


Multi-Modal Data Exploration of Medical Data with Agentic AI




Prof. Dr. Kurt Stockinger

Zurich University of Applied Sciences, Switzerland
Affiliated with University of Zurich, Switzerland

January 22, 2026, Digital Health Zurich Seminar

Outline

- OncoSupport+: Recommendations of services for **cancer patients**
- Why is it hard to explore data in **natural language**?
- What can we learn from other domains such as **banking or astrophysics**?
- How can we use **multi-agent technology** to explore **multi-modal medical data**?
 - Structured data in databases, text, images, etc.
- How advanced is the technology for applying it in a **clinical setting**?



Digital Health Zurich Project – OncoSupport+ Recommendation of Cancer Services for Patients



Universität
Zürich UZH



USZ Universitäts
Spital Zürich

Cancer patients fill out questionnaire

Bitte kreuzen Sie an, wie stark heute ihre Beschwerden sind:

	Keine	Leicht	Mittel	Stark	
Schmerz	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	reset
Übelkeit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	reset
Erbrechen	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	reset
Luftnot	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	reset
Verstopfung	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	

System recommends services based on results of questionnaire using a decision tree algorithm



Recommended services

- Psycho-oncology
- Palliative care
- Exercise/Sport program

Health care professionals can analyze results of patients

Difrancesco, S., Bauert, M. M., et al. (2025). Collaborative Design and Development of a Patient-Centered Digital Health App for Supportive Cancer Care: Participatory Study. *JMIR Human Factors*, 12(1), e73829.

OncoSupport+ - View of the Patient



Wie fühlen Sie sich heute?

Haben Sie einen Moment Zeit, um einen Fragebogen auszufüllen? Ihre Antworten helfen uns, Sie besser zu unterstützen.

Zuletzt ausgefüllt am: 23. Oktober 2025

Jetzt ausfüllen

Supportive Beratungs- und Behandlungsdienste am Comprehensive Cancer Center Zürich

Deutsch OncoSupport+ gibt keine Behandlungsempfehlungen, sondern bietet eine Orientierung über verfügbare spezialisierte Fachdienste und damit verbundene Informationen. Ihre Symptome werden in einer Pflegesprechstunde besprochen. Wenn Sie jedoch sofortige Hilfe oder Unterstützung benötigen, wenden Sie sich bitte direkt an Ihr onkologisches Behandlungsteam.

OncoSupport+ soll nicht als Grundlage für Diagnosen oder Therapieentscheidungen verwendet werden.

English This app does not provide recommendations on treatment; instead, it offers guidance on available supportive care services and related information. Your symptoms will be shared during nursing consultations. However, for immediate help or support, we ask that you contact your oncology treatment team directly.

OncoSupport+ is not intended to be used as a basis for any diagnosis or treatment decisions.

Während der letzten Woche

Fühlten Sie sich durch Schmerzen in Ihrem alltäglichen Leben beeinträchtigt?

Überhaupt nicht

Wenig

Ziemlich

Sehr

In den letzten 3 Tagen

Waren Sie wegen ihrer Erkrankung oder Behandlung besorgt oder beunruhigt?

Gar nicht

Selten

Manchmal

Meistens

Immer

OncoSupport+ - View of the Health Care Professional



Patienten Übersicht

E-Mail Adresse: patient2@example.ch

Alter: Nicht angegeben

Beantwortete Fragebögen

24. Februar 2025 um 10:57

17. Dezember 2024 um 17:36

13. Dezember 2024 um 13:09

Schlaflosigkeit

Nummer	Frage	13.12.2024, 13:09	17.12.2024, 17:36	24.2.2025, 10:57
QLQ-C30-11	Hatten Sie Schlafstörungen?			

Ernährung

Nummer	Frage	13.12.2024, 13:09	17.12.2024, 17:36	24.2.2025, 10:57
QLQ-C30-13	Hatten Sie Appetitmangel?			
NCCN-Eating-Habits	Hatten Sie in der letzten Woche, einschliesslich des heutigen Tages, Veränderungen im Essverhalten?			
EPIC SF-13	Wie gross war das Problem von Veränderungen im Körpergewicht in den letzten 4 Wochen?			

