SAPinsider SUMMITSERIES

Technology & Innovation

October 24-25, 2024 | Dallas

Developing Privacy Techniques for Secure Vector Databases and LLMs

Jide Ogunjobi
CEO, Context Data

About Me

Jide Ogunjobi

Over 15 years experience developing cutting edge Data, ML & AI platforms across global organizations and startups

AlphaSense	Director of Data Engineering & Architecture
move	Chief Technology Officer
publicis sapient	Data Engineering Manager
GRUBHUB	Lead Data Engineer
citi	Global Data Acquisition Lead
JPMorgan Chase & Co.	Data Solutions Architect

What We'll Cover

- Importance of Privacy in Vector Databases and LLMs
- Encryption Techniques for Vector Databases
  - Homomorphic Encryption
  - Secure Multi-Party Computation (SMPC)
- Privacy-Aware Query Processing Systems
- Access Control and Data Governance Strategies

Why we are talking about Vector Databases

- The Rise of Large Language Models (LLMs)
  - LLMs like GPT-4 have revolutionized natural language processing.
  - Applications span chatbots, content generation, translation, and beyond.

- Role of Vector Databases
  - Store high-dimensional embeddings for efficient similarity search.
  - Crucial for semantic retrieval and recommendation systems.

- Emergence of Privacy Concerns
  - Increasing handling of sensitive personal and business data.
  - Heightened risk of data breaches and misuse.

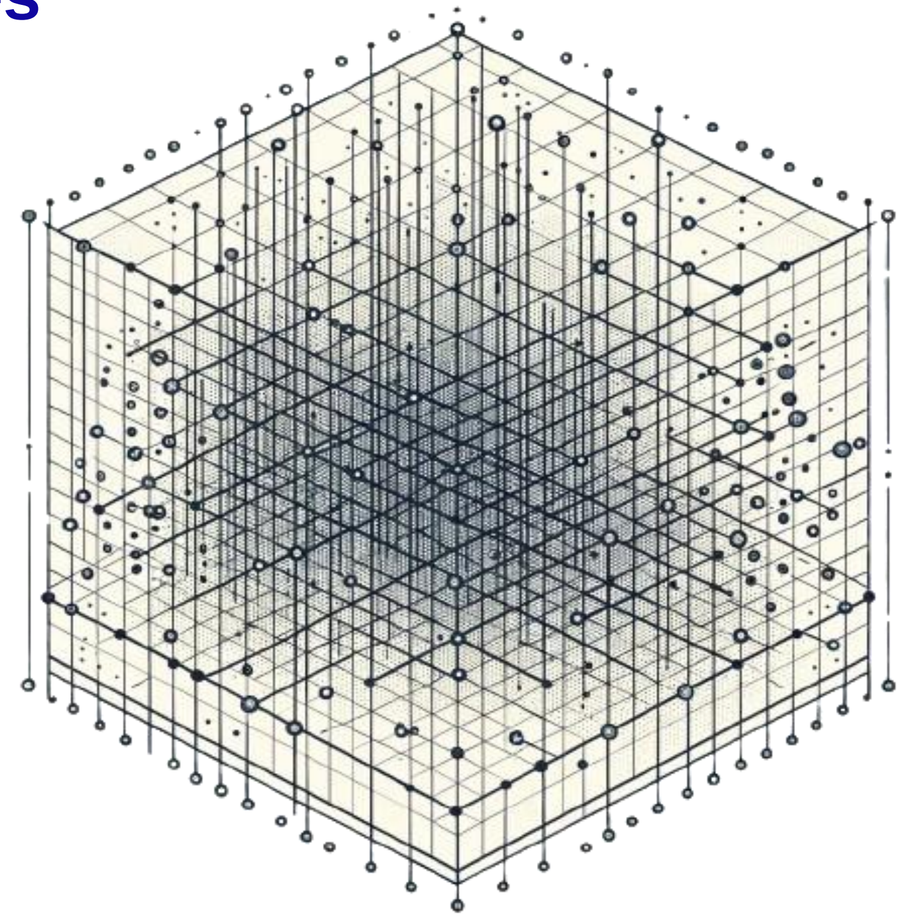

Importance of Privacy in Vector Databases and LLMs

- Sensitive Nature of Data
  - Includes personal information and proprietary business data.
  - Increasing sensitivity of data handled by LLMs
- Regulatory Compliance
  - GDPR, CCPA, and other global data protection laws demand stringent privacy measures.
- Trust and User Confidence
  - Users expect and demand secure handling of their data.
  - Trust is essential for widespread adoption of AI technologies.

Challenges in Ensuring Privacy

- High-Dimensional Data Complexity
  - Traditional encryption methods may not scale well with high-dimensional vectors.
- Performance Overhead
  - Privacy techniques can introduce latency, affecting user experience.
- Balancing Accessibility and Security
  - Data must remain usable for AI functionalities while being secured.

