



The Role of Artificial Intelligence in the SKA Era

Philipp Denzel ¹, Frank-Peter Schilling ¹ and Elena Gavagnin ^{1,2}

¹*Centre for Artificial Intelligence, Zurich University of Applied Sciences ZHAW, Technikumstrasse 71, Winterthur 8400, Switzerland*

²*Institute of Business Information Technology, Zurich University of Applied Sciences ZHAW, Theaterstrasse 17, Winterthur 8400, Switzerland*

E-mail: denp@zhaw.ch, scik@zhaw.ch, gava@zhaw.ch

The Square Kilometre Array Observatory (SKAO) will usher in an era of unprecedented data complexity and scientific opportunity in radio astronomy, producing petabyte-scale datasets and terabit-per-second streams that challenge traditional analysis paradigms. Artificial Intelligence (AI) stands at the forefront of this transformation, offering scalable, adaptive solutions to the most pressing problems in radio astronomy and astrophysics. This chapter explores the pivotal role of AI in the SKA era, from real-time operations to scientific discovery. We examine how deep learning models enable automated source detection, radio-frequency interference mitigation, anomaly detection, and parameter inference, while generative approaches accelerate sky simulations, calibration, and imaging. Reinforcement learning promises dynamic scheduling and autonomous system control, and federated learning could address the distributed nature of SKA data. Beyond performance, we emphasize the necessity of explainability, uncertainty quantification, and physics-informed inductive biases to ensure scientific integrity. By mapping SKAO's core challenges – data volume, complexity, and interpretability – onto modern AI methodologies, we review how deep learning, self-supervised frameworks, and probabilistic models can unlock new frontiers in cosmology, galaxy evolution, and time-domain astrophysics. AI is not merely an automation tool for coping with scale. It is a catalyst for discovery, redefining how we observe, model, and understand the Universe.

1 Introduction

The Square Kilometre Array Observatory (SKAO) will transform radio astronomy by delivering orders-of-magnitude gains in sensitivity, angular resolution, and survey speed across a wide range of radio frequencies. With that leap comes a scale and complexity shock: terabit-per-second raw data rates streaming off the arrays, exabyte-scale archives, and a heterogeneous landscape of observing modes (calibrated visibilities, interferometric imaging, beam-formed products) that stress every step of the scientific workflow – from calibration and imaging to inference and discovery. Traditional pipelines meticulously engineered around manual parameter tuning, human-in-the-loop quality control, and post-facto archival mining are not designed to function in real time, nor to maintain statistical optimality in the face of rapidly varying instrumental and environmental conditions. Artificial Intelligence (AI) – encompassing supervised, self-/semi-supervised, generative, probabilistic, reinforcement, physics-informed, and federated paradigms – has matured precisely as SKA-class challenges crystallize. Across precursor and pathfinder experiments and data challenges, AI has already demonstrated the ability to denoise, detect, classify, reconstruct, emulate, and infer at scale while quantifying uncertainty and respecting physical constraints (e.g., [Bonaldi and Braun, 2018](#); [Bonaldi et al., 2025](#); [Lastufka et al., 2024](#); [Riggi et al., 2023](#); [Acharya et al., 2024a](#)). The rationale for a broad AI integration within SKA is therefore twofold:

1. **Operational necessity:** automate (near) real-time decision-making and data triage, preserve faint signals under extreme compression, and stabilize calibration/imaging under direction-dependent and time-variable systematics.
2. **Scientific leverage:** extract non-Gaussian information, accelerate forward modelling, and deliver uncertainty-aware inferences at the field level; enabling robust cosmology, galaxy evolution, and time-domain astrophysics.

This chapter underscores AI as a central component to SKAO’s operational viability and scientific reach. We categorize the SKAO’s problems, list the failure modes of status-quo approaches, and propose AI-driven remedies that emphasize explainability, uncertainty quantification, and physics-aware inductive biases. We also outline Machine Learning Operations (MLOps) and governance practices required for reproducible, auditable data science at exascale. Finally, we review AI techniques for forward sky modelling, ranging from generative models for cosmological light cones to physics-informed emulation of baryonic processes. These methods are indispensable accelerators that enable rapid hypothesis testing and bridge the gap between theory and observation at SKA scales.

Building on this foundation, we now turn to a selection of specific scientific domains where AI offers transformative capabilities. The following section outlines key problem areas in SKA science, ranging from early-Universe cosmology to time-domain astrophysics, where traditional approaches face fundamental limitations, and where AI-driven methods promise significant advances in sensitivity, scalability, and interpretability.

2 Scientific problem areas

The scientific opportunities enabled by the SKAO come with analysis regimes that stress or break traditional methods: extreme data volumes and rates, highly non-Gaussian signals, complex instrument-sky systematics, and the need for real-time decision-making. In this Section, we identify problem classes in astrophysics and radio astronomy where AI, in particular, deep learning methods have the potential to materially outperform more traditional pipelines; either by improving sensitivity and robustness, reducing latency and cost, or enabling/automating analyses that are otherwise impractical. For each area, we outline

- (i) the scientific objective and limitations of status-quo approaches,
- (ii) AI-based strategies likely to deliver leaps in performance, referencing a selection of recent works implementing such strategies,
- (iii) followed by risks and challenges induced through the application of AI, e.g., uncertainty quantification, interpretability, and physics-aware inductive biases.

Technical definitions of the AI methods referenced throughout this section are collected in Appendix A for reference.

2.1 Radio interferometric imaging & sky simulations

Radio interferometric imaging is governed by the radio interferometer measurement equation (RIME), an inherently ill-posed inverse problem where minor perturbations in visibility data can produce substantial deviations in reconstructed images. For decades, the CLEAN algorithm (Wright et al., 2024) has dominated this domain, relying on the heuristic assumption of a sparse sky composed primarily of point sources. While widely adopted, CLEAN exhibits critical limitations for exascale data requirements of the SKA era: poor computational scalability with image size and source count, dependence on manual expert intervention for convergence, reduced accuracy for faint or diffuse structures, and the absence of statistically calibrated confidence measures. More modern variants such as SARA (Sparsity Averaging Reweighted Analysis; Carrillo et al., 2012) improve reconstruction fidelity but remain computationally prohibitive. Recent developments, such as Bluebird (Tolley et al., 2025), leverage eigendecomposition and GPU acceleration to reduce time-to-solution for large-scale imaging, yet fundamental challenges persist (see also Zhang et al., 2026).

To benchmark next-generation algorithms, realistic sky simulations have become indispensable. Simulation frameworks like OSKAR (Dulwich and Mort, 2017) and Karabo (Sharma et al., 2026) which produce synthetic interferometric visibilities in the uv -plane by modelling a full signal chain, enable end-to-end simulation of telescope characteristics, including beam patterns, system noise, calibration errors, and ionospheric distortions. Complementary hydrodynamical simulations – most relevant for *SKA-Mid* science at low to intermediate redshifts – provide physically motivated mock skies incorporating diffuse emission, extended structures, and multi-scale interactions, far surpassing simplistic point-source or Gaussian models in complexity. Magneto-hydrodynamics are essential for capturing magnetic-field-driven phenomena such as synchrotron radiation, Faraday rotation, and polarization. However, a persistent computational tension exists between the need

for cosmological volumes to sample rare sources and the sub-kpc resolution required to match *SKA-Mid*'s sub-arcsecond imaging capabilities. Current “zoom-in” simulations achieve the necessary resolution for phenomena like jet formation and turbulence but only within limited volumes. Moreover, sensitivities in subgrid physics tuning parameters exacerbate the computational burden, underscoring the need for efficient emulation strategies.

Interferometric deconvolution is increasingly utilizing AI-based solutions to these challenges, using deep learning architectures such as convolutional neural networks (CNNs; Appendix A.3.1), U-Nets (Appendix A.3.3), and residual networks (e.g. ResNets; Appendix A.3.2). These models aim to accelerate processing, enhance reconstruction fidelity for extended emission, and mitigate artefacts in dirty images. For instance, AIRI (AI for Regularization in radio-interferometric Imaging; Terris et al., 2022) integrates a neural denoise within an optimization framework, achieving image quality comparable to SARA at significantly reduced computational cost. Schmidt et al. (2022) explore end-to-end U-Net architectures for direct visibility-to-image reconstruction, trading extreme inference speed against risks of hallucinations and poor generalization. POLISH (Connor et al., 2022), a ResNet-based super-resolution model, designed to “polish” dirty images, improves signal-to-noise ratio (SNR) by 3 dB over CLEAN and is tailored for high-throughput surveys such as DSA-2000. The R2D2 (Residual-to-Residual DNN series of high-Dynamic range imaging; Aghabiglou et al., 2024, 2025; Tajja et al., 2025) paradigm replaces heuristic steps with iteration-specific neural networks, delivering high-dynamic-range imaging at CLEAN-like speeds. Semi-supervised frameworks such as VisRec (Wang et al., 2025a) further enhance robustness under limited labelled data. Conditional generative diffusion models (Appendix A.4.4, e.g. Drozdova et al., 2024) achieve superior source localization and flux estimation, while Martínez et al. (2024) demonstrate conditional diffusion models as emulators for realistic synthetic skies, enabling controlled morphology and flux generation for training and benchmarking.

Radio-frequency interference (RFI) mitigation and calibration traditionally major obstacles to interferometric imaging, are typically treated as a detection or segmentation task by deep learning models. RFI mitigation is critical for both *SKA-Mid* and *SKA-Low*, though the challenges differ: *SKA-Mid* contends primarily with terrestrial transmitters and satellite interference across 350 MHz–15 GHz, while *SKA-Low* operates in a crowded low-frequency band (50–350 MHz) where RFI is pervasive and spectral purity is essential for EoR science. For both, U-Nets and transformer-based architectures (Appendix A.3.4) are primarily employed, but more customized hybrid approaches have been explored, combining neural networks with wavelet-based decomposition. Notable examples include convolutional U-Nets for RFI detection (Akeret et al., 2017), residual architectures such as R-Net for MeerKAT (Sadr et al., 2020), and liquid state machines, a type of spiking neural network that processes temporal spike patterns rather than continuous activations (Pritchard et al., 2025), achieving near-perfect per-pixel accuracy. Recent work leverages foundation models like SAM (Deal and Jagannathan, 2024) for automated RFI segmentation without manual intervention.