OncoSupport+ - View of the Database

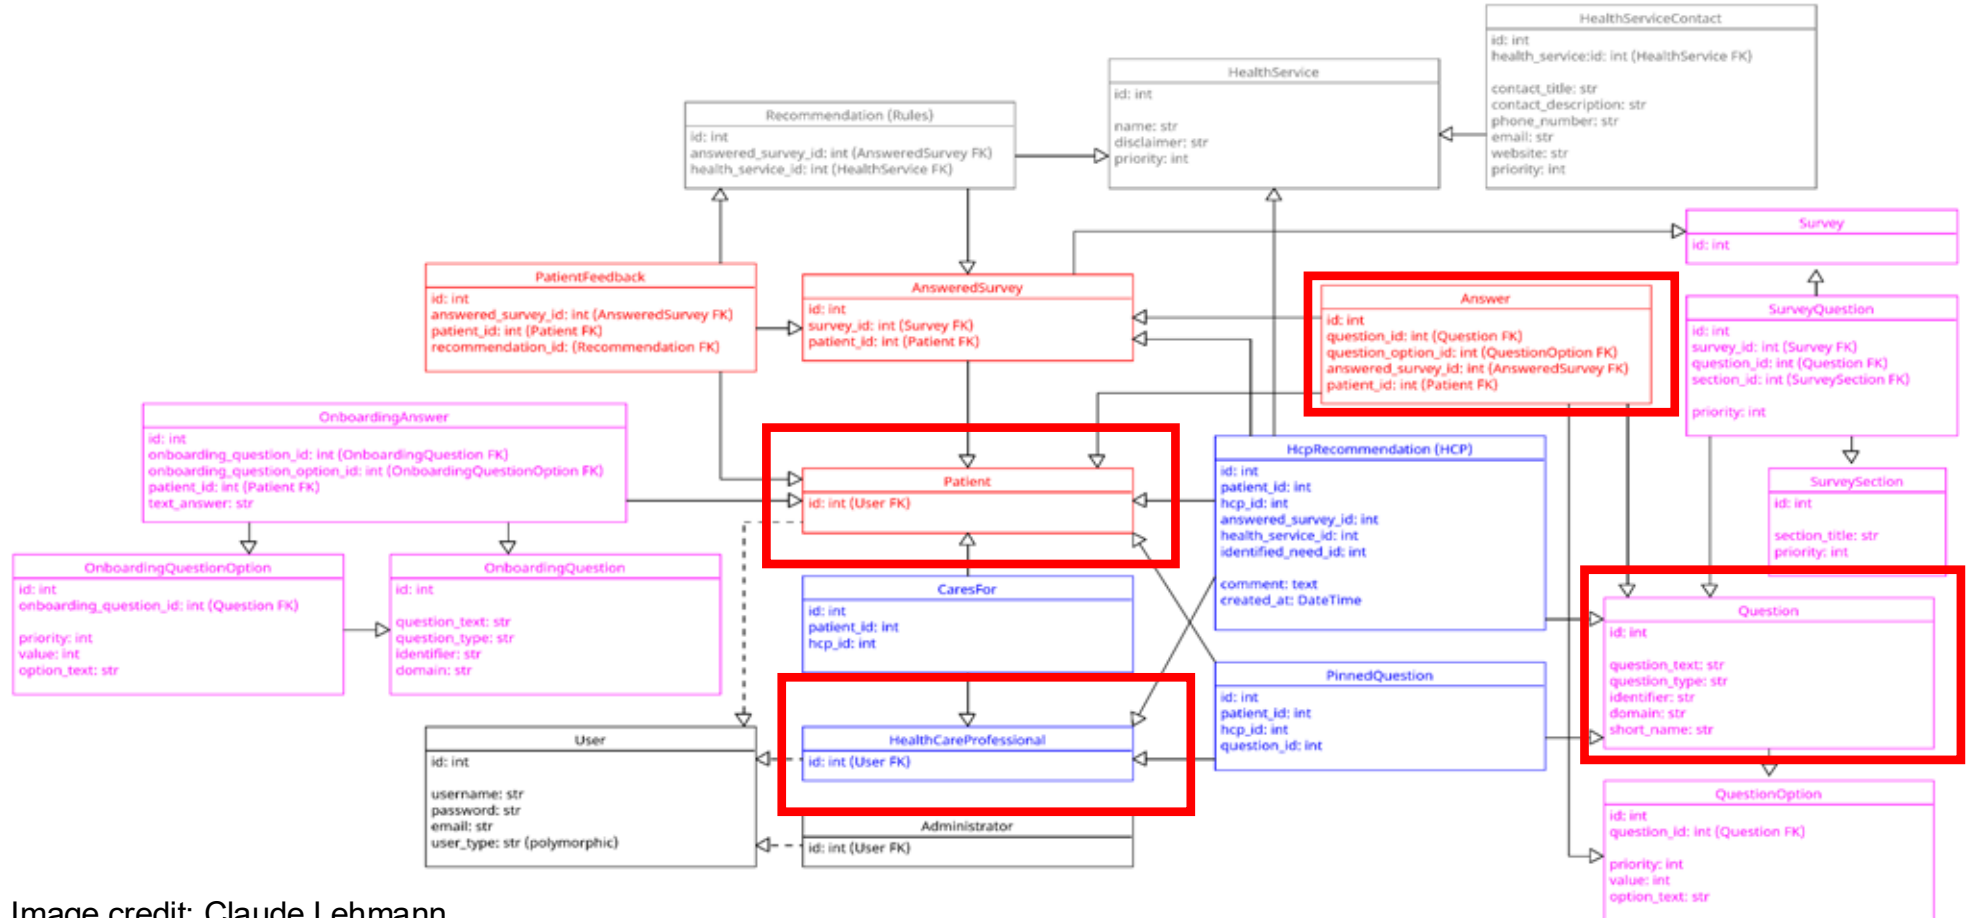
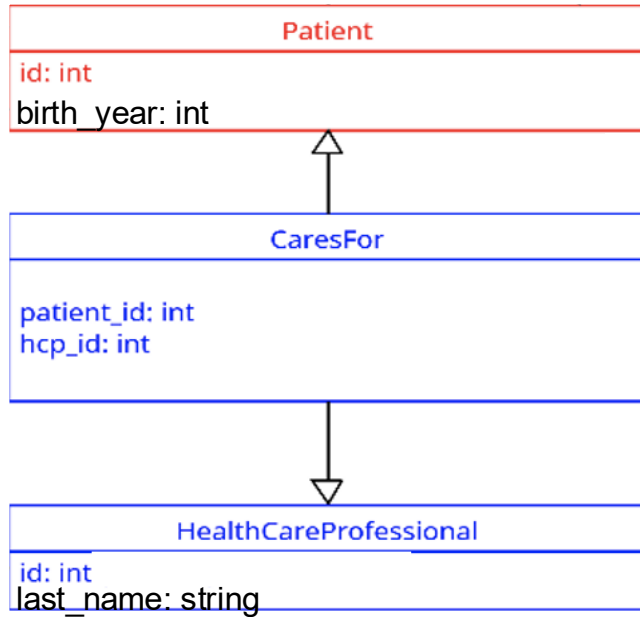


