

Chasing Cosmic Rarities: Intelligent Agents and Physical Modeling in the Time-Domain Frontier

Clecio R. Bom, Bernardo Fraga, Phelipe Darc, Gabriel Teixeira
Andre Santos, Gustavo Schwarz, Luidhy Santana-Silva
debom@cbpf.br

RAPID Response:
Hot-wiring the Next Generation of Time-Domain Science



Meet Our Team




Clécio R. Bom

Observational Cosmology
LLM
Time domain astronomy
Deep Learning



Bernardo Fraga

Classification of transient
sources
LLM
Time domain astronomy
Large sky surveys



Mariana Bittencourt

ExoplanetsData
infrastructure
Data processing



Phelipe D'arc

Sim. Based Inference.
Time Domain Astronomy.
Symbolic Regression.
Binary Black Holes



Viviane Alfradique

Cosmology
Gravitational Waves
Standard Sirens



Gabriel Teixeira

Photo-z
Type II SN
Deep Learning
Time Domain
Astronomy



Gustavo Schwarz

Programmer - DevOps
LLMs
Data infrastructure
Data processing



Luidhy Santana

Galaxy Formation
Large Sky Surveys
LLM & Deep Learning
Telescope Operations
HPC




André Santos

Observer
Data Reduction
Time domain astronomy

The General Time-Domain paradigm

(a) Current Modular Pipeline

Survey Telescope
LSST



10M Alerts/Night



Alert Brokers



Selection &
Classification



Community/Decision
-making



Follow-up prioritization/
Resource Optimization



Discover of
novelties



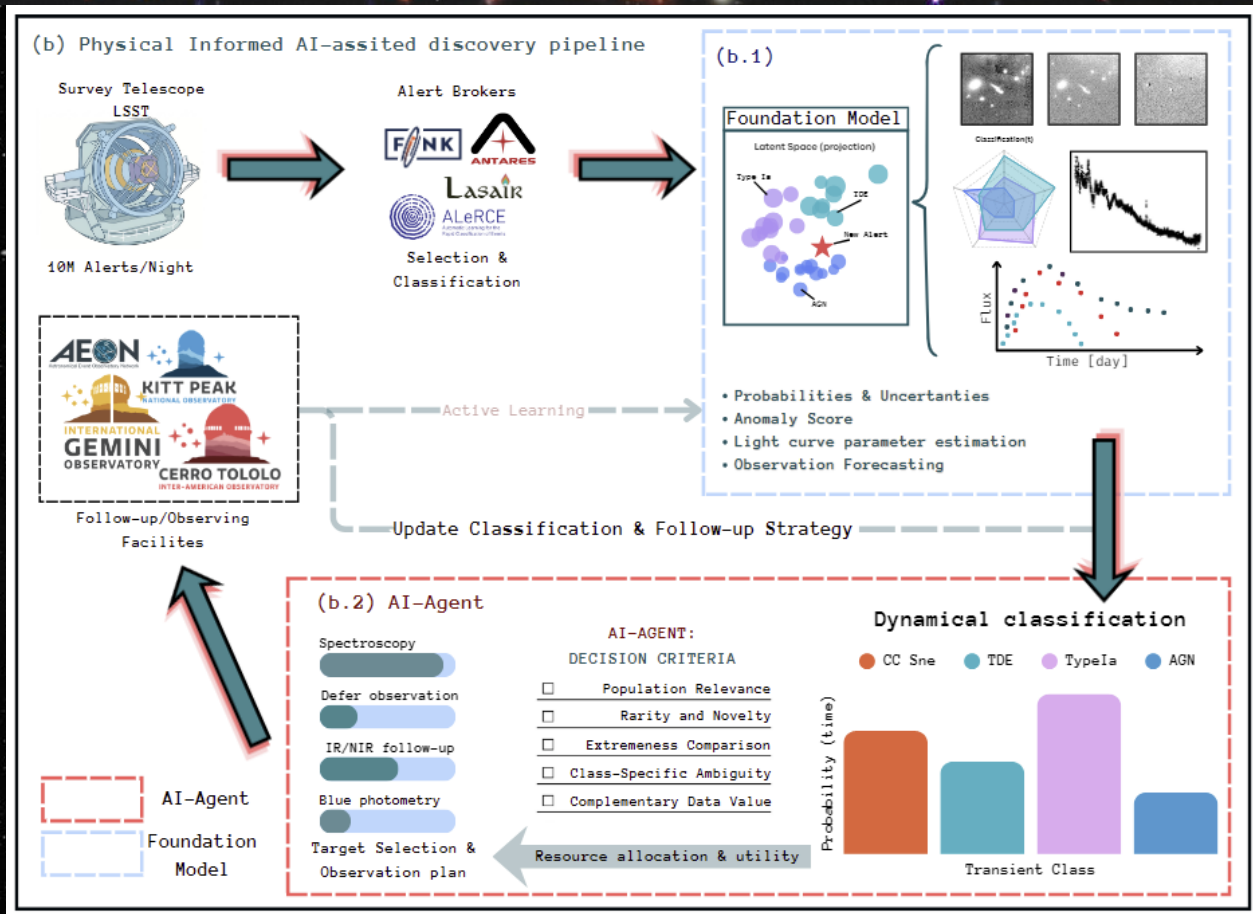
Physical Insights
& Discovery

- “Classification Centric”
- Great if you can look and make a complete assessment of the full stream

The AI-Assisted Time-Domain

Science Question driven

Reasoning and personalized recommendation, plans of effective follow-up and Science relevance at scale.



