



# Alternative Facts: Knowledge Graphs in the Age of Large Language Models

**Dr. Ora Lassila**

Principal Technologist  
Amazon **Neptune**



# Who am I?

Principal Technologist in the Amazon Neptune graph database team

Co-author of

- the original RDF specification
- the seminal paper on the Semantic Web

Recipient of the 1<sup>st</sup> ISWC “10-year award”

Former W3C Fellow & elected member of the W3C Advisory Board

Author of the KR system that flew on NASA’s “Deep Space 1” probe in 1998

Grand Prize winner, 1989 USENIX Obfuscated C Code Contest

# Game plan

1. General observations about AI and its history
2. My worries about the current state of AI
3. Knowledge graphs to the rescue?

**NB: This is not a talk about Neptune, and I only mention the Semantic Web once (OK, twice)**

# I have been studying AI for a long time...

This is the first page of my first published conference paper:

- Finnish AI Symposium 1984
- (I was ~~in kindergarten~~ a college freshman...)

I will take a "long view" discussing AI in this talk

## The ALLL Environment - USING OBJECT-ORIENTED CONCEPTS IN LISP

Orri Erling  
Keijontie 5  
00610 HELSINKI 61

Pekka P. Pirinen

Ora Lassila  
Lamminjarventie  
02880 VEIKKOLA

### 1. Introduction

This paper discusses the use of object-oriented concepts, especially in LISP. All this is seen in the context of ALLL, basically a LISP enhanced with objects, implemented on a microcomputer. The acronym "ALLL" actually stands for "A Lisp-Like Language" (pronounced as the English word "all"), a name chosen to reflect the fact that the language implements all of LISP's essential functions. However, ALLL is not just another subset of LISP, but a practical programming language of its own.

We developed ALLL (beginning in 1981) for implementing a sophisticated database system, a task which would have been extremely cumbersome using conventional programming languages. LISP had some of the qualities required, but was not considered efficient enough, in particular on a microcomputer. A LISP-based language, without the characteristic lack of speed, seemed however the optimum solution. The influence of Smalltalk led us to introduce objects to the language. A software-based virtual memory facility was developed, providing an elegant answer to the problems of file management. From these elements emerged ALLL, which has been used as a research and development tool since 1983.

# What is AI anyway?

Historically, AI has been about logic and the manipulation of symbols

- “symbolic AI”
- symbol grounding problem

More recently, this manipulation has been more statistical in nature

- “non-symbolic AI”
- miraculously successful, but unable to explain results (backlash looms)


The broad goal of AI has been (and remains) automated problem solving

- the dream of “artificial general intelligence” still eludes us...

# What is AI anyway?

Many technologies may have started out in the field of AI, but are now “mainstream”

- object-oriented programming
- dynamic memory management
- rule systems
- ontologies and graphs
- etc.



*Pictured: CPU board of Texas Instruments Explorer II, a 36-bit bespoke hardware architecture for running Lisp*

# What is AI anyway?

I entered the field during an “AI boom” (1980s)

Many technologies were quite promising, but expectations led to “over-promising”, and eventually the so-called “AI winter” (1990s)

Later, the “miraculous success” of non-symbolic methods (2000s and onwards)

*Note: Everyone now fashions themselves as an AI expert, including those former colleagues of mine who earlier were downright hostile in their opposition of anything that smelled like AI*

# What is AI anyway?

AI has suffered from scarcity of computational ability (algorithms) and capability (CPU, memory), **driving innovation** also outside the field of AI

- machine translation in the 1960s
- running Lisp programs
- ML, neural networks, etc., require **lots** of CPU and memory

Amara's Law seems to apply to AI quite well:

***“We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run.”***

– Roy Amara (American futurologist, 1925-2007)

# What is AI anyway?

I will come back to this...

AI has suffered from scarcity of computational ability (algorithms) and capability (CPU, memory), **driving innovation** also outside the field of AI

- machine translation in the 1960s
- running Lisp programs
- ML, neural networks, etc., require **lots** of CPU and memory

Amara's Law seems to apply to AI quite well:

*“We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run.”*

– Roy Amara (American futurologist, 1925-2007)

# Large Language Models

Essentially these are massive ML models which, given some text, are capable of predicting what should come next (your mileage may vary)

- “autocorrect on steroids”

Typical flow of operations:

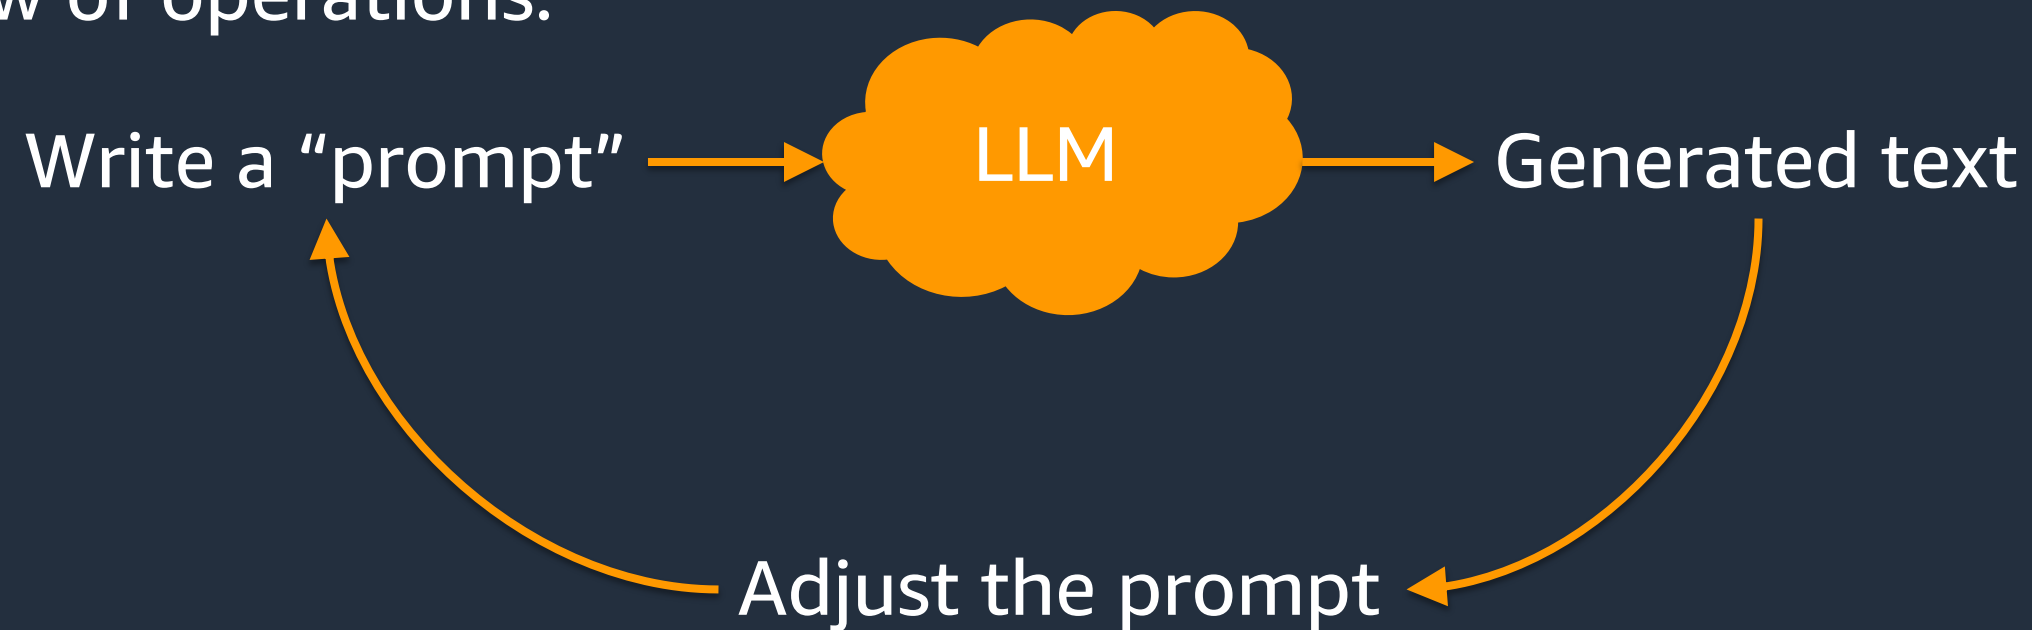


# Large Language Models

Essentially these are massive ML models which, given some text, are capable of predicting what should come next (your mileage may vary)

- “autocorrect on steroids”

Typical flow of operations:

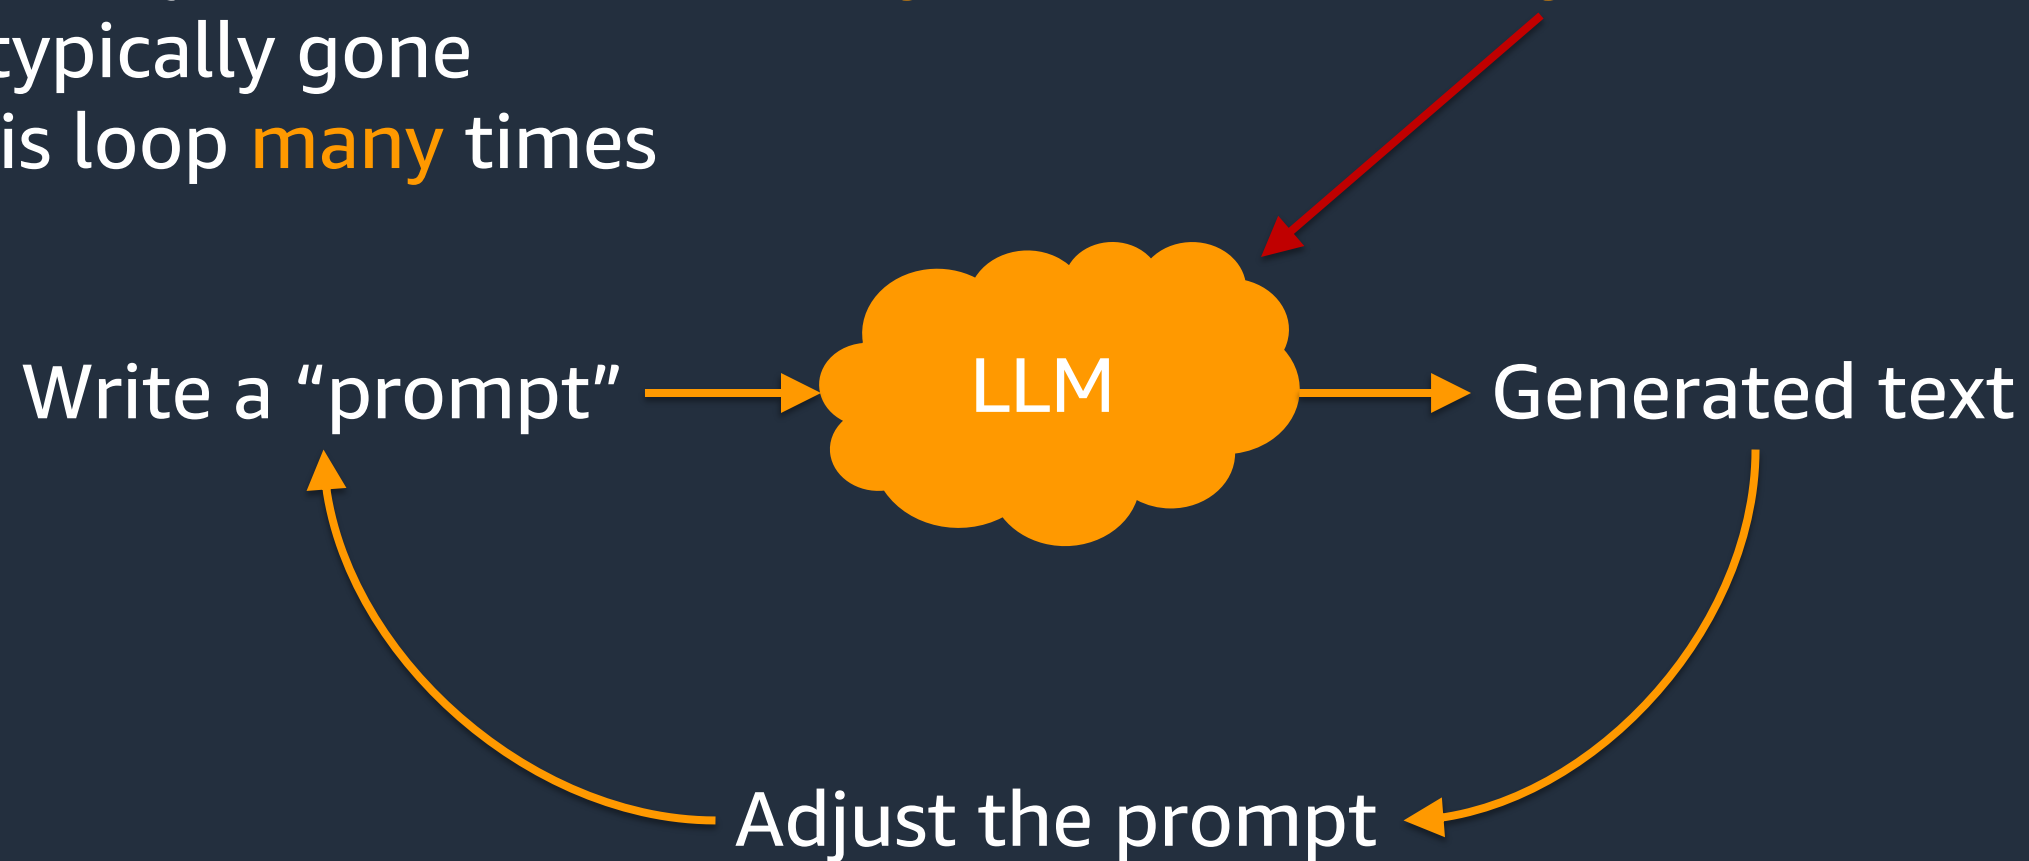


New discipline: “prompt engineering” (seriously)

# Large Language Models need a “human in the loop”

When people show you results from (say) ChatGPT, they have typically gone through this loop **many** times

*What is happening here? Why did we get the answer we got? Is it correct?*



# Large Language Models are now being productized

OpenAI: GPT, ChatGPT (available on Azure)

Google: Bard

Amazon: Bedrock, Titan

Custom (domain-constrained) LLMs (e.g., Bloomberg)

Many companies are now building UIs using LLMs

# My worries

My colleague Charles Ivie was experimenting with ChatGPT and asked it to write some Python code against the RDFLib library

Resulting code was decent Python, but ChatGPT **invented completely new API calls** (that do not exist in the library, and never have)

This is still easy to detect, though

# My worries

A more chilling example:

- fake scenario (a school class trip that never actually took place)
- fake WP article cited as “evidence”
- but... **a real person was named**



The screenshot shows the top of a Washington Post article. The header includes the site logo, a search icon, a 'Subscribe' button, and a user profile icon labeled 'ora'. Below the header is a navigation bar with categories: Tech, Help Desk, Future of Transportation, Innovations, Internet Culture, Space, and Tech Policy. The article is categorized under 'INNOVATIONS'. The main headline reads: 'ChatGPT invented a sexual harassment scandal and named a real law prof as the accused'. A sub-headline states: 'The AI chatbot can misrepresent key facts with great flourish, even citing a fake Washington Post article as evidence'. The byline is 'By Pranshu Verma and Will Oremus' and the date is 'April 5, 2023 at 2:07 p.m. EDT'.