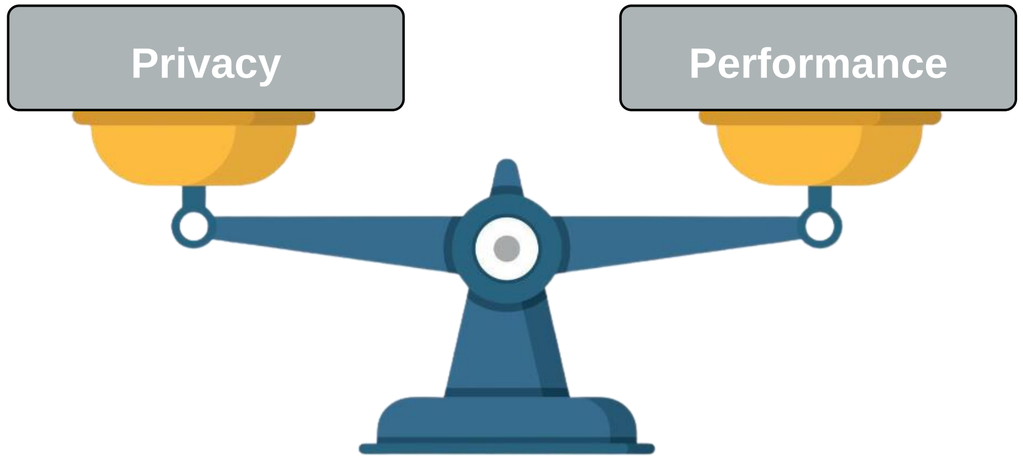

Encryption Techniques for Vector Databases

- Overview of Encryption in Databases
  - Encryption at rest, in transit, and in use.
- Homomorphic Encryption
  - Enables computations on encrypted data without decryption.
  - Maintains data confidentiality during processing.
- Secure Multi-Party Computation (SMPC)
  - Allows parties to compute a function over their inputs while keeping them private.

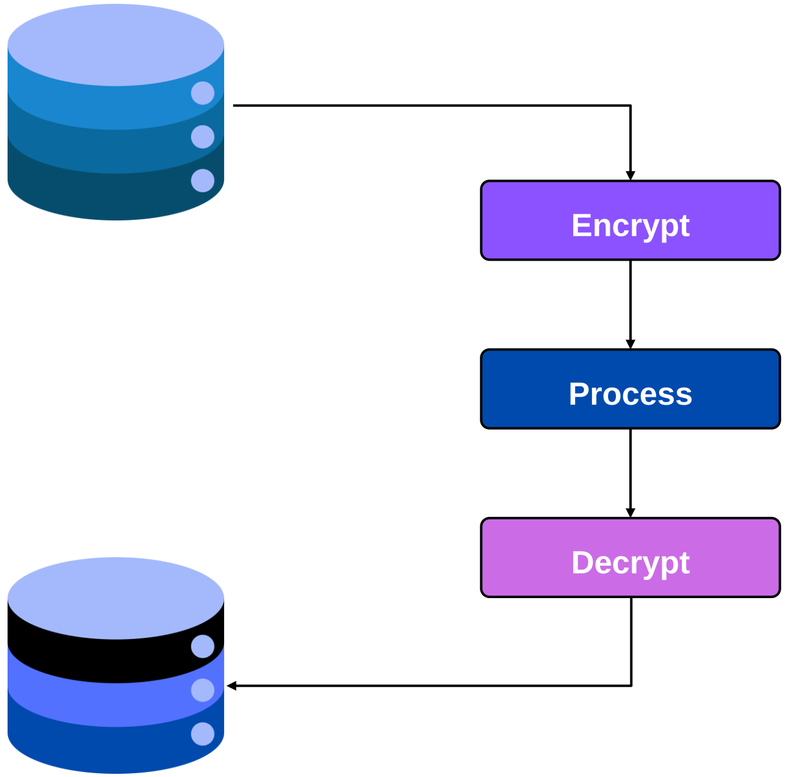

Homomorphic Encryption

Encryption that allows you to perform calculations on encrypted data without needing to decrypt it first.

- Applications in Vector Databases
  - Secure similarity searches over encrypted vectors.
  - Privacy-preserving recommendations.
- Challenges
  - Computational overhead and complex key management.
  - Larger ciphertext sizes impacting storage and transmission.

