

# SIGMUS: Semantic Integration for Knowledge Graphs in Multimodal Urban Spaces

Brian Wang  
University of California, Los Angeles  
Los Angeles, California, USA  
wangbri1@g.ucla.edu

Mani Srivastava\*  
mbs@ucla.edu  
University of California, Los Angeles  
Los Angeles, California, USA

## Abstract

Modern urban spaces are equipped with an increasingly diverse set of sensors, all producing an abundance of multimodal data. Such multimodal data can be used to identify and reason about important incidents occurring in urban landscapes, such as major emergencies, cultural and social events, as well as natural disasters. However, such data may be fragmented over several sources and difficult to integrate due to the reliance on human-driven reasoning for identifying relationships between the multimodal data corresponding to an incident, as well as understanding the different components which define an incident. Such relationships and components are critical to identifying the causes of such incidents, as well as producing forecasting the scale and intensity of future incidents as they begin to develop. In this work, we create *SIGMUS*, a system for **Semantic Integration for Knowledge Graphs in Multimodal Urban Spaces**. *SIGMUS* uses Large Language Models (LLMs) to produce the necessary world knowledge for identifying relationships between incidents occurring in urban spaces and data from different modalities, allowing us to organize evidence and observations relevant to an incident without relying on human-encoded rules for relating multimodal sensory data with incidents. This organized knowledge is represented as a knowledge graph, organizing incidents, observations, and much more. We find that our system is able to produce reasonable connections between 5 different data sources (new article text, CCTV images, air quality, weather, and traffic measurements) and relevant incidents occurring at the same time and location.

## CCS Concepts

• **Information systems** → *Spatial-temporal systems*; • **Computing methodologies** → **Spatial and physical reasoning**; *Ontology engineering*.

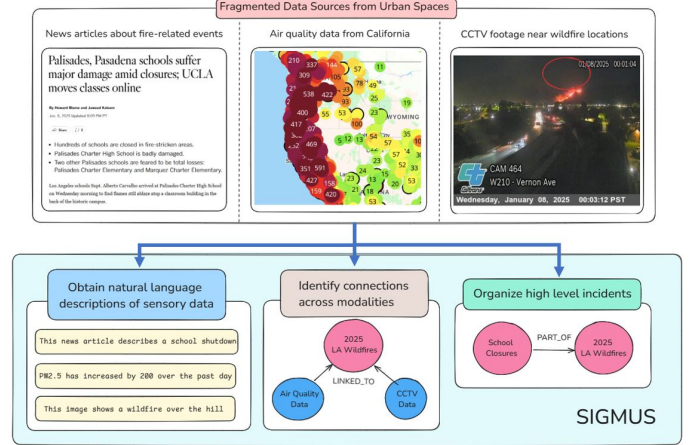
## Keywords

Urban Computing, Large Language Model, Knowledge Graphs, Multimodal Learning

## 1 Introduction

The proliferation of digital sensors embedded in modern urban infrastructure, coupled with increasingly efficient and robust connectivity technologies has introduced a growing ecosystem of publicly accessible data streams - from traffic cameras to weather stations, environmental monitors, and real-time news feeds. This ecosystem is further augmented both by human observations made through

\*The author holds concurrent appointments as Amazon Scholar and Professor at UCLA, but the work in this paper is not associated with Amazon.



**Figure 1: High level overview of *SIGMUS*, which aims to identify incidents in urban spaces and draw connections to sensory data collected in those spaces.**

community-contributed information on social platforms as well as crowd-sourced environmental monitoring through mobile devices. However, despite the abundance of such rich multimodal data in urban spaces, these data streams often reside in fragmented, siloed systems that lack mechanisms for integrated analysis and reasoning across data streams. As a result, this ecosystem remains heavily underutilized and lacks tools for understanding and monitoring major phenomena that require analysis across various sources and modalities.

One example of such major phenomena was the catastrophic Palisades Fire (and other concurrent wildfires), which devastated parts of Los Angeles in January 2025. Understanding the full impact of the wildfires requires an analysis across disparate data sources. CCTV data, air quality measurements, and news reports each provide a unique and critical insight into how the city may be affected. Such insight is valuable to emergency service providers, policymakers, health agencies, and future urban planners. However, this full potential can only be achieved when information across data sources is unified and organized in a human-readable manner.

By unifying disparate data sources, we enable urban knowledge systems to fuse independent reports across different modalities into cohesive, contextualized incidents that can be readily visualized. This allows for retroactive analysis of important city events. Furthermore, as many of these data sources publish live data, it becomes possible to create a live knowledge system suited for monitoring and early detection of incidents occurring in urban environments.

Ideally, a live knowledge system for analysis and monitoring of urban environments should describe relationships between data

from heterogeneous sources. Different sources may report data which is related by time, space, as well as the incident itself. Thus, to overcome the fragmented nature of many data sources, we propose using a knowledge graph to unify and integrate data.

However, many of these relationships within and between data sources may be challenging to infer, as they may be conditioned on the context of the data as well as some degree of world knowledge. For example, identifying that an air quality trend is related to the report of a house fire depends on both the context (e.g. time and space), but also the physics of the world (e.g. a fire should reduce air quality). Importantly, such inferences require both world knowledge and a degree of fuzzy reasoning.

Recently, Large Language Models (LLMs) have been shown to demonstrate significant world knowledge about sensory data [45, 47]. We hypothesize that LLMs, when used to reason about incidents and sensory data, provide valuable insights into the emergence of real-world incidents and their detection through sensory data. Importantly, the ability to generate meaningful relationships between incidents and different sensory sources (e.g. CCTV, traffic measurements, news articles) requires a significant degree of world knowledge. In this work, we show the effectiveness of LLMs at the qualitative reasoning required to organize and create relationships between incidents and sensory data.

The main contribution of this work, *SIGMUS*, is a system for ingesting data from a variety of data sources (across text, image and tabular data), and integrating multimodal information with identified real-world incidents. Figure 1 showcases the main objectives of our system, allowing for readable information about sensory data and associated incidents to be stored in a knowledge graph. We provide representative examples to illustrate *SIGMUS*'s capacity to identify incidents from multimodal inputs, organize them into a hierarchy, and establish semantic connections between sensory data and incidents using a knowledge graph. We also study the latency of processing multimodal information and storing it on the knowledge graph, and discuss the design decisions in our system leading aiming to lead to more efficient use of LLMs. As a case study, we examine the 2025 Los Angeles Wildfires and showcase some of the relationships identified by different incidents and their associated multimodal information.

