

Enabling user-friendly interaction and data analysis with 3D city models through LLM agents

June 20, 2026

EnergyADE 3.0 Workshop, TU Delf

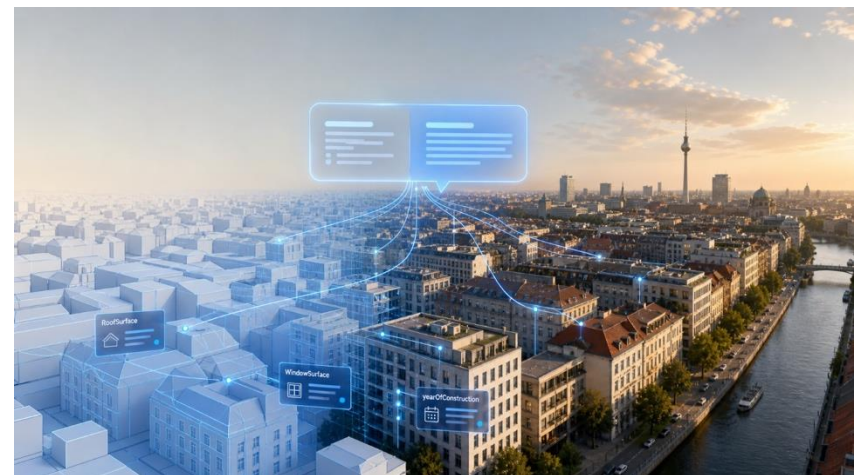
Khaoula Kanna

Thomas H. Kolbe

Technical University of Munich

Chair of Geoinformatics

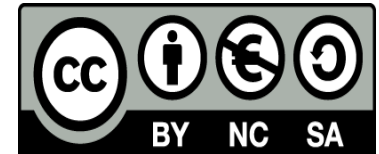
khaoula.kanna@tum.de



@ Image Generated with Gemini

License

This presentation is licensed under the [Creative Commons License CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/). According to CC BY-NC-SA 4.0 permission is granted to share this document, i.e. copy and redistribute the material in any medium or format, and to adapt it, i.e. remix, transform, and build upon the material under the following conditions:



- **Attribution:** You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
- **NonCommercial:** You may not use the material for commercial purposes.
- **ShareAlike:** If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.
- **No additional restrictions:** You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Urban Digital Twins: Rich data, Narrow access

What we have:

- ▶ **Smart city projects** often require a lot of **expertise** to analyze Urban Digital Twins data for decision-making purposes. This can be very **time-consuming** as well as very **expensive**.
- ▶ **Urban Digital Twins (UDTs)** use semantic 3D city models to represent the built environment and sensor data for real-time data analysis.

What we don't have:

- ▶ No **direct interaction** with the UDT for **non geospatial experts**.
- ▶ **No automatic way of enriching** the UDT with open or closed data.

Opportunities:

- ▶ Recent advancements on **Generative AI**, specifically LLMs has facilitated intuitive question answering which **bridged the gap between data and end-users**.

From Talking to Acting: in 4 steps

Talk to

the twin

How does a non-expert query the UDT in natural language?

Enrich

the twin

How does the twin enrich itself from external sources?

Scale across

twins

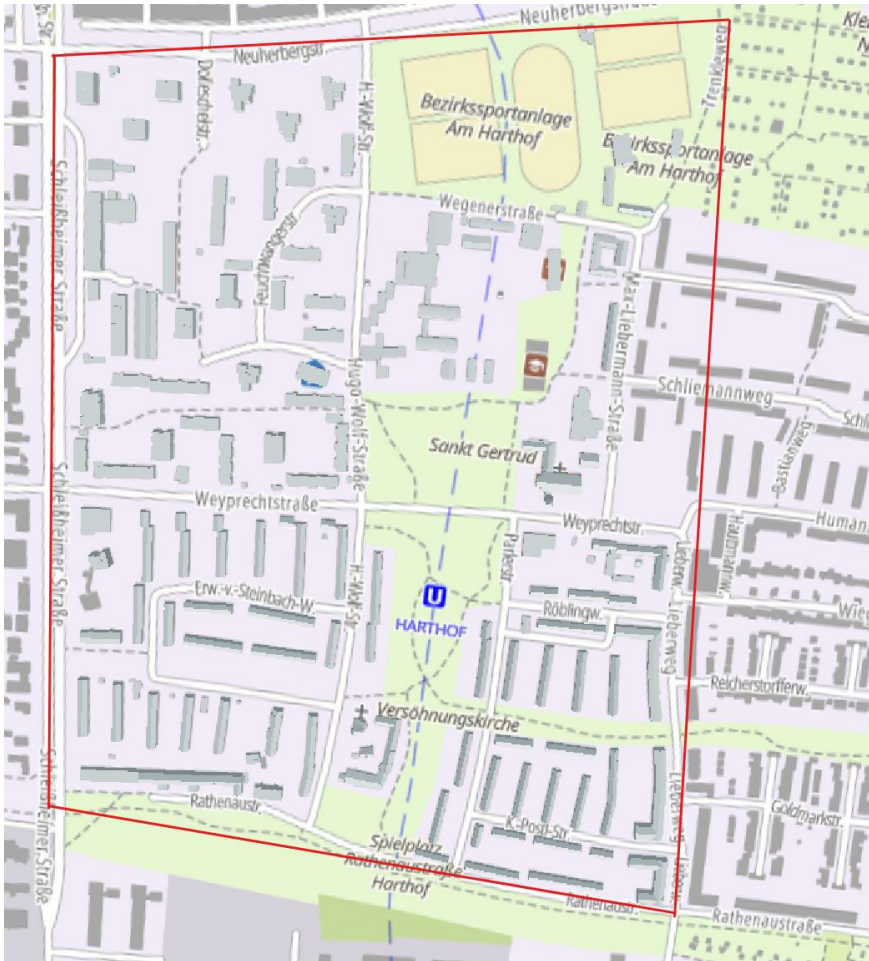
How can the UDT domain knowledge be made automatically accessible to AI agents?

Act on

the twin

How can users edit and simulate scenarios in an UDT through natural language?

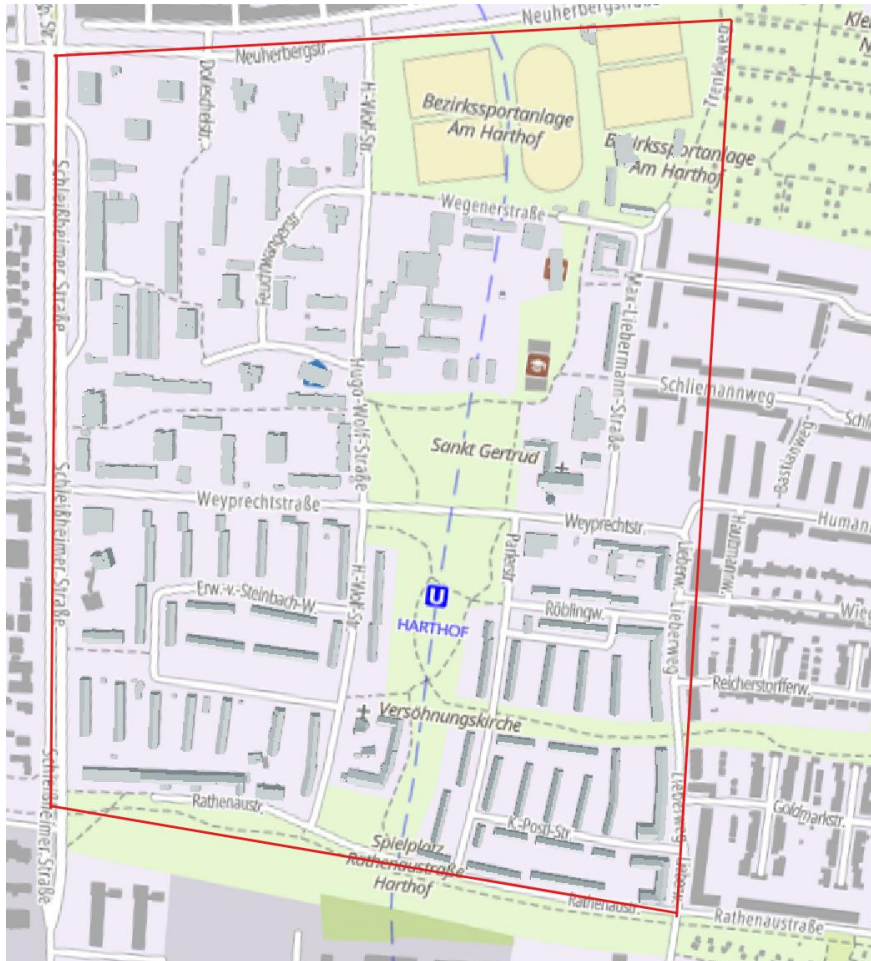
Our Use Case



Harthof district, Munich, Germany

- ▶ **An EU Funded Project called ASCEND**
- ▶ *The goal is to Create Positive Clean Energy Districts (PCEDs) across Europe to reduce the impact of climate change.*
- ▶ Making cities smarter, more sustainable and more interactive.

Our Use Case

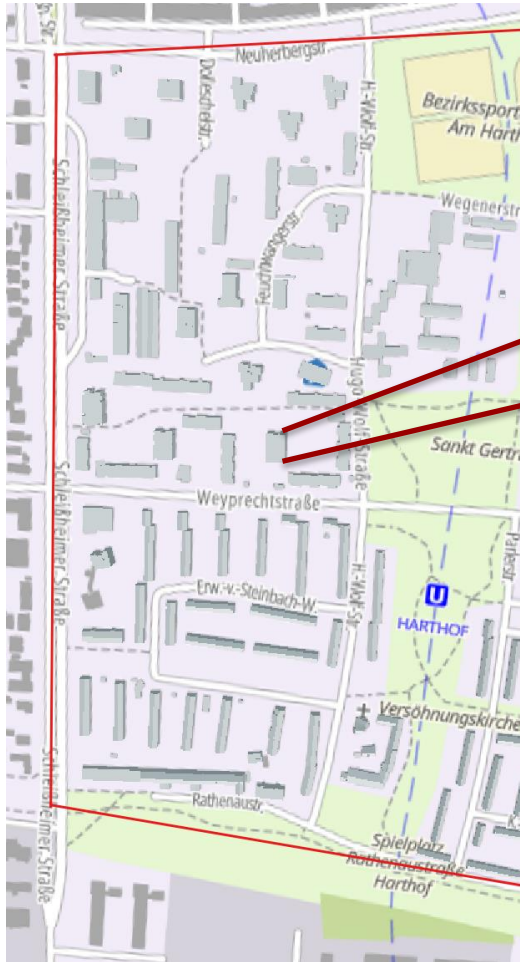


Harthof district, Munich, Germany

- ▶ **Study area:** Harthof district, Munich, Germany.
- ▶ **56 ha, 11500 inhabitants, 5500 residential units.**
- ▶ **Data Sources:**
 - CityGML LoD2 buildings stored in the **3DCityDB**.
 - Dynamic data: monthly heat demand, monthly solar irradiation potential, hourly electricity consumptions, stored in the **SensorThings API**.