# My worries

We might be entering a true gilded age of “fake news” ...

This example: AI-generated video used as a political attack ad, deliberately created to have nothing to do with actual facts

The screenshot shows a Forbes article page. At the top, the Forbes logo is displayed in white on a black background. Below the logo, the breadcrumb 'FORBES > BUSINESS' is visible. The article is marked as 'BREAKING'. The main headline is 'Republicans Launch Eerie AI-Generated Attack Ad On Biden'. The author is Sara Dorn, Forbes Staff, with a 'Follow' button next to her name. The article is categorized under 'Politics'. The publication date is 'Apr 25, 2023, 11:21am EDT'. A 'Listen to article 3 minutes' button is present, featuring a play icon and a waveform icon. The article text begins with a 'TOPLINE' section: 'The Republican National Committee debuted an apocalyptic foreshadowing of a second Biden-Harris term moments after President Joe Biden announced his reelection campaign on Tuesday in a video rebuttal generated entirely by artificial intelligence—the latest sign that computer-generated content could shift the political campaign landscape.'

# My worries

Many AI researchers and practitioners have called for LLM development to be paused or slowed down

- an open letter from the Future of Life Institute has approx. 30k signatures

They have the right set of worries...

- they call for AI to be “accurate, safe, interpretable, transparent, robust, aligned, trustworthy, loyal”

I think trying to limit innovation is a bad idea, however

- (and the proverbial cat is already out of the bag anyway...)

Instead, we need to develop technological safeguards and guardrails

- many such tools already in development (e.g., NVIDIA NeMo)

# My worries, specifics

Symbol grounding, truth, world models

“Hallucinations”

Complexity of these systems is beyond what we can understand

- hence the trial-and-error of “prompt engineering”

Cost? (= carbon footprint)

Non-symbolic AI in general tends to be “opaque”

- more calls for “responsible AI” or “accountable AI” (this is hard, though)
- backlash possible

# Knowledge graphs?

The Semantic Web (as the foundation for practical KR and modern knowledge graphs) is much what we wanted from AI before the “AI winter”

Properly constructed knowledge graphs

- employ well thought-out models/ontologies
- represent curated knowledge (verifiable facts)

