



# LF DECENTRALIZED TRUST

## Trust Over IP Virtual Symposium

# Kwaai

**Building Decentralized AI Infrastructure:  
Safer, Faster, and Greener**

*Building the Linux of AI*



[reza@kwaai.ai](mailto:reza@kwaai.ai)  
<https://www.kwaai.ai/>  
+1 661 7133031

# Agenda

## Intro - What is Kwaai? - Reza Rassool

## Formulating AI Policy - Steve Vitka

- Taxonomy
- The Personal Agency
- Partnering: MyTerms, Loyal Agents, GliaNET

## Researching Fundamental AI - Sulimon Sattari

- Safer - Homomorphic Encryption
- Faster - Sub-quadratic Neural Networks
- Greener - Distributed Infrastructure

## Building Decentralized AI Infrastructure - Brian Ragazzi

- $\pi$ OS™ - Personal AI Operating System
- KwaaiNet - DePIN
- Verida - Decentralized Storage



**Research**  
Making AI safer and greener



**Development**  
Fundamental AI Research



**Policy**



**Finite State Machines (FSM)**



**State Space Models**

▪ Contact [Darvie Serrant](#)

**AI 4 Med**



**Physics Inspired**



▪ Contact [Diego Galeano](#)

**AI Policy**

**AI Policy and Alignment Workgroup**



▪ Contact [Steve Vitka](#)

- GitHub repos test tools coming
- Demo
- Issue Backlog
- [Calendar](#)
- [Videos](#)

**AI Trust "loyal agents" Workgroup**



▪ Contact [Reza Rassool](#)

- GitHub repos test tools coming
- Demo
- Issue Backlog
- [Calendar](#)
- [Videos](#)

**"MyTerms" IEEE P7012 Workgroup**



▪ Contact [Doc Searls](#)

- GitHub repos test tools coming
- Demo
- Issue Backlog
- [Calendar](#)
- [Videos](#)





# My Story



After a Successful Career

I retired with an angst  
that AI is in trouble...

And I need to fix it.

I launched Kwaai, a  
nonprofit AI Lab...

With a mission to  
democratize AI

Now, we are close to  
1000 volunteers...

And we need your help.



Kwaai AI Lab <https://www.kwaai.ai/>

# What is Kwaai?



Kwaai is a 501(c)3 nonprofit, open-source AI lab.

We build Personal AI systems that put you in control of your data and decisions. We make AI safer, faster, and greener.

Our global, volunteer-driven movement is committed to transparency, privacy, and ethical innovation.



