ATTACKS ON GEN AI DATA

AND USING VECTOR ENCRYPTION TO STOP THEM



IRONCORE LABS



Bob Wall

CTO, IronCore Labs

bob.wall@ironcorelabs.com

@bithead_bob

@bobwall23@infosec.exchange

bobwall23



Patrick Walsh

CEO, IronCore Labs

patrick.walsh@ironcorelabs.com

@zmre

@mastodon.social@zmre

zmre

Large Language Model

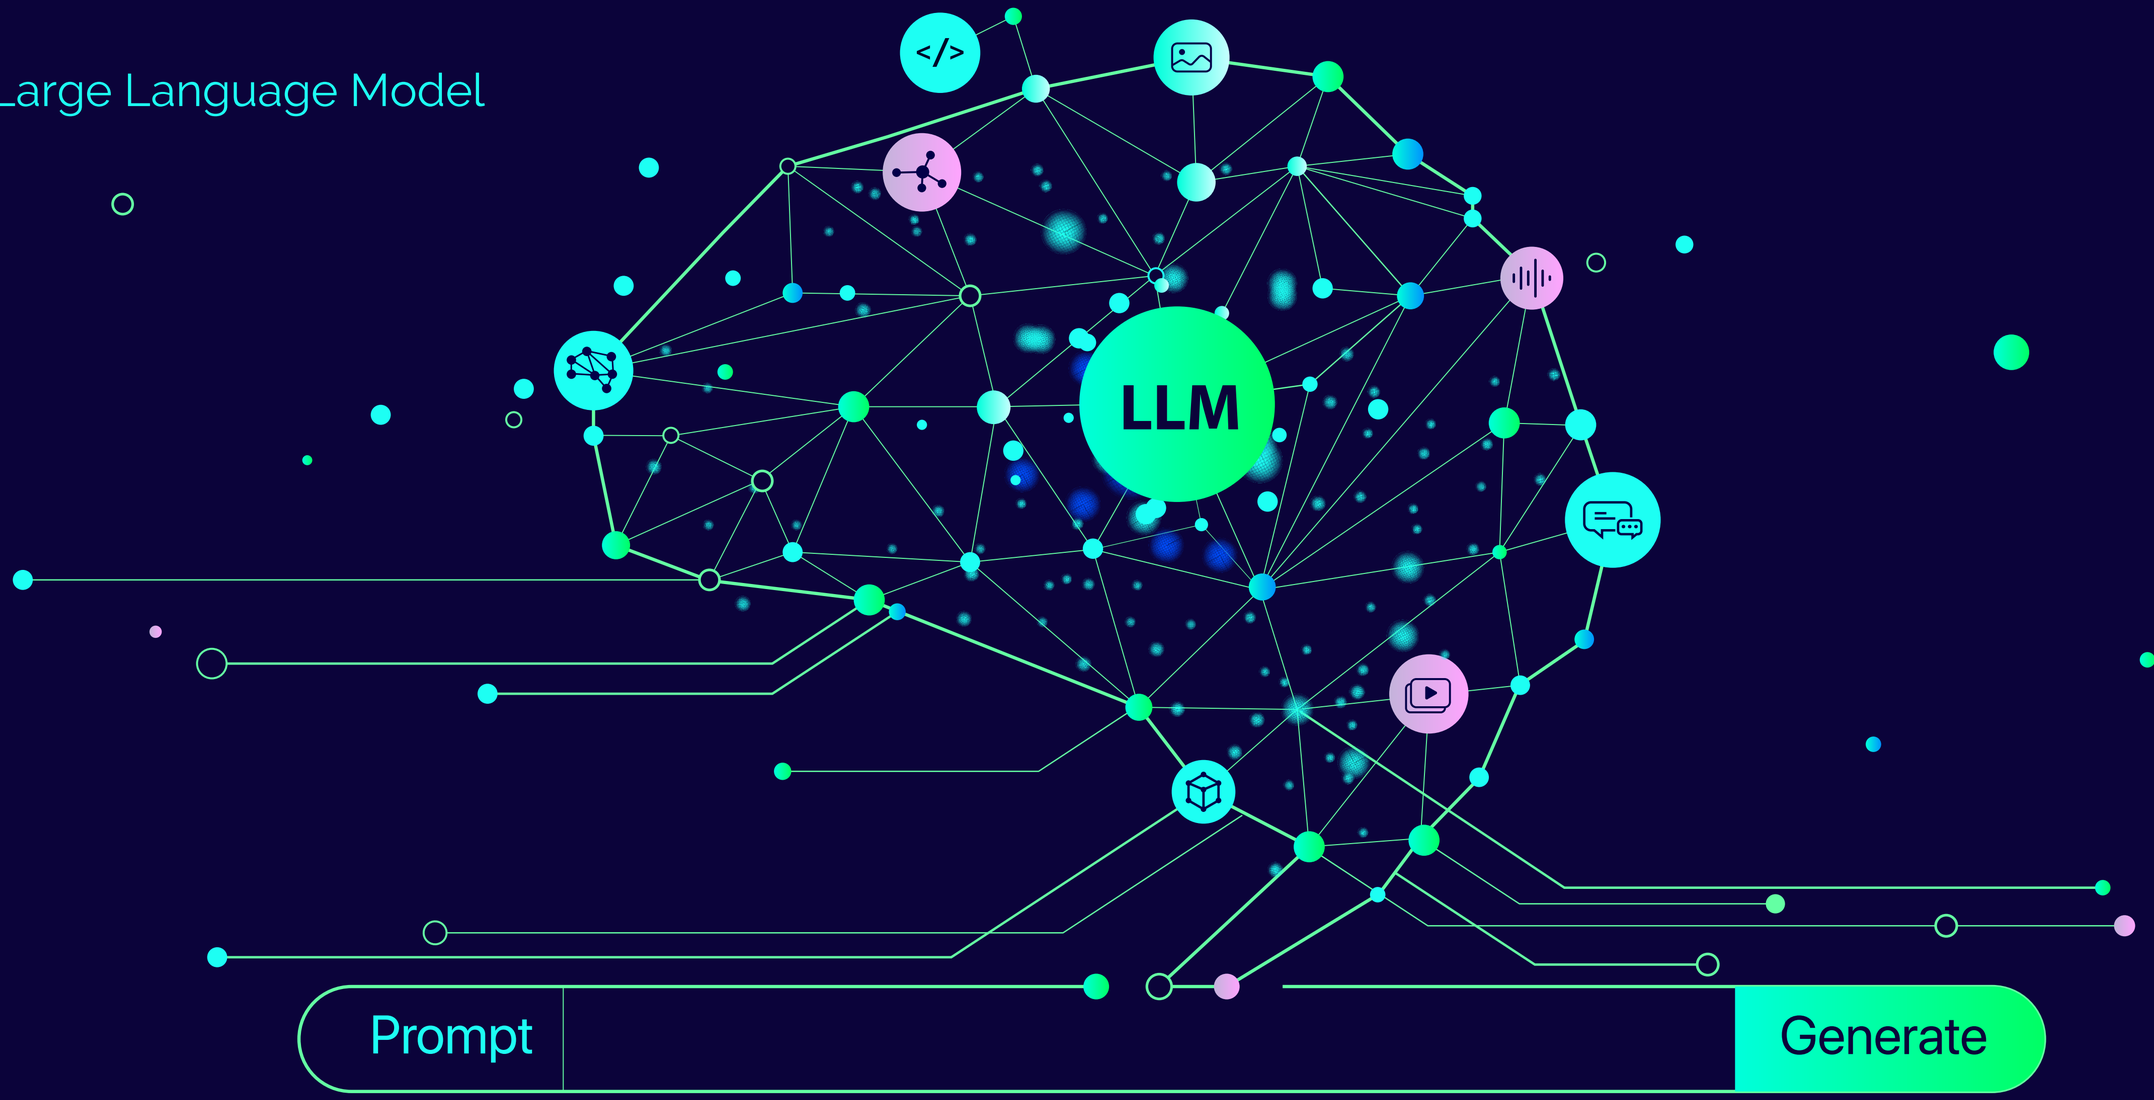

LLM **PROBLEMS**

1. Hallucinated answers

- LLMs predict plausible "completions." These are often factually incorrect.

LLM **PROBLEMS**

1. **Hallucinated answers**
   - LLMs predict plausible "completions." These are often factually incorrect.
2. **Stale training data**
   - It takes a long time to train an LLM and that means that the information they learn is often several years old.

LLM **PROBLEMS**

1. **Hallucinated answers**
   - LLMs predict plausible "completions." These are often factually incorrect.
2. **Stale training data**
   - It takes a long time to train an LLM and that means that the information they learn is often several years old.
3. **Lack of private data**
   - Asking questions about the news is one thing, but to query your own data, you need RAG.

WHAT IS RAG?

(HINT: NOT CLOTH)



R•A•G | rag

noun

1. Acronym for **Retrieval-Augmented Generation**, a technique for enhancing the accuracy, utility, and timeliness of Large Language Models (LLMs).

2. A terrible name that the authors of the paper that coined it have said they regret.

RAG WORKFLOWS



Chat App

RAG WORKFLOWS

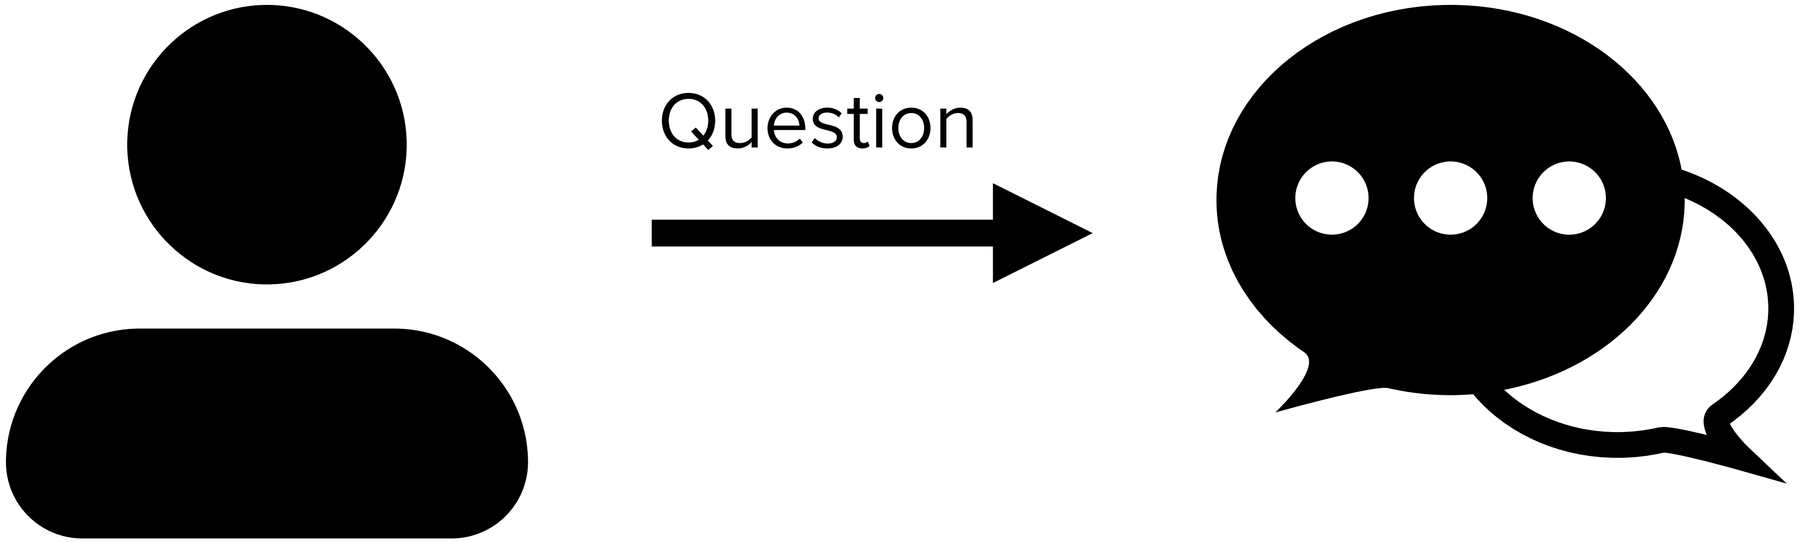

RAG WORKFLOWS



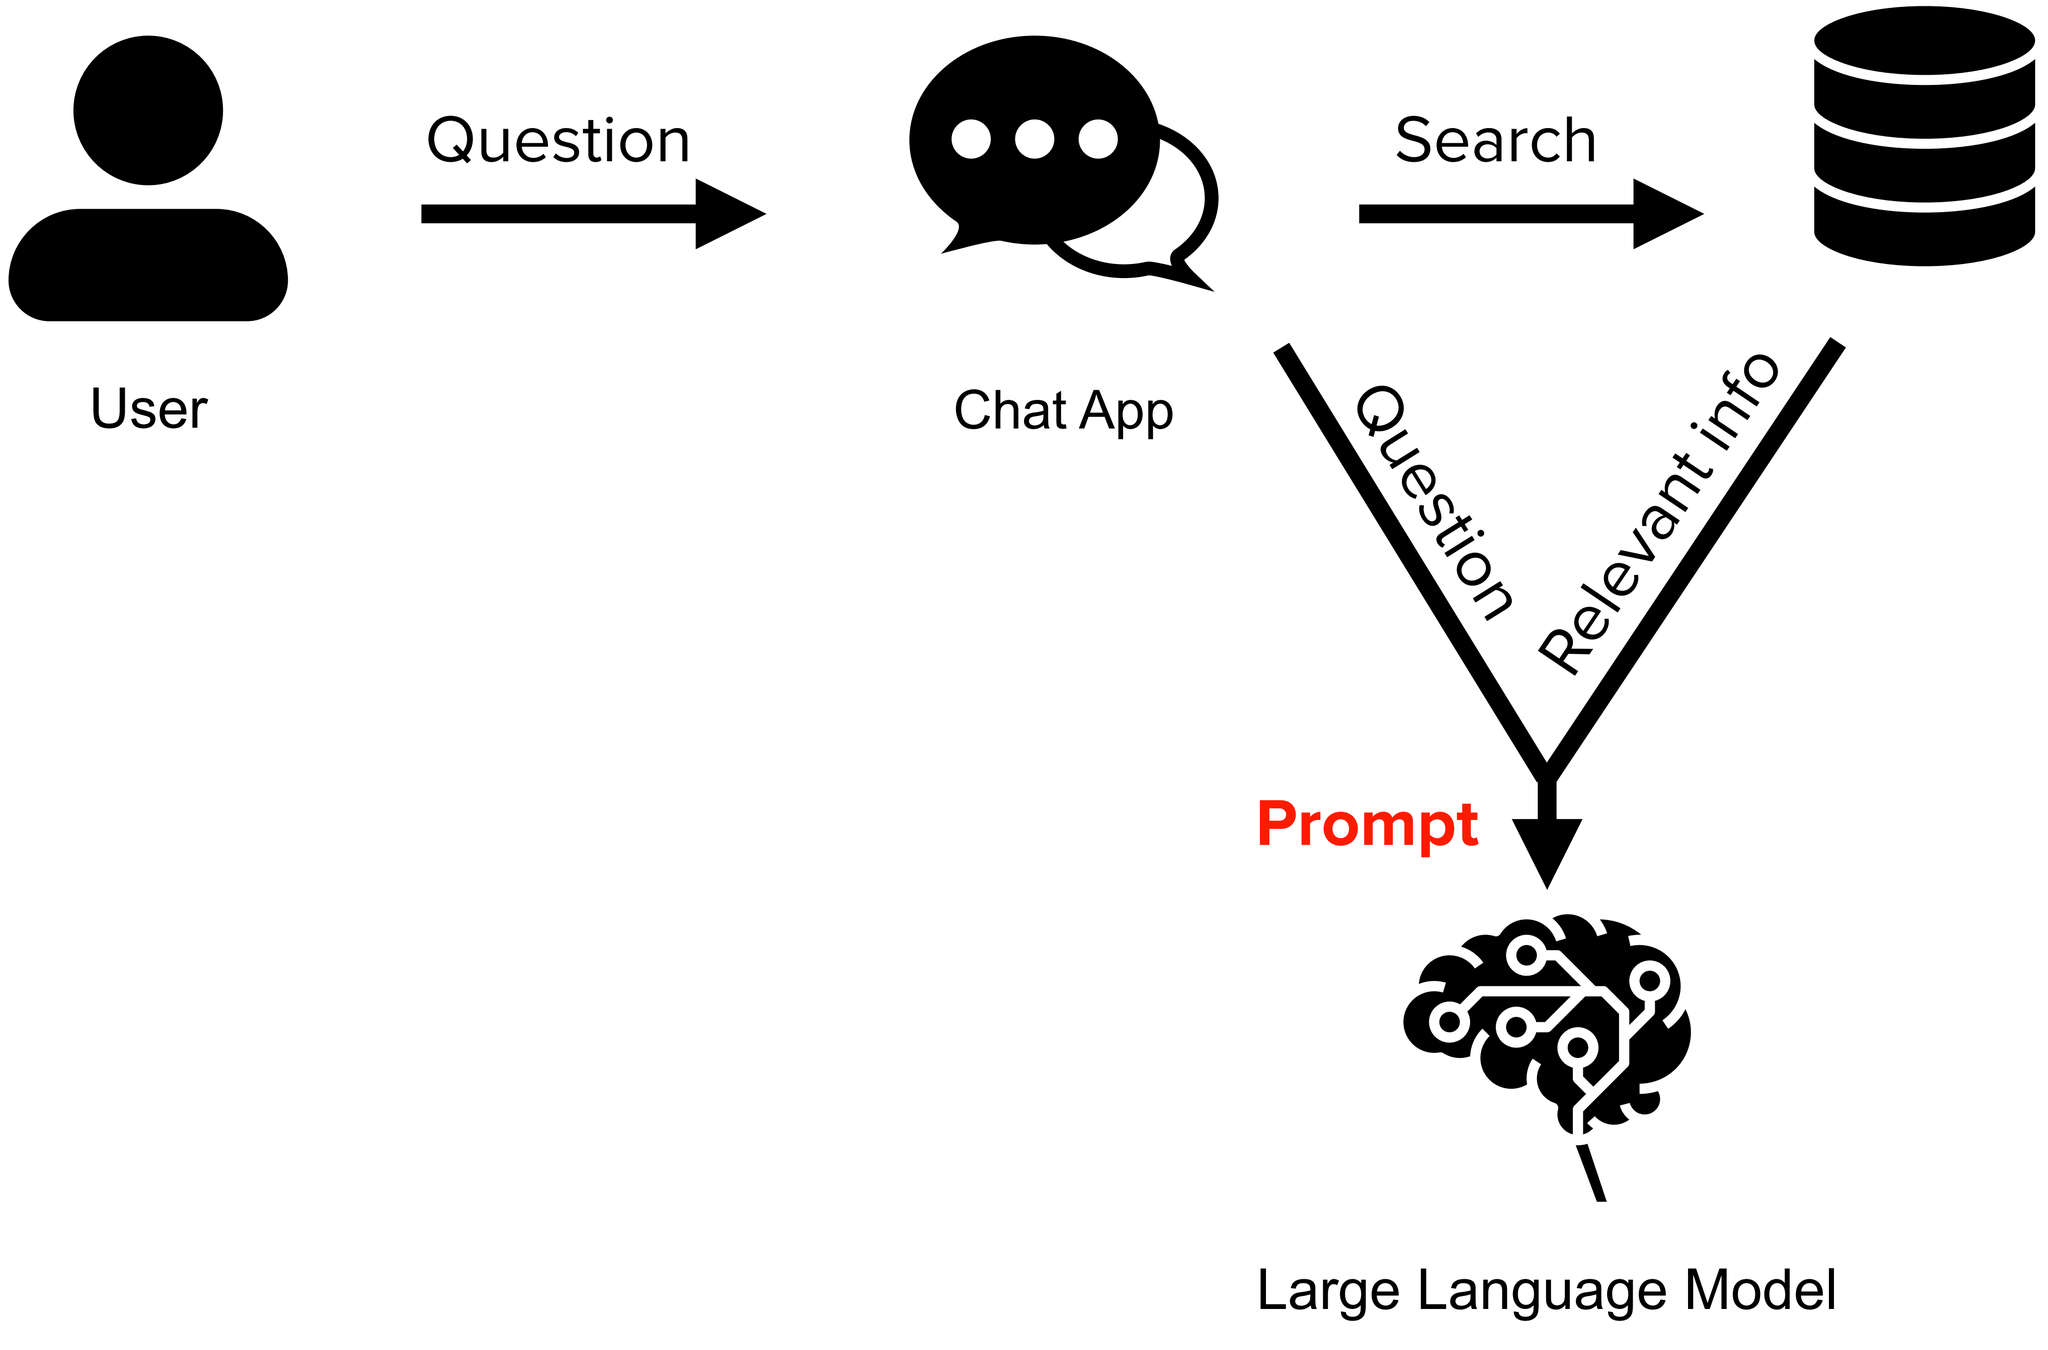
Question
Search
User
Chat App
Question
Relevant info
Prompt
Large Language Model