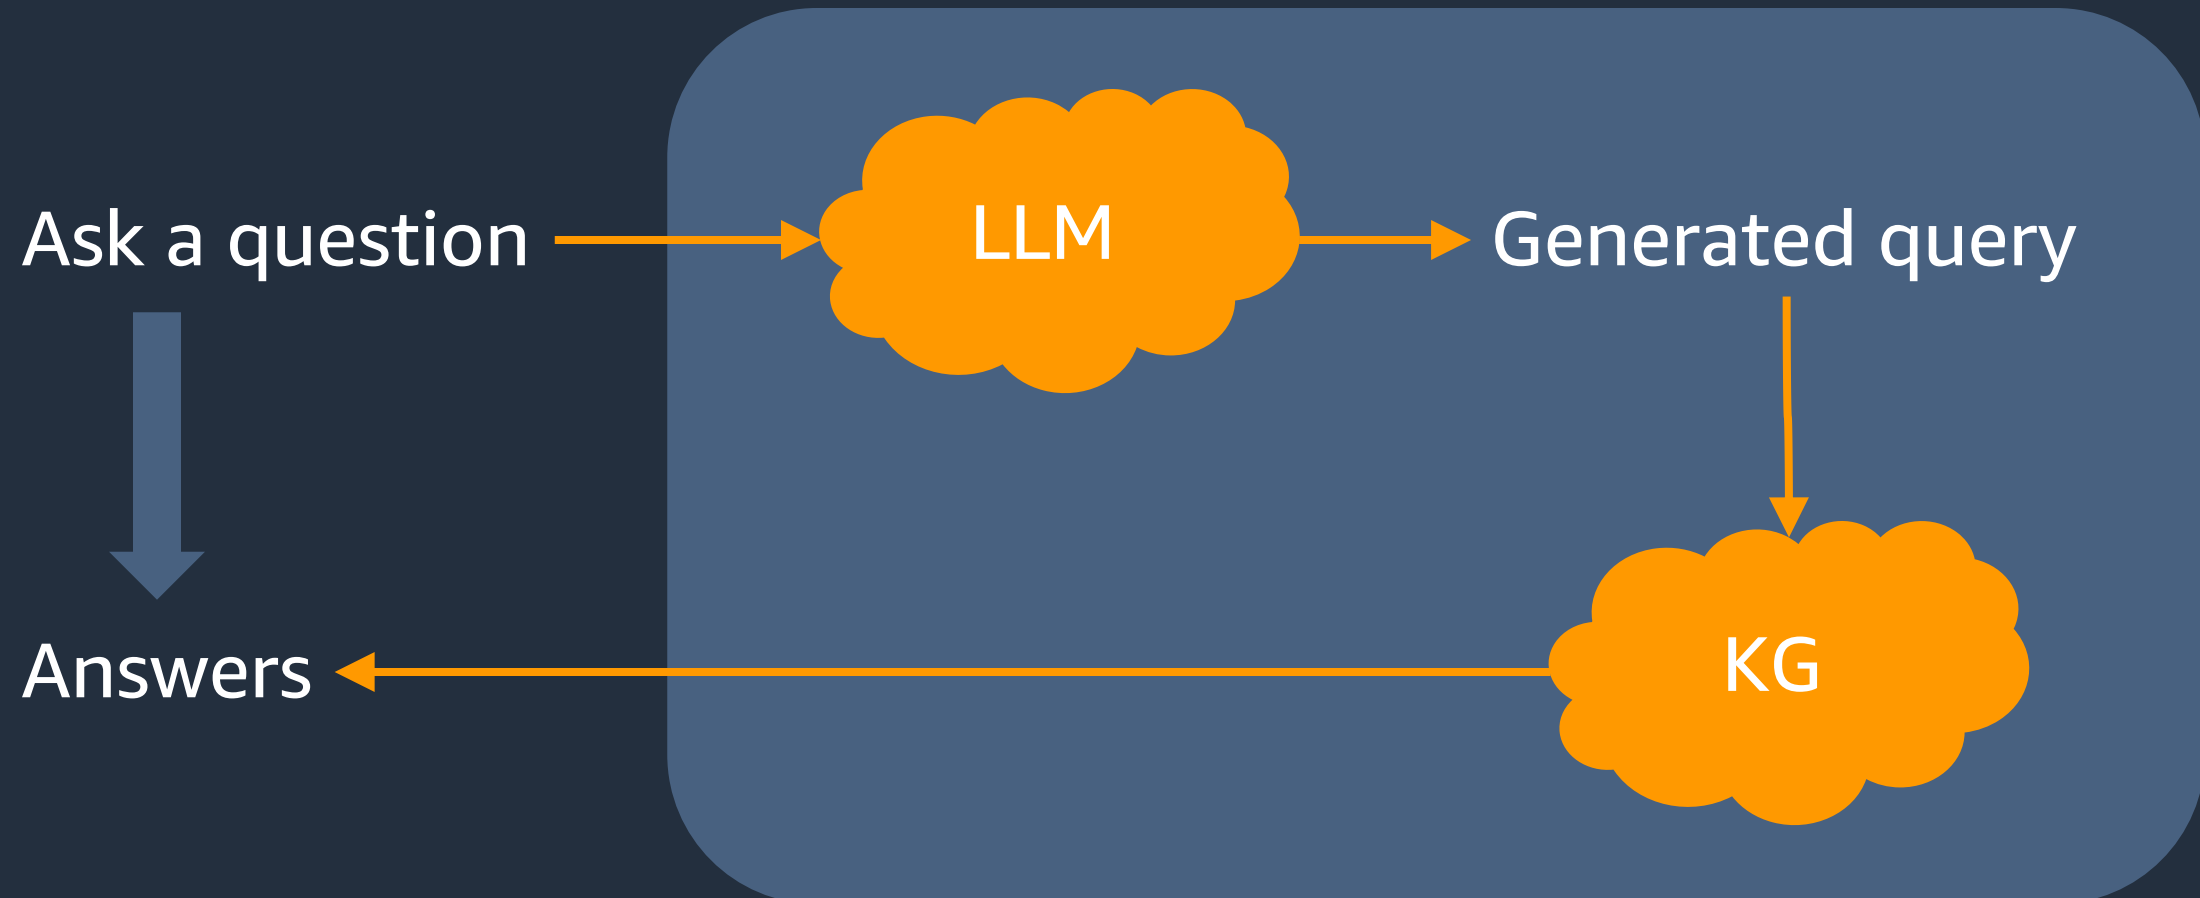
Graph databases (as means of storing KGs) have become commonplace, and graph technologies are now “mainstream”

What we need now is to understand how generative AI (and non-symbolic methods in general) can work with knowledge graphs