Cosmological hydrodynamical simulations proxy for theory, and can be used as tests against observations. Since simulations track the evolution of gas density, velocities, and temperature, which telescopes do not observe directly, it is essential to have translations that enable a comparison of simulated and observational quantities (see also Lagos et al., 2026). This is non-trivial and

requires complex models such as conditional generative models for map-to-map translation (see Appendices A.4.6, A.4.7, and A.4.4). Denzel et al. (2026) propose such a solution, primarily relevant for *SKA-Mid* science, which probes low-redshift HI and galaxy structure (see also Wang et al., 2026). The core premise is that the mappings between various simulated quantities of low-redshift galaxies such as gas to neutral hydrogen (HI) brightness or dark matter, while complex, are deterministic enough to be learned by a neural network. The study comparatively investigates the computational efficiency and accuracy of two state-of-the-art paradigms GANs and diffusion models, using physics-aware metrics to assess observational plausibility. By tuning the U-Net architecture, this work reduces computational cost, and bridges the gap between scalable, generative frameworks for forward modelling and observational reconstruction. EMBER (Bernardini et al., 2026, 2025, 2021) exemplifies this approach, accurately predicting key statistics such as mass fractions and profiles, and recovering HI-dark matter correlations across cosmic time down to smaller scales. At higher redshifts, where full hydrodynamical simulations become computationally prohibitive over the cosmological volumes required by SKA-Low, semi-numerical codes such as 21cmFAST (Mesinger et al., 2010) or BEoRN (Schaeffer et al., 2023) provide the tractable simulation backbone for forward modelling. Neural emulators trained on such codes – most notably 21cmEMU (Breitman et al., 2023) – can reproduce the 21-cm power spectrum, global signal, mean neutral hydrogen fraction $x_{\text{HI}}(z)$, and UV luminosity functions across a multi-dimensional astrophysical parameter space with median errors below 1% and a speed-up of roughly 10^4 over the underlying simulator.

Despite these advances, significant challenges remain. Deep learning-based deconvolution is inherently opaque, raising concerns about hallucinated structures and bias propagation from training data. The absence of calibrated uncertainties further complicates scientific validation. Bayesian frameworks such as quantifAI (Liaudat et al., 2023) address these issues by embedding learned priors within convex constraints, ensuring log-concave posteriors and enabling pixel-wise uncertainty estimation. Physics-informed inductive biases and rigorous benchmarking on realistic simulations are essential to mitigate domain gaps and enhance interpretability. Finally, robust MLOps practices (Appendix A.1.6), including data versioning, audit trails, and continuous monitoring, are critical for reproducibility and long-term reliability of AI-driven imaging and simulation pipelines at SKA scale.

2.2 Epoch of Reionization and 21-cm Cosmology

Detecting the redshifted 21-cm signal from the Epoch of Reionization (EoR), the Cosmic Dawn, represents one of the most challenging objectives in observational cosmology. It is a main scientific target of *SKA-Low*, which will observe the 21-cm line across the EoR in its AA4 and AA* configurations (see also Trott et al., 2026). The task involves isolating a faint, highly non-Gaussian cosmological field from spectrally smooth but instrumentally distorted foregrounds that are four to five orders of magnitude brighter, while accounting for contributions such as direction-dependent beams, calibration errors, and thermal noise. Classical approaches, polynomial foreground fitting, principal component analysis, wedge avoidance in $(k_{\perp}, k_{\parallel})$ space, and power-spectrum estimators, prioritize robustness at the expense of sensitivity. These methods often suffer from leakage due to calibration residuals or, conversely, excessive foreground suppression that erases cosmological signal components, introducing systematic biases. Furthermore, most traditional pipelines neglect

field-level inference, forfeiting morphological information critical for constraining astrophysical processes.

Foreground separation beyond linear models has emerged as a key application of AI (see also [Spinelli et al., 2026](#); [Burba et al., 2026](#)). Self-supervised and semi-supervised representation learning enables models such as variational auto-encoders (VAEs), diffusion models, or normalizing flows to learn the manifold of smooth foregrounds jointly with instrumental effects (see Appendices [A.4.3](#), [A.4.4](#), and [A.4.2](#)). These artificial neural network (ANN) architectures can be embedded within Bayesian frameworks to provide uncertainty-aware inference. Conditional generative models (Appendix [A.4.6](#)) further allow projection of visibilities or image cubes onto foreground subspace while preserving cosmological modes, by modelling the conditional distribution $p(\text{foreground} | \text{observation})$ (treating cosmological signals as residuals). [Acharya et al. \(2024a\)](#) for instance, introduce a VAE-trained kernel for Gaussian Process Regression (GPR) foreground mitigation, addressing signal loss due to misestimated covariance. Tested on LOFAR-like simulations, this approach improves recovery of both shape and power of the 21-cm signal. Subsequent work ([Acharya et al., 2024b](#)) applies this methodology to 141 hours of LOFAR data, reporting revised upper limits and reduced correlation with systematic excess noise; similar GPR approaches have been studied recently by [Ceccotti et al. \(2025\)](#); [Mertens et al. \(2025\)](#); [Liu et al. \(2025a\)](#). Deep learning-based foreground mitigation has also been explored using convolutional architectures. [Gao et al. \(2025\)](#) employ a 3D U-Net to extract EoR signals from SKA-Low simulations (AA*) contaminated by foreground residuals and thermal noise, achieving reliable power-spectrum predictions at long integration times. Similarly, [Li et al. \(2019\)](#) demonstrate that convolutional denoising auto-encoders outperform polynomial fitting and wavelet-based methods when beam effects introduce strong spectral fluctuations, recovering EoR signals with correlation coefficients exceeding 0.92. These results underscore the potential of deep learning to disentangle complex foreground structures beyond the assumptions of spectral smoothness.

Field-level inference & HI fraction estimation represents another frontier where AI offers significant advantages (see also [Acharya et al., 2026b,a](#)). Simulation-based inference (SBI; Appendix [A.2.5](#)) can infer astrophysical and cosmological parameters directly from maps or summary statistics, bypassing explicit likelihoods by leveraging amortized neural density estimation (typically through flow-based networks), capturing non-Gaussian information inaccessible to power-spectrum analyses (see Appendices [A.4.2](#) and [A.2.5](#)). [Pietschke et al. \(2025\)](#) demonstrate an application of SBI, EoRFlow, which infers the global neutral hydrogen fraction $\chi_{\text{HI}}(z)$ directly from cylindrically averaged power spectra, enabling piece-wise reconstruction of the reionization timeline; beyond χ_{HI} , SBI approaches can simultaneously constrain astrophysical parameters such as the ionizing emissivity, mean free path of ionizing photons, and X-ray heating efficiency, or parameters more generally related to star formation (cf. [Cerardi et al., 2025](#)). [Schosser et al. \(2025\)](#) extend this paradigm using invertible neural networks linked to physics-inspired latent representations, achieving rapid posterior estimation with well-calibrated summaries. Similarly, [Tiwari et al. \(2022\)](#) show with Bayesian inference accelerated by ANN emulators and SKA-Low-like noise assumptions that all-shape bi-spectra improve reionization-parameter constraints by factors of 2-4 depending on the reionization stage. Complementary approaches such as TPCNet ([Nguyen et al., 2025](#)) combine

convolutional and transformer architectures to infer cold neutral gas fractions and opacity corrections from HI spectra, outperforming CNN baselines in accuracy and robustness. Choudhury et al. (2022) use ANN to predict 21-cm power spectra from foreground-corrupted synthetic observations and then infer reionization parameters from those predicted spectra; similar works have been carried out using CNNs on 2D power spectra by La Plante and Ntampaka (2019); Mangena et al. (2020). Recently, de Salis et al. (2025) performed a comparative study of various neural network architectures for the analysis of 2D power spectra, finding multi-layer perceptron (MLP) mixer and CNN ensembling strategies to provide most flexibility and practicability. For the direct processing on 3D tomographic image cubes, Neutsch et al. (2022) introduce a 3D CNN (3D-21cmPIE-Net) to jointly infer Cosmic Dawn and EoR astrophysics alongside warm-dark-matter properties. Finally, computational costs for likelihood-based inference scale with the number of parameters and novel approaches such as that of Piras et al. (2024) can be applied in the context of EoR. Simulation-based calibration (SBC) provides a complementary diagnostic for any of these inference pipelines: by comparing posterior distributions over inferred parameters against the known simulation inputs, it validates whether posteriors are well-calibrated without requiring tractable likelihoods – a critical quality-control step before applying SBI methods to real SKA-Low observations (e.g., Schosser et al., 2025; Semelin et al., 2025; Prelogović and Mesinger, 2023).

Calibration with inductive biases is essential for mitigating wedge leakage and residual systematics. Physics-informed neural networks (PINNs; Appendix A.3.6) and regularized learning strategies incorporate physical constraints into the training objective, reducing spectral structure in gains and beams. For instance, Tripathi et al. (2025) quantify tolerance levels for calibration and antenna-position errors using ANN-Bayesian frameworks, finding that calibration errors must remain below 0.001% to avoid foreground contamination. Bianco et al. (2025) recently used binary segmentation map priors of ionized HII regions injected into skip connections of a U-Net to recover the 21-cm signal from SKA-Low’s foreground-contaminated observations. Their framework, SERENet, achieves up to 90% accuracy in late-stage signal recovery and over 93% accuracy in cylindrical power spectra. Moreover, Korber et al. (2023) use PINNs for the acceleration of radiative transfer simulations for the forward modelling of cosmic reionization. In these diverse examples, well studied and calibrated inductive biases improved predictions and enhance the efficiency of signal extraction which is needed to recover the faint 21-cm signal.

Uncertainty-aware detection is critical for scientific reliability. Bayesian neural networks (BNNs; Appendix A.3.5) and deep ensembles provide calibrated posterior predictive distributions, enabling robust parameter inference under limited training data. This was demonstrated by Mahida et al. (2025) which showed that BNN emulators outperform deterministic ANN surrogates for 21-cm power spectrum and bi-spectrum inference, yielding tighter constraints on EoR parameters and highlighting the utility of higher-order statistics.

