

**WHAT GOES  
AROUND COMES  
AROUND...  
AND AROUND...**



**Andy Pavlo**  
PGConf NYC  
October 2023

**Carnegie  
Mellon  
University**



<https://cmudb.io/pgconf23>

# DATABASES

---

A database's **data model** is the underlying structure and organization of data within the database.

The **relational model** (RM) + **SQL** have dominated the database landscape since the 1980s.

But every 10 years somebody invents a RM/SQL "killer" that addresses some deficiency...

## DATABASES



**Jo Kristian Bergum**

@jobergum

Tensor and vector databases will replace most legacy databases in this decade. A disruption fueled by natural language interfaces and deep neural representations. In other words:

Natural query languages (NQL) replace the structured query language (SQL).

2:35 AM · Apr 27, 2023 · **177.2K** Views

**39** Retweets

**32** Quotes

**330** Likes

**196** Bookmarks

g structure

ase.

e 1980s.

/SQL

## DATABASES



**Jo Kristian Bergum**  
@jobergum

Tensor and vector databases  
decade. A disruption fueled by  
neural representations. In other

Natural query languages (NQL)  
(SQL).

2:35 AM · Apr 27, 2023 · 177.2K

39 Retweets 32 Quotes 33



**Gagan Biyani**   
@gaganbiyani

SQL is going to die at the hands of an AI. I'm serious.

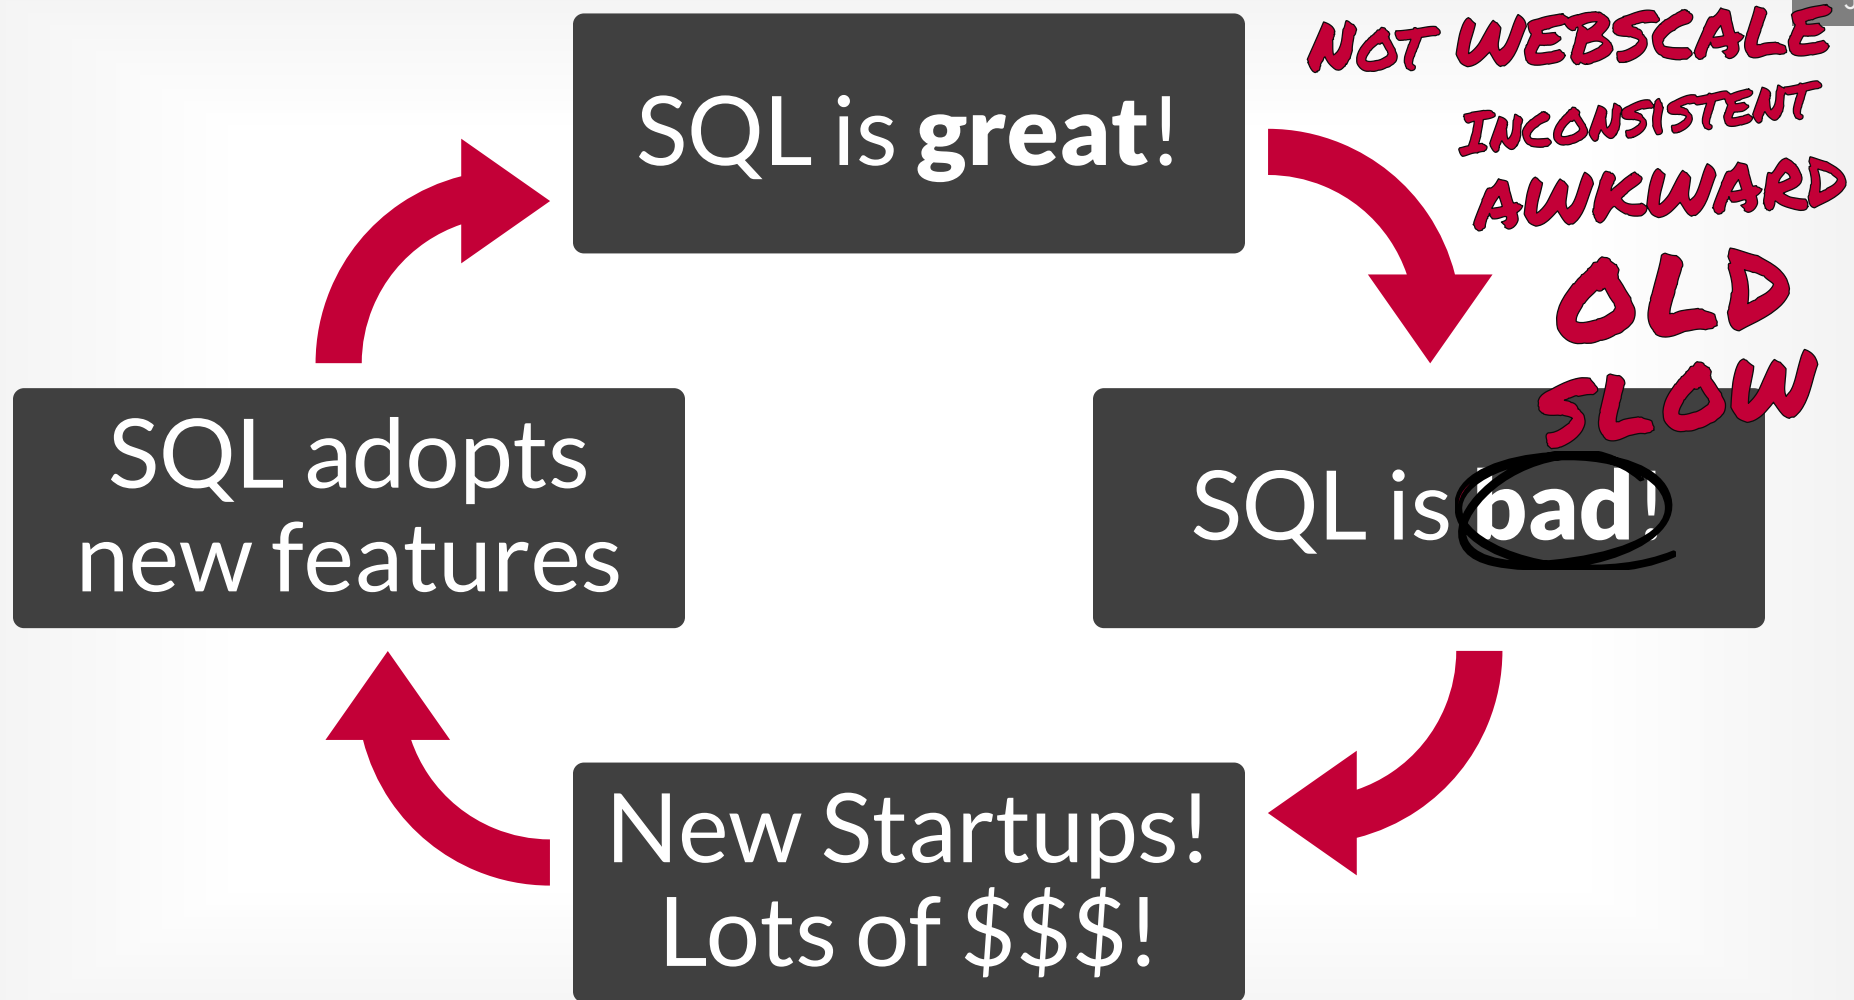
@mayowaoshin is already doing this. Takes your company's data and  
ingests it into ChatGPT. Then, you can create a chatbot for the data and  
just ask it questions using natural language.

This video demos the output.