## 2 Related Work

### 2.1 Multimodal Knowledge Graphs

Existing works on multimodal knowledge graphs typically involve both semi-structured text and complex unstructured data such as images. [35] focuses on extracting text, images, and source code from academic publications and creating their respective graphs. They apply a separate processing pipeline for each modality and attempt to align different entities into the Computer Science Ontology (CSO) [5]. [43] takes an existing knowledge graph (ConceptNet [38]) and annotates it with additional attributes for images, sound, and video. [42] describes an ontology for structuring knowledge collected by sensors in smart cities. However, it lacks a concrete implementation and doesn't collect event information necessary for smart city applications.

### 2.2 Multimodal Data Integration

The concept of unifying data sources which are diverse in type into a coherent view for monitoring or prediction has also been explored. In the web domain, a **sensor web** [24] describes an infrastructure that enables discovery and integration of sensors and their data. An example of such a system is Senseweb [31], which is a web platform that ingests data from different sensory data streams and arranges them into a map view. Some tools, such as Graph of Things [36] describe an idea similar to ours, where they aim to construct a live knowledge graph of various IoT and social media data, and creating connections between heterogeneous data sources using several ontologies. While these platforms and tools share a similar goal to ours, they lack the ability to produce inferences and correlations between information which has historically required human intervention to directly process such correlations or create flexible pipelines to represent such reasoning. It's also worth mentioning that while multimodal data fusion also brings together data from different sources, its objective is often task-specific and typically doesn't use human knowledge as explicitly as data integration might (such as the use of an ontology or linking against an external knowledge graph). Thus we won't be including related work from this field.

### 2.3 LLMs for Knowledge Graphs

There are variety of works which utilize Large Language Models (LLMs) in various roles for interacting with knowledge graphs. **LLMs for constructing knowledge graphs**

[33] iteratively creates a semantic map of events occurring across news documents, and uses the LLM to fill out actors, relationships, and states given the raw text document. [30] attempts to create an ontology from a document of terms, using some internal knowledge of the world to decide taxonomic and non-taxonomic relations. [49] provides knowledge graphs as answers to users questions by searching for relevant information and organizing it into facts represented as triples.

**LLMs expanding knowledge graphs with new facts and verifying existing facts**

[46] uses an LLM in a Retrieval-Augmented-Generation (RAG) style to query a knowledge graph and answer a user query. In particular, it converts a natural language question into the appropriate query language, and passes the response as part of the context into the LLM for question answering. [41] uses additional text from online sources to fill out empty fields in some given knowledge templates. [48] uses an existing ontology and seeks to extend it with new relationships between entities given a collection of raw documents.

[22] uses ground truth documents in order to verify triples found in knowledge graphs using LLMs. [23] works with several additional sources of obtaining evidence: web searches, LLM internal knowledge, and reference knowledge graphs.

**LLMs for querying Knowledge Graphs**

[27] stores nodes from a knowledge graph as document embeddings which are used at query time for RAG. [26] is given a graph as an input, either as a textual description or an image, along with a user query. [44] aims to improve Q/A performance by rewriting extracted triples from a KG into more descriptive, readable text.

[39] uses both a knowledge graph as well as some degree of the LLM’s internal reasoning to answer more open ended Q/A questions. [29, 34, 40?] allows for querying the knowledge graph by iteratively traversing the KG. These works treat the LLM as a planner for obtaining the necessary information to answer a question.

## 2.4 Ontologies for Multimodal Sensory Knowledge Graphs

Several ontologies were found to be relevant to our scenarios, and are particularly useful when discussing Semantic Web related platforms. The Brick schema [11] describes an RDF-written ontology, designed to represent resources in building management systems. It includes concepts covering individual sensory devices as well as locations within buildings, represented by logical relationships. The Smart Applications Reference Ontology (SAREF) [16] and its wearable version (SAREF4WEAR) [17] both describe sensor-specific characteristics, including properties of the device, measurements, and real world features to be measured. The Sensor, Observation, Sample, Actuator Ontology (SOSA) [18] provides a similar sensor-centric view, describing features of interest, properties of sensors, as well as measurements.

We found that these existing ontologies often lacked certain features necessary to organize urban sensory data (such as spatial, temporal, or incident classes). Thus, we did not adopt any existing ontology as-is. However, we have aligned a significant set of concepts with existing ontology classes and terms, which we describe later in Section 3.

While there exists some works on smart city modeling, such as [19, 21, 28], these works either do not involve the same notion of events that we wish to model, lack sensor-specificity, or target simpler structured sensor readings and not more complex information such as actors within news articles.

## 3 Background

In this section we provide an overview of terminology used in the *SIGMUS* system, as well as the ontology underlying the knowledge graph.

### 3.1 Terminology

Table 1 describes the different terms we use when organizing data in urban spaces. The notion of **incidents** is motivated by the need for larger occurrences influencing cities or states (such as natural disasters, major entertainment events, etc) to be represented in our knowledge graph, allowing us to identify more specific components contributing to them. **Reports** are an abstraction where sensory data (and its associated metadata like time, accuracy, and location), as well as human-generated knowledge (such as news and tweets) can be integrated in our knowledge graph. In our system, a report is an individual news article, image, or sample from a time series. **Observers** can be used to identify the entity publishing a particular report. For sensors, this can be a specific sensor ID, while for human-generated knowledge this can be an organization name (such as a news station) or username. **Modality** represents a specific modality of a **report**. This allows us to separate out different types of data present in a report possessing different properties. For example, a news report may use images, audio, and text. Images yield a caption

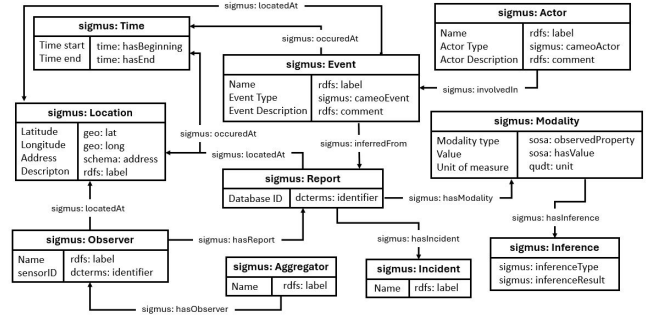


Figure 2: The *SIGMUS* ontology

property describing the image, audio describes the loudness of a scene, and text reveals the named actors involved an event. **Aggregators** are the platforms from which we obtain data from (such as X.com, Caltrans CCTV, etc), and may vary in how trustworthy its knowledge is depending on the nature of its observers (though currently we do not model these attributes). **Events** allow us to describe interactions between two parties, such as "accuse of a crime", or "protest" (where interactions are described under CAMEO codebook [7]). In our current implementation, **events** are specific to news articles, but other modalities may also reveal specific interactions between parties. **Actors** are the parties involved with an **event**, where both individual people and larger organizations of people may be named as actors.