Current-generation instruments such as LOFAR and MWA lack the sensitivity for direct EoR detection; statistical power-spectrum detection is in principle within reach, but remains hindered by calibration systematics and analysis bottlenecks that leave large fractions of acquired data unanalysed (another situation where AI-driven pipelines can provide significant leverage). This necessitates reliance on high-fidelity end-to-end simulations for model pretraining. However, sim-

to-real discrepancies, arising from differences in noise characteristics, calibration artefacts, and unmodelled systematics, remain a major challenge. Domain adaptation techniques, including adversarial training, contrastive representation learning, and fine-tuning on real observations, mitigate these gaps (Baron Perez et al., 2025; Parker et al., 2024; Riggi et al., 2024b; Slijepcevic et al., 2023; Margalef-Bentabol et al., 2020). Additionally, domain randomization during simulation, where beam shapes, gain errors, and foreground morphologies are varied, improves generalization. Combined with uncertainty-aware architectures and physics-informed priors, these strategies enable robust foreground separation and cosmological signal recovery under SKA-scale conditions. Finally, end-to-end instrument simulations remain foundational utilities for assessing SKA-scale performance of AI-based processing pipelines; the community-wide efforts for the SKAO Science Data Challenge 3a (Bonaldi et al., 2025) are exemplary demonstrations that current pipelines can recover EoR signals under realistic conditions, though confidence intervals remain underestimated, motivating cross-pipeline comparisons for bias reduction.

2.3 Radio Galaxies - Detection/Classification

Continuum surveys conducted by *SKA-Mid* (AA4 and beyond) are expected to detect hundreds of millions of radio galaxies spanning diverse redshifts, morphologies, and physical states. These sources often exhibit complex, multi-component structures, such as lobes, jets, bent tails, and diffuse haloes, that challenge traditional source-finding and classification pipelines (see also Hale and Tabatabaei, 2026; Sasmal et al., 2026).

Established automated tools such as PyBDSF (Mohan and Rafferty, 2015) and Aegean (Hancock et al., 2012, 2018) have proven reliable workhorses for detecting and cataloguing compact and moderately extended sources in precursor surveys, performing well on the well-characterized, high-SNR cases that dominate source counts (see also Pal et al., 2026). However, they rely on parametric fitting and fixed morphological assumptions that break down precisely where the science is most demanding: associating disjoint multi-component sources, distinguishing faint diffuse emission from artefacts, classifying rare or ambiguous morphologies, and generalizing robustly across heterogeneous noise and uv-coverage conditions. Manual inspection and rule-based systems are infeasible at SKA scales, particularly in low SNR regimes and for rare or poorly labelled morphologies. Consequently, automated, morphology-aware approaches capable of learning flexible, data-driven representations are essential.

Morphology-aware detection and segmentation has become a central application of deep learning in radio astronomy (see also Prandoni et al., 2026; Hardcastle et al., 2026). Architectures, such as YOLO (Appendix A.3.7), Mask R-CNN (Appendix A.3.9), SAM (Appendix A.3.8), U-Net (Appendix A.3.3), and Vision Transformers (ViTs; Appendix A.3.4) enable pixel-level segmentation, object detection (Appendix A.2.2), and classification of extended structures, outperforming classical algorithms in associating disjoint emission regions and identifying multi-lobed sources.

Riggi et al. (2023, 2024a) present a Mask R-CNN-based source finder for SKA-precursor continuum surveys like ASKAP EMU, reducing false detections in the Galactic plane and improving association of fragmented emission islands; similar studies have been carried out by Burke et al. (2019); Mostert et al. (2022); Zhang et al. (2022). Sortino et al. (2023) provides a comparative

study of object detection models such as Mask R-CNN, Detectron2, DETR, and EfficientDet, finding promising results of above 90% accuracy for radio astronomical images. The SKAO Data Challenge 1 (Bonaldi and Braun, 2018) primarily focused radio source finding, for which Cornu et al. (2024) developed YOLO-CIANNA, a customized deep-learning object detector; the model was later applied in the SDC2, ranking at the top of the scoring board (Cornu et al., 2026). It outperforms previous methods in both detection purity and the number of sources found, while maintaining high computational efficiency and strong characterization accuracy. Similarly, Tolley (2024) investigates Fanaroff–Riley (FR) morphology classification using scattering transforms, showing robustness under noise injection and training-set truncation, and achieving performance competitive with CNN-based classifiers.

Scalable survey analysis and cross-validation requires AI pipelines integrated into SKA workflows for wide-field source detection, classification, and multi-wavelength cross-identification. Transfer learning and contrastive learning approaches (see Appendices A.1.2 and A.1.3) enhance robustness across instruments and surveys. Lastufka et al. (2024) show that self-supervised learning can be trained directly on wide-field MGCLS MeerKAT continuum crops without single-galaxy cutouts, using DINO with ResNet50 and ViT backbones to learn reusable representations for SKA-era surveys. Their SSL models reach state-of-the-art FRI/FRII classification and much better compact-source count prediction, while requiring only $\sim 20k$ source-rich crops and yielding fine-tuned multi-class performance comparable to supervised and natural-image foundation models.

Vision foundation models (VFMs; Appendix A.2.3), pretrained on massive natural-image datasets using self-supervised learning (Appendix A.1.3), offer unprecedented feature extraction capabilities for radio astronomy. Models such as CLIP, SAM, and DINO can be adapted for tasks like FRI/FRII classification, AGN identification, or source count estimation, even in low-label regimes, reducing the need for extensive manual annotation in long-term citizen science projects (see also Hota et al., 2026). Their high-level semantics and low-level patterns captured during training with natural images provide a starting point in the radio data domain, reducing the computational effort (and data volumes) to learn those aspects from scratch. Lastufka et al. (2025) evaluates how vision foundation models can be adapted for optical and radio astronomy tasks, finding that these models improve classification accuracy for optical galaxy images and perform well in radio object detection, but struggle with radio galaxy classification. Going even a step further than VFMs, Riggi et al. (2025) study fine-tuning of a vision-language model (LLaVA) on a large set of radio images and captions. They achieve about 30% F1-score gains in extended source detection compared to base models, but still under-performing vision-only classifiers and showing a 20% drop on general multimodal tasks. This work highlights both the promise and current limitations of such models for radio astronomy, including the need for better multimodal alignment and higher-quality datasets. Similarly, Drozdova et al. (2025) evaluate various vision–language models for classifying radio galaxy morphologies, finding that prompt-based approaches can perform well even without astronomy-specific training, but their outputs are highly sensitive to prompt variations.

Data augmentation using generative models, GANs, VAEs, and diffusion models (see Appendices A.4.1, A.4.3, and A.4.4), further improves classifier generalization under limited data conditions. Synthetic radio galaxy images generated by these models replicate realistic morphologies,

enabling robust training for rare classes and mitigating class imbalance.

Class-imbalance is a common problem among galaxy detection and classification tasks and can lead to biases towards the well-represented classes. Appropriate sampling strategies, synthetic augmentation, and anomaly-aware training can alleviate such problems. Due to the inherently high data compression/reduction of tasks such as physical parameter inference or morphological classification (from images to a few scalars or tokens), it is difficult to retrace single decisions of models and retain scientific trust. Thus, it is important to validate results using interpretability techniques (Appendix A.1.7), such as feature attribution (through saliency maps via Grad-CAM or SmoothGrad), prototype-based explanations, or counterfactual analysis. Simulation-based calibration can validate whether posterior distributions produced by probabilistic inference methods are well-calibrated through comparison against simulated parameters. This provides a diagnostic tool without requiring ground-truth likelihoods and ultimately increases trustworthiness of AI outputs. Finally, the adoption of MLOps best practices prevents the occurrence of pain points (operational failure modes such as silent model drift, training-serving skew, and data version mismatch), contributes to the reproducibility of models, and provides means to monitor model drift and performance after deployment.

2.4 Pulsar & Transient Science

The SKA is set to revolutionize time-domain radio astronomy by enabling the discovery of fast radio bursts (FRBs), exotic pulsar systems, and other transient phenomena across a vast range of timescales and astrophysical regimes (see also Burgay et al., 2026; Miller-Jones et al., 2026; Keane et al., 2026; Shannon et al., 2026). They are primary science drivers for *SKA-Mid* (AA4), which provides the sensitivity, frequency coverage, and timing precision required for pulsar timing arrays, FRB localization, and pulsar discovery. *SKA-Low* contributes a more limited but complementary role, enabling searches for long-period pulsars and low-frequency transient phenomena in a regime inaccessible to *SKA-Mid* (see also Qiu et al., 2026). In general, the scale and cadence of SKA data, spanning petabyte-scale archives and real-time beam-formed streams, render traditional manual and heuristic approaches infeasible. Classical pipelines struggle with RFI, rare-event imbalance, and latency constraints, particularly in the context of real-time follow-up and precision timing. To fully exploit SKA’s sensitivity and temporal resolution, automated, adaptive, and scalable analysis frameworks are required.

Candidate triage and prioritization is a critical bottleneck in transient discovery. Deep learning models, including CNNs and transformers (see Appendices A.3.1 and A.3.4), can automate the ranking and filtering of transient candidates from high-volume data streams, reducing expert workload and enabling dynamic prioritization of events based on scientific relevance (cf. Sravan et al., 2022; Agarwal et al., 2020; Connor and van Leeuwen, 2018). These systems adapt to changing noise environments, and optimize telescope time allocation for follow-up resources (see also Anderson et al., 2026).

Anomaly detection for rare events leverages unsupervised and semi-supervised learning to identify low-likelihood signals without laboriously labelled datasets (see also Andersson et al., 2026b,a). Techniques such as one-class classifiers, VAEs, and deep generative models (see Appendices A.4.3

and A.2.4) have shown promise in detecting novel transients (e.g., Kuiper et al., 2025). Active learning strategies further refine these models by incorporating minimal human feedback. Andersson et al. (2025) benchmark anomaly-detection algorithms on MeerKAT light curves, demonstrating that ranking by anomaly score recovers over half of known transients within the top 10% of candidates, while active learning improves recall by up to 20 percentage points when supplemented with 2% labelled data.

Real-time pulsar search with edge processing is indispensable for SKA-scale operations. Lightweight CNN-based architectures (Appendix A.3.1) can be directly deployed on telescope back-ends or FPGA-based edge devices, enabling on-the-fly analysis during acquisition. This approach mitigates the risk of losing weak or short-lived signals during aggressive data compression and circumvents the impracticality of archival reprocessing. Fu et al. (2025) present a deep-learning–assisted pulsar-search pipeline developed and validated on current-generation telescopes (FAST, GBT, Arecibo, and MWA) that ranks folded “snapshot” candidates from de-dispersed time series, so the costly *folding stage*, in which candidate periods are used to coherently fold thousands of time-series samples into a pulse profile for visual or automated inspection, is invoked far less often in next-generation surveys. The pipeline delivers ~10–60 fold speed-ups while retaining nearly 100% recall, demonstrating strong cross-instrument generalization and direct methodological relevance for SKA pulsar surveys.