EXAMPLE PROMPTS

SHOWING THE AUGMENTED DATA

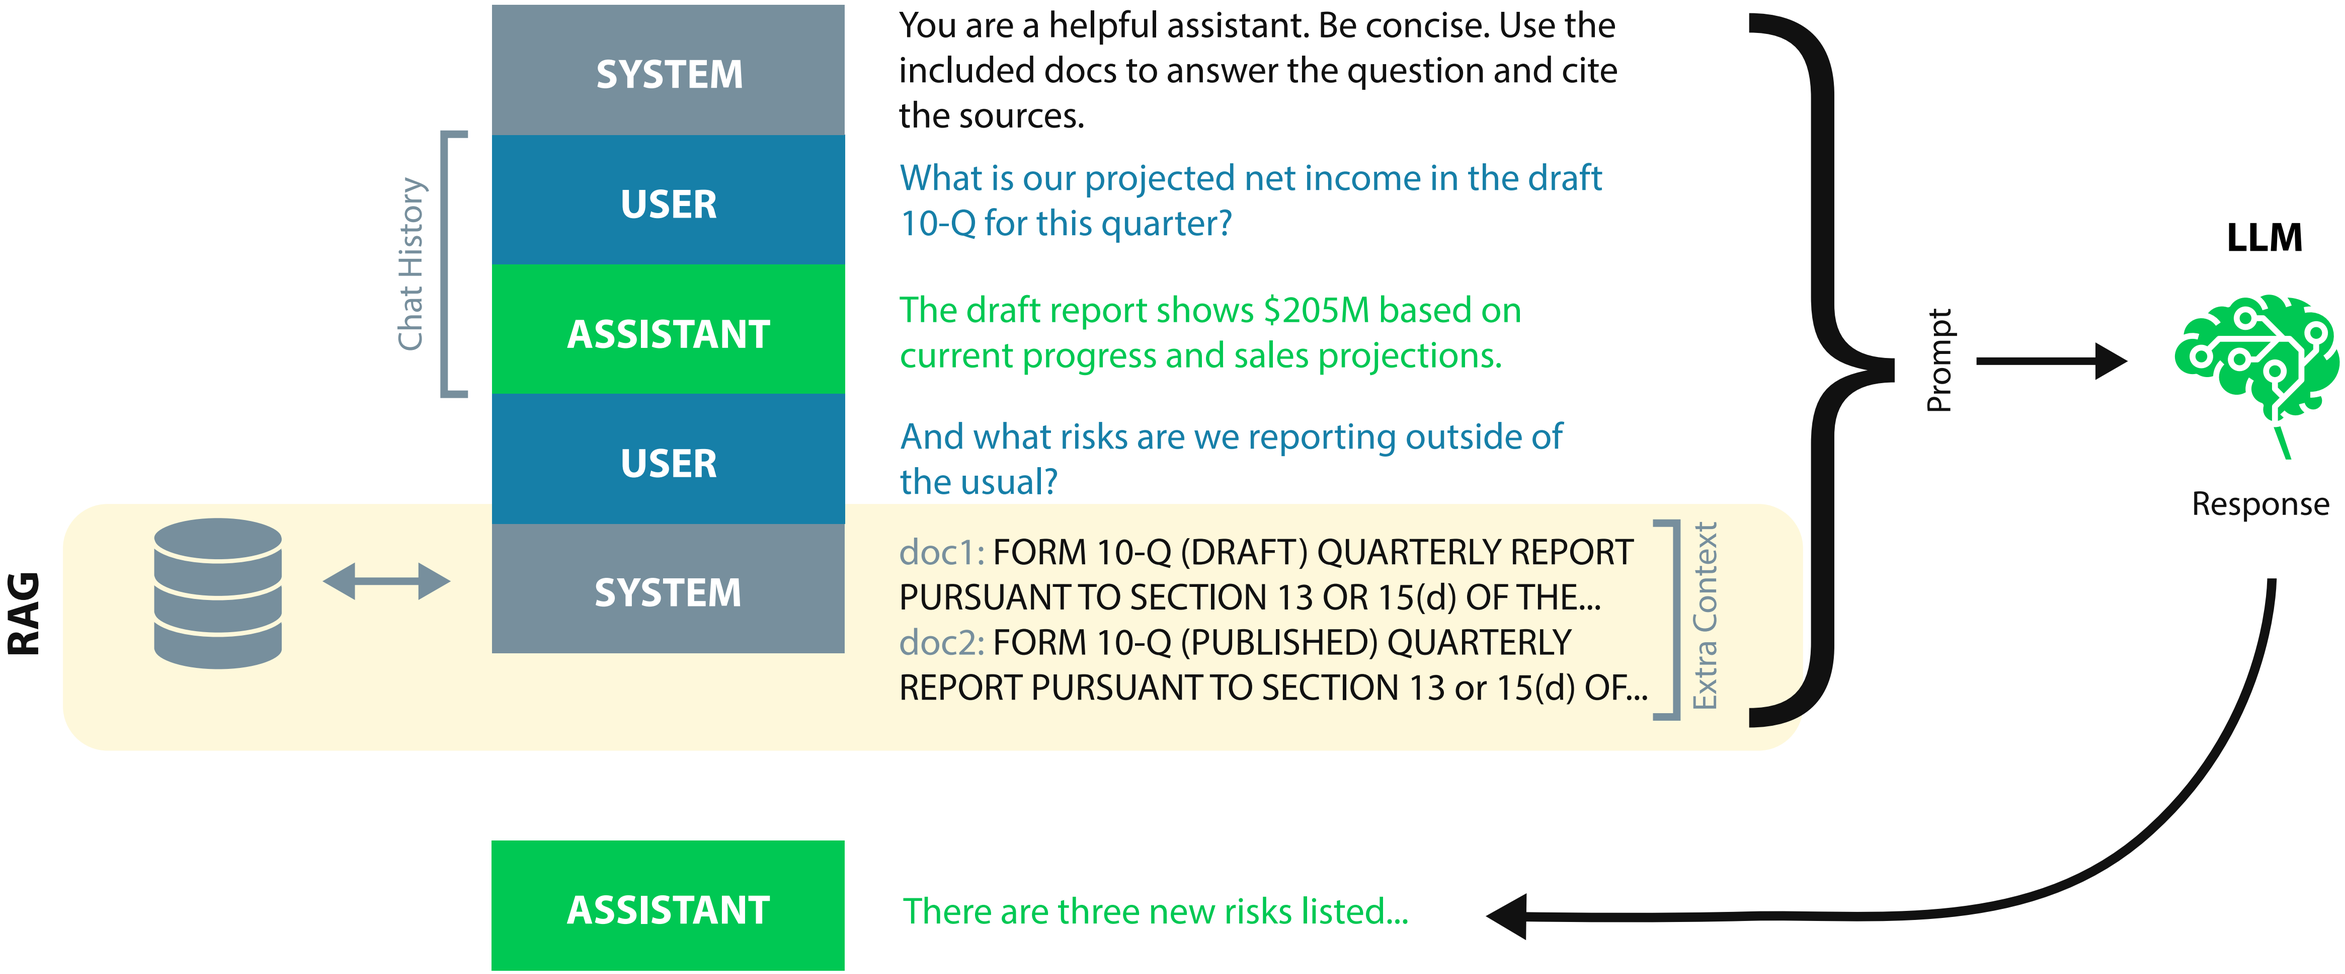

AI SEARCH

NATURAL LANGUAGE DOCUMENT RETRIEVAL





Also known as:

- Semantic search
- Neural search
- Vector search
- Natural language search
- Meaning-based search
- Intelligent search
- Intent-based search

SEMANTIC SEARCH

MEANING-BASED SIMILARITY SEARCH

CONTEXTUAL MEANING EXAMPLE:

ARM



(group)



(body part)

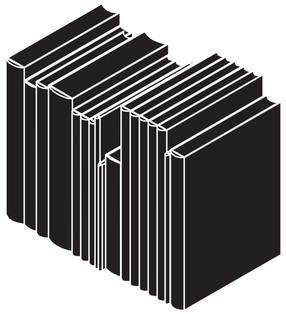

(armed with...)



(strong arm...)



(to give a weapon)

"How to build strong arms" vs. "How to strong arm"

➔ weight lifting exercises

➔ mafia pressure tactics

AI SEARCH

How it works with models

When in the Course of human events, it becomes necessary for one people to dissolve the political bands which have connected them with another, and to assume among the powers of the earth, the separate and equal station to which the Laws of Nature and of Nature's God entitle them, a decent respect to the opinions of mankind requires that they should declare the causes which impel them to the separation.

We hold these truths to be self-evident, that all men are created equal, that they are endowed by their Creator with certain unalienable. Rights, that among these are Life, Liberty and the pursuit of Happiness. That to secure these rights, Governments are instituted among Men, deriving their just powers from the consent of the governed.

Input Document

AI SEARCH

How it works with models

When in the Course of human events, it becomes necessary for one people to dissolve the political bands which have connected them with another, and to assume among the powers of the earth, the separate and equal station to which the Laws of Nature and of Nature's God entitle them, a decent respect to the opinions of mankind requires that they should declare the causes which impel them to the separation.

We hold these truths to be self-evident, that all men are created equal, that they are endowed by their Creator with certain unalienable. Rights, that among these are Life, Liberty and the pursuit of Happiness. That to secure these rights, Governments are instituted among Men, deriving their just powers from the consent of the governed.

Input Document

AI SEARCH

How it works with models

When in the Course of human events, it becomes necessary for one people to dissolve the political bands which have connected them with another, and to

assume among the powers of the earth, the separate and equal station to which the Laws of Nature and of Nature's God entitle them, a decent respect to the

opinions of mankind requires that they should declare the causes which impel them to the separation.

We hold these truths to be self-evident, that all men are created equal, that they are endowed by their

Creator with certain unalienable. Rights, that among these are Life, Liberty and the pursuit of Happiness.

That to secure these rights, Governments are instituted among Men, deriving their just powers from the consent of the governed.

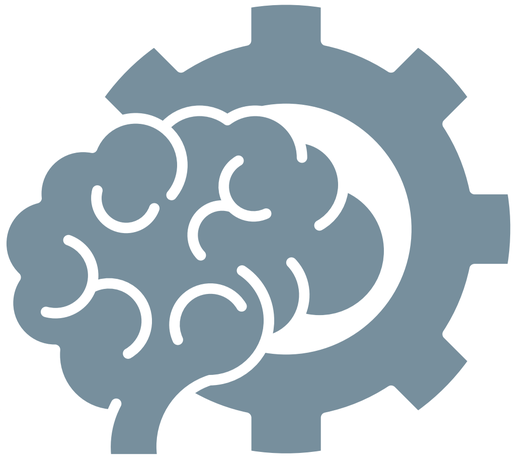

Input Document

Text
Embedding Model

AI SEARCH

How it works with models

When in the Course of human events, it becomes necessary for one people to dissolve the political bands which have connected them with another, and to assume among the powers of the earth, the separate and equal station to which the Laws of Nature and of Nature's God entitle them, a decent respect to the opinions of mankind requires that they should declare the causes which impel them to the separation.

We hold these truths to be self-evident, that all men are created equal, that they are endowed by their Creator with certain unalienable. Rights, that among these are Life, Liberty and the pursuit of Happiness.

That to secure these rights, Governments are instituted among Men, deriving their just powers from the consent of the governed.

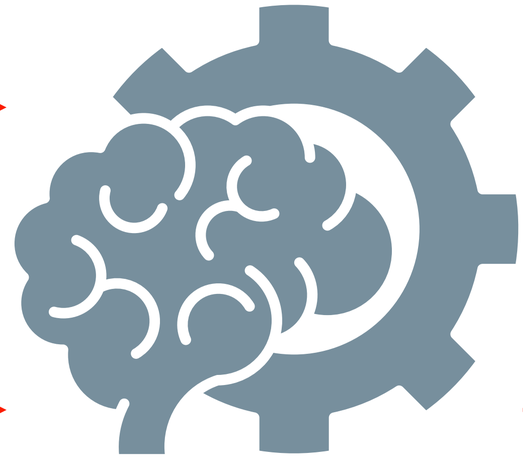

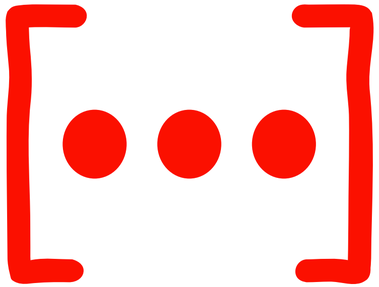

Input Document

Text Embedding Model

Vector Embeddings

VECTOR EMBEDDING

How it works with models

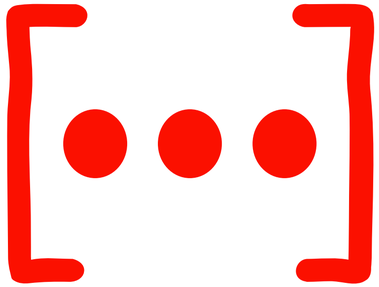

[0.10334, -0.842, 0.004, -1.736, ...]

It's a list of hundreds or thousands of small numbers.

OpenAI Model	Embedding Dimension
ada v2	1,536
text-embedding-3-small	512 or 1,536
text-embedding-3-large	256 or 1,024 or 3,072

VECTOR EMBEDDING

A VISUAL REPRESENTATION

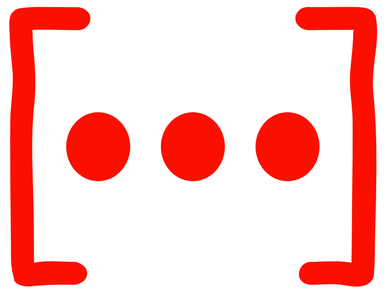

A=[3, 2, 1]

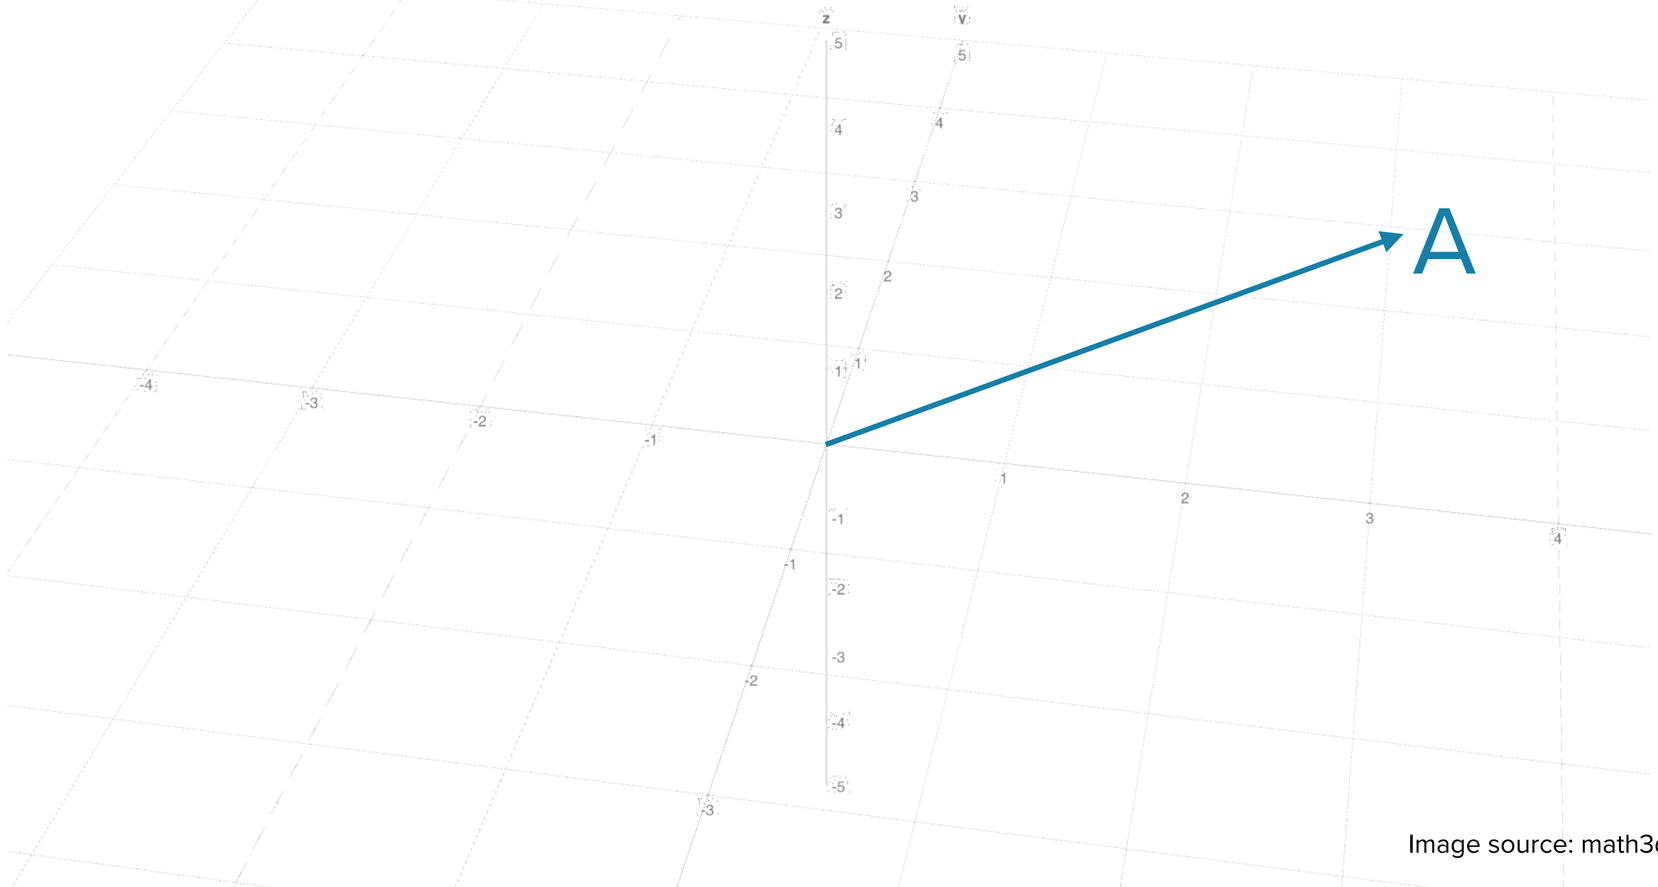

Image source: math3d.org/vectors

VECTOR EMBEDDING

A VISUAL REPRESENTATION

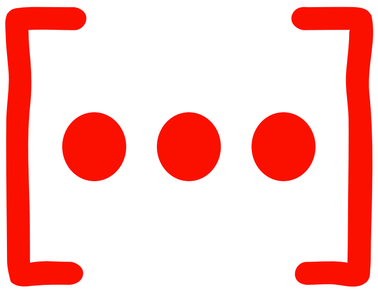

A=[3, 2, 1]

B=[4, 4, 1]

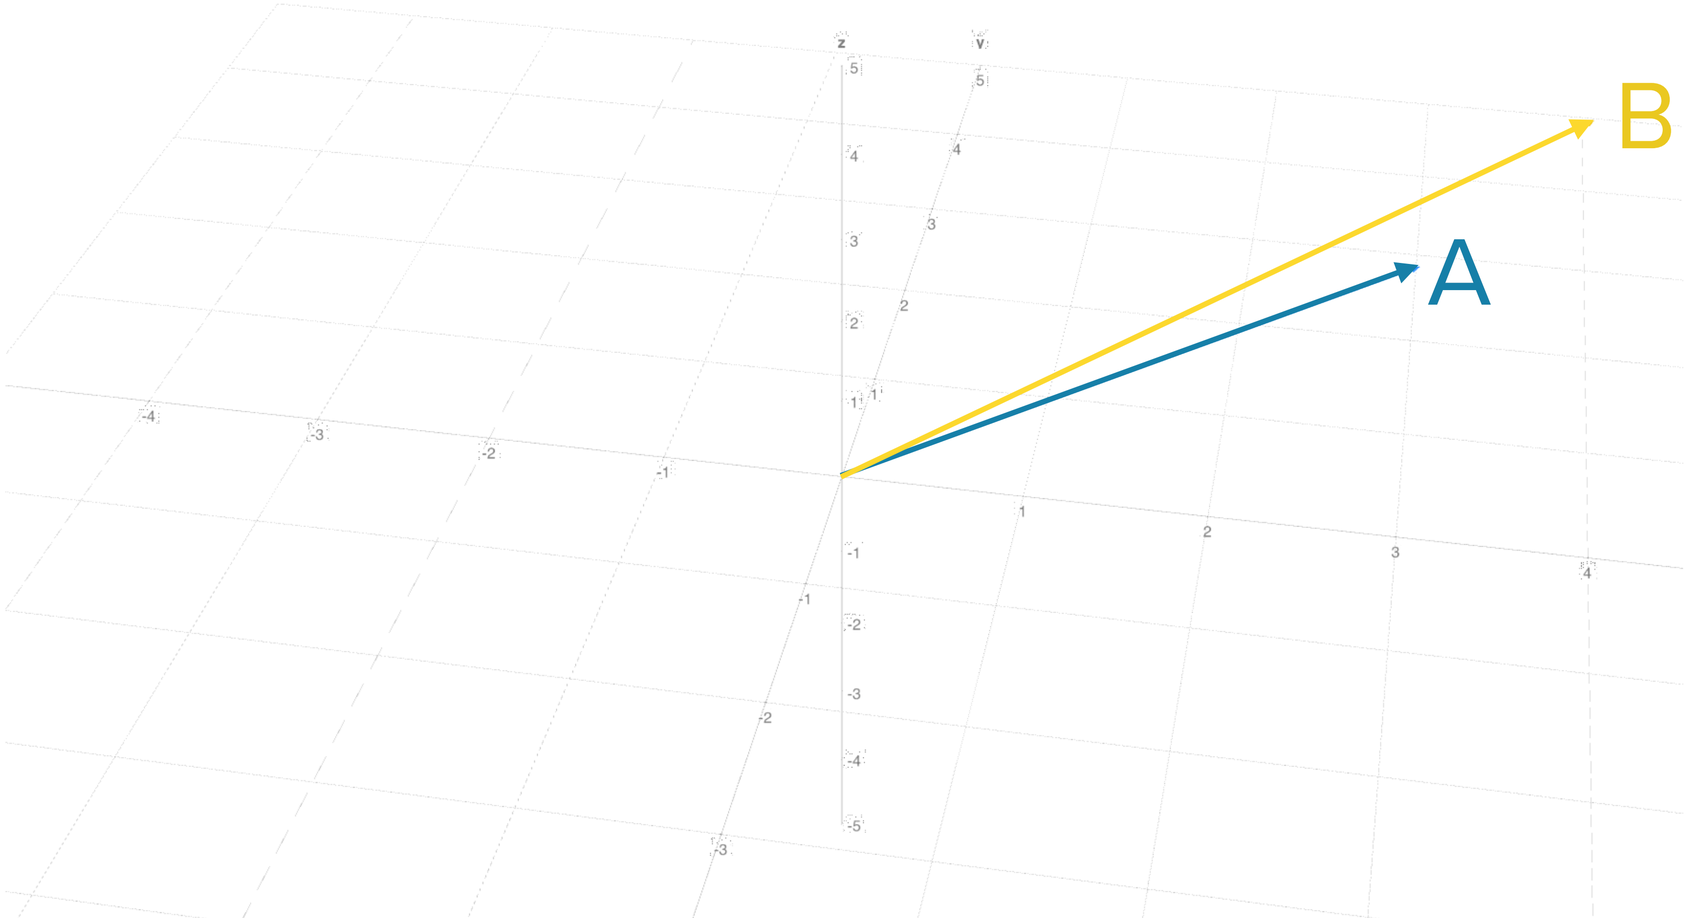

VECTOR EMBEDDING

A VISUAL REPRESENTATION

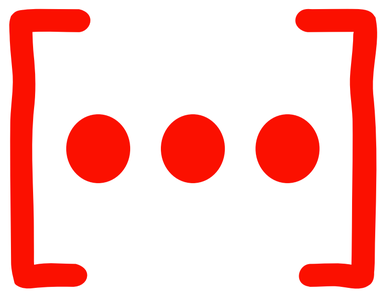

A=[3, 2, 1]

B=[4, 4, 1]

C=[-2, 1, 3]