Case Study

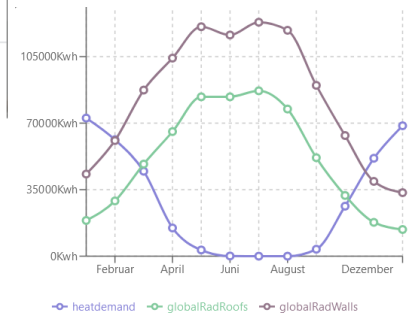
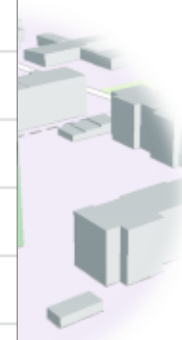
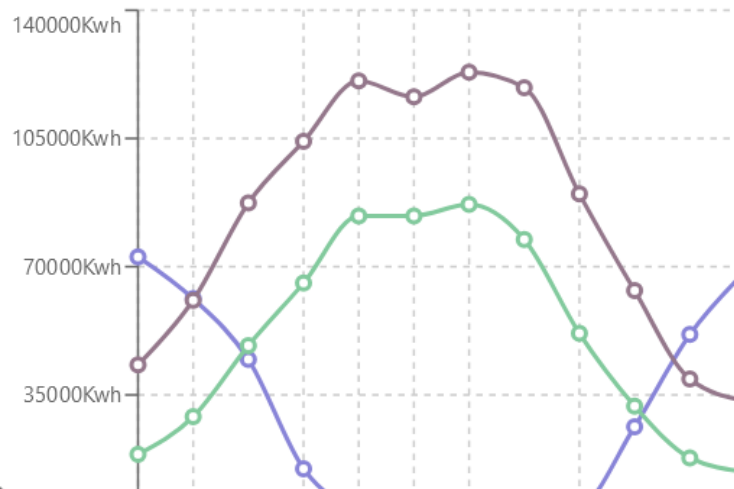
Attributes	Values
GMLID	DEBY_LOD2_4913221
building_usage	Wohnen
Year Built	1973
Building Height (m)	26.17
Street Name	Weyprechtstraße
House Number	25
Number of Inhabitants	138
Number of Building Units	54
Total Number of Floors	9



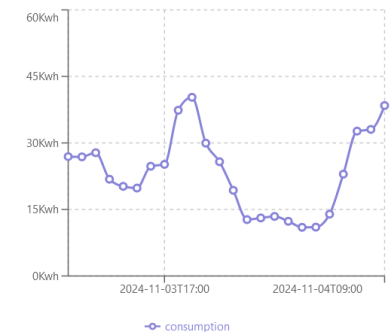
Harthof district, Munich

Yearly Energy Performance of the Building:

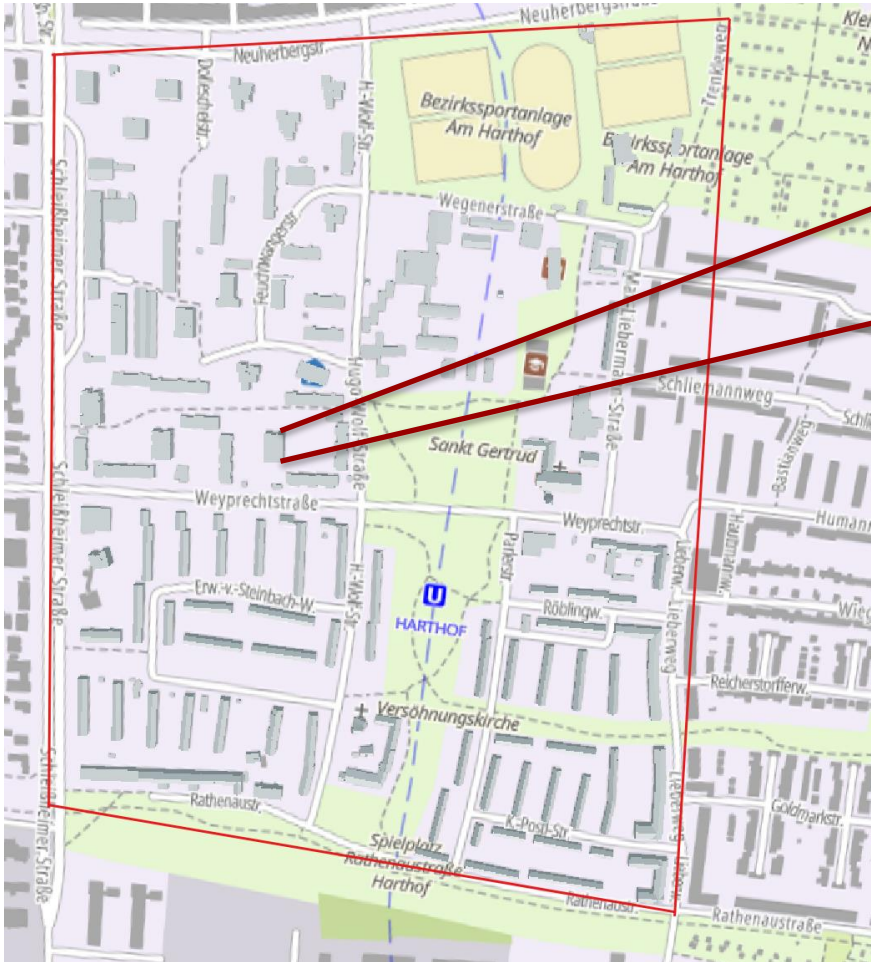
✓ Heat Demand and solar potential



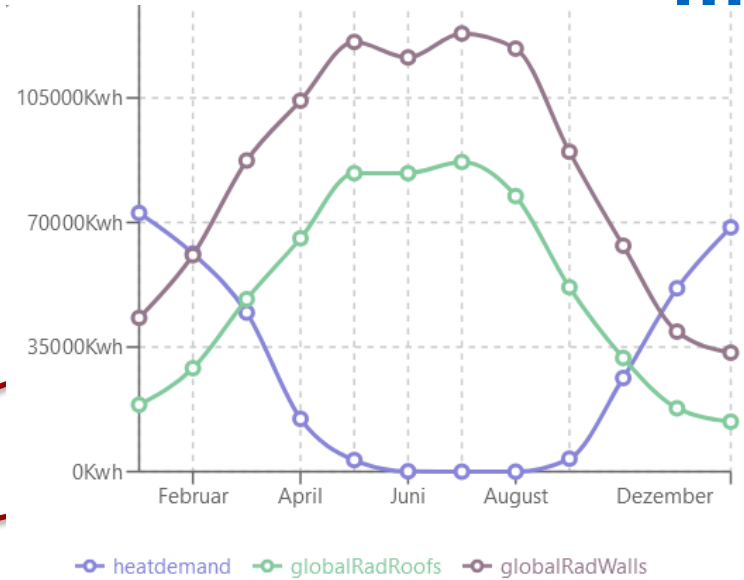
✓ Electricity Consumption for the last 24H