Precision enhancement and noise reduction represents another domain where AI excels, and the techniques discussed here are broadly applicable across the scientific areas covered in this chapter. Conditional generative models, such as GANs, diffusion models, and flow-matching architectures (Appendix A.4.7, A.4.4, or A.4.8), can denoise and super-resolve pulse profiles, improving timing accuracy for applications like pulsar timing arrays in gravitational wave detection (Li et al., 2021). These methods preserve temporal resolution while suppressing noise and instrumental artefacts. For instance, Zhang et al. (2024) reframe RFI mitigation in pulsar observations (for LOFAR/NenuFAR) as a joint detection-and-restoration problem on the full time-frequency dynamic spectrum.

Distributed learning frameworks, particularly federated learning (FL; Appendix A.1.5), enable collaborative model training across observatories without centralizing raw data, reducing bandwidth demands and enhancing privacy. This paradigm is particularly relevant for global transient detection and large-scale pulsar surveys. Unfortunately, there is a significant research gap in the application: while federated model training has remained entirely unexplored in the context of exascale radio interferometry, works by Scaife (2020) and Wang et al. (2025b) share some architectural similarities with FL.

Autonomous decision-making and scheduling can be optimized using reinforcement learning (RL; Appendix A.1.4). RL agents learn policies that balance scientific yield with operational constraints, dynamically allocating resources for follow-up observations and survey planning. Recent work also demonstrates the use of RL to automate calibration workflows in radio interferometry for current-generation instruments (cf. Kirk et al., 2024), with the explicit goal of developing transferable approaches for SK-scale autonomous operations. Such calibration agents can generalize across diverse observations, minimizing the need for manual intervention and enabling “smart calibration”

of radio telescopes (Yatawatta, 2023); so far, the SKAO has not considered autonomous scheduling and adaptive calibration as key operational requirements.

Despite these advances, AI-driven systems must be designed with safeguards to ensure scientific reliability, especially those with a high degree of autonomy. Continuous monitoring and retraining strategies are essential to avoid model drift and maintain performance under evolving noise and signal conditions. Diverse training datasets (in particular when data is synthetic) help prevent bias toward known classes, while anomaly-aware architectures improve robustness. Moreover, explainable AI techniques (see Appendix A.1.7), such as saliency maps, integrated gradients, and counterfactual analysis, enhance transparency and increase trust in results during human validation (cf. Denzel et al., 2024). Furthermore, MLOps practices (Appendix A.1.6) aid in seamless integration of the aforementioned explainable AI techniques, and monitoring models after deployment (Billeter et al., 2024).

2.5 Cosmic Magnetism

Understanding cosmic magnetism is central to explaining galaxy evolution, star formation, and cosmic ray transport (cf. Carretti and Vazza, 2025). Observables such as Zeeman splitting and Faraday rotation encode magnetic field strength and topology, yet these signals are intrinsically faint and easily masked by instrumental artefacts and foreground contamination (see also Vernstrom et al., 2026; Robishaw et al., 2026; Bourke et al., 2026). SKA-AA* will deliver an unprecedented number of Zeeman measurements in star-forming regions, and provide invaluable insights into the strength and structure of magnetic fields in the interstellar medium (ISM), however, disentangling these faint signals from Stokes V spectra remains challenging (cf. Pattle et al., 2022); AA* is the only configuration with enough sensitivity to detect magnetic field strengths of the order of a few μG (Heald et al., 2020). Similarly, extracting extra-galactic Faraday rotation signals from rotation measure (RM) catalogues requires separating Galactic contributions and measurement uncertainties, a task that classical approaches struggle to achieve at SKA scales (cf. Takahashi, 2023).

Radio polarimetry is the first step in the study of cosmic magnetism. High-fidelity imaging in the Stokes Q and U polarization parameters are essential components in the study of cosmic magnetism. Wang et al. (2023) address the fundamental data representation in deep learning models for this first step in the data processing pipeline, developing and validating their approach on current interferometric datasets with direct applicability to *SKA-Mid* polarimetric imaging. While traditional CNNs operate on Cartesian grids, their PolarRec framework operates in polar coordinates along elliptical tracks in the uv-plane, and demonstrates superior performance using transformer encoders, reducing spurious magnetic signal artefacts.

Zeeman line profile recovery can benefit substantially from machine learning. Models trained on synthetic and real Stokes V spectra learn non-linear mappings that isolate weak Zeeman signatures from instrumental distortions, enabling robust magnetic field estimation in molecular clouds. Diffusion models, in particular, offer flexibility for recovering magnetic field strength from multiple observables, including column density, dust polarization angles, and line-of-sight velocity dispersions (Appendix A.4.4). Xu et al. (2025) use such diffusion models trained on synthetic

ISM observables to estimate magnetic field strength detectable through SKA-Mid AA*. With a 3-channel diffusion model (using column density, polarization angles, and velocity dispersion) they outperform both traditional power-law fitting and Davis–Chandrasekhar–Fermi (DCF) methods, especially on previously unseen data. The classical DCF method overestimated field strength by about an order of magnitude, while the 3-channel diffusion demonstrated superior precision and robustness.

Faraday rotation disentanglement is another critical application where AI can optimize bottlenecks. Extracting detailed RM structures from radio data under complex Faraday dispersion and noise conditions can be formulated as an inverse problem solvable by generative models (see also [Carcamo et al., 2026](#); [O’Sullivan et al., 2026](#); [Loi et al., 2026](#)). Models consisting of simple U-Nets (Appendix A.3.3), conditional GANs (Appendix A.4.7), or diffusion models (Appendix A.4.4) have demonstrated superior performance compared to traditional algorithms like RMCLEAN ([Gustafsson et al., 2024](#)), with studies conducted primarily on MeerKAT data as a direct SKA-Mid precursor (e.g. [Ndiritu et al., 2021](#); [Carcamo et al., 2022](#)). [Gustafsson et al. \(2024\)](#) present a semi-supervised deep-learning model for Faraday spectrum deconvolution, achieving accurate recovery of dispersion functions across a broad RM range while scaling efficiently to large MeerKAT datasets. This approach reveals intricate magnetic structures in cluster relics and AGN, highlighting the potential of AI for high-resolution magnetism studies.

3D mapping of magnetic fields represents a frontier where AI-driven methods outperform classical techniques (see also [Tahani et al., 2026](#); [Sun et al., 2026](#); [Ma et al., 2026](#)). Conditional generative models (Appendix A.4.6) trained on magneto-hydrodynamic simulations can predict 3D field orientation, strength, and turbulence parameters from complex, noisy observational proxies such as HI emission, synchrotron maps. These models surpass velocity gradient and equi-partition methods in accuracy and robustness, even under low SNR conditions. [Hu and Lazarian \(2024, 2025\)](#) demonstrate this by training a CNNs and ViTs on synthetic data to robustly extract the plane-of-the-sky orientation, inclination angle, and magnetization from synchrotron observations. Similarly, [Hu \(2025\)](#) present a deep learning framework based on conditional ResNets to predict the full 3D magnetic field structure and turbulence properties in HI clouds from spectroscopic observations. Moreover, [Denzel et al. \(2026\)](#) demonstrates magnetic field strength recovery from multiple other quantities such as gas column density with conditional GANs and diffusion models trained on magneto-hydrodynamics simulations. While their reconstructions are technically only 2D, their approach is straight-forward to adapt to full 3D reconstruction of magnetic fields within galaxies.

Similar to many other scientific fields, counting of events with time-independent probability follow a Poisson distribution, which means the noise variance is proportional to square-root of the signal and can generally be reduced with more acquisition time. When longer integration times are impractical or signals are still intrinsically close to the SNR level, it has been demonstrated that model training on samples artificially enhanced with simple noise simulations can improve real-world performance (cf. [Oppliger et al., 2024](#)). While CNNs (Appendix A.3.1) and ViT (Appendix A.3.4) maintain high accuracy even with high noise and missing data, making them suitable for low signal-to-noise ratio (SNR) observations, their predictions are not immediately transparent. [Peek and Burkhart \(2019\)](#)

investigate the interpretability of such models, showing that learned features can be meaningfully linked to physical properties such as turbulence anisotropy and morphological curvature, thereby partially alleviating black-box concerns.

2.6 Gravitational Lensing

Gravitational lensing, particularly in the strong regime, is the key probe of dark matter and its distribution on small scales. Low-mass dark matter haloes are of special interest because their population mass function allows direct testing of competing dark matter models (cold versus warm dark matter). Lensed radio sources provide constraints on dark matter properties through magnification and differential time delay measurements; *SKA-Mid* (AA4) is poised to discover thousands of such systems (cf. McKean et al., 2015, and references therein), rendering manual identification and modelling infeasible at survey scale.

Unlike optical observations, radio lensing analysis faces unique challenges: the morphological characteristics of radio sources are harder to distinguish from lensing features due to typically lower angular resolution and SNR, and radio emission traces gas rather than stellar continuum, complicating direct transfer of optical pipelines. Consequently, AI-based work on radio gravitational lensing remains scarce compared to the optical domain, partly due to the complexity of interferometric visibility-space data and the comparatively greater availability of wide-field optical imaging.

Nevertheless, the AI architectures and inference frameworks developed extensively in the optical regime provide methodological templates for radio application, and we review both here, emphasizing the adaptations needed for interferometric data. Beyond methodology transfer, the combination of radio and optical observations of the same lensing systems offers qualitatively stronger science: optical data constrain the lens and source light distributions, while radio interferometric data provide independent astrometric constraints, probe gas-phase source morphology invisible in optical bands, and are free of dust obscuration – together breaking degeneracies such as the mass-sheet degeneracy that neither waveband can resolve alone. This argument for cross-spectrum synergy is not unique to gravitational lensing (although more apparent due to the data scarceness); it applies broadly across SKA science, where joint radio–optical–infrared analyses can disentangle physical components, validate detections, and tighten parameter constraints in ways that no single waveband can achieve independently (cf. Prandoni and Seymour, 2015).

