LLM-Driven Knowledge Graph Construction from Earth Observation Data for Extreme Events

Theodoros Aivalis[1,2,*] **Iraklis A. Klampanos**[1] **Antonis Troumpoukis**[2]

[1]University of Glasgow, UK [2]National Centre for Scientific Research "Demokritos", Greece

[*]teoaivalis@iit.demokritos.gr

October 25, 2025

Motivation: Rising Disaster Intensity and Data Complexity

- **Rising disaster intensity**
  - 403 major climate disasters in the US since 1980 $\rightarrow$ \$2.9T in damages[1]
  - $>$150 extreme weather events globally in 2024, displacing 800,000+ people[2]
- **Explosion of Earth Observation (EO) data**
  - Sentinel / Copernicus missions generate massive multimodal data daily
  - $>$11,700 active satellites; 2,800 launched in 2024 alone[3]

[1] National Centers for Environmental Information (NCEI), "U.S. Billion-Dollar Weather and Climate Disasters," 2025.

[2] The Guardian, "Extreme weather displaced more than 800,000 people globally in 2024," 2025.

[3] Live Science, "Number of active satellites in orbit surpasses 11,700," 2025.

Integrating Text, Imagery, and Metadata through Multimodal LLMs

Need for a unified and generalisable pipeline to integrate structured and unstructured disaster data.

- Transform heterogeneous disaster-related data into interpretable knowledge
- Use multimodal LLMs to extract ontology-guided triples and build Knowledge Graphs
- Enable transparent event comparisons, support disaster response decisions, and enhance interpretability across socio-economic, spatial, and temporal dimensions

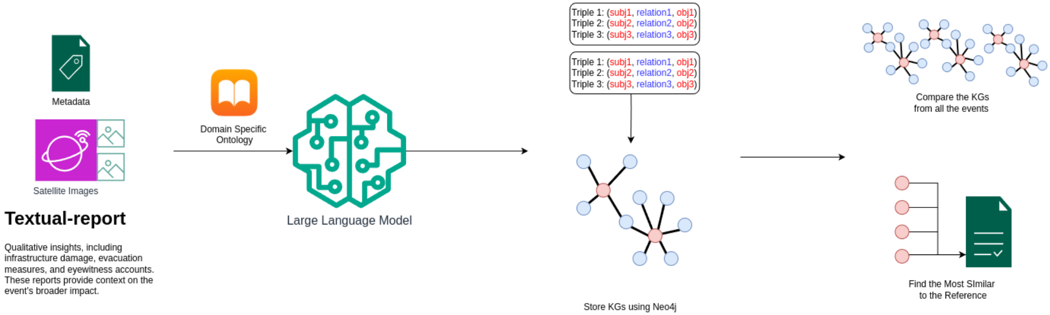

Figure: Pipeline overview for transforming heterogeneous disaster-related data into structured KGs using domain-specific ontologies and multimodal LLMs.

Case Study: Integrating Multimodal Data for Flood Event Understanding

Table: Multimodal data for the Sri Lanka flood event (December 2019).

ReliefWeb

Field (Type – Name)	Content
string – Name	Sri Lanka: Floods and Landslides - Dec 2019
date – Date	1 Dec 2019
string – Country	Sri Lanka
float – (lot, lat)	(80.7, 7.61)
string – Disaster Type	Flood
string – GLIDE ID	FL-2019-000171-LKA
text – Description	At least 3 people died following a landslide in Walapane Village. At least 570 people have been displaced in Batticaloa District and more than 4,100 people have been affected[...]

OpenStreetMap

Field (Type – Name)	Content
list – Natural Features	Peaks: 2 total; Woods: 2 total; Waters: 1 total

Microsoft Planetary Computer

Field (Type – Name)	Content
image – RGB (visual.png)	True-colour composite. Date: 11–31 Dec
image – NIR (nir.png)	Near-Infrared. Date: 11–31 Dec
image – NDWI (ndwi.png)	Normalized Difference Water Index. Date: 11–31 Dec

- Ontology encodes disaster concepts and relationships
- Based on existing standards (YAGO2geo, GeoSPARQL)
- Enables structured queries and interpretable comparisons