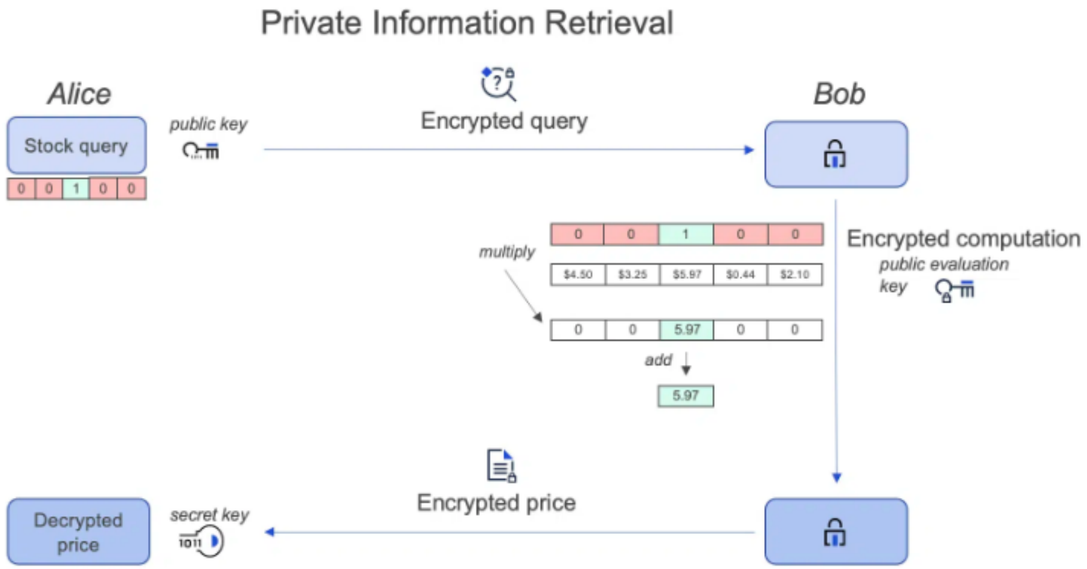

Source

Homomorphic Encryption

Partially Homomorphic Encryption

Partially Homomorphic Encryption schemes support unlimited computations of a single mathematical operation

- Either addition or multiplication but not both
- Generally more efficient and faster due to simpler algorithms

Use Cases

1). Financial Summation: Calculating the total salary expenses without exposing individual salaries

2). Polling and voting: Situations where there is a need to aggregate responses without revealing participants

Somewhat Homomorphic Encryption

Somewhat Homomorphic Encryption schemes support both addition and multiplication operations but only a limited number of times or for computations of limited complexity.

- Supports a predetermined number of homomorphic additions and multiplications.

Use Cases

1). Statistical Analysis: Researcher computing basic statistics on encrypted datasets

2). Secure Queries: Executing simple database queries over encrypted data

Fully Homomorphic Encryption

Fully Homomorphic Encryption schemes allow an unlimited number of both addition and multiplication operations on encrypted data

- No restrictions on the types or number of operations
- Computationally intensive but has been improving with ongoing research

Use Cases

1). Machine Learning: Training or deploying machine learning models on encrypted datasets

2). Secure Multi-Party Computation: Multiple party inputs while keeping those inputs private.

Homomorphic Encryption: Real World Application

Genomic data contains highly sensitive personal information that can reveal predispositions to certain diseases or conditions. Sharing this data for research purposes poses significant privacy risks.

Application of Homomorphic Encryption:

- Secure Computations: Researchers want to perform computations on genomic data to identify disease markers without accessing the raw genetic information.

- Homomorphic Encryption Use: The genomic data is encrypted using a homomorphic encryption scheme before being shared with researchers. Researchers then perform computations directly on the encrypted data.

- Results Decryption: The output of these computations is an encrypted result, which is sent back to the data owner. The data owner decrypts the result to obtain the meaningful insights.

Example

- Microsoft's Simple Encrypted Arithmetic Library (SEAL)

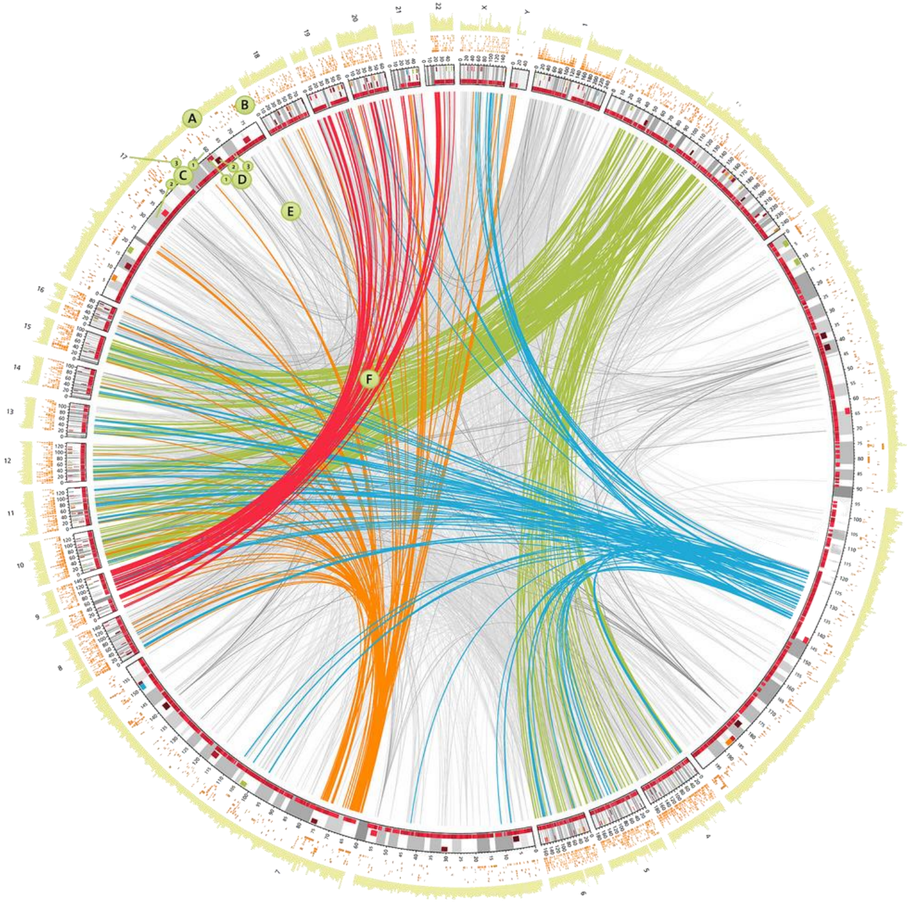

Secure Multi-Party Computation (SMPC)

Cryptographic protocol that enables multiple parties to jointly compute a function over their inputs while keeping those inputs completely private

- Use Cases in LLM Applications
  - Collaborative model training without sharing raw data.
  - Joint analysis between organizations while preserving confidentiality.

- Implementations
  - Secret Sharing: Data is split into shares distributed among parties.
  - Oblivious Transfer: Securely transferring data without revealing the sender's choice.

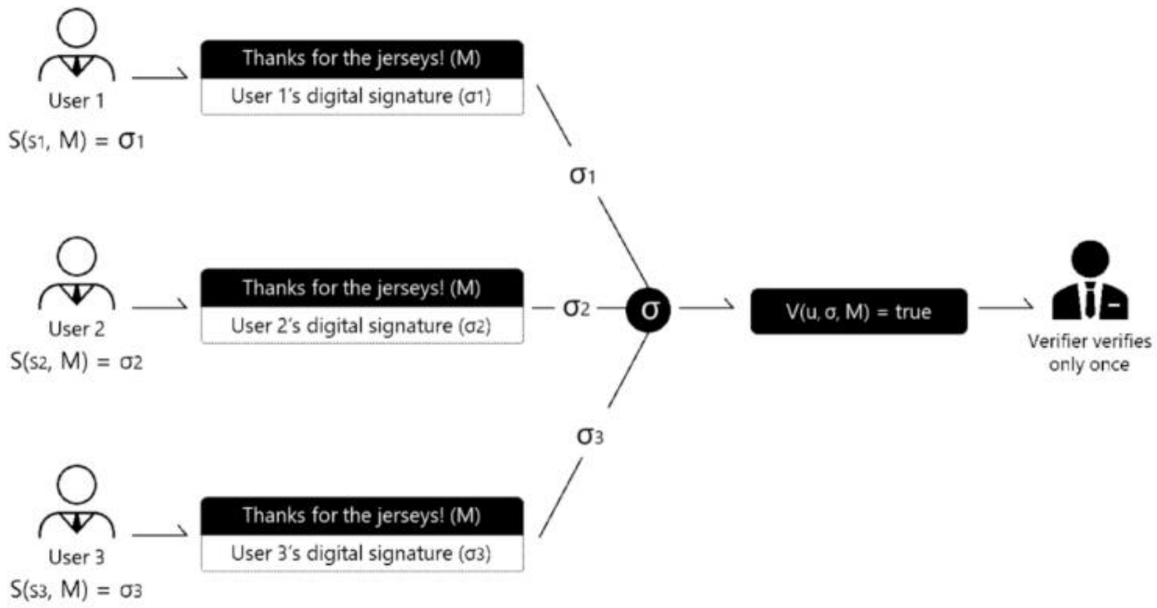

Source

How Secure Multi-Party Computation (SMPC) Works

How SMPC Achieves This

- **Data Masking:**
  - Participants hide their data using techniques that keep the actual values secret.
- **Secure Computation Protocols:**
  - Mathematical methods allow computations on the hidden data.
  - Ensures that combining the hidden data still produces the correct result.
- **No Trusted Third Party Needed:**
  - The system doesn't rely on an outside person or entity to keep data safe.
  - Security is maintained even if some participants are untrustworthy.

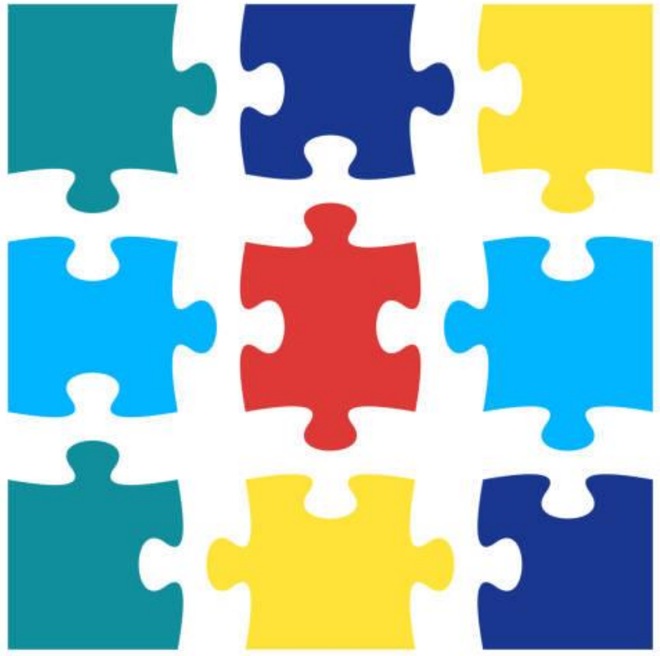

SMPC: Real World Application

Collaborative Fraud Detection in Financial Services

Financial institutions aim to detect fraudulent activities by analyzing transaction data. Sharing detailed transaction data between institutions is restricted due to privacy laws and competitive concerns.

Application of SMPC:

- Joint Fraud Detection Models:
  - Objective: Collaborate on fraud detection algorithms without exposing sensitive client data.
  - Method: Banks use SMPC to jointly compute over their encrypted transaction datasets.

Example

- Financial Crime Intelligence and Insights (FCi2)

Privacy-Aware Query Processing Systems

Systems that allow users to retrieve information from databases while ensuring that sensitive or personal data remains protected.

These systems are designed to provide useful answers to queries without exposing confidential details that should stay private.

Techniques

- Query Obfuscation Techniques
- Secure Aggregation

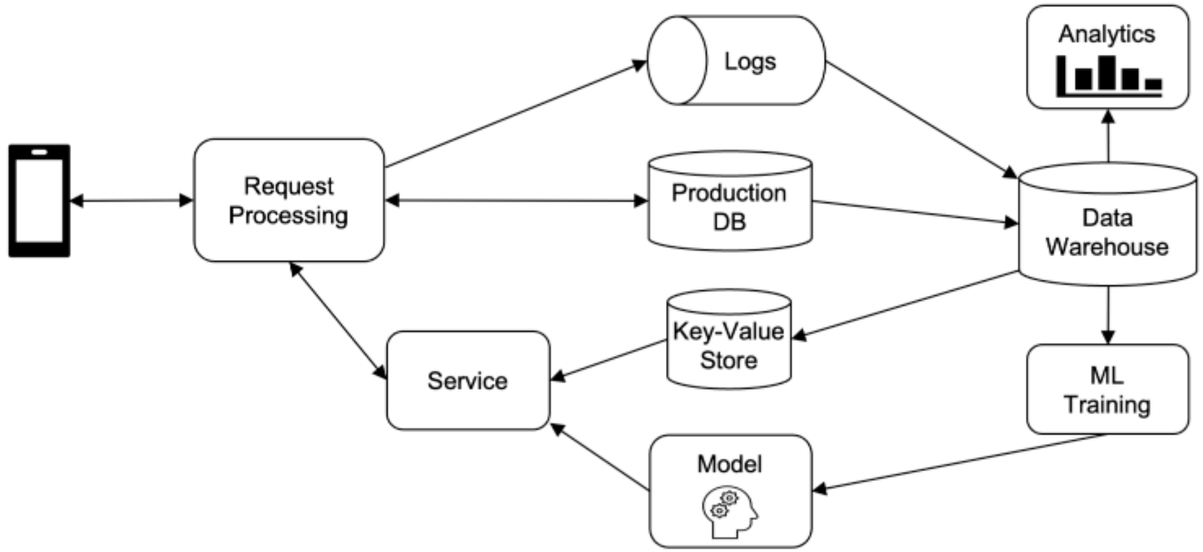

Source

How Privacy-Aware Query Processing Works

User Submits a Query:

- A person asks the system for information from a database.

Privacy Engine Evaluates the Query:

- The system checks if answering the query as-is would expose sensitive data.

Adjusting the Query or Data:

- If there's a risk, the system modifies the query or the data results to protect privacy.
- This could involve summarizing data, omitting certain details, or adding slight inaccuracies that don't affect the overall usefulness.

Delivering the Safe Results:

- The user receives the information they need without any confidential data being leaked.

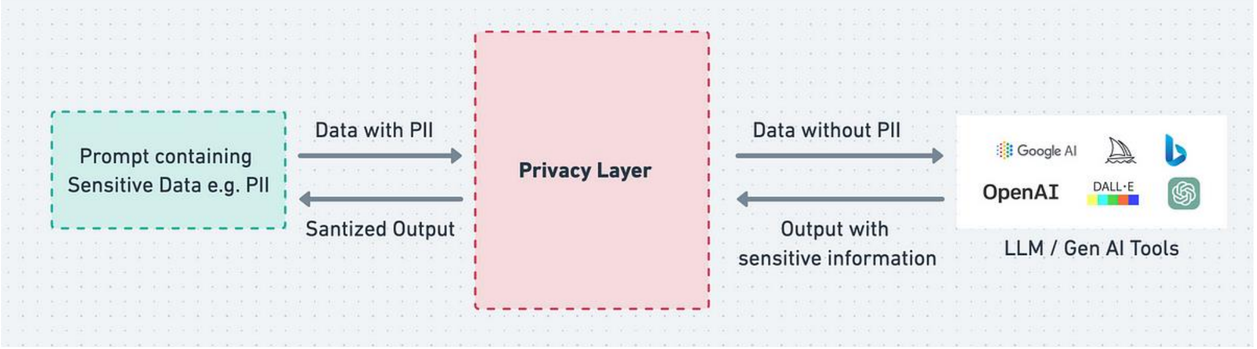

Source

Privacy-Aware Query: Real World Example

Microsoft collects diagnostic data from Windows users to identify issues and improve system performance.

Application of Privacy-Aware Query Processing:

- Differential Privacy Implementation:
  - Data Collection Controls: Users can control the level of data sharing, and Microsoft applies differential privacy to anonymize the collected data.
  - Query Processing: When querying this telemetry data to identify, for example, how often a feature is used or when errors occur, the differential privacy techniques ensure individual users cannot be identified.

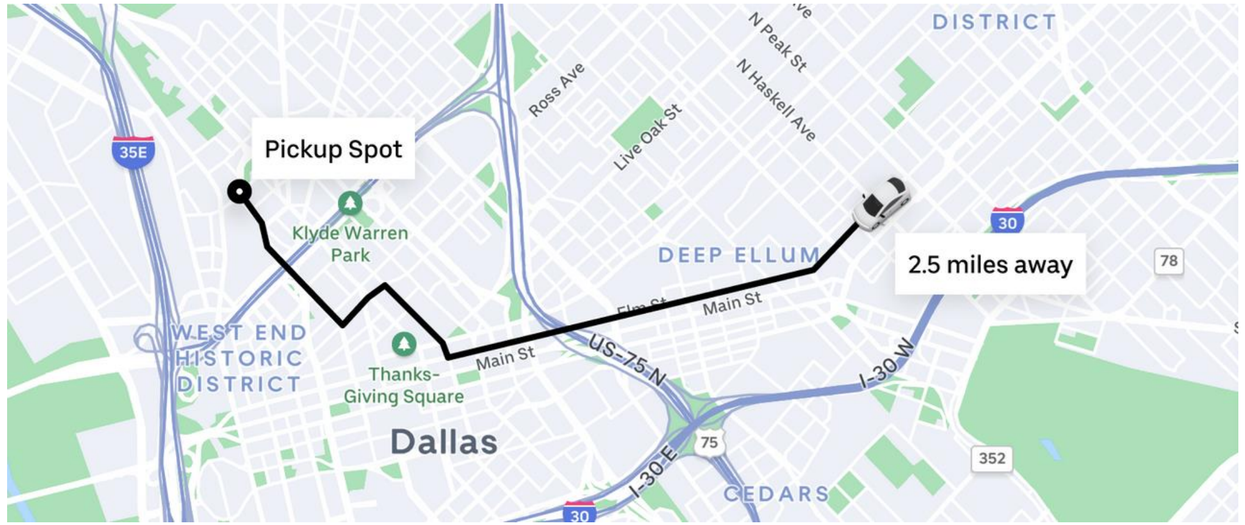

Additional Examples

- Ride-Sharing Apps: Companies like Uber use privacy-aware queries to analyze trip data without revealing individual user movements
- User Query Analysis: Search engines like DuckDuckGo emphasize privacy by not storing personal search queries.

Query Obfuscation

Modifying database queries in order to hide the true intent or sensitive information within the query.

- Methods of Query Obfuscation
  - Differential Privacy: Introducing randomness to prevent identification of individual data points.
  - k-Anonymity: Ensuring data points are indistinguishable within a group of k entities.
- Impact on Query Accuracy
  - Balancing privacy with the precision of results.
- Applications
  - Anonymous data analysis and reporting.
  - Protecting user search patterns.

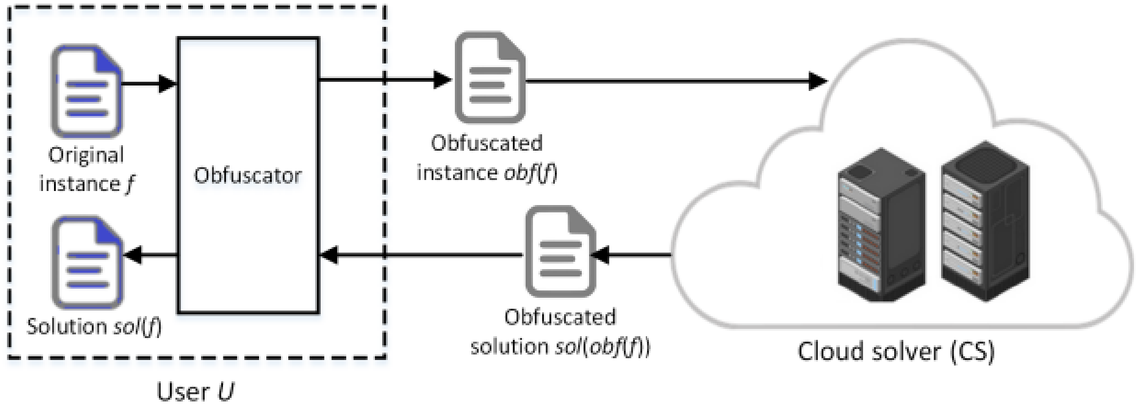

Source

Query Obfuscation: Real World Example

Obfuscation in Location-Based Services

Apps that use GPS data can inadvertently expose users' exact locations, leading to potential privacy breaches. To protect users, developers implement query obfuscation to generalize or mask precise location information.

Application of Query Obfuscation:

- Foursquare's Location Obfuscation:
  - User Check-Ins: Foursquare allows users to check in at locations, sharing this information with friends.
  - Privacy Controls: Users can choose to obfuscate their exact location by only sharing the city or neighborhood instead of the specific venue.
  - Time Delays: Introducing delays in sharing location updates to prevent real-time tracking.

Secure Aggregation Techniques

Allowing multiple parties to collaboratively compute aggregate functions over their individual data inputs without revealing those inputs to each other or any third party.

- Importance in Federated Learning
  - Training models on decentralized data without centralizing it.
- Cryptographic Methods
  - Masking Techniques: Data is masked before aggregation.
  - Homomorphic Encryption: Enables aggregation on encrypted data.
- Benefits
  - Preserves privacy without compromising collaborative benefits.
  - Reduces risk of data breaches.

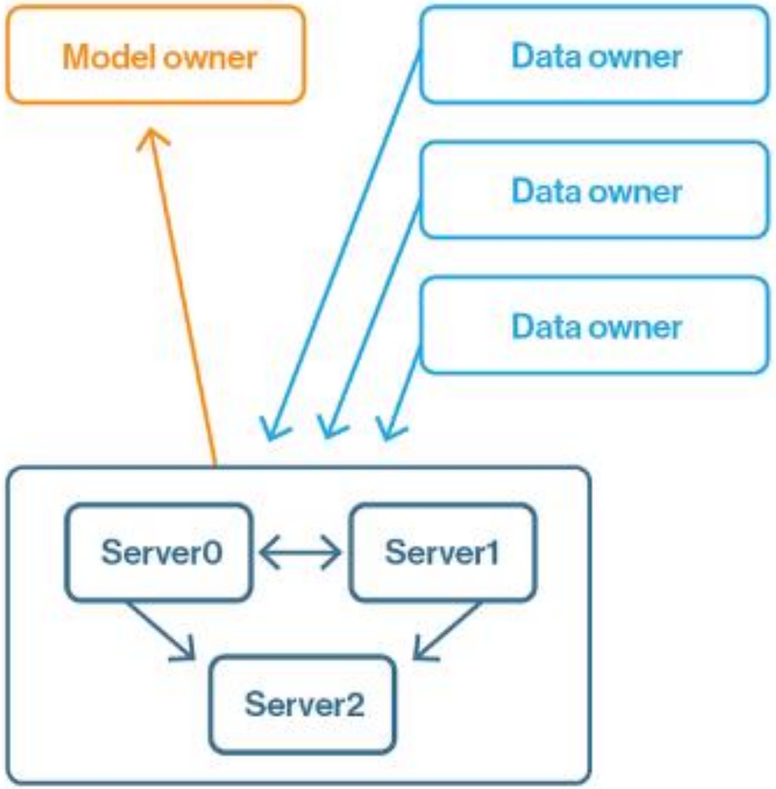

Source

Secure Aggregation: Real World Application

Virtual keyboard apps offer features like next-word prediction, autocorrect, and emoji suggestions. These features rely on machine learning models that benefit from large amounts of user data. Collecting sensitive user typing data (which may include personal messages, passwords, or other confidential information) raises significant privacy issues.

Federated Learning combined with Secure Aggregation to train models without accessing individual user data.

Additional Examples

- Apple employs similar techniques to improve Siri's voice recognition
- Collaborative Anomaly Detection in Cybersecurity: Organizations share encrypted logs or security events to collectively detect threats without revealing sensitive internal data

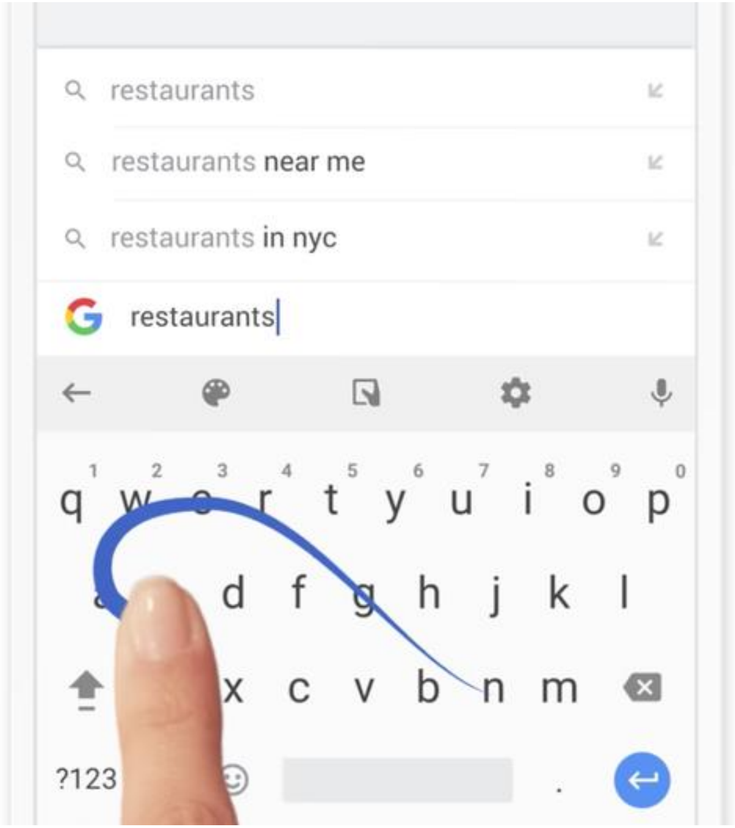

Access Control and Data Governance

- Remember a vector database is still just a database!

- Need for Robust Access Control
  - Prevent unauthorized data access.
  - Limit data exposure to necessary personnel.

- Role-Based Access Control (RBAC)
  - Assign permissions based on roles within the organization.
  - Simplifies management of user permissions.

- Data Governance Strategies
  - Policies for data lifecycle management.
  - Ensuring data quality, consistency, and compliance.

Compliance with Privacy Regulations

- Overview of Relevant Regulations
  - GDPR: Data protection and privacy in the EU.
  - CCPA: California's consumer privacy law.
  - HIPAA: Protects health information in the U.S.

- Ensuring Compliance
  - Data minimization and purpose limitation.
  - Implementing user rights like data deletion and access.

- Impact on Vector Databases
  - Need for data deletion capabilities.
  - Transparent data handling and documentation.

Balancing Privacy and Performance

- Performance Considerations
  - Privacy techniques can increase computation time and resource usage.
- Optimization Strategies
  - Algorithmic optimizations to reduce overhead.
  - Hardware acceleration using GPUs or TPUs.
- Measuring Success
  - Privacy Metrics: Quantifying the level of privacy (e.g., differential privacy's epsilon value).
  - Performance Benchmarks: Tracking latency, throughput, and scalability.

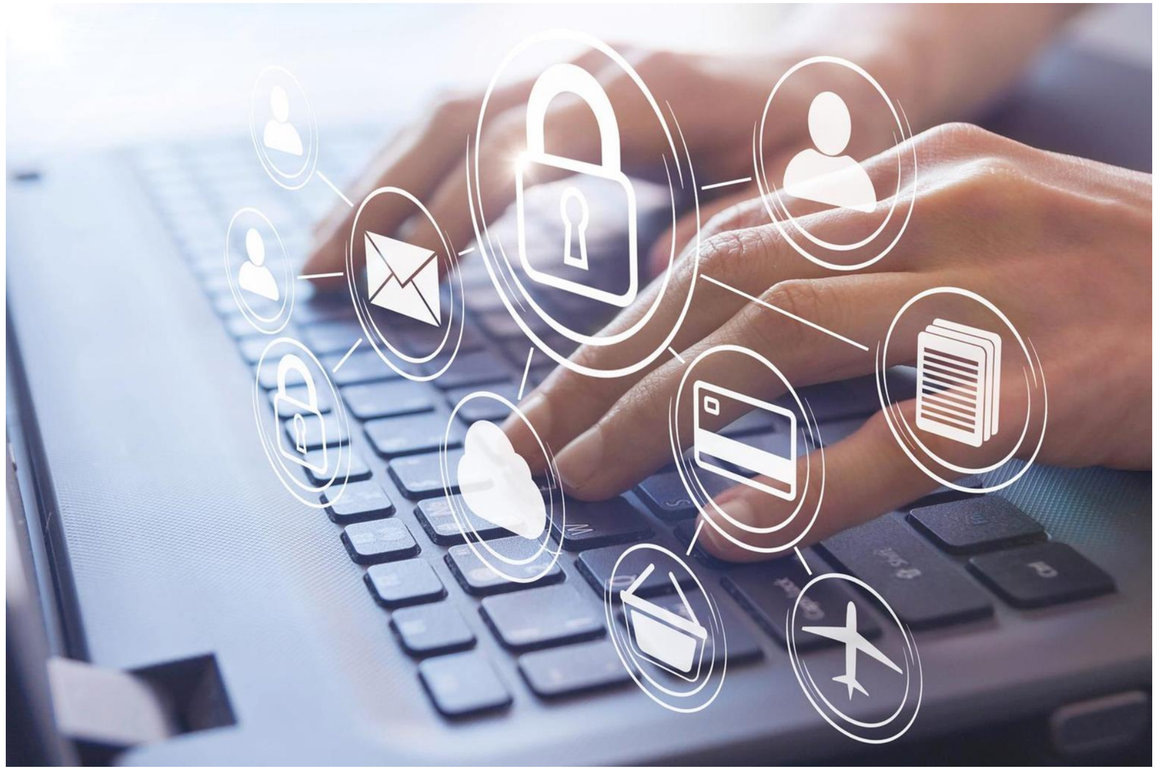

Wrap-up

Where to Find More Information

How to Use Homomorphic Encryption in the Real World

Secure Multi-Party Computation

How Meta enforces purpose limitation via Privacy Aware Infrastructure at scale

CENSOR: Privacy-preserving Obfuscation for Outsourcing SAT formulas

Key Points to Take Home

- Vector databases play a vital role in modern AI systems, particularly for LLMs, by enabling efficient similarity search and information retrieval. However, their use introduces significant privacy challenges that must be addressed.

- Advanced encryption techniques such as Homomorphic Encryption (HE) and Secure Multi-Party Computation (SMPC) are essential for securing high-dimensional data in vector databases without compromising functionality.

- Implementing techniques like query obfuscation, differential privacy, and secure aggregation can significantly enhance privacy in LLM applications while maintaining query intent and result accuracy.

- Fine-grained access control and comprehensive data governance frameworks are critical for ensuring compliance with privacy regulations and protecting sensitive information in AI systems.

- Integrating privacy considerations from the outset of system design, rather than as an afterthought, is essential for building trust in AI systems and ensuring long-term regulatory compliance.

Thank you! Any questions?

Jide Ogunjobi

https://www.linkedin.com/in/jide-o-87602512/

Please remember to complete your session evaluation.

SAPinsider

SAPINSIDER MEMBERSHIP MATTERS. YOU HAVE TO BE PART OF THE GLOBAL COMMUNITY.

Visit sapinsider.org