10:30 AM · May 18, 2023 · 2.6M Views

247 Retweets 203 Quotes 2,842 Likes 3,624 Bookmarks



## What Goes Around Comes Around

Michael Stonebraker  
Joseph M. Hellerstein

### Abstract

This paper provides a summary of 35 years of data model proposals, grouped into 9 different eras. We discuss the proposals of each era, and show that there are only a few basic data modeling ideas, and most have been around a long time. Later proposals inevitably bear a strong resemblance to certain earlier proposals. Hence, it is a worthwhile exercise to study previous proposals.

In addition, we present the lessons learned from the exploration of the proposals in each era. Most current researchers were not around for many of the previous eras, and have limited (if any) understanding of what was previously learned. There is an old adage that he who does not understand history is condemned to repeat it. By presenting “ancient history”, we hope to allow future researchers to avoid replaying history.

Unfortunately, the main proposal in the current XML era bears a striking resemblance to the CODASYL proposal from the early 1970’s, which failed because of its complexity. Hence, the current era is replaying history, and “what goes around comes around”. Hopefully the next era will be smarter.

### 1 Introduction

Data model proposals have been around since the late 1960’s, when the first author “came on the scene”. Proposals have continued with surprising regularity for the intervening 35 years. Moreover, many of the current day proposals have come from researchers too young to have learned from the discussion of earlier ones. Hence, the purpose of this paper is to summarize 35 years worth of “progress” and point out what should be learned from this lengthy exercise.

We present data model proposals in nine historical epochs:

Hierarchical (IMS): late 1960’s and 1970’s  
Network (CODASYL): 1970’s  
Relational: 1970’s and early 1980’s  
Entity-Relationship: 1970’s  
Extended Relational: 1980’s  
Semantic: late 1970’s and 1980’s  
Object-oriented: late 1980’s and early 1990’s  
Object-relational: late 1980’s and early 1990’s

<https://cmudb.io/wgaca>

# WHAT GOES AROUND COMES AROUND

## READINGS IN DB SYSTEMS, 4TH EDITION (2005)

**Hierarchical (1960s)**

**Network (1960s)**

**BCE**

**Relational (1970s)**

**Entity-Relationship (1970s)**

**Extended Relational (1980s)**

**Semantic (1980s)**

**Object-Oriented (1980s)**

**Object-Relational (1990s)**

**Semi-Structured/XML (1990s)**

## What Goes Around Comes Around

Michael Stonebraker  
Joseph M. Hellerstein

### Abstract

This paper provides a summary of 35 years of data model proposals, grouped into 9 different eras. We discuss the proposals of each era, and show that there are only a few basic data modeling ideas, and most have been around a long time. Later proposals inevitably bear a strong resemblance to certain earlier proposals. Hence, it is a worthwhile exercise to study previous proposals.

In addition, we present the lessons learned from the exploration of the proposals in each era. Most current researchers were not around for many of the previous eras, and have limited (if any) understanding of what was previously learned. There is an old adage that he who does not understand history is condemned to repeat it. By presenting "ancient history", we hope to allow future researchers to avoid replaying history.

Unfortunately, the main proposal in the current XML era bears a striking resemblance to the CODASYL proposal from the early 1970's, which failed because of its complexity. Hence, the current era is replaying history, and "what goes around comes around". Hopefully the next era will be smarter.

### 1 Introduction

Data model proposals have been around since the late 1960's, when the first author "came on the scene". Proposals have continued with surprising regularity for the intervening 35 years. Moreover, many of the current day proposals have come from researchers too young to have learned from the discussion of earlier ones. Hence, the purpose of this paper is to summarize 35 years worth of "progress" and point out what should be learned from this lengthy exercise.

We present data model proposals in nine historical epochs:

Hierarchical (IMS): late 1960's and 1970's  
Network (CODASYL): 1970's  
Relational: 1970's and early 1980's  
Entity-Relationship: 1970's  
Extended Relational: 1980's  
Semantic: late 1970's and 1980's  
Object-oriented: late 1980's and early 1990's  
Object-relational: late 1980's and early 1990's

<https://cmudb.io/wgaca>

# WHAT GOES AROUND COMES AROUND

READINGS IN DB SYSTEMS, 4TH EDITION (2005)

Hierarchical (1960s)

Network (1960s)

"Before Codd Era"

Relational (1970s)

Entity-Relationship (1970s)


Extended Relational (1980s)

Semantic (1980s)

Object-Oriented (1980s)

Object-Relational (1990s)

Semi-Structured/XML (1990s)



## What Goes Around Comes Around... And Around...

Michael Stonebraker  
Massachusetts Institute of Technology  
stonebraker@csail.mit.edu

Andrew Pavlo  
Carnegie Mellon University  
pavlo@cs.cmu.edu

### ABSTRACT

Two decades ago, one of us co-authored a paper commenting on the previous 40 years of data modelling research and development [98]. That paper demonstrated that SQL and the relational model (RM) reigned supreme for database management systems (DBMSs), and all the efforts to completely replace either the query language or the data model had failed. Instead, SQL absorbed the best ideas from these alternative approaches.

We revisit this issue and argue that little has changed since 2005. Once again there are repeated efforts to replace either SQL or the RM, and none have been successful. Instead, in the last few years the conventional wisdom in industry has swung back to relational DBMSs that use SQL. We suggest that system builders examine history before they invent more query languages or data models that are likely to fail. We also discuss the evolution of DBMS implementations and argue that the major advancements have been in RM systems.

### 1 Introduction

In 2005, one of the authors participated in writing a chapter for the *Red Book* titled “What Goes Around Comes Around” [98]. That paper examined the major data modelling movements since the 1960s. Those were:

- Hierarchical (e.g., IMS): late 1960s and 1970s
- Network (e.g., CODASYL): 1970s
- Relational: 1970s and early 1980s
- Entity-Relationship: 1970s
- Extended Relational: 1980s
- Semantic: late 1970s and 1980s
- Object-Oriented: late 1980s and early 1990s
- Object-Relational: late 1980s and early 1990s
- Semi-structured (e.g., XML): late 1990s and 2000s

Our conclusion was that the relational model with an extendable type system (i.e., object-relational) has dominated all corners, and nothing else has succeeded in the marketplace. Although many of the non-relational DBMSs that were covered in 2005 still exist today, their vendors have relegated them to legacy maintenance mode and nobody is building new applications on them. This persistence is more of a testament to the “stickiness” of

data rather than the lasting power of these systems. In other words, there still are many IBM IMS databases running today because it is expensive and risky to switch them to use a modern DBMS. But no start-up would willingly choose to build a new application on IMS.

A lot has happened in the world of databases since our 2005 survey. During this time, DBMSs have expanded from their roots in business data processing and are now used for almost every kind of data. This led to the “Big Data” era of the early 2010s and the current trend of integrating machine learning (ML) with DBMS technology.

In this paper, we analyze the last 18 years of data model and query language activity in databases. We structure our commentary into the following areas: (1) **MapReduce Systems**, (2) **Key-value Stores**, (3) **Document Databases**, (4) **Column Family / Wide-Column**, (5) **Text Search Engines**, (6) **Array Databases**, (7) **Vector Databases**, and (8) **Graph Databases**.

We contend that most systems that deviate from SQL or the RM are either already dead or are niche markets at the present time. Many systems that started out rejecting the RM with much fanfare (think NoSQL) have since changed their tune and now expose a SQL-like interface for RM databases. Meanwhile, SQL incorporated the best parts of these failed efforts to expand its support for modern applications and remain relevant. We are not optimistic that future deviations will prove productive.