Image credit: Claude Lehmann

When Were Patients of Dr. XYZ Born?



Database query language SQL
(structured query language)

```
SELECT birth_year
FROM Patient p
JOIN CaresFor c ON p.id = c.patient_id
JOIN HealthCareProfessional h on c.hcp_id = h.id
WHERE h.last_name = 'XYZ'
ORDER BY birth_year DESC;
```

Data Exploration: Show Results of Patients that had Difficulty in Taking a Walk

Need to write analysis in **database query language SQL** (structured query language)

```
WITH prepared_answers AS (
  SELECT
    a.*,
    qo.value,
    qo.option_text
  FROM answers a
  JOIN question_options qo ON a.question_option_id = qo.id
  WHERE a.question_id = 2
),
ordered_questions AS (
  SELECT
    *,
    ROW_NUMBER() OVER (
      PARTITION BY patient_id, question_id
      ORDER BY created_at
    ) AS rn,
    LAG(value, 1) OVER (
      PARTITION BY patient_id, question_id
      ORDER BY created_at
    ) AS v1,
    LAG(value, 2) OVER (
      PARTITION BY patient_id, question_id
      ORDER BY created_at
    ) AS v2
  FROM prepared_answers
),
worsening_questions AS (
  SELECT *
  FROM ordered_questions
  WHERE
    value > 1
    AND value > v1
    AND v1 > v2
)
SELECT o.patient_id, q.identifier, q.question_text, o.value, o.option_text, o.created_at
FROM ordered_questions o
JOIN worsening_questions w
  ON o.patient_id = w.patient_id
  AND o.question_id = w.question_id
  AND o.rn BETWEEN w.rn - 2 AND w.rn
JOIN questions q ON o.question_id = q.id
ORDER BY o.patient_id, o.created_at;
```

Simplified (anonymized) result

value	option_text	created_at
1	Überhaupt nicht	2024-12-13 12:10:33.544192 +00:00
2	Wenig	2024-12-13 12:19:24.369030 +00:00
4	Sehr	2024-12-13 13:22:43.885492 +00:00

Image credit: Claude Lehmann

The Data Exploration Vision with AI Agents



- Doing data exploration using the database query language **SQL is hard**
- Why not doing it in **natural language** and let **AI agents** do the work?


Challenges of Natural Language Questions #1



- Ambiguity:
 - *How many great people come from Zurich?*

Challenges of Natural Language Questions #2

- Ambiguity:
 - *How many great people come from Zurich?*
 - What is meant by Zurich?
 - A city, a canton or a company?



Challenges of Natural Language Questions #3



- Ambiguity:
 - *How many great people come from Zurich?*
 - What is meant by great?
 - Is it the *size*? In meters, in feet, in inches?
 - Is it the *achievement* of these people? How do you measure it?

What **approaches** can be used to tackle the problem?

Can we leverage experience from **domains outside health**?

SODA – Search over Data Warehouse

Our First Try in 2012: “little” NLP, no ML, no AI



SODA: Generating SQL for Business Users

Lukas Blunski
ETH Zurich, Switzerland
lukas.blunski@inf.ethz.ch

Claudio Jossen
Credit Suisse AG, Switzerland
claudio.r.jossen@credit-
suisse.com

Donald Kossmann
ETH Zurich, Switzerland
donald.kossmann@inf.ethz.ch

Magdalini Mori
Credit Suisse AG, Switzerland
magdalini.mori@credit-
suisse.com

Kurt Stockinger
Credit Suisse AG, Switzerland
kurt.stockinger@credit-
suisse.com

ABSTRACT

The purpose of data warehouses is to enable business analysts to make better decisions. Over the years the technology has matured and data warehouses have become extremely successful. As a consequence, more and more data has been added to the data warehouses and their schemas have become increasingly complex. These systems still work great in order to generate pre-canned reports. However, with their current complexity, they tend to be a poor match for non tech-savvy business analysts who need answers to ad-hoc queries that were not anticipated.