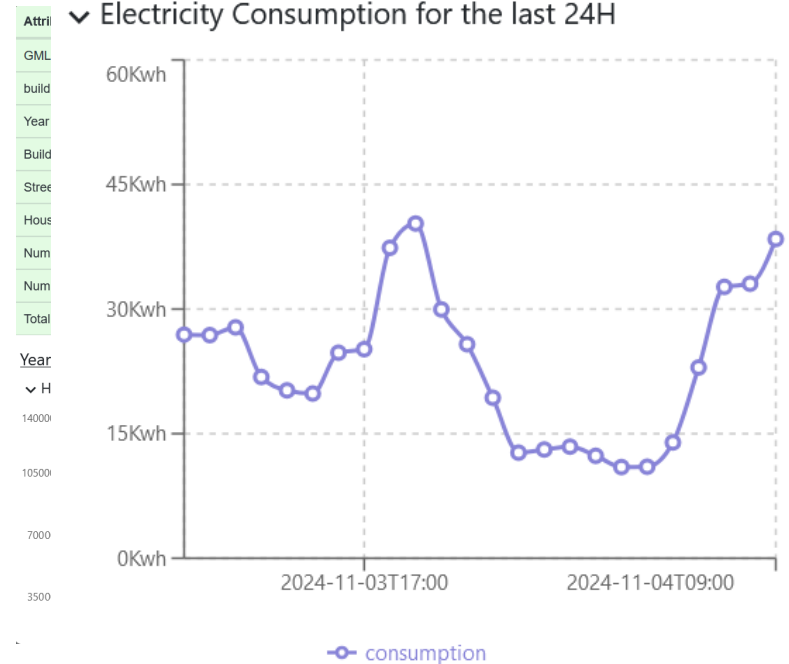
Case Study



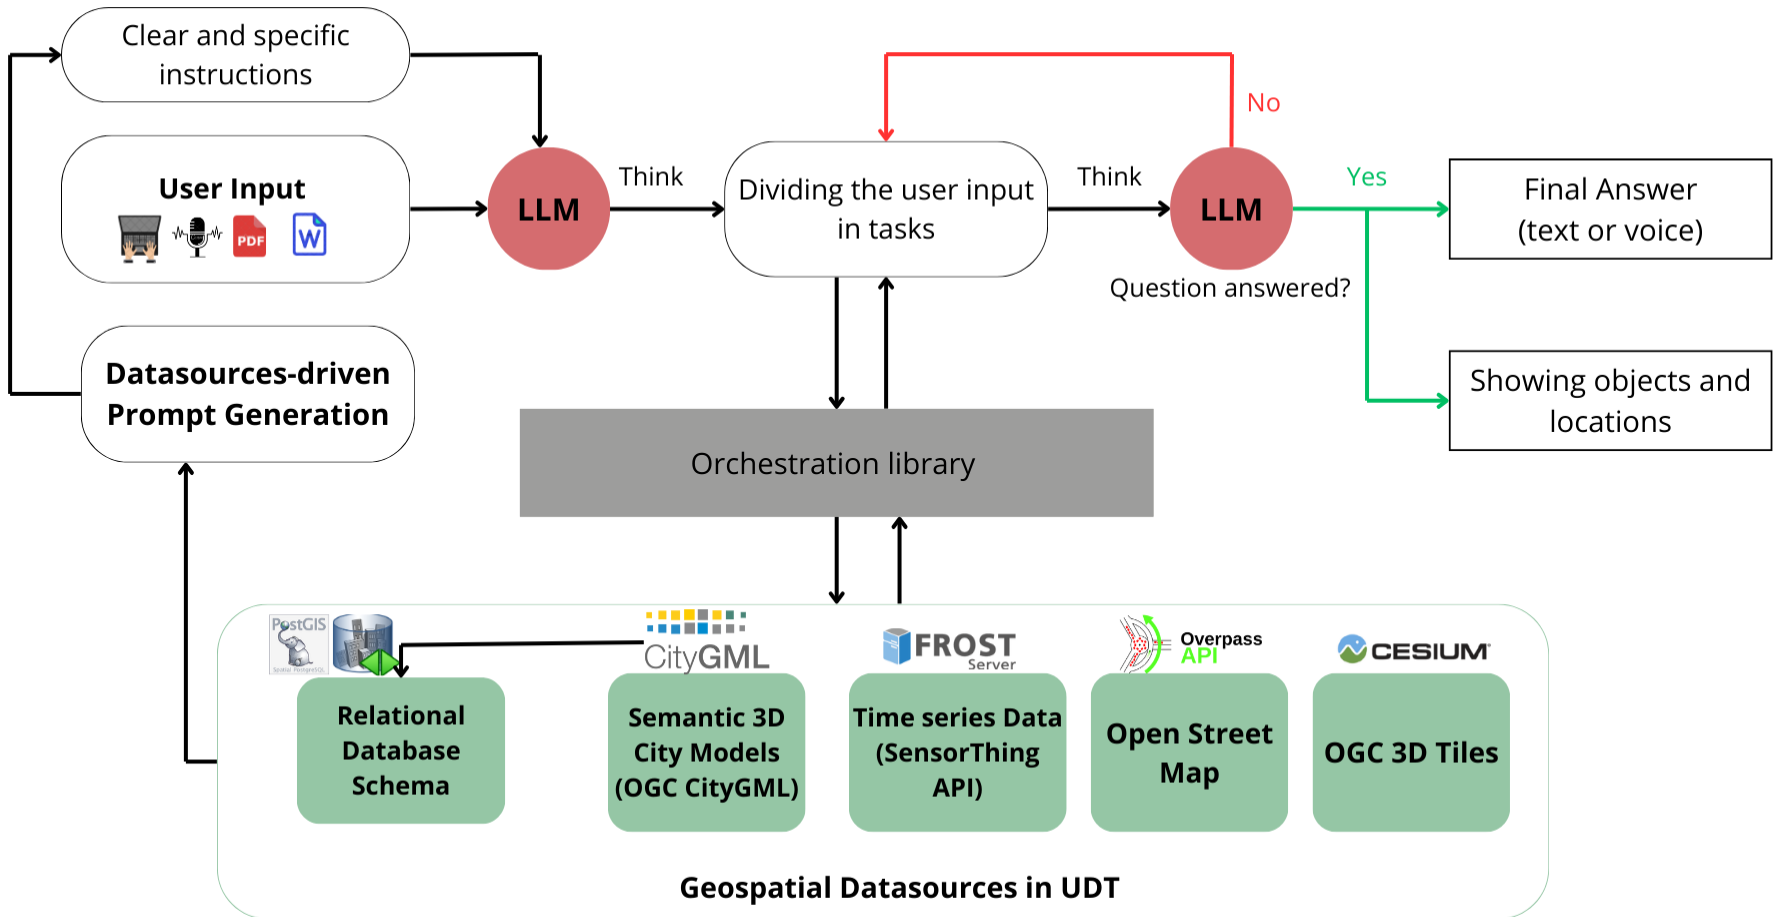
Harthof district, Munich, Germany



Electricity Consumption for the last 24H



(1) Talking to the Twin: Framework



(1) Talking to the Twin: Implementation



Who are you?



I'm your helpful assistant, here to provide information and answer questions about buildings in the Harthof district of Munich, based on the database structure you've provided. If you have any specific questions about these buildings or need assistance with related queries, feel free to ask!



How many buildings are in your database?



There are 314 buildings in the database.



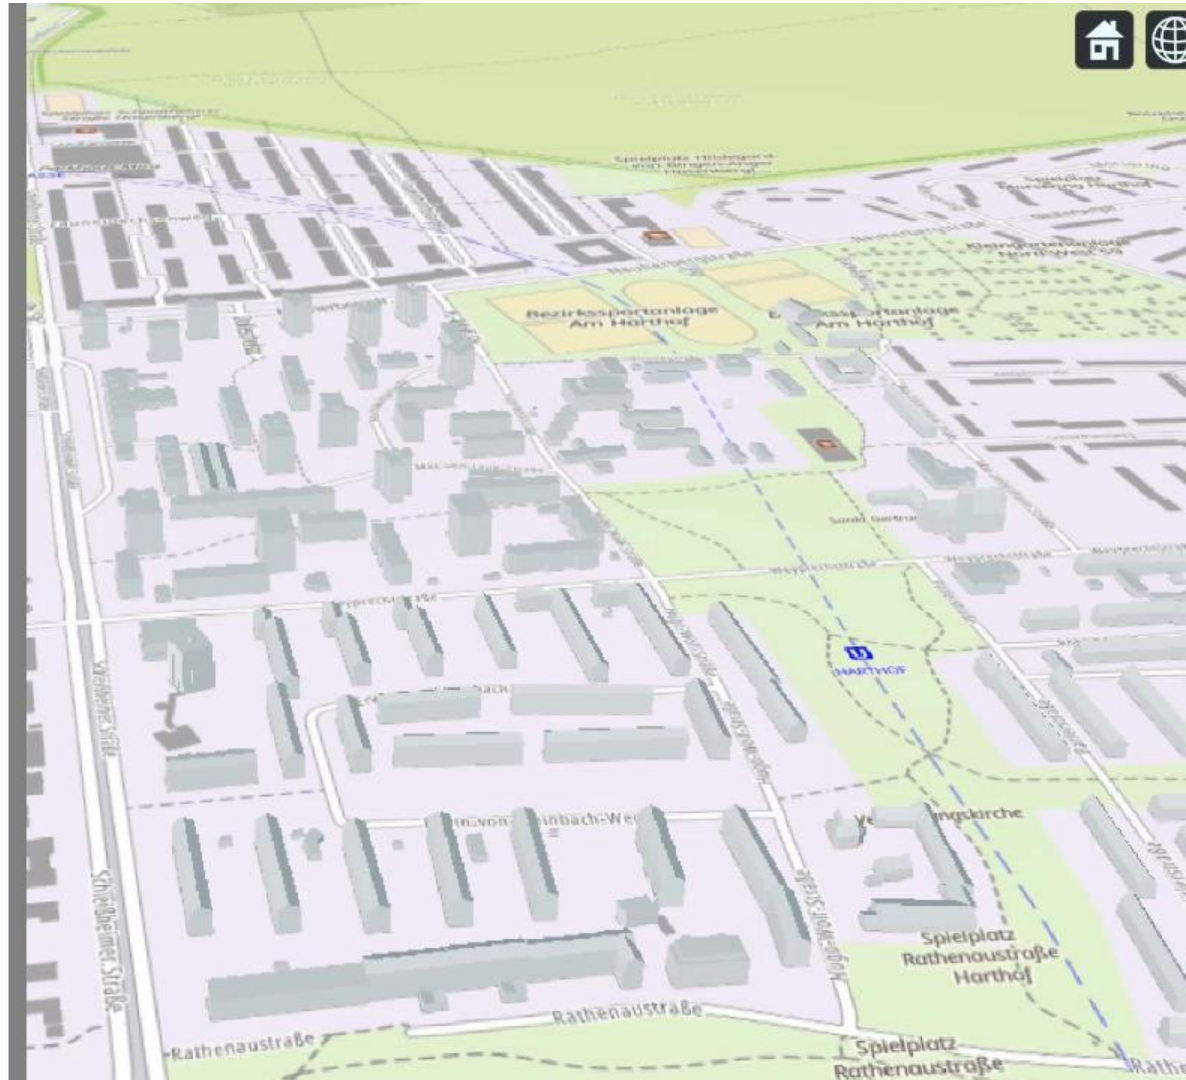
Who are the buildings' owners?