Automated lens finding and modelling are both well developed in the optical regime and provide the template for radio application. For *identification*, CNNs substantially outperform traditional methods (see Appendix A.3.1), with single-stage object detectors such as YOLO offering the throughput required at survey scale (e.g. Liu et al., 2025b), and unsupervised/anomaly-detection approaches (see Appendix A.1.2 and A.2.4) – notably auto-encoders – exploited for the rare-event regime (e.g. Cheng et al., 2020; Etsebeth et al., 2024). The one explicitly radio-interferometric demonstration to date is Rezaei et al. (2022), who tailored CNN-based lens finding to LOFAR images and recovered 95.3% of true positives at a 0.008% false-positive rate, implying that most galaxy-scale lenses in LoTSS should be detectable. For *modelling*, Bayesian neural networks and simulation-based inference stochastically infer parametric lens-model parameters (Appendix A.2.5; e.g. Legin et al., 2021), while parameter-free reconstructions using Recurrent Inference Machines

(RIM) properly explore the degenerate solution space (Morningstar et al., 2018, 2019); through explicit integration of gradient log-likelihoods, probabilistic methods such as RIM or diffusion models can enforce physical priors and observational constraints to suppress implausible solutions (cf. Remy et al., 2023; Rüstig et al., 2024). Crucially for SKA, differentiable lens-modelling frameworks such as PyAutoLens (Nightingale et al., 2018) are already compatible with radio interferometric and visibility data and can serve as a template for radio-adapted, AI-based modelling.

Due to the inherently ill-posed nature of the gravitational lensing problem the identification and modelling of gravitational lensing require precise physical constraints and should be interpretable to assess model plausibility. This is in direct conflict with the “black box” nature of deep learning. The rarity of the phenomenon also leads to another critical risk of supplementing or replacing actual observational datasets with simulations and therefore widening the sim-to-real domain gap for neural networks. Furthermore, morphological mimics pose a significant challenge; intrinsic radio structures such as bent-tail galaxies or chance alignments of star-forming clumps can closely resemble lensing arcs, potentially contaminating catalogues in the absence of redshift information. There is also the risk of training set bias, where supervised models favour systems with large Einstein radii or high SNR ratios, effectively blinding the survey to lower-mass lenses or compact configurations that are crucial for dark matter studies.

These risks can be mitigated using a variety of strategies: physics-informed architectures that explicitly incorporate the radio interferometric measurement equation in to the loss function or architecture (such as in RIM) ensure that any reconstructed lens model adheres to mathematical consistency with sampled visibilities, thereby reducing artefacts. Optical-to-radio domain and sim-to-real gaps can be mitigated using unsupervised domain adaptation techniques (Swierc et al., 2024). They allow networks to learn feature representations that are invariant to the differences between simulations and real data and have been shown to improve inference accuracy by orders of magnitude when applied to noise properties that differ from the training set. To address computational bottlenecks and one-shot reliability issues when dealing with large datasets, a hierarchical inference pipeline could be considered: in a first phase smaller, fast networks such as CNNs are used for initial candidate discovery and artefact rejection (as in Hezaveh et al., 2017), calibrated for high coverage and low false-positive rate, followed by a more rigorous, physics-informed approach such as SBI for final parameter estimation of high-probability candidates. Finally, considering the comparably large number of lenses already detected with other instruments such as Euclid or JWST, a multi-wavelength approach provides validation opportunities and could provide even stronger constraints for lens models to break/alleviate the mass-sheet degeneracy. Moreover, since radio emission and optical emission trace different physical components of galaxies, a true lens candidate should exhibit consistent astrometry across bands, while mimics more often do not. Many deep learning frameworks provide the flexibility to integrate multiple data domains for lens candidate analysis.

3 Conclusion

The Square Kilometre Array Observatory (SKAO) represents a paradigm shift in radio astronomy, not only in terms of sensitivity and survey speed but also in the scale and complexity of its data ecosystem. This chapter has outlined how Artificial Intelligence (AI) is uniquely positioned




Figure 1: A semantic cross correlation heatmap, qualitatively illustrating the relative importance of AI-based methods in SKA problem areas. Note that this figure is derived from the references discussed in this chapter and reflects an indicative rather than exhaustive selection of topics. However, the frequency of each AI-based method is representative when compared against matched papers from arXiv.

to address these challenges, enabling real-time operations, accelerating scientific workflows, and unlocking new discovery modes across cosmology, galaxy evolution, and time-domain astrophysics. From deep learning–based interferometric imaging and generative sky simulations to uncertainty-aware inference and physics-informed architectures, AI has demonstrated transformative potential in domains where classical methods fail under SKA-scale conditions (cf. Figure 1).

Yet, the integration of AI into SKA science is not without caveats. The opacity of deep learning models, risks of hallucinated structures, and domain gaps between simulations and observations underscore the need for rigorous validation and interpretability. Embedding physical constraints, adopting Bayesian frameworks for uncertainty quantification, and implementing robust MLOps practices – including data versioning, model registries, and continuous performance monitoring across the distributed SKA data infrastructure – will be essential to maintain scientific integrity at AA4 scale and beyond. These safeguards must evolve alongside increasingly autonomous systems, ensuring that AI remains a tool for discovery rather than a source of bias or error.

Looking ahead, several research frontiers merit attention. First, scalable simulation-based inference (Cranmer et al., 2020; Cerardi et al., 2025) will become indispensable for extracting non-Gaussian information from SKA datasets, bridging theory and observation through amortized neural density estimation. Second, foundation models and self-supervised learning (Parker et al., 2024; Lastufka et al., 2025) promise to reduce reliance on labelled data, enabling cross-survey generalization and multimodal integration with optical, infrared, and gravitational-wave observations. Third, distributed and federated learning (Kairouz et al., 2021) will support collaborative science across global SKA nodes without centralizing raw data, addressing bandwidth constraints. It should be noted, however, that the SKAO data management framework is currently built around a network of centralized SKA Regional Centres (SRCs), and federated learning is not part of the SKAO operations

road map at this stage (SKA Observatory, 2025; SRCNet, 2025). The relevance of federated approaches will depend on how data access policies and SRC governance evolve, particularly with respect to proprietary data periods and the eventual opening of archival datasets to the broader community. Finally, reinforcement learning and adaptive scheduling (Kirk et al., 2024; Yatawatta, 2023) could optimize telescope operations in real time, balancing scientific yield against resource limitations.

The SKA era will not merely scale existing workflows; it will redefine the scientific method in radio astronomy. AI has already contributed to advances across the board in SKAO science. From the acceleration of forward models, over the automation of real-time analyses, to the dynamic compression of unprecedented data volumes, AI provides flexible and efficient solutions that lighten the load on computational pipelines and field experts alike. The evidence suggests that AI is no longer merely an optimization tool but a prerequisite for unlocking the physics of the Universe in the radio spectrum, offering necessary breakthroughs in scalability, fidelity, and the serendipitous discovery of “unknown unknowns”.

A AI Methods Glossary

The following provides a concise reference summary of the main AI concepts and methods cited in Section 2, with a focus on learning paradigms, common tasks and learning objectives, neural network architectures, as well as deep generative models. Where applicable, the SKA-relevant use case for each method is noted; of course, this application area relevance is non-exclusive.

A.1 Learning Paradigms and Development Patterns

A.1.1 Deep Learning

Deep learning is a sub-field of machine learning that trains multi-layer artificial neural networks to learn hierarchical representations of data, enabling end-to-end learning of complex tasks such as image recognition, speech recognition, machine translation, and scientific discovery (LeCun et al., 2015; Goodfellow et al., 2016). It relies on large datasets, gradient-based optimization (back-propagation), and architectural innovations (convolutional, recurrent, transformer-based networks) to automatically extract features at multiple levels of abstraction and has driven many recent advances in AI.

SKA relevance: the foundational paradigm underlying all deep learning methods applied across the scientific areas in Section 2.

A.1.2 Supervised and unsupervised Learning

Supervised learning trains models on labelled examples (input–output pairs) so they learn to map inputs to known targets, optimizing a loss that measures prediction error; it excels at tasks like classification and regression when annotated data are available (e.g., image labels, paired translations). Unsupervised learning discovers structure from unlabelled data by modelling data distributions, finding clusters, or learning useful representations (e.g., clustering, density estimation, auto-encoders), and is useful when labels are scarce. Semi-supervised and self-supervised methods

bridge the two by leveraging large unlabelled datasets plus limited labels to improve performance (Mitchell, 1997).

SKA relevance: supervised learning underpins source detection and classification pipelines; unsupervised and semi-supervised approaches are critical where labelled radio data is scarce, as is common for rare morphologies (Section 2.3) and EoR signal extraction (Section 2.2).

A.1.3 Self-supervised learning (SSL)

Self-supervised learning trains models using automatically generated supervision from the data itself by defining pretext tasks (e.g., predicting masked inputs, contrasting augmented views, or solving context/order prediction) so models learn rich representations without manual labels; these learned features can be fine-tuned or transferred to downstream tasks and have driven major advances in vision, language, and multimodal learning (e.g., masked language modelling, contrastive learning, and masked image modelling).

SKA relevance: SSL enables learning from the vast unlabelled datasets that SKA-era surveys will produce, reducing annotation bottlenecks for radio galaxy classification and continuum surveys (Section 2.3).

A.1.4 Reinforcement learning (RL)

Reinforcement learning (RL) studies how an agent learns to make sequential decisions by interacting with an environment to maximize cumulative reward: the agent observes states, takes actions according to a policy, and updates that policy using feedback (rewards and state transitions) to improve long-term performance; RL spans model-free methods (Q-learning, policy gradients, actor-critic), model-based methods, and modern deep RL that combines neural function approximators with classic algorithms to solve high-dimensional control, games, and robotics problems (Sutton et al., 1998; Mnih et al., 2015; Murphy, 2025).

SKA relevance: RL is being explored for autonomous telescope scheduling, calibration workflow optimisation, and adaptive follow-up for transient events (Section 2.4).

A.1.5 Federated learning

Federated learning is a distributed machine-learning paradigm where model training is performed across many client devices (e.g., mobile phones, edge sensors) that keep raw data locally and only send model updates (gradients or weights) to a coordinating server for aggregation, preserving data locality and reducing the need to centralize sensitive datasets (McMahan et al., 2023; Kairouz et al., 2021).

SKA relevance: federated learning is a candidate paradigm for training shared models across distributed SKA data centres in South Africa and Australia without centralizing petabyte-scale raw data. So far, this is not on the SKAO operations road-map.