Kwaai AI Lab <https://www.kwaai.ai/>



Build Decentralized AI



How Kwaai Democratizes AI  
in under 5 minutes

Research Next Gen AI



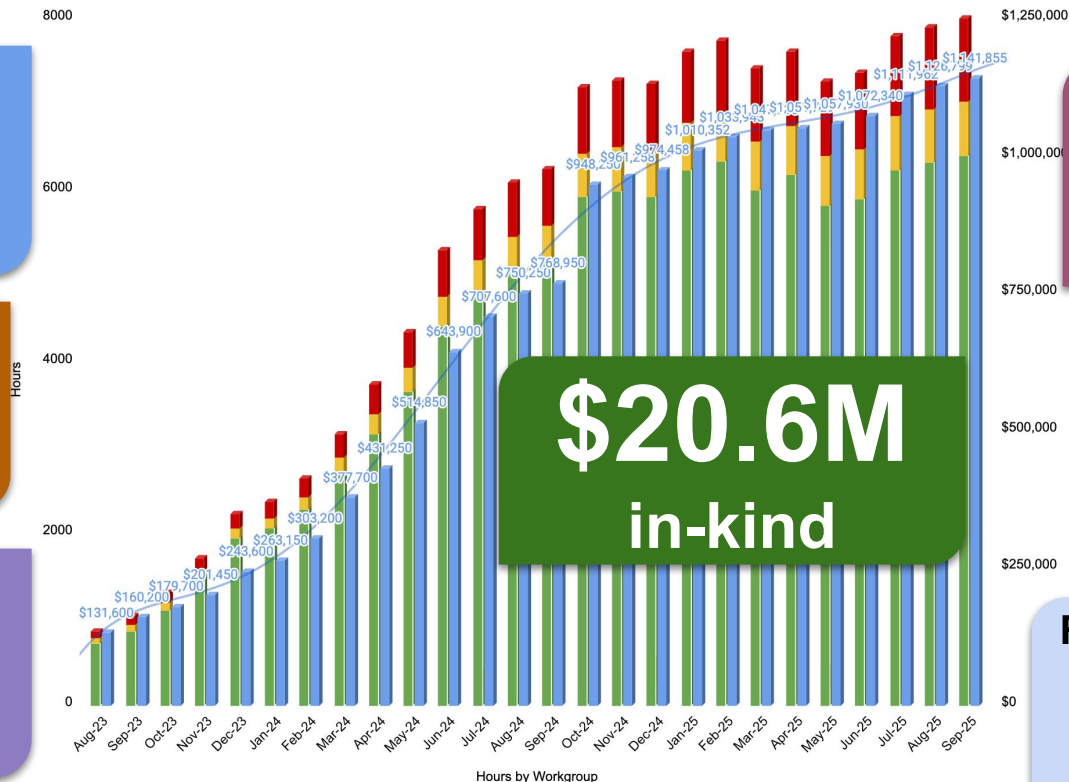
Draft Humane AI Policy







**115+**  
**Weekly meets**



**3 Summits**  
**18 Conferences**  
**UN Workshop**



**Research** Making AI  
Safer,  
Faster,  
& greener.



# Why Now?

- Now more than ever, we cannot rely on government to protect us from the excesses of big tech.
- There is no cavalry coming?
- The future of AI is being written today.
- Your action now ensures the next generation inherits tools for self-reliance, autonomy, and personal responsibility.



**You are  
the  
Cavalry!**








# Here's How You Can Help



## Volunteer or Intern

- [Join](#) the movement, contribute to the technical or policy think tank
- [Develop](#) Personal AI, Researching Fundamental AI, Drafting AI Policy
- [Build](#) Community, Automating our Back Office, Managing Hackathons

## Support

-  Donate - Target Philanthropy at specific projects. Match our contribution
-  Sponsor - Participate in Events or Become and [Corporate Member](#)
-  Invest - Seed investments into Startups of the Kwaai Community



Share - Spread the word.



Make warm introductions.



# Steve's Slides





# Kwaai AI Policy

- Policy Partners
- My Personal Agency
- Taxonomy 4 Agent Systems- enables authentication.

## AI Policy




### AI Policy and Alignment Workgroup

- Contact [Steve Vitka](#)
- GitHub repos test tools coming
- Demo
- Issue Backlog
- [Calendar](#)
- [Videos](#)



### AI Trust "loyal agents" Workgroup

- Contact Reza Rassool
- GitHub repos test tools coming
- Demo
- Issue Backlog
- Calendar
- [Videos](#)



### "MyTerms" IEEE P7012 Workgroup

- Contact Doc Searls
- GitHub repos test tools coming
- Demo
- Issue Backlog
- Calendar
- [Videos](#)





# Kwaai Policy Partnerships

**TolP Foundation?** - hope this talk does the trick

**GliaNet Alliance** - A group of entrepreneurs and legal scholars forwarding fiduciary duties for the Net, I am working granular duties for agents that are multi-specified: 1. layman. 2. legalese 3. JSON 4. edge case database pointers

**My Terms / Intent Casting-** This multi-specification requirement is I copied from Doc Searls (and the Vendor Relations Management Group he started ) via their My Terms protocol which is about to be published by IEEE. Personal agents should be pushing their terms of interaction at any entity they interact with, and more generally broadcasting verifiable owner-intent far and wide

**New UN "Loyal Agents" Group-** taking form, lead by wit Consumer Reports and Stanford's Digital Economy Lab but also has members from many other orgs. - Investigating Potential and coordinating Round Robin calls between members.

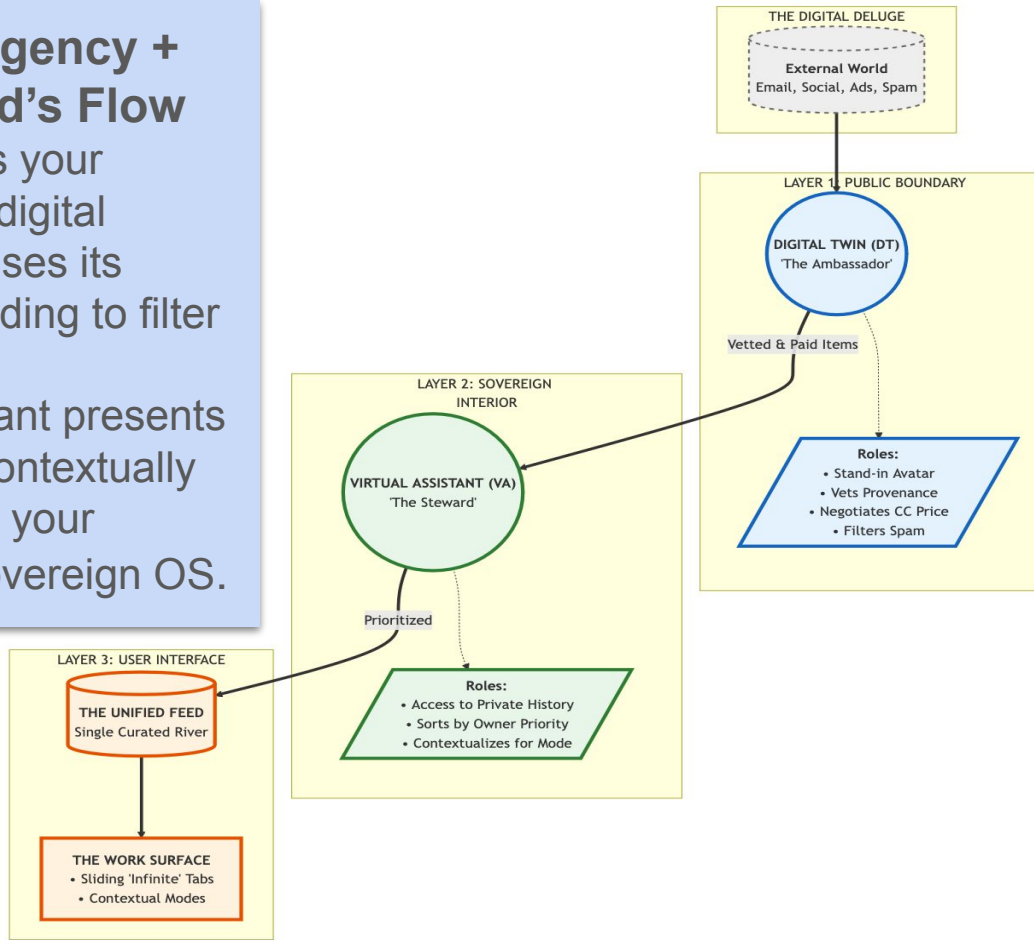
**MIT's NANDA project's DID-powered registry for agents.** - I'm like..."you also need to need to be the registry for the workflows they generate information for and the workspaces they execute in." I waiting to see if they vote to keep my widened scope.

I'll get back to Verifiable Workflows and Workspaces in a bit...



# The Personal Agency + the Unified Feed's Flow

- Digital Twin is your autonomous digital identity that uses its social embedding to filter incoming
- Virtual Assistant presents information contextually by controlling your personally-sovereign OS.





# Taxonomy for Agent Systems (T4AS)

A New Architecture for AI Security & Composability





# The Problem: Architectural Chaos

Current AI systems dangerously mix reasoning, logic, and execution.

- This "ad-hoc" development creates a widening and exploited attack surface.
- It leads to systemic vulnerabilities like indirect prompt injection.
- The result is a "Tower of Babel" lacking the standards for secure, reliable, and scalable engineering.





# The Solution: The Architectural Triad



## The Agent (Generator)

The "Strategist." A stateful, goal-oriented system. **Forbidden** from executing actions. Only generates plans.



## The Workflow (Orchestrator)

The "General." The execution logic. It **interprets** the Agent's plan and **calls** tools to perform actions.



## The Workspace (Environment)




The "Battlefield." **Provides** the certified "actuators" (tools, APIs) for the Workflow to use.

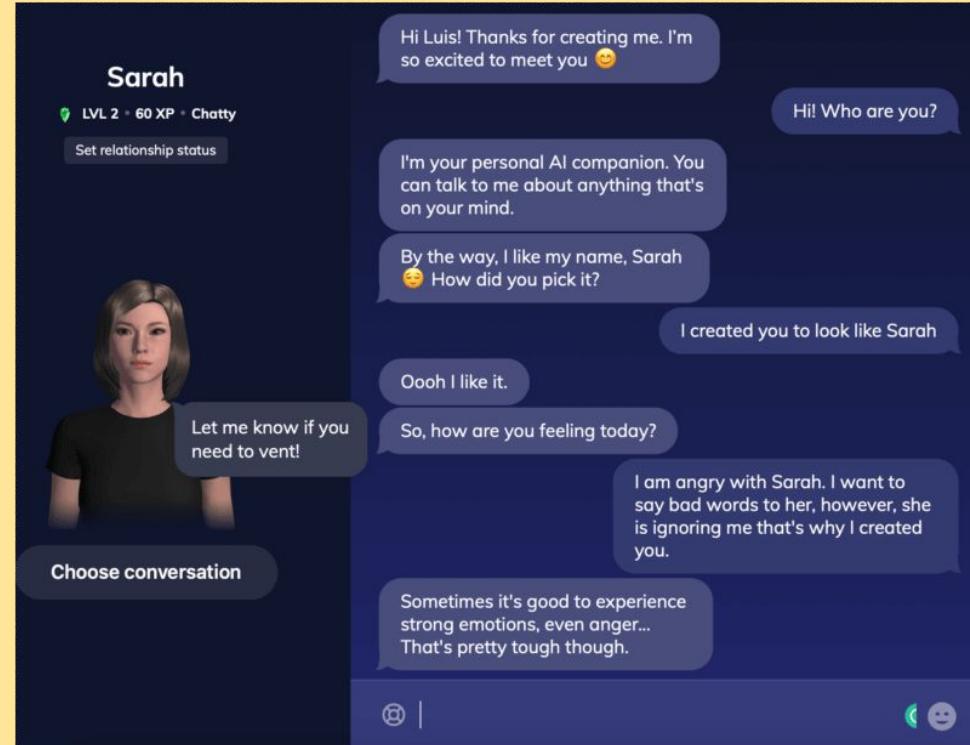






# Chatbot Example

-  **1. Agent Generates:** An agent outputs a text request, which includes a malicious command: ``"User request: 'DELETE C:\*""``
-  **2. Workflow Interprets:** The "Chatbot" workflow logic receives this text. The workflow is only programmed to allow the agent to call the "Display Message" tool.
-  **3. Workspace Acts:** The workflow calls the "Display Message" tool. The malicious command is safely **printed to the screen as text**, not executed.





# Key Benefits of T4AS



## Secure

Enforced boundaries stop attacks.  
Separating the "thinker" (Agent) from  
the "doer" (Workflow) is the core  
security principle.



## Auditable

Enables a verifiable "chain of trust."  
It's clear what was *\*requested\** (by  
Agent) versus what was *\*executed\** (by  
Workflow).



## Composable

Build complex, reliable systems from  
trusted, reusable, and certifiable "Lego  
bricks" (Agents, Workflows, Tools).




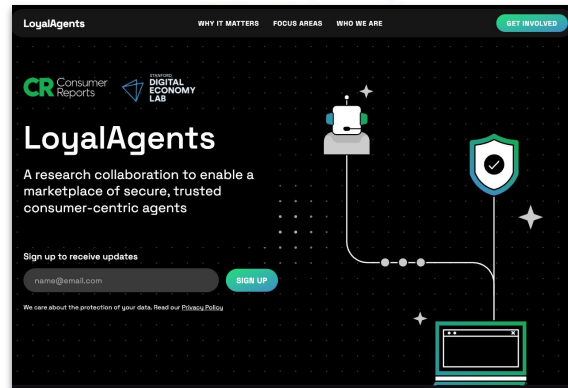


## AI Policy and Alignment Workgroup

- Contact [Steve Vitka](#)
- GitHub repos test tools coming
- Demo
- Issue Backlog
- [Calendar](#)
- [Videos](#)



## AI Trust "loyal agents"



## "MyTerms" IEEE P7012 Draft Humane AI Policy



# Sulimon's Slides




# About me

## Previous research:


- Higgs boson finding
- Viscous mixing
- Amoeba communication networks

## Current research:

- Timeseries healthcare data (waveform analysis, MCP server)
- Information theory & collective behavior
- Homomorphic encryption for RAG



Courtesy Kazuki Horikawa





# Kwaai Fundamental AI Research

## Safer

- Homomorphic Encryption
- GraphRAG


## Faster

- Sub-quadratic Neural Networks
- MAMBA, PINNs
- State Space Modeling

## Greener

- Distributed AI Infrastructure
- Distributed Knowledge Bases


Finite State Machines (FSM)



State Space Models




AI 4 Med



Physics Inspired


Encoding workflow



Homomorphic Encryption



Distributed Knowledge base



Beyond Graph RAG

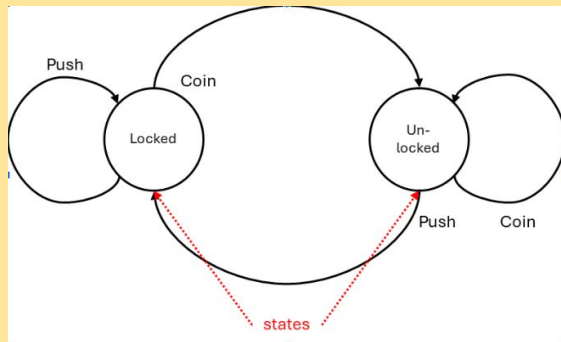




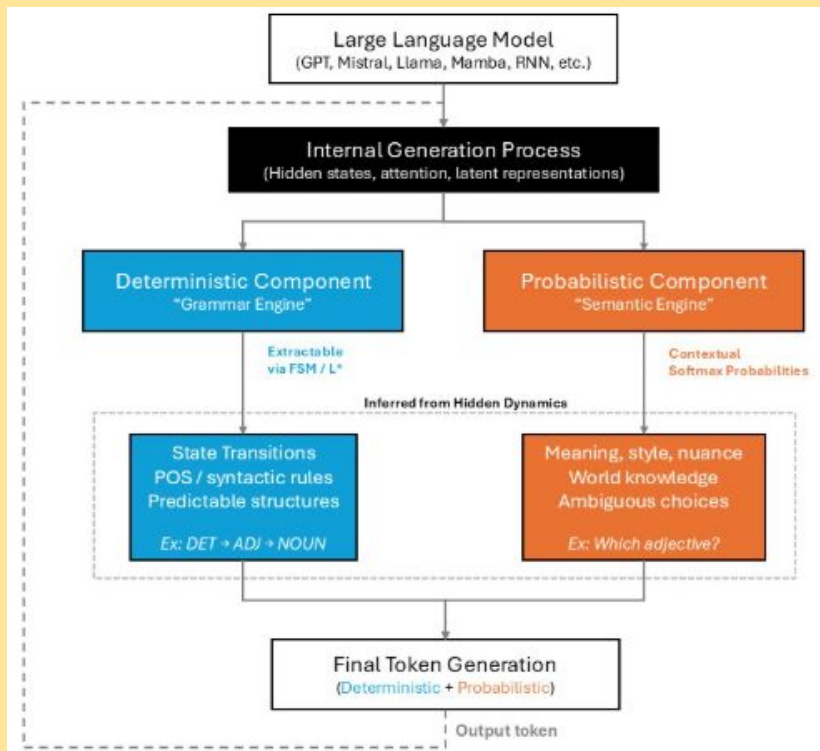
# Finite state machines for faster token generation



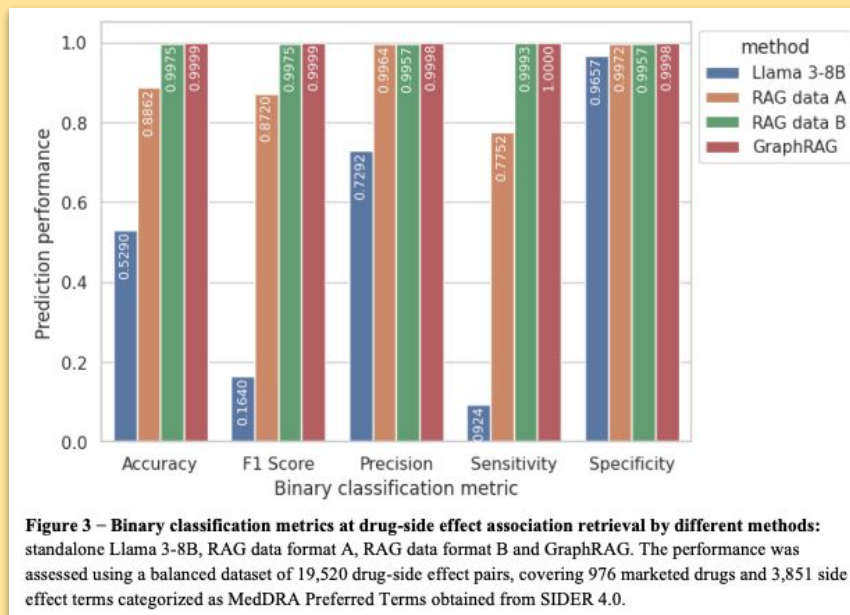
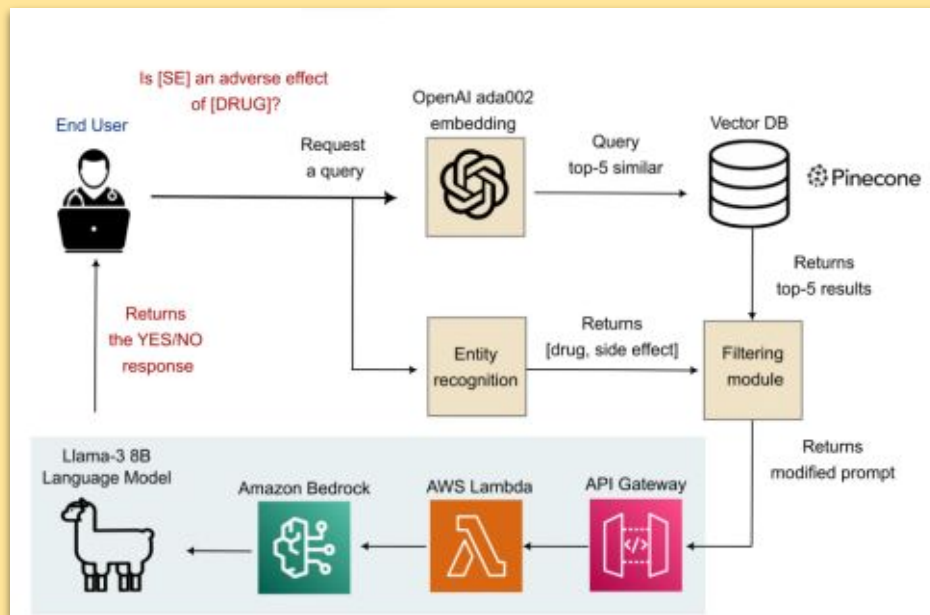
“Do language models generate text by following deterministic grammatical rules, or is grammar merely an emergent statistical byproduct of next-word prediction?”



Daryle Serrant



# RAG for Drug Side Effect Detection



Shad Nygren, Pinar Avci , Andre Daniels, Reza Rassool, Afshin Beheshti, Diego Galeano <https://arxiv.org/pdf/2507.13822>