Image Properties	
hasLandCover	includes
showsFeature	traversedBy
hasVegetationDensity	containsFeature
hasVegetationHealth	adjacentTo
hasVegetationMoisture	isPartOf
hasMoistureLevel	mayIndicate
capturedInSeason	likelyRepresents
showsSignsOf	experiencing

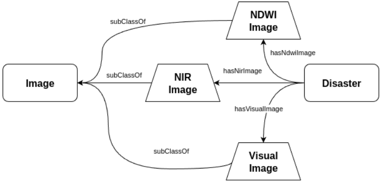

Disaster Properties				
hasCountry	hasGlide	schoolsAffected	agriculturalLandAffected	hasWetlands
hasISO3	peopleAffected	businessesDamaged	needsShelter	hasWoods
hasDisasterType	displacedPeople	roadsDamaged	humanitarianNeeds	hasScrubs
hasStatus	fatalitiesReported	infrastructureDamaged	causedBy	hasGrasslands
hasDate	injuriesReported	economicDamage	weatherCause	hasSprings
hasLatitude	housesFlooded	powerOutage	affectedArea	hasTrees
hasLongitude	housesCollapsed	waterSupplyCut	hasWaters	hasMountainRanges

Ontology schema structuring multimodal disaster data into a Knowledge Graph.

Extracting Structured Triples from Text Descriptions – Example

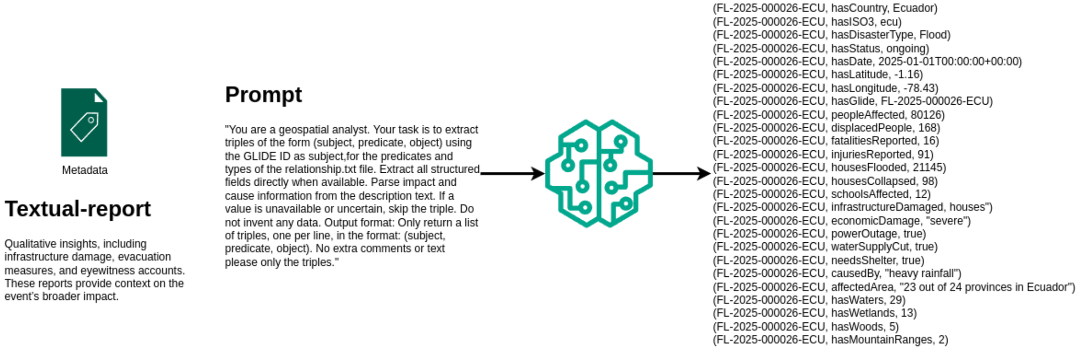

Example of triple extraction from textual disaster reports and metadata using ontology-guided prompts.

Extracting Structured Triples from Satellite Imagery – Example

NIR Image

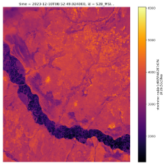

NDWI Image

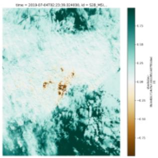

RGB Image

Prompt

"You are a geospatial analyst. Your task is to extract semantic triples of the form (subject, predicate, object) from the satellite image. The image is an **RGB/NIR/NDWI** image. The focus is on detecting flood presence and surface water dynamics. Use only the relationships found in the file **visual.txt/nir.txt/ndwi.txt**. Do not infer or assume unobservable details. Describe only what is visually supported by the image. Only output one triple per line in the format: (subject, predicate, object). No extra comments or interpretation."

(FL-2017-000054-BFA, hasvisualimage, FL-2017-000054-BFA-visual)
(FL-2017-000054-BFA-visual, hasLandCover, urban)
(FL-2017-000054-BFA-visual, showsFeature, roads)
(FL-2017-000054-BFA-visual, hasVegetationDensity, low)
(FL-2017-000054-BFA-visual, capturedInSeason, summer)
(FL-2017-000054-BFA-visual, showsSignsOf, development)
(FL-2017-000054-BFA-visual, includes, buildings)
(FL-2017-000054-BFA-visual, traversedBy, roads)
(FL-2017-000054-BFA-visual, containsFeature, green spaces)
(FL-2017-000054-BFA-visual, mayIndicate, industrial areas)