A.1.6 Machine Learning Operations (MLOps)

MLOps (Machine Learning Operations) is the set of practices, tools, and cultural patterns that combine software engineering, DevOps, and data science to reliably deploy, monitor, and maintain

machine learning systems in production; it covers the full model life-cycle—data versioning and pipelines, reproducible training, model validation, CI/CD for models, scalable serving, monitoring for data and model drift, and governance (including reproducibility, explainability, and compliance)—with aims to reduce time-to-production, improve reliability, and enable continuous improvement of ML services (Sculley et al., 2015; Breck et al., 2017; Huyen, 2022).

SKA relevance: MLOps frameworks are essential for deploying, monitoring, and maintaining AI pipelines across the full SKAO data life-cycle, from real-time telescope back-ends to science-archive post-processing.

A.1.7 Explainable/interpretable AI/ML

Interpretable and explainable AI/ML (Guidotti et al., 2018; Molnar, 2025) concerns methods and practices that make model behaviour, predictions, and learned representations understandable to humans—either by using inherently interpretable models (rules, generalized additive models, sparse linear models) or by applying post-hoc explanation techniques (feature attributions, saliency maps, counterfactuals, surrogate models)—so stakeholders can trust, validate, debug, and meet regulatory or ethical requirements. In practice, explainable and interpretable AI are often used interchangeably (we follow this practice throughout this chapter); however, there is a noteworthy distinction. *Interpretable AI* refers to models that are inherently transparent by design – you can inspect their structure directly to understand why a prediction was made. Simple examples include linear regression, decision trees, or sparse rule-based systems: the model itself is the explanation. *Explainable AI* refers to post-hoc techniques applied to models that are not inherently transparent.

SKA relevance: explainability techniques such as saliency maps and counterfactual analysis are used to validate source detection, morphology classification (Section 2.3), and EoR inference models (Section 2.2), building scientific trust in AI outputs.

A.2 Tasks and learning objectives

A.2.1 Image Classification

Image classification is a fundamental task in computer vision where an AI model, typically a Convolutional Neural Network (CNN), analyses an entire image’s pixel data and assigns it a single, predefined categorical label based on its visual content (Toennies, 2024). Unlike object detection, which localizes multiple objects, image classification provides a high-level summary of the image’s primary subject, a capability often benchmarked using large-scale datasets like ImageNet (Deng et al., 2009).

SKA relevance: applied to radio galaxy morphology classification (FRI/FRII and beyond) and pulsar candidate scoring in SKA-Mid and SKA-Low surveys (Section 2.3).

A.2.2 Object Detection & Image Segmentation

Object detection is a fundamental task within computer vision that enables simultaneous identification as well as localization (by means estimating a bounding box) of objects within images. By contrast, image segmentation is a more granular and sophisticated task than object detection, as it involves classifying every single pixel in an image to belong to a specific object class or background.

This process effectively creates precise boundaries or masks for objects, providing a pixel-level understanding of the scene. Popular architectures include Mask R-CNN (Appendix A.3.9), YOLO (Appendix A.3.7) or SAM (Appendix A.3.8).

SKA relevance: core task for radio source finding and multi-component association in SKA-Mid continuum surveys; also used for RFI segmentation across both telescopes (Sections 2.1 and 2.3).

A.2.3 Vision foundation models (VFMs)

Vision foundation models are large, pretrained visual models learned from massive, diverse image (and often image–text) datasets using self-supervised or weakly supervised objectives; they produce versatile, transferable representations that enable a wide range of downstream tasks (classification, detection, segmentation, image captioning, and multimodal reasoning) via fine-tuning or prompting, and have spurred architectures combining vision transformers, contrastive and masked modelling, and multimodal alignment, e.g., CLIP (Radford et al., 2021) or DINO (Caron et al., 2021).

SKA relevance: VFMs pretrained on optical data are being fine-tuned for radio source classification and detection (Section 2.3), offering strong baselines in low-label regimes relevant for SKA-era citizen science and automated pipelines.

A.2.4 Anomaly Detection

Anomaly detection identifies patterns in data that deviate from expected behaviour (outliers, novelties, or rare events) to flag errors, fraud, faults, or unusual phenomena; methods range from simple statistical thresholds and distance- or density-based techniques (e.g., k-NN, isolation forest, local outlier factor) to reconstruction- or representation-based approaches using auto-encoders, one-class support vector machines, and deep generative models (variational auto-encoders, GANs, normalizing flows), as well as time-series-specific methods and hybrid/ensemble systems that emphasize robustness, interpretability, and calibrated scoring for real-world monitoring and alerting (Chandola et al., 2009; Pang et al., 2021).

SKA relevance: used for transient and FRB discovery in SKA-Low and SKA-Mid data streams (Section 2.4), and for identifying rare lensing candidates (Section 2.6) and unusual source morphologies in continuum surveys (Section 2.3).

A.2.5 Simulation-based inference (SBI)

Simulation-based inference (also called likelihood-free inference) refers to methods for performing statistical inference when the likelihood function is intractable or unavailable but one can forward-simulate data from a generative model; by comparing simulated and observed data via summary statistics, discrepancy measures, or learned neural surrogates, these methods estimate posterior distributions, model parameters, or perform model comparison. Approaches include Approximate Bayesian Computation (ABC), which accepts parameter draws that produce simulations close to observations; synthetic likelihood and density-ratio or conditional density estimation methods that learn likelihood or posterior surrogates; and modern neural approaches (neural posterior estimation, neural likelihood estimation, and neural ratio estimation) that scale simulation-based inference to

high-dimensional data using flexible neural networks and amortized inference (Beaumont, 2019; Cranmer et al., 2020).

SKA relevance: SBI enables parameter inference for EoR cosmology (Section 2.2), gravitational lensing (Section 2.6), and pulsar physics (Section 2.4) directly from SKA observables without requiring tractable likelihoods – essential when forward models are complex simulations.

A.3 Neural network architectures

A.3.1 Convolutional Neural Networks (CNNs)

Convolutional neural networks (CNNs) (Lecun et al., 1998; Krizhevsky et al., 2017) are a class of deep learning models designed to process grid-structured data (especially images) by using convolutional layers that apply learned localized filters across spatial dimensions to extract translation-equivariant features; stacking convolutions, non-linearities, pooling, and normalization produces hierarchical representations that capture edges, textures, parts, and high-level concepts, enabling state-of-the-art performance in image classification, object detection, segmentation, and related tasks.

SKA relevance: the most widely used architecture across all six scientific areas reviewed, from RFI detection and source finding to EoR signal extraction.

A.3.2 Residual Network (ResNet)

A Residual Network (ResNet) is a deep convolutional architecture that uses identity-skip (residual) connections to allow training of substantially deeper networks by reformulating layers as learning residual functions (He et al., 2016a,b). Each block learns the difference between the desired mapping and the input, thereby mitigating vanishing gradients and enabling easier optimization. ResNets achieved state-of-the-art performance on image classification and became a foundational backbone for many vision tasks and architectures.

SKA relevance: used as backbone architecture in many models such as POLISH for interferometric image super-resolution (Section 2.1) and in 3D magnetic field reconstruction from spectroscopic SKA-Mid observations (Section 2.5).

A.3.3 U-Net

U-Net is a convolutional encoder–decoder architecture with symmetric contracting and expanding paths connected by skip connections, originally developed for biomedical image segmentation; the encoder progressively captures context via down-sampling convolutions while the decoder restores spatial resolution using up-convolutions, and the lateral skip connections transfer fine-grained localization information from corresponding encoder layers to the decoder, enabling precise pixel-wise predictions from relatively few training images (Ronneberger et al., 2015).

SKA relevance: the workhorse architecture for interferometric deconvolution, RFI segmentation (Section 2.1), EoR foreground separation (Section 2.2), and many deep generative modelling approaches (Appendix A.4).

A.3.4 Transformers

Transformers are a neural network architecture that uses self-attention mechanisms to model relationships between all elements of a sequence in parallel, replacing recurrence and enabling efficient scaling to very long contexts (Vaswani et al., 2017; Devlin et al., 2019; Brown et al., 2020). By computing attention-weighted combinations of token representations, transformers learn contextualized embeddings that power state-of-the-art results in machine translation, language modelling, and many other domains. For the image modality, Vision Transformers (ViT) (Dosovitskiy et al., 2020) have been developed that tokenize images in patches. The model then uses a self-attention mechanism on these patches to capture global relationships across the entire image, providing a powerful alternative to traditional Convolutional Neural Networks (CNNs).

SKA relevance: most widely used architecture when data is sufficient, e.g., for uncertainty-aware EoR parameter inference (Section 2.2), radio source classification and detection (Section 2.3), and gravitational lens modelling (Section 2.6), where calibrated posteriors are essential for scientific validity.

A.3.5 Bayesian Neural Network (BNN)

Bayesian neural networks (BNNs) (MacKay, 1992; Neal, 2012) place probability distributions over a network’s weights (or outputs) to capture epistemic uncertainty and enable principled Bayesian inference for prediction, calibration, and decision-making; by combining prior beliefs with likelihoods from data, BNNs produce posterior distributions that quantify uncertainty about model parameters and predictions, which is useful for robust learning, active learning, and safety-critical applications. Exact Bayesian inference is intractable for large networks, so practical methods use approximations such as variational inference, Monte Carlo dropout (Gal and Ghahramani, 2016), deep ensembles, and Markov Chain Monte Carlo (MCMC) (Robert and Casella, 2011) to approximate posteriors and predictive distributions.

SKA relevance: used for uncertainty-aware EoR parameter inference (Section 2.2) and gravitational lens modelling (Section 2.6), where calibrated posteriors are essential for scientific validity.

A.3.6 Physics-informed neural networks (PINNs)

Physics-informed neural networks (PINNs) incorporate known physical laws—typically expressed as differential equations, boundary/initial conditions, and conservation constraints—directly into the training loss so neural networks approximate solutions that satisfy both data and governing equations (Raissi et al., 2017a,b, 2019; Karniadakis et al., 2021). By penalizing residuals of the PDE (and related constraints) at collocation points, PINNs enable mesh-free solution of forward and inverse problems, parameter estimation, and surrogate modelling for complex physical systems while often reducing required labelled data and improving generalization.

SKA relevance: applied to EoR calibration error quantification (Section 2.2) and radiative transfer emulation for SKA-Low forward modelling (Section 2.1).

