MULTI-STEP REASONING FOR DEEP-FIELD GALAXY DETECTION WITH UNCERTAINTY-AWARE HUMAN FEEDBACK

Anonymous authors
Paper under double-blind review

ABSTRACT

Source detection in deep-field astronomical imaging poses a formidable challenge characterized by extreme dynamic ranges in scale and flux, low signal-to-noise ratios, ambiguous morphological boundaries, and severe crowding. A fundamental yet overlooked bottleneck is the entanglement of distinct uncertainty sources: aleatoric uncertainty, arising from stochastic background noise, sensor limits, and intrinsic photon shot noise; and epistemic uncertainty, stemming from model underconfidence and out-of-distribution morphology. Disentangling these two uncertainty modes is essential to diagnose whether detection failures are physically bounded by observational limits or algorithmically constrained by representation capacity. To resolve this, we propose a multi-step reasoning framework that couples fine-grained uncertainty disentanglement with a human-in-the-loop calibration loop for robust galaxy detection. Following background stabilization and faint-structure enhancement, a Continuous Thought Machine (CTM) dynamically navigates heterogeneous signal-strength regimes across iterative reasoning steps, progressively integrating multi-stage spatial representations. Crucially, the framework quantifies and decouples aleatoric from epistemic uncertainty maps, isolating ambiguous physical blending from algorithmic failure modes. These spatial uncertainty representations are mapped to an interactive interface that aligns machine uncertainty with expert perceptual intuition, enabling targeted human feedback to iteratively refine model boundaries. Evaluated on simulated deep-field benchmarks, our framework achieves a recall of 88.99% (a 6.30 percentage point improvement over the standard SExtractor baseline of 82.69%). Finally, we demonstrate the real-world utility and cross-instrument generalizability of our framework on ultra-deep observations from the James Webb Space Telescope (JWST) and the Euclid space mission.

1 INTRODUCTION

Deep-field astronomical imaging has become a fundamental approach for studying galaxy formation, evolution, and large-scale cosmic structuresWilliams et al. (2000); Jeon & Bromm (2019). With the development of advanced space telescopes such as the James Webb Space Telescope (JWST)Gardner et al. (2006) and the Euclid missionTroja et al. (2022), astronomical surveys are producing unprecedented volumes of high-resolution observations containing extremely faint and diverse galaxy populationsCowie et al. (1995). Reliable source detection serves as the foundation for subsequent scientific analyses, including morphological characterization and photometric measurementCoe et al. (2006); Siddiquee et al. (2025); Sterken & Manfroid (2012); Romanishin (2006). However, detecting galaxies in deep-field images remains challenging due to large variations in source brightness, low signal-to-noise ratios, complex galaxy morphologies, and severe source blending.

Traditional astronomical source detection methods, such as SExtractor, rely on background estimation, threshold-based detection, and deblending to identify astronomical sourcesBertin & Arnouts (1996); Teeninga et al. (2016).Although these methods offer efficient and interpretable detection pipelines,they often struggle to detect faint galaxies in crowded fields, where weak source signals can be confused with background fluctuations or blended with nearby objects. Recent advances in

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

deep learning have enabled more effective astronomical image analysis by learning hierarchical features directly from dataGonzález et al. (2018); Andreon et al. (2000). Convolutional neural networks and transformer-based architectures have shown promising results in astronomical object detection and segmentationAnanthamoorthy et al. (2024); Raab (2002); Andreon et al. (2000); Odewahn et al. (1992); Pan et al. (2024); Biswas et al. (2025). While these approaches have improved detection performance, uncertainty is often represented by a single confidence score, providing limited insight into the underlying causes of prediction failuresKhalifa et al. (2017).

A critical but under-investigated issue in deep-field source detection is that prediction uncertainty originates from fundamentally different sources. Aleatoric uncertainty is induced by inherent observational limitations, including photon noise, detector noise, and background fluctuations, which cannot be eliminated even with improved models. In contrast, epistemic uncertainty reflects limitations in the learned representation, such as insufficient training data, unfamiliar galaxy morphology, or model overconfidence in ambiguous regionsLoredo (2004); Singh & Pandey (2025); Cheng et al. (2026); Tamames-Rodero et al. (2025); Hatfield (2022); Perreault Levasseur et al. (2017). Without disentangling these two uncertainty modes, it remains difficult to determine whether a missed detection is caused by unavoidable observational ambiguity or by insufficient model capability. Such ambiguity limits the interpretability and reliability of automated astronomical analysis systems.