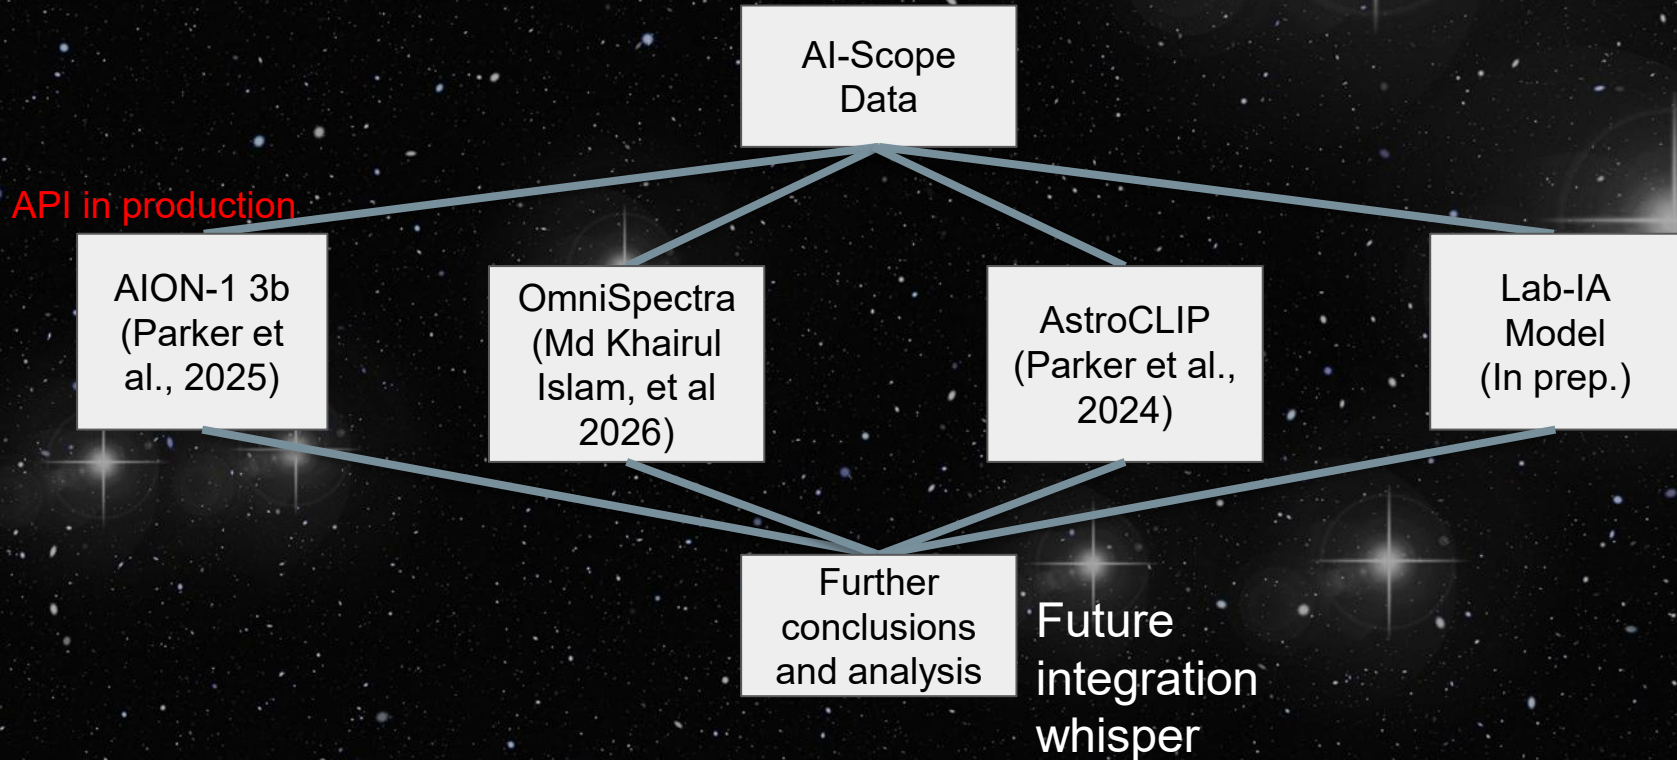
Foundation Models as Infrastructure AI

A single **large model** pretrained on **massive, diverse data**
Learns general representations
Reused for many downstream tasks with little or no retraining

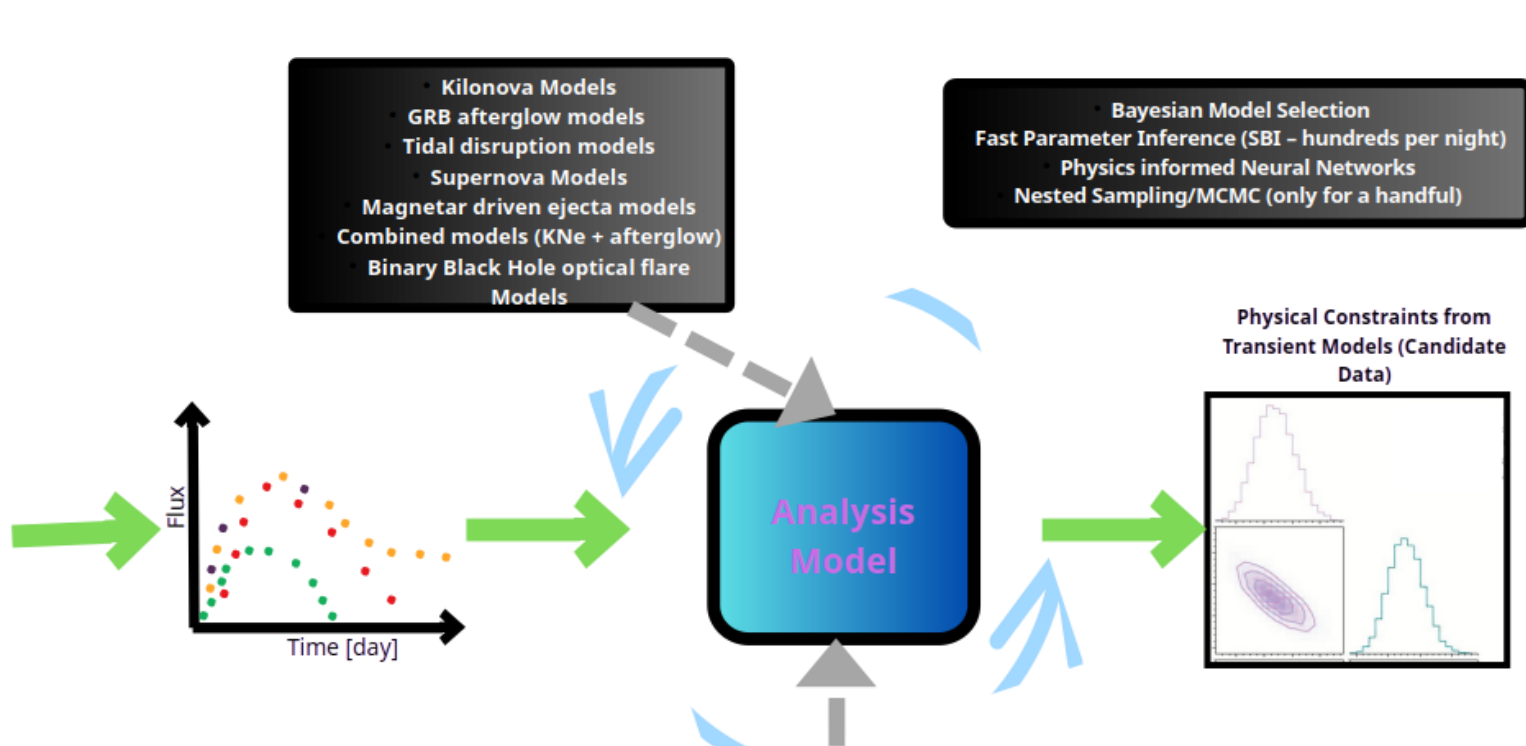
Why?

Replace many task-specific models
Exploit **unlabeled data** (self-supervised learning)
Enable **transfer learning** across domains
Better **scaling behavior** with more data & parameters

Foundational Models as API service with AI-Scope



WHISPER



If I observe one more epoch with this Telescope would I be able to resolve the correct model? On which band, on what Limiting Magnitude? How many follow-up do I need? The system can recomend spectroscopy.

Light Curves

Multi-band
photometry

Host Images

Galaxy morphology
& environment

Metadata

Coordinates,
redshift, alerts

WHISPER


Posterior distribution
+ Metrics

Trained on
Personalized
Dataset



Graph
knowledge
representation
over the whole
history of arxiv

AI-AGENT



AI-ASSITED DISCOVERY
PIPELINE

LLM Integration Layer

Natural Language Interface: Scientists query the pipeline in plain English: "Show me all TDE candidates above $z=0.3$ observed in the last 24h." LLM translates to structured API calls.

Reasoning & Justification: LLM generates human-readable explanations for every Agent decision: why spectroscopy was chosen over deferral, which decision criteria dominated.

Adaptive Science Goals: Community encodes observing objectives as natural-language GOALS (e.g., "prioritize rare SNe-I over AGN this season").

Anomaly Narrative: When the LLM Model flags a "rare" event, the LLM synthesizes context from literature and internal catalogs to generate a discovery Report.

Agentic Tool Use: LLM directly calls observatory APIs, broker endpoints, and scheduling systems — acting as the orchestration layer that translates model outputs into telescope commands. **(Under development and testing in Small Telescope)**




Redback (nikhil-sarin)

- Kilonova Models
- GRB afterglow models
- Tidal disruption models
- Supernova Models
- Magnetar driven ejecta models
- Combined models (KNe + afterglow)
- Binary Black Hole optical flare Models



WHISPER



AUTOMATED ANALYSIS FOR TIME-DOMAIN ASTRONOMY

Approximate Bayesian Computation (ABC)

Speed:
Constraining
Power:



Simulation Based Inference (SBI - NPE/NLE)

Speed:
Constraining
Power:



MCMC/Nested Sampling (Dynesty)

Speed:
Constraining
Power:



WHISPER

How WHISPER will modelling these peculiar sources?

Kilonova Spectral Inverse Modelling with Simulation-based Inference: An Amortized Neural Posterior Estimation Analysis

CrossMark

P. Darc¹, C. R. Bom¹, B. Fraga¹, and C. D. Kilpatrick²

¹Centro Brasileiro de Pesquisas Físicas, Rua Dr. Xavier Sigaud 150, CEP 22290-180, Rio de Janeiro, RJ, Brazil; phelipedarc@cbpf.br

²Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, Evanston, IL 60201, USA

Received 2024 March 13; revised 2024 May 18; accepted 2024 May 29; published 2024 August 7

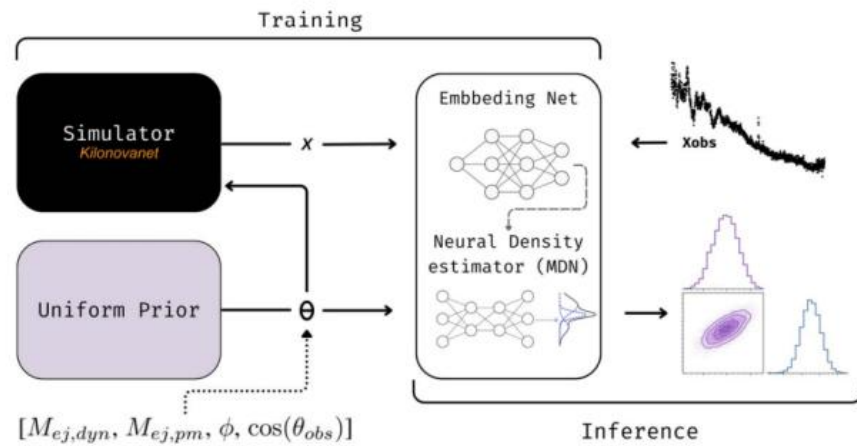


Figure 2. Inference pipeline using ANPE—amortized inference enables rapid inference for every new spectrum observed. In the training phase, a neural-network-based density estimator (MDN) learns the probabilistic relationship between the model’s parameters and simulated spectra, at high computational cost. Then, in the inference phase, the trained density estimator takes the observed real data as input and infers the posterior distribution of the parameters, at low computational cost: less than a second.

WHISPER

How WHISPER will modelling these peculiar sources?

Multi-Wavelength Analysis of Kilonova Associated with GRB 230307A: Accelerated Parameter Estimation and Model Selection Through Likelihood-Free Inference

Phelipe Darc*

Centro Brasileiro de Pesquisas Físicas (CBPF)
Rua Xavier Sigaud, 150, Urca, Rio de Janeiro, Brazil.

Clecio R. Bom

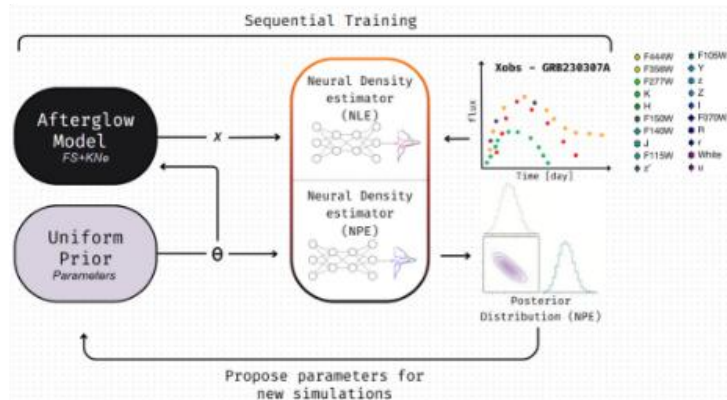


Centro Brasileiro de Pesquisas Físicas (CBPF)
Rua Xavier Sigaud, 150, Urca, Rio de Janeiro, Brazil.
Centro Federal de Educação Tecnológica Celso Suckow da Fonseca (CEFET-RJ),
Av. Maracanã, 229, Maracanã, Rio de Janeiro, Brazil.

Gabriel S. M. Teixeira

Centro Brasileiro de Pesquisas Físicas (CBPF)
Rua Xavier Sigaud, 150, Urca, Rio de Janeiro, Brazil.

Charles Kilpatrick

Center for Interdisciplinary Exploration and Research in Astro-physics (CIERA)
and Department of Physics and Astronomy
North-western University, Evanston, IL 60208, USA.



WHISPER

How WHISPER will modelling these peculiar sources?

The progenitor of the GRB 230307A: a decision through a Bayesian analysis

Viviane Alfradique, Rodrigo da Mata, Clécio R. Bom 

*Centro Brasileiro de Pesquisas Físicas,
Rua Dr. Xavier Sigaud 150,
22290-180 Rio de Janeiro, RJ, Brazil*

TABLE I. Summary statistical results.

	$\ln(\mathcal{Z}^1/\mathcal{Z}^{\text{ref}})$	Maximum $\ln(\mathcal{L})$	BIC	Dimensions	KN model	Reference
BNS system (2-component kilonova)	-23.07	-108.12	382.86	17		[29]
BNS system (Merger-Nova)	-98.45	-182.50	525.12	15		[30]
BNS system (General Merger-Nova)	-98.05	-181.06	515.82	13		[31]
NS+WD system (magnetar spin-down)	-	-87.19	323.97	14		[16]
SBH+WD system (TDE - 4/3 polytropes stars)	-8398.99	-8458.49	17126.20	15		[32]
NS+BH system (2-component kilonova - Ejecta rel)	-3849.79	-3932.98	8027.80	15		[29]



Our results show a statistical preference for a BNS or NS–WD progenitor producing a kilonova powered by a magnetar and ${}^{56}\text{Ni}$ decay, characterized by a ${}^{56}\text{Ni}$ mass of $\sim 4 \times 10^{-4} M_{\odot}$ and an ejecta mass of $0.06 M_{\odot}$. Furthermore, under the assumption of a BNS origin within this model, we infer binary component masses of $m_1 = 1.81_{-0.61}^{+0.46} M_{\odot}$ and $m_2 = 1.61_{-0.41}^{+0.65} M_{\odot}$,

How AI-Scope can accelerate Transient Research?

- The LLMs report evaluates population level features and call attentions to interesting and outliers

Probing the origin of the kilonova candidate GRB 230307A: analysis of host galaxy and offset

CLECIO R. BOM,¹ DAVI C. RODRIGUES,^{2,1} ARIANNA CORTESI,^{3,4} AMANDA E. ARAUJO-CARVALHO,⁴ DANIEL RUSCHEL-DUTRA,⁵
GIULIANO IORIO,⁶ LUIDHY SANTANA-SILVA,¹ CHARLES D. KILPATRICK,⁷ FABRICIO FERRARI,⁸ LUIS LOMELÍ-NÚÑEZ,⁴
THOMAS HARVEY,⁹ DUNCAN AUSTIN,⁹ CHRISTOPHER J. CONSELICE,⁹ NATHAN ADAMS,⁹ AND ROBERTO CID FERNANDES⁵

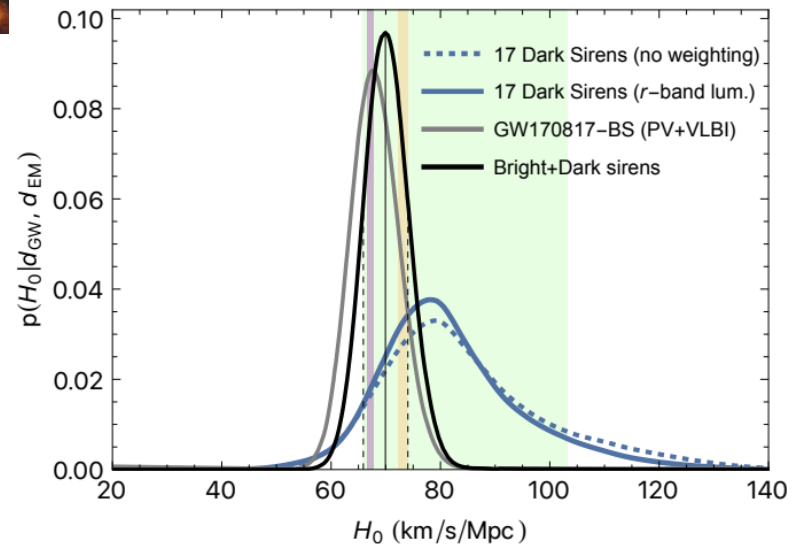


The time to arrive at the merger position departing from the disk compared to the total kick of binary system, given a galaxy mass model and BNS population simulation overlap only 0.1% within the 2σ region of the galaxy mass model. Accepted to **ApJ Letters 2026**.

How AI-Scope can accelerate Research?

Improved constraint on the Hubble constant from dark sirens with LIGO/Virgo/KAGRA O4a

Viviane Alfradique¹,* Clécio R. Bom², Gabriel Teixeira³, and André Santos⁴
 Centro Brasileiro de Pesquisas Físicas,
 Rua Dr. Xavier Sigaud 150,
 22290-180 Rio de Janeiro, RJ, Brazil



Event(s)	Method(s)	H_0 (km/s/Mpc)	σ_{H_0} (km/s/Mpc)
O1-O3 – 47 dark sirens ¹	Catalog (K-band luminosity-weighted) + BNS/NSBH/BBH population (PLP model with fixed parameters)	67^{+13}_{-12}	12.5 (18%)
O1-O4a – 142 dark sirens ²	Catalog (K-band luminosity-weighted) + BNS/NSBH/BBH population (FullPop-4.0 model)	$81.6^{+21.5}_{-15.9}$	18.7 (23%)
O1-O4a – 142 dark sirens ³	Catalog (K-band luminosity-weighted) + BNS/NSBH/BBH population (BPL+3peak model)	$82.5^{+16.8}_{-14.3}$	15.6 (19%)
O1-O4a – 142 dark sirens ²	BNS/NSBH/BBH population (FullPop-4.0 model)	$76.4^{+23.0}_{-18.1}$	20.6 (27%)
O1-O4a – 137 dark sirens ⁴	BBH population (semiparametric B-spline model)	$57.8^{+21.9}_{-20.6}$	21.3 (37%)
O1-O4a - 15 dark sirens ⁵	Catalog (no luminosity weighting)	$70.4^{+13.6}_{-11.7}$	12.6 (18%)
GW170817 ⁶	Bright (v_p corrected)	$68.8^{+17.3}_{-7.6}$	12.5(18%)
GW170817 ⁷	Bright (v_p +VLBI)	$68.3^{+4.6}_{-4.5}$	4.6(7%)
O1-O4a – 17 dark sirens ⁸	Catalog (no luminosity weighting; GW likelihood Gaussian approximation, eq. 5)	$72.6^{+11.6}_{-12.6}$	11.4 (16%)
O1-O4a – 17 dark sirens ⁸	Catalog (r -band luminosity-weighted; GW likelihood Gaussian approximation, eq. 5)	$72.8^{+12.6}_{-14.3}$	13.5 (18%)
O1-O4a – 17 dark sirens ⁸	Catalog (no luminosity weighting; full GW likelihood, eq. 6)	$78.8^{+14.6}_{-12.2}$	13.4 (17%)
O1-O4a – 17 dark sirens ⁸	Catalog (r -band luminosity-weighted; full GW likelihood, eq. 6)	$78.2^{+12.0}_{-11.0}$	11.5 (15%)
O1-O4a – 17 dark + 1 bright siren ⁸	Bright (host) + Catalog (r -band luminosity-weighted; full GW likelihood, eq. 6)	$75.2^{+8.0}_{-7.9}$	8.0(11%)
O1-O4a – 17 dark + 1 bright siren (EM) ⁸	Bright (v_p +VLBI) + Catalog (r -band luminosity-weighted; full GW likelihood, eq. 6)	$69.9^{+4.1}_{-4.0}$	4.1 (6%)

AI-SCOPE

- Alert Streams, GCN notices & circulars, brokers (FINK) powered by LLMs.
- Probability maps (LIGO Skymaps) or transient positions + uncertainties (GRB, FXT).
- Catalogues (DES, DELVE, DESI, SPLUS, open LSST, GAIA, ZTF).
- Processing and analysis of images and spectra

A restrict version is open online. If you are interested to unleash the full (fun) functionalities let me know



<https://ai-scope.cbpf.br/>

A screenshot of the AI-SCOPE Portal website. The page has a dark purple and blue background with a grid pattern and a glowing orange arc at the bottom. The top navigation bar includes the CBPF logo (Centro Brasileiro de Pesquisas Físicas), links for FAQ, Documentation, and Tools, and a Sign in button. The main heading is "AI-SCOPE Portal" followed by "Smart Cosmic Exploration with AI". Below this is a paragraph describing the portal's purpose: "The Smart Cosmic Exploration with AI (AI-SCOPE) is designed to maximize the efficiency of astronomical data distribution by integrating state-of-the-art technologies and artificial intelligence." At the bottom of the main content area are two buttons: "Start Querying" and "Explore Database".

CBPF
Centro Brasileiro de Pesquisas Físicas

FAQ Documentation Tools

Sign in

AI-SCOPE Portal

Smart Cosmic Exploration with AI

The Smart Cosmic Exploration with AI (AI-SCOPE) is designed to maximize the efficiency of astronomical data distribution by integrating state-of-the-art technologies and artificial intelligence.

Start Querying Explore Database

Our Resources



Processing & Memory

- **CPU Cores: 1.128**
Platform: Intel Xeon 5th Generation
Total RAM memory (TB): 23,5
Type: DDR5 4800 MHz ECC/REG

Storage & GPU

- **SSD 1 & 2: 1 TB PCI-E 4.0 NVMe (Operating System, & Applications)**
SSD 3: 100 TB PCI-E 4.0 NVMe
SSD 4 & 5 (next months): 1 Pb PCI-E 4.0 NVMe
~146 (GPUS, Nvidia H100 / A100 / RTX 6000 Ada)
CUDA Cores: 1.422.336 , for 2026 including 2x H800 planned



ATLAS MACHINE



Home Performance Metrics System Status



HPC cluster composed of 21 (twenty-one) racks, including 17 (seventeen) compute racks, 1 (one) rack for control servers, 2 (two) racks for the parallel file system (FS) with a total capacity of 1.3 PB, and 1 (one) rack housing InfiniBand EDR switches. Each compute rack contains 8 (eight) servers, totaling 136 (one hundred and thirty-six) compute nodes. Each compute server is equipped with eight V100 GPUs. The nominal power consumption of the cluster is 610 kVA, and its peak performance reaches 8.9 PFlops (Rpeak).

**Train from Scratch a 70B parameter LLM
~1-2.5 month.**

1088 GPUS + additional 1.3 Pb Storage

List	Rank	System	Vendor	Total Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
11/2025	295	Bull 4029GP-TVRT, Xeon Gold 6240 18C 2.6GHz, NVIDIA Tesla V100, Infiniband EDR	EVIDEN	91,936	4.38	8.85	546.58



AGENTIC ASTRONOMER

TVS LLMS and Foundational Models

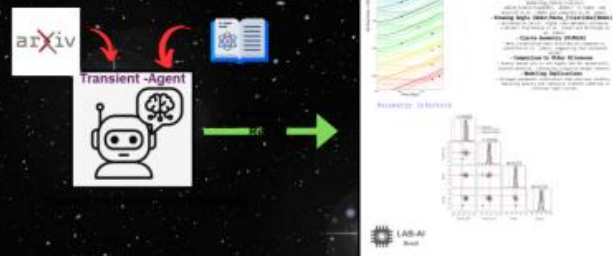
Coordinators: Clecio R. Bom

Slack channel: [#tvs-llms-foundational-models](#)

Telecons: TBD

Goals

- Develop TVS related Agents and LLMs to assist TVS community
- Integration between informed AI agents, LLMS and recommendation systems with Brokers
- Built community efforts in AI Agents and Generative AI in TVS



Current Status

Portal for development of AI agents integrating public alerts + public available datasets already-up


Summary

- AI can do more than classification.
- Agentic AI can encapsulate contextual information efficiently
- There are many more modules going on (e.g. Cosmology).
- AI can help people get informed decisions on what is the best transient for follow-up based on the possible scientific outcome, and optimize spectro and photometric telescope time.
- Local tunned Infrastructure for astronomy, Local Run and development of LLMs, fine-tunning, Rags.

Demonstration vídeo:

<https://www.youtube.com/watch?v=WojdkXPqhPw> - ABC and LLM report (LLMs detects a fail fit due to a weak choice of modelling right away)

<https://youtu.be/vIzQ3tLftxY> - Foundation Model API and LLM coding for Astro



Chasing Cosmic Rarities: Intelligent Agents and Physical Modeling in the Time-Domain Frontier

Clecio R. Bom, Bernardo Fraga, Phelipe Darc, Gabriel Teixeira
Andre Santos, Gustavo Schwarz, Luidhy Santana-Silva
debom@cbpf.br

RAPID Response:
Hot-wiring the Next Generation of Time-Domain Science