(FL-2024-000024-DZA, hasnirimage, FL-2024-000024-DZA-nir)
(FL-2024-000024-DZA-nir, hasLandCover, vegetation)
(FL-2024-000024-DZA-nir, hasLandCover, water)
(FL-2024-000024-DZA-nir, hasLandCover, builtSurfaces)
(FL-2024-000024-DZA-nir, showsFeature, brightAreas)
(FL-2024-000024-DZA-nir, showsFeature, darkRegions)
(FL-2024-000024-DZA-nir, hasVegetationHealth, healthyVegetation)
(FL-2024-000024-DZA-nir, hasMoistureLevel, highMoisture)
(FL-2024-000024-DZA-nir, showsSignsOf, floodingImpact)
(FL-2024-000024-DZA-nir, containsFeature, brightAreas)
(FL-2024-000024-DZA-nir, containsFeature, darkRegions)
(FL-2024-000024-DZA-nir, isPartOf, landscapePatterns)
(FL-2024-000024-DZA-nir, mayIndicate, flooding)
(FL-2024-000024-DZA-nir, likelyRepresents, affectedVegetation)

(FF-2021-000098-MNG-ndwi, hasndwiimage, FF-2021-000098-MNG-ndwi)
(FF-2021-000098-MNG-ndwi, showsFeature, waterBody)
(FF-2021-000098-MNG-ndwi, showsFeature, drySurface)
(FF-2021-000098-MNG-ndwi, hasLandCover, bareSoil)
(FF-2021-000098-MNG-ndwi, hasLandCover, built-UpAreas)
(FF-2021-000098-MNG-ndwi, hasVegetationMoisture, lowMoisture)
(FF-2021-000098-MNG-ndwi, hasMoistureLevel, lowMoisture)
(FF-2021-000098-MNG-ndwi, showsSignsOf, floodPresence)
(FF-2021-000098-MNG-ndwi, includes, coastalArea)
(FF-2021-000098-MNG-ndwi, containsFeature, turbulentWater)
(FF-2021-000098-MNG-ndwi, adjacentTo, landmass)
(FF-2021-000098-MNG-ndwi, likelyRepresents, Flooded Areas)

Example of triple extraction from satellite imagery (RGB, NIR, NDWI) using ontology-guided prompts.

Visualising Disasters through a Unified Knowledge Graph

Graph view of a single event (left) and the full dataset (right).

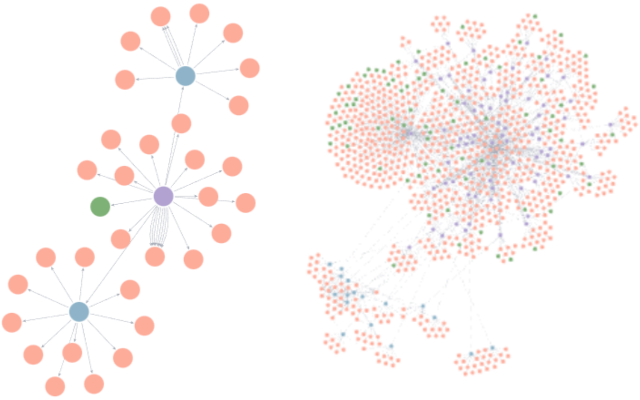

Evaluating the Quality of the Extracted Triples

- **Text-based triples:** Cosine similarity between descriptions and extracted triples.
- **Image-based triples:** Evaluated against ground-truth triples using Precision, Recall, and F1.

Metric	Value
Matching similarity	0.8915
Mean non-matching	0.8285
Min non-matching	0.7759
Max non-matching	0.8821
Std. non-matching	0.0153

Metric	RGB	NDWI	NIR
Precision	0.55	0.60	0.78
Recall	0.56	0.64	0.83
F1 score	0.55	0.61	0.80

Evaluation of extracted triples for the text modality (left) and the image modality (right).

- Results show high semantic alignment for text (cosine $\approx$ 0.89) and best image performance for NIR triples (F1 $\approx$ 0.80).