This paper describes the design, implementation, and experience of the SODA system (Search over DATA Warehouse). SODA bridges the gap between the business needs of analysts and the technical complexity of current data warehouses. SODA enables a Google-like search experience for data warehouses by taking keyword queries of business users and automatically generating executable SQL. The key idea is to use a graph pattern matching algorithm that uses the metadata model of the data warehouse. Our results with real data from a global player in the financial services industry show that SODA produces queries with high precision and recall, and makes it much easier for business users to interactively explore highly-complex data warehouses.

A second observation that can be made in modern data warehouses is that there is a growing gap between the high-level (conceptual) view of business users and the low-level (physical) perspective of database administrators. Business users still think of the data in star schemas with fact tables in the center and dimension tables as satellites [13]. Database administrators need to integrate many such star schemas of different kinds of business users with varying information needs into a single physical schema. Their job is to optimize the data warehouse, thereby minimizing cost (i.e., \$) and meeting all performance goals (i.e., response time and throughput). At the same time, they must manage the data and the schema.

Given these differences in goals, it is not surprising that the conceptual world of business users and the physical world of database administrators is very different. For instance, database administrators may implement a simple business concept such as *Customer* using many different tables, thereby partitioning the data horizontally and vertically. Furthermore, database administrators may store information from different business entities in a single table if that helps improve performance or manageability. Database administrators may also implement *inheritance* and *generalization* in different ways, depending on the query workload that they anticipate. As an extreme example, database administrators may use cryptic naming schemes for table and column names, thereby helping them with

The SODA Intuition: Combine Database Technology with Information Retrieval

- Input query in natural language:

- Find the top 10 costumers with securities of Bank A and rank them by the trading volume.

- Assumptions:

- No complex natural language processing, only key word search.
- Only those entities can be extracted that are part of the schema graph or the domain ontology.
 - Unknown entities would result in a processing error.

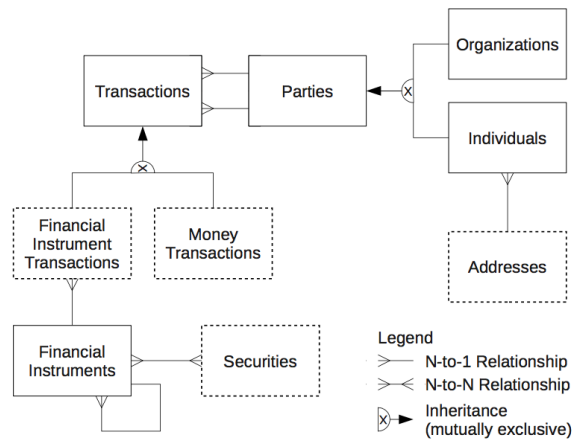
- Query consists of:

- **Business objects** (exist in schema graph or domain ontology as entities and attributes)
- **Operators** (e.g. top n, rank, below, all, etc.)
- **Values** (exist in base data, e.g. Bank A, 10)

- Hence, our input query is interpreted as follows:

- Find the **top 10 costumers** with **securities** of **Bank A** and **rank** them by the **trading volume**.

Data Structure: Inverted Indexes over Base Data and Meta Data




Keyword	Table	Column
Zurich	Addresses	City
Zurich	Addresses	Canton
Zurich	Organizations	Name
Stockinger	Individuals	LastName
Bond	Financial Instruments	Type
Bond	Individuals	LastName

INODE – Intelligent Open Data Exploration

<http://www.inode-project.eu/>



- Users should **interact with data** in a more dialectic and intuitive way similar to a **dialog with a human**
- Services for exploration of open data sets that help users:
 - **Link** and leverage **multiple datasets**
 - Access and **search data using natural language**, using examples and using analytics
 - **Get guidance from the system** in understanding the data and formulating the right queries
 - Explore data and discover new insights through **visualizations**
 - Focus on **Astrophysics, Cancer Biomarker Research** and **Research & Innovation Policy Making**



INODE – Intelligent Open Data Exploration

A classic unified, comprehensive platform that provides interactive access to open datasets through natural language queries in the fields of Cancer Biomarker Research, Strategic and Innovation Policy Making and Astrophysics, for a wide range of users from larger scientific communities to public.