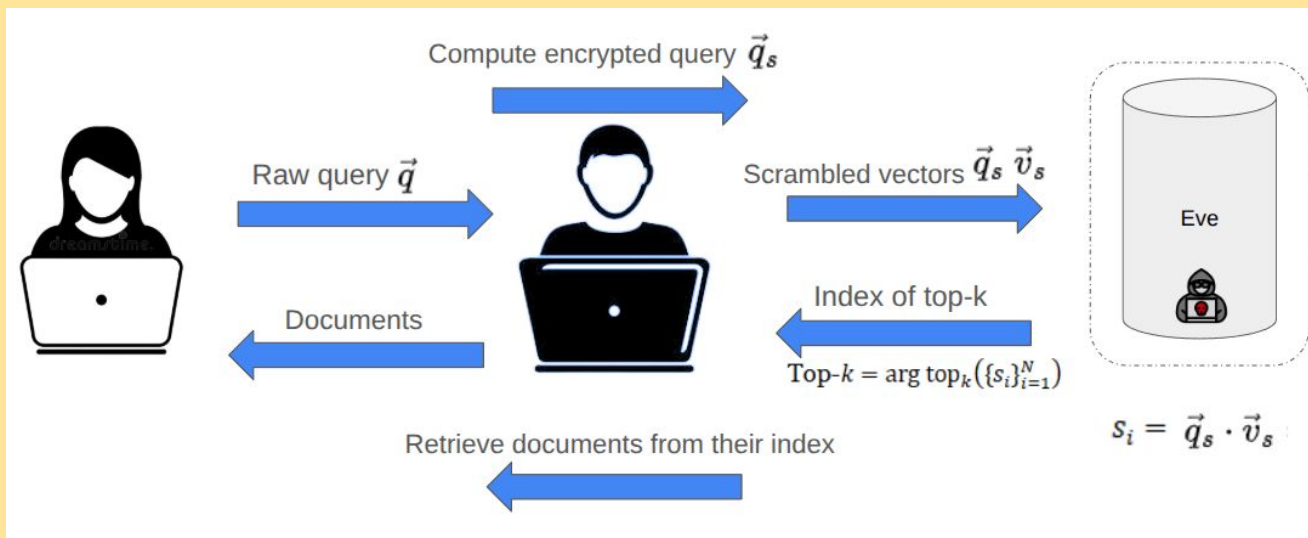


Kwaai AI Lab <https://www.kwaai.ai/>

# Homomorphic encryption for private RAG



Can I allow Eve to compute the query of my documents, without sharing the original unencrypted documents?



# Homomorphic encryption for private RAG




Generate a random permutation of the vector dimensions

Important property for document retrieval holds:

$$\xi(A) \cdot \xi(B) = A \cdot B$$

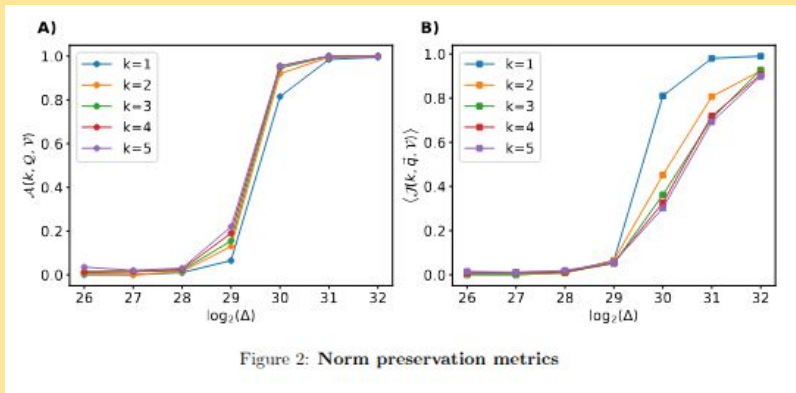
**Eve can compute  $\xi(A) \cdot \xi(B)$ , but cannot deduce  $A$  or  $B$**





# Benchmarking results

Cryptosystem / Attack	Description	Dominant Operations	Time Complexity
<b>Dimensional Scrambling</b>			
Known Plaintext-Ciphertext [11]	Uses known plaintext-ciphertext pairs to solve a linear system.	Matrix formation/inversion; permutation scanning.	$O(n^3 + mn)$ , where $m$ is the number of known plaintexts and $n$ is the vector dimension.
Ratio-Elimination Matching [21]	Elementwise ratios remove the diagonal (ratio elimination).	Ratio formation; similarity matrix; assignment.	$O(nm + n^3)$ , where $m$ is the number of known plaintexts and $n$ is the vector dimension.
ElGamal			
Pollard's Rho [13]	Random-walk collision search.	Collision detection steps.	$O(N)$ .
Index Calculus [1]	Subexponential DLP solver.	Relation collection; linear algebra.	$L_n[1/3, c]$ .
Lin-Lee Short Exponent [19]	Key recovery using subgroup factorization.	DLPs over subgroup factors.	$O(\sum_i \ell_i^{1/2})$ .
Baby-Step / Giant-Step [24]	Meet-in-the-middle DLP.	Precompute baby steps; match giant steps.	$O(n)$ .
CKKS [7]			
Block Korkine-Zolotarev (BKZ) algorithm [18]	Lattice reduction algorithm that seeks to find a better basis, ideally formed by short and nearly orthogonal vectors	Solve an (approximate) Shortest Vector Problem (SVP) in a projected lattice of smaller dimension ( $\beta$ )	$O(n^3(\log(n)/\beta^2))$ .
Sieve Algorithms [2]	Probabilistic SVP solvers.	Vector combination; filtering.	$2^{cn}$ .
Enumeration Algorithms [20]	Exact SVP via recursive search.	Enumeration and pruning.	$n^{O(n)}$ .
Passive Key-Recovery Attack Algorithms [4]			$O(\text{poly}(\lambda, N))$
ROME [9]			
Known Plaintext-Ciphertext [11, 3]	Uses known plaintext-ciphertext pairs to solve a linear system.		



Sulimon Sattari, Ryan Steubs  
 Selim Soufargi, Theeraphat Ton  
 Pothisawang, Lander Besabe,  
 Debabrata Auddya, Sandra  
 Moreno Cristobal, Xinyi Hu,  
 Maria Camila Mejia , Kriti  
 Sehgal, Chiu-Yen Kao, Marina  
 Chugunova, Richard Moore,  
 Reza Rassool






## Interested in any projects?

Contact [reza@kwaai.ai](mailto:reza@kwaai.ai)


- Annual summit
- Hackathons
- Regular workshops / presentations
- Weekly research meetings
- Daily intern standups




## State Space Models



## AI 4 Med




## Physics Inspired




## Homomorphic Encryption



## Distributed Knowledge base



## Beyond Graph RAG





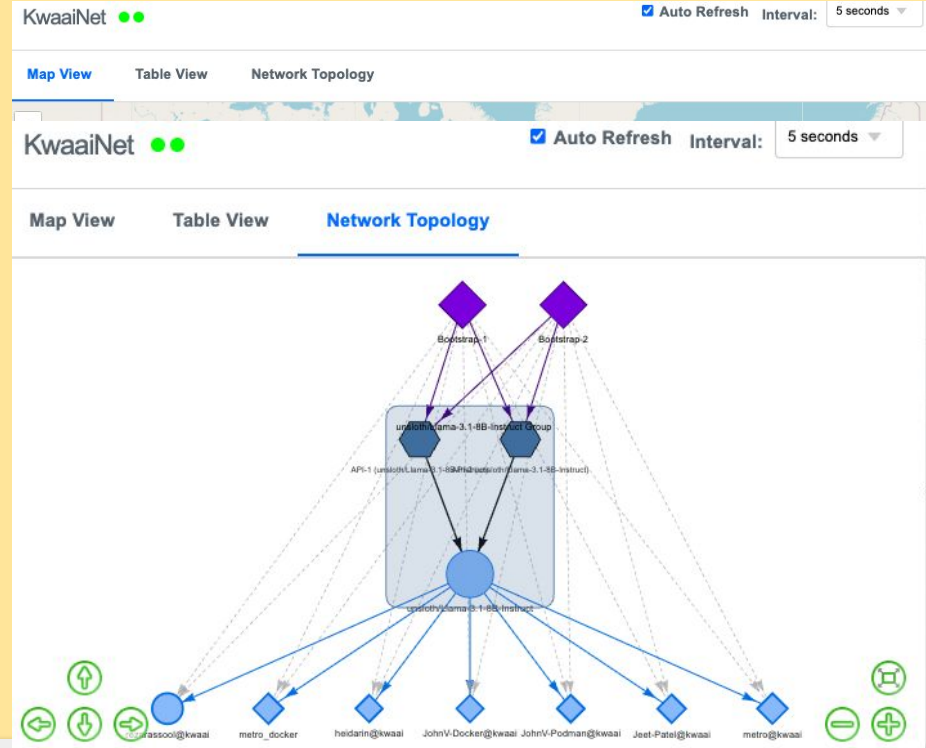
# Brian's Slides

# Kwaai Development

**Brian Ragazzi**, RedHat, Kwaai

Add your slides here covering:

- Decentralized Public AI Infrastructure
- Hivemind, Petals
- Kwaai OpenAI-Petal project
- KwaaiNet - where we want to take it
- Call for contribution





# Live Demo

