



HAL
open science

Encapsulating textual contents into a MOC data structure for advanced applications

Giuseppe Greco, Thomas Boch, Pierre Fernique, Manon Marchand, Mark G. Allen, François-Xavier Pineau, Matthieu Baumann, Marco Molinaro, Roberto de Pietri, Marica Branchesi, et al.

► To cite this version:

Giuseppe Greco, Thomas Boch, Pierre Fernique, Manon Marchand, Mark G. Allen, et al.. Encapsulating textual contents into a MOC data structure for advanced applications. *Astron.Comput.*, 2026, 54, pp.101014. <10.1016/j.ascom.2025.101014>. <hal-05301085>

HAL Id: hal-05301085

<https://hal.science/hal-05301085v1>

Submitted on 5 Nov 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Encapsulating Textual Contents into a MOC data Structure for Advanced Applications

Giuseppe Greco^a, Thomas Boch^b, Pierre Fernique^b, Manon Marchand^b, Mark Allen^b, Francois-Xavier Pineau^b, Matthieu Baumann^b, Marco Molinaro^c, Roberto De Pietri^{d,e}, Marica Branchesi^{f,g}, Steven Schramm^{h,i}, Gergely Dályaj^j, Elahe Khalouei^k, Barbara Patricelli^l, Giulia Stratta^m

^aINFN Sezione di Perugia, via A. Pascoli, 23c, I-06123 Perugia, Italy

^bUniversité de Strasbourg, CNRS, Observatoire astronomique de Strasbourg, UMR 7550, 67000 Strasbourg, France

^cINAF - Osservatorio Astronomico di Trieste, 34143 Trieste, Italy

^dINFN, Gruppo Collegato di Parma, Parco Area delle Scienze, 43124 Parma, Italy

^eUniversità di Parma, Parco Area delle Scienze, 43124 Parma, Italy

^fGran Sasso Science Institute (GSSI), L'Aquila, Italy

^gINFN - Laboratori Nazionali del Gran Sasso, 67100 Assergi, Italy

^hDPNC, Section de Physique, Université de Genève, Geneva, Switzerland

ⁱGravitational Wave Science Center (GWSC), Université de Genève, Geneva, Switzerland

^jL2IT, Laboratoire des 2 Infinis - Toulouse, Université de Toulouse, CNRS/IN2P3, UPS, F-31062 Toulouse Cedex 9, France

^kAstronomy Research Center, Research Institute of Basic Sciences, Seoul National University, 08826 Seoul, South Korea

^lPhysics Department, University of Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy

^mINAF - OAS Bologna, Via P. Gobetti 93/3, 40129 Bologna (BO), Italy

Abstract

Context. The Multi-Order Coverage map (MOC) is a widely adopted standard promoted by the International Virtual Observatory Alliance (IVOA) to support data sharing and interoperability within the Virtual Observatory (VO) ecosystem. This hierarchical data structure efficiently encodes and visualizes irregularly shaped regions of the sky, enabling applications such as cross-matching large astronomical catalogs, visualizing multi-wavelength and multi-messenger surveys, and facilitating collaborative research through seamless interoperability in big-data-driven exploration.

Aims. This study aims to explore potential enhancements to the MOC data structure by encapsulating textual descriptions and semantic embeddings into sky regions. Specifically, we introduce "Textual MOCs", in which textual content is encapsulated, and "Semantic MOCs" that transform textual content into semantic embeddings. These enhancements are designed to enable advanced operations such as similarity searches and complex queries and to integrate with generative artificial intelligence (GenAI) tools to improve context-aware interactions and response accuracy in astronomical data analysis, and support agent-based applications.

Method. We experimented with Textual MOCs by annotating detailed descriptions directly into the MOC sky regions, enriching the maps with contextual information suitable for interactive learning tools. For Semantic MOCs, we converted the textual content into semantic embeddings, numerical representations capturing textual meanings in multidimensional spaces, and stored them in high-dimensional vector databases optimized for efficient retrieval.

Results. The implementation of Textual MOCs enhances user engagement by providing meaningful descriptions within sky regions, facilitating the development of effective game-based learning. Semantic MOCs enable sophisticated query capabilities, such as similarity-based searches and context-aware data retrieval, enhancing astronomical data analyses. Integration with multimodal generative AI systems allows for more accurate and contextually relevant interactions supporting both spatial, semantic and visual operations for advancing astronomical data analysis capabilities. Through straightforward examples, we discuss the fundamentals of this new experimental implementation.

Keywords: Virtual Observatory, Multi Order Coverage map, IVOA Standards, Textual MOC, Semantic MOC

Introduction

The International Virtual Observatory Alliance (IVOA)¹ is an international organization that develops and promotes open technical standards to enable seamless interoperability of astronomical data and tools within the Virtual Observatory (VO, Berriman et al., 2021). Founded in 2002, IVOA is a free and open consortium that brings together scientific communities,

research institutions, and software developers. It serves as a collaborative platform where members debate and define guidelines for data management, exchange, and analysis, ensuring that astronomical data is accessible and usable across different systems and platforms (Allen et al., 2004).

The Multi-Order Coverage Map (MOC) is a data structure recommended by the IVOA organization that provides a standardized method for mapping and representing sky regions based on the HEALPix sky tessellation (Górski et al., 2005), promoting interoperability across astronomical datasets

¹<https://ivoa.net/>

(Schaaff, 2007).

MOCs are highly efficient for representing complex sky regions at multiple levels of resolution, enabling rapid logical operations, seamless data processing, and fast querying capabilities. These advantages are further supported by IVOA data providers, who facilitate MOC-based queries, enabling swift access and retrieval of relevant astronomical data (Fernique et al., 2014, 2019). As a common framework for mapping and querying astronomical datasets, MOC acts as a technical standard within the IVOA ecosystem, enhancing the FAIR (Findability, Accessibility, Interoperability, and Reusability) principles, which are crucial for realizing the vision of a fully integrated and collaborative Virtual Observatory (O’Toole and Tocknell, 2024).

Due to these properties, MOCs have been widely adopted in various fields of astronomical research (Petter et al., 2023; Carvajal et al., 2021; Smareglia et al., 2019; Smith et al., 2017), proving particularly significant in the emerging field of multi-messenger astronomy, especially for handling gravitational-wave sky localizations (Greco et al. 2019, 2022a; 2022b, Singer et al., 2022, Coughlin et al., 2023, Hosseinzadeh et al., 2024). Furthermore, MOCs are extensively utilized in interactive data visualization tools, such as Aladin Lite (Boch and Fernique, 2014) and Aladin Desktop (Bonnarel et al., 2000), and have been integrated into online platforms like GWOPS (Xu et al., 2020), ESASky (Merín et al., 2017), GLADenet (Brozzetti et al., 2024), Treasure Maps (Wyatt et al., 2020), GWsky², and Fink³ (Möller et al., 2021), among others.

These interactive visualizations are enhanced by the ability to overlay a wide range of images or catalogs using the IVOA HiPS (Hierarchical Progressive Surveys) standard. HiPS is a widely used format in astronomy for representing and visualizing large datasets, such as images, catalogs and three-dimensional data cubes. HiPS technology divides spatial data into hierarchical tiles with progressive resolution, allowing efficient zooming and exploration of finer details as needed (Fernique et al., 2015).

MOC and HiPS are designed to be highly interoperable with each other. This interoperability allows users to seamlessly integrate MOC maps with HiPS data within visualization tools. Users can display MOCs alongside HiPS layers and customize graphic properties (e.g., transparency, outlines, color), enabling a more personalized and interactive exploration of astronomical data and facilitating skymap comparisons (Greco et al., 2022b).