Querying the Graph: Analysing Human and Economic Impact

Top 10 floods (peopleAffected)

Region	Date	PeopleAffected
China	1989-07-01T...	100,000,000
China	1994-06-01T...	73,360,000
China	1999-06-20T...	60,000,000
China	2010-05-06T...	29,000,000
China	1990-06-01T...	26,000,000
Bangladesh	1998-07-17T...	24,000,000
India	2018-11-16T...	23,000,000
Pakistan	2010-07-22T...	20,000,000
India	2015-06-09T...	13,700,000

Top 10 floods (fatalities)

Region	Date	Fatalities
Libya	2023-09-10T...	3,922
Philippines	1991-11-05T...	2,155
India	1994-06-24T...	2,001
Sudan	2011-08-05T...	2,000
China	1989-07-01T...	2,000
Pakistan	2010-07-22T...	1,985
China	2012-07-30T...	1,250
China	2011-09-25T...	1,250
China	2004-02-20T...	1,250

Top 10 floods (BusinessesDamaged)

Region	Date	BusinessesDamaged
China	2013-04-30T00:0...	1,400,000
Bangladesh	2012-06-27T00:0...	230,000
Niger	2017-06-16T00:0...	16,000
Niger	2017-06-16T00:0...	9,800
Bangladesh	1998-07-17T00:0...	6,500
Bangladesh	2023-08-07T00:0...	1,430
China	1997-05-01T00:0...	1,000
Iran (Islamic Repu...	2020-01-11T00:0...	719
Viet Nam	2017-06-05T00:0...	425

Top 10 floods (AgriculturalLandAffected)

Region	Date	AgricultureAffected
Syrian Arab Republic	2019-03-31T00:00:...	86,000,000
China	1989-07-01T00:00:...	8,500,000
India	1998-08-20T00:00:...	7,000,000
China	1994-06-01T00:00:...	4,500,000
Pakistan	2011-07-30T00:00:...	2,100,000
Pakistan	2010-07-22T00:00:...	2,000,000
India	2010-12-06T00:00:...	1,500,000
Pakistan	2013-08-04T00:00:...	1,500,000
Thailand	1995-10-01T00:00:...	1,336,000

Example KG queries: ranking floods by human, economic, and temporal impact.

Querying the Graph: Spatial and Environmental Patterns in Flood Events

Floods in Europe - Region (via lot-lat)

Region	Date	Latitude	Longitude
Armenia	1997-06-01T...	40.61	44.66
Azerbaijan	1997-06-01T...	40.42	47.72
Romania	1997-06-01T...	45.943	24.967
Poland	1997-07-01T...	52.12	19.3
Czechia	1997-07-01T...	49.74	15.34
Moldova	1997-08-01T...	47.19	28.57
Romania	1998-06-01T...	45.943	24.967
Türkiye	1998-08-12T...	39.06	35.18
Slovakia	1998-09-14T...	48.71	19.49

All floods happened in Afghanistan

Region	DisasterName
Afghanistan	Afghanistan: Flash Floods - ...
Afghanistan	Afghanistan: Floods - Jun 19...
Afghanistan	Afghanistan: Floods - May 1...
Afghanistan	FF-2006-000163-AFG
Afghanistan	FF-2007-000091-AFG
Afghanistan	FF-2012-000118-AFG
Afghanistan	FF-2013-000088-AFG
Afghanistan	FF-2014-000060-AFG
Afghanistan	FF-2019-000018-AFG

Features of Thailand

Region	Waters	Woods	FeatureType
Thailand	6	3	hasWaters
Thailand	6	3	hasWoods

Countries similar to Thailand

Country	Waters	Woods	Wetlands	Trees	Scrubs	Grasslands	Springs	Mountain...
Italy	5	4	0	10	14	0	0	0
Monten...	8	2	0	0	1	3	24	0
Saint Vi...	8	4	0	0	1	0	0	1
Thailand	6	3	0	0	0	0	0	0
the Gren...	8	4	0	0	1	0	0	1

Example KG queries: ranking floods by regional characteristics.

Conclusions & Future Work

Conclusions

- Introduced a framework for constructing ontology-guided KGs from multimodal disaster data using multimodal LLMs
- Enabled transparent event representations and interpretable similarity comparisons
- Supports applications in disaster monitoring, early warning, and post-disaster analysis

Future Work

- Integrate additional data sources: social media, sensor networks, governmental databases
- Employ multimodal LLMs fine-tuned for EO tasks
- Build interactive tools with domain experts for real-world disaster management

Thank you!

Dataset (Hugging Face)

Linkedin

GitHub: https://github.com/teoaivalis/XtremeKG

teoaivalis@iit.demokritos.gr