Project Information

INODE
Grant agreement ID: 863410

Start date
1 November 2019

End date
30 April 2023

Funded under
EXCELLENT SCIENCE - Research Infrastructures

Total cost
€ 5 732 000

EU contribution
€ 5 732 000

Coordinated by
ZÜRCHER HOCHSCHULE FÜR ANGEWANDTE WISSENSCHAFTEN
Switzerland

Participants (8)

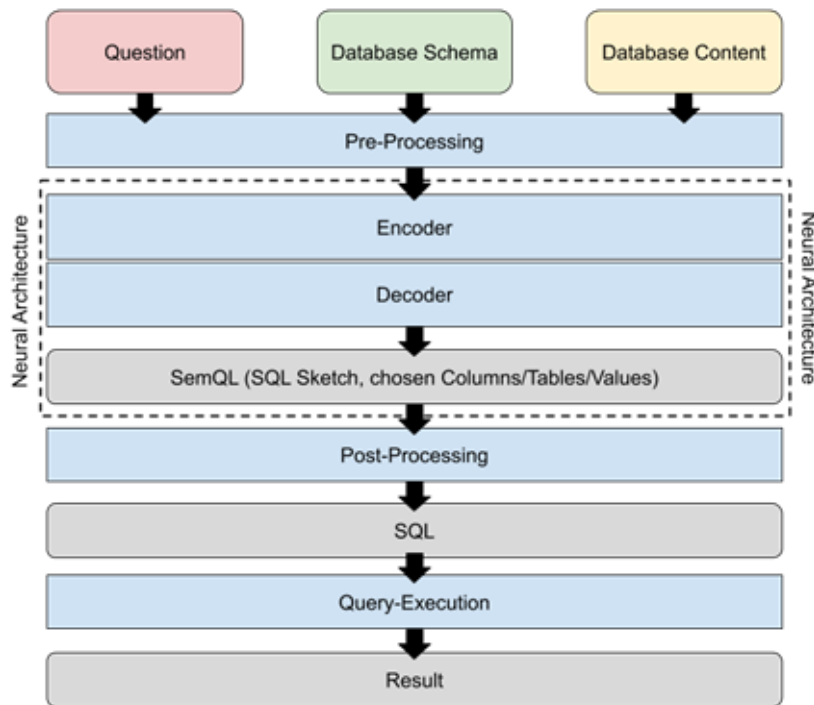
- ATHINA-EREYINTKO KENTRO KANOTOMAS STIS TECHNOLOGIES TIS PLIROFORIAS, TON EPIKONINON KAI TIS ENDSIS (Greece)
- MAX PLANCK GESELLSCHAFT ZUR FÖRDERUNG DER WISSENSCHAFTEN EV (Germany)
- FRAUNHOFER GESELLSCHAFT ZUR FÖRDERUNG DER ANGEWANDTEN FORSCHUNG E.V. (Germany)
- SIBIS ACADEMIC SL (Spain)
- CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS (France)
- LIBERA UNIVERSITA DI BOLZANO (Italy)
- SIB INSTITUT SUISSE DE BIOINFORMATIQUE (Switzerland)
- INFUJ TECHNOLOGIES PRIVATE COMPANY (Greece)



ValueNet: A Transformer-Based Neural Network Architecture

- **Generate SQL** given a natural language question – end to end
- At its core a **neural network** – consisting of an encoder / decoder architecture
- Generates an **intermediate language** – **SemQL** – which abstracts technical details
- SemQL is **deterministically transformed** to SQL, or any other query language (e.g. SPARQL)
- Uses state of the art **pre-trained transformers** to understand the natural language question.

Today we would call this RAG (Retrieval Augmented Generation)
Trick: Smart pre-processing and post-processing