The buildings' owners in the database are:

- Münchner Wohnen
- LHM
- Kirche

Send a message



(1) Talking to the Twin: How it works

(1) Generate the SQL query from user input and the database schema



Show me all the buildings in Weyprechtstraße that have more than 7 storeys.



```
SELECT DISTINCT co. gmlid
FROM cityobject co
JOIN building b ON b.id = co.id
JOIN address_to_building ab ON b.id = ab.building_id
JOIN address a ON ab.address_id = a.id
WHERE a.name = "Weyprechtstraße"
AND co.storesys_above_ground > 7
AND co.objectclass_id = 26;
```

(1) Call run_sql(query) function
 output: [{"gmlid": "DEBY_LOD2_4913221"}, {"gmlid": "DEBY_LOD2_4913228"}]

(2) Call highlight_results(ids) function

81%

F1 with GPT-4o

on 71 multilingual queries from real stakeholders



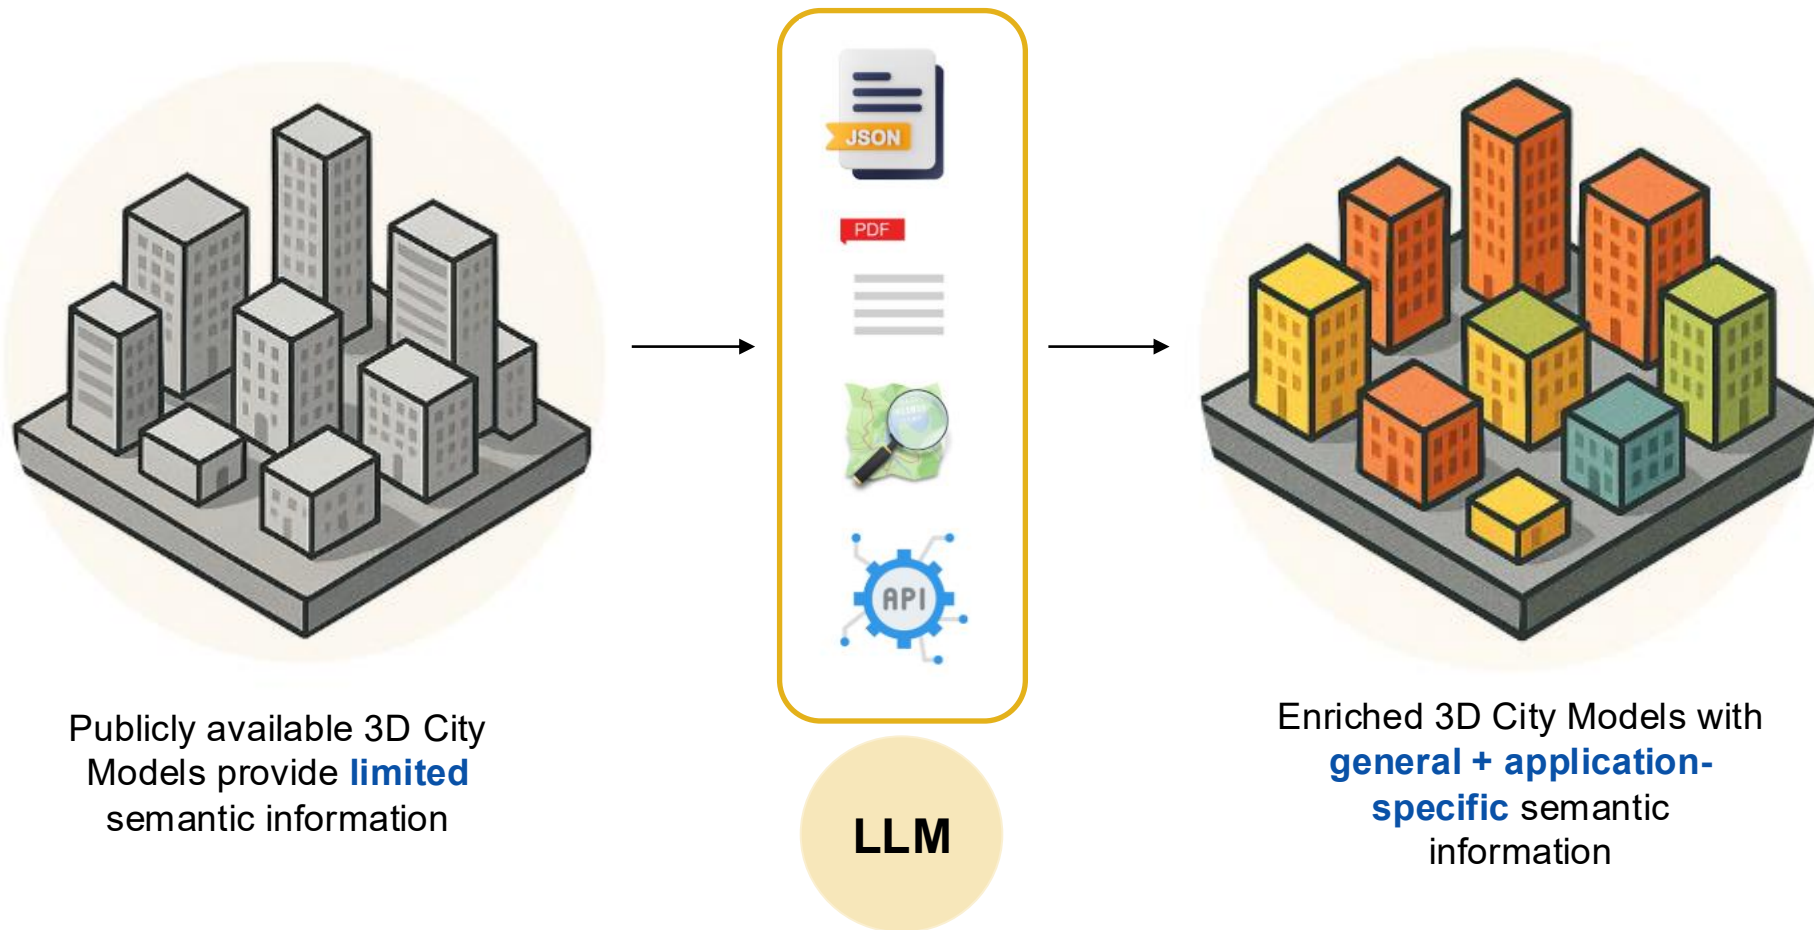
Show me all the buildings in Weyprechtstraße that have more than 7 storeys.



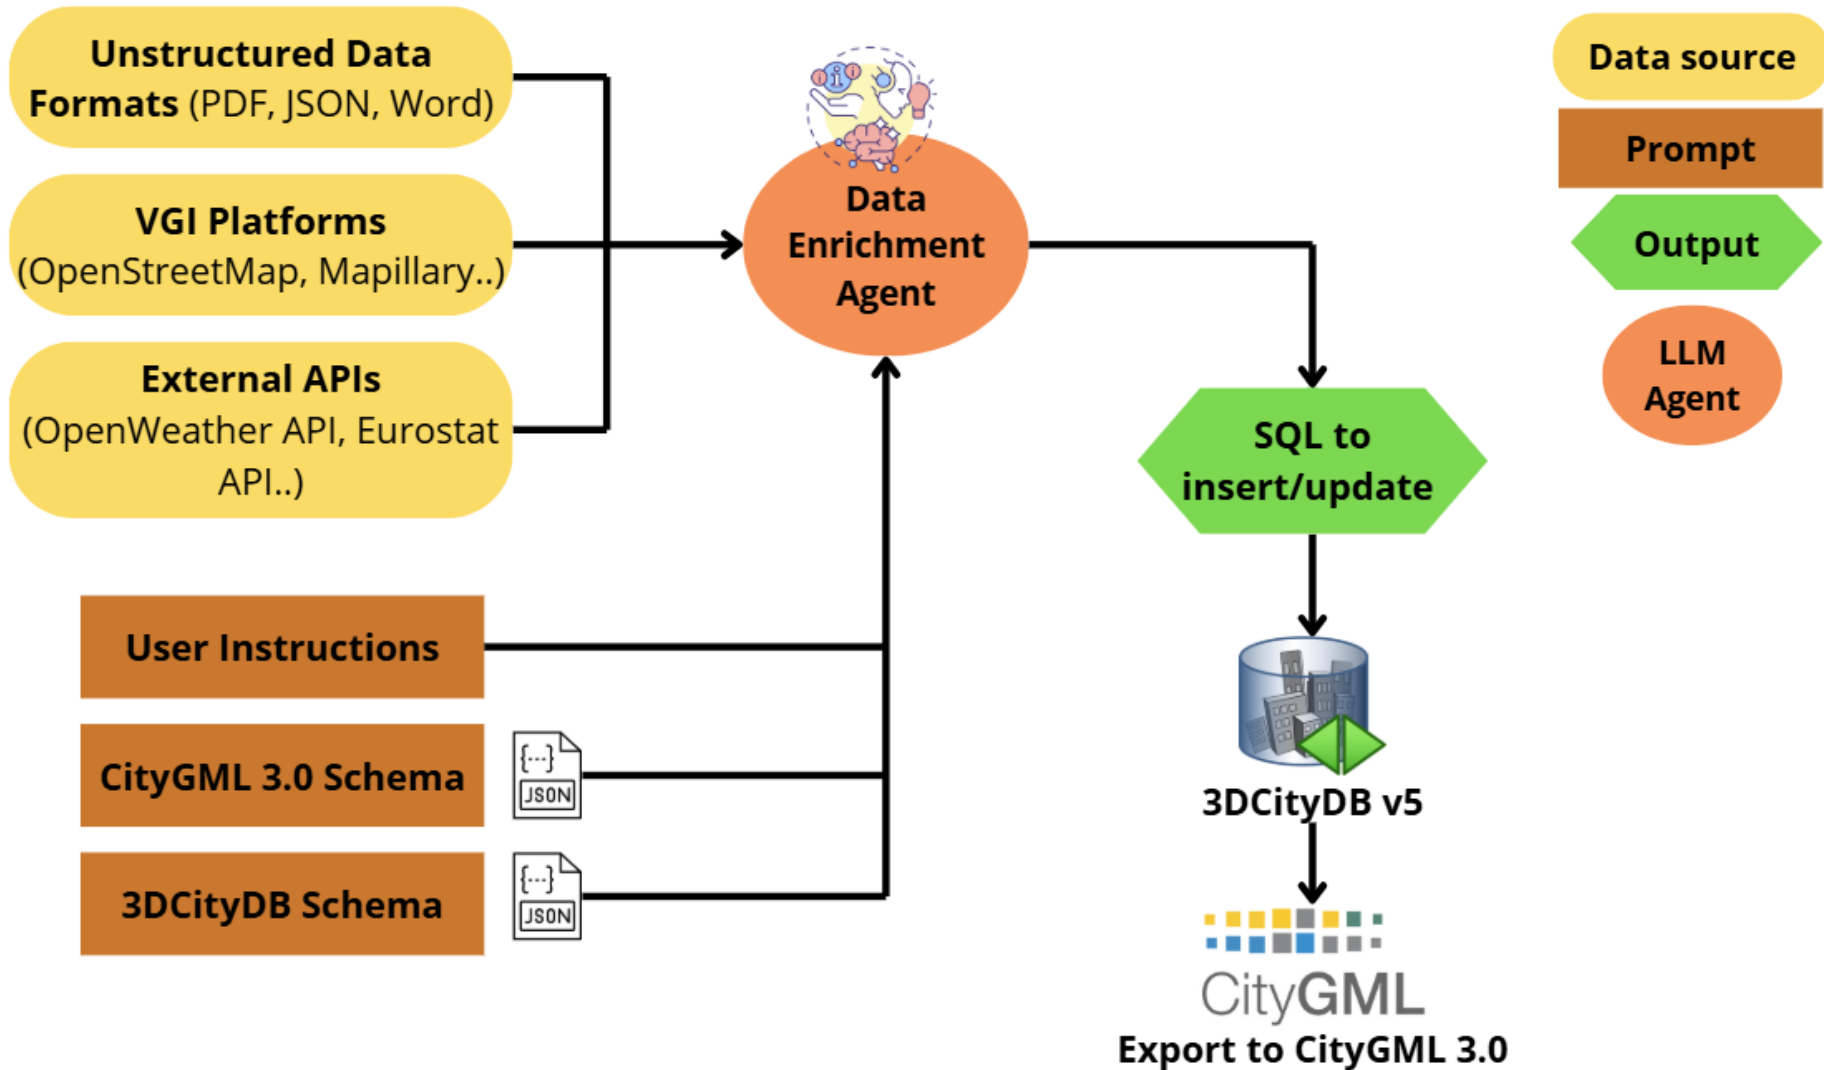
Here are the buildings in Weyprechtstraße with more than 7 storeys: [{"gmlid": "DEBY_LOD2_4913221"}] and [{"gmlid": "DEBY_LOD2_4913228"}]. You can now view these buildings.