# Knowledge graphs?

There are already promising results in letting LLMs write graph queries

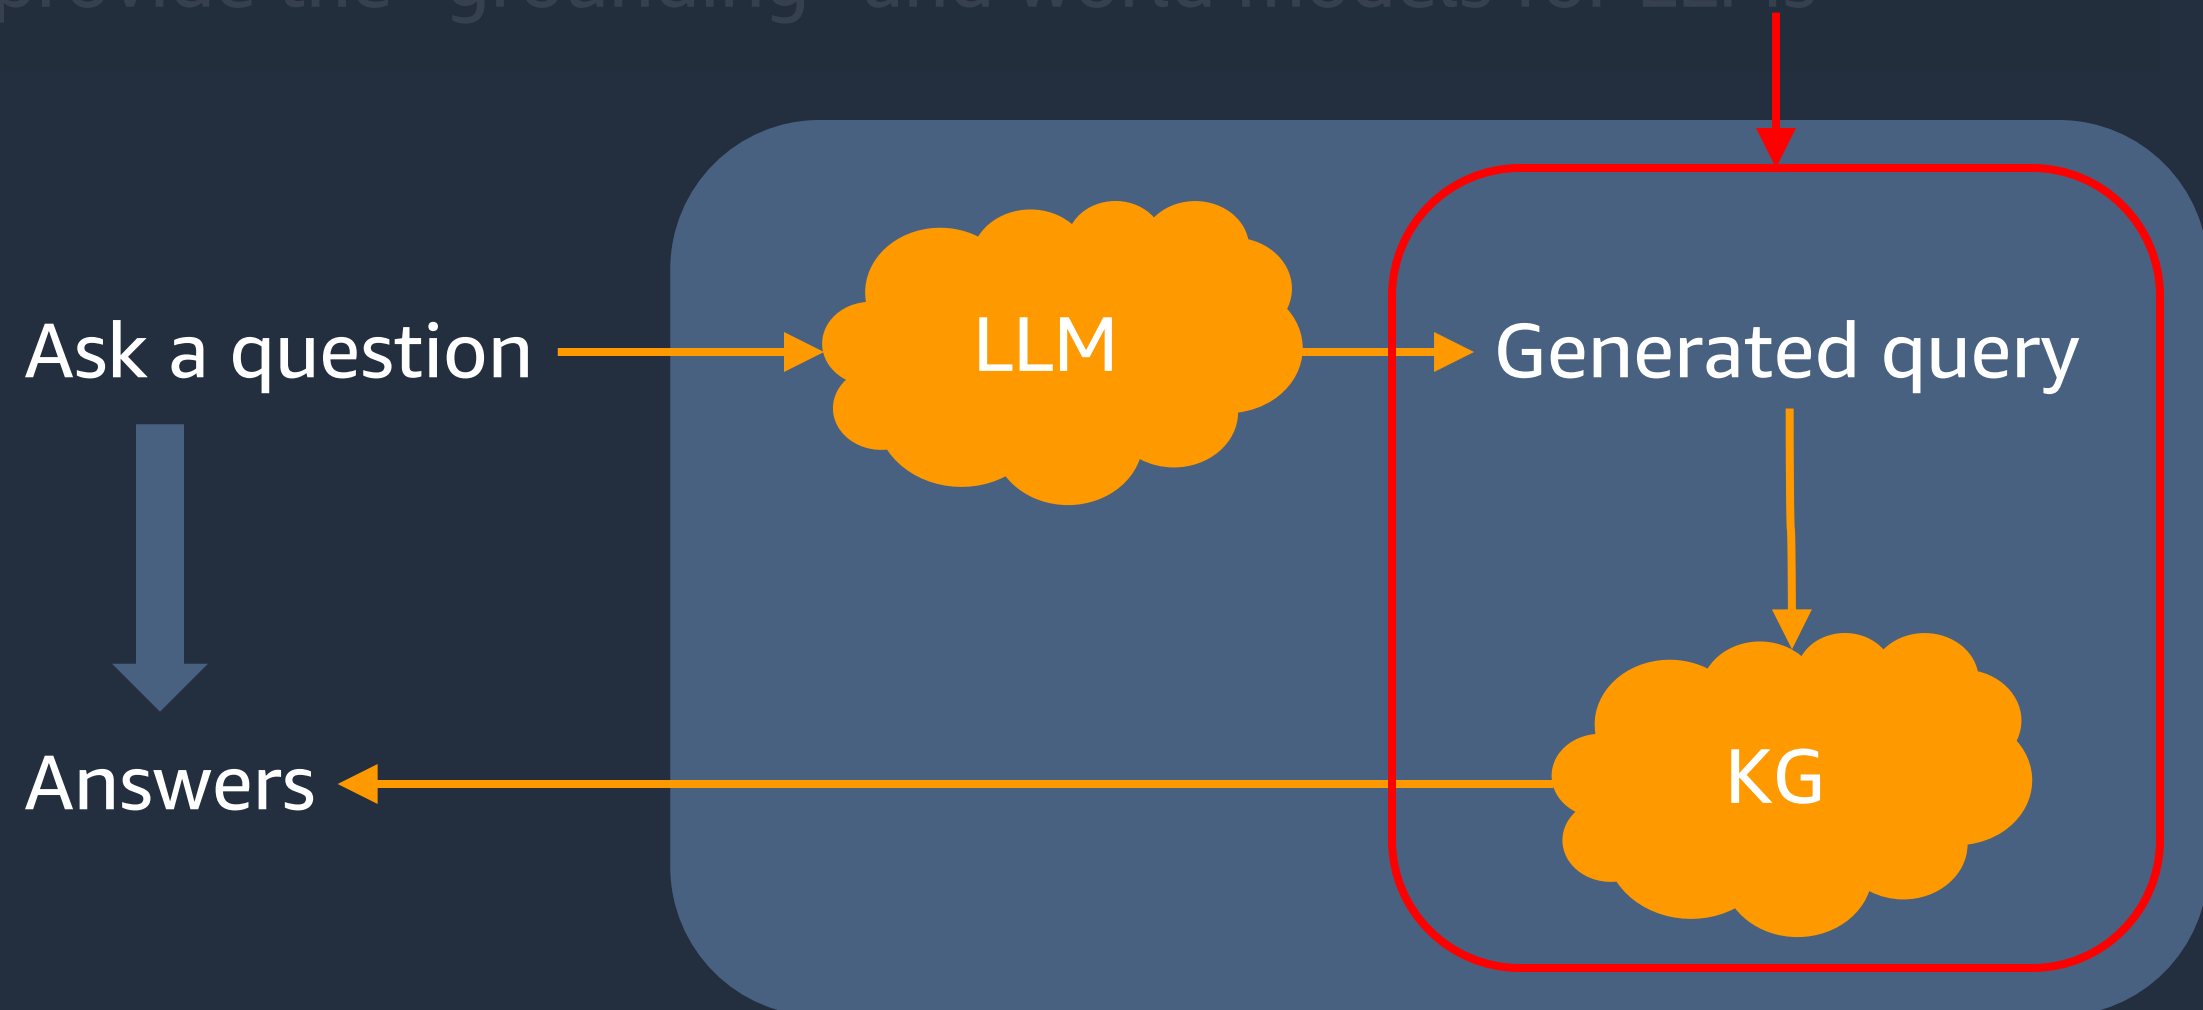
KGs could provide the “grounding” and world models for LLMs



# Knowledge graphs?

There are already promising results in letting LLMs write graph queries  
KGs could provide the “grounding” and world models for LLMs

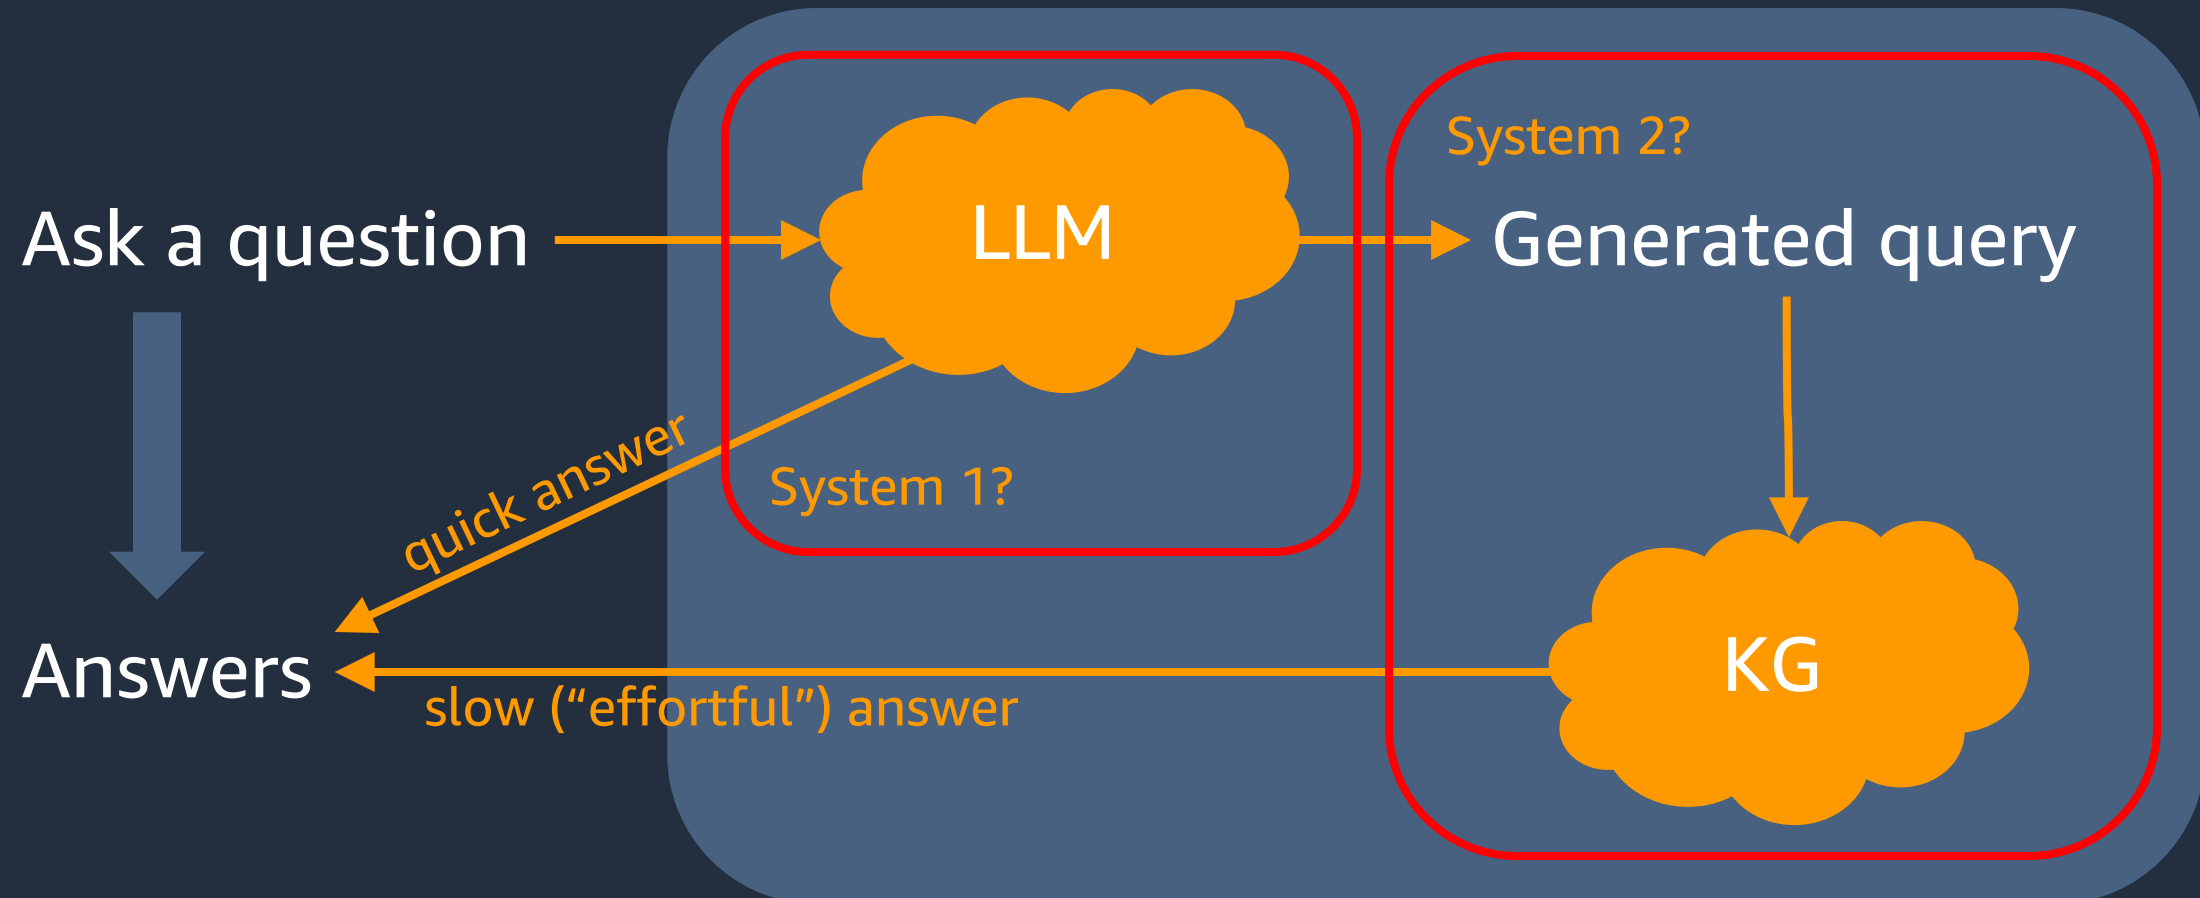
*Not just queries, but also (arbitrarily complex) ontological reasoning*



# Knowledge graphs?

“System 1” and “System 2” [Stanovich & West, Kahneman]...?

Also, “architecture” of intelligence and consciousness [Sloman, Dennett, ...]



# What “scarcities” do we now experience with ML & LLMs?

At least these:

- correct, truthful results
- accountability
- explainability

In the past, unaddressed scarcity resulted in unkept promises and backlashes against the whole AI industry

What is AI anyway?

I will come back to this...

AI has suffered from scarcity of computational ability (algorithms) and capability (CPU, memory), **driving innovation** also outside the field of AI

- machine translation in the 1960s
- running Lisp programs
- ML, neural networks, etc., require **lots** of CPU and memory

Amara's Law seems to apply to AI quite well:

*“We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run.”*

– Roy Amara (American futurologist, 1925-2007)

© 2022, Amazon Web Services, Inc. or its Affiliates.

aws

# Knowledge graphs & LLMs

LLMs (& ML in general)

LLMs + KGs

Truth, world models

Nope

Yes

Hallucinations

Yes

I hope not...

Curated knowledge

Maybe

Yes

Timely results

Some day

Yes

Explainable results

Nope

Yes

“Accountability”

Nope

Yes

# My “big” observation

AI systems are getting more sophisticated and more capable, and are more and more functioning autonomously

What we really need is a better understanding of how to operate such autonomous systems, how to limit their behavior appropriately, how to make them safe, etc.

Can we learn from human behavior?

- in human-to-human interaction, we tend to build a track record, and we build trust gradually

# Summary

AI has a history of “bloated expectations”

LLMs: amazing results, but caution is called for

Knowledge graphs could provide the necessary grounding for LLMs

We have much to learn about how to operate autonomous systems

# Thank you!

Contact:

- [ora@amazon.com](mailto:ora@amazon.com)
- Twitter: [@oralassila](https://twitter.com/oralassila)