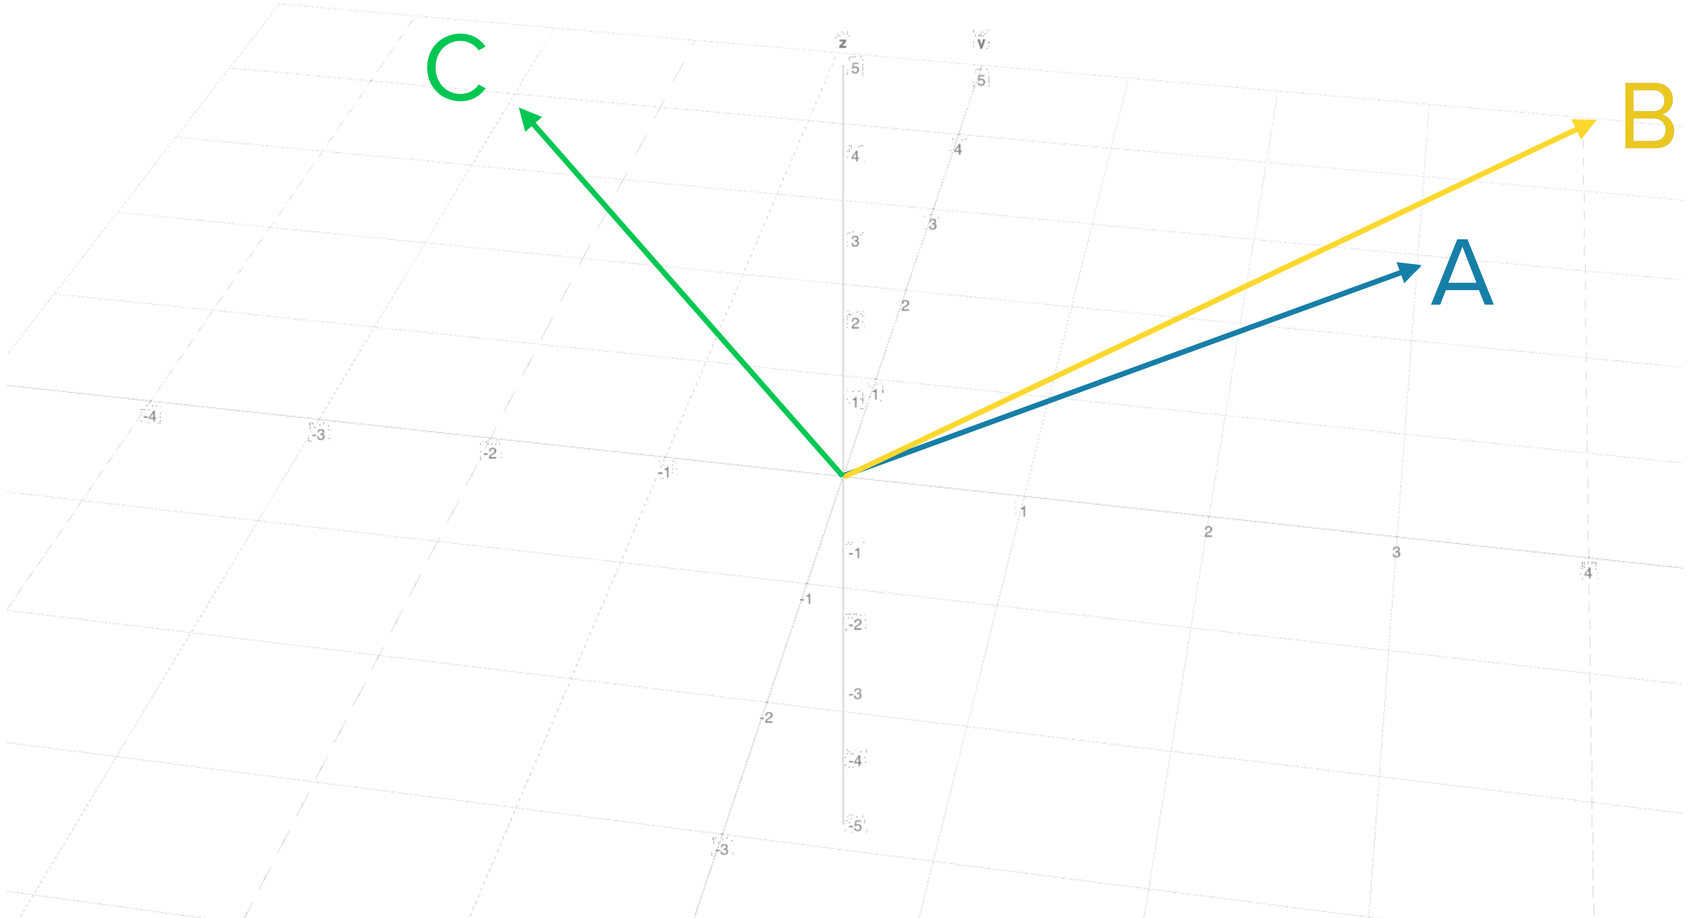

VECTOR EMBEDDING

Calculating distances / nearness

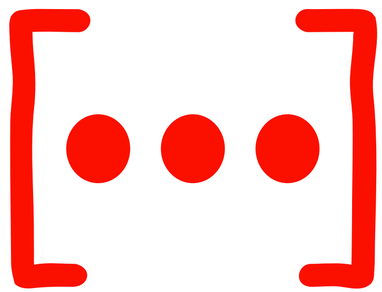

A=[3, 2, 1]

B=[4, 4, 1]

C=[-2, 1, 3]

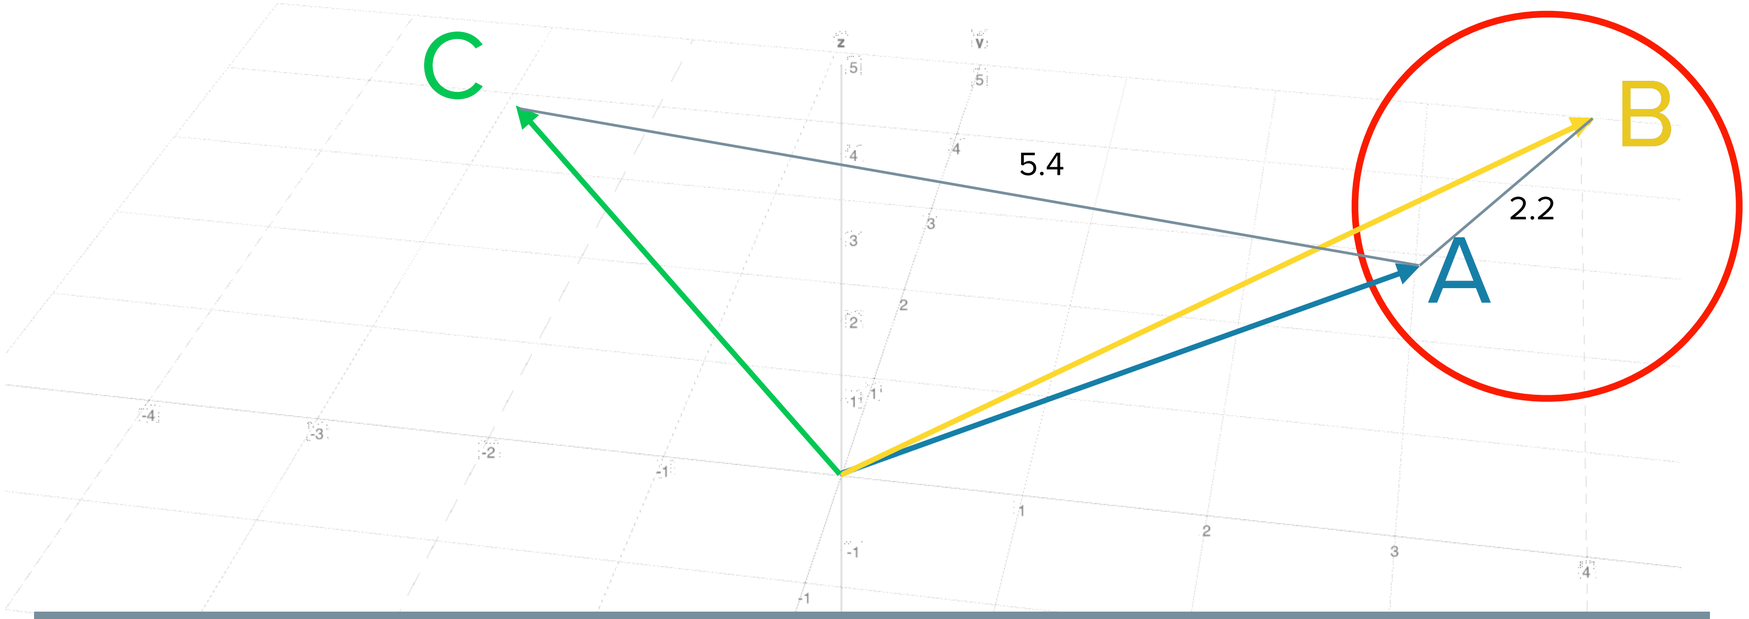

$$\text{euclidean_dist}(A, B) = \sqrt{(4-3)^2 + (4-2)^2 + (1-1)^2} = \sqrt{5} = 2.2$$

$$\text{euclidean_dist}(A, C) = \sqrt{(-2-3)^2 + (1-2)^2 + (3-1)^2} = \sqrt{30} = 5.4$$

VECTOR EMBEDDING

Calculating distances / nearness

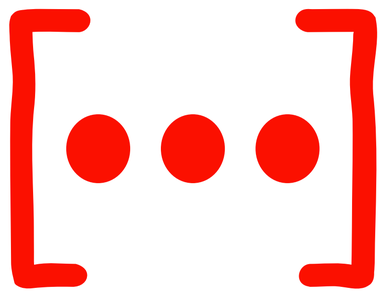

A=[3, 2, 1]

B=[4, 4, 1]

C=[-2, 1, 3]

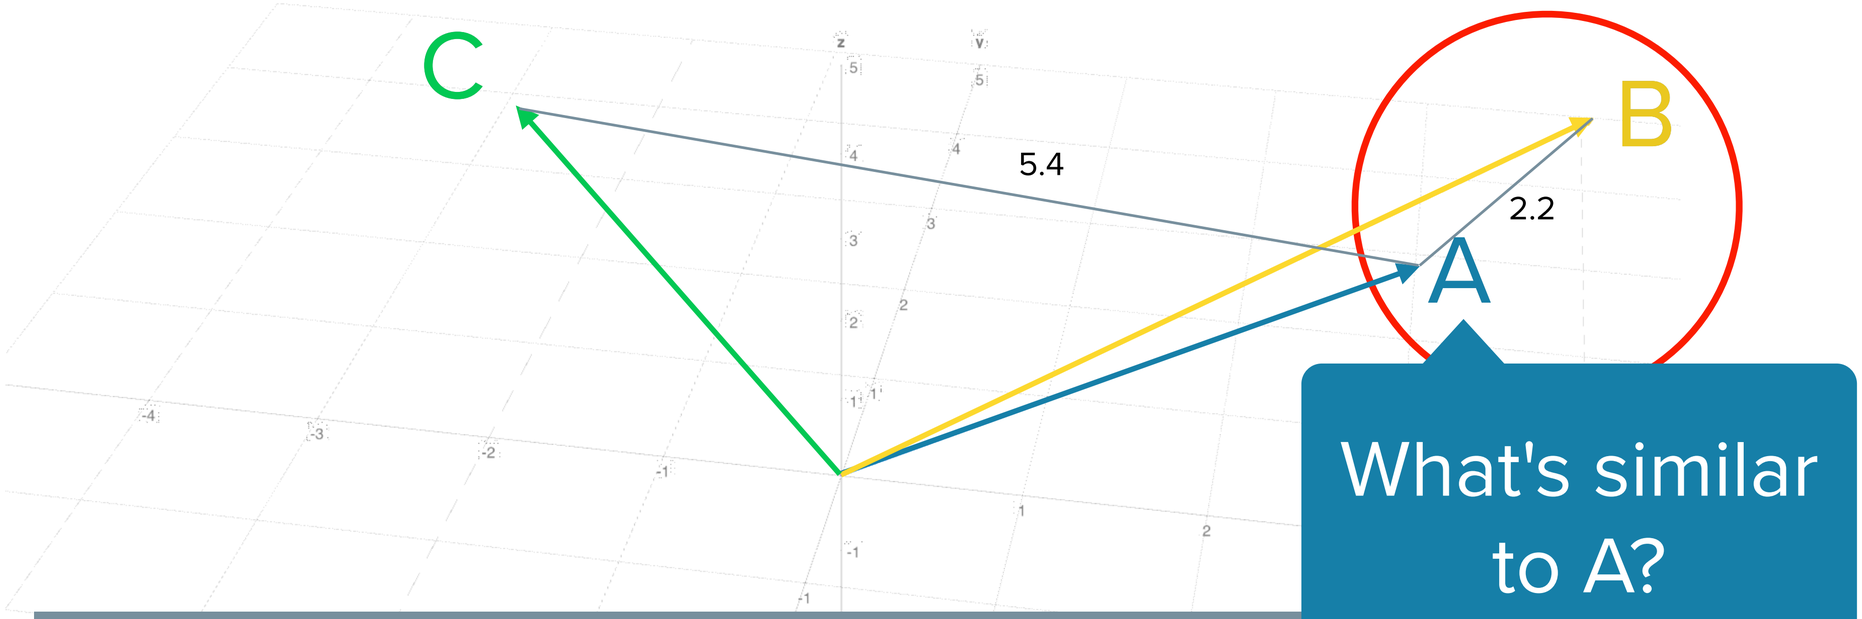

$$\text{euclidean_dist}(A, B) = \sqrt{(4-3)^2 + (4-2)^2 + (1-1)^2} = \sqrt{5} = 2.2$$

$$\text{euclidean_dist}(A, C) = \sqrt{(-2-3)^2 + (1-2)^2 + (3-1)^2} = \sqrt{30} = 5.4$$

VECTOR DATABASES

Main operations

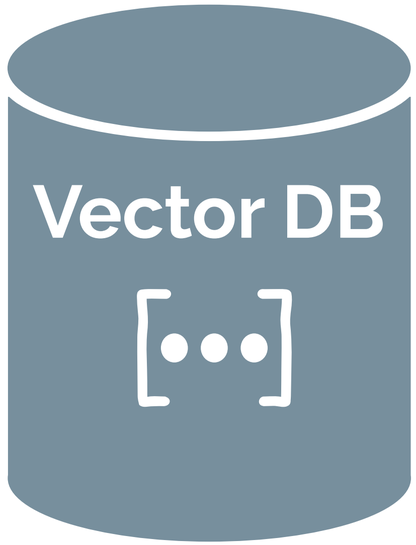

- Nearest neighbor search (kNN)
- Approximate nearest neighbor search (ANN)
- k-Means clustering

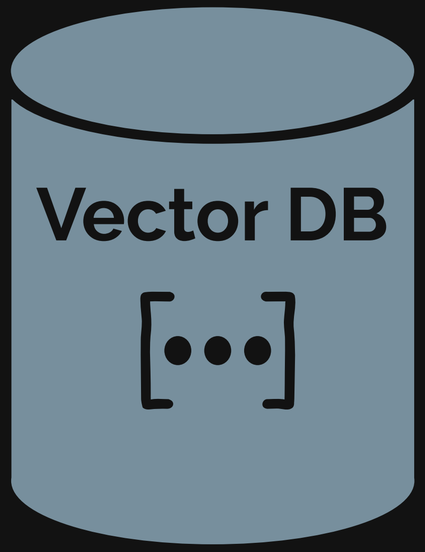

GROWING LANDSCAPE

LOTS OF DEDICATED DATABASES, BUT ALSO SUPPORT GETTING ADDED TO LEGACY DATABASES

Weaviate
drant
OpenSearch
ORACLE
redis
Chroma
elastic
vespa
pgvector
Pinecone
KX
SQLite-vss
LanceDB
milvus
MongoDB.
neo4j

The 50 Biggest SaaS Companies

On U.S. stoc

Last updated: Jul

	Company (Ticker)
1.	Salesforce, Inc. (CRM)
2.	Adobe Inc. (ADBE)
3.	Intuit Inc. (INTU)
4.	ServiceNow, Inc. (NOW)
5.	CrowdStrike Holdings, Inc. (CRWD)
6.	Shopify Inc. (SHOP)
7.	Workday, Inc. (WDAY)
8.	Palantir Technologies Inc. (PLTR)
9.	Autodesk, Inc. (ADSK)
10.	Atlassian Corporation (TEAM)
11.	Snowflake Inc. (SNOW)
12.	Datadog, Inc. (DDOG)
13.	Block, Inc. (SQ)
14.	HubSpot, Inc. (HUBS)
15.	Veeva Systems Inc. (VEEV)
16.	Zscaler, Inc. (ZS)
17.	Cloudflare, Inc. (NET)
18.	Samsara Inc. (IOT)

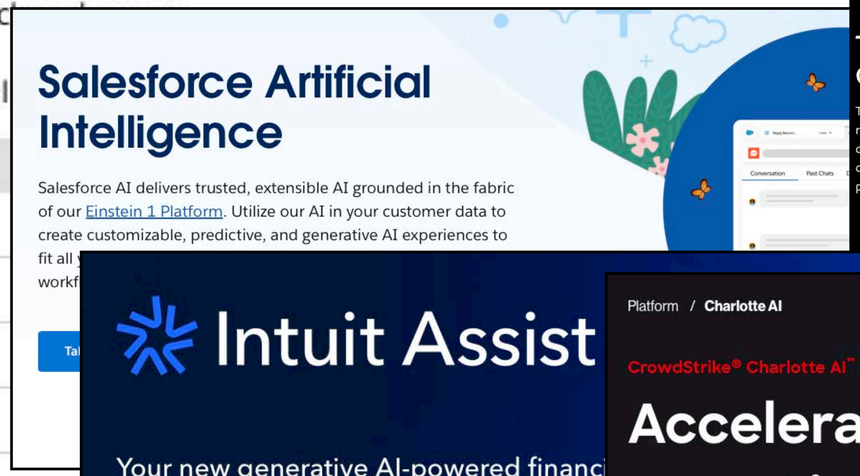

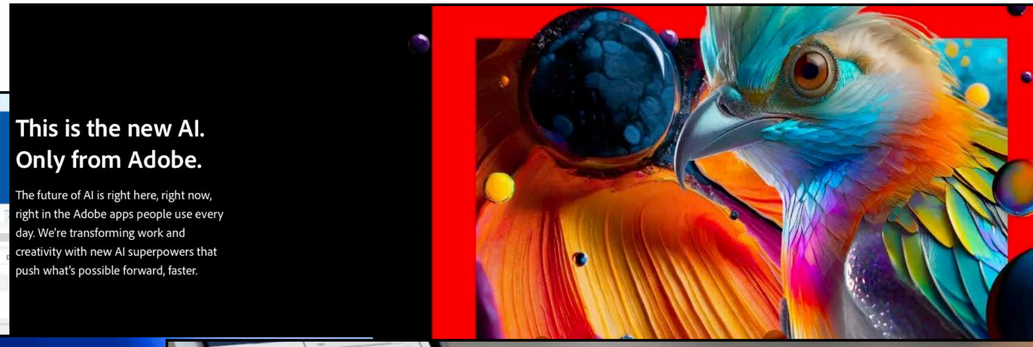

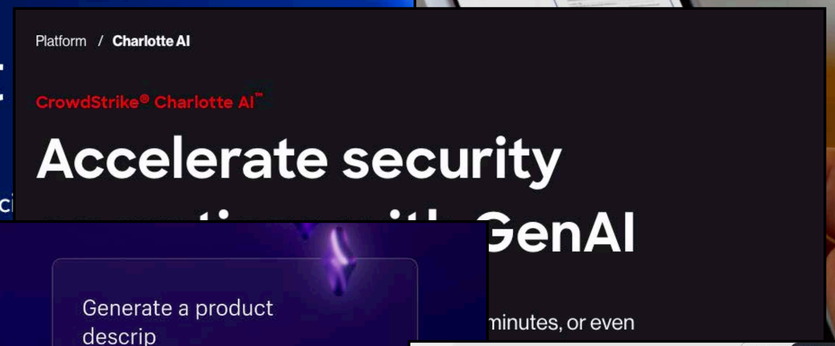

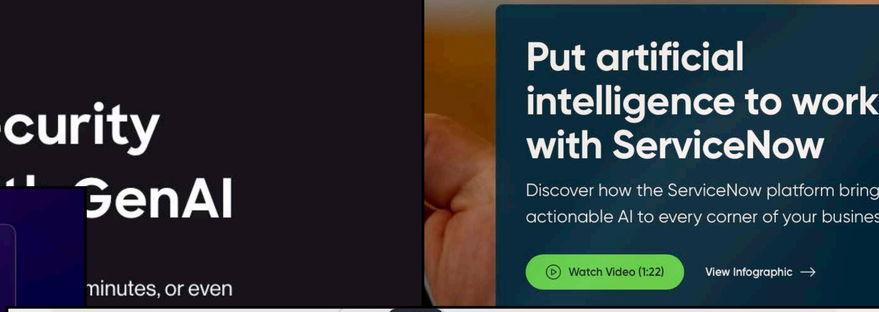

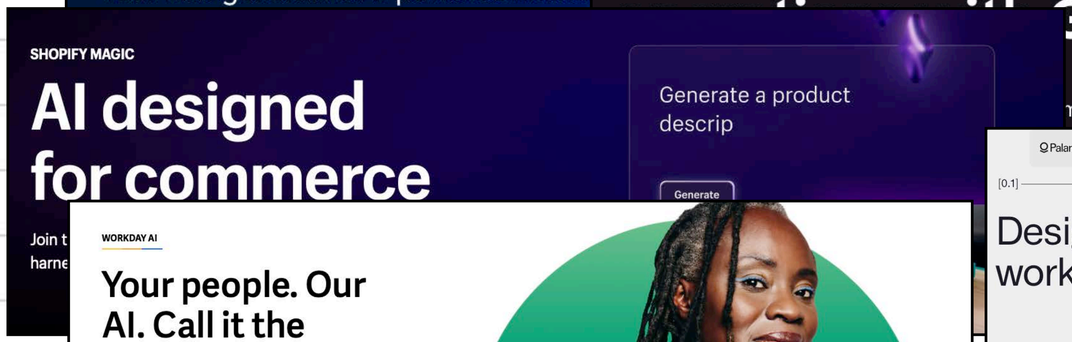

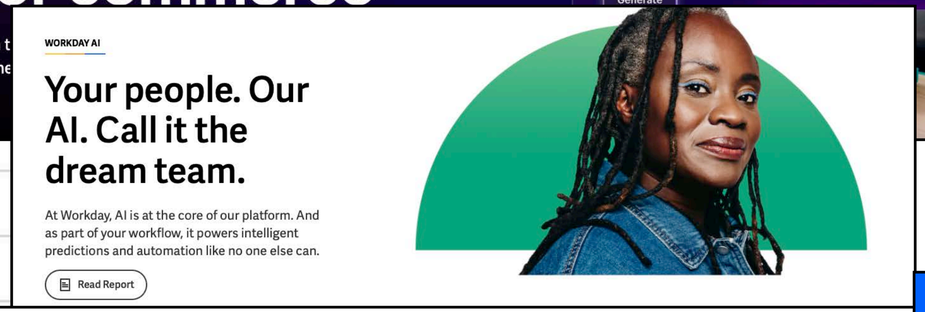

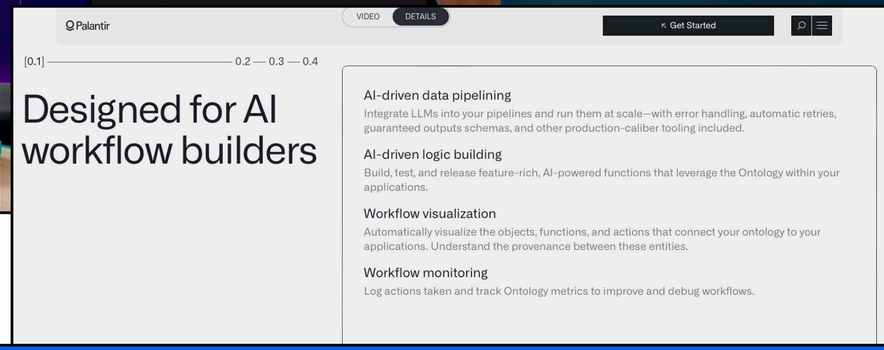





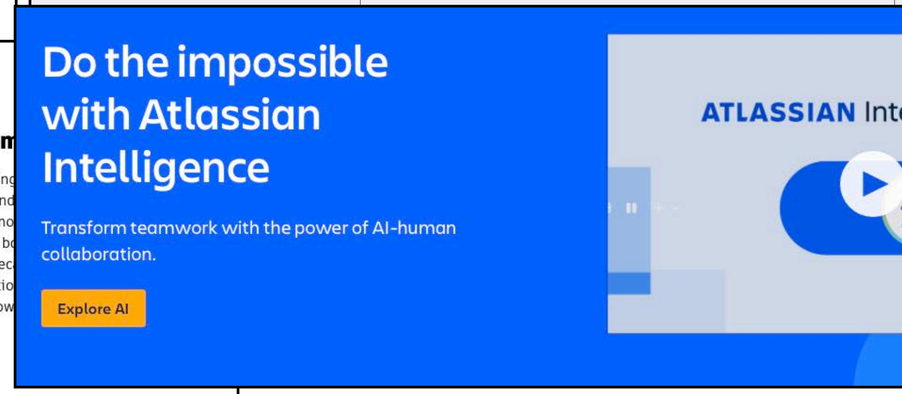

Top Cybersecurity Companies

The following table compares our top 20 providers:

Vendor	# of eSecurity Planet Top Product Lists	Overall Gartner Peer Insights Score	Ove
Palo Alto: Best Protection Against Network, Endpoint, & Remote Asset Attack	14	4.6	
Fortinet: Best for Network Security Perimeter Protection	15	4.7	
Cisco: Best for Integrated Network Security	16	4.4	
CrowdStrike: Best for Endpoint Security and Services	7	4.7	
Zscaler: Best for Cloud Security	5	4.5	3.9
IBM: Best for Advanced Encryption	14	4.4	4.1
Trend Micro: Best for Small Businesses	13	4.6	
Okta: Best for Access Management	4	4.5	
OneTrust: Best for Privacy and Compliance	2	4.3	
Rapid7: Best for Integrated Vulnerability & Threat Management	9	4.3	
Proofpoint: Best for End User Data Security	7	4.4	
Tenable: Best for Vulnerability Management	8	4.5	
KnowBe4: Best for Security Awareness Training	2	4.6	
Darktrace: Best for AI-Powered Security	1	4.7	
Check Point: Best for	11	4.4	3.7

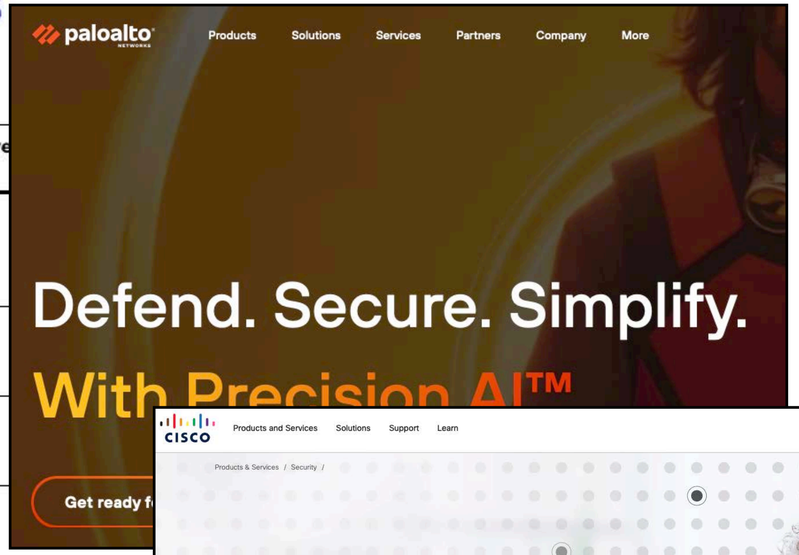

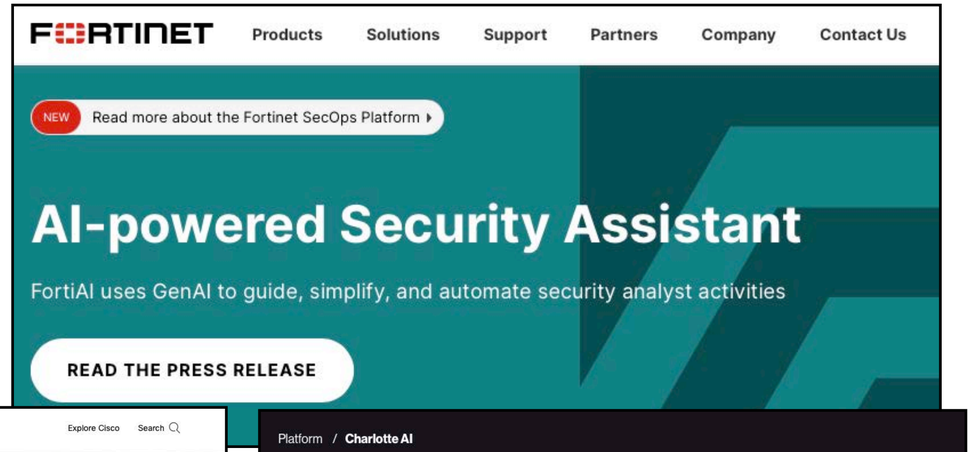

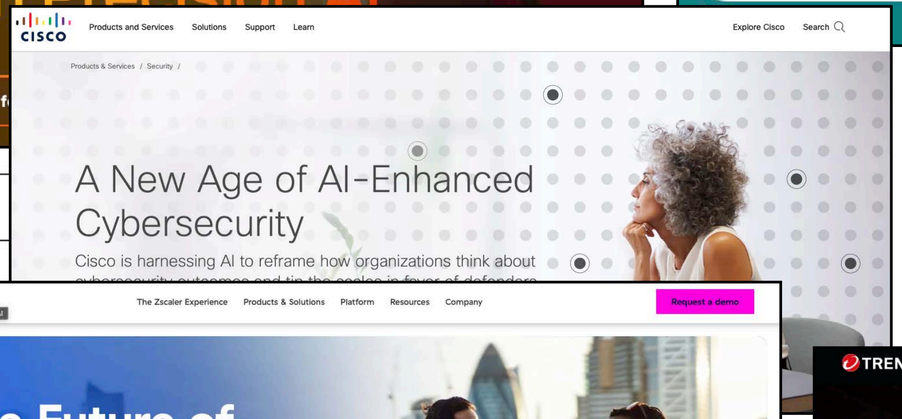

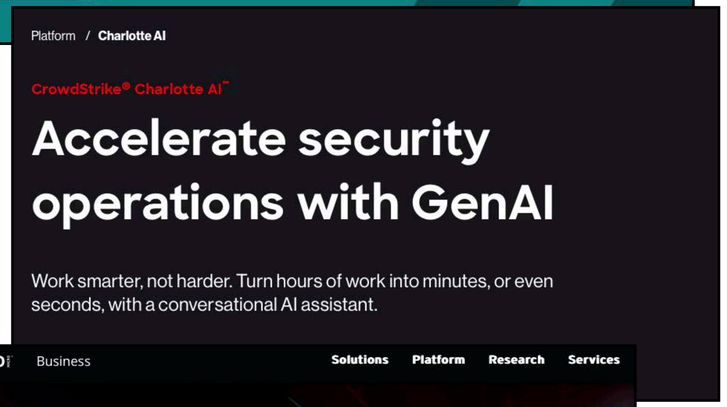

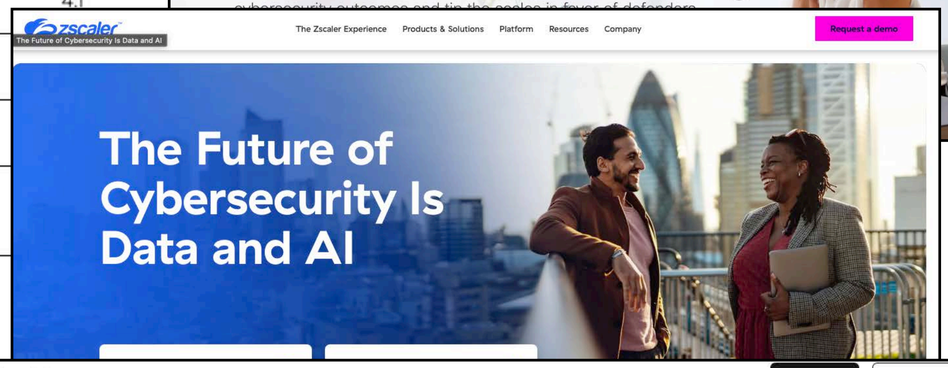

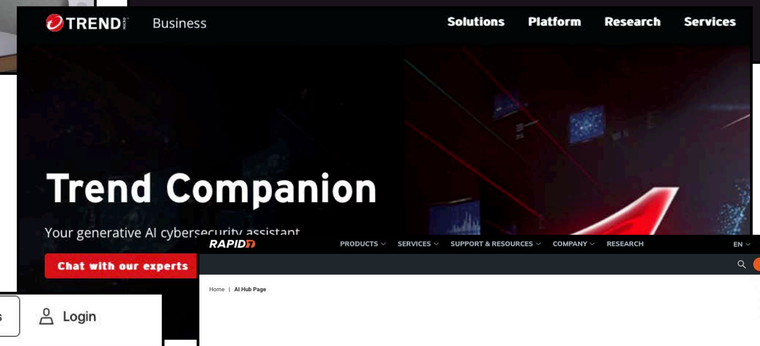

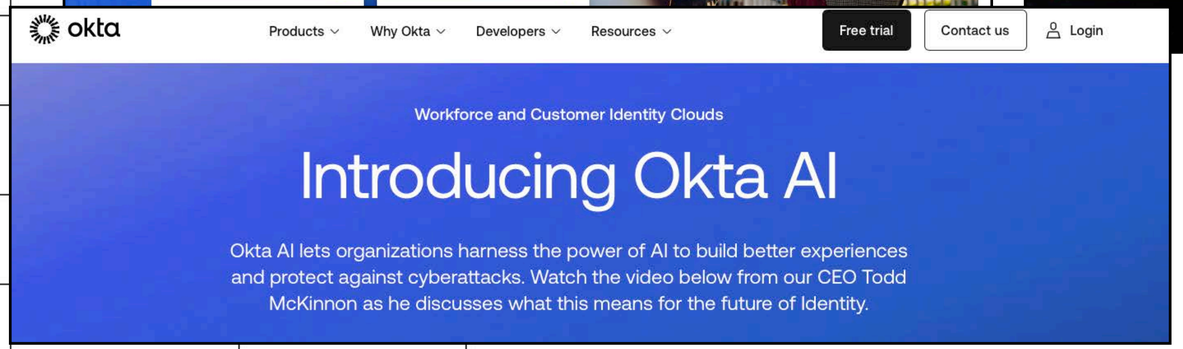

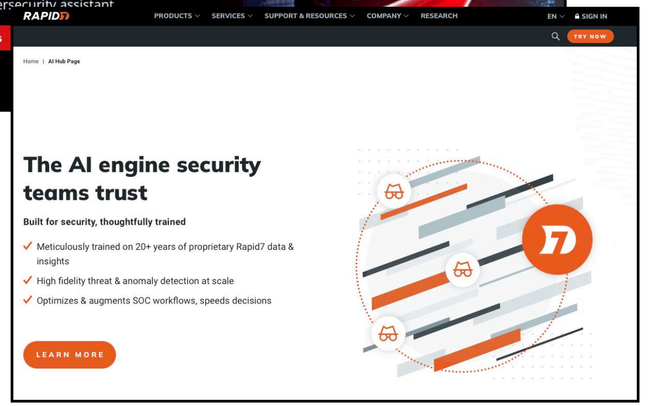



More data == more useful AI

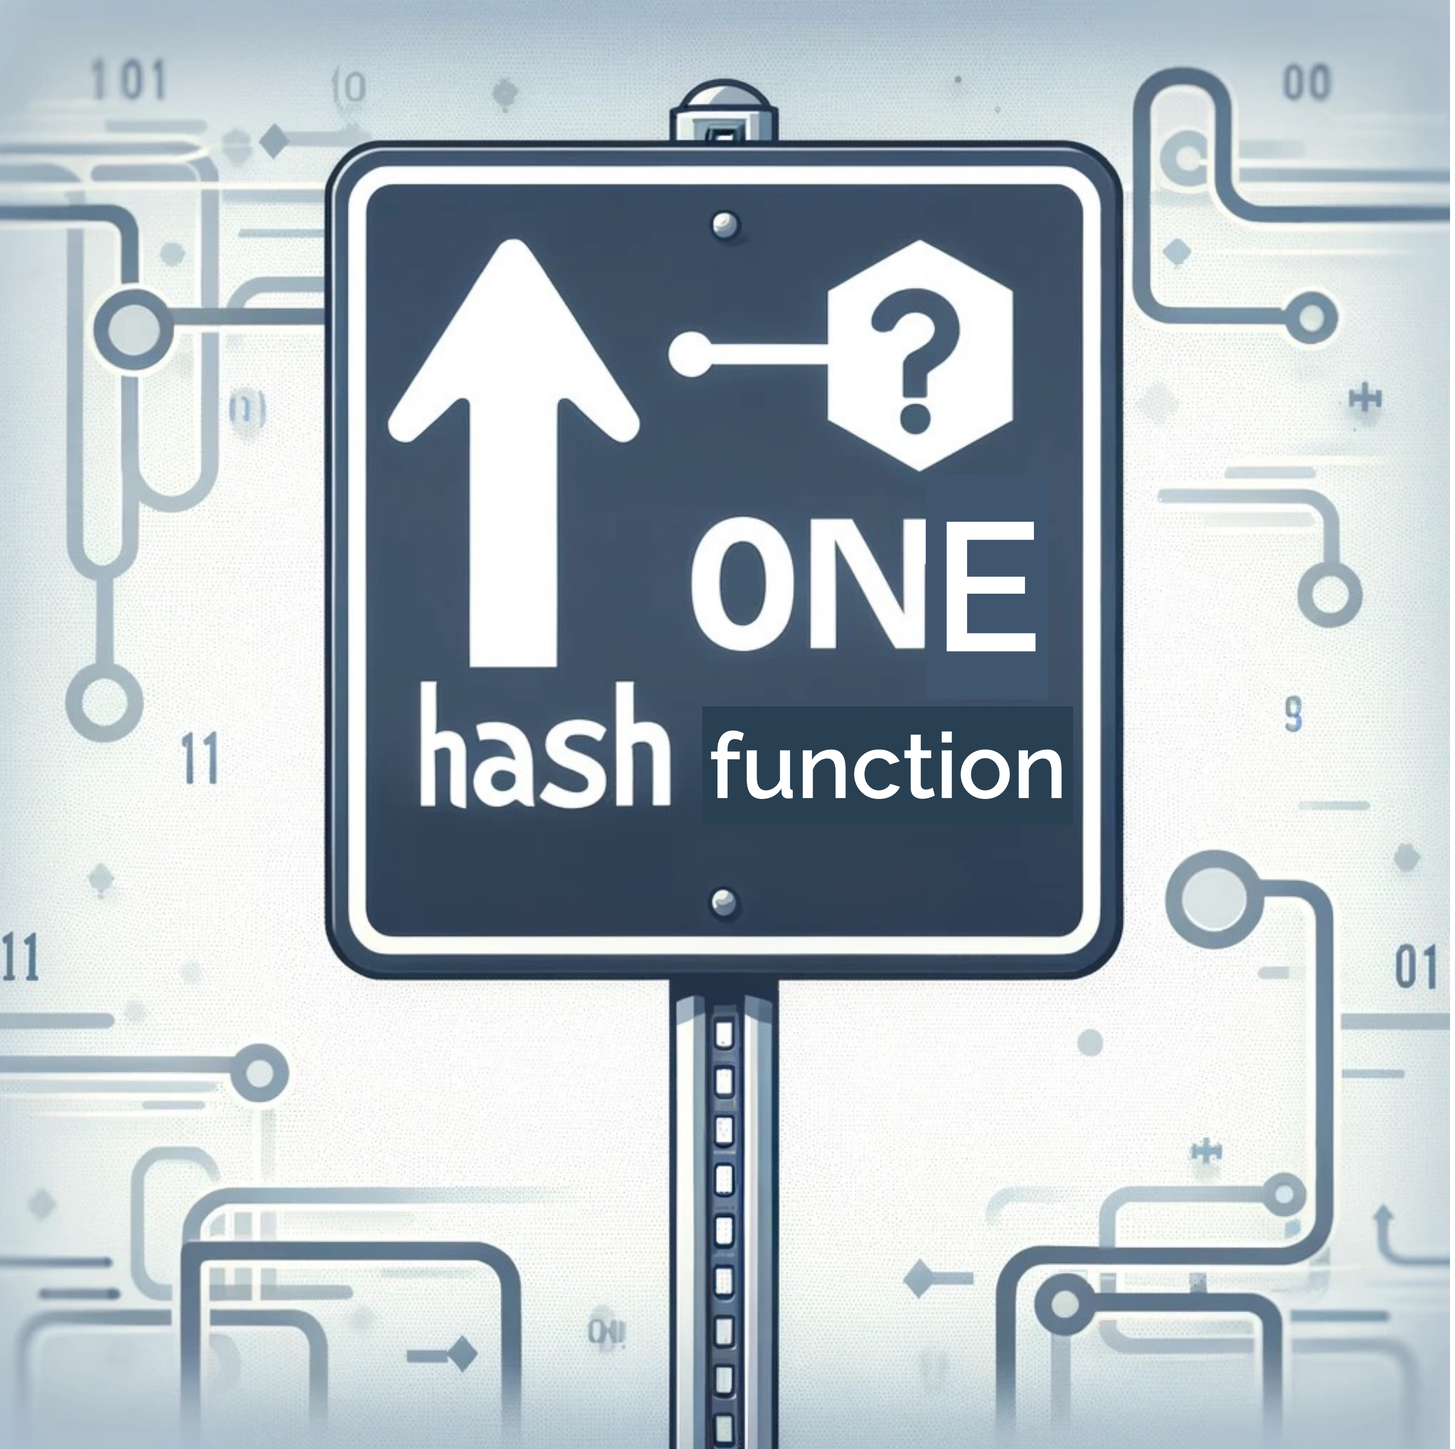

Vector embeddings:
Like a one-way hash?

(Even if so, dictionary attacks, rainbow tables, etc... not invulnerable.)

EMBEDDING VECTORS

CONVERT FROM TEXT TO VECTOR OF NUMBERS

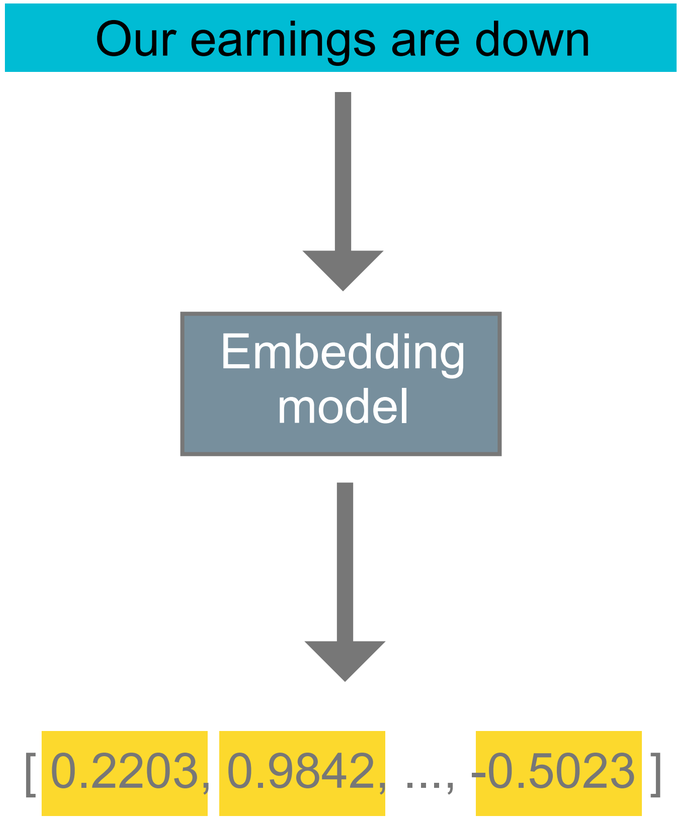

EMBEDDING VECTORS

Convert from text to vector of numbers

Our earnings are down

↓

Embedding model

↓

[0.2203, 0.9842, ..., -0.5023]

↓



```
❯ cat meow_attention.gif
GIF89az??21/31/32/IC5ID5ID6JE7KF8LG:MH;NI;NI<NJ<OJ=PK>PL@QMASNAVRFWSG
[O.]P+^Q+_Q+^Q,_R+_R,`R+aS*aS+`R,bS+bT*bT,eV*?r&?r&?s'?s(?s&?s'?s&?t'
?t&?t'?t&?u&?u'?u&??????????????????????????????????????????????????DãCä
DåE¤FåGg>ÎAÎhDÎhFŦFFŖCğCŖD2E2G2H??????????????????£?¦?è?éĘFɪI????????à?ã?ũь
I??Ü?Ý?g?ţ?Ú¨#?#?"?"?#?%?"?"?#?"?"?#??*??+??*??+??+??,??-??!?
```

EMBEDDING **INVERSIONS**

TURN VECTORS BACK INTO APPROXIMATIONS OF SOURCE DATA

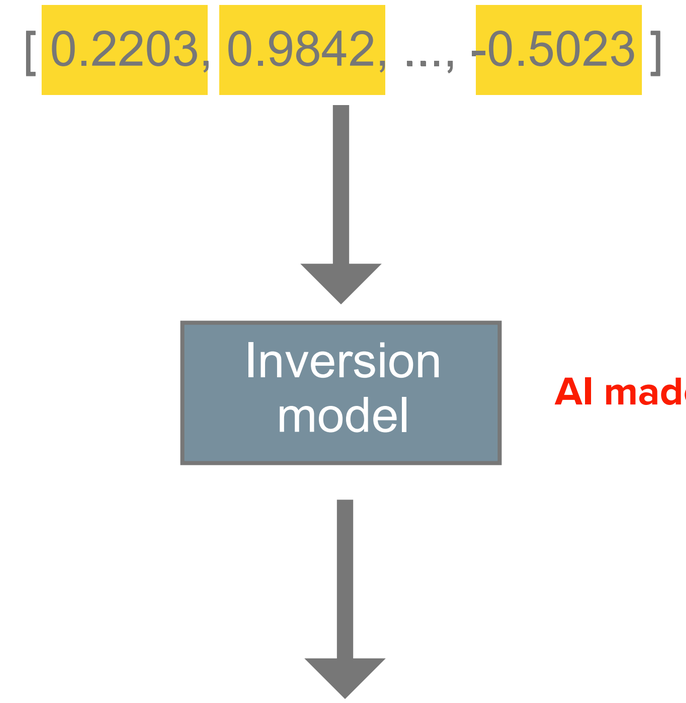

AI made it, and AI can unmake it

(who will AI come for next!?)

vs. "our earnings are down"

EMBEDDING **INVERSION**

EXTREMELY ACCURATE RECOVERIES

Google Scholar

vector embedding inversion attack

SIGN IN

Articles

About 25,300 results (**0.09** sec)

My profile

My library

Any time
Since 2024
Since 2023
Since 2020
Custom range...

Sort by relevance
Sort by date

Any type
Review articles

include patents
include citations

Create alert

Information leakage in **embedding** models

[PDF] arxiv.org

C Song, A Raghunathan - Proceedings of the 2020 ACM SIGSAC …, 2020 - dl.acm.org

… these **attacks**. We show that **embedding inversion attacks** and attribute inference **attacks** … attribute leakage from the **embedding vector**, we compare the **attack** performance between …

Save Cite Cited by 194 Related articles All 5 versions

Invernet: An **inversion attack** framework to infer fine-tuning datasets through word **embeddings**

[PDF] aclanthology.org

I Hayet, Z Yao, B Luo - Findings of the Association for …, 2022 - aclanthology.org

… pretrained **embed**ding gets access to the fine-tuned **embedding** and … In this study, we propose a novel **embedding inversion** … b, we obtain n binary context **vectors** of the target word from …

Save Cite Cited by 4 Related articles All 8 versions

Variational model **inversion attacks**

[PDF] neurips.cc

KC Wang, Y Fu, K Li, A Khisti… - Advances in Neural …, 2021 - proceedings.neurips.cc

… 2018], where we treated the latent **vector** as 1x1 images. We removed the original squeezing layers that reduced the image size to account for this change. We also experimented with …

Save Cite Cited by 89 Related articles All 7 versions

EMBEDDING **INVERSION**

Extremely accurate recoveries

Paper: Text Embeddings Reveal (Almost) As Much As Text

"...is able to **recover 92%** of 32-token text inputs exactly... and also show that our model can **recover important personal information** (full names)"

The remaining 8% is still equivalent to the input.