### 3.2 The *SIGMUS* ontology

Having identified the main terms in our system, we now introduce the relationships between each of the terms, with each term being a class in our ontology). Figure 2 describes how each class is related, as well as the main attributes present in for each class.

As our ontology focuses on creating common representations across different modalities (such as image, text, and tabular data), we designed classes with properties which can be shared across those data. In addition, we also introduce new terms which allow us to represent high impact phenomena that influence cities and states, while also being able to attach different types of contributing factors and knowledge associated with such phenomena.

Many of the properties are reused from existing ontologies:

- *dcterms: identifier* is used as a reference to a particular resource represented as some unique key or string, taken from the Dublin Core Terms ontology [8].
- *rdfs: label* and *rdfs: comment* are used to describe a human readable string and a description, respectively. These are taken from the RDF Schema [14].
- *time* properties are taken from the OWL Time ontology [32], allowing us to represent the start and end time of events.
- *sosa* and *qudt* are both part of the Semantic Sensor Network Ontology [18], which we use to describe the properties and units of sensory measurements. For more complex data (like images and audio), the **Value** field is treated as a filepath to the data rather than directly encoding it in the knowledge graph instance.

There are also several new fields:

Term	Definition
Incident	An occurrence, natural or manmade, that necessitates a response to protect life or property. The word “incident” includes planned events as well as emergencies and/or disasters of all kinds and sizes. This definition is taken from the Federal Emergency Management Agency (FEMA) [10]. In our system, incidents may be linked to multiple <b>reports</b> , and may occur as part of other incidents.
Report	An observation which may be structured or annotated in some way, produced from <b>observers</b> . Annotations involve labels such as timestamps, locations, or other labels which allow us to link the data to real world entities
Observer	An entity, either a sensor or individual that directly perceives, records and/or detects features of the environment. These organize information into <b>reports</b> which are then collected by <b>Aggregators</b> .
Modality	A subset of a report represented by a unique type of data, such as image, audio, or text. A report may have multiple modalities of information present.
Aggregator	Is an entity that collects information from multiple <b>observers</b> and presents them in the form of <b>reports</b> . These sources may be sensory or human, and may represent information in multiple modalities, such as textual, visual, auditory, and tabular.
Data Source	We use this term interchangeably with <b>Aggregator</b> , but when discussing our knowledge graph we prefer aggregator as it aligns better with the rest of our terminology.
Event	Describes interactions between actors, with various terms relating to violent conflict, political actions, or other types of cooperative behaviors. We use the Conflict and Mediation Event Observations (CAMEO) codebook [7] to describe our events. <b>Incidents</b> can be viewed as events that become significant in terms of people involved or damage caused, making them of interest to a city or state.
Actor	A real world entity, person, or organization which has a name. We use the Conflict and Mediation Event Observations (CAMEO) codebook [7] to describe our actors. However, while CAMEO is typically focused on larger groups (e.g. political parties, unions, armed groups), we allow for individual people to be identified as actors.

**Table 1: List of terms used in our system and their definitions**

- The *sigmus: Inference* class introduces the *inferenceType* property, which describes the algorithm or ML task that was used to process the data (such as “Trend Analysis”, “Image Captioning”). The result of this processing is stored in the *inferenceResult* field as a human-readable string.
- The *sigmus: cameoEvent* property represents the codes from CAMEO [7] that can be mapped to specific event terms and descriptions (e.g. assassinations) or actors (e.g. government).

## 4 Methods

### 4.1 Data sources and Ingestion

For our choices of data sources, had several criteria:

- Publishes data on a regular basis such that important incidents can be monitored
- Capable of capturing incidents in a unique way, making it useful for integration
- Data can be obtained at low or no cost, reducing operational expenses of the *SIGMUS* system in the long term.

To that end, we chose 10 different data sources, across 3 modalities: image, text, and tabular. A summary of each data source is described in Table 2.

During the **Data Ingestion** step shown in Figure 3, we have modality-specific processes for ingesting, filtering and organizing into our data warehouse. This is necessary as each data source presents different methods of publishing their own data. Our data warehouse provides both a common interface for interacting with all the ingested data, as well as durability for such data. We implement the **Data Ingestion** for each modality using methods such as

scraping, web API access, email alerts, or via a 3rd party capable of querying the data source at a lower cost (such as IFTTT for X.com). After obtaining the data, we perform some degree of processing in order to obtain the described attributes in Table 2, and filtering data involving Los Angeles County. We choose Los Angeles County to demonstrate the {sysname} system as it often has incidents occurring (such as sporting and cultural events, emergencies such as wildfires and earthquakes, and various political/social activities).

While the **Attributes** column in Table 2 lists most of the modality-specific information our system ingests, it doesn’t list metadata such as location, direction, time which is also used in our system. The **In. Frequency** column describes the frequency at which we poll the data source in order to obtain new data. Note that even for event-based sources (such as tweets or alerts), we interface with them via a 3rd party (such as an email server) which stores the data for us to poll. The **In. Latency** column describes the worst case latency between the time data is published to the time our system ingests the data. In some cases this is marked as *none*, since the polling process obtains the most recent published measurement.

Our data warehouse stores the ingested data for durability before integrating different sources together into the live knowledge graph. This data warehouse uses both time series storage (to allow for specialized to time-indexed data) while the blob storage allows for better scaling of unstructured data.

### 4.2 Knowledge Graph and Integration

The live knowledge graph acts as a working memory for integrating different data sources together and providing a visual interface for users to identify events and their connections. To create this