Encoding of Question & Database Schema

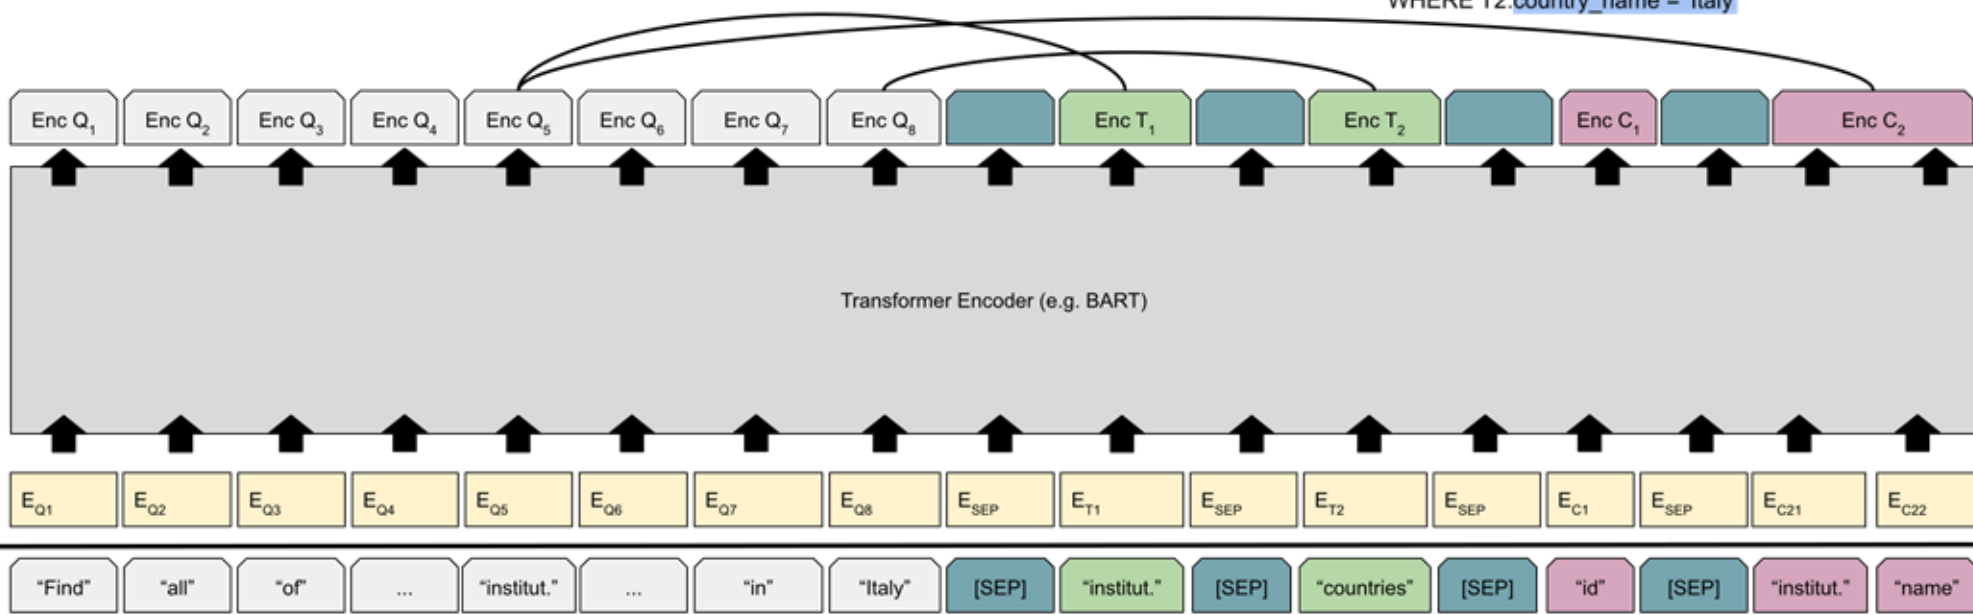
Question:

Find all of the **institutions** located in **Italy**.

Query:

```
SELECT T1.institutions_name
FROM institutions AS T1
JOIN countries AS T2 ON T1.country_id = T2.unics_id
WHERE T2.country_name = 'Italy'
```

Attentions



Natural language question

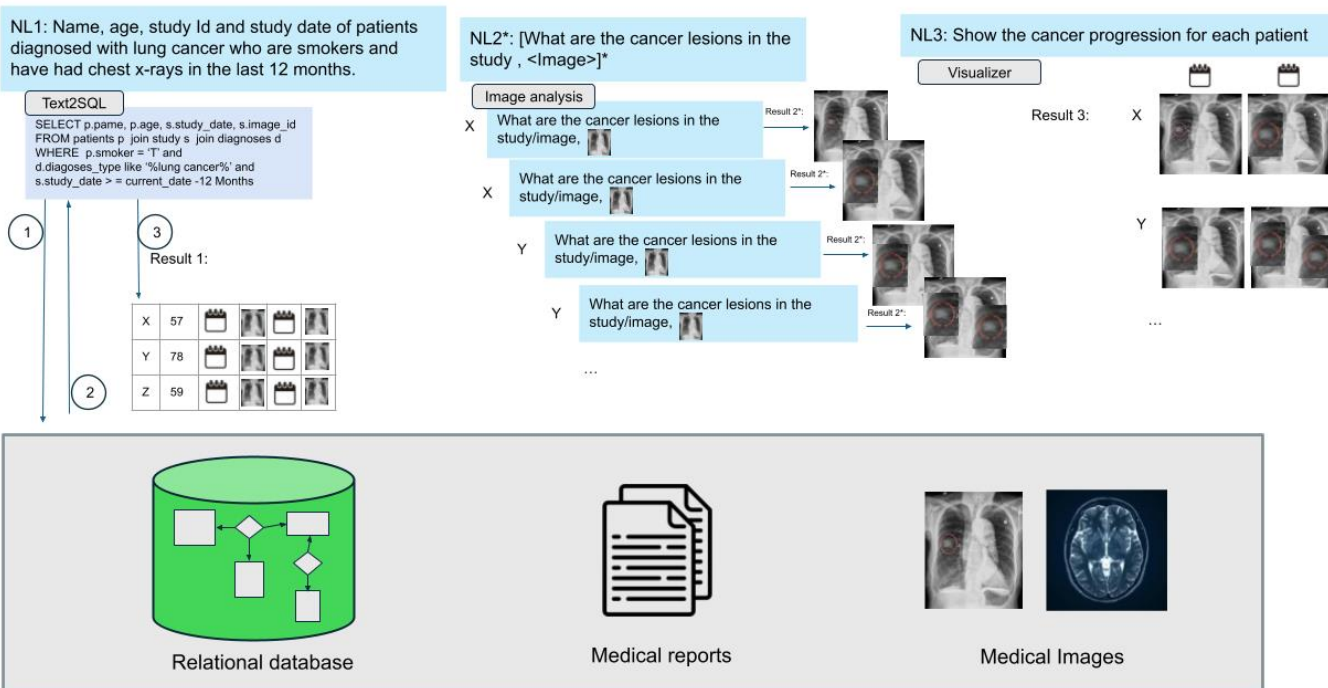
Tables and columns of corresponding SQL query