To address these challenges, we propose an uncertainty-aware multi-step reasoning framework based on the Continuous Thought Machine (CTM) for robust galaxy detection in deep-field astronomical imagesGonzález et al. (2018); Hao et al. (2024); Gozeten et al. (2026). Unlike conventional feedforward architectures that directly map images to predictions, our framework performs iterative reasoning through multiple latent states, enabling progressive refinement of spatial representations across heterogeneous signal regimes. Furthermore, we introduce an uncertainty disentanglement strategy that separately models aleatoric and epistemic uncertainty, generating spatial uncertainty maps that reveal different sources of detection ambiguity. These uncertainty representations are further integrated into a human-in-the-loop calibration mechanism, allowing experts to provide targeted feedback on uncertain regions and iteratively refine model predictionsHu et al. (2025); Fu et al. (2024).

We evaluate our framework on simulated deep-field benchmarks and further validate its generalization ability on real observations from JWST and Euclid. Experimental results demonstrate that our method achieves a recall of 88.99% on galaxy-cluster observations, improving over the standard SExtractor baseline (82.69%) by 6.30 percentage points. Our contributions are summarized as follows: First, we introduce a Continuous Thought Machine-based multi-step reasoning framework Fu et al. (2022); Yang et al. (2026); Gozeten et al. (2026) for deep-field galaxy detection, enabling progressive feature refinement under challenging observational conditions.

In addition, we incorporate both aleatoric and epistemic uncertainty into the framework to account for observational ambiguity and prediction inconsistency, providing spatial uncertainty maps to identify regions requiring further inspection. Finally, we design a human-in-the-loop interactive website that incorporates uncertainty-guided expert feedback to support model refinement and improve the reliability of galaxy detection.

2 DATA

We evaluate the proposed framework on both simulated and real deep-field astronomical observations.

2.1 CSST SIMULATED DATASET

We use simulated deep-field images from the China Space Station Telescope (CSST) survey. The dataset contains galaxies with diverse flux levels, morphologies, source densities, and signal-to-noise conditions, including both isolated and crowded sources. Object-level annotations are provided by the simulation catalog, from which pixel-level instance masks are generated for model training and evaluation. The controlled simulation setting enables quantitative comparison under different observational conditions, particularly for faint and blended sources. Representative examples at different noise levels are shown in Fig. 1a.

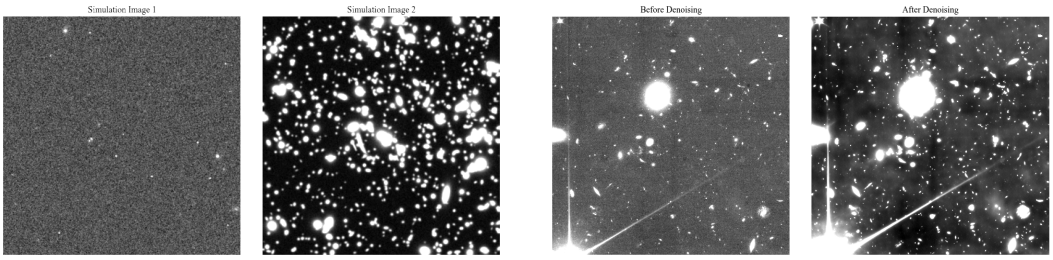

(a) CSST simulated deep-field images. (b) JWST observations before and after denoising.

Figure 1: Representative examples of the datasets used in our experiments.