Although there has not been much change in the RM’s fundamentals, there has been a dramatic change in RM system implementations. In the second part of this paper, we discuss advancements in DBMS architectures that address changing application and hardware landscapes: (1) **Columnar Systems**, (2) **Cloud Databases**, (3) **NewSQL Systems**, (4) **Hardware Accelerators**, and (5) **Blockchain Databases**.

Some of these have caused profound changes to successful DBMS implementations; others are merely trends based on faulty premises or bad ideas. But importantly, none of them have caused a new data model to emerge that supplants the RM.

In summary, the time period since the 2005 survey has left SQL and the RM more dominant than ever. We

# WHAT GOES AROUND COMES AROUND... AND AROUND...

UNDER SUBMISSION (2023)

Key-Value (1990s)

MapReduce (2000s)

Document/JSON (2000s)

Column-family (2000s)

Graph (2000s)

Text Search (1960s)

Array (1990s)

Vector (2020s)



## What Goes Around Comes Around... And Around...

Michael Stonebraker  
Massachusetts Institute of Technology  
stonebraker@csail.mit.edu

Andrew Pavlo  
Carnegie Mellon University  
pavlo@cs.cmu.edu

### ABSTRACT

Two decades ago, one of us co-authored a paper commenting on the previous 40 years of data modelling research and development [98]. That paper demonstrated that SQL and the relational model (RM) reigned supreme for database management systems (DBMSs), and all the efforts to completely replace either the query language or the data model had failed. Instead, SQL absorbed the best ideas from these alternative approaches.

We revisit this issue and argue that little has changed since 2005. Once again there are repeated efforts to replace either SQL or the RM, and none have been successful. Instead, in the last few years the conventional wisdom in industry has swung back to relational DBMSs that use SQL. We suggest that history before they invent models that are likely to fail. The evolution of DBMS implementations and advancements have been in

data rather than the lasting power of these systems. In other words, there still are many IBM IMS databases running today because it is expensive and risky to switch them to use a modern DBMS. But no start-up would willingly choose to build a new application on IMS.

A lot has happened in the world of databases since our 2005 survey. During this time, DBMSs have expanded from their roots in business data processing and are now used for almost every kind of data. This led to the "Big Data" era of the early 2010s and the current trend of integrating machine learning (ML) with DBMS technology.

In this paper, we analyze the last 18 years of data model and query language activity in databases. We

### 1 Introduction

In 2005, one of the authors wrote a chapter for the *Red Book* titled "Around" [98]. That paper catalogued movements since the

- Hierarchical (e.g., IMS)
- Network (e.g., CODAS)
- Relational: 1970s and
- Entity-Relationship: 1980s
- Extended Relational: 1990s
- Semantic: late 1970s
- Object-Oriented: late 1980s
- Object-Relational: late 1990s
- Semi-structured (e.g., XML)

Our conclusion was that the extensible type system (ET) had not penetrated all corners, and it was not the marketplace. Although DBMSs that were covered by vendors had relegated to the sidelines, and nobody is building new persistence is more of a

# WHAT GOES AROUND COMES AROUND... AND AROUND...

UNDER SUBMISSION (2023)

Key-Value (1990s)

MapReduce (2000s)

Document/JSON (2000s)

Column-family (2000s)

- List at least 3 weak points, numbered W1, W2, W3, ...:
- W1. Writing completely obscures the potential value of the work, the potentially interesting opinions and reflections.
  - W2. Writing is completely unrealistic
  - W3. Writing is dishonest
  - W4. Writing is harmful to research community.

# TALK OUTLINE

---

**Key-Value Stores (1990s)**

**MapReduce Systems (2000s)**

**Document/JSON Databases (2000s)**

**Column-family / Wide-Column (2000s)**

**Graph Databases (2000s)**

**Text Search Engines (2000s)**

**Array Databases (1990s)**

**Vector Databases (2020s)**

# TALK OUTLINE

---

Key-Value Stores (1990s)

MapReduce Systems (2000s)

Document/JSON Databases (2000s)

Column-family / Wide-Column (2000s)

Graph Databases (2000s)

Text Search Engines (2000s)

Array Databases (1990s)

Vector Databases (2020s)

**TLDR:** RM+ SQL remains the best approach for most applications.

# KEY-VALUE STORES

Associative array that maps a key to a value.

→ Value is typically an untyped byte array that the DBMS cannot interpret.

(key, value)



## Distributed KV Stores:

- Shared-nothing DBMSs for caching + session data.
- Provide higher/predictable performance instead of a more complex query language and features.



## Embedded Storage Managers:

- Low-level API systems that run in the same address space as a higher-level application.

# KEY-VALUE STORES

---

Some distributed KV stores realized that expressive APIs are important and evolved into document stores.

- If value is opaque, applications must implement more complex logic / types.
- Better to start with a RM DBMS than to contort a KV DBMS to use a more complex data model (e.g., Postgres hstore).

## Discussion:

- Embedded KV storage managers make it easier to create full-featured DBMSs.
- Very few commercial success stories for KV storage managers.

# MAPREDUCE SYSTEMS

Distributed batch-oriented programming and execution model for analyzing large data sets.

Data model decided by user-written functions.

→ **Map**: UDF that performs computation + filtering

→ **Reduce**: Analogous to **GROUP BY** operation.

```
SELECT map() FROM crawl_table GROUP BY reduce();
```



MAPR



## MapReduce Frameworks:

→ Internal implementation at Google (2003).

→ Yahoo! created the open-source version Hadoop (2005).

DOI:10.1145/1629175.1629198

**MapReduce advantages over parallel databases include storage-system independence and fine-grain fault tolerance for large jobs.**

BY JEFFREY DEAN AND SANJAY GHEMAWAT

# MapReduce: A Flexible Data Processing Tool

MAPREDUCE is a programming model for processing and generating large data sets.<sup>4</sup> Users specify a map function that processes a key/value pair to generate a set of intermediate key/value pairs and a reduce function that merges all intermediate values associated with the same intermediate key. We built a system around this programming model in 2003 to simplify construction of the inverted index for handling searches at Google.com. Since then, more than 10,000 distinct programs have been implemented using MapReduce at Google, including algorithms for large-scale graph processing, text processing, machine learning, and statistical machine translation. The Hadoop open source implementation

of MapReduce has been used extensively outside of Google by a number of organizations.<sup>10,11</sup>

To help illustrate the MapReduce programming model, consider the problem of counting the number of occurrences of each word in a large collection of documents. The user would write code like the following pseudocode:

```
map(String key, String value):
    // key: document name
    // value: document contents
    for each word w in value:
        EmitIntermediate(w, "1");

reduce(String key, Iterator values):
    // key: a word
    // values: a list of counts
    int result = 0;
    for each v in values:
        result += ParseInt(v);
    Emit(AsString(result));
```

The map function emits each word plus an associated count of occurrences (just '1' in this simple example). The reduce function sums together all counts emitted for a particular word.

MapReduce automatically parallelizes and executes the program on a large cluster of commodity machines. The runtime system takes care of the details of partitioning the input data, scheduling the program's execution across a set of machines, handling machine failures, and managing required inter-machine communication. MapReduce allows programmers with no experience with parallel and distributed systems to easily utilize the resources of a large distributed system. A typical MapReduce computation processes many terabytes of data on hundreds or thousands of machines. Programmers find the system easy to use, and more than 100,000 MapReduce jobs are executed on Google's clusters every day.