A.3.7 You Only Look Once (YOLO)

The You Only Look Once (YOLO) family of models revolutionized object detection by formulating it as a single-stage regression problem, enabling simultaneous prediction of bounding boxes and class probabilities with a single pass through the neural network, which drastically increased processing speed (Redmon et al., 2016).

SKA relevance: adapted for radio source detection (YOLO-CIANNA) and optical gravitational lens searches with direct applicability to SKA-Mid continuum surveys (Section 2.3).

A.3.8 Segment Anything (SAM)

Segment Anything (SAM), a foundational model for computer vision, introduced the concept of promptable segmentation, allowing it to generate high-quality segmentation masks for any object in an image based on simple inputs like clicks, bounding boxes, or text (Kirillov et al., 2023). Trained on a massive dataset of over one billion masks, SAM exhibits impressive zero-shot generalization capabilities, enabling it to accurately segment objects and novel scenes it has never encountered during training.

SKA relevance: applied to automated RFI segmentation in interferometric data for both SKA-Mid and SKA-Low (Section 2.1).

A.3.9 Mask R-CNN

Mask R-CNN (He et al., 2018) is a deep learning model used for object detection and instance segmentation in images. It extends the Faster R-CNN architecture (Ren et al., 2016) by adding a branch for predicting segmentation masks, allowing it to identify and delineate individual objects within an image accurately.

SKA relevance: used for multi-component radio source detection and segmentation in SKA-precursor ASKAP EMU data, with direct applicability to SKA-Mid AA4 continuum surveys (Section 2.3).

A.4 Deep Generative Models

In recent years, there has been a strong increase in research in the domain of deep learning-based generative models (Bond-Taylor et al., 2022), in particular for generating high-quality images. Such models showcase impressive generative capabilities, ranging from very high levels of detail to the diversity of the generated samples. The most important types of such models are Generative Adversarial Networks (GANs) (Goodfellow et al., 2014, 2020), Normalizing Flows (NFs) (Dinh et al., 2014; Rezende and Mohamed, 2015; Dinh et al., 2016), Variational Auto-encoders (VAEs) (Kingma and Welling, 2013), and Diffusion Models (DMs) (Cao et al., 2024; Yang et al., 2023). In addition, there are also Auto-regressive (Van Den Oord et al., 2016) and Energy-based (LeCun et al., 2006; Song and Kingma, 2021) models.

A.4.1 Generative Adversarial Networks (GANs)

GANs (Park et al., 2021; Gui et al., 2021) learn to generate data examples x from latent variables z , using an adversarial loss that encourages the generated samples to be indistinguishable from

real examples in the training data set. GANs consist of two networks, a data generator, and a discriminator which learns to separate generated from real data.

SKA relevance: used for sky simulation, map-to-map translation between simulated and observable quantities (SKA-Mid HI science; Section 2.1), and data augmentation for rare radio morphologies (Section 2.3).

A.4.2 Normalizing Flows

Normalizing flows (Kobyzev et al., 2020; Papamakarios et al., 2021) learn a probability model by transforming a simple distribution into a more complicated one using a deep network. Normalizing flows can both sample from this distribution and evaluate the probability of new examples. However, they require specialized architecture: each layer must be invertible.

SKA relevance: used as the density-estimator backbone in simulation-based inference for EoR parameter estimation (Section 2.2) and gravitational lens modelling (Section 2.6).

A.4.3 Variational Auto-encoders (VAEs)

VAEs (Doersch, 2016) combine elements of both variational inference and auto-encoding architectures. They learn to encode input data into a latent space distribution and then decode it back to reconstruct the original input, while also optimizing a variational lower bound to ensure the latent space is structured and continuous, allowing for smooth generation of new data points.

SKA relevance: used for EoR foreground mitigation via VAE-trained GPR kernels (Section 2.2), and as anomaly detectors for transient discovery in SKA-Low and SKA-Mid data (Section 2.4).

A.4.4 Diffusion Models (DMs)

Diffusion models are highly flexible probabilistic models that operate in two distinct stages, the forward and reverse diffusion processes. In the forward process, noise is incrementally added to the input, such that at the end of the process the data is transformed into pure Gaussian noise. In the reverse or denoising process, the original input data is gradually recovered from the noise. Diffusion Models fall into three broad categories: Denoising Diffusion Probabilistic Models (DDPM) (Sohl-Dickstein et al., 2015; Ho et al., 2020), Score-based Generative Models (SGMs) or Noise Conditioned Score Networks (NCSNs) (Song and Ermon, 2019, 2020), and Stochastic Differential Equations (SDEs) (Karras et al., 2022; Song et al., 2021, 2020).

SKA relevance: applied across interferometric imaging (Section 2.1), EoR foreground separation (Section 2.2), magnetic field estimation (Section 2.5), and gravitational lens modelling (Section 2.6) – among the most versatile generative architectures for SKA science.

A.4.5 Comparison of generative models

Comparing the characteristics of GANs, VAEs, NFs and DMs (Xiao et al., 2021), GANs can efficiently generate high quality samples, but their mode coverage is limited. Conversely, NFs and VAEs lack in sampling quality, however with very good mode coverage. To this end, DMs are able to compensate for these deficiencies and offer good mode coverage and high sampling quality at the same time. DMs have shown impressive performance and outperformed GANs in image

synthesis (Dhariwal and Nichol, 2021). Nevertheless, due to their iterative nature, which causes a slow sampling process, they are computationally expensive.

A.4.6 Conditional Generation

Generative models can generate new samples such as images in an unconditional setting, while being trained in an unsupervised way. Perhaps more interesting for scientific applications is conditional image generation. The conditioning can be based on attributes such as class labels, text prompts or other images. The image-conditioned setting leads to the task of image-to-image (I2I) translation. Here, the model is conditioned on an input image y , such that for a probabilistic generative model the probability for each data point $p(x)$ becomes $p(x|y)$. The model can learn a transformation from a noisy, low-resolution, or otherwise degraded input image to a higher quality output image, often using a training data set of paired images. In computer vision, this could be denoising, super-resolution, style transfer, recolouring, etc.

SKA relevance: the image-to-image translation paradigm underlies map-to-map translation between simulation and observational domains, essential for bridging hydrodynamical simulations with SKA-Mid observations (Section 2.1).

A.4.7 Pix2pix and CycleGAN

In (Isola et al., 2017) (Pix2pix), the problem of I2I translation was tackled using a conditional GAN derived from the DCGAN (Radford et al., 2015) architecture: A U-Net (Ronneberger et al., 2015) architecture is used as a generator, which receives the noise input concatenated with the conditioning image, and a patch discriminator. In (Zhu et al., 2017) (CycleGAN), this approach was extended to unpaired training data. It uses two mapping functions $G : X \rightarrow Y$ and $F : Y \rightarrow X$ with associated adversarial discriminators, as well as a cycle consistency loss to enforce $F(G(X)) \sim X$ (and vice versa).

SKA relevance: conditional GAN frameworks including Pix2pix and CycleGAN have been applied to radio galaxy image synthesis and map-to-map translation tasks relevant for SKA-Mid forward modelling (Section 2.1 and 2.2).

A.4.8 Flow Matching models

Flow matching models (Albergo and Vanden-Eijnden, 2023; Albergo et al., 2023; Liu et al., 2022; Lipman et al., 2023) represent an innovative approach in generative modelling, merging the strengths of Continuous Normalizing Flows (CNFs) and diffusion models to address key limitations of both methodologies. Unlike traditional flow-based models that explicitly parametrize invertible mappings, flow matching focuses on directly estimating the velocity fields that transport one distribution to another, often by solving an ordinary differential equation (ODE) that characterizes the flow dynamics. This approach enables efficient and flexible generation by matching the probability flow paths without requiring explicit invertibility, thereby reducing computational complexity and training difficulty. Recent developments have extended flow matching to various modalities, including images and audio (Lipman et al., 2024; Holderrieth and Erives, 2025).

SKA relevance: flow matching architectures are applicable to pulse profile denoising for SKA pulsar

timing arrays (Section 2.4), and as computationally efficient alternatives to diffusion models for sky simulation and image reconstruction tasks (Section 2.1).

References

- A. Acharya et al. *Mon Not R Astron Soc*, 527(3):7835–7846, Jan. 2024a. ISSN 0035-8711. doi: 10.1093/mnras/stad3701. URL <https://doi.org/10.1093/mnras/stad3701>.
- A. Acharya et al. *Mon Not R Astron Soc Lett*, 534(1):L30–L34, Oct. 2024b. ISSN 1745-3925. doi: 10.1093/mnrasl/slae078. URL <https://doi.org/10.1093/mnrasl/slae078>.
- A. Acharya et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026a. arXiv search: Report number AASKAII/Acharya01.
- A. Acharya et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026b. arXiv search: Report number AASKAII/Acharya02.
- D. Agarwal et al. *Monthly Notices of the Royal Astronomical Society*, 497(2):1661–1674, 09 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1856. URL <https://doi.org/10.1093/mnras/staa1856>.
- A. Aghabiglou, C. S. Chu, A. Dabbech, and Y. Wiaux. *The Astrophysical Journal Supplement Series*, 273(1):3, 2024. doi: 10.3847/1538-4365/ad46f5. URL <http://dx.doi.org/10.3847/1538-4365/ad46f5>.
- A. Aghabiglou et al. *The Astrophysical Journal Supplement Series*, 280(2):63, 2025. doi: 10.3847/1538-4365/adfbcd. URL <http://dx.doi.org/10.3847/1538-4365/adfbcd>.
- J. Akeret, C. Chang, A. Lucchi, and A. Refregier. *Astronomy and Computing*, 18:35–39, 2017. doi: 10.1016/j.ascom.2017.01.002. URL <http://dx.doi.org/10.1016/j.ascom.2017.01.002>.
- M. S. Albergo and E. Vanden-Eijnden. Building normalizing flows with stochastic interpolants, 2023. URL <https://arxiv.org/abs/2209.15571>.
- M. S. Albergo, N. M. Boffi, and E. Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL <https://arxiv.org/abs/2303.08797>.
- G. E. Anderson et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/GemmaAnderson01.
- A. Andersson et al. *Mon Not R Astron Soc*, 538(3):1397–1414, Apr. 2025. ISSN 0035-8711. doi: 10.1093/mnras/staf336. URL <https://doi.org/10.1093/mnras/staf336>.
- A. Andersson et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026a. arXiv search: Report number AASKAII/AlexAndersson01.
- A. Andersson et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026b. arXiv search: Report number AASKAII/AlexAndersson02.
- N. Baron Perez, M. Brüggén, G. Kasieczka, and L. Lucie-Smith. *A&A*, 699:A302, 2025. doi: 10.1051/0004-6361/202554735. URL <https://doi.org/10.1051/0004-6361/202554735>.
- M. A. Beaumont. *Annual review of statistics and its application*, 6(1):379–403, 2019. doi: 10.1146/annurev-statistics-030718-105212. URL <https://doi.org/10.1146/annurev-statistics-030718-105212>.
- M. Bernardini et al. *MNRAS*, 509(1):1323–1341, 2021. doi: 10.1093/mnras/stab3088. URL <http://dx.doi.org/10.1093/mnras/stab3088>.
- M. Bernardini et al. *Monthly Notices of the Royal Astronomical Society*, 538(2):1201–1215, 02 2025. ISSN 0035-8711. doi: 10.1093/mnras/staf341. URL <https://doi.org/10.1093/mnras/staf341>.