(2) Enriching the Twin: Framework



(2) Enriching the Twin: Framework



(2) Enriching the Twin: Building Use Case using Energy Certificate

ENERGIEAUSWEIS für Wohngebäude

gemäß den §§ 16 ff. Energieeinsparverordnung (EnEV) vom ¹ 18. November 2013

Gültig bis: **08.06.2028**

Registriernummer ² BY-2018-001939337

1

Gebäude

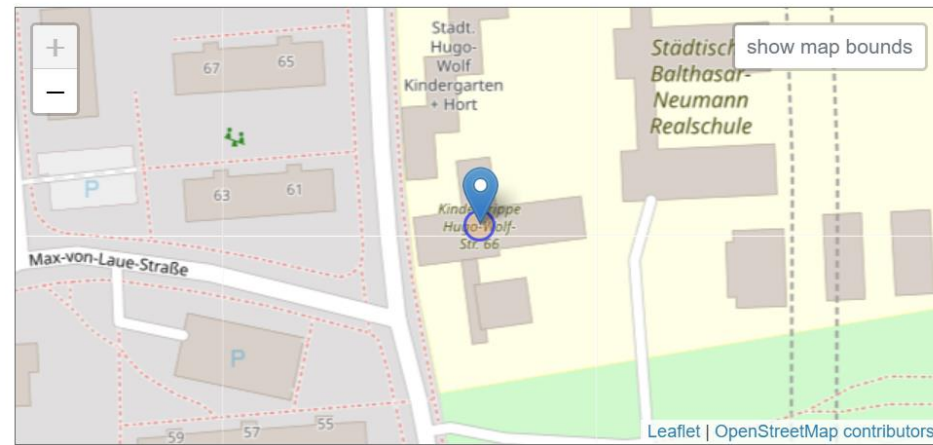
Gebäudetyp	freistehendes Mehrfamilienhaus	
Adresse	[REDACTED]	
Gebäudeteil	WE 0529 / GE 9	
Baujahr Gebäude ³	1972	
Baujahr Wärmeerzeuger ^{3,4}	2003	
Anzahl Wohnungen	24	
Gebäudenutzfläche (A _N)	1.581,4 m ²	<input type="checkbox"/> nach § 19 EnEV aus der Wohnfläche ermittelt
Wesentliche Energieträger für Heizung und Warmwasser ³	Kraft-Wärme-Kopplung, regenerativ	
Erneuerbare Energien	Art: Fernwärme SWM	Verwendung: Heizung und Warmwasser
Art der Lüftung / Kühlung	<input checked="" type="checkbox"/> Fensterlüftung <input type="checkbox"/> Lüftungsanlage mit Wärmerückgewinnung <input type="checkbox"/> Anlage zur Kühlung <input type="checkbox"/> Schachtlüftung <input type="checkbox"/> Lüftungsanlage ohne Wärmerückgewinnung	
Anlass der Ausstellung des Energieausweises	<input type="checkbox"/> Neubau <input type="checkbox"/> Modernisierung (Änderung / Erweiterung) <input type="checkbox"/> Sonstiges (freiwillig) <input checked="" type="checkbox"/> Vermietung / Verkauf	

► Energy Certificate of the building.

(2) Enriching the Twin: Building Use Case using OSM

Kinderkrippe Hugo-Wolf-Str. 66 [link to this page](#)

Name	Kinderkrippe Hugo-Wolf-Str. 66 (name)
Type	amenity:kindergarten
Last Updated	2024-11-05T15:45:19.588645+00:00
Search Rank	30
Address Rank	30 (house / building)
Importance	0.00008462272068355866
Coverage	Point
Centre Point (lat,lon)	48.206395,11.568172
OSM	node 1590583759
Place Id	117125866 (on this server)
Computed Postcode	80937 (how?)
Address Tags	München (city) Hugo-Wolf-Straße (street) DE (country) 80937 (postcode) 66 (house number)
Extra Tags	3 years 10 months (max_age) 2 months (min_age) yes (nursery) Landeshauptstadt München (operator) kita.hugowolfstr.66@muenchen.de (contact:email) +49 (0)89 31215967 (contact:phone) Mo-Th 06:30-17:00; Fr 06:30-16:00 (opening_hours) https://stadt.muenchen.de/service/info/hugo-wolf-str-66/1050559/ (contact:website)



- ▶ Example of available OSM data for this building.

(2) Enriching the Twin: Building Use Case - Result

id	A-Z name	0-9 val_double	A-Z val_string	val_timestamp	A-Z val_uri
57	GMLID	[NULL]	DEBY LOD2 4967557	[NULL]	[NULL]
58	PV_state	[NULL]	n/a - RBS	[NULL]	[NULL]
59	building_gross_floor_area	602	[NULL]	[NULL]	[NULL]
61	building_usage	[NULL]	Daycare	[NULL]	[NULL]
63	district_heating	[NULL]	not connected	[NULL]	[NULL]
66	ownership_type	[NULL]	[NULL]	[NULL]	[NULL]
67	refurbishment_state	[NULL]	[NULL]	[NULL]	[NULL]
68	refurbishment_type	[NULL]	[NULL]	[NULL]	[NULL]
69	usable_and_living_area	[NULL]	[NULL]	[NULL]	[NULL]
71	yearly_ENEV	[NULL]	[NULL]	[NULL]	[NULL]
72	dateOfConstruction	[NULL]	[NULL]	1:00:00.000 +0100	[NULL]
78	function	[NULL]	[NULL]	[NULL]	[NULL]
19,912	contact_phone	[NULL]	[NULL]	[NULL]	[NULL]
19,913	contact_website	[NULL]	[NULL]	[NULL]	https://stadt.muenchen.de/ser
19,914	contact_email	[NULL]	kita.nugowollstr.00@muenchen.de	[NULL]	[NULL]

78 - 82%

Accuracy with GPT-4o

On 20 buildings

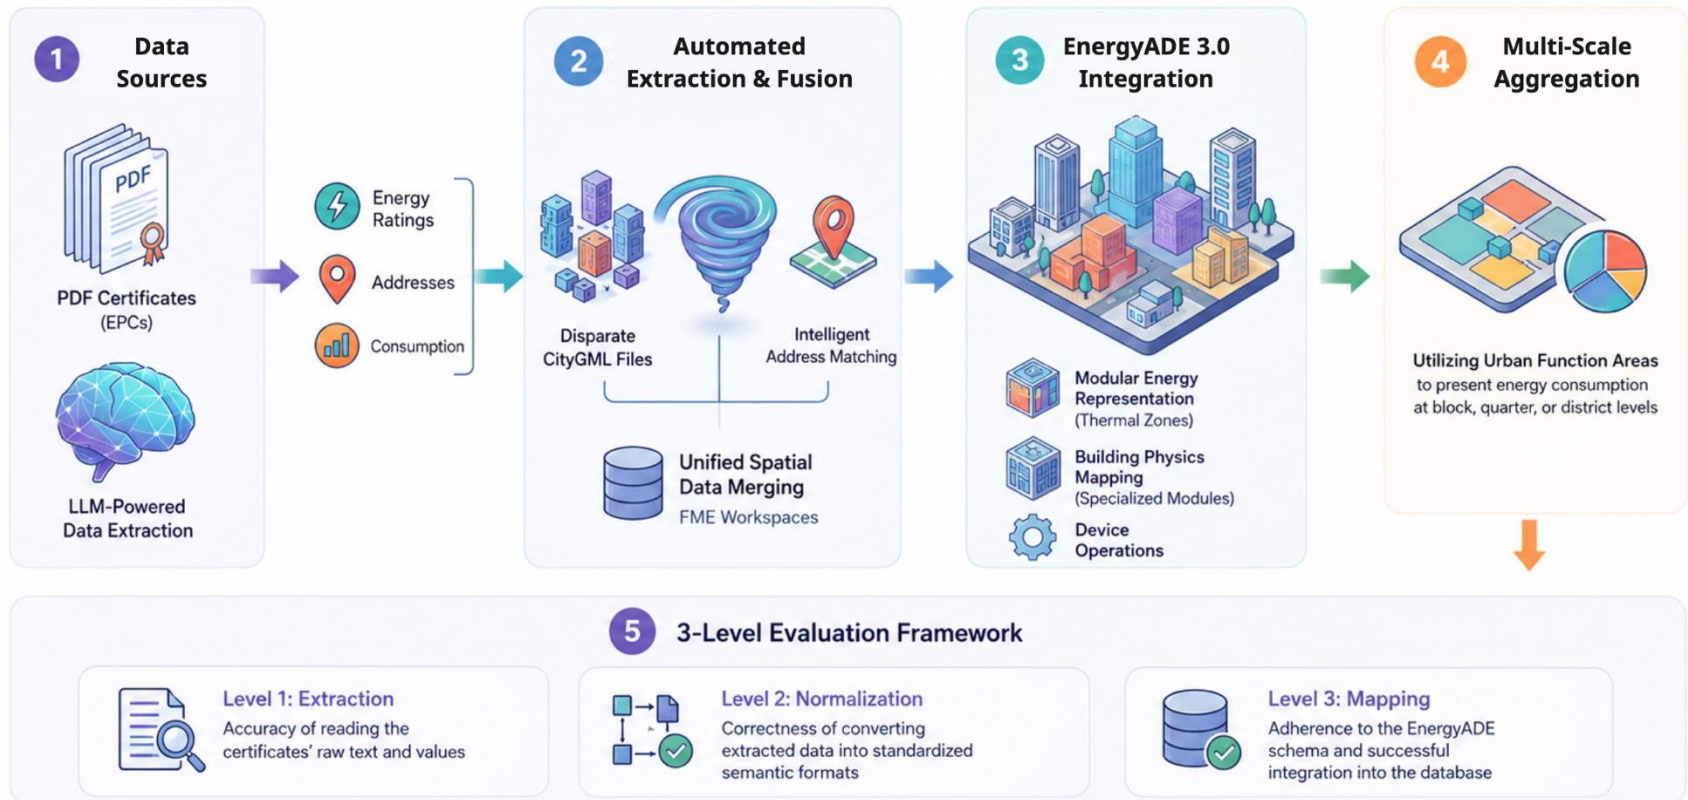
Energy
Indicate
(DF)

SM

- ▶ The **mapping** between **3DCityDB** and the external data sources is done through the **address**.

(2) Enriching the Twin: an on-going Master Thesis

AI-Driven Energy Data Integration: From EPCs to 3D City Models



(2) Enriching the Twin: an on-going Master Thesis

SELECT nbp.* FROM citydb.ng3_building_partitic | Enter a SQL expression to filter results (use Ctrl+Space)

id	objectclass_id	infiltration_rate	infiltration_rate_uom	is_heated	num_of_building_units	heat_capacity
657,483	11,081	0.7	1/h	1	45	[NULL]

Gebäude

Gebäudetyp	freistehendes Mehrfamilienhaus	
Adresse	[REDACTED]	
Gebäudeteil	WE 0529 / GE 1	
Baujahr Gebäude ³	1965	
Baujahr Wärmeerzeuger ^{3,4}	2004	
Anzahl Wohnungen	45	
Gebäudenutzfläche (A _N)	3.362,3 m ² <input type="checkbox"/> nach § 19 EnEV aus der Wohnfläche ermittelt	
Wesentliche Energieträger für Heizung und Warmwasser ³	Kraft-Wärme-Kopplung, regenerativ	
Erneuerbare Energien	Art: Fernwärme SWM	Verwendung: Heizung und Warmwasser
Art der Lüftung / Kühlung	<input checked="" type="checkbox"/> Fensterlüftung <input type="checkbox"/> Lüftungsanlage mit Wärmerückgewinnung <input type="checkbox"/> Anlage zur Schachtlüftung <input type="checkbox"/> Lüftungsanlage ohne Wärmerückgewinnung <input type="checkbox"/> Anlage zur Kühlung	
Anlass der Ausstellung des Energieausweises	<input type="checkbox"/> Neubau <input type="checkbox"/> Modernisierung (Änderung / Erweiterung) <input type="checkbox"/> Sonstiges (freiwillig)	
	<input checked="" type="checkbox"/> Vermietung / Verkauf	

4.2 Lüftungsverluste

Lüftungswärmeverluste	n = 0,70 h ⁻¹	1996,04 W/K	32,3 %
-----------------------	--------------------------	-------------	--------

4.1 spezifische Transmissionswärmeverluste (Fortsetzung)

Nr.	Bauteil	Orientierung Neigung	Fläche A m ²	U _t -Wert W/(m ² K)	Faktor F _x	F _x * U * A	
						W/K	%
10	Rollladenkasten N	N 90,0°	0,49	3,000	1,00	1,46	0,0
11	Rollladenkasten O	O 90,0°	1,18	3,000	1,00	3,54	0,1
12	Rollladenkasten S	S 90,0°	2,02	3,000	1,00	6,05	0,1
13	Rollladenkasten W	W 90,0°	1,73	3,000	1,00	5,19	0,1
			ΣA =	3737,96		Σ(F _x * U * A) =	3802,95

Wärmebrückenzuschlag ΔU

ΔU _{VB} =	0,10 W/(m ² K)	ΔU _{VB} * A =	373,80 W/K	6,1 %
--------------------	---------------------------	------------------------	------------	-------

Show SQL | Enter a SQL expression to filter results (use Ctrl+Space)

id	AZ type	AZ type_codespace	is_protected
577,050	Wohngebäude	[NULL]	[NULL]

id	objectclass_id	u_value	u_value_uom
657,482	11,051	0.1	W/(K*m^2)

Get in contact with **Mohammed Malaik Hashmi** for more info!

What both frameworks need

Talk to

the twin

How does a non-expert query the UDT in natural language?

Enrich

the twin

How does the twin enrich itself from external sources?

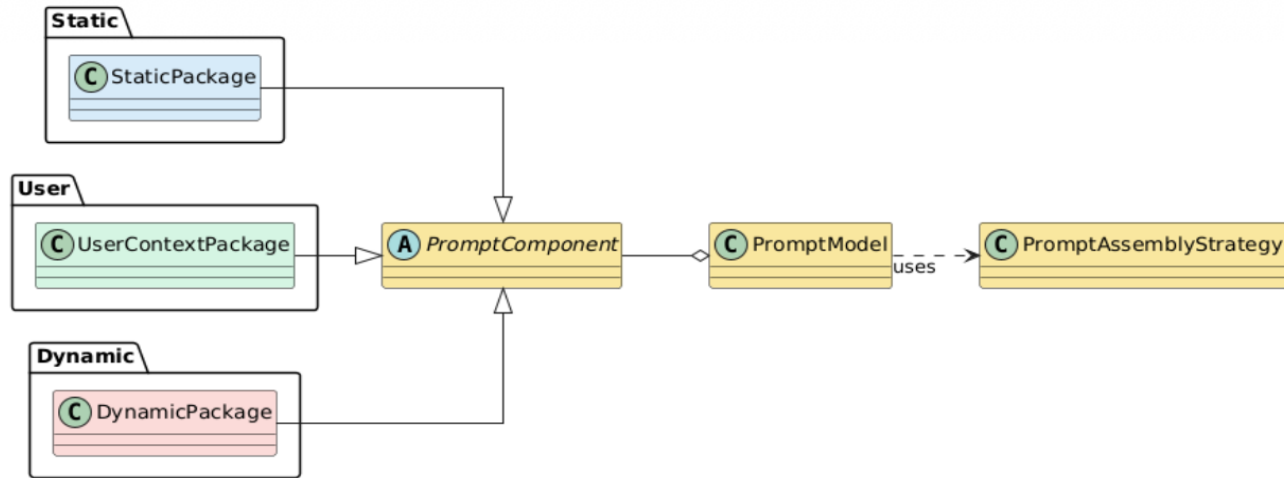


A manually written Prompt !

Issues?

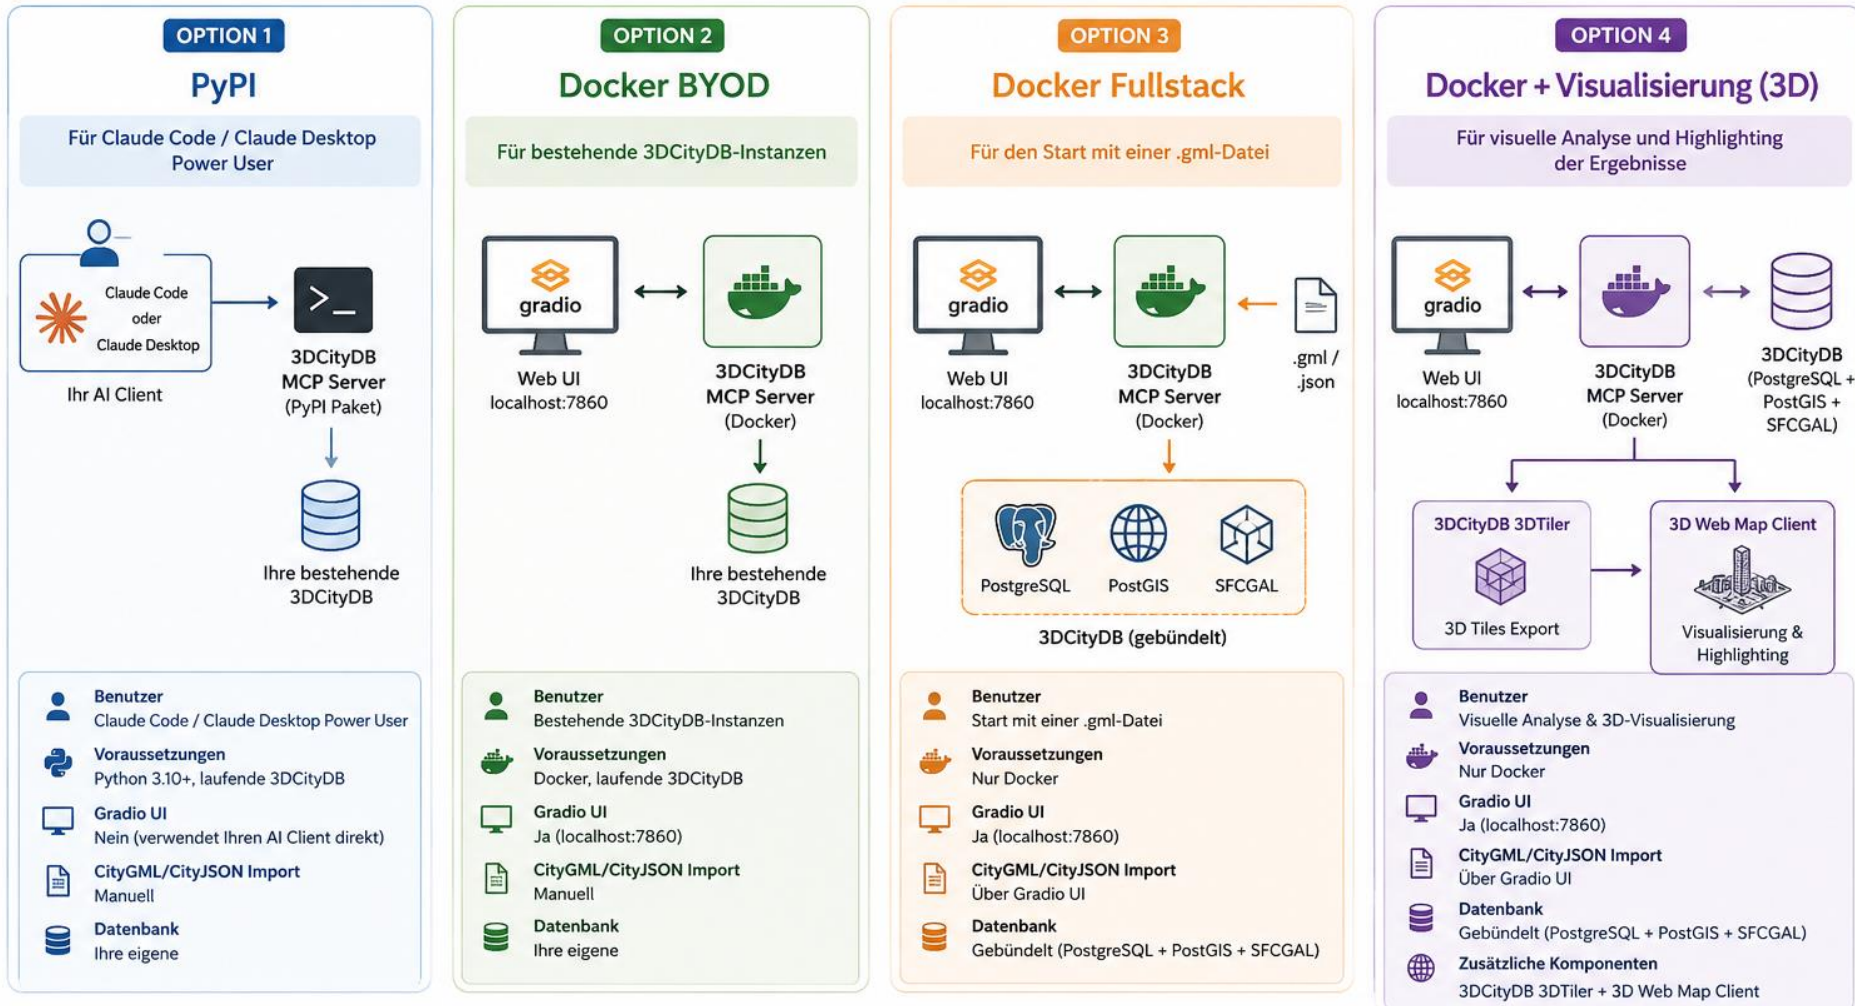
- The prompt can only be written by an **expert**
- The prompt needs to be **maintained** and **updated** each time the data in the UDT changes.

(3) Scale across Twins: Framework



- ▶ There is an **MCP Server for OGC SensorThings API**. But none for **CityGML**.
- ▶ The challenge lies in the scale and variability of this knowledge: full CityGML schema defines **hundreds of classes** and **thousands of properties**, yet any given dataset instantiates **only a subset**
- ▶ Our framework addresses this by treating the 3DCityDB as both **data store and metadata source**, automatically scanning its structure at runtime.

3DCityDB MCP Server Implementierung



3DCityDB MCP Server : Demo

3DCityDB-MCP

localhost:7860

full

Context window (Ollama)

Tokens available to the model.
128K recommended for complex queries.

32K (32,768)

New conversation

Context window

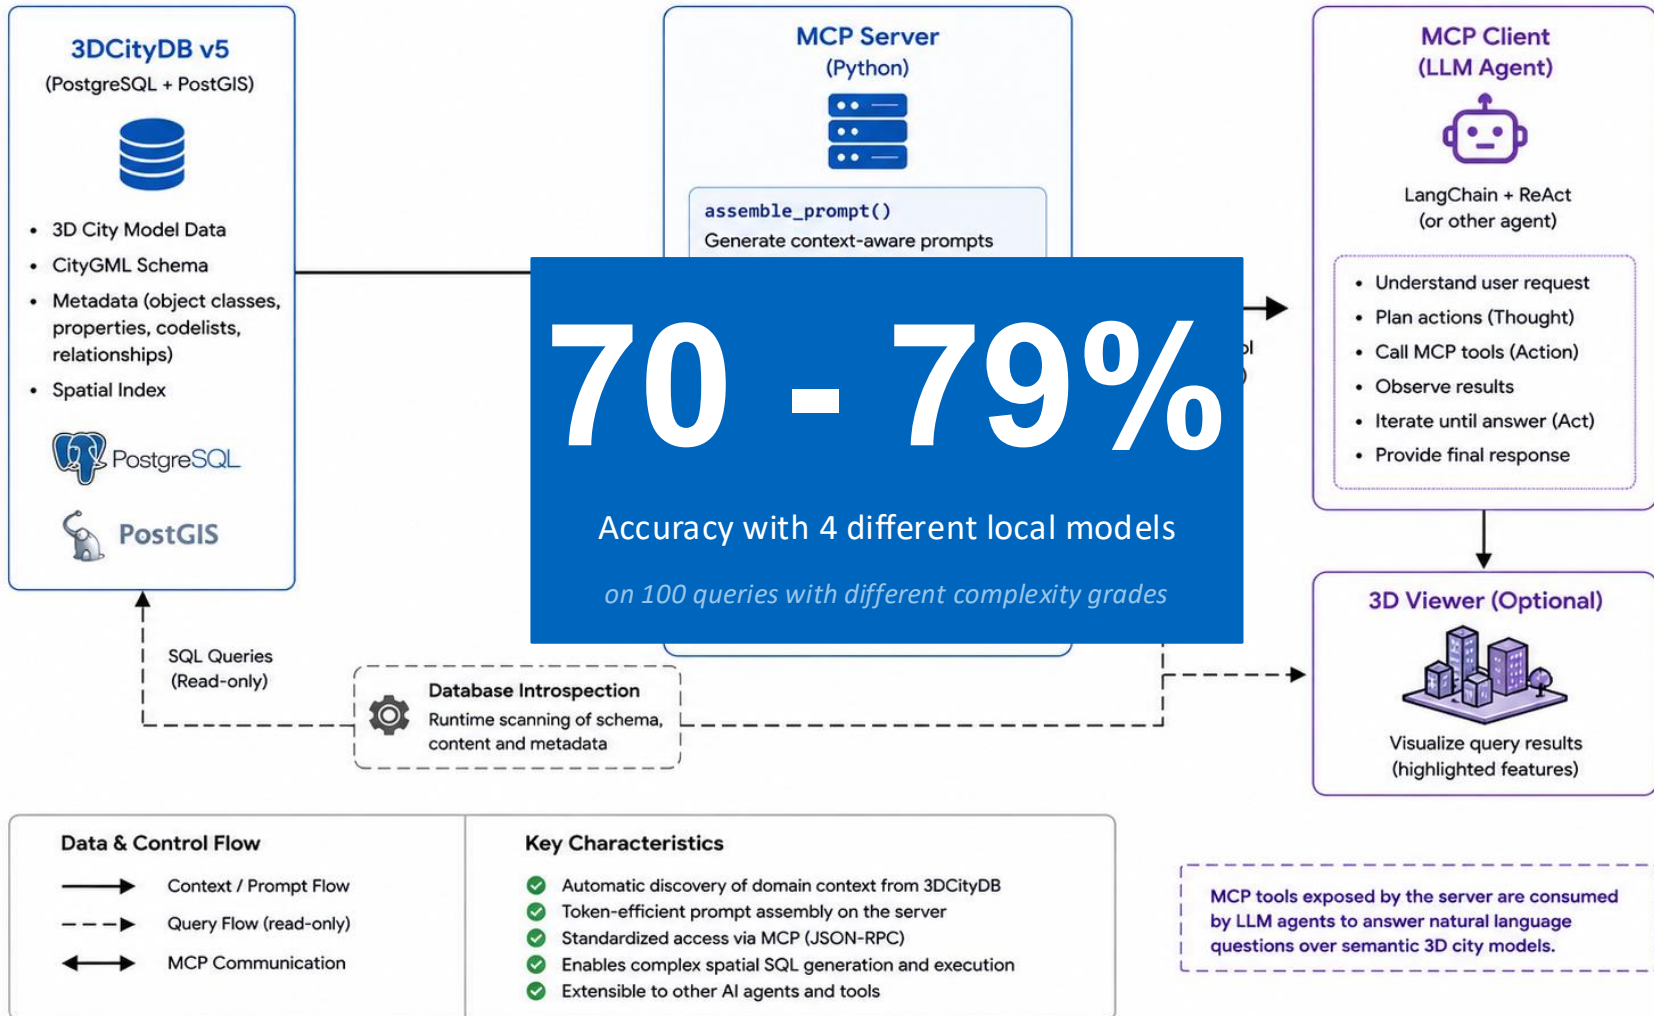
5,407 / 32,768 tokens (17%)

DB MCP Provider full
server (ollama) (auto)

