

DECIDE: Decision-making Engine for Catastrophe Insights and Disaster Evaluation

Iker Pérez-del-Olmo^{1,*}, Ibai Agesta Pérez-Ilzarbe¹, Montse Cuadros¹, Mikel Zabala¹, Unai Guardia¹, Iñigo Amonarriz¹, Giovanni Paolini¹ and Izar Azpiroz¹

¹Fundación Vicomtech, Basque Research and Technology Alliance (BRTA)

Abstract

Natural hazards, including floods, wildfires, and landslides, pose severe threats to global infrastructure and human life, resulting in substantial economic losses. While recent advancements in Earth Observation (EO) and large vision-language models (VLMs) offer new avenues for geospatial analysis, bridging the gap between complex remote sensing data and actionable decision-making remains a challenge. This paper introduces DECIDE (Decision-making Engine for Catastrophe Insights and Disaster Evaluation), a web-based platform designed to enhance disaster monitoring through the integration of georeferenced interactive tools and an AI-powered chat assistant. By synthesizing remote sensing products with foundation models, DECIDE provides intuitive access to historical disaster data and real-time situational awareness, and in the near future, EO-derived thematic risk maps will also be integrated. The system facilitates rapid damage assessment and infrastructure loss estimation, supporting emergency responders and authorities in making high-stakes decisions. Furthermore, the DECIDE framework showcases how the integration of conversational AI can streamline the transition from research to production environments. This work underscores the potential of interactive, AI-driven dashboards to transform environmental research and disaster risk management.

Keywords

Natural Hazards, Earth Observation, Vision-Language Models, Natural Language Processing, Information Retrieval, Agents

1. Introduction

Natural hazards such as floods, wildfires, landslides, and frost cause significant economic losses worldwide and threaten human life and infrastructure. Floods alone affect more than 20% of the global population, while wildfires and landslides impact key sectors such as forestry, agriculture, tourism, and urban environments. Recent advances in automated Earth Observation (EO) change detection and large vision-language models enable new possibilities for analysing geospatial data and improving disaster monitoring. In this work, we propose DECIDE, a Decision-making Engine for Catastrophe Insights and Disaster Evaluation. DECIDE is a web-based platform that integrates georeferenced interactive tools with an interactive conversational agent to support the analysis of natural disasters before, during, and after their occurrence. By combining remote sensing products with language models, the system enables intuitive access to historical disaster data, thematic risk maps, and geospatial analyses. This approach facilitates damage assessment, supports infrastructure loss estimation, and improves situational awareness, ultimately helping authorities and emergency responders make faster and better-informed decisions. The implementation of this conversational interface, which communicates with generated remote sensing technologies will enable the transfer of this technology to any project where interactive dashboards are a requirement. One of the main novelties introduced in this system will be the implemen-

SEPLN 2026: 42nd International Conference of the Spanish Society for Natural Language Processing, León, Spain, 22-25 September 2026.

*Corresponding author.

✉ iperezd@vicomtech.org (I. Pérez-del-Olmo); iagesta@vicomtech.org (I. A. Pérez-Ilzarbe); mcuadros@vicomtech.org (M. Cuadros); mzabala@vicomtech.org (M. Zabala); uguardia@vicomtech.org (U. Guardia); iamonnarriz@vicomtech.org (I. Amonarriz); gpaolini@vicomtech.org (G. Paolini); iazpiroz@vicomtech.org (I. Azpiroz)

0009-0008-4582-2311 (I. Pérez-del-Olmo); 0009-0005-1751-0386 (I. A. Pérez-Ilzarbe); 0000-0002-3620-1053 (M. Cuadros); 0009-0006-1155-5302 (M. Zabala); 0009-0000-0064-0305 (U. Guardia); 0009-0002-6714-9523 (I. Amonarriz); 0000-0002-4309-281X (G. Paolini); 0000-0003-0401-8139 (I. Azpiroz)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

tation of an AI-driven dialogue system. This tool combines remote sensing and AI models to generate personalized thematic maps, technology that could be updated with cutting-edge models thanks to the AI chat assistant and foundation model derived products. Potential markets for the DECIDE system include public sector agencies involved in civil protection, emergency response, and urban planning, seeking innovative tools for enhanced risk management. Key partners include Basque Government agencies and Civil Protection services specialized in disaster management (e.g., DAEM), along with European research networks like Horizon Europe clusters, which support knowledge exchange and innovation.