- M. Bernardini et al. *Monthly Notices of the Royal Astronomical Society*, page stag204, 03 2026. ISSN 0035-8711. doi: 10.1093/mnras/stag204. URL <https://doi.org/10.1093/mnras/stag204>.
- M. Bianco et al. *Monthly Notices of the Royal Astronomical Society*, 541(1):234–250, 2025. doi: 10.1093/mnras/staf973. URL <http://dx.doi.org/10.1093/mnras/staf973>.
- Y. Billeter et al. In *2024 11th IEEE Swiss Conference on Data Science (SDS)*, pages 37–40, 5 2024. doi: 10.1109/sds60720.2024.00013. URL <http://dx.doi.org/10.1109/SDS60720.2024.00013>.
- A. Bonaldi and R. Braun. *CoRR*, 2018. URL <http://arxiv.org/abs/1811.10454v1>.
- A. Bonaldi et al. *Monthly Notices of the Royal Astronomical Society*, 543(2):1092–1119, 2025. doi: 10.1093/mnras/staf1466. URL <http://dx.doi.org/10.1093/mnras/staf1466>.
- S. Bond-Taylor, A. Leach, Y. Long, and C. G. Willcocks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7327–7347, 2022. doi: 10.1109/TPAMI.2021.3116668. URL <https://doi.org/10.1109/TPAMI.2021.3116668>.
- T. L. Bourke et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Bourke01.
- E. Breck et al. In *2017 IEEE International Conference on Big Data (Big Data)*, pages 1123–1132. IEEE, Dec. 2017. doi: 10.1109/bigdata.2017.8258038. URL <https://api.semanticscholar.org/CorpusID:6244440>.
- D. Breitman et al. *Monthly Notices of the Royal Astronomical Society*, 527(4):9833–9852, 2023. doi: 10.1093/mnras/stad3849. URL <http://dx.doi.org/10.1093/mnras/stad3849>.
- T. B. Brown et al. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546. URL <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>.
- J. Burba et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Burba01.
- M. Burgay et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Joshi01.
- C. J. Burke et al. *Monthly Notices of the Royal Astronomical Society*, 490(3):3952–3965, 2019. doi: 10.1093/mnras/stz2845. URL <http://dx.doi.org/10.1093/mnras/stz2845>.
- H. Cao et al. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–20, 2024. doi: 10.1109/TKDE.2024.3361474. URL <https://doi.org/10.1109/TKDE.2024.3361474>.
- M. Cárcamo, A. M. M. Scaife, E. L. Alexander, and J. P. Leahy. *Monthly Notices of the Royal Astronomical Society*, 518(2):1955–1974, 2022. doi: 10.1093/mnras/stac3031. URL <http://dx.doi.org/10.1093/mnras/stac3031>.
- M. Carcamo et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Carcamo01.
- M. Caron et al. Emerging properties in self-supervised vision transformers, 2021. URL <https://arxiv.org/abs/2104.14294>.
- E. Carretti and F. Vazza. *Universe*, 11(5):164, 2025. doi: 10.3390/universe11050164. URL <http://dx.doi.org/10.3390/universe11050164>.
- R. E. Carrillo, J. D. McEwen, and Y. Wiaux. *Monthly Notices of the Royal Astronomical Society*, 426(2):1223–1234, 2012. doi: 10.1111/j.1365-2966.2012.21605.x. URL <http://dx.doi.org/10.1111/j.1365-2966.2012.21605.x>.

- [org/10.1111/j.1365-2966.2012.21605.x](https://doi.org/10.1111/j.1365-2966.2012.21605.x).
- E. Ceccotti et al. *Mon Not R Astron Soc*, page staf1629, Sept. 2025. ISSN 0035-8711. doi: 10.1093/mnras/staf1629. URL <https://doi.org/10.1093/mnras/staf1629>.
- N. Cerardi et al. *CoRR*, 2025. URL <http://arxiv.org/abs/2511.11568v1>.
- V. Chandola, A. Banerjee, and V. Kumar. *ACM computing surveys (CSUR)*, 41(3):1–58, 2009. doi: 10.1145/1541880.1541882. URL <https://doi.org/10.1145/1541880.1541882>.
- T.-Y. Cheng et al. *Monthly Notices of the Royal Astronomical Society*, 494(3):3750–3765, 2020. doi: 10.1093/mnras/staa1015. URL <http://dx.doi.org/10.1093/mnras/staa1015>.
- M. Choudhury, A. Datta, and S. Majumdar. *Mon Not R Astron Soc*, 512(4):5010–5022, June 2022. ISSN 0035-8711. doi: 10.1093/mnras/stac736. URL <https://doi.org/10.1093/mnras/stac736>.
- L. Connor and J. van Leeuwen. *The Astronomical Journal*, 156(6):256, 2018. doi: 10.3847/1538-3881/aae649. URL <http://dx.doi.org/10.3847/1538-3881/aae649>.
- L. Connor, K. L. Bouman, V. Ravi, and G. Hallinan. *Monthly Notices of the Royal Astronomical Society*, 514(2):2614–2626, 2022. doi: 10.1093/mnras/stac1329. URL <http://dx.doi.org/10.1093/mnras/stac1329>.
- D. Cornu et al. *Astronomy & Astrophysics*, 690:A211, 2024. doi: 10.1051/0004-6361/202449548. URL <http://dx.doi.org/10.1051/0004-6361/202449548>.
- D. Cornu et al. *A&A*, 707:A203, 2026. doi: 10.1051/0004-6361/202557257. URL <https://doi.org/10.1051/0004-6361/202557257>.
- K. Cranmer, J. Brehmer, and G. Louppe. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020. doi: 10.1073/pnas.1912789117. URL <https://doi.org/10.1073/pnas.1912789117>.
- E. de Salis et al. *Exploring the Early Universe with Deep Learning*, pages 426–438. Lecture Notes in Computer Science. Springer Nature Switzerland, 2025. doi: 10.1007/978-3-032-05176-9_33. URL http://dx.doi.org/10.1007/978-3-032-05176-9_33.
- D. Deal and P. Jagannathan. *Research Notes of the AAS*, 8(10):257, 2024. doi: 10.3847/2515-5172/ad84fa. URL <http://dx.doi.org/10.3847/2515-5172/ad84fa>.
- J. Deng et al. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. doi: 10.1109/cvpr.2009.5206848. URL <https://doi.org/10.1109/cvpr.2009.5206848>.
- P. Denzel et al. In *2024 11th IEEE Swiss Conference on Data Science (SDS)*, pages 84–91, 2024. doi: 10.1109/SDS60720.2024.00020. URL <https://doi.org/10.1109/SDS60720.2024.00020>.
- P. Denzel, Y. Billeter, F.-P. Schilling, and E. Gavagnin. *Monthly Notices of the Royal Astronomical Society*, 546(4):stag155, 02 2026. ISSN 0035-8711. doi: 10.1093/mnras/stag155. URL <https://doi.org/10.1093/mnras/stag155>.
- J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. In J. Burstein, C. Doran, and T. Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423/>.
- P. Dhariwal and A. Nichol. In *Proceedings of the 35th International Conference on Neural Information Processing Systems, NIPS '21*, Red Hook, NY, USA, 2021. Curran Associates

- Inc. ISBN 9781713845393. URL <https://proceedings.neurips.cc/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf>.
- L. Dinh, D. Krueger, and Y. Bengio. *arXiv preprint arXiv:1410.8516*, 2014. doi: 10.48550/arXiv.1410.8516. URL <https://doi.org/10.48550/arXiv.1410.8516>.
- L. Dinh, J. Sohl-Dickstein, and S. Bengio. *arXiv preprint arXiv:1605.08803*, 2016. doi: 10.48550/arXiv.1605.08803. URL <https://doi.org/10.48550/arXiv.1605.08803>.
- C. Doersch. *arXiv preprint arXiv:1606.05908*, 2016. doi: 10.48550/arXiv.1606.05908. URL <https://doi.org/10.48550/arXiv.1606.05908>.
- A. Dosovitskiy et al. *arXiv preprint arXiv:2010.11929*, 2020. doi: 10.48550/arXiv.2010.11929. URL <https://doi.org/10.48550/arXiv.2010.11929>.
- M. Drozdova et al. *Astronomy & Astrophysics*, 683:A105, 2024. doi: 10.1051/0004-6361/202347948. URL <http://dx.doi.org/10.1051/0004-6361/202347948>.
- M. Drozdova et al. *CoRR*, 2025. URL <http://arxiv.org/abs/2509.02615v2>.
- F. Dulwich and B. Mort. *OSKAR 2.7.0*. Zenodo, 2017. doi: 10.5281/ZENODO.3758449. URL <https://zenodo.org/record/3758449>.
- V. Etsebeth, M. Lochner, M. Walmsley, and M. Grespan. *Monthly Notices of the Royal Astronomical Society*, 529(1):732–747, 2024. doi: 10.1093/mnras/stae496. URL <http://dx.doi.org/10.1093/mnras/stae496>.
- Q. Fu et al. *Mon Not R Astron Soc*, page staf1952, Nov. 2025. ISSN 0035-8711. doi: 10.1093/mnras/staf1952. URL <https://doi.org/10.1093/mnras/staf1952>.
- Y. Gal and Z. Ghahramani. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1050–1059, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/gal16.html>.
- L.-Y. