


Did Transformer models begin Terminator era?

Tomasz Pelczarski
Data&AI Professional,
Microsoft



THE
TERMINATOR

Math teachers protest against calculator use 1966



Should we fear?

- Should we ban AI?
- What is the risk?
- Can we control it?

ChatGPT and the Death of Education

Harvard Independent



AI by Microsoft

- How it has been created - Transparency Note
- How it should be used - Responsible AI standards
- Data, privacy, and security
- Implementation cheat-sheet
- Q&A

How Azure OpenAI works



Learn

Documentation

Training

Certifications

Q&A

Code Samples

Assessments

Shows

Events

Sign in

Azure OpenAI Service Documentation

Overview

What is Azure OpenAI?

Quotas and limits

Models

Pricing [↗]

What's new

Azure OpenAI FAQ

Quickstarts

Concepts

How-to

Tutorials

Responsible AI

Overview

Transparency note

Limited access

Code of conduct

Data, privacy, and security

Reference

Resources

Download PDF

Learn / Azure / Cognitive Services / Azure OpenAI Service /



Transparency Note for Azure OpenAI Service

Article • 05/19/2023 • 3 contributors

Feedback

In this article

[What is a Transparency Note?](#)

[The basics of the Azure OpenAI Models](#)

[Capabilities](#)

[Use cases](#)

[Show 4 more](#)

What is a Transparency Note?

An AI system includes not only the technology, but also the people who use it, the people who are affected by it, and the environment in which it's deployed. Creating a system that is fit for its intended purpose requires an understanding of how the technology works, what its capabilities and limitations are, and how to achieve the best performance.

Microsoft's Transparency Notes are intended to help you understand how our AI technology works, the choices system owners can make that influence system performance and behavior, and the importance of thinking about the whole system, including the technology, the people, and the environment. You can use Transparency

Additional resources

Documentation

[Code of Conduct for the Azure OpenAI Service](#)

Code of Conduct for Azure OpenAI Service

[Limited access to Azure OpenAI Service - Azure Cognitive Services](#)

This document details the limited access policy for Azure OpenAI Service

[Azure OpenAI Service - Azure OpenAI](#)

Learn how to use prompt engineering to optimize your work with Azure OpenAI Service.

[Show 5 more](#)



はい C'est vrai!
 Tak! **40+** Si.
 languages
 Efectivamente. **हाँ**
 You bet.

\$0

We gather it.
 We aggregate it.
 You utilize it.
And it's all free.

HOW BIG?
 WE'RE TALKING
BIG
 PETABYTES BIG

our story »

RAW DATA
METADATA
TEXT DATA

billions
 OF PAGES

trillions
 OF LINKS

say what?

7
 YEARS OF DATA

WebText

[Edit](#)

Introduced by Radford et al. in [Language Models are Unsupervised Multitask Learners](#)

WebText is an internal OpenAI corpus created by scraping web pages with emphasis on document quality. The authors scraped all outbound links from Reddit which received at least 3 karma. The authors used the approach as a heuristic indicator for whether other users found the link interesting, educational, or just funny.

WebText contains the text subset of these 45 million links. It consists of over 8 million documents for a total of 40 GB of text. All Wikipedia documents were removed from WebText since it is a common data source for other datasets.

Benchmarks

[Edit](#)

No benchmarks yet. [Start a new benchmark](#) or [link an existing one](#).

Papers

"I'm not the cleverest man in the world, but like they say in French: **Je ne suis pas un imbecile (I'm not a fool)**."

In a now-deleted post from Aug. 16, Soheil Eid, Tory candidate in the riding of Joliette, wrote in French: "**Mentez mentez, il en restera toujours quelque chose**," which translates as, "**Lie lie and something will always remain**."

"I hate the word 'perfume,'" Burr says. 'It's somewhat better in French: 'parfum.'


If listened carefully at 29:55, a conversation can be heard between two guys in French: "**-Comment on fait pour aller de l'autre coté? -Quel autre coté?**", which means "**- How do you get to the other side? - What side?**".

If this sounds like a bit of a stretch, consider this question in French: **As-tu aller au cinéma?**, or **Did you go to the movies?**, which literally translates as Have-you to go to movies/theater?

"**Brevet Sans Garantie Du Gouvernement**", translated to English: "**Patented without government warranty**".

Table 1. Examples of naturally occurring demonstrations of English to French and French to English translation found throughout the WebText training set.

Usage



Pros and Cons

What is cool

1. Very natural.
2. Very simple to use and easy to access – [Edge Internet Explorer](#).
3. Save large amount of time in creation of artifacts

What are the watch out

1. Verification becomes critical.
2. Answers may express bias, misinformation or inappropriate content.
3. Customer relationships - can not replace human empathy and support – feedback channels.

Legal works in EU and the US

- EU Parliament

Generative foundation models **should ensure transparency about the fact the content is generated by an AI system**, not by humans. These specific requirements and obligations do not amount to considering foundation models as high risk AI systems, but should guarantee that the objectives of this Regulation to ensure a high level of protection of fundamental rights, health and safety, environment, democracy and rule of law are achieved.

[Microsoft Word - Consolidated document CA_IMCOLIBE_AI_ACT.docx \(europa.eu\)](#)

- The US

- National Institute of Standards and Technology's (NIST) AI Risk Management Framework (RMF), with a final release on January 26, 2023.
- Food and Drug Administration (FDA)
- Federal Trade Commission (FTC) - advertising

[The EU and U.S. diverge on AI regulation: A transatlantic comparison and steps to alignment \(brookings.edu\)](#)

Azure OpenAI

Top Capabilities and Use Cases



Content generation

Call center analytics: automatically generate responses to customer inquiries

Generate personalized UI for your website



Summarization

Call center analytics: summary of customer support conversation logs

Subject matter expert document: summarization (e.g. Financial reporting, analyst articles)

Social media trends summarization



Code generation

Convert natural language to SQL (or vice versa) for telemetry data

Convert natural language to query proprietary data models

Code documentation



Semantic search

Search reviews for a specific product/service

Information discovery and knowledge mining

Examples of multiple model use cases

End to end call center analytics: classification, sentiment, entity extraction, summarization and email generation

Customer 360: hyper-personalisation using timely summarization of customer queries & trends, search, and content generation

Business process automation: search through structured & unstructured documentation, generate code to query data models, content generation

| Microsoft Azure Cloud

Runs on trust

Your data is your data

Data is stored encrypted in *your Azure subscription*

Your data from any fine-tuning is not used to train the foundation AI models

Azure OpenAI Service provisioned in *your Azure subscription*

Model fine tuning stays in *your Azure subscription* and never moves into the foundation AI models

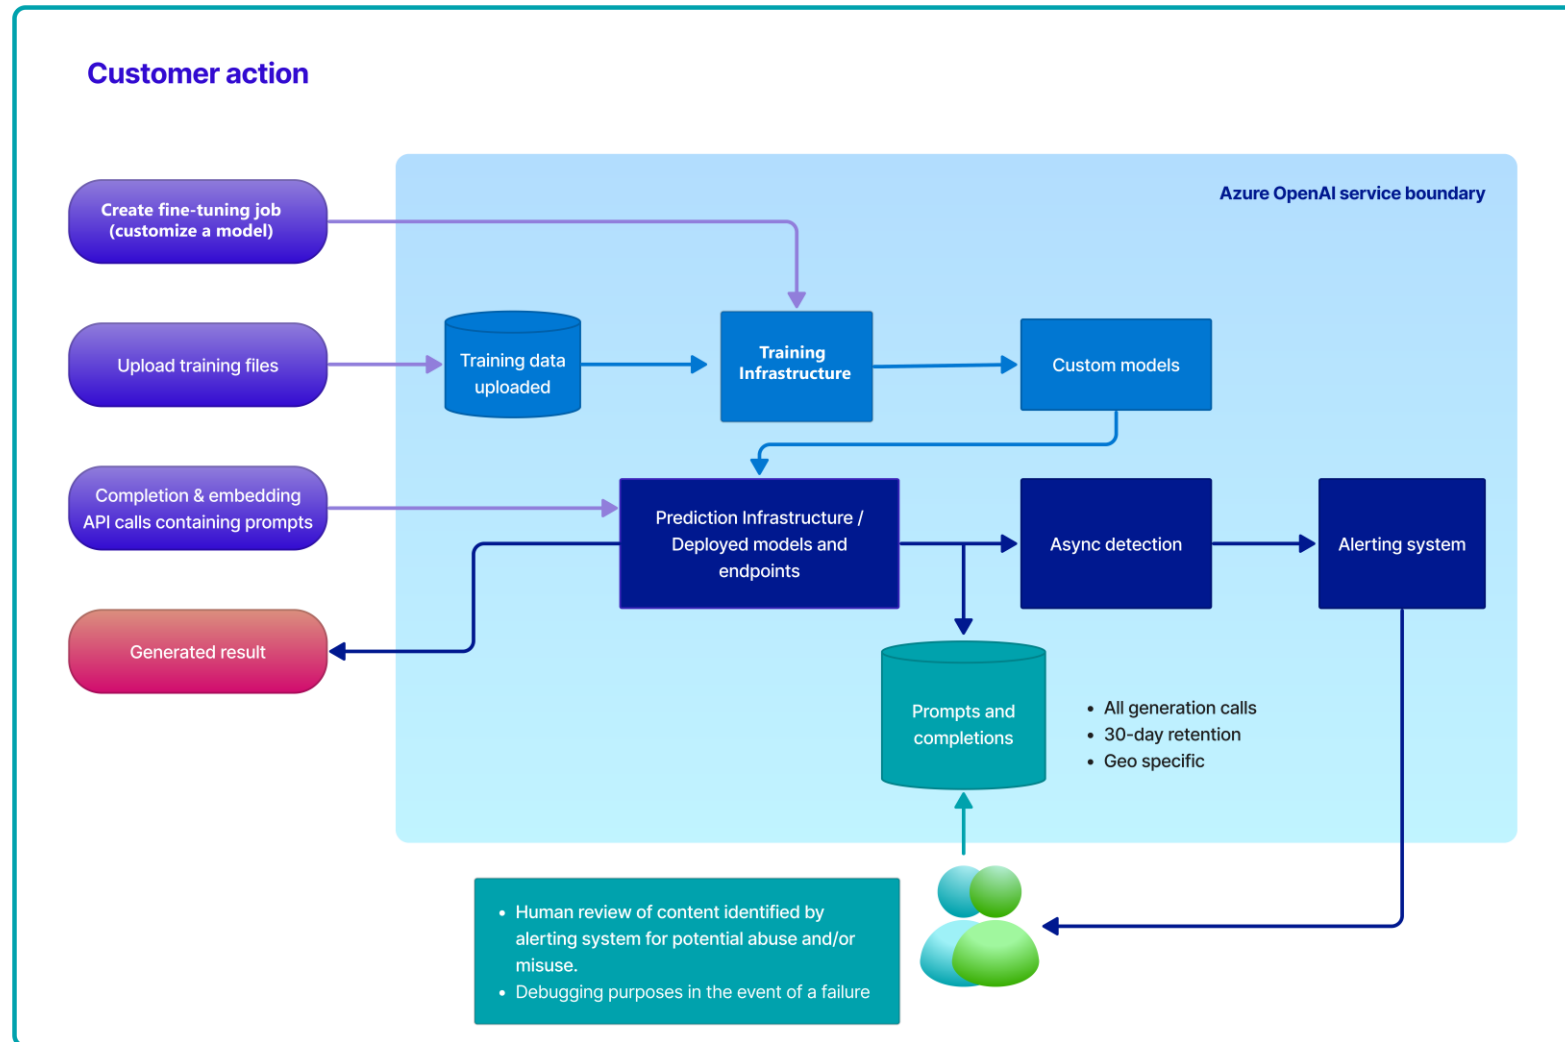
Your data is protected by the most comprehensive enterprise compliance and security controls

Encrypted with Customer Managed Keys

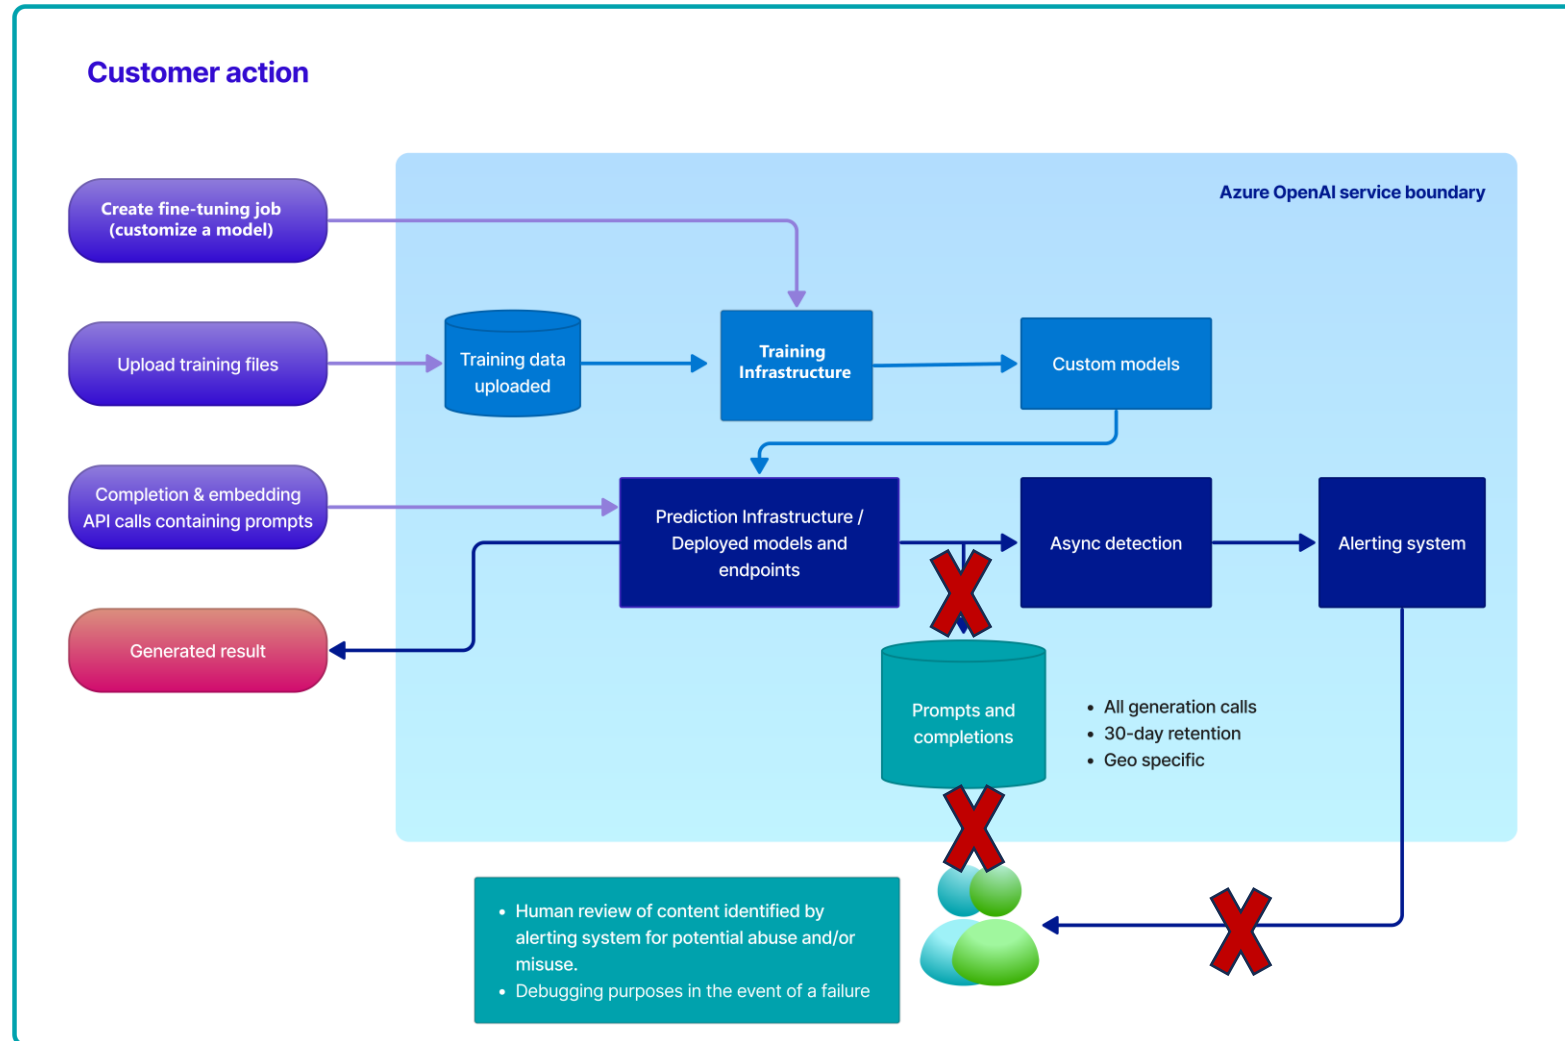
Private Virtual Networks, Role Based Access Control

Soc2, ISO, HIPPA, CSA STAR Compliant

Human-in-the-loop review



Opting - out



Responsible AI implementation cheat-sheet



Policy

Responsible AI Standard

The Microsoft Responsible AI Standard is our internal playbook for responsible AI. It shapes the way in which we create AI systems, by guiding how we design, build, and test them.

[Get the Responsible AI Standard >](#)

[Get the Responsible AI Reference Guide >](#)



Management Tool

Responsible AI Impact Assessment Template

The Responsible AI Impact Assessment Template is the product of a multi-year effort to define a process for assessing the impact an AI system may have on people, organizations, and society.

[Download the Impact Assessment Template >](#)




Guideline

Responsible AI Impact Assessment Guide

This resource provides activities and guidance for teams working through the Responsible AI Impact Assessment Template to help frame and support conversations about Responsible AI.

[Read the Impact Assessment Guide >](#)



Communication

Transparency Notes

Transparency Notes allow us to communicate the intended uses, capabilities, and limitations of our AI platform systems to customers, building trust and enabling our customers to build more responsible AI products and services on top of our platforms.

[Explore Transparency Notes >](#)

The cheat sheet

[Download the Impact Assessment Template >](#)

[GitHub - microsoft/ResponsibleAIAccelerator: Repo to hold examples of responsible model assessment for a variety of different verticals such as healthcare and financial services](#)

Microsoft Responsible AI Impact Assessment Template

Technology readiness assessment

2.5 Indicate with an "X" the description that best represents the system regarding this intended use.

Select one	Technology Readiness
	The system includes AI supported by basic research and has not yet been deployed to production systems at scale for similar uses.
	The system includes AI supported by evidence demonstrating feasibility for uses similar to this intended use in production systems.
	This is the first time that one or more system component(s) are to be validated in relevant environment(s) for the intended use. Operational conditions that can be supported have not yet been completely defined and evaluated.
	This is the first time the whole system will be validated in relevant environment(s) for the intended use. Operational conditions that can be supported will also be validated. Alternatively, nearly similar systems or nearly similar methods have been applied by other organizations with defined success.
	The whole system has been deployed for all intended uses , and operational conditions have been qualified through testing and uses in production.

Task complexity

2.6 Indicate with an "X" the description that best represents the system regarding this intended use.

Select One	Task Complexity
	Simple tasks , such as classification based on few features into a few categories with clear boundaries. For such decisions, humans could easily agree on the correct answer, and identify mistakes made by the system. For example, a natural language processing system that checks spelling in documents.
	Moderately complex tasks , such as classification into a few categories that are subjective. Typically, ground truth is defined by most evaluators arriving at the same answer. For example, a natural language processing system that autocompletes a word or phrase as the user is typing.
	Complex tasks , such as models based on many features, not easily interpretable by humans, resulting in highly variable predictions without clear boundaries between decision criteria. For such decisions, humans would have a difficult time agreeing on the best answer, and there may be no clearly incorrect answer. For example, a natural language processing system that generates prose based on user input prompts.



Microsoft Responsible AI Impact Assessment Template

FOR EXTERNAL RELEASE

June 2022

The Responsible AI Impact Assessment is the product of a multi-year effort at Microsoft to define a process for assessing the impact a system may have on people, organizations, and society. We are releasing our Impact Assessment Template externally to share what we have learned, invite feedback from others, and contribute to discussion about building better norms and practices around AI.

We invite your feedback on our approach at <https://aka.ms/ResponsibleAIQuestions>

Future of AI

TED Ideas worth spreading


WATCH DISCOVER ATTEND



How AI could save (not destroy) education

Full interview: "Godfather of artificial intelligence" talks impact and potential of AI

Terminator can be GOOD 😊



Thank you!

Tomasz Pelczarski

tomaszpe@microsoft.com