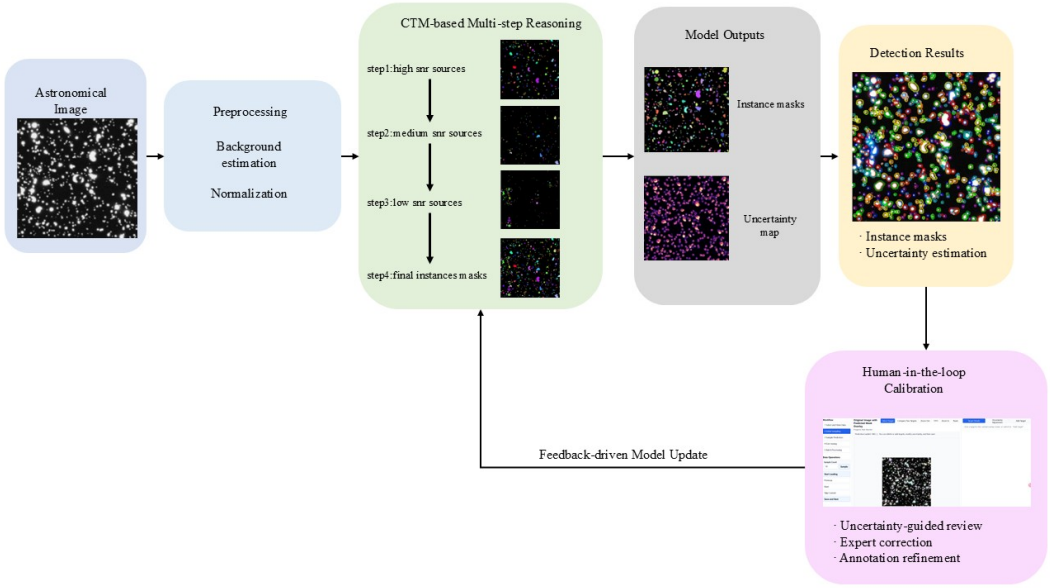

Figure 2: Human-in-the-loop framework for uncertainty-aware galaxy source detection and expert-guided refinement.

2.2 JWST Real Observations

We further evaluate the proposed framework on real deep-field observations from the James Webb Space Telescope (JWST). In contrast to the controlled CSST simulations, JWST images contain more complex background structures,instrumental effects, and highly diverse galaxy morphologies, providing a more realistic test environment for source detection.

Before model training and evaluation, we apply ASTERIS to suppress observational noise while preserving faint astronomical structures Guo et al. (2026).An example of the observations before and after ASTERIS processing is shown in Fig. 1b.

3 Methodology

3.1 Framework Overview

Deep-field galaxy detection is challenging due to faint signals, complex morphologies, and severe source blending. To address these challenges, we propose an uncertainty-aware multi-step reasoning framework based on the Continuous Thought Machine (CTM), as illustrated in Fig. 2.

The proposed framework consists of three main components: (1) an SNR-guided CTM reasoning module that progressively detects sources with different signal strengths, (2) an uncertainty disen-

tanglement module that characterizes observational and model-related uncertainty, and (3) a human-in-the-loop interactive interface for uncertainty visualization and expert inspection.

Given an astronomical image I, the framework first extracts spatial features together with signal-to-noise information. The model then performs progressive reasoning from relatively high-SNR sources to more challenging low-SNR sources. At each reasoning step, the current predictions and uncertainty information are fed back into the model, allowing ambiguous regions to be progressively refined.

3.2 Continuous Thought Machine-based Multi-step Reasoning

Unlike conventional feed-forward networks that generate predictions in a single inference pass, our framework performs galaxy detection through multiple reasoning steps. To accommodate the large variation in source visibility, galaxy sources are organized into three signal-strength stages, corresponding to relatively high-, medium-, and low-SNR sources.

At reasoning step t, the CTM iteratively refines its latent state as:

$$H_t = f_\theta(H_{t-1}, Z), \tag{1}$$

where H_t denotes the latent state after the t-th reasoning step,H_{t-1} is the previous reasoning state, and Z represents the encoded image features extracted by the backbone network.

The reasoning process proceeds progressively from relatively prominent sources to increasingly faint sources. Predictions obtained at earlier stages provide contextual information about regions that have already been sufficiently explained, allowing subsequent stages to concentrate on unresolved or ambiguous regions. Through this iterative refinement process, the model gradually improves its representation of difficult sources, particularly faint and blended galaxies.

3.3 Uncertainty Estimation

In addition to detection results, our framework estimates two complementary types of uncertainty: aleatoric uncertainty and epistemic uncertainty.

Aleatoric uncertainty mainly reflects ambiguity inherent in astronomical observations, including low signal-to-noise ratios, background fluctuations,and source blending. In contrast, epistemic uncertainty reflects inconsistency in model predictions across different reasoning stages.

We represent the uncertainty estimation as two complementary components:

$$\mathcal{U} = \{U_a, U_e\}, \tag{2}$$

where U_a and U_e denote the aleatoric uncertainty and epistemic uncertainty, respectively. These spatial uncertainty maps provide pixel-level localization of ambiguous regions, enabling a more interpretable assessment of prediction reliability beyond a single scalar confidence score.

Furthermore, uncertainty is estimated throughout the multi-step reasoning process, enabling us to observe how uncertain regions evolve as the model progressively refines its predictions.

3.4 Human-in-the-loop Interactive Interface

To facilitate expert inspection of uncertain detections, we develop a web-based human-in-the-loop interface that jointly visualizes galaxy detection results and spatial uncertainty maps, as shown in Fig. 3.

For each astronomical image, the proposed human-in-the-loop module provides the detected galaxy instances together with the corresponding spatial uncertainty maps. Instead of relying solely on prediction scores, the uncertainty information highlights regions with ambiguous detections,such as faint sources, overlapping structures, and uncertain boundaries.

The interaction process consists of three sequential stages:

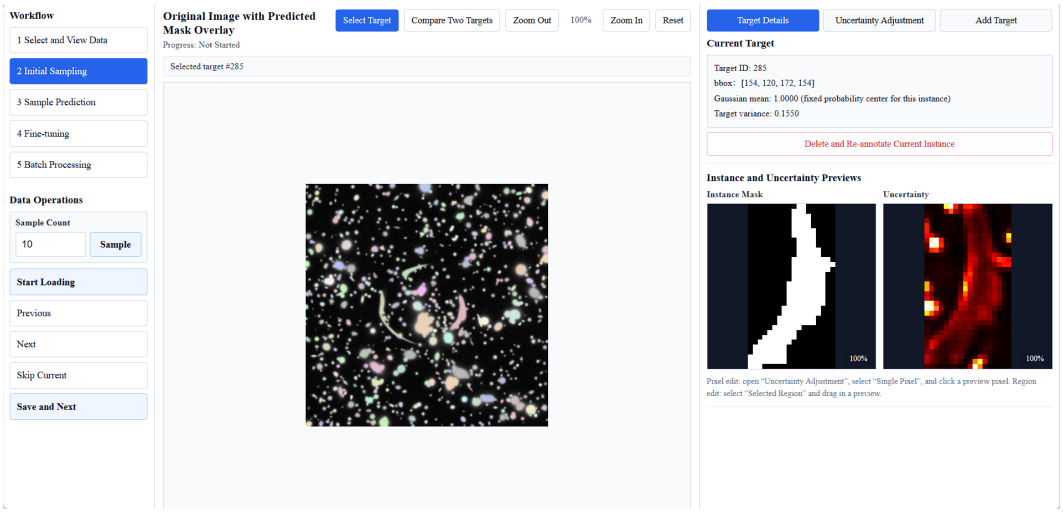

Figure 3: Overview of the proposed human-in-the-loop interactive web interface for galaxy source detection, uncertainty visualization, and expert feedback.

(1) Uncertainty-guided detection:The model generates galaxy instance predictions and spatial uncertainty maps, which are used to identify potentially unreliable regions.

(2) Interactive visualization and inspection:The detected instances and uncertainty maps are presented to experts,allowing efficient examination of ambiguous predictions and challenging astronomical sources.

(3) Expert feedback and model refinement:Experts provide corrections or annotations for uncertain regions. The collected feedback is then used for further model refinement, forming an iterative interaction loop between automated detection and expert knowledge. The interface provides an intuitive way to connect model uncertainty with expert inspection while reducing the need for exhaustive manual review of all detected sources.

4 Experiments

4.1 Comparison with SExtractor

We evaluate the detection performance of the proposed method on the CSST simulated datasets and compare it with SExtractor, a widely used astronomical source detection algorithm. A qualitative example from the CSST-Cluster dataset is presented in Fig. 4, which shows the detection results at different reasoning stages. To illustrate the internal reasoning process within a single stage,we select Stage 2 as a representative example.Fig. 5 shows the evolution of instance predictions across successive thought ticks within this stage.

Table 1 reports the quantitative results. Since the available annotations do not exhaustively cover all visible sources, valid detections without corresponding annotations may be counted as false positives, potentially leading to underestimated precision. Therefore, we primarily use recall to assess the ability of each method to recover annotated galaxy sources.

On the CSST-Galaxy dataset, the proposed method achieves a precision of 96.76% and a recall of 90.38%, compared with 84.02% and 77.02% for SExtractor, respectively. On the more crowded CSST-Cluster dataset, our method achieves a recall of 88.99%, compared with 82.69% for SExtractor, representing an improvement of 6.30 percentage points. These results indicate that the proposed framework recovers a larger fraction of annotated sources across both datasets, including crowded galaxy fields.

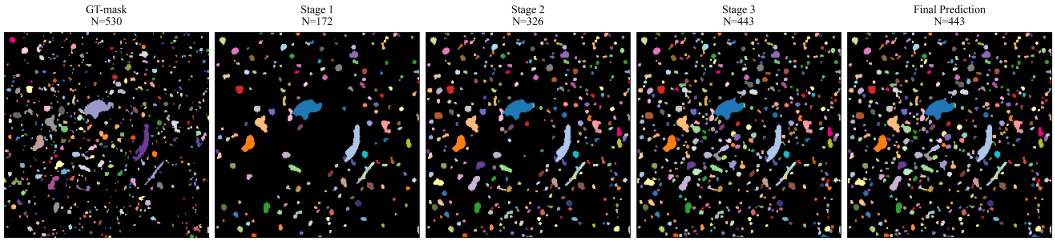

Figure 4: Qualitative example of galaxy instance detection at different stages on the CSST-Cluster dataset.The figure show the ground-truth mask, predictions from Stages 1–3, and the final merged prediction.

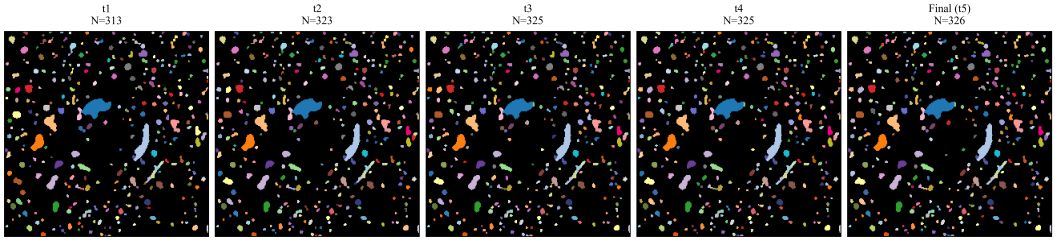

Figure 5: Intermediate instance segmentation results at successive thought ticks within Stage 2 on the CSST-Cluster dataset.

4.2 Multi-step Reasoning and Uncertainty Visualization

To further analyze the behavior of the proposed framework, we visualize the intermediate predictions generated during the multi-step reasoning process,together with the final detection result and the corresponding uncertainty map.

As shown in Fig. 6, the detection results are progressively refined across different reasoning steps. Early reasoning steps mainly identify relatively prominent sources, while subsequent steps further recover additional sources from the remaining ambiguous regions. The final prediction integrates the information accumulated throughout the reasoning process. The uncertainty map provides complementary spatial information about the reliability of the final prediction. Higher uncertainty is generally observed around faint, blended, or ambiguous sources, highlighting regions that may require further inspection. Similar behavior can be observed on both the CSST simulated data and real JWST observations.

4.3 Application on JWST Observations

Finally, we evaluate the proposed framework on real JWST deep-field observations processed by ASTERIS Guo et al. (2026).

Since the reference annotations for the JWST observations are constructed based on SExtractor detections, a direct quantitative comparison with SExtractor would not constitute an independent evaluation. Therefore, for the JWST dataset, we only report the detection performance of the proposed model with respect to the constructed reference annotations.

As reported in Table 1, the proposed method achieves a recall of 89.64% on the JWST observations. It should be noted that the relatively low precision is partly affected by the incompleteness of the reference annotations. The proposed model detects additional source candidates that are not included in the SExtractor-based reference labels, and these unmatched detections are consequently counted as false positives under the current evaluation protocol.

As shown in Fig. 7, the proposed model produces additional source candidates beyond those included in the SExtractor-based reference annotations. Although these unmatched predictions are counted as false positives under the current evaluation protocol, several of them correspond to visu-

Table 1: Detection performance of SExtractor and the proposed model on the CSST and JWST datasets.

Dataset	Method	Precision (%)	Recall (%)
CSST-Galaxy	SExtractor	84.02	77.02
	Our Model	96.76	90.38
CSST-Cluster	SExtractor	89.84	82.69
	Our Model	84.34	88.99
JWST	Our Model	58.14	89.64

Note: The JWST reference annotations are constructed based on SExtractor detections. Therefore, SExtractor is not included as an independent baseline in the current JWST evaluation.

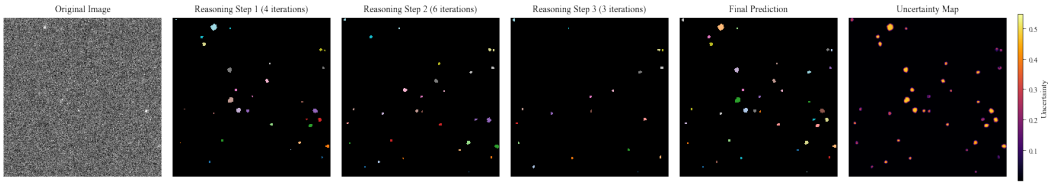

(a) Progressive reasoning results on a CSST simulated image.

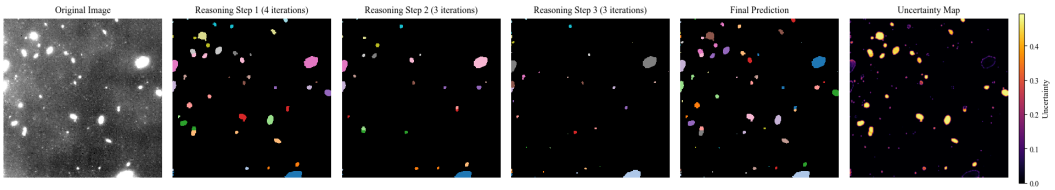

(b) Progressive reasoning results on a real JWST observation.

Figure 6: Visualization of the multi-step reasoning process on CSST simulated data and real JWST observations. From left to right, each example shows the original image, intermediate predictions at different reasoning steps, the final detection result, and the corresponding uncertainty map.

ally identifiable source-like structures in the original JWST images. This partly explains the relatively low measured precision.

The qualitative results further show that the proposed framework can recover galaxy sources under realistic observational conditions, while the corresponding uncertainty maps highlight ambiguous, faint, and blended regions for further inspection.

5 Conclusion

In this work, we propose an uncertainty-aware multi-step reasoning framework for deep-field galaxy detection based on the Continuous Thought Machine.The framework integrates progressive reasoning, uncertainty estimation, and a human-in-the-loop interactive interface to provide both source detection results and spatial uncertainty information.

Experiments on the CSST simulated datasets demonstrate the effectiveness of the proposed method under different source-density conditions. On the CSST-Galaxy dataset, our model achieves a precision of 96.76% and a recall of 90.38%, compared with 84.02% and 77.02% for SExtractor, respectively. On the more crowded CSST-Cluster dataset, the proposed method increases the recall from 82.69% to 88.99%, although with a moderate decrease in precision from 89.84% to 84.34%. These results indicate that the proposed multi-step reasoning strategy is particularly effective in recovering a larger fraction of sources in crowded fields.

We further evaluate the framework on real JWST deep-field observations processed by ASTERIS, where the model achieves a recall of 89.64% with respect to the constructed reference annotations.

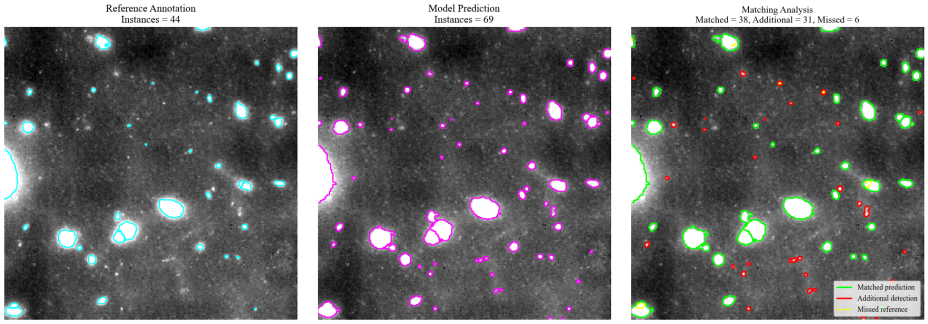

(a) JWST observation example 1.

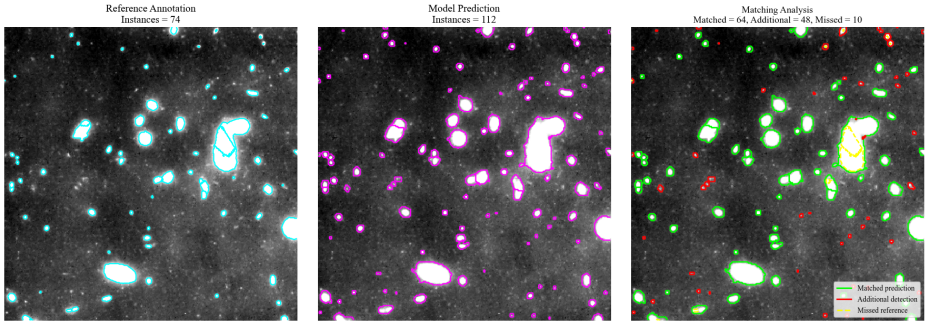

(b) JWST observation example 2.

Figure 7: Matching analysis between the proposed model predictions and SExtractor-derived reference annotations on real JWST observations. Green contours indicate matched detections, red contours denote additional detections not covered by the reference annotations, and yellow dashed contours indicate missed reference sources.

The visualized reasoning process shows progressive refinement across multiple steps, while the uncertainty maps provide complementary spatial information for identifying ambiguous regions and supporting further expert inspection.

Overall, the results demonstrate the potential of uncertainty-aware multi-step reasoning for reliable and interpretable galaxy source detection in future large-scale astronomical surveys.

AI USE STATEMENT

Generative AI tools were used in this work to assist with language polishing,code development and debugging, and the refinement of technical descriptions.They were also used to discuss implementation details and to improve the clarity of figures and experimental presentation.

All AI-assisted code was manually reviewed, tested, and validated by the authors.All experimental results, numerical values, figures, and conclusions reported in this paper were obtained and verified by the authors. Generative AI tools were not used to fabricate experimental data, results, or citations.

The authors reviewed and revised all AI-assisted content and take full responsibility for the final content of this work, including the text, claims, code, and other artifacts produced with the aid of generative AI.

REFERENCES

B Ananthamoorthy, Debbijoy Bhattacharya, P Sreekumar, and Swathi B. Detection of faint sources by the ultraviolet imaging telescope onboard astrosat using poisson distribution of background. *The Astronomical Journal*, 168(1):22, 2024.

Stefano Andreon, Giorgio Gargiulo, Giuseppe Longo, Roberto Tagliaferri, and Nicola Capuano. Wide field imaging—i. applications of neural networks to object detection and star/galaxy classification. *Monthly Notices of the Royal Astronomical Society*, 319(3):700–716, 2000.

Emmanuel Bertin and Stephane Arnouts. Sextractor: Software for source extraction. *Astronomy and astrophysics supplement series*, 117(2):393–404, 1996.

Ujan Biswas, Azizul Abedin Azmi, Modan Mohan Sarker, Tanzila Afrin, Faysal Ahmed, and Md Samiul Islam. A vision transformer adaptation of simclrv2 for self-supervised galaxy detection. In *2025 IEEE International Conference on Signal Processing, Information, Communication and Systems (SPICSCON)*, pp. 1–6. IEEE, 2025.

Kai Cheng, Ruoqi Wang, and Qiong Luo. Estimating uncertainty in galaxy morphology classification. *arXiv preprint arXiv:2608.08398*, 2026.

Dan Coe, Narciso Benítez, Sebastián F Sánchez, Myungkook Jee, Rychard Bouwens, and Holland Ford. Galaxies in the hubble ultra deep field. i. detection, multiband photometry, photometric redshifts, and morphology. *The Astronomical Journal*, 132(2):926–959, 2006.

Lennox L Cowie, Esther M Hu, and Antoinette Songaila. Faintest galaxy morphologies from hst wfpc2 imaging of the hawaii survey fields. *arXiv preprint astro-ph/9507055*, 1995.

Ming-Xiang Fu, Yu Song, Jia-Meng Lv, Liang Cao, Peng Jia, Nan Li, Xiang-Ru Li, Ji-Feng Liu, A-Li Luo, Bo Qiu, et al. A versatile framework for analyzing galaxy image data by incorporating human-in-the-loop in a large vision model. *Chinese Physics C*, 48(9):095001, 2024.

Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting for multi-step reasoning. *arXiv preprint arXiv:2210.00720*, 2022.

Jonathan P Gardner, John C Mather, Mark Clampin, Rene Doyon, Matthew A Greenhouse, Heidi B Hammel, John B Hutchings, Peter Jakobsen, Simon J Lilly, Knox S Long, et al. The james webb space telescope. *Space Science Reviews*, 123(4):485–606, 2006.

Roberto E González, Roberto P Munoz, and Cristian A Hernández. Galaxy detection and identification using deep learning and data augmentation. *Astronomy and computing*, 25:103–109, 2018.

Alperen Gozeten, Muhammed Ildiz, Xuechen Zhang, Hrayr Harutyunyan, Ankit Singh Rawat, and Samet Oymak. Continuous chain of thought enables parallel exploration and reasoning. In *International Conference on Learning Representations*, volume 2026, pp. 82881–82914, 2026.