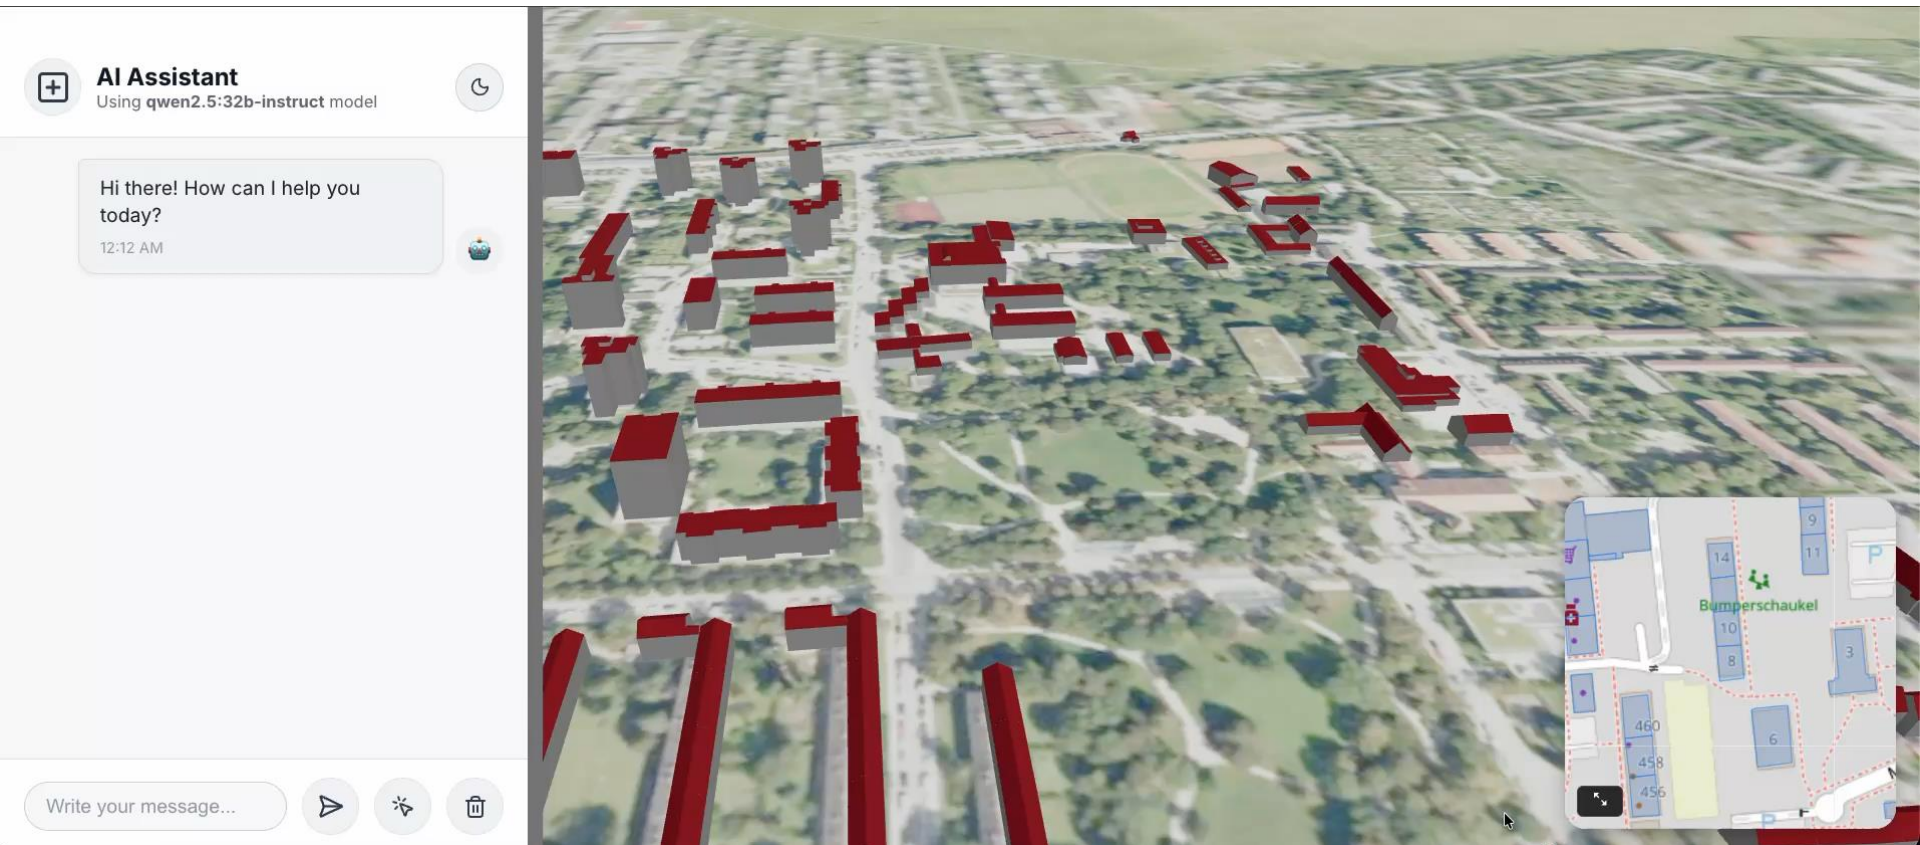
API keys are read from environment variables:
`ANTHROPIC_API_KEY`,
`OPENAI_API_KEY`,
`OLLAMA_BASE_URL`

Mit Gradio erstellt · Einstellungen

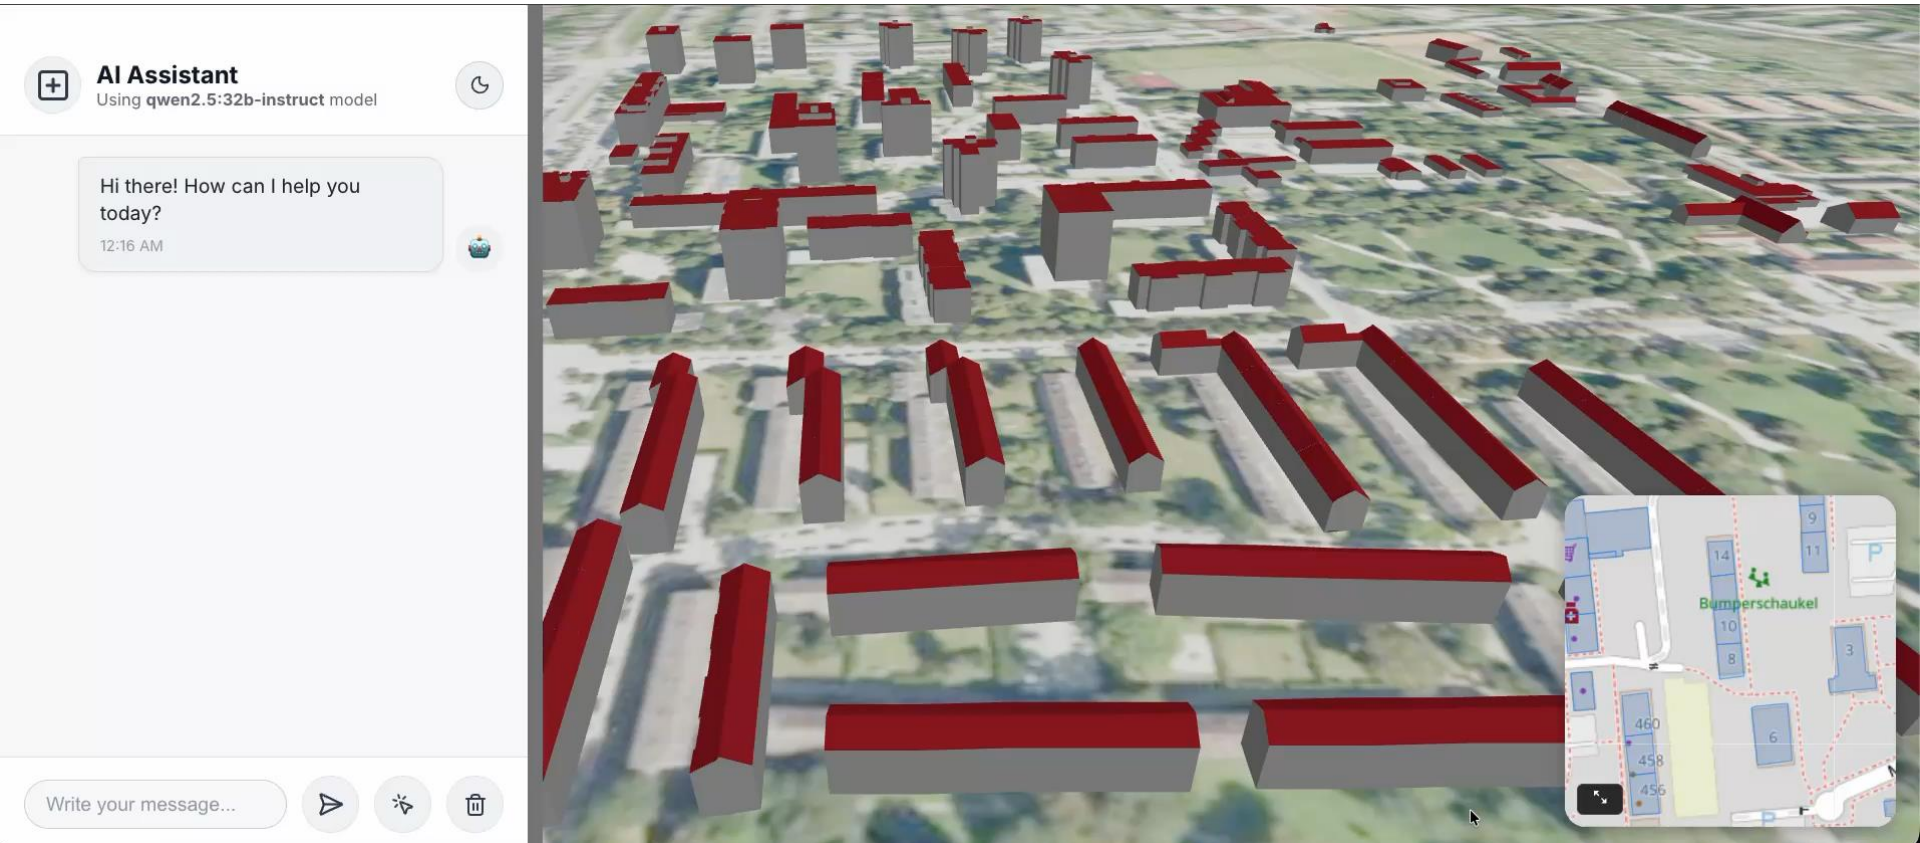
(3) Scale across Twins: Implementation



(4) Act on the Twin: Scenario generation for Planning tasks



(4) Act on the Twin: Scenario generation for Planning tasks



Conclusion

- ▶ Urban Digital Twins are **rich**, but their **data models are a challenge**. For most stakeholders, the **understanding of the data model** *is* the bottleneck.
- ▶ **LLMs already know programming, including SQL**. What they lack is structured, dataset-specific knowledge of UDT.
- ▶ With **dynamic and solid knowledge** of the UDT, the LLMs can **interact, enrich** and **act** on the UDT.