## Compared to Parallel Databases

The query languages built into parallel database systems are also used to

ILLUSTRATION BY NATHAN WATZ

ed prog  
data s  
er-wri  
comput  
UP BY  
table

wor  
on at  
en-so

DOI:10.1145/1629175.1629197

**MapReduce complements DBMSs since databases are not designed for extract-transform-load tasks, a MapReduce specialty.**

BY MICHAEL STONEBRAKER, DANIEL ABADI, DAVID J. DEWITT, SAM MADDEN, ERIK PAULSON, ANDREW PAVLO, AND ALEXANDER RASIN

# MapReduce and Parallel DBMSs: Friends or Foes?

THE MAPREDUCE<sup>7</sup> (MR) PARADIGM has been hailed as a revolutionary new platform for large-scale, massively parallel data access.<sup>16</sup> Some proponents claim the extreme scalability of MR will relegate relational database management systems (DBMS) to the status of legacy technology. At least one enterprise, Facebook, has implemented a large data warehouse system using MR technology rather than a DBMS.<sup>14</sup>

Here, we argue that using MR systems to perform tasks that are best suited for DBMSs yields less than satisfactory results,<sup>17</sup> concluding that MR is more like an extract-transform-load (ETL) system than a

DBMS, as it quickly loads and processes large amounts of data in an ad hoc manner. As such, it complements DBMS technology rather than competes with it. We also discuss the differences in the architectural decisions of MR systems and database systems and provide insight into how the systems should complement one another.

The technology press has been focusing on the revolution of "cloud computing," a paradigm that entails the harnessing of large numbers of processors working in parallel to solve computing problems. In effect, this suggests constructing a data center by lining up a large number of low-end servers, rather than deploying a smaller set of high-end servers. Along with this interest in clusters has come a proliferation of tools for programming them. MR is one such tool, an attractive option to many because it provides a simple model through which users are able to express relatively sophisticated distributed programs.

Given the interest in the MR model both commercially and academically, it is natural to ask whether MR systems should replace parallel database systems. Parallel DBMSs were first available commercially nearly two decades ago, and, today, systems (from about a dozen vendors) are available. As robust, high-performance computing platforms, they provide a high-level programming environment that is inherently parallelizable. Although it might seem that MR and parallel DBMSs are different, it is possible to write almost any parallel-processing task as either a set of database queries or a set of MR jobs.

Our discussions with MR users lead us to conclude that the most common use case for MR is more like an ETL system. As such, it is complementary to DBMSs, not a competing technology, since databases are not designed to be good at ETL tasks. Here, we describe what we believe is the ideal use of MR technology and highlight the different MR and parallel DBMS markets.

ILLUSTRATION BY NATHAN WATZ

# MAPREDUCE SYSTEMS

---

People remembered that procedural query languages are (usually) a bad idea.



MR vendors put SQL engines on top of Hadoop.

Hadoop technology/services market crashed.

Google announced they were dropping MR in 2014.

## Discussion:

- Companies kept HDFS but replaced Hadoop compute layer with relational query engines.
- Aspects of MR frameworks carried into distributed DBMSs (disaggregated compute/storage, shuffle phase).
- Rise of Hadoop alternatives (that eventually added SQL).





# DOCUMENT DATABASES

Represent a database as a collection of document objects that contain a hierarchy of field/value pairs.

- Each document field is identified by a name.
- A field's value is either a scalar type, array of values, or another document.
- Applications do not predefine schema.

```
{<field>: <scalar|[values]|{document}>}
```



## NoSQL Document-oriented Systems:

- Non-standard / procedural query languages
- Defined by what they lack instead of what they provide.

# DOCUMENT DATABASES

Document model is the same as previous models with many of the same problems.

- **Object-Oriented** (1980s)
- **Semi-Structured / XML** (1990s).

Core idea is denormalization ("pre-joining"):

- Avoid object-relational impedance mismatch between application code and DBMS data model.
- Avoid need for joins / multiple queries to retrieve data related to an object (N+1 SELECT Problem).

VERSANT

ObjectStore



MarkLogic

# DOCUMENT DATABASES

---

Almost every major NoSQL DBMS relearned (most) of the lessons from the 1970s:

- SQL APIs are a good idea.
- Schemas + integrity constraints are a good idea.
- Transactions are a good idea.
- Logical/physical data independence is a good idea.

# DOCUMENT DATABASES

Almost every major NoSQL DBMS relearned (most) of the lessons from the 1970s:

- SQL APIs are a good idea.
- Schemas + integrity constraints are a good idea.
- Transactions are a good idea.
- Logical/physical data independence is a good idea.



→ PartiQL



cassandra → CQL



→ AQL



Couchbase → SQL++

# DOCUMENT DATABASES

Almost every major NoSQL  
(most) of the lessons from the

- SQL APIs are a good idea.
- Schemas + integrity constraints
- Transactions are a good idea.
- Logical/physical data independence



# DOCUMENT DATABASES

Almost every major NoSQL DBMS relearned (most) of the lessons from the 1970s:

- SQL APIs are a good idea.
- Schemas + integrity constraints are a good idea.
- Transactions are a good idea.
- Logical/physical data independence is a good idea.

## Discussion:

- SQL:2016 introduced JSON types + operators.
- The intellectual distance between relational+JSON DBMSs and document+SQL DBMSs has shrunk.



→ PartiQL



→ CQL



→ AQL



→ SQL++

# COLUMN-FAMILY / WIDE-COLUMN

Reduction of the document data model that only supports one level of nesting.

- A record's value can only be a scalar or an array of scalars.
- Deficiencies are the same as the document model.

```
{<field>: <scalar|[values]>}
```



## Column-Family Systems:

- First implementation was Google's BigTable (2004)
- Copied by several Internet start-ups.

# GRAPH DATABASES

Direct multigraph structure that supports key/value labels for nodes and edges.

→ Property Graph vs. Resource Description Framework (RDF)

**Node** (id, {key: value}\*)

**Edge** (node\_id<sub>1</sub>, node\_id<sub>2</sub>, {key: value}\*)



## Property Graph DBMSs:

→ Provide graph-oriented traversal APIs.

→ Inefficient schemaless storage.



# GRAPH DATABASES

---

Graph model is the same as the **network model** from CODASYL (1970s) with same issues.

Advancements in algorithms and systems will diminish the perceived advantage of specialized graph DBMSs.

- Worst-case Optimal Joins
- Vectorized Query Execution
- Factorized Query Processing

## Discussion:

- SQL:2023 introduced SQL/PGQ (based on Neo4j's Cypher)  
Subset of the emerging GQL standard.
- Studies show that RM DBMSs outperform graph DBMSs.

# GRAPH DATABASES

Graph model is the same as the new CODASYL (1970s) with same issues

Advancements in algorithms and the perceived advantage of specializations

- Worst-case Optimal Joins
- Vectorized Query Execution
- Factorized Query Processing

## Discussion:

- SQL:2023 introduced SQL/PGQ (but not a subset of the emerging GQL standard)
- Studies show that RM DBMSs outperform

### DuckPGQ: Efficient Property Graph Queries in an analytical RDBMS

Daniel ten Wolde  
CWI  
The Netherlands  
djt@cw.nl

Tavneet Singh  
CWI  
The Netherlands  
tavneet.singh@cw.nl

Gábor Szárnyas  
CWI  
The Netherlands  
gabor.szarnyas@cw.nl

Peter Boncz  
CWI  
The Netherlands  
boncz@cw.nl

#### ABSTRACT

In the past decade, property graph databases have emerged as a growing niche in data management. Many native graph systems and query languages have been created, but the functionality and performance still leave much room for improvement. The upcoming SQL:2023 will introduce the Property Graph Queries (SQL/PGQ) sub-language, giving relational systems the opportunity to standardize graph queries, and provide mature graph query functionality.

We argue that (i) competent graph data systems must build on all technology that makes up a state-of-the-art relational system, (ii) the graph use case requires the addition to that of a many-source/destination path-finding algorithm and compact graph representations, and (iii) incites research in practical worst-case optimal joins and factorized query processing techniques.

We outline our design of DuckPGQ that follows this recipe, by adding efficient SQL/PGQ support to the popular open-source “embeddable analytics” relational database system DuckDB, also originally developed at CWI. Our design aims at minimizing technical debt using an approach that relies on efficient vectorized UDFs. We benchmark DuckPGQ showing encouraging performance and scalability on large graph data sets, but also reinforcing the need for future research under (iii).

#### 1 INTRODUCTION

Graph Database systems have emerged as a growing niche in data management, with many property graph systems [7] such as Neo4j, TigerGraph, Dgraph, Titan and AWS Neptune becoming available, all using different query languages (i.e., Cypher, GSQL, GraphQL, Gremlin, SPARQL [2]). Property graphs are directed graphs consisting of vertex and edge elements, where elements may have labels and associated key/value properties. Property graph systems are quite young, and performance of analytical queries on large graphs has been observed to be significantly lower than relational database systems, on graph queries that can also be formulated as SQL [16].

In RDBMS designs, there have been significant performance improvements in the past decade, with analytical systems such as Snowflake and Databricks adopting principles like skipable columnar storage with lightweight compression [24] (also popular in open-source formats such as Parquet and ORC), efficient load-balanced multi-core parallelism using “morsel-driven” scheduling [15] and efficient query execution techniques [14], either using

vectorized query execution or Just-In-Time low-level compilation of queries into executable programs.

The upcoming SQL:2023 introduces the SQL/PGQ (Property Graph Queries) sub-language [8], which allows (1) to define graph views over relational tables and (2) to formulate graph pattern matching and path-finding operations using a SQL syntax. These features narrow the functionality gap between RDBMSs and native graph systems, and unify the feature space with a common graph query sub-language, as PGQ is also a subset of the upcoming ISO Graph Query Language GQL [8] that native graph systems intend to adopt. GQL will add graph updates, querying multiple graphs and queries that return a graph result, rather than a binding table.

**SQL/PGQ by example.** If we have relational tables `student` and `college` and connecting tables `know` and `enrol`, we can define a property graph `ag` consisting of `know` vertices connected to each other by edges with label `know` and to `college` vertices via `student` edges:<sup>1</sup>

```
CREATE PROPERTY GRAPH PG
  VERTEX TABLES(
    Student PROPERTIES(id,name,birthdate) LABEL Person,
    College PROPERTIES(id,college) LABEL Student,
  )
  EDGE TABLES(
    know SOURCE Person KEY(id) DESTINATION Person KEY(id)
    PROPERTIES(crateDate,wagCount),
    enrol SOURCE Student KEY(id) DESTINATION College KEY(id)
    PROPERTIES(classYear)
```

In the below `SELECT` query the `MATCH` will bind variable `a` to all vertices that satisfy a label-test `Person` and have property `name = 'Ana'`. The comma separating the two pattern expressions implies a conjunction<sup>2</sup> with matching variable bindings: it requires `a` to also have an edge labeled `student` towards a `College` `c`:

```
SELECT study.college, study.pid FROM GRAPH_TABLE (PG,
  MATCH (a:Person WHERE a.name='Ana'),
  (a)-[:student]->(c:College))
  COLUMNS (c.college, ELEMENT_ID(a) AS pid) study
```

The `MATCH` clause produces a conceptual binding table with each row holding denoted bindings and one column for each variable. These bindings denote elements (e.g., a vertex or edge); the `COLUMNS` clause retrieves scalar values from the rows. The example retrieves the property `c.college` and the implicit element identifier<sup>3</sup> of `a`, as the columns of a temporary `GRAPH_TABLE` named `study` in the `FROM` clause.

<sup>1</sup>The table name is the default label. DuckPGQ allows an additional `LABEL` list of max length 64, and a `BIGINT LABEL FROM col` specifier column. Elements only have a label from the list if their corresponding list is set. This allows e.g., to express class membership with inheritance in labels. DuckPGQ will not support having the same label in multiple tables, as element patterns must always bind to a single table.

<sup>2</sup>Inside path expressions, the `|` will `UNION` pattern bindings, and `|>` stands for `JOIN`. `ELEMENT_ID()` is implementation-dependent; in DuckPGQ it returns a rowid.

This paper is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0). Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution, provided that you attribute the original work to the authors and CIDR 2023, 15th Annual Conference on Innovative Data System Research (CIDR '23), January 8-11, 2023, Amsterdam, The Netherlands.

# GRAPH DATABASES

Graph model is the same as the new  
CODASYL (1970s) with same idea

\* 10 points by apavlo on Dec 30, 2021 | parent | next [-]

> Databases in 2030: SQL DB finally succumbs to Graph DB as #1  
Graph databases will *not* overtake relational databases in 2030 by marketshare.  
Bookmark this comment. Reach out to me in 2030. If I'm wrong, I will replace my official CMU photo with one of me wearing a shirt that says "Graph Databases Are #1". I will use that photo until I retire, get fired, or a former student stabs me.

## Discussion:

- SQL:2023 introduced SQL/PGQ (b)
- Subset of the emerging GQL standard
- Studies show that RM DBMSs outp

### DuckPGQ: Efficient Property Graph Queries in an analytical RDBMS

Daniel ten Wolde  
CWI  
The Netherlands  
djt@cw.nl

Tavneet Singh  
CWI  
The Netherlands  
tavneet.singh@cw.nl

Gábor Szárnyas  
CWI  
The Netherlands  
gabor.szarnyas@cw.nl

Peter Boncz  
CWI  
The Netherlands  
boncz@cw.nl

#### ABSTRACT

In the past decade, graph databases have

been or Just-In-Time low-level compilation

programs.

SQL/PGQ [8], which allows (1) to define graph

edges and (2) to formulate graph pattern

operations using a SQL syntax. These

the feature space with a common graph

PGQ is also a subset of the upcoming ISO

QL [8] that native graph systems intend

graph updates, querying multiple graphs

graph result, rather than a binding table.

We have relational tables student and

know and enrol, we can define a prop-

erties vertices connected to each other

to college vertices via studentat edges.<sup>1</sup>

me, birthdate) LABEL Person,

(age))

(8) DESTINATION Person KEY(id)

(date, wsgCount),

(id) DESTINATION College KEY(id)

(par)

(LABEL studentat)

the MATCH will bind variable a to all

Person and have property name

two pattern expressions implies a

variable bindings: it requires a to also

towards a College c:

```
SELECT study.college, study.pid FROM GRAPH_TABLE (PG,
MATCH (a:Person WHERE a.name='Jon'),
(a)-[:studentat]->(c:College),
COLUMNS (c.college, ELEMENT_ID(a AS pid)) study;
```

The MATCH clause produces a conceptual binding table with each row holding matched bindings and one column for each variable. These bindings denote elements (e.g., a vertex or edge); the COLUMNS clause retrieves scalar values from those. The example retrieves the property c.college and the implicit element identifier<sup>2</sup> of a, as the columns of a temporary GRAPH\_TABLE named study in the FROM clause.

<sup>1</sup>The table name is the default label. DuckPGQ allows an additional LABEL list of max length 64, and a RIGHT LABEL FROM col specifier column. Elements only have a label from the list if their corresponding list is set. This allows e.g., to express class membership with inheritance in labels. DuckPGQ will not support having the same label in multiple tables, as element patterns must always bind to a single table.

<sup>2</sup>Inside pattern expressions, the j will UNION pattern bindings, and i stands for UNION.

<sup>3</sup>ELEMENT\_ID() is implementation-dependent: in DuckPGQ it returns a rowid.

performance of analytical queries on large graphs systems, on graph queries that can also be formulated as SQL [16].

In RDBMS designs, there have been significant performance improvements in the past decade, with analytical systems such as Snowflake and Databricks adopting principles like skipable columnar storage with lightweight compression [24] (also popular in open-source formats such as Parquet and ORC), efficient load-balanced multi-core parallelism using "morsel-driven" scheduling [15] and efficient query execution techniques [14], either using

This paper is published under the Creative Commons Attribution 4.0 International License. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution, provided that they attribute the original work to the authors and CIDR 2023, 15th Annual Conference on Innovative Data Systems Research (CIDR '23), January 8-11, 2023, Amsterdam, The Netherlands.

# TEXT SEARCH ENGINES

---

Systems that extract structure (e.g., meta-data, indexes) from text data and support queries over that content.

- Tokenize documents into "bag of words" and then build inverted indexes over those tokens.
- No data model because text data is inherently unstructured.

Core ideas pioneered by Cornell's SMART (1965).

 elasticsearch

 splunk>

 vespa

 Solr

## Text Search Engines:

- Quickly parse, index, and store large documents.
- Built-in support for noise/salient words + synonyms.

# TEXT SEARCH ENGINES

---

Leading RM DBMSs include full-text search indexes but their adoption is stymied by non-model reasons.

- Non-standard SQL operations / syntax.
- Text data is large but not high importance. DBMS storage is always more expensive than generic storage.

## **Discussion:**

- Maintaining a separate text search DBMS should be unnecessary but lots of people still do it.
- All DBMS vendors are augmenting inverted-index text search with vector-based similarity search...

# ARRAY DATABASES

Collection of data where each element is identifiable by one or more dimension offsets.

- Vectors (1D), Matrices (2D), Tensors (+3D)
- Dimensions do not have to align with integer grids.

$(\text{dimension}_1, \text{dimension}_2, \dots [\text{values}])$

**rasdaman**  
raster data management

**kx**

**SciDB**

**[tile]DB**

## Array DBMSs:

- Specialized storage managers and execution engines.
- Sparse vs. Dense Arrays



A female Jeopardy! contestant with long dark hair, wearing a dark green ribbed turtleneck, stands behind a wooden podium. She is holding a blue buzzer in her right hand. The podium features a blue screen displaying the amount "\$7,400" in white. Above the screen is a control panel with several red buttons. The background consists of a blue curtain and a curved blue wall. To the right, a portion of a wooden ladder is visible.

**\$7,400**

# ARRAY DATABASES

---

Supporting arrays as first-class data types violates the original RM vision. But this is a good example of RM evolving to meet the needs of applications.

## **Discussion:**

- SQL:2023 added multi-dimensional arrays (SQL/MDA).
- Array data access patterns do not follow row-oriented or columnar patterns. Likely requires new execution engine.



# VECTOR DATABASES

Document DBMSs with specialized indexes for (approximate) similarity search on 1D arrays.  
→ Vectors represent embedding of corresponding object.

```
{vector: [values],  
  metadata: {key: value}*}
```

 Pinecone

 Weaviate

 milvus


 drant

## Vector DBMSs:

- Accelerate approximate nearest neighbor search via indexes.
- Not meant to be primary / database-of-record storage.


# VECTOR DATABASES

id	name	year	lyrics
Id1	<a href="#">Enter the Wu-Tang</a>	1993	<text>
Id2	<a href="#">Run the Jewels 2</a>	2015	<text>
Id3	<a href="#">Liquid Swords</a>	1995	<text>
Id4	<a href="#">We Got It from Here</a>	2016	<text>



## *Embeddings*

Id1 → [0.32, 0.78, 0.30, ...]  
Id2 → [0.99, 0.19, 0.81, ...]  
Id3 → [0.01, 0.18, 0.85, ...]  
Id4 → [0.19, 0.82, 0.24, ...]  
⋮



HNSW, IVFFlat  
Meta Faiss, Spotify Annoy,  
Microsoft DiskANN

# VECTOR DATABASES

id	name	year	lyrics
Id1	<a href="#">Enter the Wu-Tang</a>	1993	<text>
Id2	<a href="#">Run the Jewels 2</a>	2015	<text>
Id3	<a href="#">Liquid Swords</a>	1995	<text>
Id4	<a href="#">We Got It from Here</a>	2016	<text>

## Query

Find albums with lyrics about  
**running from the police**



OpenAI



Hugging Face

Transformer

## Embeddings

Id1 → [0.32, 0.78, 0.30, ...]

Id2 → [0.99, 0.19, 0.81, ...]

Id3 → [0.01, 0.18, 0.85, ...]

Id4 → [0.19, 0.82, 0.24, ...]

⋮

[0.02, 0.10, 0.24, ...]

*Ranked List of Ids*

**Vector  
Index**

HNSW, IVFFlat

Meta Faiss, Spotify Annoy,

Microsoft DiskANN

# VECTOR DATABASES

id	name	year	lyrics
Id1	<a href="#">Enter the Wu-Tang</a>	1993	<text>
Id2	<a href="#">Run the Jewels 2</a>	2015	<text>
Id3	<a href="#">Liquid Swords</a>	1995	<text>
Id4	<a href="#">We Got It from Here</a>	2016	<text>



OpenAI



Hugging Face

**Transformer**

## *Embeddings*

Id1 → [0.32, 0.78, 0.30, ...]

Id2 → [0.99, 0.19, 0.81, ...]

Id3 → [0.01, 0.18, 0.85, ...]

Id4 → [0.19, 0.82, 0.24, ...]

⋮



**Vector  
Index**

HNSW, IVFFlat  
Meta Faiss, Spotify Annoy,  
Microsoft DiskANN

# VECTOR DATABASES

id	name	year	lyrics
Id1	<a href="#">Enter the Wu-Tang</a>	1993	<text>
Id2	<a href="#">Run the Jewels 2</a>	2015	<text>
Id3	<a href="#">Liquid Swords</a>	1995	<text>
Id4	<a href="#">We Got It from Here</a>	2016	<text>



OpenAI



Hugging Face


Transformer

## Embeddings

Id1 → [0.32, 0.78, 0.30, ...]

Id2 → [0.99, 0.19, 0.81, ...]

Id3 → [0.01, 0.18, 0.85, ...]

Id4 → [0.19, 0.82, 0.24, ...]

⋮

## Query

Find albums with lyrics about  
**running from the police**  
 and released after **2005**

[0.02, 0.10, 0.24, ...]

year &gt; 2005


Vector Index

HNSW, IVFFlat

Meta Faiss, Spotify Annoy,

Microsoft DiskANN

# VECTOR DATABASES

id	name	year	lyrics
Id1	<a href="#">Enter the Wu-Tang</a>	1993	<text>
Id2	<a href="#">Run the Jewels 2</a>	2015	<text>
Id3	<a href="#">Liquid Swords</a>	1995	<text>
Id4	<a href="#">We Got It from Here</a>	2016	<text>

## Query

Find albums with lyrics about  
**running from the police**  
and released after **2005**



OpenAI



Hugging Face

Transformer

## Embeddings

Id1 → [0.32, 0.78, 0.30, ...]

Id2 → [0.99, 0.19, 0.81, ...]

Id3 → [0.01, 0.18, 0.85, ...]

Id4 → [0.19, 0.82, 0.24, ...]

⋮

[0.02, 0.10, 0.24, ...]

year &gt; 2005

  
Vector  
IndexHNSW, IVFFlatMeta Faiss, Spotify Annoy,Microsoft DiskANN

# VECTOR DATABASES

---

The vector model is not a substantial deviation from existing models that requires new DBMS architectures. Vector DBMSs offer better integration with AI tooling ecosystem (e.g., OpenAI, LangChain).

## **Discussion:**

- Every major DBMS will provide native vector index support in the near future.
- The time from "ChatGPT Buzz" (Q4'22) to existing DBMSs announcing support for vectors (Q3'23) is telling.

# VECTOR DATABASES

Substantial deviation from architectures.

announc

**Timescale**

## How We Made PostgreSQL a Better Vector Database

25 Sep 2023  
24 min read

AI

Contributors  
Arthar Sewrathan  
Matvey Arye  
Samuel Gichohi  
Maheedhar PV

Share  
Twitter LinkedIn YouTube

Blog Categories  
All posts  
AI  
Announcements  
Cloud  
Developer Q&A  
Engineering  
General  
Grafana  
Observability  
PostgreSQL  
Product Updates

Dataset: 1 Million OpenAI vector embeddings

Architecture	Search Time (ms)
Timescale vector (DiskANN)	1252.20
pgvector HNSW	898.33
Weaviate	364.25
pg_embedding HNSW	270.76

Raouf Chebri  
Developer Advocate

**ClickHouse**

**SingleStore Blog**

**ROCKSET**

**DATASTAX**

**Google Cloud**

**NEON**

**Introdu Search: Cassandra DB Developer Build Ge Applica**

June 7, 2023

FILED IN: TECH

In the age of AI, Apache Cassandra, a distributed database, can provide high performance for many AI applications. The introduction of general capabilities are needed.

**AI Search to run with**

April 18, 2023  
John Soltau

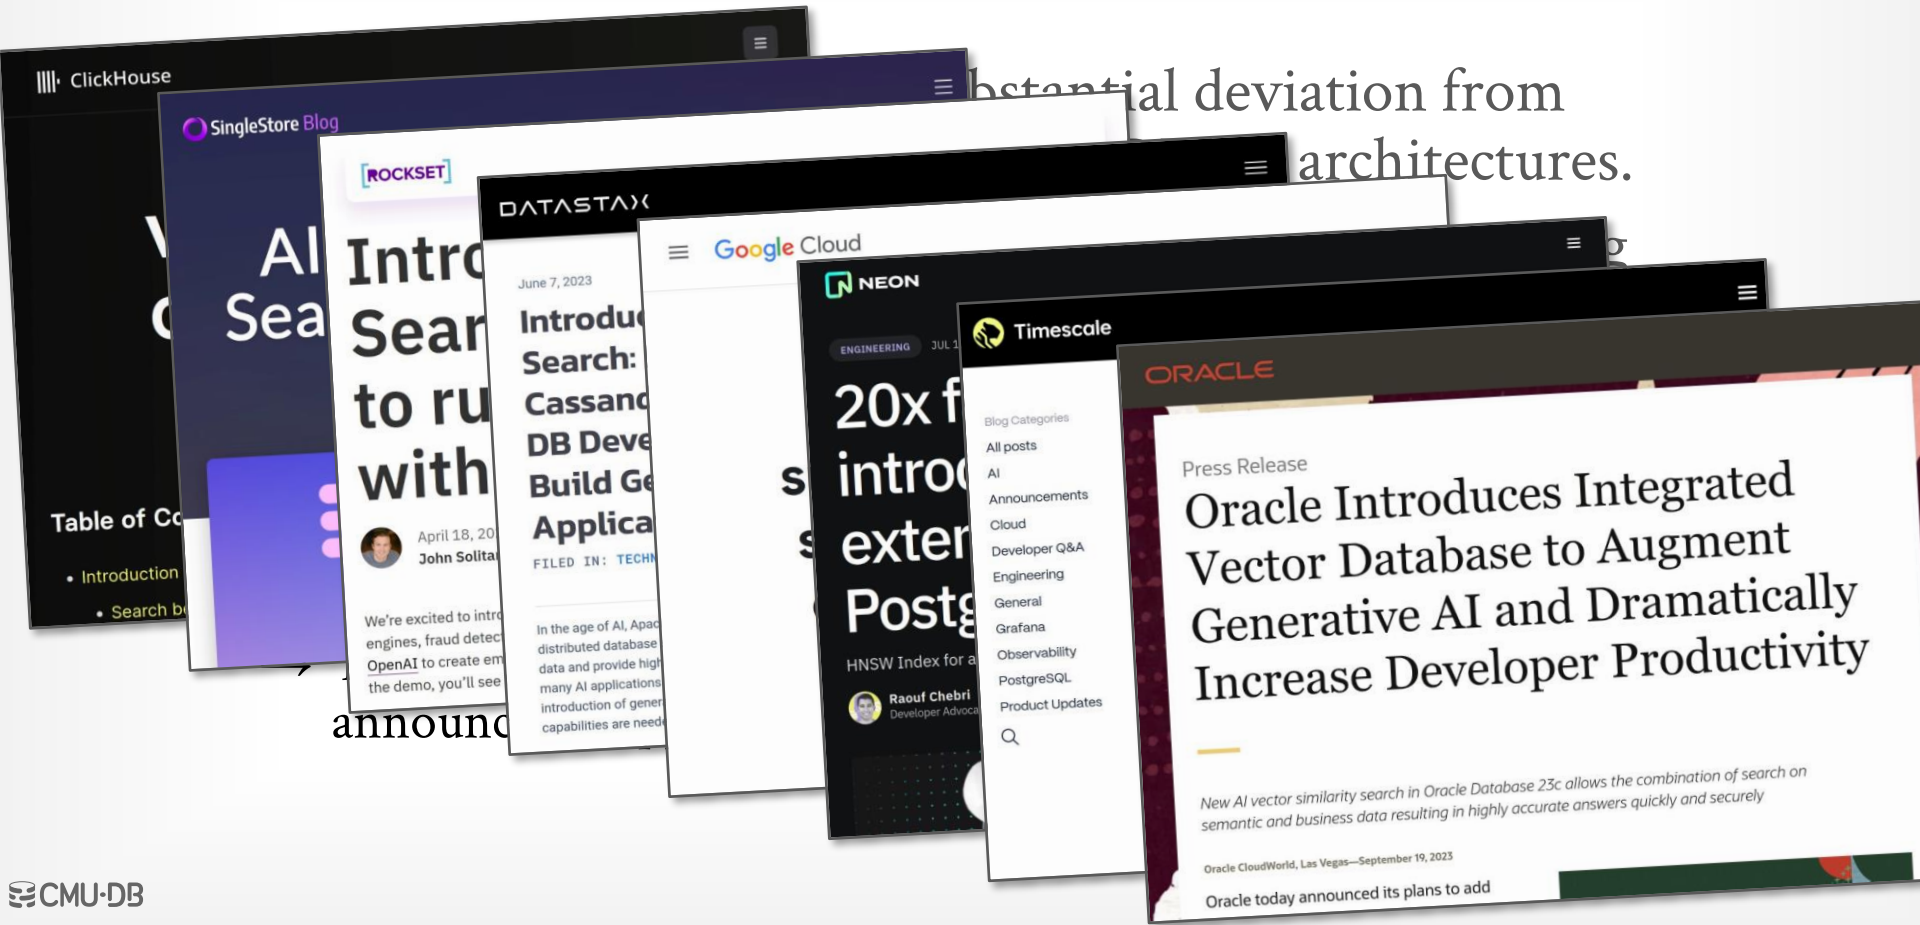
We're excited to introduce a new search engine, fraud detection, and OpenAI to create embeddings. In the demo, you'll see

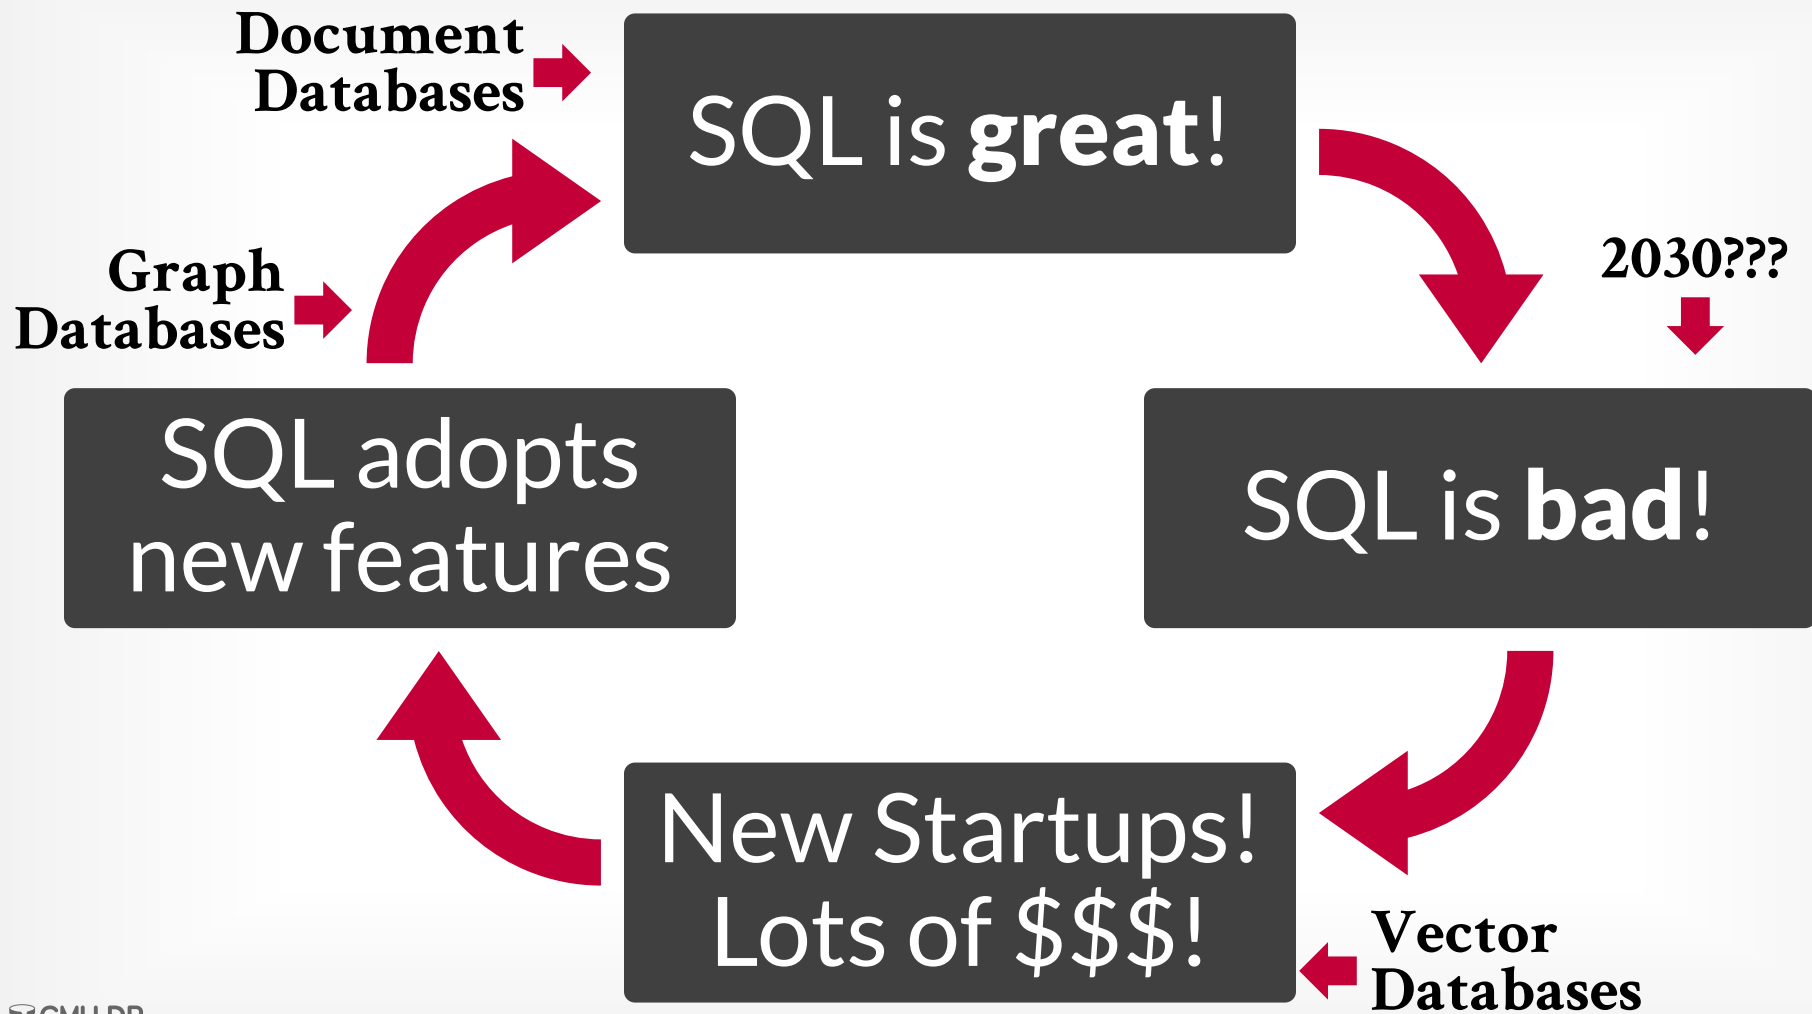
**Table of Contents**

- Introduction
- Search capabilities



# VECTOR DATABASES





# RELATIONAL IS NOT PERFECT

---

Many non-relational DBMSs provide a better "out-of-the-box" experience than relational DBMSs.

→ Pandas / Jupyter notebooks are still more popular.

Relational DBMS developers should strive to make their systems easier to use and adaptive.

→ Cloud DBaaS hide much of the provisioning / configuration for high availability and durability.

**AI/ML helps tuning + optimization.**



# PARTING THOUGHTS

---

People will continue to make the same mistakes in future DBMS projects.

The demarcation lines of DBMS categories will continue to blur over time as specialized systems expand the scope of their domains.

The relational model and declarative query languages promote better data engineering.



# END

**Email:** pavlo@cs.cmu.edu

**Twitter:** @andy\_pavlo

**Mastodon:** @andy\_pavlo@discuss.systems

**LinkedIn:** linkedin.com/in/andy-pavlo

<https://cmudb.io/pgconf23>