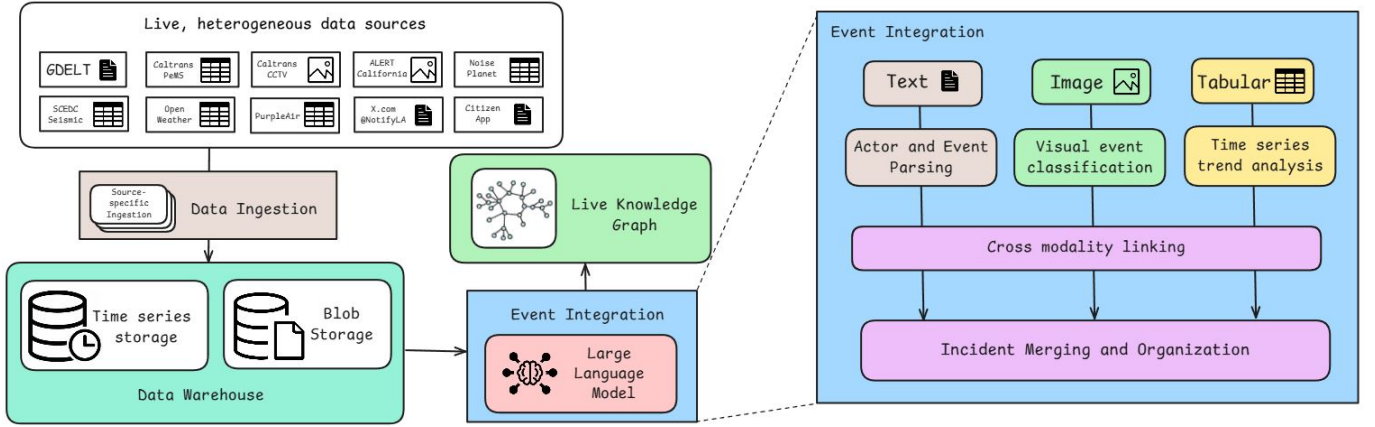


Figure 3: Technical architecture for ingesting knowledge into a live knowledge graph for urban analysis.

Data Source	Attributes	Data Publish Method	In. Frequency	In. Latency	Data Filters
GDELT events [9]	news article URLs, named actors, CAMEO event codes, Geo names	Zip files of CSVs in 15 minute intervals	15min	15min	Mentions of Los Angeles
Caltrans PeMS [25]	Traffic incidents, vehicle speeds, vehicle occupancy	Zip files of CSVs updated daily	1 day	1 day	incidents and sensors in LA
Caltrans CCTV [2]	images from highway cameras	Live video streams	15min	None	Cameras in LA county
ALERT California [1]	images from wildfire cameras	Live video streams	30min	None	Cameras in LA county
Noise Planet [13]	noise levels in db(A)	Zip files updated daily	1 day	1 day	Measurements from LA
SCEDC Seismic [20]	seismic waveform data from earthquake monitoring stations	Zip files of waveform data updated daily	1 day	1 day	Seismic stations in LA county
OpenWeather [6]	wind speed, precipitation, and weather descriptions	API-based access, updated every 10 minutes	1 hour	None	Weather stations in LA county
PurpleAir [15]	PM2.5	API-based access, updated every 2 minutes	1 hour	None	Outdoor sensors in LA county
X.com	Tweet author, emergency type, location, tweet text	Social network postings (event-based)	15min	15min	LA emergency service accounts
Citizen App [3]	Text alerts, alert category, location of alert	App alerts (event-based)	15min	15min	Alerts about locations in LA county

Table 2: Summary of data extracted from different data sources

knowledge graph, we first must integrate our different data sources together and create connections based on our ontology.

Each general modality type (image, text, and tabular) have different associated processes with them.

**Actor and Event Parsing** uses LLMs to identify any real world organizations and persons involved in the text, as well as classifying these into the corresponding CAMEO term [7]. In addition, it must categorize the event occurring between the actors using CAMEO event terms [7], as well as identifying any alerts or emergencies using the Common Alerting Protocol (CAP) [4], if any. An high level example of this process is shown in Figure 4.

In addition to establishing the actors and events, we also attempt to identify any noteworthy incidents. As major incidents will

have an associated name (i.e. Palisades Fire 2025) created by news agencies, we identify incidents primarily using the news modality (GDELT). Later on, we use the incidents as a way of connecting different modalities together.

This process of identifying actors can be fairly time intensive, particularly as the LLM is prompted with a list of possible names which can be referring to the same actor (as well as the context, such as location, association with other actors, or other relevant news information). As the number of actors present in our knowledge graph increases, this process can become incredibly inefficient especially as the number of input tokens increase. Other problems may also arise as the input begins to reach the maximum input length of the LLM model. Thus, we utilize a modified word edit



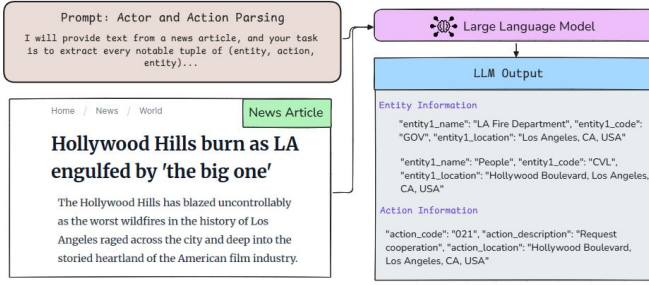


Figure 4: Example of actor and event parsing

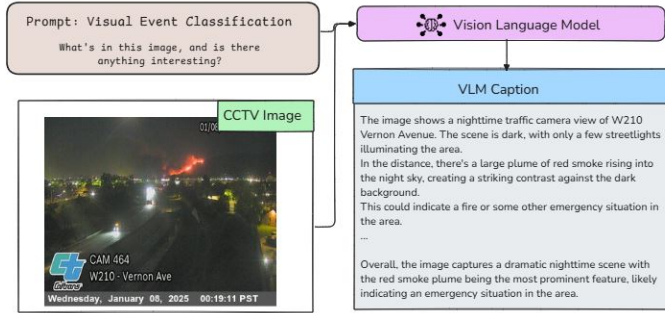


Figure 5: Example of visual event classification

distance which allows for relaxed penalties when matching part of names. Using this, we obtain some top-k similar set of names which are used to query the model.

**Visual Event Classification** uses a Vision-Language Model (VLM) to caption images and identify interesting occurrences, if any. An example of this process is shown in Figure 5. Our main objective is to transform visual data into a human readable format which can both be stored on the knowledge graph, as well as utilized in downstream LLM reasoning.

**Time Series Analysis** simply involves statistical analysis of values as well as a historical comparison at different intervals (e.g. 1 hour, 1 day, 1 week). The result of this analysis is also stored on the knowledge graph, and allows for downstream LLM reasoning to identify anomalies or correlate values across modalities.

Having produced additional annotations for text, image, and tabular data, we create reports (i.e. from an image, a news article, a sensory reading). After generating our reports, we can begin to link information across modalities via incidents.

**4.2.1 Cross Modality Linking.** We connect reports from different modalities by linking them to the same shared incident. In this work, we determine connections on every incoming report for each data source (aside from GDELT events, as that data source is focused on generating incidents).