In this article, we report on our experimentation with a possible extension of the MOC standard to incorporate textual information, enabling simultaneous spatial and semantic operations. This new data structure, called "Textual MOCs," allows for the association of detailed descriptions with spatial data, enriching the context of astronomical information. This work is motivated by the possibility that these Textual MOCs can be transformed into "Semantic MOCs" by converting the textual content into semantic embeddings—numerical vectors that encapsulate the meaning of the text within high-dimensional

spaces (Mikolov et al., 2013; Pennington et al., 2014; Devlin et al., 2018). This transformation process unlocks significant potential for training generative AI models (GenAI), including large language models (LLM)-based systems (Goodfellow et al., 2014; Brown et al., 2020; Thoppilan et al., 2022; Zhang et al., 2022). Semantic MOCs have the potential to significantly enhance query comprehension enabling advanced applications, such as retrieval-augmented frameworks (Lewis et al., 2020; Guu et al., 2020; Karpukhin et al., 2020), multi-agent orchestration (Tran et al., 2025), and multimodal approaches (Munikoti et al., 2024). These references represent just a few examples from the extensive literature on the topic.

Each section in which the article is structured is complemented by a Jupyter notebook (Kluyver et al., 2016) with sample code, available for download to facilitate practical exploration and implementation of these techniques⁴.

For large-scale operational applications, we refer to more detailed studies to be presented in future work, as this article primarily establishes the foundational concepts and outlines the key ideas of this approach.

1. From MOCs to Textual MOCs

In this section, we provide a brief overview of the basic concepts for creating MOC maps. Over time, the MOC structure has evolved to incorporate both spatial and temporal dimensions, leading to the development of the Space-Time MOC (ST-MOC) (Fernique et al., 2021). However, in this work, we focus exclusively on MOCs with spatial indexing and will generally refer to (Space) MOCs in all examples. The textual/semantic extension of ST-MOCs is not presented here for brevity; as future work, we plan to employ it in agent-based applications that verify the availability of temporal information and enable joint spatio-temporal operations.

1.1. MOC data structure and serializations

The MOC data structure is based on the HEALPix sky tessellation algorithm (Górski et al., 2005). HEALPix partitions the sphere into equal-area pixels, making it highly suitable for representing astronomical data on the celestial sphere. In this algorithm, the sky is divided hierarchically into equal-area pixels, starting from twelve base pixels at the coarsest resolution level. Each subsequent level of resolution, indicated by the HEALPix parameter *order*, divides each pixel into four smaller ones.

Each MOC cell is defined by two properties:

- **Hierarchy Level** (*order*): Represents the level of refinement or resolution of the tessellation.
- **Pixel Index** (*npix*): Uniquely identifies a pixel on the sky at a given resolution.

The finest level of refinement in the MOC hierarchy is determined by the *order* parameter, allowing for an efficient representation of arbitrarily shaped regions of the sky with varying

²<https://virgo.pg.infn.it/maps/>

³<https://fink-portal.org/gw>

⁴<https://github.com/ggreco77/TextualMOC>

levels of detail. The HEALPix orders typically range from 0 to 29, allowing for a mean cell resolution from 58.63 degrees down to 393.2 μas .

To facilitate interoperability and ease of use across different platforms and applications, Multi-Order Coverage (MOC) supports multiple serialization formats. These formats enable the encoding and decoding of MOC data structures for storage, transmission, and visualization purposes. Two serialization formats are supported: (i) FITS (Flexible Image Transport System) and (ii) JSON (JavaScript Object Notation), as defined in the IVOA Recommendation *MOC: Multi-Order Coverage map Version 2.0 (2022-07-27)* (Fernique et al., 2022). Although JSON is not a normative serialization format in the MOC standard, it is widely used in practice and has already been employed in applications such as the tutorial by Greco & Marchand to compute the visibility of a MOC from a ground-based observatory as a function of airmass⁵. In this work, we adopt JSON serialization when extensions are required, with the intention of extending our approach to the other serialization formats in future developments.

To store a MOC in a FITS file, each cell is converted into a single integer:

$$\text{uniq} = 4 \times 4^{\text{order}} + \text{npix} \quad (1)$$

A MOC, when serialized in JSON format, follows the following syntax:

Snippet 1: JSON structure of a Space MOC

```
{
  "order_n": [npix_i | where npix_i uniquely
              defined integers],
  // ... continue for other orders ...
}
```

In this JSON structure, shown in Snippet 1, each `order_n` represents a different resolution level, with `npix_i` indicating the pixel indices covering specific sky regions at that order.

1.2. Textual MOC data structure

For the creation of Textual MOCs, we decide to use the JSON serialization format. JSON is a lightweight, human-readable data interchange format that is language-independent and features a simple syntax, making it ideal for transmitting data over the web. Its self-describing nature and minimal structure make JSON an excellent choice for efficiently encoding and sharing information.

Snippet 2 illustrates the JSON structure of a Textual MOC. In addition to the first two entries (`order_n` and `npix_i`) discussed in Subsection 1.1, text-based MOCs contain the following new elements:

text Provides a textual description of the content and scientific relevance of the celestial region or dataset.

metadata Contains additional information such as the author, creation date, and other key-value pairs for metadata enrichment.

annotated_cells Maps specific MOC orders and pixels to user-defined annotations, allowing labels to be attached to individual sky cells.

While the `text` key provides a global annotation across the entire MOC coverage, the `annotated_cells` key facilitates the assignment of labels to specific MOC cells. In multi-agent workflows, the system can be instructed to identify particular objects referenced in the text, retrieve their coordinates from VO services such as SIMBAD, and annotate the corresponding cells. In addition, the MOC resolution can be increased to achieve more detailed coverage when required.

A possible use case can be considered when, within the localization region of a gravitational wave event, a set of candidate objects is identified. These candidates can be inserted into the keyword `annotated_cells` to enable specific searches in databases such as Vera Rubin, or to check for coincident high-energy events reported by satellite observatories. This type of annotation enables automated searches, for example on kilonova events or short gamma-ray bursts, because the textual information embedded in a MOC allows the search to be properly contextualized within CBC (Compact Binary Coalescence) gravitational-wave events.

Snippet 2: JSON structure of a Textual MOC with text, metadata, and cell annotations

```
{
  "order_n": [npix_i | where npix_i uniquely
              defined integers],
  // ... continue for other orders ...
  "text": "Your textual description here",
  "metadata": {
    "key1": "value1",
    "key2": "value2",
    // ... continue metadata entries ...
  },
  "annotated_cells": {
    "order_n": {
      "npix_i": "annotation_i"
      // ... continue for other pixels ...
    }
    // ... continue for other orders ...
  }
}
```

In the accompanying tutorial, we introduce basic Python functions to encapsulate textual information within MOCs, demonstrating how to directly include this data in the JSON representation. These functions utilize the `mocpy` (Boch, 2019), `matplotlib` (Hunter, 2007), and `astropy` (Ginsburg et al., 2019) libraries, along with Python's native `json` library⁶.

⁵https://cds-astro.github.io/tutorials/7_Multi-messenger_astronomy__Planning_observations_rapidly_with_MOCs.html

⁶<https://docs.python.org/3.12/library/json.html>

2. The Role of MocServer in Creating Textual MOCs

In this section, we recall the main features of MocServer, a server specifically dedicated to MOCs. In our context, this is particularly valuable as it allows us to create textual MOCs with metadata curated by archival scientists, ensuring both accuracy and reliability. We present an example application showcasing both the embedded textual content within the MOC and its spatial component simultaneously. MOCs can describe the spatial coverage of any dataset; when generated from images, the corresponding images in HiPS format can be retrieved. To render the spatial component, we use `ipyaladin`, a Jupyter/IPython widget that enables interactive visualization of HiPS and MOCs directly within notebook environments (Boch and Desroziers, 2020).

2.1. MocServer

A dedicated directory service, MocServer⁷, is specifically designed for the storage, curation, and preservation of MOC data. MocServer ensures that MOC datasets are securely maintained, while also facilitating easy access and retrieval for further analysis and use.

The MocServer is developed by the Strasbourg Astronomical Data Centre (CDS) in 2015, designed to efficiently manage collections of spatial and temporal data encoding with MOC data structures. Its primary objective is to quickly provide a list of relevant astronomical resources for a given region of the sky, with response times in milliseconds. MocServer stores dataset names, each linked with a MOC (Multi-Order Coverage) spatial coverage map and detailed metadata.

This service is widely utilized by various tools, such as Aladin Desktop⁸, Aladin Lite⁹, and the CDS portal¹⁰, which leverage MocServer’s capabilities to provide quick access to a wealth of astronomical data. As of 2023, MocServer contains more than 32,000 entries, underscoring its significance within the astronomical community.

This service can be accessed remotely through an HTTP API. Additionally, a Python module is available on GitHub as part of the `astroquery` (Ginsburg et al., 2019), providing a convenient interface for integrating this functionality into Python-based workflows.

The result of such a query can go beyond an identifier-only list, extending to the properties themselves and effectively offering a "metadata directory service." The use of MocServer keys is based on the IVOA model and relies extensively on its vocabulary. A comprehensive list of metadata keywords managed by the MocServer, along with their occurrence counts and an example of the associated value for each keyword, can be obtained through the API. This is achieved by appending the parameter `get=example` to the MocServer URL, as shown in the following example: <https://alaska.cds.unistra.fr/MocServer/query?get=example>.

Furthermore, we emphasize that a lightweight and flexible local infrastructure for MOC-based coverage queries can be built with `textttMOC-Set`¹¹, which stores only the MOC regions without the associated metadata provided by the MocServer. This limitation can be readily addressed by storing metadata separately (e.g. SQLite, PostgreSQL, MongoDB, or other lightweight database systems) and retrieving them via the list of MOC identifiers returned by MOC-Set queries. A detailed analysis of these technologies lies beyond the scope of this work; here we simply demonstrate the broad support of the MOC standard within the VO ecosystem and show how the proposed extension can be seamlessly integrated into it.

2.2. Textual MOC powered by MocServer

We have experimented with the MocServer as a web tool for creating Textual MOCs because it provides access to a wide range of curated datasets with reliable references and detailed metadata. This access allows researchers to incorporate authoritative data into their MOCs, enhancing them with scientifically accurate and comprehensive information.

Figure 1 illustrates how a Multi-Order Coverage (MOC) map can be enriched with textual content retrieved from the MocServer, enabling simultaneous visualization of both textual and spatial components.

As an illustrative example, we retrieve the MOC for the dataset `CDS/P/Euclid/Q1/color` via MocServer, which provides both the spatial coverage map (MOC) and its associated metadata. This dataset originates from the Euclid mission provided by the ESA (European Space Agency). In this case, the metadata field `obs_description` is utilized to create the textual MOC.

3. Astronomical Platforms for Textual MOC Generation

Other astronomical platforms can be utilized for the creation of Textual MOCs, such as the General Coordinates Network (GCN)¹², the Astronomer’s Telegram (ATel)¹³, the Transient Name Server (TNS) (Gal-Yam, 2021), the Astrophysical Multimessenger Observatory Network (AMON) (Ayala Solares et al., 2020) and follow-up pointings in the Gravitational Wave Treasure Map (Wyatt et al., 2020).

As a user case, consider displaying 90% of a credible region of a gravitational-wave sky localization simultaneously with an associated alerts of the LIGO, Virgo and KAGRA collaborations (LVK)¹⁴ via The General Coordinates Network (GCN). The GCN Circular text is retrieved through the site’s functionality, which provides access to it in both text and JSON formats. The method for generating MOCs of sky localization area of gravitational-wave events is described in detail by (Greco et al., 2022b). In this work, we extend that approach by integrating textual data with spatial regions through Textual MOCs, combining the corresponding GCN Circular(s) with the corresponding sky localization MOC map.

⁷<https://alaska.cds.unistra.fr/MocServer/query>

⁸<https://aladin.cds.unistra.fr/AladinDesktop/>

⁹<https://aladin.cds.unistra.fr/AladinLite/>

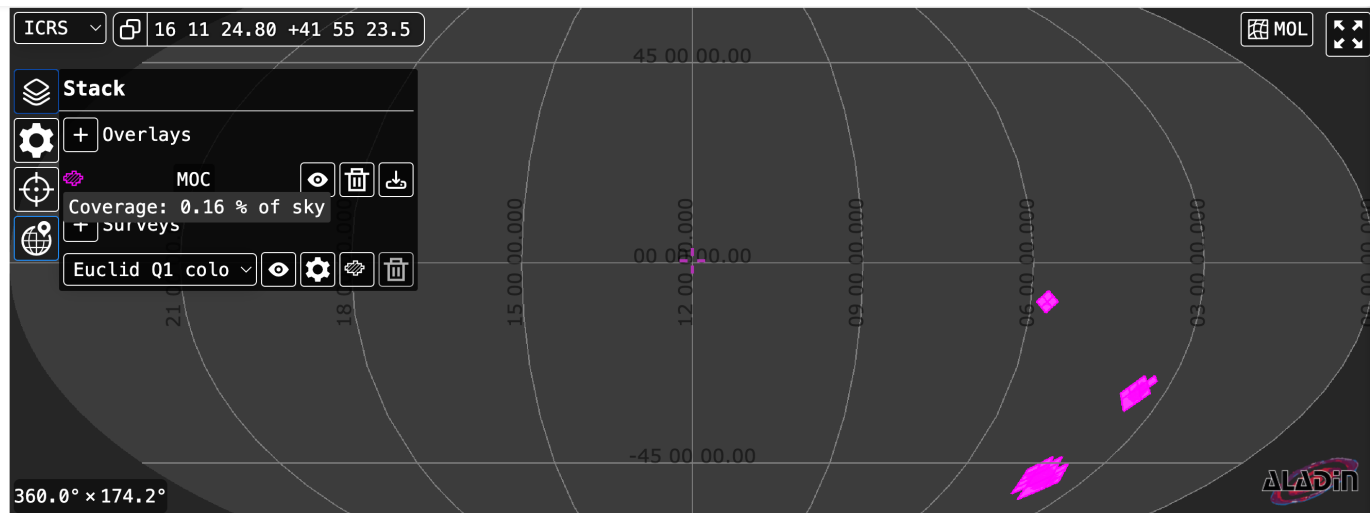
¹⁰<http://cdsportal.u-strasbg.fr/>

¹¹<https://pypi.org/project/moc-set/>

¹²<https://gcn.nasa.gov/>

¹³<https://www.astronomerstelegam.org/>

¹⁴<https://emfollow.docs.ligo.org/userguide/>



Text: The Euclid Q1 data release encompasses $\approx 65 \text{ deg}^2$ and contains the VIS and NIR data of one pass on each of the three Euclid deep fields (EDF-North, EDF-South, and EDF-Fornax), as well as data on Lynds Dark Nebula LDN1641 in the Orion star-formation region. The release aims to showcase the scientific potential of Euclid data, help the community prepare for the upcoming DR1, and support early science projects spanning topics from cosmology to the Solar System. In the 3 Euclid deep fields, this HiPS has been produced using bands NISP.H, NISP.Y and VIS for creating the color, and VIS for the luminance. In the Orion star-formation region, this HiPS has been produced combining infrared bands NISP.H, NISP.J and NISP.Y

Figure 1: An illustration showing how a Multi-Order Coverage (MOC) map can be enriched with textual content retrieved from MOCServer, allowing the simultaneous visualization of both spatial and textual components. The MOC for the dataset CDS/P/Euclid/Q1/color (Euclid Quick Data Release 1) from the Euclid space survey mission is displayed, with metadata (specifically the `obs_description` field) used to generate a textual MOC. The interactive visualization is provided by ipyaladin, a Jupyter/IPython widget, showing both MOCs and their corresponding image datasets in HiPS format.

4. Enhancing Data Exploration with Textual MOCs and Aladin Lite

In Section 2, we have explored the use of ipyaladin (a Jupyter/IPython widget) to simultaneously display the spatial and textual content of MOCs. Here, we utilize Aladin Lite in combination with Textual MOCs to develop web applications that enhance interactivity.

We use version 3 of Aladin Lite, which introduces a major overhaul of the platform by leveraging GPU acceleration through WebGL enhancing performance and interactivity (Bauermann et al., 2022).

4.1. Contextualization and Interactivity

As a minimal example, we show how to integrate and utilize Textual MOC data structure within a web page using Aladin Lite. As shown in Figure 2, the HTML structure is designed to include essential elements such as a container for the Aladin Lite viewer, where the MOC map is displayed, and an informative pop-up that appears when the cursor hovers over a MOC region. The information depicted in the pop-up is extracted from the textual component of the MOC, enabling simultaneous access to both spatial and descriptive data.

Additionally, it manages sound effects, providing auditory feedback when the mouse enters or exits a MOC region, thereby enhancing the application’s overall interactivity. Users can then interact with the Aladin Lite viewer, explore the map, view contextual information in the pop-up, and control sound effects via the dedicated menu. This implementation effectively combines

visual, textual, and auditory elements, creating an immersive and interactive experience for exploring data within a Textual MOC structure.

Both MOCs and HiPS in Aladin Lite utilize level-of-detail (LOD) techniques, dynamically adjusting the resolution based on zoom level. Lower resolution is used for large fields of view to ensure smooth interaction, while higher resolution is progressively loaded as users zoom in, optimizing both performance and data visualization. This flexibility supports efficient comparison of datasets and simultaneous analysis of multiple layers.

4.2. Learning Game and Public Engagement

Building upon the previous example, we developed an interactive game demonstrator that guides users in locating the Large Magellanic Cloud (LMC) within the sky. The game challenges users to navigate through a sky map where multiple regions are covered by semitransparent MOCs. These MOCs represent various areas of the sky, and the user’s objective is to identify the region corresponding to the LMC. When the cursor hovers over the magenta-colored MOC that overlays the LMC, a pop-up displays a message confirming the discovery, and the user’s score is updated. Auditory feedback is also incorporated to signal correct or incorrect selections, enhancing the interactive aspect of the game, see Figure 3.

This implementation showcases the effective use of Textual MOCs in an interactive learning setting. Users can visually explore the sky, guided by real-time feedback provided by the MOCs. The combination of spatial and textual data not only

promotes an engaging activity but also serves as a practical game-based learning, highlighting the application of astronomical data in an interactive format. The examples can be reproduced using the tutorial associated with the article.

5. MOCs in the Generative AI Stack: From Embeddings to Multimodal RAG and Agents

Here we discuss how these concise textual annotations within MOCs can now be utilized by GenAI models. To enable this, the text must be converted into numerical vectors, as the processing method in GenAI relies on embeddings—vectors (lists) of floating-point numbers that capture semantic information and meaning.

We demonstrate how the Textual MOCs can be utilized within multimodal Generative AI (GenAI) frameworks, enabling simultaneous semantic, spatial and image-based visual recognition tasks. We show how to transform the textual component into semantic embeddings, store them in vector databases, and perform basic operations such as retrieval and classification. The HiPS standard, adopted by the IVOA for storing and visualizing astronomical images, will be used to integrate image URLs into semantic MOCs. This integration will be implemented through the `hips2fits`¹⁵ server, allowing seamless access to corresponding images. With this approach, multimodal models—especially those optimized for visual recognition, image reasoning, caption generation, and question answering on images—can be effectively utilized.

5.1. Semantic MOC

Textual MOCs are converted into semantic embeddings, which are compact numerical representations that encapsulate the meaning of the text within a multidimensional geometric space (Mikolov et al., 2013; Pennington et al., 2014; Devlin et al., 2018). Various approaches exist for generating semantic embeddings, enabling the numerical representation of text, images, audio, or structured data. The choice of method depends on the data type, the model used, and the specific application.

To ensure reproducibility in the generation of Semantic MOCs, we provide a Python function that converts the text field of a MOC into a numerical embedding and stores it within the same data structure. The function employs the `LangChain` framework with the `OllamaEmbeddings` integration¹⁶, allowing researchers to select among different embedding models depending on the accuracy and computational cost requirements. The function `embedding_from_custom_text` reads the `text` key from the MOC JSON object, generates the embedding vector using the specified model (default: `nomic-embed-text`), and stores both the resulting vector and the model name under the dedicated keys (`embedding`, `embedding_model`). This guarantees a straightforward and traceable conversion from

¹⁵<https://alaska.cds.unistra.fr/hips-image-services/hips2fits>

¹⁶https://python.langchain.com/docs/integrations/text_embedding/ollama/

Textual to Semantic MOC, since the applied model is explicitly documented in the assigned keys. Different embedding models can be employed depending on the desired trade-off between accuracy and efficiency (for example, `all-MiniLM-L6-v2`¹⁷, `nomic-embed-text`¹⁸, `mistral-embed` models¹⁹, among others), which can be consulted on the providers' official documentation for further details. This compact representation enables efficient storage and indexing of embeddings in modern vector databases.

Snippet 3: JSON structure of a Semantic MOC with text transformed into semantic embeddings.

```
{
  "order_n": [npix_i | where npix_i uniquely
              defined integers],
  // ... continue for other orders ...
  "text": "Your textual description here",
  "metadata": {
    "key1": "value1",
    "key2": "value2",
    // ... continue metadata entries ...
  },
  "annotated_cells": {
    "order_n": {
      "npix_i": "annotation_i",
      // ... continue for other pixels ...
    }
    // ... continue for other orders ...
  },
  "embedding": [text vectorization],
  "embedding_model": "selected embedding
                    model",
  "image": "https://alaska.cds.unistra.fr
           /.../...",
}
```

To specify semantic MOCs, we introduce three new entries in the serialization of textual MOCs, as illustrated in Snippet 3.

The `embedding` field contains the numerical vector representation of the text, generated through a chosen text vectorization technique, encoding the semantic meaning in a compact form suitable for computational analysis.

The `embedding_model` entry specifies the embedding model used to generate the vector representation. The choice of model determines the characteristics of the embedding and its applicability to different data types.

The `image` key links to an astronomical image. In this scenario, the `hips2fits`²⁰ server is used to retrieve astronomical images via direct URL parsing, thereby ensuring fast and consistent data access. The `hips2fits` service enables the dynamic generation of FITS image cutouts of arbitrary size, resolution, and field of view from a specified HiPS (Hierarchical Progress-

¹⁷<https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

¹⁸<https://huggingface.co/nomic-ai/nomic-embed-text-v1>

¹⁹<https://docs.mistral.ai/capabilities/embeddings/>

²⁰<https://alaska.cds.unistra.fr/hips-image-services/hips2fits>

Textual MOC Application with Aladin Lite

This application demonstrates a simple example of using Textual MOCs in an interactive page powered by Aladin Lite. Hovering over the magenta-colored MOC will open a popup at the bottom displaying the information contained in the textual MOC. A button allows you to enable sounds that will be played when the cursor enters or exits the region defined by the MOC.




Figure 2: An illustrative example demonstrating the integration of the Textual MOC data structure into a web page using Aladin Lite. The figure illustrates the Aladin Lite viewer with the MOC map, an informative popup activated by the cursor, and a menu for managing sound effects. Interactive demo: https://ggreco77.github.io/TextualMOC/AladinGame/aladin_game_d1/

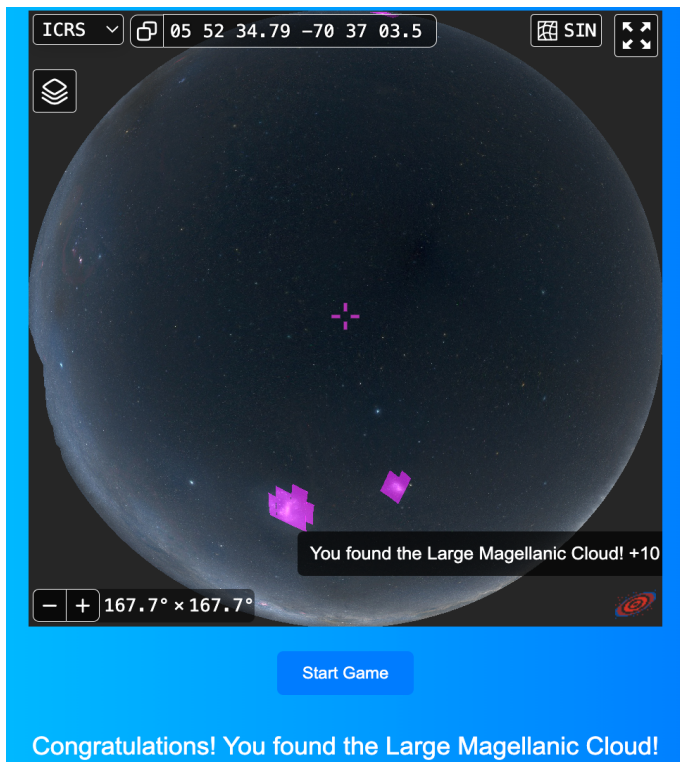


Figure 3: An interactive game using Textual MOCs within Aladin Lite helps users locate the Large Magellanic Cloud (LMC). The magenta-colored MOC marks the position of the LMC, and informational popups provide feedback as the user interacts with the map. Auditory cues enhance the user experience by signaling correct or incorrect selections. Interactive demo: https://ggreco77.github.io/TextualMOC/AladinGame/aladin_game_d2/

sive Survey) dataset. Each image is properly accompanied by a comprehensive, source-specific bibliography.

In the next section, we show how these new fields enable the use of MOCs in generative AI models, not only text-based but also those with visual capabilities.

5.2. Vector database, RAG, and Visual Models

In the context of generative AI and LLMs, an increasingly widespread approach to querying document collections is the Retrieval-Augmented Generation (RAG) paradigm (Lewis et al., 2020; Guu et al., 2020; Karpukhin et al., 2020). RAG combines embedding-based semantic retrieval with automatic text generation via advanced LLMs. In this framework, each document is transformed into a numerical vector—an embedding—that captures its semantic content using pre-trained models such as Sentence-BERT (Reimers and Gurevych, 2019). These vectors are stored in a vector database, which is optimized to quickly return the documents most similar to a query based on semantic similarity (Taipalus, 2024). Upon receiving a query, the system encodes it as an embedding, retrieves the most relevant documents, and supplies them as context to a generative model. This approach enables the generation of responses grounded in concrete sources, thereby reducing so-called “hallucinations”—unfounded assertions produced by the model (Shuster et al., 2021). By employing specialized vector storage solutions, the pipeline remains scalable and can be refreshed in near real-time, supporting low-latency queries over extensive text corpora while ensuring content stays up to date.

In our approach, input documents are represented as Textual MOCs and processed through a conventional Retrieval-

Augmented Generation (RAG) workflow. We enhance this pipeline with a dedicated spatial filter in the vector database enabling hybrid queries that jointly consider semantic relevance and spatial location.

Specifically, from a user query, the system first encodes the Textual MOCs into embedding vectors (Step 1 - see Section 5.1) and indexes them in a vector database optimized for similarity search (Step 2). During the retrieval step (Step 3), a set of similarity scores sim_k and their corresponding record indices idx_k is returned: each sim_k quantifies the relevance of the k -th textual MOC to the query, while idx_k points to its original entry in the database. These indices then allow the system to fetch the associated spatial MOC objects (Step 4) and use them as context to generate grounded responses via a language model (Step 5). Concurrently, the selected MOCs are overlaid on an all-sky image using Aladin Lite (Step 6), providing an immediate visual representation of the queried region. Finally, the FITS image URLs (`image`) are streamed to a vision-enabled LLM (Step 7) for automated analysis of visual features. The complete pipeline—from textual query to sky map and visual analysis—is summarized in Figure 4.

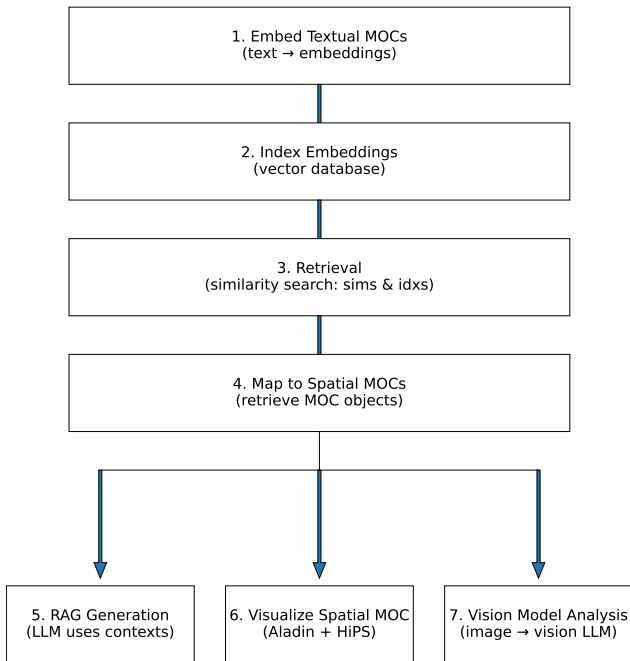


Figure 4: RAG Processing Framework with spatial MOC integration and vision analysis. From a user query, (1) we embed textual MOCs into vector representations and (2) index them in a vector database. In (3) a retrieval step returns a set of similarity scores (sim_k) and their corresponding record indices (idx_k), where each sim_k quantifies the relevance of the k -th MOC to the query and idx_k points to the original MOC entry. These indices allow us to (4) fetch the matching spatial MOC objects, (5) generate context-grounded answers with a language model, (6) overlay the selected MOCs on all-sky imagery via Aladin Lite, and (7) stream the FITS image URLs (`image`) to a vision LLM for automated visual feature analysis.

In the basic tutorial, we present a proof-of-concept system implementing a multimodal Retrieval-Augmented Generation (RAG) pipeline that combines spatial and semantic retrieval with visual enrichment via automated image analysis. We construct ten spatial Multi-Order Coverage (MOC) footprints centered on well-known galaxies; for each object we provide a ~ 150 -word textual summary and a link to an image obtained via `hips2fits`. The pipeline integrates lightweight local models for embeddings, text generation, and vision, supporting end-to-end retrieval and interpretation. The system supports two complementary query modes. (i) Semantic mode accepts natural-language queries (e.g., “Check whether the database contains the host galaxy of the gravitational-wave event GW170817. If yes, provide a concise summary of the corresponding document.”). Object descriptions are indexed in a vector search index built from sentence embeddings; the RAG step selects the most relevant documents, and a local LLM produces the answer. To ensure transparency, the LLM output enforces a trailing `USED:` line that enumerates only the documents actually cited. The MOC footprints of the `USED` documents are overlaid in `ipyaladin`, and their corresponding images are downloaded and supplied to a vision model for morphological classification (spiral/elliptical/irregular) and detection of tidal features. (ii) Positional mode accepts sky coordinates (RA, Dec) and tests containment in available MOCs using the `mocpy` library. Matching footprints are rendered in `ipyaladin`, and the associated images are analyzed by the same vision model. The proof-of-concept is not intended as a comprehensive validation of embedding, vision, or textual models; rather, it demonstrates that lightweight, locally hosted models can be chained to deliver coherent, inspectable results (retrieval \rightarrow spatial overlay \rightarrow visual reasoning). Finally, the system can be designed to scale toward automatic result collection, alert dissemination, and real-time specialist chatbots, using an extension of a well-known and well-documented IVOA standard (MOC)—enriched with textual contents and image references—while preserving full operational compatibility.

6. Conclusions and Future Outlook

In this paper, we explore several potential applications of a widely adopted IVOA standard for encoding complex sky regions—namely, MOC. Our primary goal is to investigate how this data structure can be efficiently integrated into modern generative AI architectures. We introduce an experimental data structure, which we refer to as Textual MOCs when raw text annotations are applied directly, or Semantic MOCs when that text is converted into semantic embeddings via pre-trained models and encapsulated within the same MOC—making it more suitable for storage in vector databases. More specifically, the textual portion, encapsulated through specific keys within the spatial MOC, is then converted into semantic embeddings, which can be employed, for example, in RAG applications for the joint retrieval of MOC information, semantic data, and spatial coordinates. Moreover, Virtual Observatory servers (e.g., `hips2fits`) allow the generation of URLs for images from various astronomical surveys and, within the same workflow, direct these

requests to vision models for further analysis. In this work, we implement JSON serialization of the MOC; future developments may consider employing FITS serialization by embedding the additional metadata within the header or by structuring ASCII tables with MOCs in one column and the corresponding metadata in additional columns.

Textual/Semantic MOCs are currently a practical tool for embedding textual information into spatial data. They have the potential to evolve into a formalized model aligned with the IVOA proposed standards. When integrated into the broader IVOA Resource Metadata framework (VOResource) (Demleitner et al., 2025), Textual and Semantic MOCs can fully comply with FAIR principles (Findable, Accessible, Interoperable, Reusable), allowing tools and services to access and combine spatial and textual information seamlessly and efficiently.

We note that Textual/Semantic MOCs inherently produce structured outputs, making them ideal building blocks for agent-based systems. In this paradigm, each MOC can function as an autonomous software agent—encapsulating its embeddings, textual metadata, and behavior (e.g., relevance scoring, summary generation, visualization)—and responding to incoming queries or environmental triggers with discrete actions.

Looking ahead, we plan to conduct dedicated simulations of these data-driven agents in the context of third-generation interferometers (Einstein Telescope (Punturo et al., 2010), Cosmic Explorer (Reitze et al., 2019)) and space-based missions (e.g., LISA, THESEUS, ATHENA). In such scenarios, multi-messenger inputs—from neutrinos and GRBs to gravitational waves—must be efficiently orchestrated across varied time windows, with rapid, agent-triggered alerts for follow-up observations. By treating each MOC as a self-contained agent, we achieve a naturally distributed, scalable, and robust architecture capable of meeting the demands of next-generation multi-messenger astronomy.

Acknowledgements

The authors acknowledge support from PNRR ETIC (IR0000004), “Einstein Telescope Infrastructure Consortium” — CUP I53C21000420006 (MISSIONE 4, COMPONENTE 2, INVESTIMENTO 3.1); from the Astrophysics Centre for Multi-messenger Studies in Europe (ACME), funded by the European Union’s Horizon Europe Research and Innovation Programme (Grant Agreement No. 10113192); from the ESCAPE project (European Science Cluster of Astronomy & Particle Physics ESFRI Research Infrastructures), funded by the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 824064); and from the Italian National Institute for Nuclear Physics (INFN) and the Italian Space Agency (ASI) under the ASI-INFN agreement No. 2021-43-HH.0.

References

Allen, M.G., Derriere, S., Bonnarel, F., Boch, T., Fernique, P., Dolensky, M., Louys, M., Richards, A.M., 2004. Virtual

observatory standards in action, in: Quinn, P.J., Bridger, A. (Eds.), *Optimizing Scientific Return for Astronomy through Information Technologies*, International Society for Optics and Photonics. SPIE. pp. 153 – 162. doi:10.1117/12.551513.

Ayala Solares, H.A., Coutu, S., Cowen, D., DeLaunay, J.J., Fox, D.B., Keivani, A., Mostafá, M., Murase, K., Oikonomou, F., Seglar-Arroyo, M., Tešić, G., Turley, C.F., 2020. The astrophysical multimessenger observatory network (amon): Performance and science program. *Astroparticle Physics* 114, 68–76. URL: <https://www.sciencedirect.com/science/article/pii/S0927650519301227>, doi:<https://doi.org/10.1016/j.astropartphys.2019.06.007>.

Baumann, M., Boch, T., Pineau, F.X., Fernique, P., Bot, C., Allen, M., 2022. Aladin Lite v3: Behind the Scenes of a Major Overhaul, in: Ruiz, J.E., Pierfederici, F., Teuben, P. (Eds.), *Astronomical Data Analysis Software and Systems XXX*, p. 7.

Berriman, G.B., the IVOA Executive Committee, the IVOA Technical Coordination Group, the IVOA Community, 2021. The International Virtual Observatory Alliance (IVOA) in 2020, in: Ruiz, J.E., Federici, P. (Eds.), *ADASS XXX, ASP, San Francisco*.

Boch, T., 2019. MOCpy, a Python Library to Manipulate Spatial Coverage Maps, in: Molinaro, M., Shortridge, K., Pasian, F. (Eds.), *Astronomical Data Analysis Software and Systems XXVI*, p. 487.

Boch, T., Desroziers, J., 2020. ipyaladin: Enabling Aladin Lite in Jupyter Notebooks, in: Ballester, P., Ibsen, J., Solar, M., Shortridge, K. (Eds.), *Astronomical Data Analysis Software and Systems XXVII*, p. 117.

Boch, T., Fernique, P., 2014. Aladin Lite: Embed your Sky in the Browser, in: Manset, N., Forshay, P. (Eds.), *Astronomical Data Analysis Software and Systems XXIII*, p. 277.

Bonnarel, F., Fernique, P., Bienaymé, O., Egret, D., Genova, F., Louys, M., Ochsenbein, F., Wenger, M., Bartlett, J.G., 2000. The ALADIN interactive sky atlas. A reference tool for identification of astronomical sources. *Astronomy & Astrophysics Supplement Series* 143, 33–40. doi:10.1051/aas:2000331.

Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al., 2020. Language models are few-shot learners, in: *Advances in neural information processing systems*, pp. 1877–1901.

Brozzetti, M.L., Dálya, G., Greco, G., Bawaj, M., Matcovich, T., Branchesi, M., Boch, T., Baumann, M., Cutini, S., De Pietri, R., Khalouei, E., Fernique, P., Punturo, M., Vocca, H., 2024. GLADENet: A progressive web app for multi-messenger cosmology and electromagnetic follow-ups of

- gravitational-wave sources. *A&A* 684, A44. doi:[10.1051/0004-6361/202348073](https://doi.org/10.1051/0004-6361/202348073).
- Carvajal, R., Matute, I., Afonso, J., Amarantidis, S., Barbosa, D., Cunha, P., Humphrey, A., 2021. Exploring New Redshift Indicators for Radio-Powerful AGN. *Galaxies* 9, 86. doi:[10.3390/galaxies9040086](https://doi.org/10.3390/galaxies9040086), [arXiv:2111.00778](https://arxiv.org/abs/2111.00778).
- Coughlin, M.W., Bloom, J.S., Nir, G., Antier, S., du Laz, T.J., van der Walt, S., Crellin-Quick, A., Culino, T., Duev, D.A., Goldstein, D.A., Healy, B.F., Karambelkar, V., Lilleboe, J., Shin, K.M., Singer, L.P., Ahumada, T., Anand, S., Bellm, E.C., Dekany, R., Graham, M.J., Kasliwal, M.M., Kostadinova, I., Kiendrebeogo, R.W., Kulkarni, S.R., Jenkins, S., LeBaron, N., Mahabal, A.A., Neill, J.D., Parazin, B., Peloton, J., Perley, D.A., Riddle, R., Rusholme, B., van Santen, J., Sollerman, J., Stein, R., Turpin, D., Wold, A., Amat, C., Bonnefon, A., Bonnefoy, A., Flament, M., Kerkow, F., Kishore, S., Jani, S., Mahanty, S.K., Liu, C., Llinares, L., Makarison, J., Ollieric, A., Perez, I., Pont, L., Sharma, V., 2023. A Data Science Platform to Enable Time-domain Astronomy. *ApJS* 267, 31. doi:[10.3847/1538-4365/acdee1](https://doi.org/10.3847/1538-4365/acdee1), [arXiv:2305.00108](https://arxiv.org/abs/2305.00108).
- Demleitner, M., Plante, R., Benson, K., Graham, M., Greene, G., Harrison, P., Lemson, G., Linde, T., Rixon, G., 2025. VOResource: an XML Encoding Schema for Resource Metadata Version 1.2. IVOA Recommendation 16 April 2025.
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K., 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv e-prints*, [arXiv:1810.04805](https://arxiv.org/abs/1810.04805)doi:[10.48550/arXiv.1810.04805](https://doi.org/10.48550/arXiv.1810.04805), [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
- Fernique, P., Allen, M.G., Boch, T., Oberto, A., Pineau, F.X., Durand, D., Bot, C., Cambrésy, L., Derriere, S., Genova, F., Bonnarel, F., 2015. Hierarchical progressive surveys. Multi-resolution HEALPix data structures for astronomical images, catalogues, and 3-dimensional data cubes. *A&A* 578, A114. doi:[10.1051/0004-6361/201526075](https://doi.org/10.1051/0004-6361/201526075), [arXiv:1505.02291](https://arxiv.org/abs/1505.02291).
- Fernique, P., Boch, T., Donaldson, T., Durand, D., O’Mullane, W., Reinecke, M., Taylor, M., 2014. MOC - HEALPix Multi-Order Coverage map Version 1.0. IVOA Recommendation 02 June 2014. doi:[10.5479/ADS/bib/2014ivoa.spec.0602F](https://doi.org/10.5479/ADS/bib/2014ivoa.spec.0602F), [arXiv:1505.02937](https://arxiv.org/abs/1505.02937).
- Fernique, P., Boch, T., Donaldson, T., Durand, D., O’Mullane, W., Reinecke, M., Taylor, M., 2019. MOC - HEALPix Multi-Order Coverage map Version 1.1. IVOA Recommendation 07 October 2019. doi:[10.5479/ADS/bib/2019ivoa.spec.1007F](https://doi.org/10.5479/ADS/bib/2019ivoa.spec.1007F).
- Fernique, P., Nebot, A., Durand, D., Baumann, M., Boch, T., Greco, G., Donaldson, T., Pineau, F.X., Taylor, M., O’Mullane, W., Reinecke, M., Derrière, S., 2022. MOC: Multi-Order Coverage map Version 2.0. IVOA Recommendation 27 July 2022. doi:[10.5479/ADS/bib/2022ivoa.spec.0727F](https://doi.org/10.5479/ADS/bib/2022ivoa.spec.0727F).
- Fernique, P., Nebot, A., Durand, D., Baumann, M., Boch, T., Greco, G., Donaldson, T., Pineau, F.X., Taylor, M., O’Mullane, W., Reinecke, M., Derrière, S., 2021. MOC: Multi-Order Coverage map Version 2.0. IVOA IVOA Proposed Recommendation 01 November 2021. <https://ivoa.net/documents/MOC/20211101/index.html>.
- Gal-Yam, A., 2021. The TNS alert system, in: American Astronomical Society Meeting Abstracts, p. 423.05.
- Ginsburg, A., Sipőcz, B.M., Brasseur, C.E., Cowperthwaite, P.S., Craig, M.W., Deil, C., Guillochon, J., Guzman, G., Liedtke, S., Lian Lim, P., Lockhart, K.E., Mommert, M., Morris, B.M., Norman, H., Parikh, M., Persson, M.V., Robitaille, T.P., Segovia, J.C., Singer, L.P., Tollerud, E.J., de Val-Borro, M., Valtchanov, I., Woillez, J., The Astroquery collaboration, a subset of the astropy collaboration, 2019. astroquery: An Astronomical Web-querying Package in Python. *AJ* 157, 98. doi:[10.3847/1538-3881/aafc33](https://doi.org/10.3847/1538-3881/aafc33), [arXiv:1901.04520](https://arxiv.org/abs/1901.04520).
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets, in: *Advances in neural information processing systems*, pp. 2672–2680.
- Górski, K.M., Hivon, E., Banday, A.J., Wandelt, B.D., Hansen, F.K., Reinecke, M., Bartelmann, M., 2005. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal* 622, 759–771. doi:[10.1086/427976](https://doi.org/10.1086/427976), [arXiv:astro-ph/0409513](https://arxiv.org/abs/astro-ph/0409513).
- Greco, G., Branchesi, M., Chassande-Mottin, E., Coughlin, M.W., Stratta, G., D’Álya, G., Hemming, G., Rei, L., Brocato, E., Fernique, P., Boch, T., Derriere, S., Baumann, M., Genova, F., Allen, M., 2019. Working with Gravitational-Wave sky localizations: new methods and implementations, in: *New Era of Multi-Messenger Astrophysics — PoS (Astetrics2019)*, p. 031. doi:[10.22323/1.357.0031](https://doi.org/10.22323/1.357.0031).
- Greco, G., Branchesi, M., Punturo, M., Vocca, H., Travasso, F., De Pietri, R., Fernique, P., Boch, T., Baumann, M., Pineau, F.X., Derriere, S., Nebot, A., Allen, M., 2022a. Capability for Encoding Gravitational-wave Sky Localizations with the Multi Order Coverage Data Structure: Present and Future Developments, in: Ruiz, J.E., Pierfederici, F., Teuben, P. (Eds.), *Astronomical Data Analysis Software and Systems XXX*, p. 131.
- Greco, G., Punturo, M., Allen, M., Nebot, A., Fernique, P., Baumann, M., Pineau, F.X., Boch, T., Derriere, S., Branchesi, M., Bawaj, M., Vocca, H., 2022b. Multi order coverage data structure to plan multi-messenger observations. *Astronomy and Computing* 39, 100547. URL: <https://www.sciencedirect.com/science/>

[article/pii/S2213133722000026](https://doi.org/10.1016/j.ascom.2022.100547), doi:<https://doi.org/10.1016/j.ascom.2022.100547>.

- Guu, K., Lee, K., Tung, Z., Pasupat, P., Chang, M.W., 2020. REALM: Retrieval-Augmented Language Model Pre-Training. arXiv e-prints , arXiv:2002.08909doi:[10.48550/arXiv.2002.08909](https://doi.org/10.48550/arXiv.2002.08909), arXiv:2002.08909.
- Hosseinzadeh, G., Paterson, K., Rastinejad, J.C., Shrestha, M., Daly, P.N., Lundquist, M.J., Sand, D.J., Fong, W.f., Bostroem, K.A., Hall, S., Wyatt, S.D., Gibbs, A.R., Christensen, E., Lindstrom, W., Nation, J., Chatelain, J., McCully, C., 2024. SAGUARO: Time-domain Infrastructure for the Fourth Gravitational-wave Observing Run and Beyond. *ApJ* 964, 35. doi:[10.3847/1538-4357/ad2170](https://doi.org/10.3847/1538-4357/ad2170), arXiv:2310.08624.
- Hunter, J.D., 2007. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering* 9, 90–95. doi:[10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55).
- Karpukhin, V., Oğuz, B., Min, S., Lewis, P., Wu, L., Edunov, S., Chen, D., Yih, W.t., 2020. Dense Passage Retrieval for Open-Domain Question Answering. arXiv e-prints , arXiv:2004.04906doi:[10.48550/arXiv.2004.04906](https://doi.org/10.48550/arXiv.2004.04906), arXiv:2004.04906.
- Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., Willing, C., 2016. Jupyter notebooks – a publishing format for reproducible computational workflows, in: Loizides, F., Schmidt, B. (Eds.), *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, IOS Press. pp. 87 – 90.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Kuttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., Riedel, S., Kiela, D., 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks, in: *Advances in neural information processing systems*, pp. 9459–9474.
- Merín, B., Salgado, J., Giordano, F., Baines, D., Sarmiento, M.H., Martí, B.L., Racero, E., Gutiérrez, R., Pollock, A., Rosa, M., Castellanos, J., González, J., León, I., Ortiz de Landaluce, I., de Teodoro, P., Nieto, S., Lennon, D.J., Arviset, C., de Marchi, G., O’Mullane, W., 2017. ESA Sky: A New Astronomy Multi-Mission Interface, in: Lorente, N.P.F., Shortridge, K., Wayth, R. (Eds.), *Astronomical Data Analysis Software and Systems XXV*, p. 495.
- Mikolov, T., Chen, K., Corrado, G., Dean, J., 2013. Efficient Estimation of Word Representations in Vector Space. arXiv e-prints , arXiv:1301.3781doi:[10.48550/arXiv.1301.3781](https://doi.org/10.48550/arXiv.1301.3781), arXiv:1301.3781.
- Möller, A., Peloton, J., Ishida, E.E.O., Arnault, C., Bachelet, E., Blaineau, T., Boutigny, D., Chauhan, A., Gangler, E., Hernandez, F., Hrivnac, J., Leoni, M., Leroy, N., Moniez, M., Pateyron, S., Ramparison, A., Turpin, D., Ansari, R., Allam, Tarek, J., Bajat, A., Biswas, B., Boucaud, A., Bregeon, J., Campagne, J.E., Cohen-Tanugi, J., Coleiro, A., Dornic, D., Fouchez, D., Godet, O., Gris, P., Karpov, S., Nebot Gomez-Moran, A., Neveu, J., Plaszczyński, S., Savchenko, V., Webb, N., 2021. FINK, a new generation of broker for the LSST community. *MNRAS* 501, 3272–3288. doi:[10.1093/mnras/staa3602](https://doi.org/10.1093/mnras/staa3602), arXiv:2009.10185.
- Munikoti, S., Stewart, I., Horawalavithana, S., Kvinge, H., Emerson, T., E Thompson, S., Pazdernik, K., 2024. Generalist Multimodal AI: A Review of Architectures, Challenges and Opportunities. arXiv e-prints , arXiv:2406.05496doi:[10.48550/arXiv.2406.05496](https://doi.org/10.48550/arXiv.2406.05496), arXiv:2406.05496.
- O’Toole, S., Tocknell, J., 2024. FAIR standards for astronomical data, in: Hugo, B.V., Van Rooyen, R., Smirnov, O.M. (Eds.), *Astronomical Data Analysis Software and Systems XXXI*, p. 265.
- Pennington, J., Socher, R., Manning, C., 2014. GloVe: Global vectors for word representation, in: Moschitti, A., Pang, B., Daelemans, W. (Eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Doha, Qatar. pp. 1532–1543. URL: <https://aclanthology.org/D14-1162>, doi:[10.3115/v1/D14-1162](https://doi.org/10.3115/v1/D14-1162).
- Petter, G.C., Hickox, R.C., Alexander, D.M., Myers, A.D., Geach, J.E., Whalen, K.E., Andonie, C.P., 2023. Host Dark Matter Halos of Wide-field Infrared Survey Explorer-selected Obscured and Unobscured Quasars: Evidence for Evolution. *ApJ* 946, 27. doi:[10.3847/1538-4357/acb7ef](https://doi.org/10.3847/1538-4357/acb7ef), arXiv:2302.00690.
- Punturo, M., Abernathy, M., Acernese, F., et al., 2010. The Einstein Telescope: a third-generation gravitational wave observatory. *Classical and Quantum Gravity* 27, 194002. doi:[10.1088/0264-9381/27/19/194002](https://doi.org/10.1088/0264-9381/27/19/194002).
- Reimers, N., Gurevych, I., 2019. Sentence-bert: Sentence embeddings using siamese bert-networks, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics. pp. 3982–3992.
- Reitze, D., Adhikari, R.X., Ballmer, S., Barish, B., Barsotti, L., Billingsley, G., Brown, D.A., Chen, Y., Coyne, D., Eisenstein, R., Evans, M., Fritschel, P., Hall, E.D., Lazarini, A., Lovelace, G., Read, J., Sathyaprakash, B.S., Shoemaker, D., Smith, J., Torrie, C., Vitale, S., Weiss, R., Wipf, C., Zucker, M., 2019. Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO, in: *Bulletin of the American Astronomical Society*, p. 35. arXiv:1907.04833.
- Schaaff, A., 2007. Data in astronomy: From the pipeline to the virtual observatory, in: Weske, M., Hacid, M.S., Godart, C. (Eds.), *Web Information Systems Engineering – WISE 2007*

- Workshops, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 52–62.
- Shuster, K., Poff, S., Chen, M., Kiela, D., Weston, J., 2021. Retrieval augmentation reduces hallucination in conversation. CoRR abs/2104.07567. URL: <https://arxiv.org/abs/2104.07567>, doi:10.48550/arXiv.2104.07567, arXiv:2104.07567.
- Singer, L.P., Parazin, B., Coughlin, M.W., Bloom, J.S., Crellin-Quick, A., Goldstein, D.A., van der Walt, S., 2022. HEALPix Alchemy: Fast All-Sky Geometry and Image Arithmetic in a Relational Database for Multimessenger Astronomy Brokers. AJ 163, 209. doi:10.3847/1538-3881/ac5ab8, arXiv:2112.06947.
- Smareglia, R., Molinaro, M., Butora, R., Schisano, E., Carey, D., Elia, D., Molinari, S., 2019. Archive, Discover and Match Compact and Diffuse Objects on the Galactic Plane in the VIALACTEA Project, in: Molinaro, M., Shortridge, K., Pasian, F. (Eds.), *Astronomical Data Analysis Software and Systems XXVI*, p. 116.
- Smith, M.W.L., Ibar, E., Maddox, S.J., Valiante, E., Dunne, L., Eales, S., Dye, S., Furlanetto, C., Bourne, N., Cigan, P., Ivison, R.J., Gomez, H., Smith, D.J.B., Viaene, S., 2017. The Herschel-ATLAS Data Release 2, Paper I. Submillimeter and Far-infrared Images of the South and North Galactic Poles: The Largest Herschel Survey of the Extragalactic Sky. ApJS 233, 26. doi:10.3847/1538-4365/aa9b35, arXiv:1712.02361.
- Taipalus, T., 2024. Vector database management systems: Fundamental concepts, use-cases, and current challenges. Cognitive Systems Research 85, 101216. URL: <https://www.sciencedirect.com/science/article/pii/S1389041724000093>, doi:https://doi.org/10.1016/j.cogsys.2024.101216.
- Thoppilan, R., De Freitas, D., Hall, J., Shazeer, N., Kulshreshtha, A., Cheng, H.T., Jin, A., Bos, T., Baker, L., Du, Y., Li, Y., Lee, H., Zheng, H.S., Ghafouri, A., Menegali, M., Huang, Y., Krikun, M., Lepikhin, D., Qin, J., Chen, D., Xu, Y., Chen, Z., Roberts, A., Bosma, M., Zhao, V., Zhou, Y., Chang, C.C., Krivokon, I., Rusch, W., Pickett, M., Srinivasan, P., Man, L., Meier-Hellstern, K., Ringel Morris, M., Doshi, T., Delos Santos, R., Duke, T., Soraker, J., Zevenbergen, B., Prabhakaran, V., Diaz, M., Hutchinson, B., Olson, K., Molina, A., Hoffman-John, E., Lee, J., Aroyo, L., Rajakumar, R., Butryna, A., Lamm, M., Kuzmina, V., Fenton, J., Cohen, A., Bernstein, R., Kurzweil, R., Aguera-Arcas, B., Cui, C., Croak, M., Chi, E., Le, Q., 2022. LaMDA: Language Models for Dialog Applications. arXiv e-prints , arXiv:2201.08239doi:10.48550/arXiv.2201.08239, arXiv:2201.08239.
- Tran, K.T., Dao, D., Nguyen, M.D., Pham, Q.V., O’Sullivan, B., Nguyen, H.D., 2025. Multi-Agent Collaboration Mechanisms: A Survey of LLMs. arXiv e-prints , arXiv:2501.06322doi:10.48550/arXiv.2501.06322, arXiv:2501.06322.
- Wyatt, S.D., Tohuvavohu, A., Arcavi, I., Lundquist, M.J., Howell, D.A., Sand, D.J., 2020. The Gravitational Wave Treasure Map: A Tool to Coordinate, Visualize, and Assess the Electromagnetic Follow-up of Gravitational-wave Events. The Astrophysical Journal 894, 127. doi:10.3847/1538-4357/ab855e, arXiv:2001.00588.
- Xu, Y., Xu, D., Cui, C., Fan, D., Zhu, Z., Yu, B., Li, C., Han, J., Mi, L., Li, S., He, B., Tao, Y., Yang, H., Yang, S., 2020. GWOPS: A VO-technology Driven Tool to Search for the Electromagnetic Counterpart of Gravitational Wave Event. PASP 132, 104501. doi:10.1088/1538-3873/aba69f, arXiv:2009.03497.
- Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., Dewan, C., Diab, M., Li, X., Lin, X.V., Mihaylov, T., Ott, M., Shleifer, S., Shuster, K., Simig, D., Singh Koura, P., Sridhar, A., Wang, T., Zettlemoyer, L., 2022. OPT: Open Pre-trained Transformer Language Models. arXiv e-prints , arXiv:2205.01068doi:10.48550/arXiv.2205.01068, arXiv:2205.01068.