The even bigger challenge:

How can we do **multi-modal data exploration**?


Multi-Modal Exploration of Medical Data

Show me the progression of cancer lesions over the last 12 months of patients with lung cancer who are smokers.

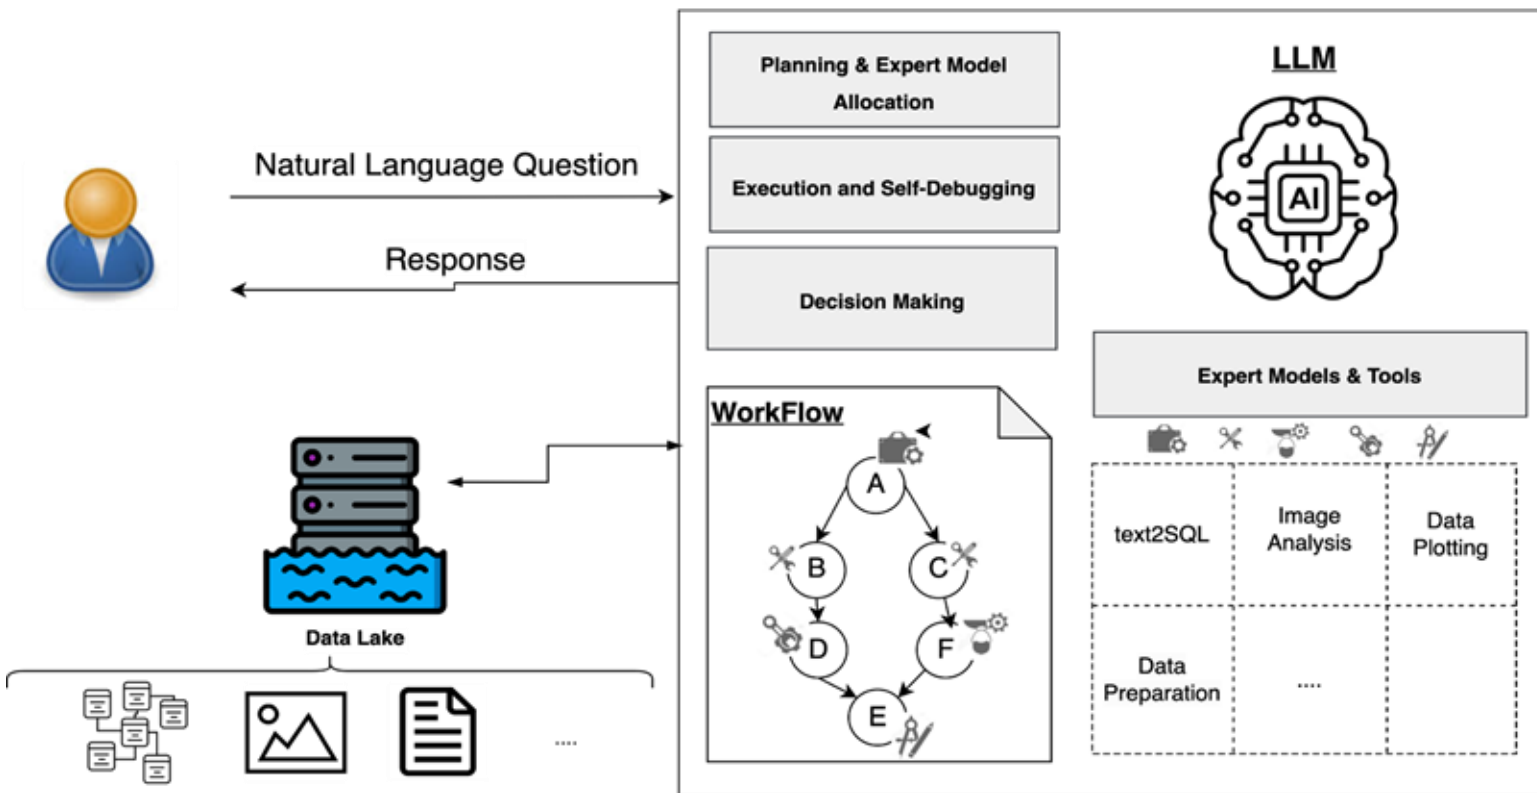


Major Challenges

- Translation from text to SQL
- Question answering over text documents
- Visual question answering