Once the report is processed depending on its modality (as described in the above sections), it is inserted into our knowledge graph. After it is inserted into our knowledge graph, we query the knowledge graph for all incidents (as well as some additional related context, such as related time and geographic information for that incident). We feed the given report's information (e.g. image

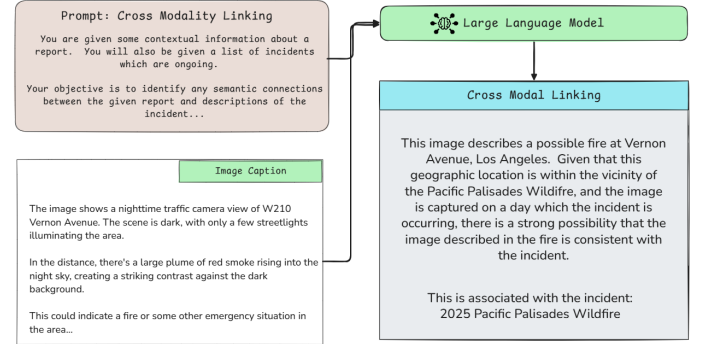


Figure 6: Example of cross modal linking

caption, trends) into an LLM, along with the list of the incidents and contexts. An example of this is shown in Figure 6

**4.2.2 Incident Merging and Organization.** In order to both disambiguate different names which refer to the same incident, we use a RAG approach. Similar to the problems arising in actor parsing, there may be many incidents, both ongoing and historical which may be relevant to a new incident. In order to efficiently query the model, we also provide a similarity ranking of incidents outside of LLMs using a vector database and text embedding model to store embeddings of an incident and associated description of the incident. Given a new incident and its news text, we search through existing embeddings and rank them by similarity. Finally, we use the top-k similar incidents and obtain the text of their original news article. We feed this as context for the LLM to perform a selection over the top-k similar incidents to decide if the given incident can be organized under the same name.

## 5 Evaluation

We evaluate *SIGMUS* using two LLMs. One LLM is a DeepSeek-R1 70B model, which we operate on an H100 GPU server hosted on our local premises. Another is the cloud-based OpenAI GPT-4o model. The reason for these two different models is mainly due to quality of the output, with the GPT-4o performing significantly better than the DeepSeek model on the cross-modality linking task. For that reason, we rely on the GPT-4o model for the cross-modality linking task, while for all other tasks we utilize DeepSeek-R1. For Visual Event classification, we rely on a Visual Large Language Model called NVILA [37], which also runs on the local H100 server. During the Incident Merging and Organization task, we utilize a vector database called Marqo [12], which allows us to find similar incidents based on their textual embeddings. Our data warehouse uses TimescaleDB for its time series storage, and SeaweedFS for the blob storage. All ingested data is organized and visualized in Neo4j, the knowledge graph used in *SIGMUS*. In our evaluation, data storage services (i.e. time series, blob, vector, KG) operate on the same desktop PC.

It is also worth noting that we mainly use the 2025 Los Angeles Wildfires as our case study. As not all modalities are relevant or have data from this period of time, we only focus on a subset of the data sources described in Table 2. More specifically, we use GDELT events, PeMS, CCTV, OpenWeather, and PurpleAir. For

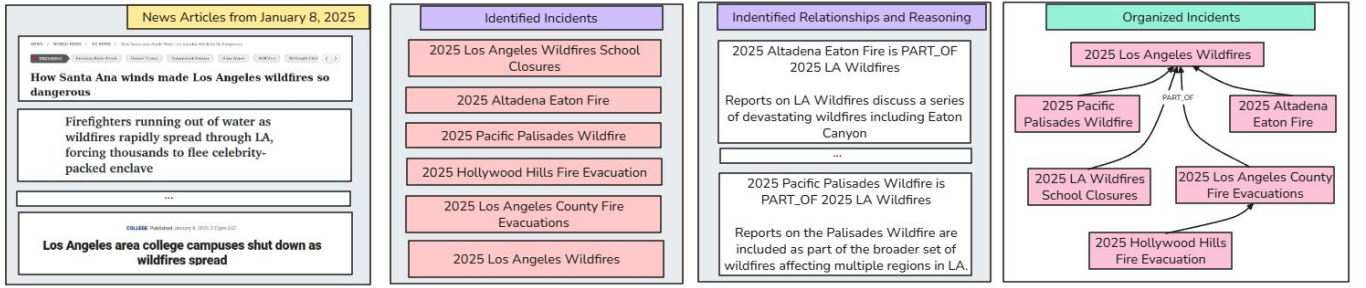


Figure 7: Incident Merging and Organization

Data Source	Processing Latency (s)	KG insert latency (s)
GDELT events	39.3795	30.8118
Caltrans PeMS	0.0105	0.0387
Caltrans CCTV	5.905	0.0247
ALERT California	6.4577	0.0507
Air Quality	0.0043	0.0207
Weather	0.0144	0.0267

Table 3: Latency of ingesting a single report from each data source

the remaining sources, we either lacked data from that period of time, or they are not as relevant to the incident itself (i.e. seismic information).

Table 3 describes the average time, in seconds, to ingest a single report from each data source and add it to the knowledge graph. While these times prove to be quite slow when performing historical analysis, the goal of *SIGMUS* is to execute and ingest data in a live manner - thus, these reported times are not batched operations.

As mentioned in section 3, a single report is treated as an individual news article, image, or row in a table. Note that depending on the modality, the processing and knowledge graph latency may be affected. For example, GDELT events (being a text modality) involves calling an LLM to identify actors, events, and incidents (which makes the bulk of the processing latency). However, when inserting such data into a knowledge graph, we have to disambiguate names which refer to the same entity. These are two kinds of disambiguation processes which occur when we insert items into the knowledge graph: Actor disambiguation and Incident disambiguation. Note that because event names are canonical (generated from [7]), we do not perform disambiguation for these.

**Actor Disambiguation** To disambiguate similar actors, *SIGMUS* first identifies a set of top-k similar actors based on word edit distance of their names, and prompts the LLM to identify whether these top-k actors (and their context, such as news report topic, location, etc) are likely to be the same actor. On average, this process of actor merging takes 8.1101 seconds on average, with each report possibly involving multiple actors.

**Incident Disambiguation** When inserting incidents into the knowledge graph, sysname identifies the top similar existing incidents based on both its label (e.g. "2025 Pacific Palisades Wildfire") as well as the original text from the article. This process of ranking the top

similar incidents is done via vector search in Marqo [12], which produces embeddings of each incident and allows for similarity searching given a query. After obtaining the top-k similar incidents, our system prompts the LLM to identify if the inserted incident already exists in the top-k. In addition, the LLM also identifies if the inserted incident related to another incident, and we mainly focus on "IS\_PART\_OF" as a possible relationship between incidents. This process of inserting incidents takes 7.8677 seconds on average.