Source: https://arxiv.org/abs/2310.06816

EMBEDDING **INVERSION**

LEVEL OF DIFFICULTY: LOW



- The attacks sound complicated -- and in some cases they are.
- **But anyone can run them because they're open source.**

EMBEDDING **INVERSION**

LEVEL OF DIFFICULTY: LOW

The Hidden Dangers of Face Embeddings: Unmasking the Privacy Risks

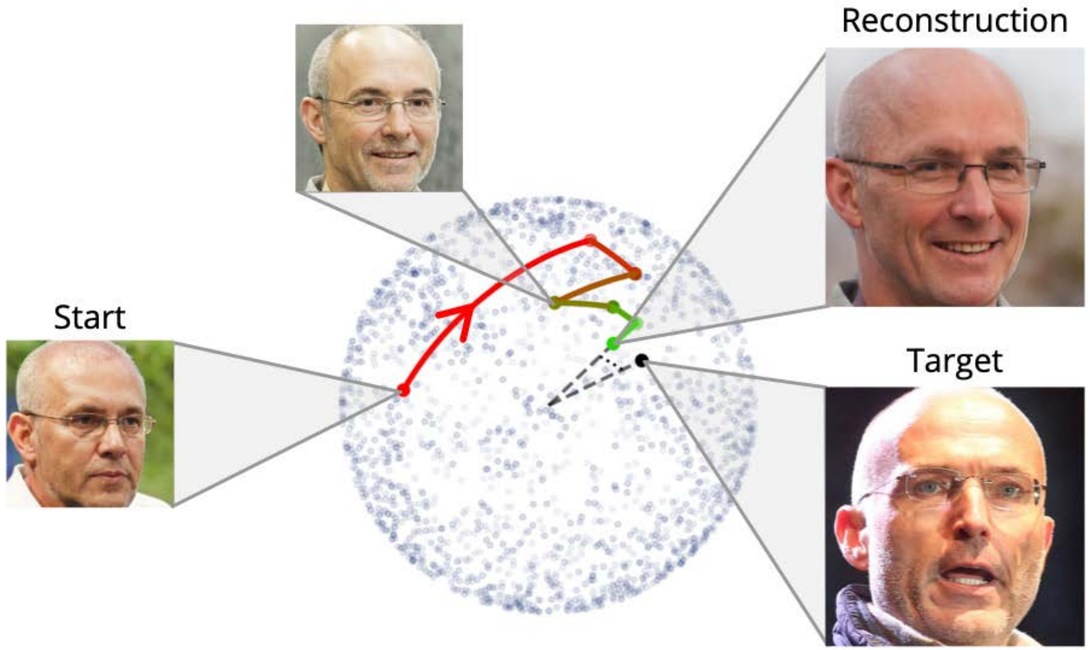

Source: https://ironcorelabs.com/blog/

EMBEDDING **INVERSION**

LEVEL OF DIFFICULTY: LOW

The Hidden Dangers of Face Embeddings: Unmasking the Privacy Risks

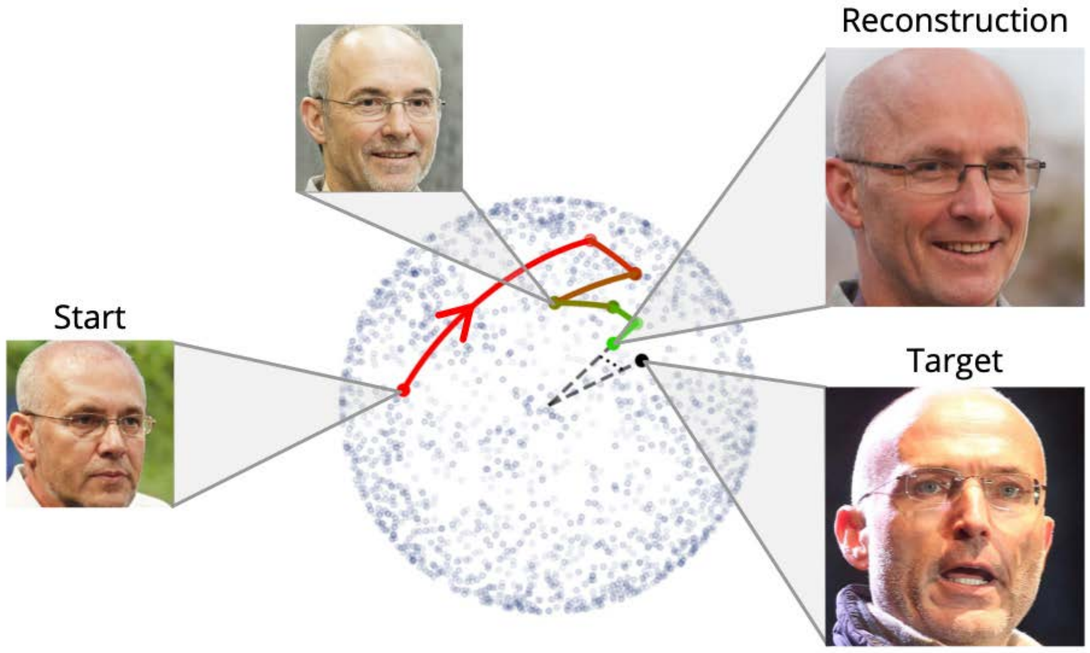

There and Back Again: An Embedding Attack Journey

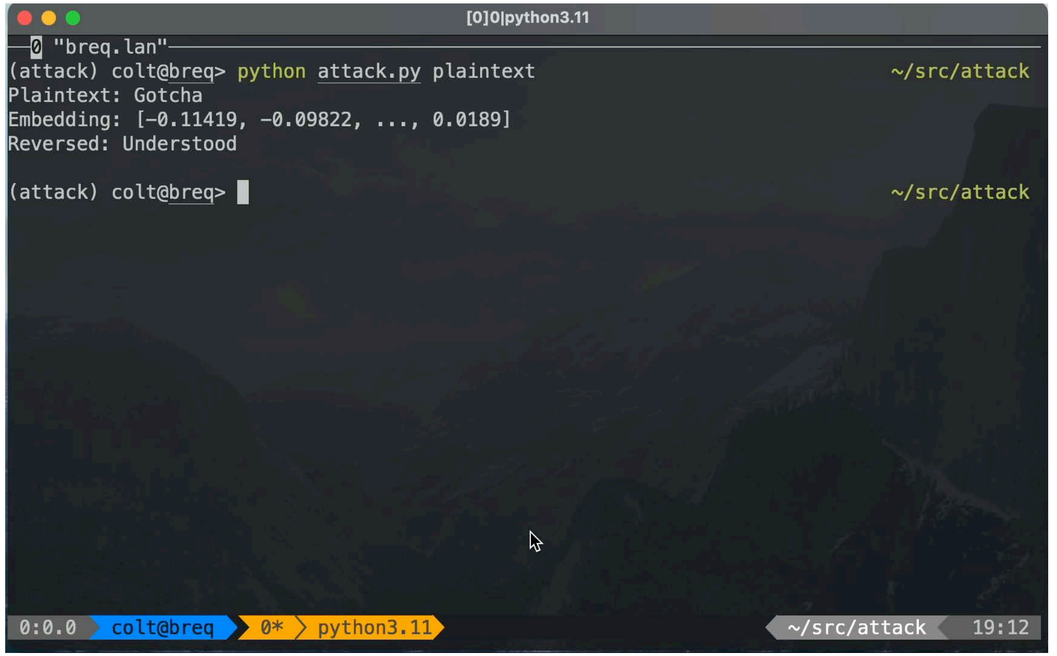

Source: https://ironcorelabs.com/blog/

VECTOR DB **SECURITY**

So immature...

	Level for score of 13 out of possible 70
Security Maturity Level:	Weak
	0-15: weak, 15-30: reasonable, 30+: exceptional

- Default: no auth required
- Backdoor ports
- False security claims like E2EE
- Weak / no RBAC
- No production access restriction for engineers


Vector DB
[•••]

STEALING **DATA**

BUT YOU HAVE TO FIND IT FIRST

BEFORE

Finding information and distilling it in a traditional database takes time and skill to reverse engineer, hunt, and to extract the valuable info.



STEALING **DATA**

BUT YOU HAVE TO FIND IT FIRST



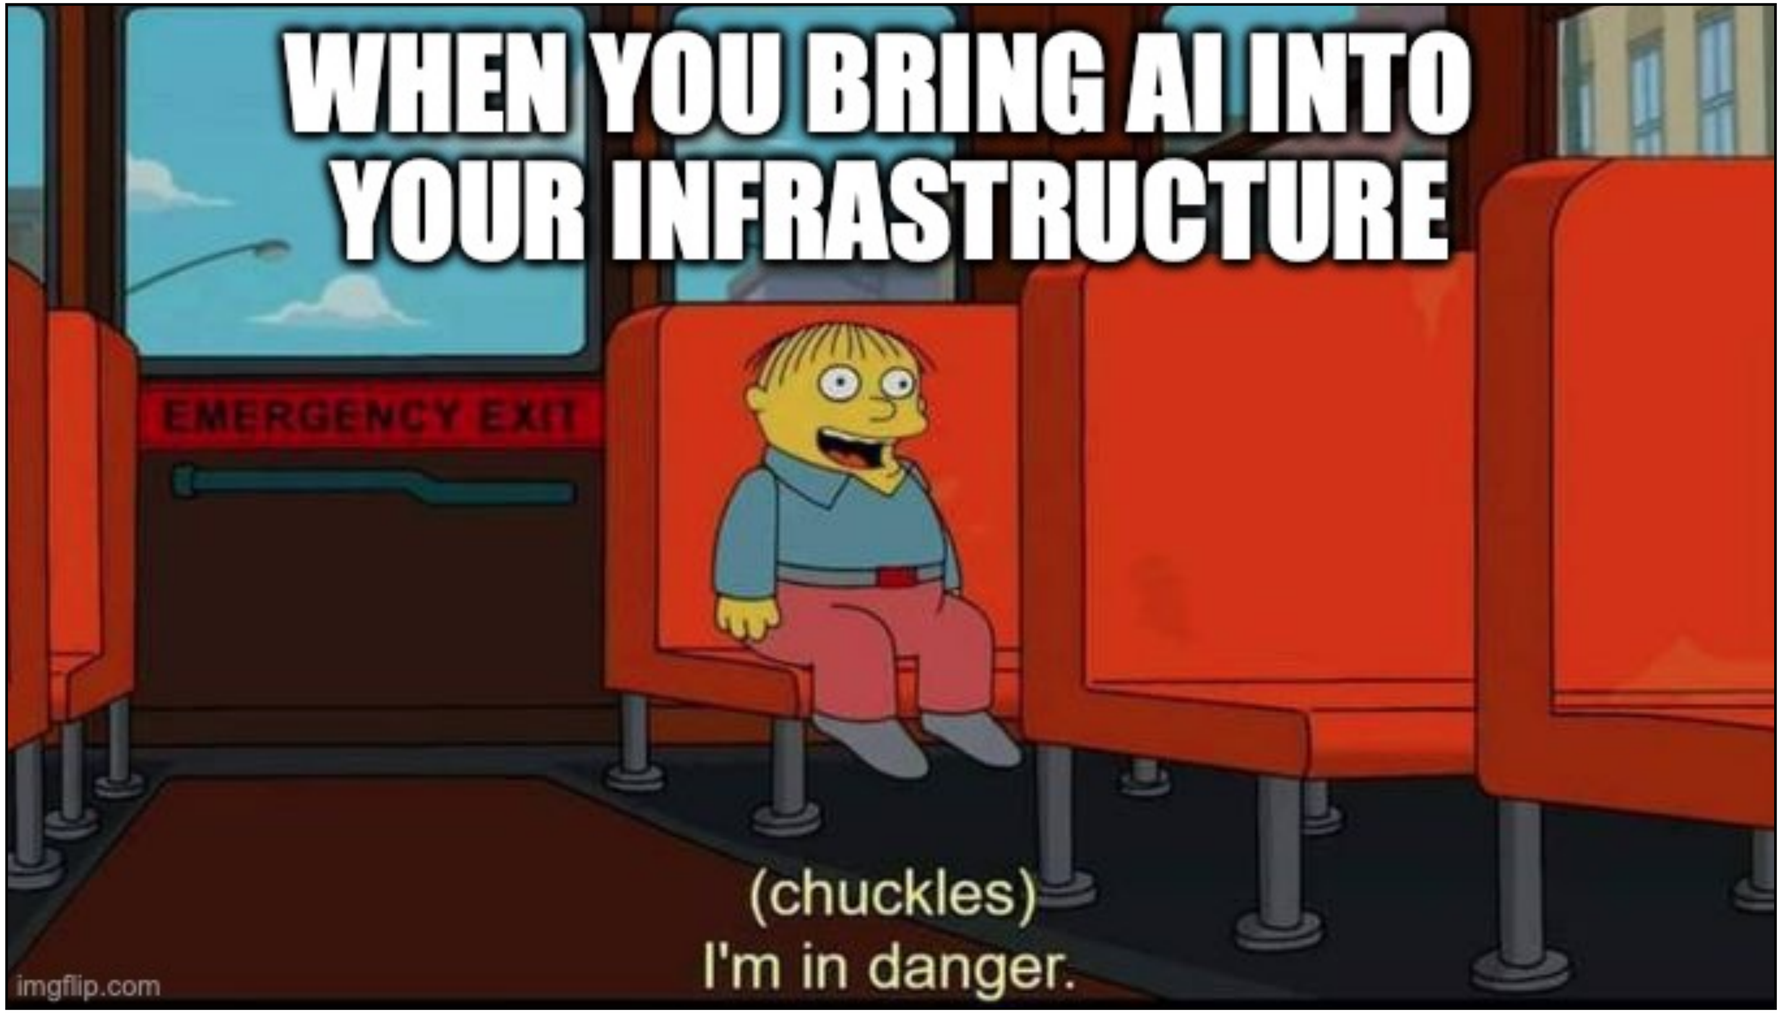
WHEN YOU BRING AI INTO
YOUR INFRASTRUCTURE
EMERGENCY EXIT
(chuckles)
I'm in danger.
imgflip.com

