents/National-Artificial-Intelligence-Initiative-Overview.pdf

NIST - Artificial Intelligence Risk Management Framework (AI RMF 1.0)

An engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy

https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

NIST AI RMF Playbook: https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook

Executive Order on Responsible AI Innovation (10/30/23)

Safety & Security: Establishes new standards for AI systems, mandates safety test sharing, and addresses critical infrastructure and biosecurity risks.

Privacy Protection: Urges Congress to pass data privacy laws, promotes development of privacy-preserving AI technologies, and reinforces federal privacy guidelines.

Equity & Civil Rights: Provides guidance against AI-driven discrimination and develops fair practices in criminal justice.

Consumer & Worker Advocacy: Protects against AI-related harms in healthcare, education, and labor, promoting responsible AI applications.

Innovation & Competition: Encourages AI research, supports small developers, and maintains U.S. leadership in AI.

Global Leadership: Expands international collaborations on AI safety and standards, and promotes responsible AI development worldwide.

Government AI Use: Sets standards for AI in government, enhances procurement, and focuses on workforce training.

Virginia - Executive Directive on AI (9/20/23)

Key Areas of Focus:

Legal Protections: Review AI's legal implications, including privacy and intellectual property.

Policy Standards: Identify standards for state agencies' effective AI use.

IT Safeguards: Establish cybersecurity and other safeguards to mitigate security and privacy risks.

Education: Ensure competitive AI training while preventing misuse in classrooms.

https://www.governor.virginia.gov/media/governorvirginiagov/governor-of-virginia/pdf/ed/Executive-Directive-No.-5---Recognizing-the-Risks-and-Seizing-the-Opportunities-of-Artificial-Intelligence.pdf

Virginia Tech Policies, Standards and Guidelines

Acceptable Use

- 7000 - Acceptable Use and Administration of Computer and Communication Systems
- 4082 - Appropriate Use of Electronic Personnel and Payroll Records

Security and Data Protection

- 7010 - Policy for Securing Technology Resources and Services
- 7035 - Privacy Policy for Employees' Electronic Communications
- Virginia Tech Risk Classifications
- Guidelines for Data Stewards

etc.

https://it.vt.edu/resources/policies.html

Risk Classification Labels

These definitions establish the Virginia Tech Risk Classification labels. Risk Classification serves primarily the "Confidentiality" and "Integrity" security objectives of the CIA triad.

HIGH RISK

Data and IT resources are classified as High-Risk if:

1. Protection of the data is required by law/regulation/contractual obligation, and

2. Virginia Tech is required to self-report to a government agency and/or provide notice to the individual if the data is inappropriately accessed; or

3. The loss of confidentiality, integrity, or availability of the data or system could have a significant adverse to catastrophic impact on our mission, safety, finances, or reputation.

MODERATE RISK

Data and IT resources are classified as Moderate-Risk if they are not considered to be High-Risk, and:

1. The data is not generally available to the public, or

2. The loss of confidentiality, integrity, or availability of the data or system could have a mild to moderate adverse impact on our mission, safety, finances, or reputation.

LOW RISK

Data and IT resources are classified as Low-Risk if they are not considered to be Moderate or High-Risk, and:

1. The data is intended for public disclosure, or

2. The loss of confidentiality, integrity, or availability of the data or system would have no adverse impact on our mission, safety, finances, or reputation.

https://it.vt.edu/content/dam/it_vt_edu/policies/Virginia-Tech-Risk-Classifications.pdf

Recommendations for faculty

1. Become familiar with generative AI tools
2. Consider the Honor Code and its applicability to generative AI tools
3. Avoid being drawn into a confrontational mindset regarding these tools
4. Set clear expectations for your students regarding the use of generative AI/ChatGPT - add guidance to course documents and discuss with students
5. Explore potential changes to your course design and/or assessment strategies

https://tlos.vt.edu/resources/generative-ai.html

Limitations

1. Understanding Context and Nuance: While they have improved significantly in understanding context, they can still struggle with complex or nuanced scenarios, often missing subtle implications or emotional undertones.
2. Dependency on Training Data: Their responses are based on patterns learned from the training data. If certain topics or perspectives are underrepresented in the training data, the model may provide incomplete or biased responses.
3. Handling Ambiguity: They can have difficulties with ambiguous queries where multiple interpretations are possible, sometimes choosing an interpretation randomly or based on more commonly seen patterns.
4. Lack of Creativity and Original Thought: While they can generate content that seems creative, it's essentially a recombination of existing ideas and styles. They lack genuine creativity or the ability to form original thoughts.
5. No Real Understanding or Consciousness: They simulate understanding but don't actually comprehend the information they process. This lack of real understanding can lead to errors or nonsensical responses.

Concerns

1. Misinformation and Reliability: Chatbots can inadvertently generate and spread misinformation, especially if they rely on unreliable sources or misinterpret data

2. Bias and Ethical Concerns: They can perpetuate biases present in their training data, which raises ethical concerns regarding fairness, stereotyping, and representation.

   Privacy and Data Security: The use of personal data in training these models raises privacy concerns. Additionally, there's the risk of them generating responses that inadvertently reveal sensitive information.

3. Dependence and Reduced Human Interaction: Over reliance on chatbots for communication and problem-solving could lead to reduced human interaction and potentially impact social skills and emotional intelligence.

4. Manipulation and Misuse: There's a risk of chatbots being used for malicious purposes, such as spreading propaganda, scamming, or manipulating users.

5. Impact on Employment: Automation of tasks using chatbots could lead to job displacement in certain sectors, raising economic and social concerns.

6. Ethical and Societal Implications:

MIDJOURNEY

Vegan Buddha Bowl, Colorful selection of fresh vegetables and grains, Zen-like simplicity, Serene and Balanced, Natural and Minimalist styling, Soft diffused lighting

woodblock print, cardinal in summer

GQ cover photo of Abraham Lincoln, stylish modern white shirt, short hair, handsome, fashion Photography, studio lighting

Speech synthesis

- Accurate replication of voice patterns
- Ethical considerations for privacy & authenticity

Based on the signs provided:

Today is Wednesday. You cannot park from 11 a.m. to 1 p.m. due to street cleaning.
It is currently 10:01 a.m., so you should avoid parking now since the no-parking period starts in 59 minutes.
There's a "No Stopping Mon thru Fri 7:30 a.m. thru 4 p.m." sign which applies on school days. However, there's an exception for "15 minutes Mon thru Fri 7:30 to 8 a.m. for school business," but that time has passed.
There's also a "1 Hour Parking NON-SCHOOL DAYS Sun 12 midnight thru Fri 6 p.m." sign, but today being Wednesday, if it's a school day, you can't park until after 4 p.m.
After 4 p.m., there's a "1 Hour Parking Mon thru Fri 4 p.m. thru 6 p.m." sign, allowing you to park for an hour.
Given this information, your next opportunity to park here on a Wednesday would be after 1 p.m. However, if it's a school day, you'll need to wait until after 4 p.m. If you park after 4 p.m., you have a 1-hour parking window until 6 p.m.

It is currently 10:01am on Wednesday. When can I park here next?

Explain

The cartoon illustrates a humorous juxtaposition. There's a character that looks like a salt shaker, standing atop a platform with a banner reading "SLUGS FOR SALT". Below the platform, several slugs are shown appearing as if they are supporting or cheering for the salt shaker.

The humor comes from the fact that in real life, slugs are adversely affected by salt. When salt is applied to slugs, they can dehydrate and potentially die. Thus, the image humorously and ironically portrays the slugs as advocating for something that is dangerous to them. The cartoon can be a metaphorical representation of situations where individuals or groups unknowingly support causes or entities detrimental to their well-being.

Video

- Text to video applications create new videos from scratch
- Video and audio cloning combine to allow for functions such as translation (and deep fakes)

Examples of Potential Use Cases

- Information Retrieval and Management
- Analytical Assistance
- Interactive Simulation
- Cognitive Extension
- Process Automation
- Instruction and Tutoring
- Content Creation and Modification
- Emotional and Social Interaction

Image created using DALLE 3 with prompt that listed terms on this slide

Evolving Characteristics/Features

- Autonomy
- Integration
- Continuity
- Scalability
- Multimodality

Image created using Midjourney with prompt: "Illustrate an AI system gaining Autonomy, Integration, Continuity, Scalability, Multimodality"

Evaluation/Governance Criteria

- Privacy and Security
- Accessibility
- Cost (Fiscal, Environmental, Social)
- Human On/In the Loop
- Veracity
- Ethical Alignment
- Transparency

Image created using DALL-E with prompt: “a beautiful non-textual visualization of the concepts of Privacy and Security”

Academic Integrity

- Some students are using tools like ChatGPT inappropriately
- AI “detection” technology is not currently reliable enough to implement
- Students must learn to use AI responsibly
- Faculty should consider how learning assessment is affected by generative AI

Image created using Midjourney with prompt: “a student stands at a digital crossroads, one path leading to brightly lit digital honesty, and the other leading to rocky digital dishonesty.”

Considerations for Distributed Computing

- Become familiar with available chatbots
- Encourage and support exploration by faculty
- Explore introductory courses on generative AI
- Learn more about prompt design/engineering
- Consider API-driven implications (e.g.Retrieval Augmented Generation (RAG))
- Participate in conversations about governance, privacy and security

Thank you!