## 6 Discussion

### 6.1 Future Evaluation

While much of our evaluation describes qualitative results for different modalities of data and some limited performance analysis of *SIGMUS*, there is still much evaluation to be done in the future. First is determining a metric for evaluating the quality of connections made across modalities. While for some cases this is fairly straightforward (e.g. CCTV image of a fire is captures information about a specific known fire incident) while for other cases it may be more complex (poor air quality at a particular location may be affected by a faraway incident such as smoke from a wildfire, or a local cooking action). Verifying connectivity in these more complex cases requires significant user effort in examining both historical and neighborhood information to reach a conclusion.

There is also the issue of consistent outputs across different models. Each LLM may produce a different set of connections within our knowledge graph, and it may be useful in future work to use voting strategies among a set of LLMs to determine the most likely set of connections. A similar issue arises with the inputs to LLMs - small changes in the text prompt also affect performance, thus requiring some degree of prompt tuning to achieve the best performance.

### 6.2 Case studies

In this work we only examined the 2025 Los Angeles Wildfires. While this incident itself is quite interesting (with multiple sub-incidents and relevance to multiple different data sources), our system may also be applied to other incidents. In the future, we may consider major entertainment events (such as sports or music), or cases more relevant to domestic security, should they arise. In addition, our geographic scope has thus far been limited to Los Angeles as our main urban space of interest, but other cities may present more interesting case studies as well. As there is a plethora



Report	Identified Connection	Connection Reasoning
	This report is associated with incident 2025 Pacific Palisades Wildfire	The report describes an image from a Caltrans traffic camera at "W210 - Vernon Ave." showing a busy highway in California with a noticeable haze or smoke in the air. This haze or smoke could be from wildfires, as mentioned in the report, and aligns with the Pacific Palisades Wildfire incident ongoing on January 8, 2025. The geographical location in the vicinity and the timing of the report at 11:33:13 PST connects it to the Pacific Palisades area and the ongoing wildfire there, providing a visual snapshot of the impact of the wildfire, evident through the haze in the air due to the incident.
	This report is associated with incident 2025 Pacific Palisades Wildfire	The report describes a nighttime traffic scene with a bright orange glow resembling a massive fire, visible in the background. It is timestamped "January 08, 2025 at 05:15:03 AM," which coincides with the timeline of the 2025 Pacific Palisades Wildfire. The camera is located at "W210 - Vernon Ave.," which is a region in Southern California, not far from the Pacific Palisades area, which is experiencing a significant wildfire incident on the same day. The report captures visual elements consistent with a large fire event, indicating a connection to the wildfire incident.
Air Quality at location [34.047226, -118.51949] PM2.5 measured as: 1068.5	This report is associated with incident 2025 Los Angeles Wildfires School Closure	The report shows a high air quality value of 1068.5 with an increasing trend over the past 15 minutes at the location [34.047226, -118.51949] on January 08, 2025 at 10:00:02 PM. This matches the context of incident 3, which is the 2025 Los Angeles Wildfires School Closures, involving similar air quality readings at the same time and location. The increasing air quality value is indicative of the impact from nearby wildfires, which is consistent with school closures due to poor air conditions.
Air Quality at location [34.142403, -118.11669] PM2.5 measured as: 178.4	This report is associated with incident 2025 Altadena Eaton Fire	The report is connected to the 2025 Eaton Fire, as it captures air quality data from a location [34.142403, -118.11669] that is quite close to Altadena, CA, one of the affected areas in the Eaton Fire incident. This proximity in location, combined with the elevated pollution levels indicated by the "Mean value" of 178.4, suggests a direct impact of the wildfire on the local air quality. The report's timestamp of January 08, 2025 at 11:00:02 PM closely follows the Eaton Fire activity, which was significant around 6:00 PM on the same day, indicating that the air quality changes are likely due to the ongoing wildfire.
Report	Identified Connection	Connection Reasoning
Weather at location Beverly Hills with lat/long [34.0736, -118.4004] Temp: 60.39F Description: Smoky Humidity 20% wind speed: 10.8m/s	This report is associated with incident 2025 Pacific Palisades Wildfire	The report from Beverly Hills indicates the presence of smoke with a temperature of 60.39°F and wind speeds of 10.8 m/s on January 08, 2025, at 03:03:25 PM. Coupled with the low humidity of 20%, these conditions are likely indicative of wildfires affecting air quality. The Pacific Palisades Wildfire is reported within close proximity at a similar timeframe, with incidents of wildfire and heavy winds leading to evacuations in nearby areas. The smoke detected in Beverly Hills can reasonably be attributed to the ongoing Pacific Palisades Wildfire.
Weather at location Pasadena with lat/long [34.1478, -118.1445] Temp: 60F Description: Overcast Humidity 40% wind speed: 6.17m/s	This report is associated with incident 2025 Altadena Eaton Fire	The report is relevant to the 2025 Eaton Fire incident near Pasadena. The report is capturing weather data from Pasadena, US, exactly coinciding with the time of the incident, January 08, 2025, at 06:03:25 PM. The recorded wind speed of 6.17 m/s and the overcast cloud weather condition are significant as they can influence fire behavior and spread, making this data relevant to understanding conditions surrounding the Eaton Fire incident at the time.
Traffic data at location [34.146752, -118.029146] Occupancy: 0.2511 Speed: 12.4mph	This report is associated with incident 2025 Altadena Eaton Fire	The report was taken near the reported location of the Eaton Fire on January 8, 2025, at 8:00 AM PST, at coordinates [34.146752, -118.029146]. Both show's an increase in trend and low values for occupancy and speed [0.2511 and 12.4 respectively], which may indirectly indicate road closures or restrictions due to the Eaton Fire. The geographic proximity to the reported location of the Eaton Fire supports the connection between the report and the reported incident.
Traffic data at location [34.1544, -118.649487] Occupancy: 0.2325 Speed: 26.1mph	This report is associated with incident 2025 Los Angeles Wildfires	The report describes traffic conditions with an occupancy mean value of 0.2325 and a speed mean value of 26.1 at coordinates [34.1544, -118.649487] on January 08, 2025, at 08:55 AM. These conditions might be affected by the ongoing Los Angeles Wildfires incident, known for causing traffic disruptions due to road closures and detours. Although the report's coordinates are not an exact match, the described speed and occupancy suggest unusual traffic behaviors, which align with the disruptions expected from the nearby wildfires.

Figure 8: Examples of linking multimodal reports with incidents

of public data sources collecting information for Los Angeles, we found it easier to start off in this location.