RISKS SUMMARY

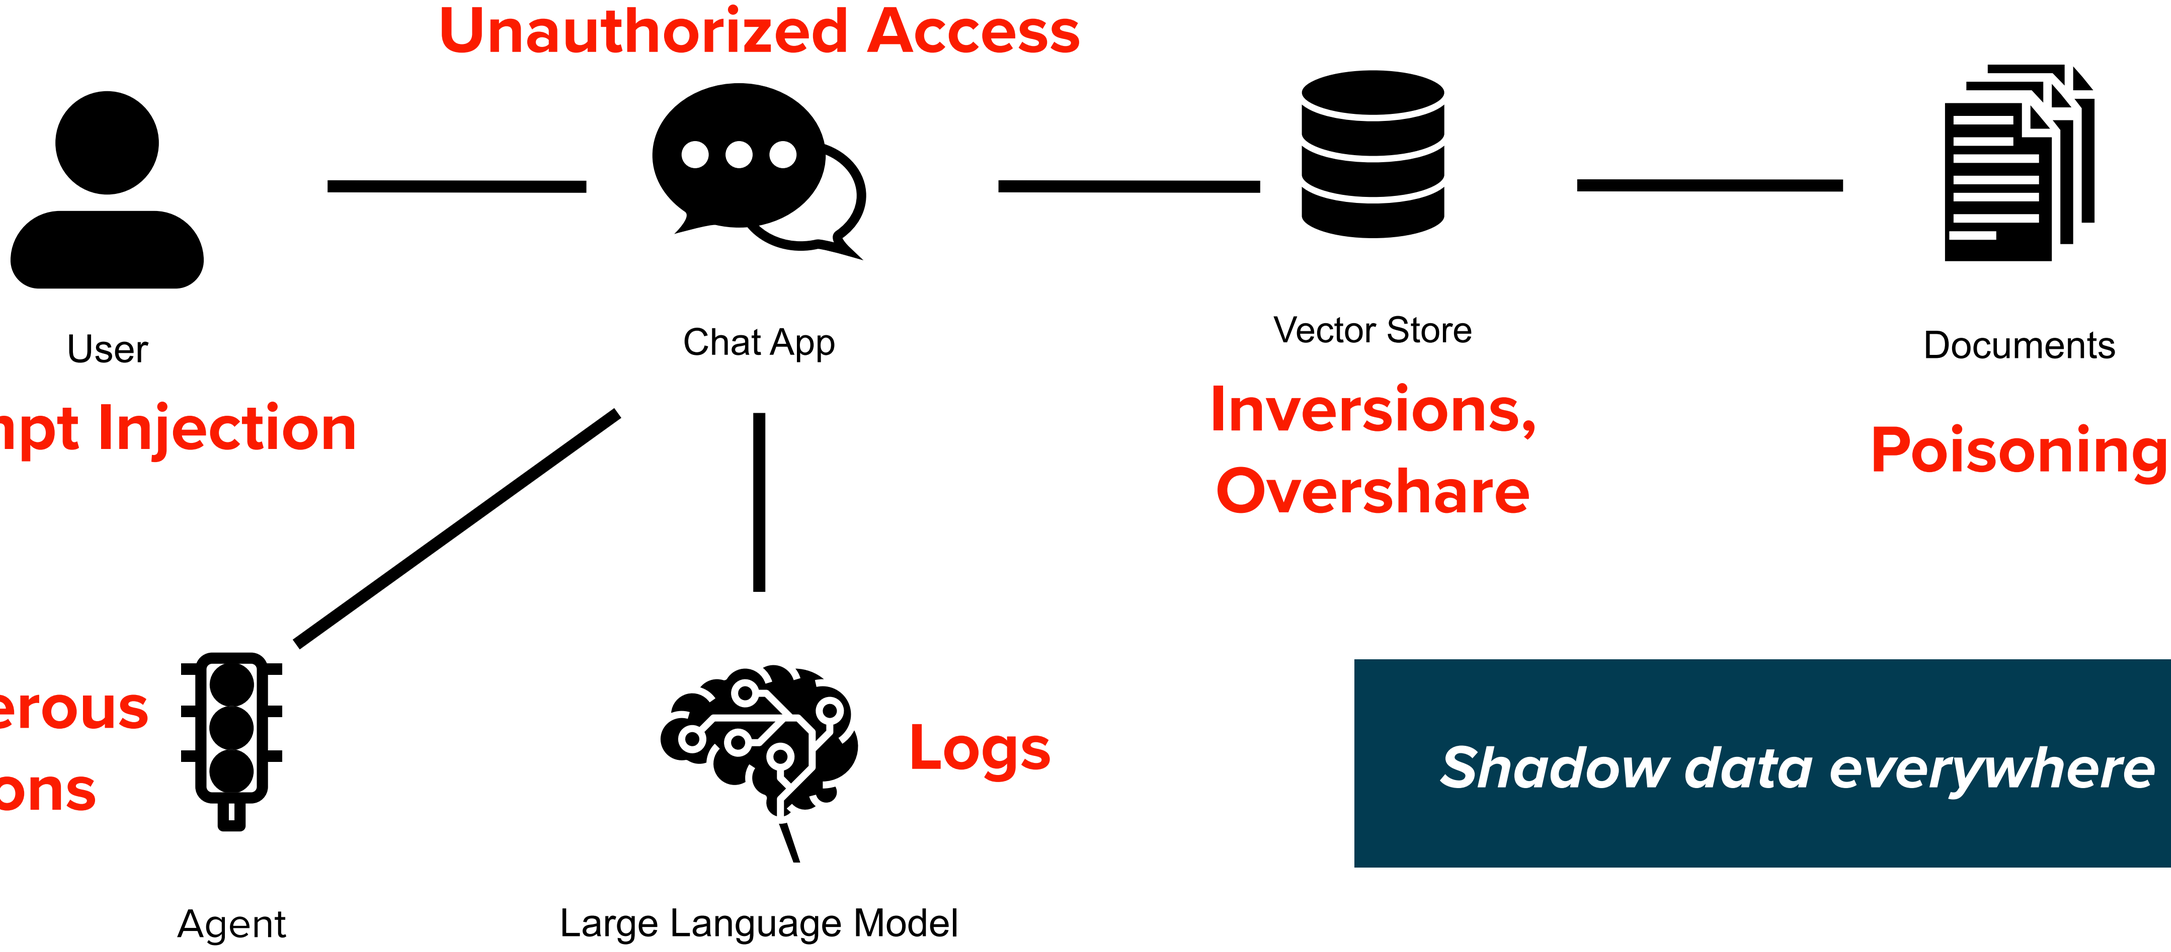

Shadow data everywhere

FIXING VECTORS

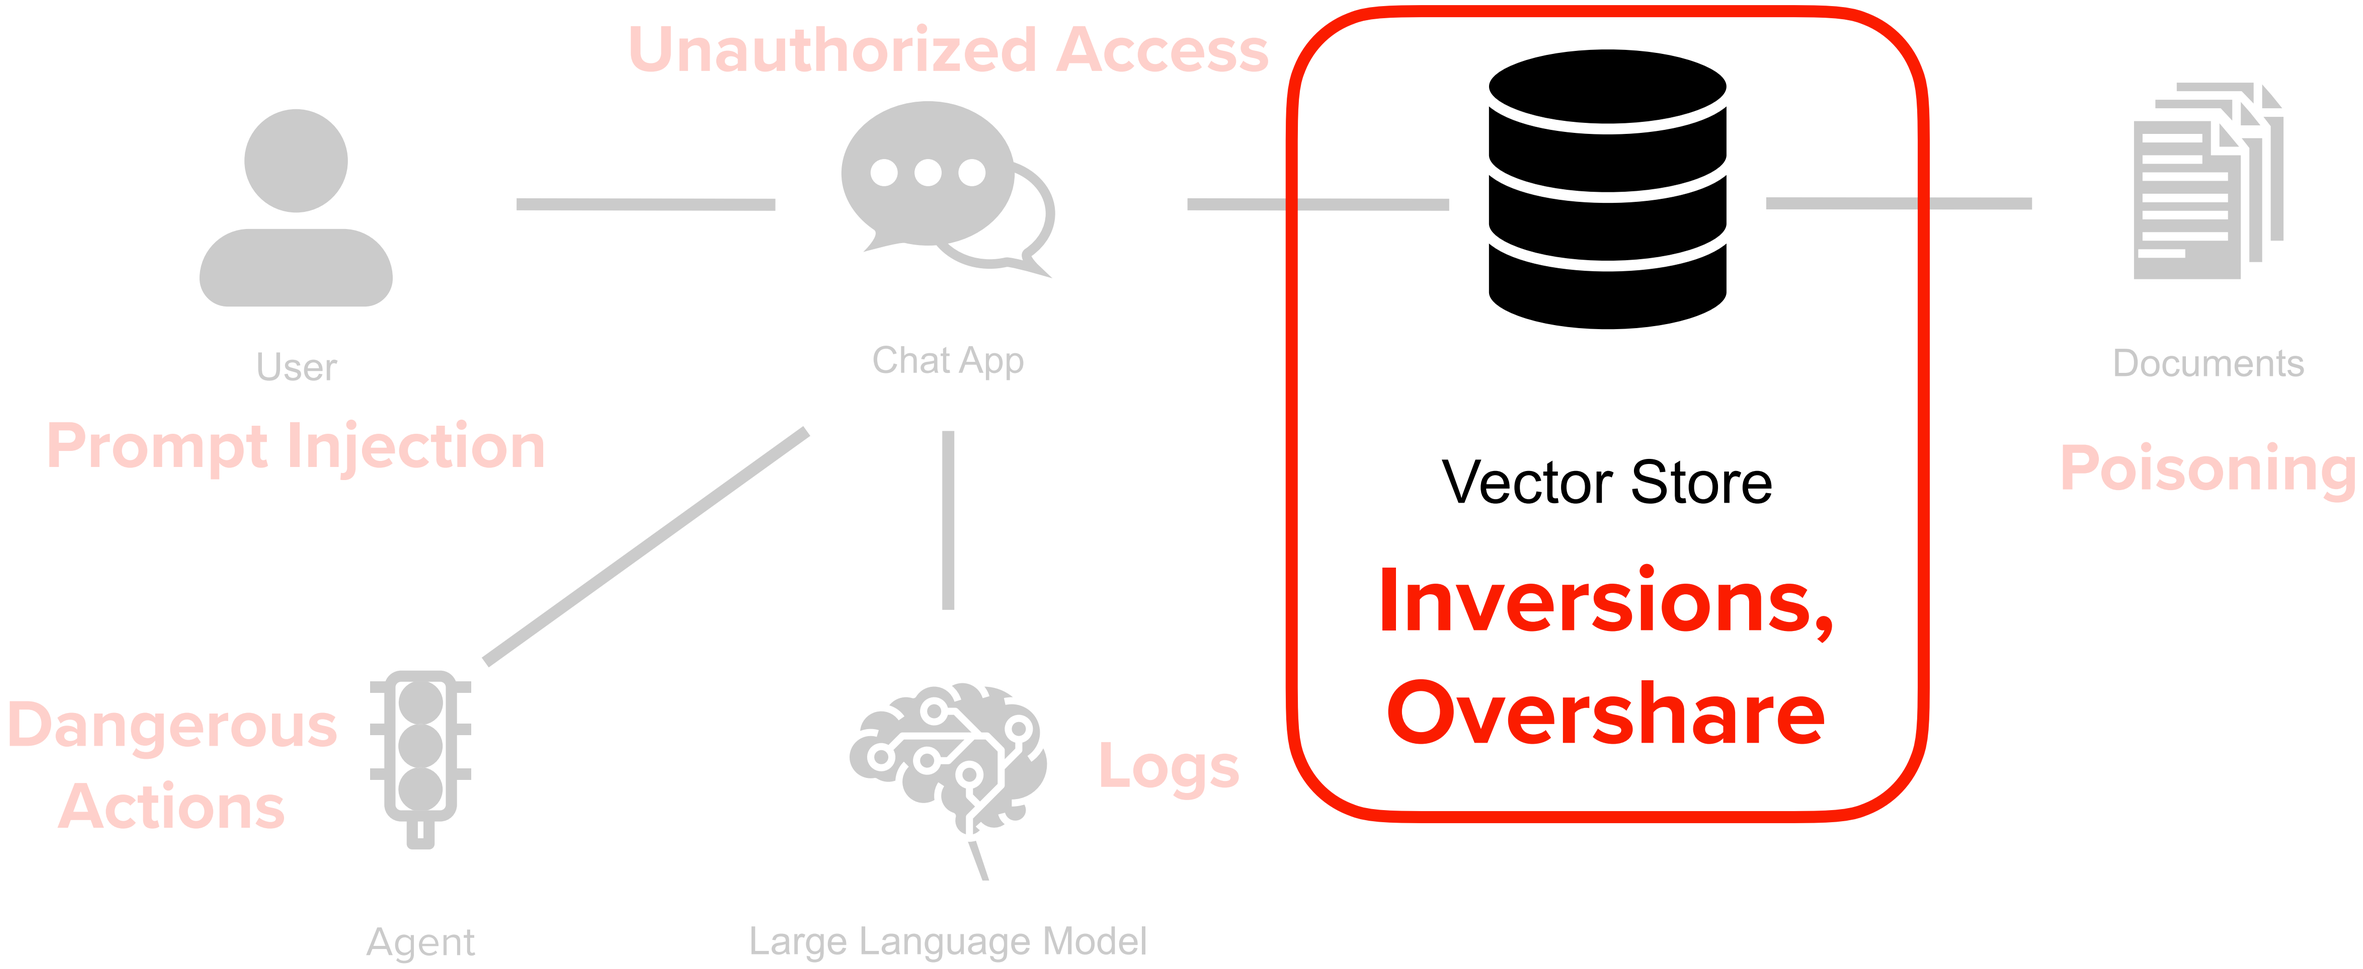

VECTOR ENCRYPTION



Vector Store

Approaches

- Fully Homomorphic Encryption (FHE)
- Partially Homomorphic Encryption

VECTOR ENCRYPTION

Fully Homomorphic Encryption of Vectors



Vector Store

Efficient Homomorphic Encryption on Integer Vectors and Its Applications

Hongchao Zhou and Gregory Wornell
Dept. Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: {hongchao,gww}@mit.edu

***Abstract*—Homomorphic encryption, aimed at enabling computation in the encrypted domain, is becoming important to a wide and growing range of applications, from cloud computing to distributed sensing. In recent years, a number of approaches to fully**

Efficient Integer Vector Homomorphic Encryption

Angel Yu, Wai Lok Lai, James Payor

Massachusetts Institute of Technology
77 Massachusetts Avenue, Cambridge, MA 02139, USA

Efficient and Secure *k*NN Classification over Encrypted Data using Vector Homomorphic Encryption

Haomiao Yang*, Weichao He†, Jie Li‡, Hongwei Li§
†*School of Computer Science and Engineering & Center for Cyber Security*
University of Electronic Science and Technology of China
Chengdu, China
Email: {haomyang, hongweili}@uestc.edu.cn
‡*Xi'an Institute of Space Radio Technology*
Xi'an, China
Email: {wechaohe, xnkjlj}@163.com

***Abstract*—The *k*-nearest neighbor (*k*NN) classification has been widely adopted in data mining applications. In the age of big data, *k*NN classification process has to be outsourced to the cloud. However, as data may contain sensitive information, outsourcing data services directly to public clouds inevitably raises privacy concerns. To ensure the privacy of data, it is a well-known method to encrypt them prior to uploading to the cloud, which also brings great challenges to effective *k*NN classification. Homomorphic encryption (HE) allows operations on**

to the data records in the training set. Then the query record is classified by a majority vote of its nearest neighbors. Therefore, it is one of the most fundamental functionalities for big data analytics, and can widely be used in many data mining applications. Especially in the big data age, data centers may host the *k*NN classification services. As a feasible solution, outsourcing *k*NN classification process to the commercial public clouds is an economical option for small enterprises.

2021 IEEE 24th International Conference on Computational Science and Engineering (CSE)

A K-nearest neighbor classifier based on homomorphic encryption scheme

1st Zhenzhou Guo
School of Computer Science
Shenyang Aerospace University
Shenyang, China
guozhenzhou@sau.edu.cn

2nd Weifeng Jin
School of Computer Science
Shenyang Aerospace University
Shenyang, China
jwf_1996@163.com

3th Xintong Li
School of Computer Science
Shenyang Aerospace University
Shenyang, China
1725593953@qq.com

4th Han Qi
School of Computer Science
Shenyang Aerospace University
Shenyang, China
qihan@sau.edu.cn

5th Changqing Gong
School of Computer Science
Shenyang Aerospace University
Shenyang, China
gongchangqing@sau.edu.cn

VECTOR ENCRYPTION

Partially Homomorphic Encryption of Vectors



Vector Store

Confidentiality-Preserving Image Search: A Comparative Study Between Homomorphic Encryption and Distance-Preserving Randomization

WENJUN LU[1], AVINASH L. VARNA[2], (Member, IEEE), AND MIN WU[3],

[1]Google, Mountain View, CA 95014, USA
[2]Intel, Pheonix, AZ 85226, USA
[3]Department of Electrical and Computer Engineering and Institute for Advanced Computer Studies, University of M

Corresponding author: W. Lu (luwj02@gmail.com)

ABSTRACT Recent years have seen increasing popularity of storing and data using online services. Preserving confidentiality of online persona functionalities thus becomes an important and pressing research issue. In t of content-based search of image data archived online while preserving cont has different settings from those typically considered in the secure computati in rank-ordered search, and has a different security-efficiency requirement. such as homomorphic encryption, can potentially be used in this application and communication complexity. Alternatively, efficient techniques based o search indexes have been proposed recently to enable similarity comparison paper focuses on comparing these two major paradigms of techniques, nar based techniques and feature/index randomization-based techniques, for c search. We develop novel and systematic metrics to quantitatively evaluate

Check for updates

Approximate Distance-Comparison-Preserving Symmetric Encryption

Georg Fuchsbauer[1], Riddhi Ghosal[2(✉)], Nathan Hauke[3], and Adam O'Neill[4]

[1] TU Wien, Vienna, Austria
georg.fuchsbauer@tuwien.ac.at
[2] UCLA, Los Angeles, USA
riddhi@cs.ucla.edu
[3] Georgetown University, Washington D.C., USA
nah52@georgetown.edu
[4] Manning College of Information and Computer Sciences, University of Massachusetts, Amherst, USA
adamo@cs.umass.edu

Abstract. We introduce *distance-comparison-preserving* symmetric encryption (DCPE), a new type of property-preserving encryption that preserves relative distance between plaintext vectors. DCPE is naturally suited for nearest-neighbor search on encrypted data. To boost security,

VEC

PART



Vector Store

Approximate Distance-Comparison-Preserving Symmetric Encryption

Georg Fuchsbauer[1], Riddhi Ghosal[2(✉)], Nathan Hauke[3], and Adam O'Neill[4]

[1] TU Wien, Vienna, Austria
georg.fuchsbauer@tuwien.ac.at
[2] UCLA, Los Angeles, USA
riddhi@cs.ucla.edu
[3] Georgetown University, Washington D.C., USA
nah52@georgetown.edu
[4] Manning College of Information and Computer Sciences, University of Massachusetts, Amherst, USA
adamo@cs.umass.edu

Abstract. We introduce *distance-comparison-preserving* symmetric encryption (DCPE), a new type of property-preserving encryption that preserves relative distance between plaintext vectors. DCPE is naturally suited for nearest-neighbor search on encrypted data. To boost security, we divert from prior work on Property Preserving Encryption (PPE) and ask for *approximate comparison*, which is natural given the prevalence of approximate nearest neighbor (ANN) search. We study what security approximate DCPE can provide and how to construct it.

VECTOR ENCRYPTION

DISTANCE-COMPARISON-PRESERVING ENCRYPTION (DCPE)



Vector Store

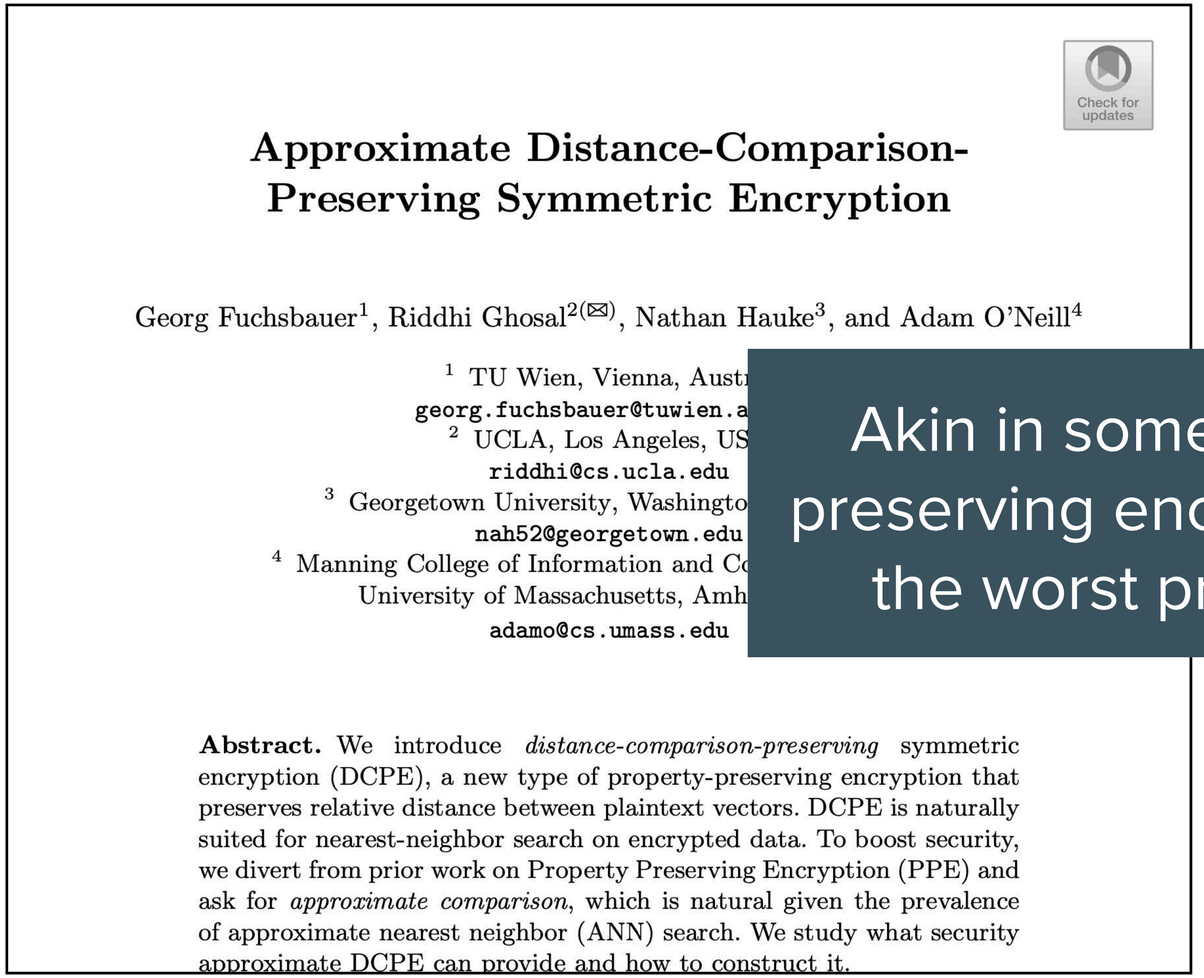

Check for updates

Approximate Distance-Comparison-Preserving Symmetric Encryption

Georg Fuchsbauer[1], Riddhi Ghosal[2(✉)], Nathan Hauke[3], and Adam O'Neill[4]

[1] TU Wien, Vienna, Austr
georg.fuchsbauer@tuwien.a
[2] UCLA, Los Angeles, US
riddhi@cs.ucla.edu
[3] Georgetown University, Washingto
nah52@georgetown.edu
[4] Manning College of Information and C
University of Massachusetts, Amh
adamo@cs.umass.edu

Abstract. We introduce *distance-comparison-preserving* symmetric encryption (DCPE), a new type of property-preserving encryption that preserves relative distance between plaintext vectors. DCPE is naturally suited for nearest-neighbor search on encrypted data. To boost security, we divert from prior work on Property Preserving Encryption (PPE) and ask for *approximate comparison*, which is natural given the prevalence of approximate nearest neighbor (ANN) search. We study what security approximate DCPE can provide and how to construct it.

Akin in some ways to order preserving encryption, but fixes the worst problems there.

HOW IT **WORKS**

AT A HIGH LEVEL: INDEX AND INGEST



Input Text → Embedding Model → Vector Embedding → **Encrypt** → Encrypted Embedding

Store



Vector Datastore

Note: we suggest also encrypting metadata in appropriate ways to protect the related info. But we'll focus on vectors here.

HOW IT **WORKS**

At a high level: Querying

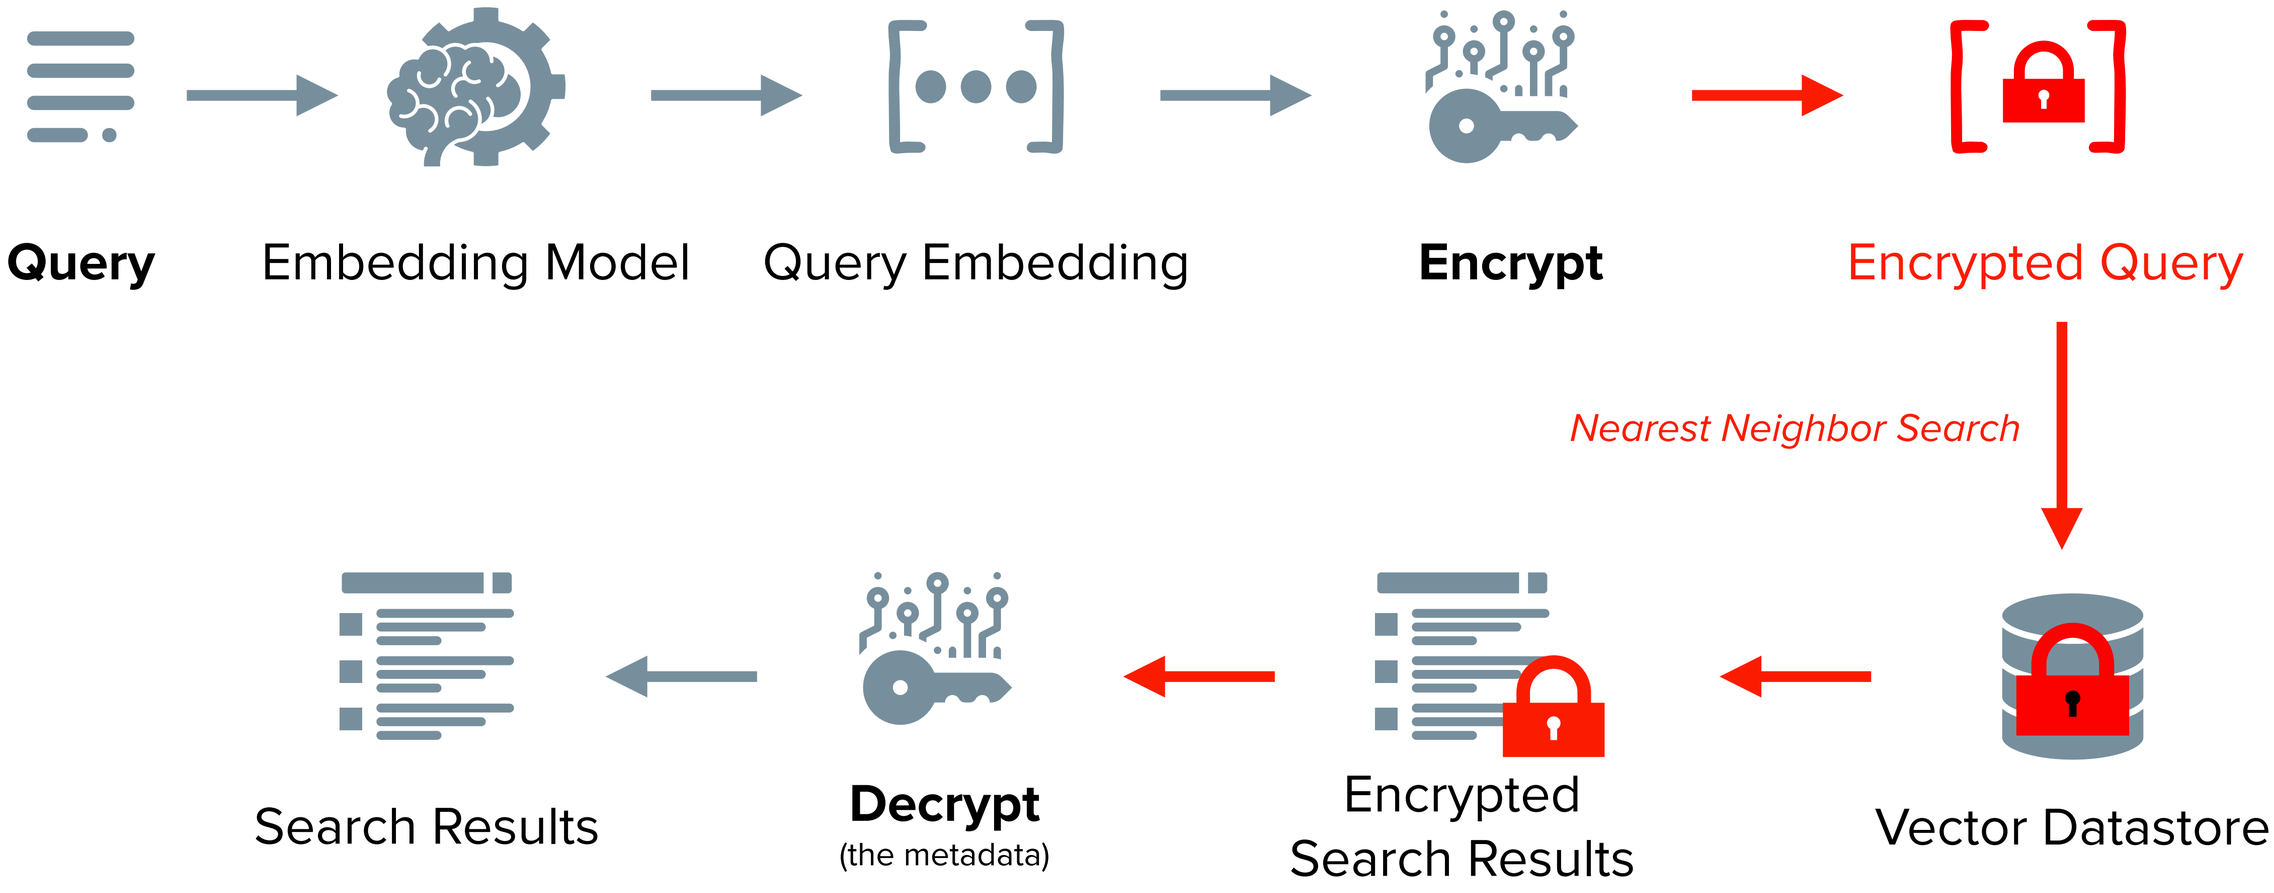

HOW IT **WORKS**

At a high level: private data and queries

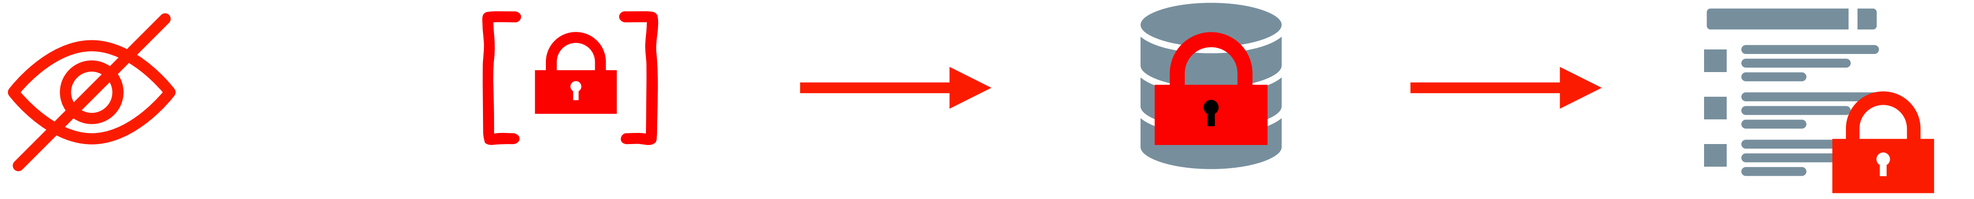

Encrypted Query

Vector Datastore

Encrypted Search Results

Server blind to data and queries

Index can't be inverted

Results retain high accuracy

Negligible latency impact

HOW IT **WORKS**

Distance-Comparison-Preserving Encryption (DCPE)

Approximate-Distance-Comparison-Preserving (β-DCP) Function: For $\beta \in \mathbb{R}^+$, a function $f : \mathcal{X} \to \mathcal{Y}$ is said to be β-DCP if

$$\forall x, y, z \in \mathcal{X} : \mathsf{dist}(x, y) < \mathsf{dist}(x, z) - \beta \implies \mathsf{dist}(f(x), f(y)) < \mathsf{dist}(f(x), f(z)).$$

NOTE: dist() could be ANN and error could be $\beta + c$

HOW IT **WORKS**

Scale and Perturb (SAP)

Algorithm 2. The SAP scheme.

procedure $\mathcal{K}eyGen_{\mathsf{SAP}}()$
$s \xleftarrow{\$} \mathsf{S}$
$K \xleftarrow{\$} \{0,1\}^k$
return (s, K)

procedure $\mathcal{E}nc_{\mathsf{SAP}}((s, K), \mathbf{m})$
$n \xleftarrow{\$} \{0,1\}^\ell$
$\mathsf{coin}_1 \| \mathsf{coin}_2 \leftarrow \mathsf{PRF}(K, n)$
$\mathbf{u} \leftarrow \mathcal{N}(0, I_d; \mathsf{coin}_1)$
$x' \leftarrow \mathcal{U}(0, 1; \mathsf{coin}_2)$
$x \leftarrow \frac{s\beta}{4}(x')^{\frac{1}{d}}$; $\lambda_{\mathbf{m}} \leftarrow \frac{\mathbf{u}x}{\|\mathbf{u}\|}$
$\mathbf{c} \leftarrow s\mathbf{m} + \lambda_{\mathbf{m}}$
return $(\mathbf{c}, n)$

procedure $\mathcal{D}ec_{\mathsf{SAP}}((s, K), (\mathbf{c}, n))$
$\mathsf{coin}_1 \| \mathsf{coin}_2 \leftarrow \mathsf{PRF}(K, n)$
$\mathbf{u} \leftarrow \mathcal{N}(0, I_d; \mathsf{coin}_1)$
$x' \leftarrow \mathcal{U}(0, 1; \mathsf{coin}_2)$
$x \leftarrow \frac{s\beta}{4}(x')^{\frac{1}{d}}$
$\lambda_{\mathbf{m}} \leftarrow \frac{\mathbf{u}x}{\|\mathbf{u}\|}$
$\mathbf{m} \leftarrow \frac{\mathbf{c} - \lambda_{\mathbf{m}}}{s}$
return $\mathbf{m}$

HOW IT **WORKS**

Scale and Perturb (SAP)

Algorithm 2. The SAP scheme.

procedure $\mathcal{K}eyGen_{\mathsf{SAP}}()$
$s \xleftarrow{\$} \mathsf{S}$
$K \xleftarrow{\$} \{0,1\}^k$
return (s, K)

procedure $\mathcal{E}nc_{\mathsf{SAP}}((s, K), \mathbf{m})$
$n \xleftarrow{\$} \{0,1\}^\ell$
$\mathsf{coin}_1 \| \mathsf{coin}_2 \leftarrow \mathsf{PRF}(K, n)$
$\mathbf{u} \leftarrow \mathcal{N}(0, I_d; \mathsf{coin}_1)$
$x' \leftarrow \mathcal{U}(0, 1; \mathsf{coin}_2)$
$x \leftarrow \frac{s\beta}{4}(x')^{\frac{1}{d}}$; $\lambda_{\mathbf{m}} \leftarrow \frac{\mathbf{u}x}{\|\mathbf{u}\|}$
$\mathbf{c} \leftarrow s\mathbf{m} + \lambda_{\mathbf{m}}$
return $(\mathbf{c}, n)$

procedure $\mathcal{D}ec_{\mathsf{SAP}}((s, K), (\mathbf{c}, n))$
$\mathsf{coin}_1 \| \mathsf{coin}_2 \leftarrow \mathsf{PRF}(K, n)$
$\mathbf{u} \leftarrow \mathcal{N}(0, I_d; \mathsf{coin}_1)$
$x' \leftarrow \mathcal{U}(0, 1; \mathsf{coin}_2)$
$x \leftarrow \frac{s\beta}{4}(x')^{\frac{1}{d}}$
$\lambda_{\mathbf{m}} \leftarrow \frac{\mathbf{u}x}{\|\mathbf{u}\|}$
$\mathbf{m} \leftarrow \frac{\mathbf{c} - \lambda_{\mathbf{m}}}{s}$
return $\mathbf{m}$

- Plaintext vector is multiplied by scaling factor in secret key (***s***) then each attribute is adjusted up/down by some amount.
- There's a per-vector random "IV" (***n***) so no two vectors are encrypted the same. This is needed to decrypt.
- Perturbation (***λ***) amount depends on ***β*** with higher bit security for higher ***β***.

HOW IT **WORKS**

Scale and Perturb (SAP)

Algorithm 2. The SAP scheme.

procedure $\mathcal{K}eyGen_{\mathsf{SAP}}()$
$s \xleftarrow{\$} \mathsf{S}$
$K \xleftarrow{\$} \{0,1\}^k$
return (s, K)

procedure $\mathcal{E}nc_{\mathsf{SAP}}((s, K), \mathbf{m})$
$n \xleftarrow{\$} \{0,1\}^\ell$
$\mathsf{coin}_1 \| \mathsf{coin}_2 \leftarrow \mathsf{PRF}(K, n)$
$\mathbf{u} \leftarrow \mathcal{N}(0, I_d; \mathsf{coin}_1)$
$x' \leftarrow \mathcal{U}(0, 1; \mathsf{coin}_2)$
$x \leftarrow \frac{s\beta}{4}(x')^{\frac{1}{d}}$; $\lambda_{\mathbf{m}} \leftarrow \frac{\mathbf{u}x}{\|\mathbf{u}\|}$
$\mathbf{c} \leftarrow s\mathbf{m} + \lambda_{\mathbf{m}}$
return $(\mathbf{c}, n)$

procedure $\mathcal{D}ec_{\mathsf{SAP}}((s, K), (\mathbf{c}, n))$
$\mathsf{coin}_1 \| \mathsf{coin}_2 \leftarrow \mathsf{PRF}(K, n)$
$\mathbf{u} \leftarrow \mathcal{N}(0, I_d; \mathsf{coin}_1)$
$x' \leftarrow \mathcal{U}(0, 1; \mathsf{coin}_2)$
$x \leftarrow \frac{s\beta}{4}(x')^{\frac{1}{d}}$
$\lambda_{\mathbf{m}} \leftarrow \frac{\mathbf{u}x}{\|\mathbf{u}\|}$
$\mathbf{m} \leftarrow \frac{\mathbf{c} - \lambda_{\mathbf{m}}}{s}$
return $\mathbf{m}$

Plus a column shuffle to make analysis harder

- Plaintext vector is multiplied by scaling factor in secret key (***s***) then each attribute is adjusted up/down by some amount.
- There's a per-vector random "IV" (***n***) so no two vectors are encrypted the same. This is needed to decrypt.
- Perturbation (***λ***) amount depends on ***β*** with higher bit security for higher ***β***.

HOW IT **WORKS**

Distance-Comparison-Preserving Encryption (DCPE)

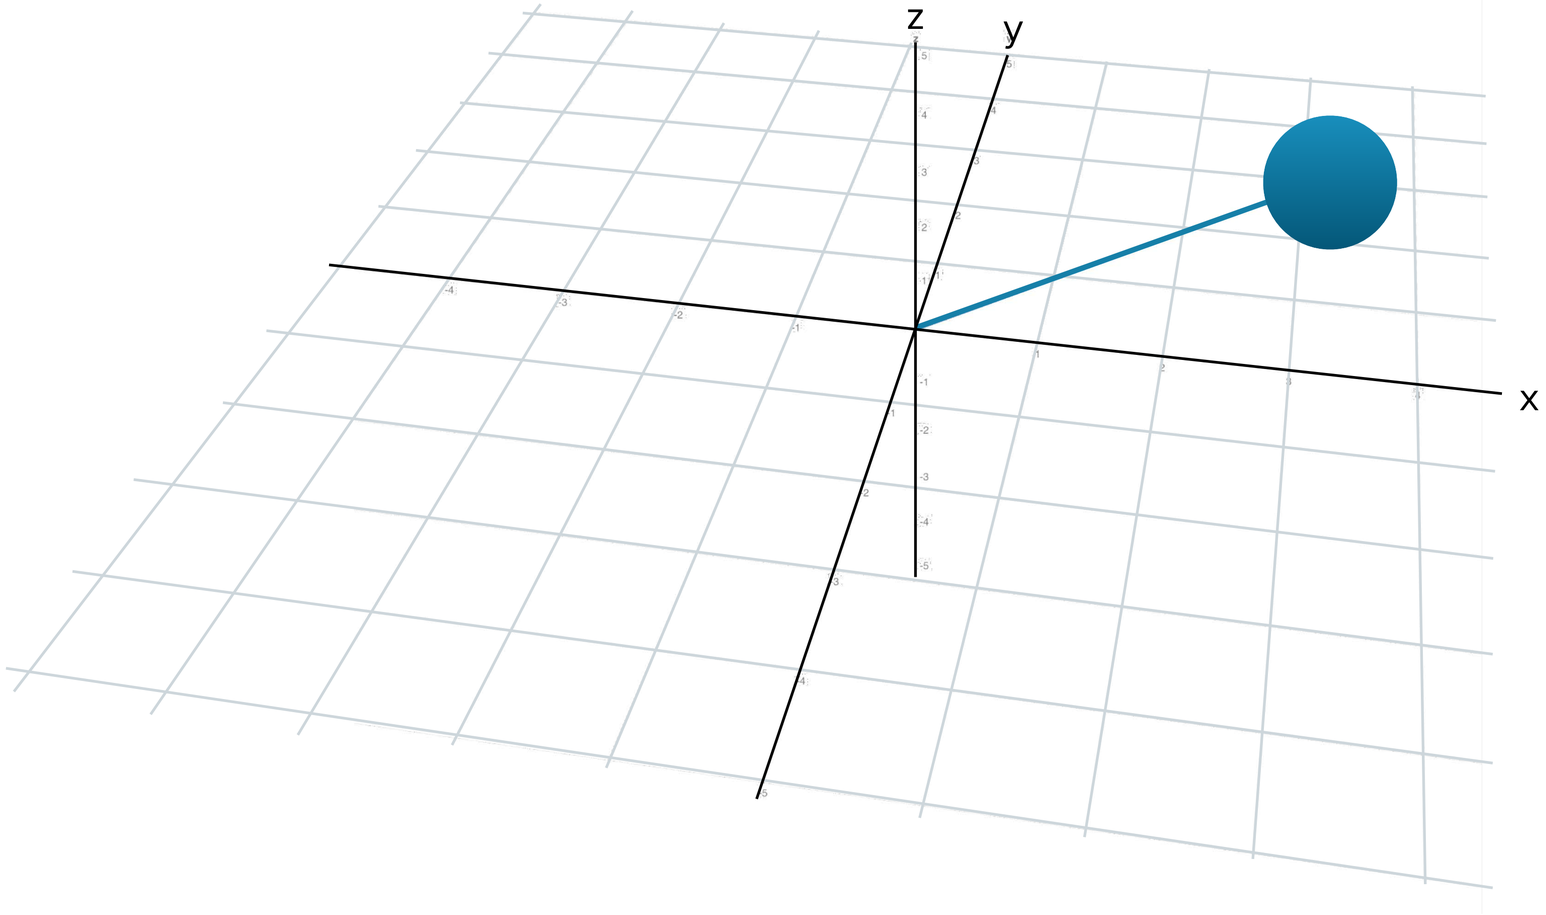

HOW IT **WORKS**

Distance-Comparison-Preserving Encryption (DCPE)

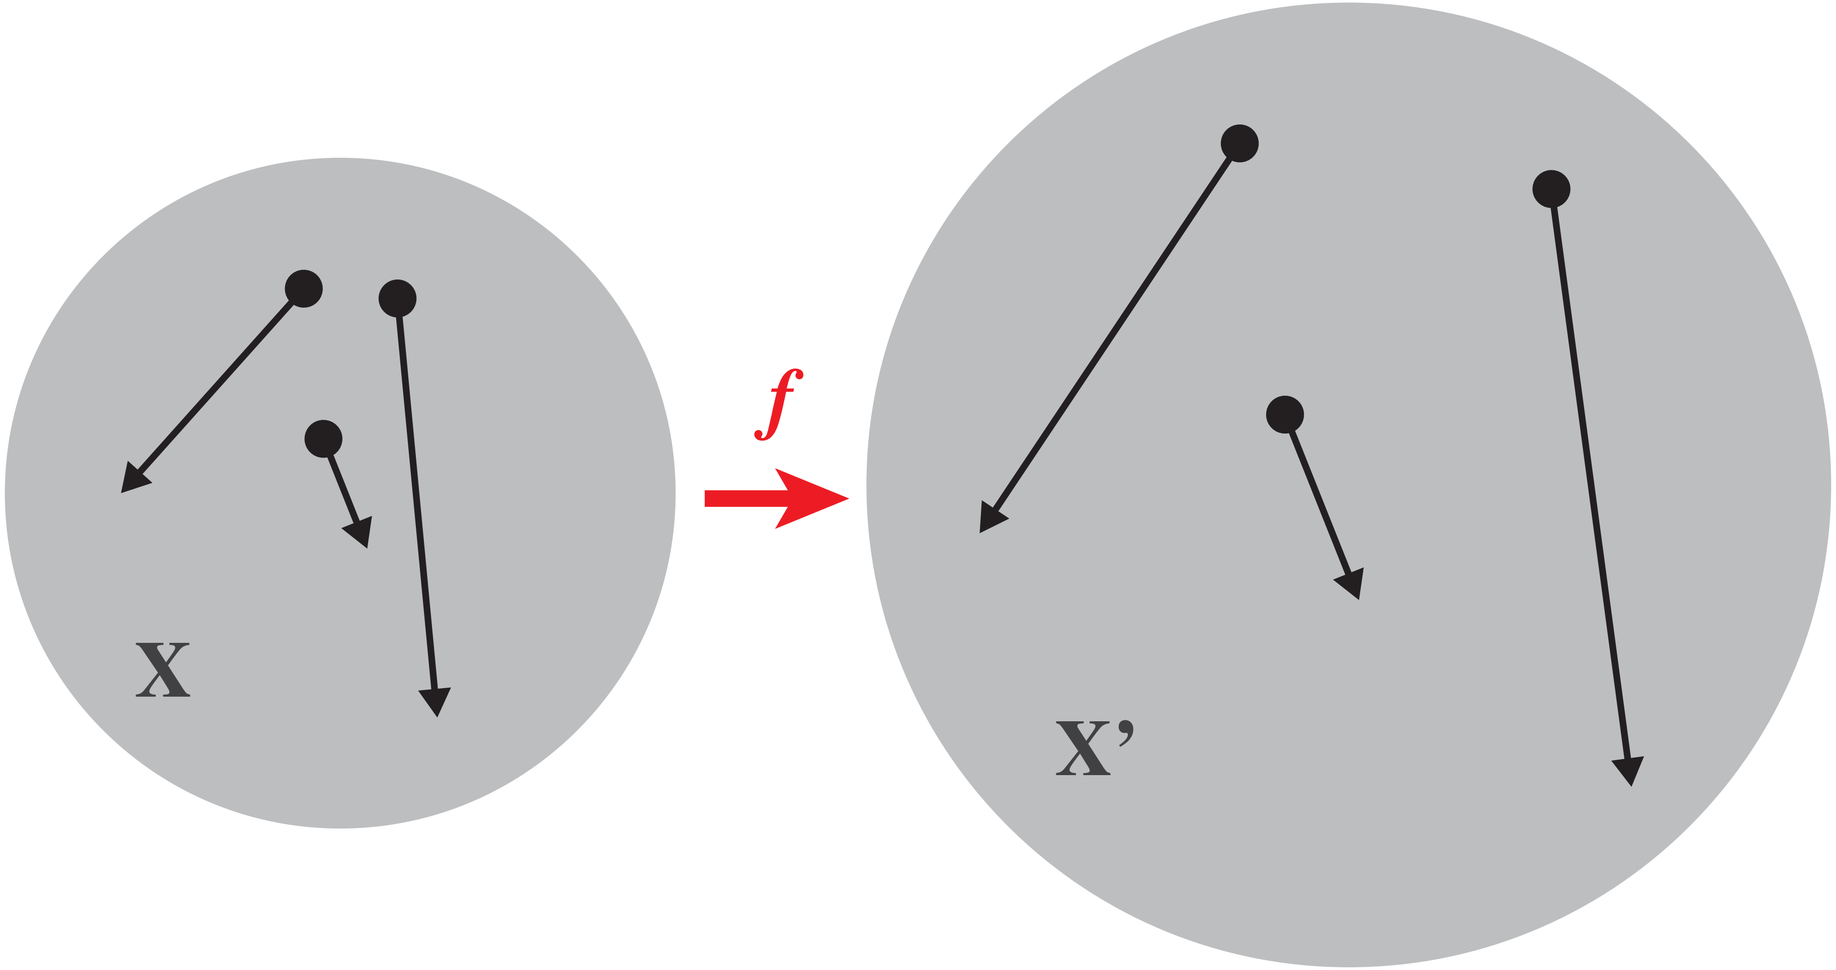

HOW IT **WORKS**

Distance-Comparison-Preserving Encryption (DCPE)

```
  +          +          -          +          +          -          +
[0.00818394, 0.03740552, -0.059338, 0.04153062, 0.066233,  -0.0868215, 0.0082078, ...]
```

f Encrypt (note: we disabled the attribute shuffle for this)



```
  +          -          +          -          +          -          +
[-126650.46, -20264.875, 36461.188, -1050889.5, 855246.1, -1388306.6, 442813.88, ...]
```