This paper is organized as follows: Section 2 reviews related work, Section 3 describes the resources used in the system, Section 4 outlines the proposed methodology, and Section 5 presents the conclusions and future work.

2. Related work

The acceleration of climate change has established resilience and adaptation as critical priorities on the international agenda. According to the 2030 Climate Target Plan [1] and the EU Strategy on Adaptation to Climate Change [2], protecting natural capital and safeguarding citizens' health through cutting-edge technology is no longer optional but mandatory. Despite these efforts, the UNEP Adaptation Gap Report 2020 [3] highlights a significant maturity gap in systems designed to protect against climate-driven disasters such as droughts and floods.

At a regional level, this challenge is addressed through strategic frameworks like the RIS3 priority for Ecosystems and Sustainable Cities [4] and the Euskadi 2030 Science, Technology, and Innovation Plan (STIP) [5]. These initiatives emphasize the development of eco-innovation technologies and Decision Support Systems (DSS) as central pillars for disaster prevention and recovery.

Natural Language Processing (NLP) within the disaster management domain has undergone a profound transformation. Early breakthroughs were driven by Natural Language Understanding (NLU) models based on the Transformer architecture [6], such as BERT [7], which excelled at classification and information extraction. Subsequent iterations like Sentence-BERT [8] and encoder-decoder models such as T5 [9] further optimized semantic search and Information Retrieval (IR) from technical reports and drug-protein interaction datasets.

The current paradigm, however, has shifted toward Large Language Models (LLMs). Unlike purely discriminative models, LLMs enable sophisticated text generation and reasoning. This capability is transformative for emergency response: it allows for the synthesis of complex civil protection protocols into concise, actionable advice. Recent research confirms that integrating LLMs with knowledge graphs significantly enhances Emergency Decision-making [10], providing a level of contextual awareness and communicative flexibility that was previously unattainable.

The “State of the Art” in Earth Observation (EO) is currently defined by the emergence of Geospatial Foundation Models (FMs). Initial research into self-supervised learning, such as GASSL [11] and SatMAE [12], demonstrated the potential for early detection of environmental changes. More recent architectures like OFA-Net [13], DOFA [14], and Prithvi [15] have specialized these capabilities for disaster scenarios, including rapid flood mapping and wildfire detection [16].

The most disruptive frontier lies in Vision-Language alignment. Models such as TEOChat [17] represent the transition toward multimodal assistants. These systems allow end-users to interact with temporal satellite data through natural language, effectively bridging the gap between raw geospatial data and human decision-making. By combining LLM-driven text generation with precise thematic mapping, these platforms—exemplified by digital twin initiatives [18]—provide a unified interface for post-disaster damage assessment and proactive risk mitigation.

3. Resources

This section outlines the primary data resources utilized in the DECIDE platform. To create the local knowledge base, the Retrieval-Augmented Generation (RAG) [19] system integrates official catastrophe management plans from the Basque Country region, specifically focusing on flood and forest fire emergency protocols: “Plan especial de emergencias ante el riesgo de inundaciones de la comunidad autónoma del País Vasco”¹ and “Plan especial de emergencias por riesgo de incendios forestales de la Comunidad Autónoma Vasca”².

For real-time and historical meteorological data, the system interfaces with Basque Meteorology Agency (Euskalmet)³, which supplies critical environmental parameters ranging from severe weather alerts to localized sensor readings, which are essential to nourish the autonomous agent tasked with monitoring natural hazards.

4. Methodology

The following technologies, tools, and methods provide a robust framework for developing a dynamic, web-based Decision Support System (DSS) [20] integrated with an autonomous AI assistant. This assistant is designed to process information from both external APIs—providing real-time or historical data—and local sources.

To integrate external data, a dedicated middleware layer was developed using FastAPI⁴. This layer serves as an intermediary between the Euskalmet meteorological service and the AI assistant, standardizing the data retrieval process.

For local data handling, information is managed by a Retrieval-Augmented Generation (RAG) solution (see Section 4.2 and Figure 3), which retrieves the most relevant context from a database compiled from official documentation.

A critical capability of this framework is its ability to accurately identify user intent from natural language queries and route them to the appropriate data source. This orchestration is managed by an autonomous agent [21] that dynamically selects the most suitable tool from a custom suite to fulfill the user’s request.

4.1. Agentic orchestration

The proposed agent is architected as an autonomous decision-making entity capable of perceiving its environment (through available tools), processing contextual information, and dynamically planning and executing actions to fulfill user objectives. Unlike traditional static systems, this agent implements a cyclic, modular architecture that emulates human-like reasoning and iterative problem-solving. This workflow, illustrated in Figure 1, is triggered upon every user query and comprises several interconnected modules:

- **Context Management:** This module is responsible for interpreting the user’s initial textual input. The query is enriched and combined with the current conversation context and real-time temporal data. This context acts as a dynamic memory that orchestrates the workflow, stores the global state, and persists data across the current execution run.
- **Reasoning Engine (cognitive core):** Acting as the agent’s logical nucleus, this module analyzes the current state and decomposes complex user requests into actionable steps. It evaluates the available tools and determines the optimal sequence to answer the query. The core

¹https://www.euskadi.eus/contenidos/informacion/planes_inundaciones/es_doc/adjuntos/Plan_Inundaciones_discapacidades_2021_SIN_ANEXOS-es.pdf

²https://www.euskadi.eus/contenidos/informacion/planes_incendios/es_doc/adjuntos/Plan-emergencia-incendios-forestales_2025-es.pdf

³<https://www.euskalmet.euskadi.eus/el-tiempo/euskadi/>

⁴<https://fastapi.tiangolo.com/>




Figure 1: Agent's flowchart.

relies on a 8-billion-parameter Qwen 3.5 model [22] explicitly configured for function calling (`is_function_calling_model=True`). Crucially, the model's generation temperature is strictly set to 0.0 to ensure deterministic, highly predictable, and syntactically correct outputs when generating tool arguments, thereby minimizing hallucinations during the planning phase. Once the Reasoning Engine outputs a plan, the system evaluates the response payload, routing the workflow into one of three distinct paths:

- If the LLM determines that external information or actions are required, it outputs an array of functions to execute.
 - If the LLM assesses that it already possesses sufficient context to fulfill the user's request, it directly synthesizes the final response.
 - If the model fails to generate a valid plan or encounters an unresolvable constraint, an exception path is triggered to notify the user.
- **Execution Module (tools interface):** This module serves as the agent's operational gateway to the external environment, acting as a secure dispatcher for the actions orchestrated by the Reasoning Engine. The agent is equipped with a specialized toolkit categorized into two primary functional areas:
 - **External API Integrations:** A suite of functions designed to retrieve real-time and forecasted environmental data. This includes meteorological alerts, oceanographic conditions, tidal information, and localized or regional weather data, as well as granular sensor readings from meteorological stations. To ensure robustness, these functions utilize pre-configured default parameters for technical variables that the user may not specify in their initial query. This design allows for seamless execution of general requests while remaining flexible; if the user provides a more detailed or insightful query, the system dynamically overrides these defaults with the user's specific parameters.
 - **Local Knowledge Retrieval:** A RAG tool configured to query a specialized local vector database. This tool provides access to official catastrophe management protocols and emergency response plans, ensuring the agent's advice is grounded in authoritative documentation.

During execution, the module iterates through the requested tools, dynamically mapping them to their corresponding Python functions and schemas. Crucially, the execution is enclosed within a safe-handling mechanism; if a tool fails or throws an exception, the system does not crash. Instead, it captures the error message and formats it as a tool response, allowing the LLM to

understand the failure and adapt its subsequent planning. The raw outputs generated by these tools—whether successful data retrieval or error messages—are systematically appended to the contextual memory.

Once the Action Module has executed all the tools scheduled in the current iteration, the generated outputs are not immediately forwarded to the user. Instead, the updated conversational memory, now enriched with the external data, is fed back into the Reasoning Engine. The LLM evaluates this updated state to determine the next phase of the operation. If the collected data reveals gaps or requires further investigation, the model will generate a new set of tool calls, initiating another execution loop. Conversely, if the reasoning engine concludes that the acquired information is sufficient to resolve the user’s initial query, it halts the tool-calling cycle, synthesizes the complex raw data into a coherent, human-readable response, and delivers the final plain text to the user.

4.2. Retrieval Augmented Generation

The Retrieval-Augmented Generation function allows users to retrieve specific domain information with low latency and near-zero hallucination. A robust architecture has been developed to process and extract information from regional catastrophe plans. As illustrated in Figure 2, the RAG logic is divided into three primary modules.




Figure 2: RAG’s flowchart.

Ingestion module: The objective of this module is to create a vector database containing the information from the official catastrophe documents. This is a one-time execution process; once the data is loaded, it remains persistently accessible. The documents are split into page-sized chunks. Each chunk is processed independently to create embeddings using the *jina-embeddings-v4* [23] model, which converts both text and images within the chunk into a 2048 dimensional vector. Finally, a Qdrant ⁵

⁵<https://qdrant.tech/>

vector database is created to store these representations.

Information Retrieval Module: This module identifies the most relevant sections of the documentation with respect to a user's query. To achieve this, an embedding of the query is generated using the same encoder *jina-embeddings-v4* utilized in the ingestion phase. The system then selects the K chunks whose embeddings are closest to the query's embedding. In this implementation, K is set to 10.

Inference module: The objective of this module is to generate a response that satisfies the query based on the relevant information obtained from the retrieval module. The retrieved information is analyzed using a multimodal Qwen [24] family model, specifically *Qwen3-VL-8B-Instruct*⁶. This model is tasked with answering the query using only the provided context extracted from the official emergency plans. By grounding the response strictly in this retrieved documentation, the architecture effectively mitigates model hallucinations.




Figure 3: Response from DECIDE when asked for information in the official catastrophe management plans.

⁶<https://huggingface.co/Qwen/Qwen3-VL-8B-Instruct>

4.3. Environmental data acquisition

To provide the DECIDE agent with real-time and historical environmental data, a dedicated middleware API was developed. Rather than having the agent query Euskalmet directly, this intermediate layer abstracts and standardizes access to external data. It translates complex meteorological services into simple, unified tool calls that the agent can easily interpret and execute.

4.3.1. Endpoints and Data Typology

The middleware, built using FastAPI, exposes a set of POST endpoints designed to function as standard tools for the agent. The available endpoints and their corresponding data outputs are:

- Alerts: weather warnings for a specific area and date.
- Calendar: astronomical info like sunrise, sunset, and moon phases.
- Ocean: marine forecasts (waves, visibility, water temperature).
- Station readings: hourly sensor data from a specific station.
- Station data: daily aggregated data from a station.
- Summarized readings: daily summary for a specific sensor at a station.
- Tides: tide information by geographic area.
- Weather (location): detailed weather for a specific place (see Figure 4).
- Weather (region): broader, regional weather conditions.

The raw data retrieved from Euskalmet is pre-processed and filtered to extract only the most relevant mission-critical features. This refined data is then serialized into JSON format and injected into the agent's context, providing a clean and structured basis for autonomous decision-making without exceeding token limits.

4.3.2. Geospatial Resolution

A key capability of the DECIDE framework is its support for natural language location entities. Rather than relying exclusively on exact geographic coordinates, the agent can process common location names extracted from user prompts or external reports. To achieve this, the system integrates the Mapbox API⁷ for geocoding, converting textual location names into accurate latitude and longitude coordinates.

Several endpoints accept these coordinates alongside a search radius. To determine whether a specific meteorological event or sensor station falls within this target area, the system automatically handles spatial queries using projected coordinate systems. Specifically, it projects coordinates into UTM Zone 30N (EPSG:25830), ensuring metric-accurate distance calculations for the region.

The spatial filtering accommodates two distinct scenarios:

- Point-based data (e.g., weather stations): The system verifies if the exact coordinates of the meteorological station fall within the defined circular search radius around the target location.
- Area-based phenomena (e.g., regional alerts): If the meteorological event encompasses a wider geographical area represented by a polygon (a set of points), the system performs a polygon intersection. It evaluates whether the circular bounding area of the target location overlaps with the polygon defining the meteorological phenomenon.

This spatial filtering allows the agent to intuitively query data for disaster-affected locations using natural language, while the middleware handles the geospatial calculations in the background.

⁷<https://www.mapbox.com/>




Figure 4: Response from DECIDE when asked for meteorological data in a concrete location.

4.3.3. Authentication and Performance

Authentication with the official Euskalmet service is handled securely through JWT tokens [25] signed with RS256 using a private RSA key [26]. To maximize performance and reduce response latency, especially when fetching data over multi-day ranges or multiple zones, all HTTP requests are executed in parallel using a thread pool. Results are aggregated and returned as a single JSON response, which the agent can immediately process and integrate into its final response to the user.

4.4. Earth Observation data module

This module will focus on collecting and preprocessing data from various sources, including remote sensing tools like Sentinel-1 ⁸, Sentinel-2 ⁹, and Copernicus EMS ¹⁰. This data will be transformed into a georeferenced, data cube format that is suitable for an AI model. Optical indices (NDVI, NBR, and mDVI) enable the characterization of the phenological and hydrological status of forest and natural ecosystems. The analysis of their absolute values, anomalies, and temporal dynamics is essential for identifying areas vulnerable to fires and floods. Recently, a generative model was created to estimate NDVI in cloudy areas using Sentinel-2 data [27]; this methodology is being extrapolated to improve the temporal resolution of other indices of interest.

A dataset of paired images from the Sentinel-1 and Sentinel-2 missions (cloud-free) is currently being compiled, with a maximum time lag of 24 hours. The study areas are located in the Basque Country and were selected for their representativeness of the region's hydrological and forest ecosystems. This dataset will enable the training of generative models designed to estimate optical indices of interest, thereby ensuring a six-day temporal resolution regardless of atmospheric conditions.

These developed functionalities—ranging from generative models to the processing of vegetation indices and thematic mapping—will be exposed via a FastAPI service. In accordance with the endpoint specifications described in Section 4.3.1, this implementation will enable the agent to autonomously invoke those capabilities based on the operational context.

5. Conclusions and Future Work

This paper presents DECIDE, decision-making Engine for Catastrophe Insights and Disaster Evaluation. The system proposes the integration of Large Language Models with Earth Observation data to obtain a robust and intuitive Decision Support System for disaster management. By utilizing a RAG-based architecture and a deterministic reasoning engine, the system can mitigate hallucinations, ensuring that all tactical advice is strictly grounded in official emergency protocols and real-time environmental data. Furthermore, the implementation of generative models to overcome cloud-cover limitations ensures operational continuity, providing responders with reliable vegetation and hydrological indices regardless of atmospheric conditions. Moving forward, future work will prioritize the integration of Vision-Language Models for direct, autonomous interpretation of satellite imagery and the development of predictive modules for wildfire and flood propagation. Finally, the framework will undergo field validation with the Directorate of Emergency Management (DAEM) to refine its performance in high-stress scenarios while scaling the system to support broader European civil protection networks.

Acknowledgments

The DECIDE system was fully supported by Vicomtech through an internal funding program designed to promote interdisciplinary collaboration and accelerate the development of innovative research outcomes.

Declaration on Generative AI

During the preparation of this work, the authors used Gemini in order to Grammar and spelling check. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

⁸<https://sentinels.copernicus.eu/copernicus/sentinel-1>

⁹<https://sentinels.copernicus.eu/copernicus/sentinel-2>

¹⁰<https://emergency.copernicus.eu/>

References

- [1] European Commission, Stepping up Europe's 2030 climate ambition - Investing in a climate-neutral future for the benefit of our people, Technical Report, European Commission, 2020. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52020DC0562>.
- [2] European Commission, Forging a climate-resilient Europe - the new EU Strategy on Adaptation to Climate Change, Technical Report, European Commission, 2021. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52021DC0082>.
- [3] United Nations Environment Programme, Adaptation gap report 2020, 2021. URL: <https://www.unep.org/resources/adaptation-gap-report-2020>.
- [4] Eusko Jaurlaritza, Euskadi ris3, 2024. URL: <https://www.euskadi.eus/informacion/ris3/web01-a2kulind/es/>.
- [5] Eusko Jaurlaritza, EUSKADI 2030 - Science, Technology and Innovation Plan, Technical Report, 2021. URL: https://www.euskadi.eus/contenidos/informacion/pcti_euskadi_2030/es_def/adjuntos/EUSKADI-2030-STIP.pdf.
- [6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Advances in neural information processing systems* 30 (2017).
- [7] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019, pp. 4171–4186.
- [8] N. Reimers, I. Gurevych, Sentence-BERT: Sentence embeddings using siamese BERT-networks, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.
- [9] X. Sui, W. Wang, J. Zhang, Text mining drug-protein interactions using an ensemble of BERT, sentence BERT and T5 models, *bioRxiv* (2021).
- [10] M. Chen, Z. Tao, W. Tang, T. Qin, R. Yang, C. Zhu, Enhancing emergency decision-making with knowledge graphs and large language models, *International Journal of Disaster Risk Reduction* 113 (2024) 104804.
- [11] K. Ayush, et al., Geography-aware self-supervised learning, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 10181–10190.
- [12] Y. Cong, et al., Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery, in: *Advances in Neural Information Processing Systems*, volume 35, 2022, pp. 197–211.
- [13] Z. Xiong, et al., One for all: Toward unified foundation models for earth vision, *arXiv preprint* (2024).
- [14] Z. Xiong, et al., Neural plasticity inspired multimodal foundation model for earth observation, *arXiv preprint* (2024).
- [15] D. Szwarcman, S. Roy, P. Fraccaro, O. E. Gíslason, B. Blumenstiel, R. Ghosal, P. H. De Oliveira, J. L. de Sousa Almeida, R. Sedona, Y. Kang, et al., Prithvi-eo-2.0: A versatile multi-temporal foundation model for earth observation applications, *IEEE Transactions on Geoscience and Remote Sensing* (2025).
- [16] S. Lu, et al., AI foundation models in remote sensing: A survey, *arXiv preprint arXiv:2408.03464* (2024).
- [17] J. A. Irvin, E. R. Liu, J. C. Chen, I. Dormoy, J. Kim, S. Khanna, et al., TeoChat: A large vision-language assistant for temporal earth observation data, *arXiv preprint arXiv:2410.06234* (2024).
- [18] Vicomtech, Gemelo digital chs, 2024. URL: <https://www.vicomtech.org/es/idi-tangible/casos-exito/caso/gemelo-digital-chs>.
- [19] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al., Retrieval-augmented generation for knowledge-intensive nlp tasks, *Advances in neural information processing systems* 33 (2020) 9459–9474.
- [20] J. D. Little, Models and managers: The concept of a decision calculus, *Management science* 16 (1970) B–466.

- [21] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. R. Narasimhan, Y. Cao, React: Synergizing reasoning and acting in language models, in: The eleventh international conference on learning representations, 2022.
- [22] Q. Team, Qwen3. 5: Towards native multimodal agents, URL: <https://qwen.ai/blog> (2026).
- [23] M. Günther, S. Sturua, M. K. Akram, I. Mohr, A. Ungureanu, B. Wang, S. Eslami, S. Martens, M. Werk, N. Wang, et al., jina-embeddings-v4: Universal embeddings for multimodal multilingual retrieval, in: Proceedings of the 5th Workshop on Multilingual Representation Learning (MRL 2025), 2025, pp. 531–550.
- [24] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang, et al., Qwen technical report, arXiv preprint arXiv:2309.16609 (2023).
- [25] M. Jones, J. Bradley, N. Sakimura, Json web token (jwt), Technical Report, 2015.
- [26] R. L. Rivest, A. Shamir, L. Adleman, A method for obtaining digital signatures and public-key cryptosystems, Communications of the ACM 21 (1978) 120–126.
- [27] P. Soglia, P. Gonzalez, I. Azpiroz, U. Otamendi, M. Quartulli, S. Salata, Estimating ndvi from sar images using conditional generative adversarial networks, in: IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium, IEEE, 2023, pp. 5222–5225.