## 7 Conclusion

In this work, we introduced *SIGMUS*, a system that leverages Large Language Models (LLMs) to integrate and reason over multimodal urban data streams for the purpose of incident detection and knowledge organization. By unifying fragmented sensory inputs—including text, image, and tabular data—into a structured, semantically rich knowledge graph, *SIGMUS* enables a more comprehensive understanding of complex urban events such as wildfires, environmental hazards, or public safety incidents.

Through our case study of the 2025 Los Angeles wildfires, we demonstrated the system's ability to qualitatively associate disparate data sources and derive contextually relevant relationships between them. Our findings suggest that LLMs can serve as powerful tools for bridging the gap between raw sensory input and actionable urban intelligence, particularly when world knowledge and contextual reasoning are required.

This work highlights the promise of LLM-powered knowledge systems in urban computing, especially for real-time monitoring and historical analysis of city-scale phenomena. Future directions include developing more case studies and incorporating more robust evaluation strategies for the quality and correctness of inferred relationships.

## Acknowledgments

This research was sponsored in part by the Sandia National Laboratories award 2169310 and the DEVCOM ARL award W911NF1720196.

## References

- [1] [n. d.]. ALERTCalifornia. <https://alertcalifornia.org/>
- [2] [n. d.]. Caltrans Streaming Video Locations. <https://cwwp2.dot.ca.gov/vm/streamlist.htm>
- [3] [n. d.]. Citizen: Keeping you safe & informed. <http://www.citizen.com>
- [4] [n. d.]. Common Alerting Protocol. <https://docs.oasis-open.org/emergency/cap/v1.2/CAP-v1.2-os.html>
- [5] [n. d.]. CSO - Portal. <https://cso.kmi.open.ac.uk/home>
- [6] [n. d.]. Current weather and forecast - OpenWeatherMap. <https://openweathermap.org/>
- [7] [n. d.]. Data: CAMEO. <https://eventdata.parusanalytics.com/data.dir/cameo.html>
- [8] [n. d.]. DCMI Metadata Terms. <https://www.dublincore.org/specifications/dublin-core/dcmi-terms/>
- [9] [n. d.]. The GDELT Project. <https://www.gdelproject.org/>
- [10] [n. d.]. Glossary | FEMA.gov. <https://www.fema.gov/about/glossary>
- [11] [n. d.]. Introduction — Brick Ontology Documentation. <https://docs.brickschema.org/intro.html>
- [12] [n. d.]. Marqo | Smarter Search, Higher Conversion. <https://www.marqo.ai/>
- [13] [n. d.]. Noise-Planet - Scientific tools for environmental noise assessment. <https://noise-planet.org/>
- [14] [n. d.]. RDF Schema 1.1. <https://www.w3.org/TR/rdf-schema/>
- [15] [n. d.]. Real-time Air Quality Monitoring by PurpleAir. <https://www2.purpleair.com/>
- [16] [n. d.]. SAREF Portal. <https://saref.etsi.org/>
- [17] [n. d.]. SAREF4WEAR: an extension of SAREF for Wearables. <https://saref.etsi.org/saref4wear/v1.1.1/>
- [18] [n. d.]. Semantic Sensor Network Ontology. <https://www.w3.org/TR/vocab-ssn/>
- [19] [n. d.]. Smart City Ontology 2.0 - URENIO | Intelligent Cities - Smart Cities - Innovation Ecosystems. <https://urenio.org/smart-city-ontology/>
- [20] [n. d.]. Southern California Earthquake Data Center at Caltech. <https://scedc.caltech.edu/>
- [21] 2025. Azure/opendigitaltwins-smartcities. <https://github.com/Azure/opendigitaltwins-smartcities> original-date: 2020-09-01T23:19:08Z.
- [22] Daniel Adam and Tomás Kliegr. 2024. Traceable LLM-based validation of statements in knowledge graphs. doi:10.48550/arXiv.2409.07507 arXiv:2409.07507 [cs].
- [23] Jack Boylan, Shashank Mangla, Dominic Thorn, Demian Gholipour Ghalandari, Parsa Ghaffari, and Chris Hokamp. 2024. KGValidator: A Framework for Automatic Validation of Knowledge Graph Construction. doi:10.48550/arXiv.2404.