OPEN SOURCE LIBRARY

IronCore Alloy: https://github.com/ironcorelabs/ironcore-alloy/

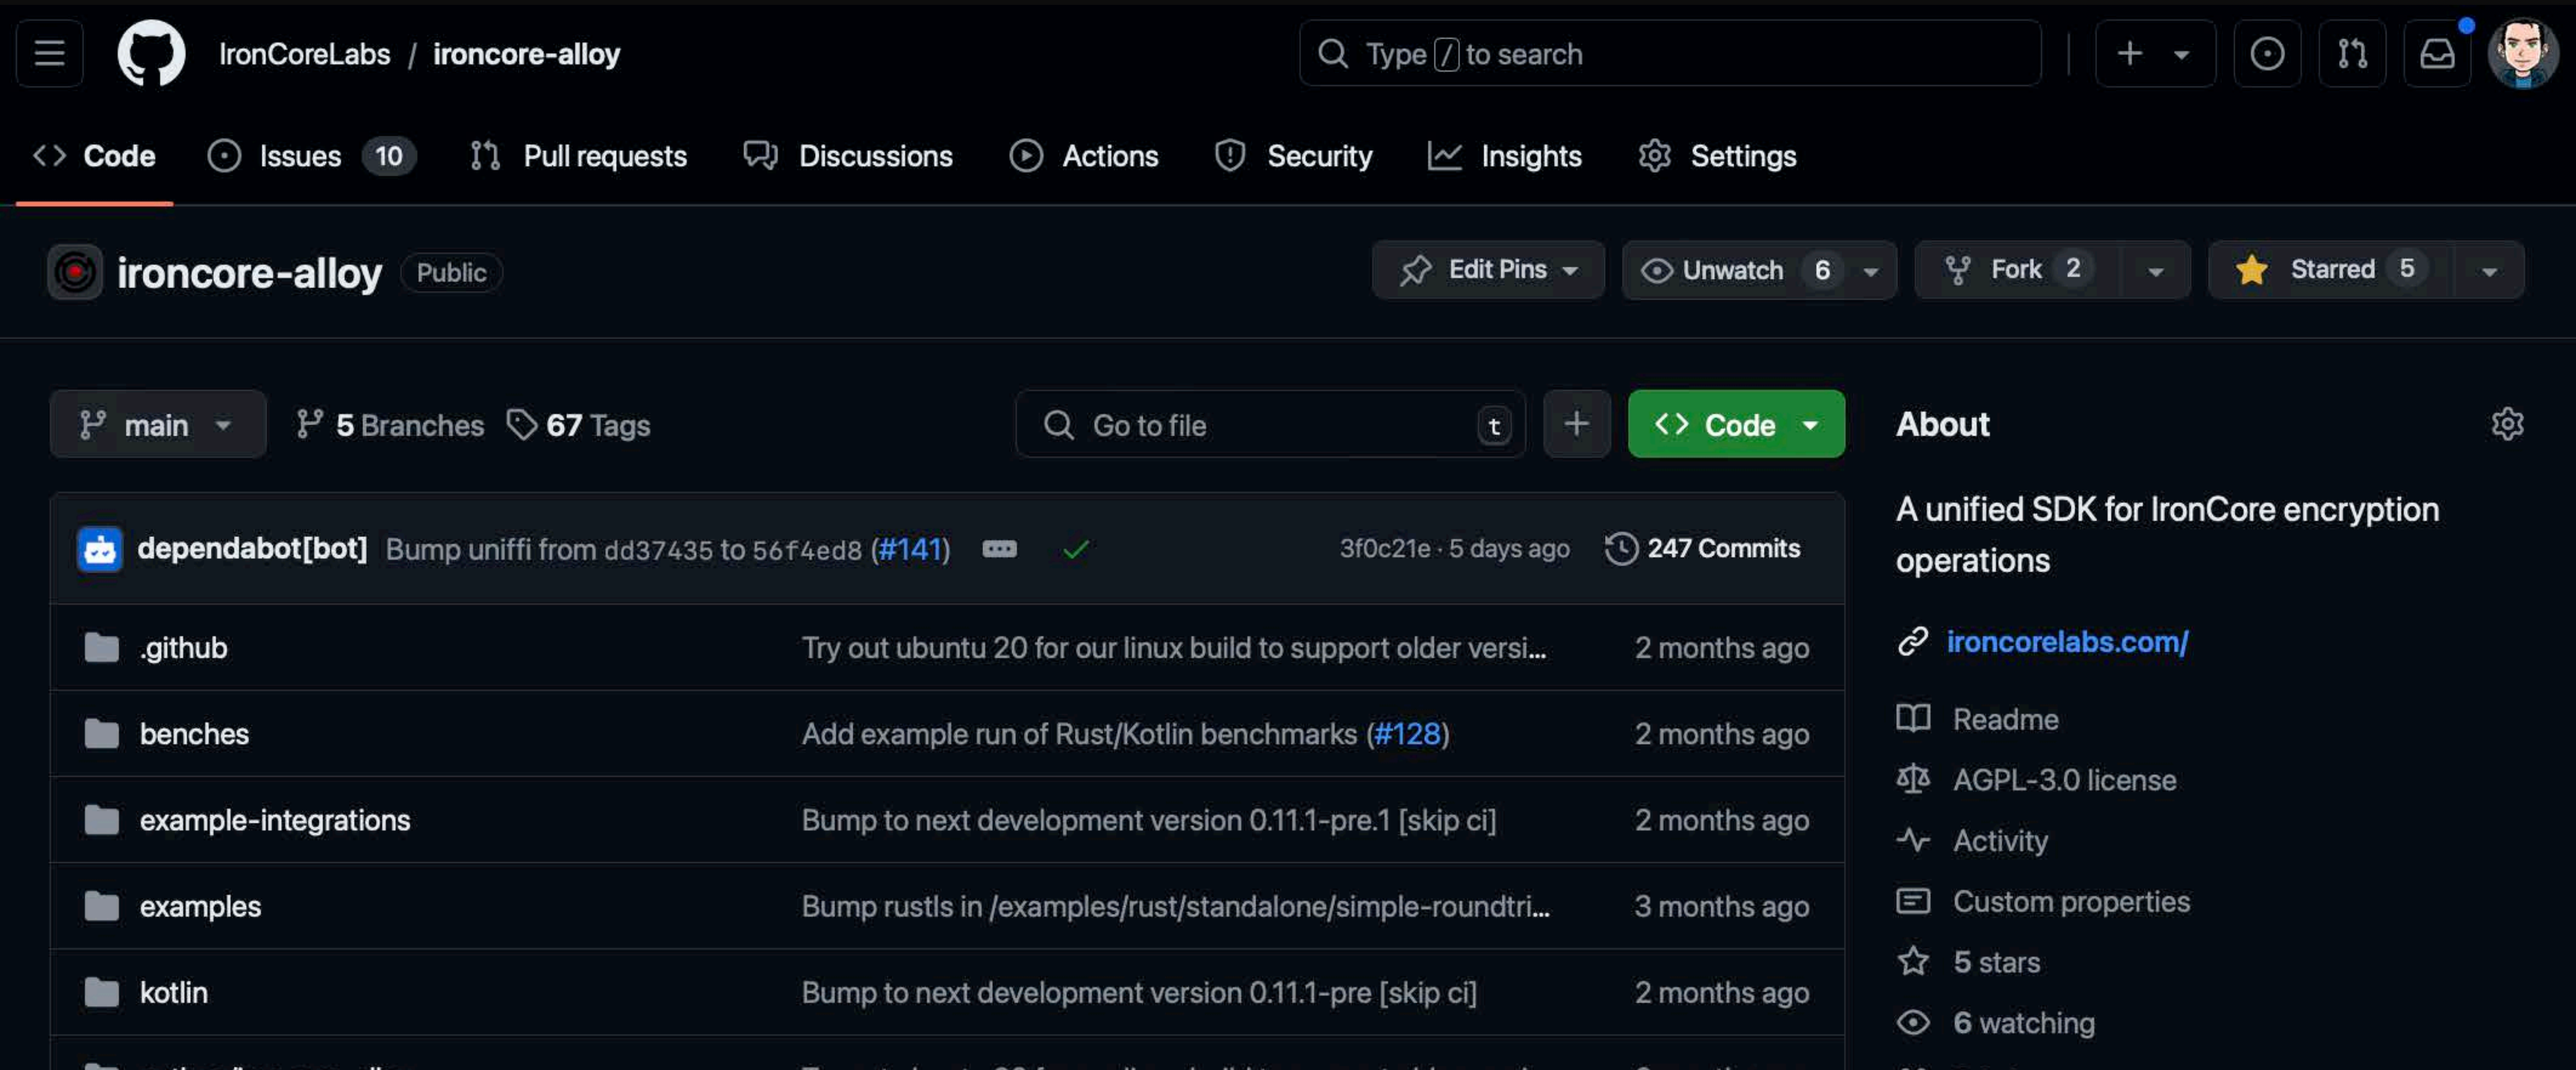

OPEN SOURCE LIBRARY

IronCore Alloy: https://github.com/ironcorelabs/ironcore-alloy/

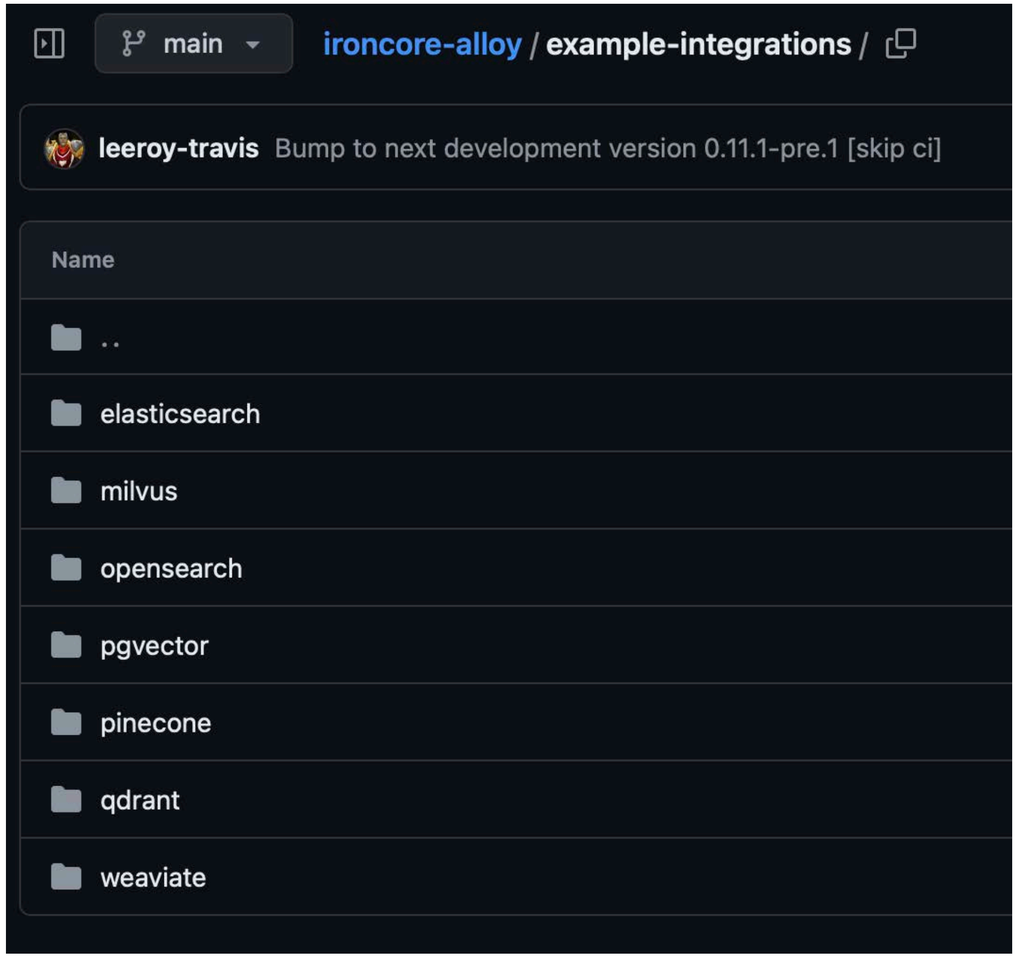


```
question_emb = model.encode(row["Question"]).tolist()
plaintext_vector = alloy.PlaintextVector(question_emb,
  secret_path,
  derivation_path)
encrypted_vector = await alloy_client.vector().encrypt(
  plaintext_vector,
  metadata)
```

OPEN SOURCE LIBRARY

IronCore Alloy: https://github.com/ironcorelabs/ironcore-alloy/

Dataset	Model	Approximation Factor	Precision Loss	Recall Loss
nfcorpus	all-miniLM-L6-v2 (dot)	1.5	4%	6%
scifact	all-miniLM-L6-v2 (cos_sim)	1.5	3%	3%
nfcorpus	msmarco-distilbert-base-tas-b (dot)	1.5	2%	5%
nfcorpus	paraphrase-MiniLM-L3-v2 (dot)	1.5	8%	13%
nfcorpus	text-embedding-ada-002 (dot)	1.5	7%	8%

ATTACK ENCRYPTED VECTORS

Using an inversion model trained on the embedding model outputs

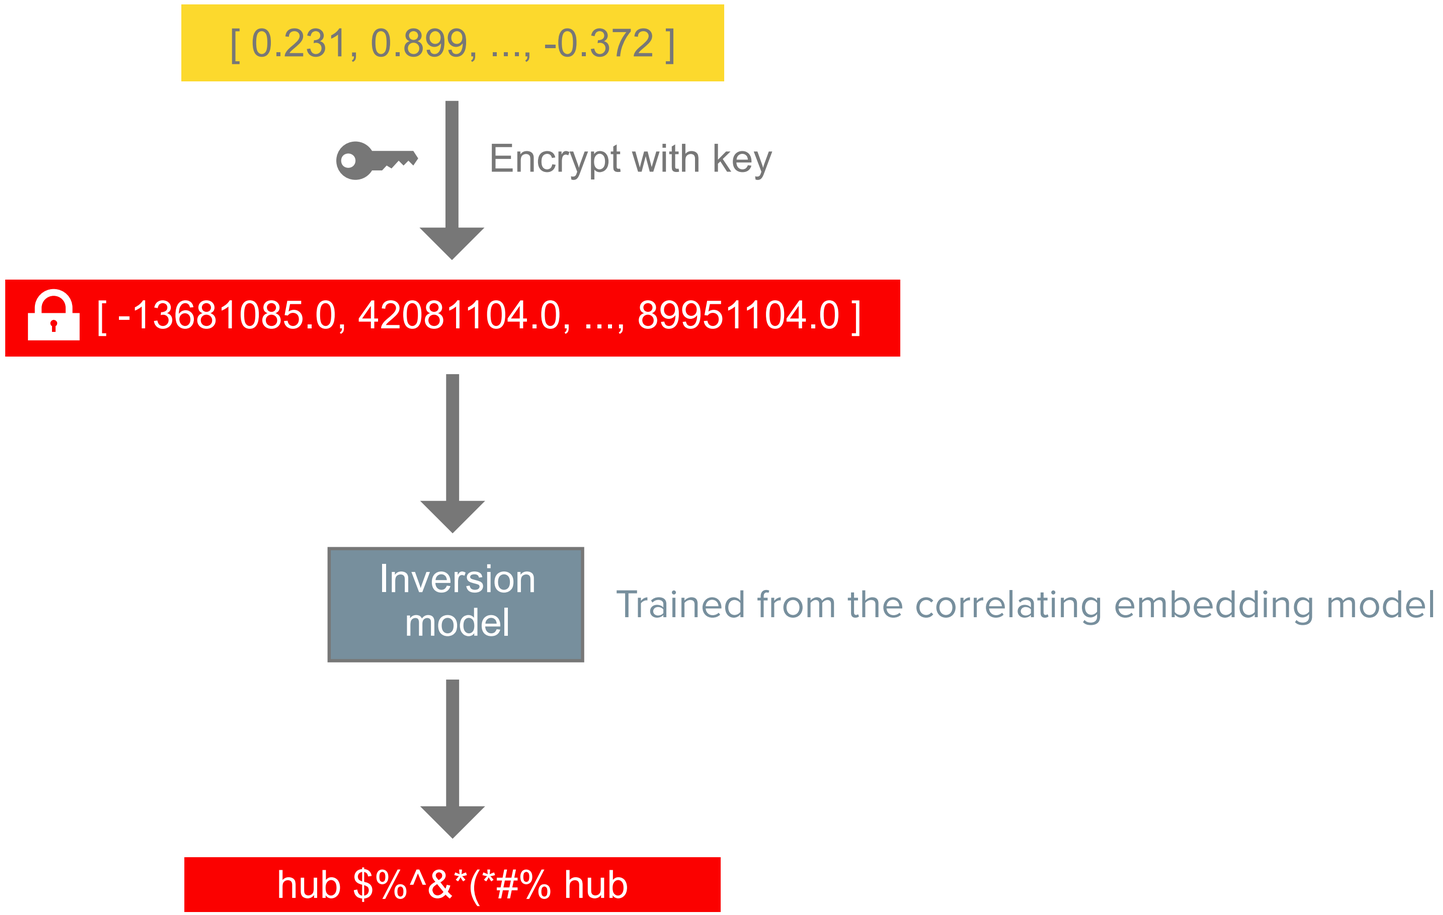

CHOSEN PLAINTEXT ATTACK

USING AN INVERSION MODEL **TRAINED ON THE ENCRYPTED DATA**

Assume: We can query an oracle and for any given plaintext, we can get back the corresponding ciphertext.

Given: We use these corresponding values to *train an inversion model* that's *specific to the key* used by the oracle.

Question: What results do we get?

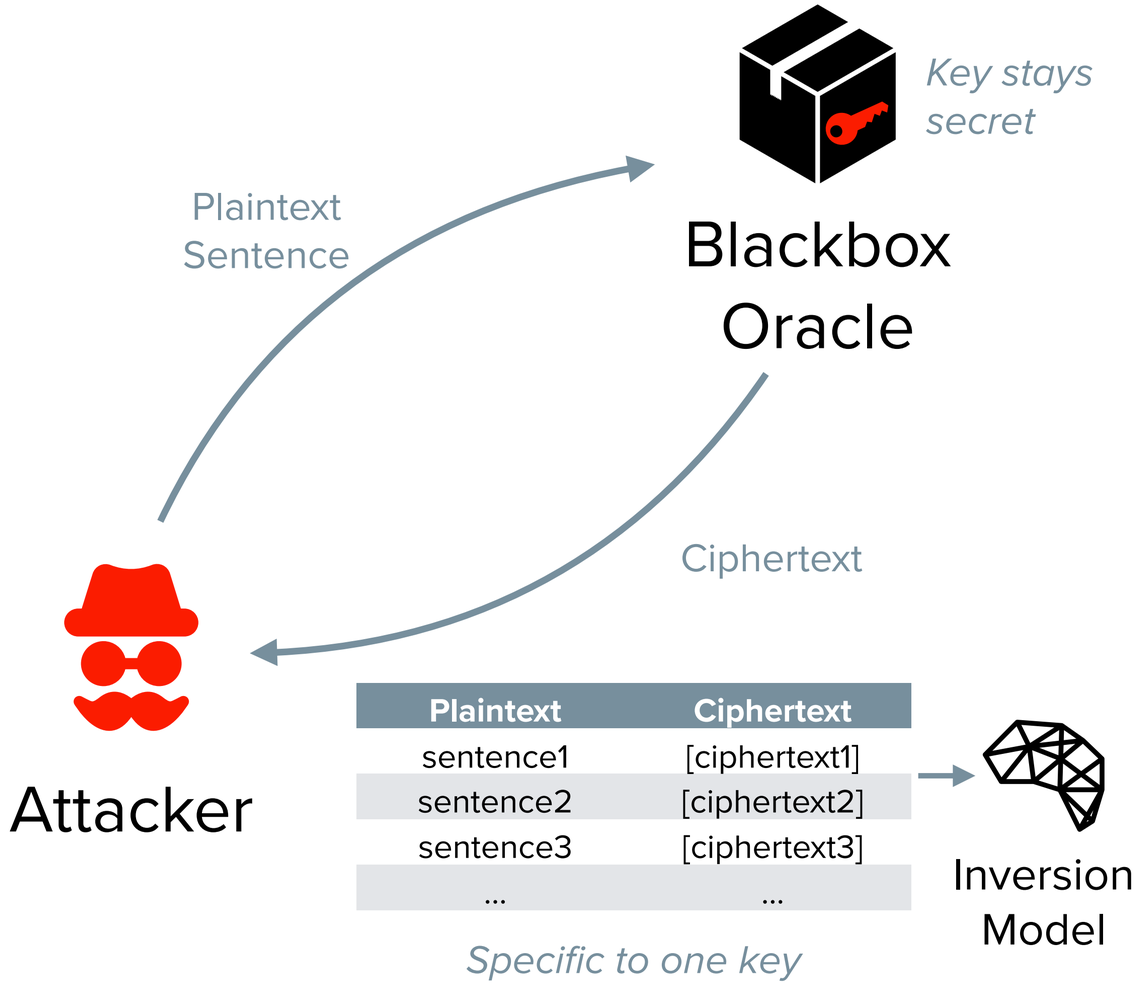

CHOSEN PLAINTEXT ATTACK

USING AN INVERSION MODEL **TRAINED ON THE ENCRYPTED DATA**

Answer: *It depends* on the approximation factor, but the jitteriness of the encrypted values makes *the results poor* -- somewhere between "right general idea or theme of text with wrong specifics" to "laughably far off the mark".

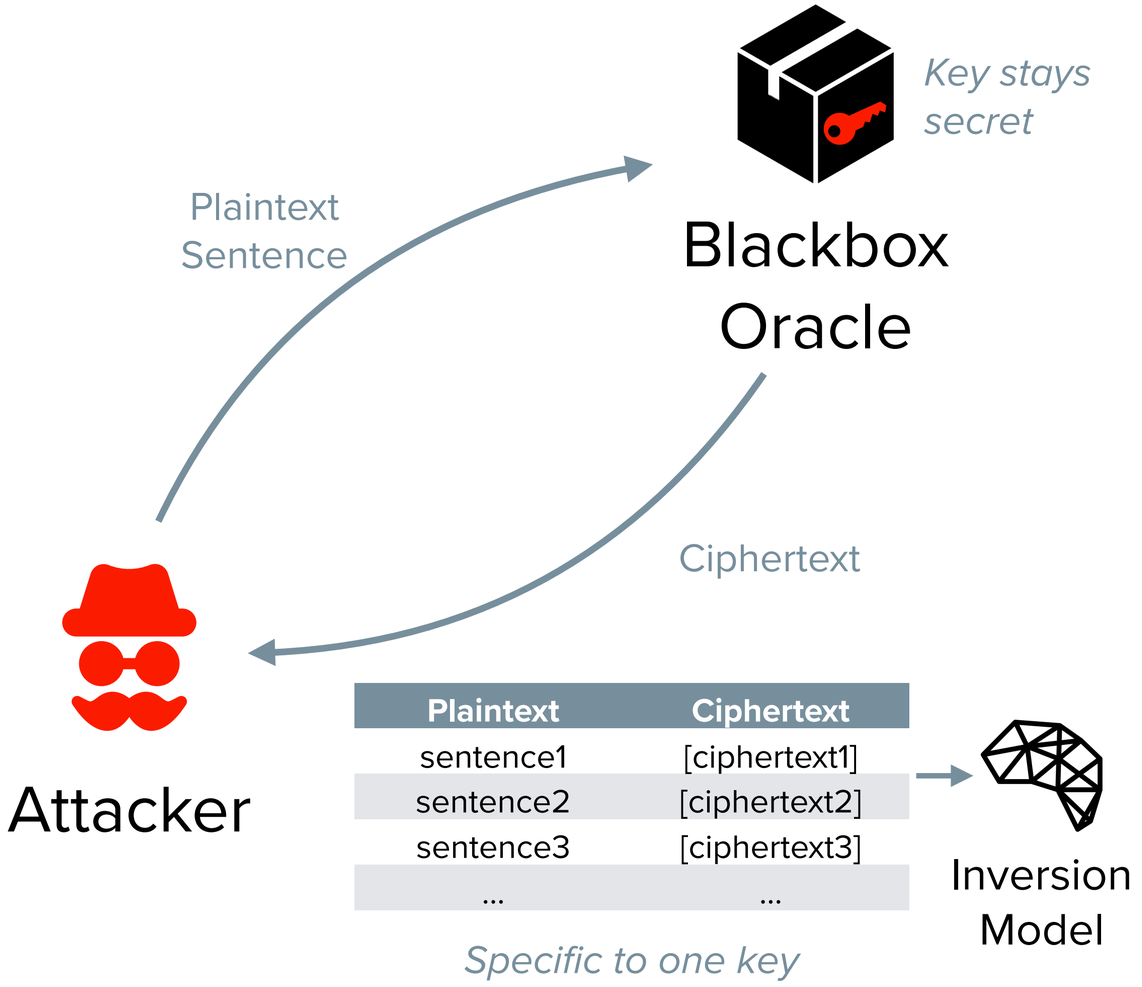

CHOSEN PLAINTEXT ATTACK

BIT SECURITY AND KEY RECOVERY ANALYSIS

Question: Given an associated plaintext vector and encrypted vector **and the IV** for that vector, can we figure out the key?

Answer: No

The key is protected by the PRF(K, n) function. Even knowing n, we'd have to exhaustively check all possible values of K in a brute force attack to find the key. **In ironcore-alloy, K is 256 bits.**

$$
\begin{aligned}
&\textbf{procedure } \mathcal{D}ec_{\mathsf{SAP}}((s, K), (\mathbf{c}, n)) \\
&\quad \mathsf{coin}_1 \| \mathsf{coin}_2 \leftarrow \mathsf{PRF}(K, n) \\
&\quad \mathbf{u} \leftarrow \mathcal{N}(0, I_d; \mathsf{coin}_1) \\
&\quad x' \leftarrow \mathcal{U}(0, 1; \mathsf{coin}_2) \\
&\quad x \leftarrow \tfrac{s\beta}{4}(x')^{\frac{1}{d}} \\
&\quad \lambda_{\mathbf{m}} \leftarrow \tfrac{\mathbf{u}x}{\|\mathbf{u}\|} \\
&\quad \mathbf{m} \leftarrow \tfrac{\mathbf{c} - \lambda_{\mathbf{m}}}{s} \\
&\quad \text{return } \mathbf{m}
\end{aligned}
$$

= HMAC-SHA256(key, IV)

Plaintext	Ciphertext	IV
plainvector1	[ciphertext1]	number1
plainvector2	[ciphertext2]	number2
plainvector3	[ciphertext3]	number3
...	...	...

Data specific to one key

SUMMARY 1/3

1. RAG workflows hide shadow copies of sensitive data and **sprinkle them around new systems**.

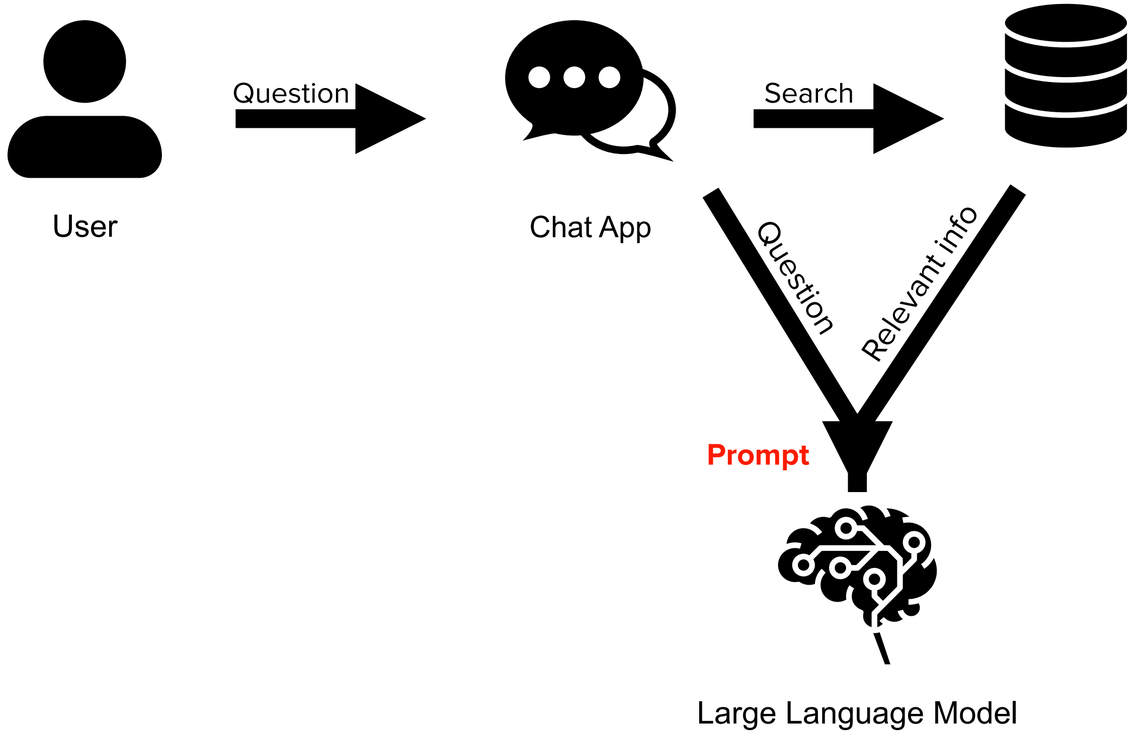

SUMMARY 2/3

2. Vectors in vector databases can be **inverted back to their original input** with high fidelity using low effort.

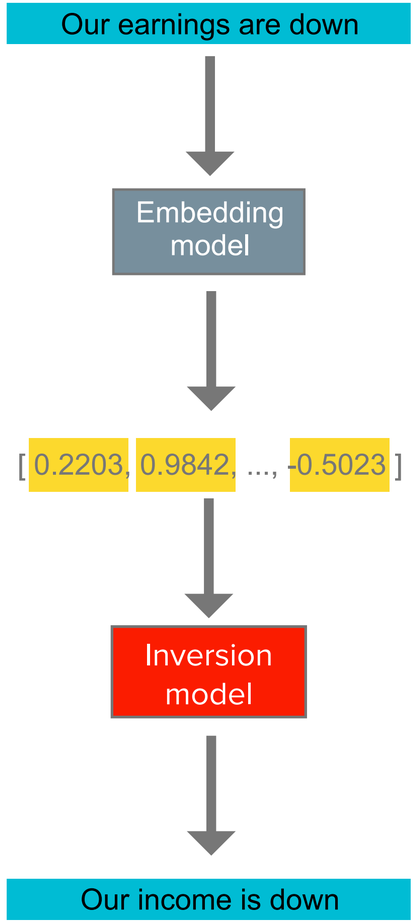

SUMMARY 3/3

3. Approximate distance preserving encryption protects vectors in any vector database and **prevents inversion with minimal downsides**.



Any vector-capable database (even when hosted by 3rd-party)

THANK YOU

Patrick Walsh, CEO, IronCore Labs

Bob Wall, CTO, IronCore Labs

https://ironcorelabs.com/resources/


bob.wall@ironcoreelabs.com


@bithead_bob


@bobwall23@infosec.exchange


bobwall23


patrick.walsh@ironcoreelabs.com


@zmre


@mastodon.social@zmre


zmre