Yuduo Guo, Hao Zhang, Mingyu Li, Fujiang Yu, Yunjing Wu, Yuhan Hao, Song Huang, Yongming Liang, Xiaojing Lin, Xinyang Li, et al. Deeper detection limits in astronomical imaging using self-supervised spatiotemporal denoising. *Science*, 392(6797):eady9404, 2026.

Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. Training large language models to reason in a continuous latent space. *arXiv preprint arXiv:2412.06769*, 2024.

Peter Hatfield. Quantification of unknown unknowns in astronomy and physics. *arXiv preprint arXiv:2207.13993*, 2022.

Dong Hu, Guangzhong Zhou, Jingda Wu, and Chao Huang. Trust-calibrated human-in-the-loop reinforcement learning for safe and efficient autonomous navigation. *IEEE Internet of Things Journal*, 12(18):39029–39041, 2025.

Myoungwon Jeon and Volker Bromm. Signature of the first galaxies in jwst deep field observations. *Monthly Notices of the Royal Astronomical Society*, 485(4):5939–5950, 2019.

Nour Eldeen M Khalifa, Mohamed Hamed N Taha, Aboul Ella Hassanien, and IM Selim. Deep galaxy: Classification of galaxies based on deep convolutional neural networks. *arXiv preprint arXiv:1709.02245*, 2017.

Thomas J Loredo. Accounting for source uncertainties in analyses of astronomical survey data. In *AIP Conference Proceedings*, volume 735, pp. 195–206. American Institute of Physics, 2004.

SC Odewahn, EB Stockwell, RL Pennington, Roberta M Humphreys, and WA Zumach. Automated star/galaxy discrimination with neural networks. In *Digitised Optical Sky Surveys: Proceedings of the Conference on 'Digitised Optical Sky Surveys', Held in Edinburgh, Scotland, 18–21 June 1991*, pp. 215–224. Springer, 1992.

Jia-Shu Pan, Yuan-Sen Ting, and Jie Yu. Astroconformer: The prospects of analysing stellar light curves with transformer-based deep learning models. *Monthly Notices of the Royal Astronomical Society*, 528(4):5890–5903, 2024.

Laurence Perreault Levasseur, Yashar D Hezaveh, and Risa H Wechsler. Uncertainties in parameters estimated with neural networks: Application to strong gravitational lensing. *The Astrophysical Journal Letters*, 850(1):L7, 2017.

Herbert Raab. Detecting and measuring faint point sources with a ccd. In *Proceedings of Meeting on Asteroids and Comets in Europe (MACE)(Visnjam)*, pp. 1–12. Citeseer, 2002.

W Romanishin. An introducation to astronomical photometry using ccds, 2006.

Md Fairuz Siddiquee, Md Mehedi Hasan, Shifat E Arman, Md Shahedul Islam, and AKM Azad. Exploring celestial classification: Astrophysical features-guided machine learning for spectral and morphological analysis. *Astronomy and Computing*, pp. 101048, 2025.

Raahul Singh and Ashutosh Pandey. Starkindler: An uncertainty aware objective for photometric redshift estimation. *arXiv preprint arXiv:2512.22566*, 2025.

Christiaan Sterken and Jean Manfroid. *Astronomical photometry: a guide*. Springer Science & Business Media, 2012.

Víctor Tamames-Rodero, Andrés Moya, Luis Manuel Sarro, and Roberto Javier López-Sastre. Unveiling the power of uncertainty: A journey into bayesian neural networks for stellar dating. *Astronomy and Computing*, 52:100957, 2025.

Paul Teeninga, Ugo Moschini, Scott C Trager, and Michael HF Wilkinson. Statistical attribute filtering to detect faint extended astronomical sources. *Mathematical Morphology-Theory and Applications*, 1(1):100–115, 2016.

Antonino Troja, Isaac Tutusaus, and Jenny G Sorce. Euclid in a nutshell. *arXiv preprint arXiv:2211.09668*, 2022.

Robert E Williams, Stefi Baum, Louis E Bergeron, Nicholas Bernstein, Brett S Blacker, Brian J Boyle, Thomas M Brown, C Marcella Carollo, Stefano Casertano, Riccardo Covarrubias, et al. The hubble deep field south: formulation of the observing campaign. *The Astronomical Journal*, 120(6):2735–2746, 2000.

Tong Yang, Yu Huang, Yingbin Liang, and Yuejie Chi. Multi-head transformers provably learn symbolic multi-step reasoning via gradient descent. *Advances in Neural Information Processing Systems*, 38:116194–116285, 2026.