- 15923 arXiv:2404.15923 [cs].
- [24] Arne Bröring, Johannes Echterhoff, Simon Jirka, Ingo Simonis, Thomas Everding, Christoph Stasch, Steve Liang, and Rob Lemmens. 2011. New generation sensor web enablement. *sensors* 11, 3 (2011), 2652–2699.
- [25] State of California. [n.d.]. PeMS Data Source | Caltrans. <https://dot.ca.gov/programs/traffic-operations/mpr/pems-source>
- [26] Yihe Deng, Chenchen Ye, Zijie Huang, Mingyu Derek Ma, Yiwen Kou, and Wei Wang. [n.d.]. GraphVis: Boosting LLMs with Visual Knowledge Graph Integration. ([n.d.]).
- [27] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitan, Robert Osazuwa Ness, and Jonathan Larson. 2025. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. doi:10.48550/arXiv:2404.16130 arXiv:2404.16130 [cs].
- [28] Paola Espinoza-Arias, Maria Poveda-Villalón, Raúl García-Castro, and Oscar Corcho. 2019. Ontological Representation of Smart City Data: From Devices to Cities. *Applied Sciences* 9, 1 (Jan. 2019), 32. doi:10.3390/app9010032 Number: 1 Publisher: Multidisciplinary Digital Publishing Institute.
- [29] Chao Feng, Xinyu Zhang, and Zichu Fei. 2023. Knowledge Solver: Teaching LLMs to Search for Domain Knowledge from Knowledge Graphs. doi:10.48550/arXiv:2309.03118 arXiv:2309.03118 [cs].
- [30] Hamed Babaei Giglou, Jennifer D’Souza, and Sören Auer. 2023. LLMs4OL: Large Language Models for Ontology Learning. doi:10.48550/arXiv:2307.16648 arXiv:2307.16648 [cs].
- [31] William I Grosky, Aman Kansal, Suman Nath, Jie Liu, and Feng Zhao. 2007. Senseweb: An infrastructure for shared sensing. *IEEE multimedia* 14, 4 (2007), 8–13.
- [32] Jerry R. Hobbs and Feng Pan. 2004. An ontology of time for the semantic web. *ACM Transactions on Asian Language Information Processing* 3, 1 (March 2004), 66–85. doi:10.1145/1017068.1017073
- [33] Pavan Holur, Shreyas Rajesh, David Chong, and Vwani Roychowdhury. 2024. Creating an AI Observer: Generative Semantic Workspaces. doi:10.48550/arXiv.2406.04555 arXiv:2406.04555 [cs].
- [34] Jinhao Jiang, Kun Zhou, Zican Dong, Keming Ye, Xin Zhao, and Ji-Rong Wen. 2023. StructGPT: A General Framework for Large Language Model to Reason over Structured Data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 9237–9251. doi:10.18653/v1/2023.emnlp-main.574
- [35] Amar Viswanathan Kannan, Dmitriy Fradkin, Ioannis Akrotirianakis, Tugba Kulahcioglu, Arquimedes Canedo, Aditi Roy, Shih-Yuan Yu, Malawade Arnav, and Mohammad Abdullah Al Faruque. 2020. Multimodal Knowledge Graph for Deep Learning Papers and Code. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. ACM, Virtual Event Ireland, 3417–3420. doi:10.1145/3340531.3417439
- [36] Danh Le-Phuoc, Hoan Nguyen Mau Quoc, Hung Ngo Quoc, Tuan Tran Nhat, and Manfred Hauswirth. 2016. The graph of things: A step towards the live knowledge graph of connected things. *Journal of Web Semantics* 37 (2016), 25–35.
- [37] Zhijian Liu, Ligeng Zhu, Baifeng Shi, Zhuoyang Zhang, Yuming Lou, Shang Yang, Haocheng Xi, Shiyi Cao, Yuxian Gu, Dacheng Li, Xiuyu Li, Yunhao Fang, Yukang Chen, Cheng-Yu Hsieh, De-An Huang, An-Chieh Cheng, Vishwesh Nath, Jinyi Hu, Sifei Liu, Ranjay Krishna, Daguang Xu, Xiaolong Wang, Pavlo Molchanov, Jan Kautz, Hongxu Yin, Song Han, and Yao Lu. 2025. NVILA: Efficient Frontier Visual Language Models. doi:10.48550/arXiv.2412.04468 arXiv:2412.04468 [cs].
- [38] Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 31.
- [39] Yuan Sui, Yufei He, Zifeng Ding, and Bryan Hooi. 2025. Can Knowledge Graphs Make Large Language Models More Trustworthy? An Empirical Study Over Open-ended Question Answering. doi:10.48550/arXiv.2410.08085 arXiv:2410.08085 [cs].
- [40] Jiashuo Sun, Chengjin Xu, Luminyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel M. Ni, Heung-Yeung Shum, and Jian Guo. 2024. Think-on-Graph: Deep and Responsible Reasoning of Large Language Model on Knowledge Graph. doi:10.48550/arXiv.2307.07697 arXiv:2307.07697 [cs].
- [41] Sabrina Toro, Anna V. Anagnostopoulos, Sue Bello, Kai Blumberg, Rhiannon Cameron, Leigh Carmody, Alexander D. Diehl, Damion Dooley, William Duncan, Petra Fey, Pascale Gaudet, Nomi L. Harris, Marcin Joachimiak, Leila Kiani, Tiago Lubiana, Monica C. Munoz-Torres, Shawn O’Neil, David Osumi-Sutherland, Aleix Puig, Justin P. Reese, Leonore Reiser, Sofia Robb, Troy Ruemping, James Seager, Eric Sid, Ray Stefancsik, Magalie Weber, Valerie Wood, Melissa A. Haendel, and Christopher J. Mungall. 2024. Dynamic Retrieval Augmented Generation of Ontologies using Artificial Intelligence (DRAGON-AI). *Journal of Biomedical Semantics* 15, 1 (Oct. 2024), 19. doi:10.1186/s13326-024-00320-3 arXiv:2312.10904 [cs].
- [42] Atiya Usmani, M. Jaleed Khan, John G. Breslin, and Edward Curry. 2023. Towards Multimodal Knowledge Graphs for Data Spaces. In *Companion Proceedings of the ACM Web Conference 2023*. ACM, Austin TX USA, 1494–1499. doi:10.1145/3543873.3587665
- [43] Xin Wang, Benyuan Meng, Hong Chen, Yuan Meng, Ke Lv, and Wenwu Zhu. 2023. TIVA-KG: A Multimodal Knowledge Graph with Text, Image, Video and Audio. In *Proceedings of the 31st ACM International Conference on Multimedia*. ACM, Ottawa ON Canada, 2391–2399. doi:10.1145/3581783.3612266
- [44] Yike Wu, Nan Hu, Sheng Bi, Guilin Qi, Jie Ren, Anhuan Xie, and Wei Song. 2023. Retrieve-Rewrite-Answer: A KG-to-Text Enhanced LLMs Framework for Knowledge Graph Question Answering. doi:10.48550/arXiv.2309.11206 arXiv:2309.11206 [cs].
- [45] Huatao Xu, Liying Han, Qirui Yang, Mo Li, and Mani Srivastava. 2024. Penetrative AI: Making LLMs Comprehend the Physical World. In *Proceedings of the 25th International Workshop on Mobile Computing Systems and Applications (HotMobile ’24)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3638550.3641130
- [46] Yao Xu, Shizhu He, Jiabei Chen, Zihao Wang, Yangqiu Song, Hanghang Tong, Guang Liu, Kang Liu, and Jun Zhao. 2024. Generate-on-Graph: Treat LLM as both Agent and KG in Incomplete Knowledge Graph Question Answering. doi:10.48550/arXiv.2404.14741 arXiv:2404.14741 [cs].
- [47] Bufang Yang, Siyang Jiang, Lilin Xu, Kaiwei Liu, Hai Li, Guoliang Xing, Hongkai Chen, Xiaofan Jiang, and Zhenyu Yan. 2024. DrHouse: An LLM-empowered Diagnostic Reasoning System through Harnessing Outcomes from Sensor Data and Expert Knowledge. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 4 (Nov. 2024), 153:1–153:29. doi:10.1145/3699765
- [48] Antonio Zaitoun, Tomer Sagi, Szymon Wilk, and Mor Peleg. 2023. Can Large Language Models Augment a Biomedical Ontology with missing Concepts and Relations? doi:10.48550/arXiv.2311.06858 arXiv:2311.06858 [cs].
- [49] Yuqi Zhu, Xiaohan Wang, Jing Chen, Shuofei Qiao, Yixin Ou, Yunzhi Yao, Shumin Deng, HuaJun Chen, and Ningyu Zhang. 2024. LLMs for knowledge graph construction and reasoning: recent capabilities and future opportunities. *World Wide Web* 27, 5 (Aug. 2024), 58. doi:10.1007/s11280-024-01297-w