M²EX: Multi-Modal Data Exploration in Natural Language via LLM Agents (LLM = Large Language Model)



Explanation Data Exploration with M²EX

Was patient 18061894 prescribed acetaminophen, and did a chest x-ray show any technical assessments until 12/2103?

i Planning & Expert Model Allocation

workflow plan	intent_tables	db_schema
t_0	Was patient 18061894 prescribed acetaminophen?	db_schema
t_1	Did a chest x-ray show any technical assessments for patient 18061894 until 12/2103?	db_schema
t_2	Was patient 18061894 prescribed acetaminophen?	$\$t_0$
t_3	Retrieve study_id and image_id for patient 18061894's chest x-ray studies until 12/2103.	$\$t_1$
t_4	Did the chest x-ray show any technical assessments?	$\$t_3$


Expert Models & Tools

Data Lake

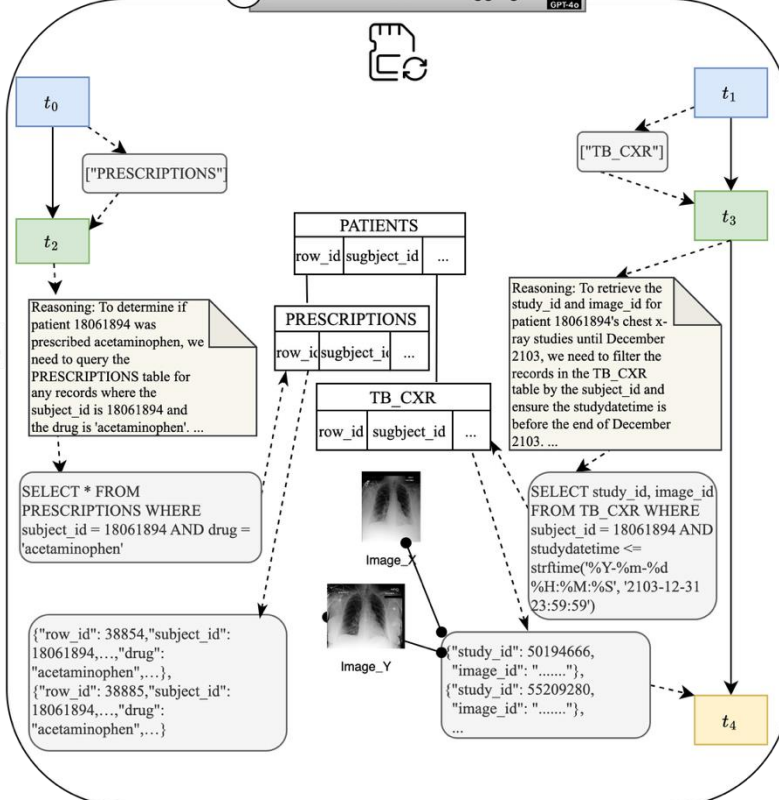
intent_table



image_analysis



ii Execution and Self-Debugging





iii Decision Making

Thought: I have successfully retrieved the necessary information. Patient 18061894 was prescribed acetaminophen, and the technical assessment of the chest x-ray was successful. Therefore, I can provide a final response.


{'Summary': 'Patient 18061894 was prescribed acetaminophen and the chest x-ray showed technical assessments until 12/2103.', 'details': 'Patient 18061894 was prescribed acetaminophen between 2103-12-26 and 2103-12-30. Several chest x-ray studies were performed, and they had technical assessments.', 'source': 'The PRESCRIPTIONS and TB_CXR tables from the database.', 'inference': 'yes', 'extra explanation': 'The chest x-ray studies had technical assessments.'}

Evaluation of Our Agentic AI System

- **Medical benchmark dataset:**
 - Electronic health records EHRXQA
 - Multi-modal question answering data with relational data + chest X-ray images
- **Baseline system:**
 - Neural SQL: LLM integrated with a visual question answering system (NeurIPS 2024)
- **LLM and agent technology:**
 -  (Claude and Anthropic Claude Sonnet 4.5 and Qwen3 VL 8B Instruct)
 - 

Performance of M²EX vs. State-of-the-Art

System		Scope			Output Type		Overall (100)	Generated	Replanning
		Image Single-1 (30)	Image Single-2 (30)	Image+Table Single (40)	Binary (50)	Categorical (50)		Plan	
NeuralSQL	zero-shot	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	N/A	No
	few-shot ($n = 10$)	26.67%	20.00%	47.50%	48.00%	18.00%	33.00%		
M ² EX	zero-shot	23.33%	43.33%	77.50%	74.00%	28.00%	51.00%	98%	Yes



Conclusion

- Multi-modal data exploration using **agentic AI shows promising results**
- **Improvements required** to achieve higher accuracy and better explanations
- Next step: validate the approach in a **clinical setting**



[Dr. Jonathan Fürst](#)



[Dr. Farhad Nooralahzadeh](#)



[Yi Zhang](#)

Core Agentic AI Team of M²EX

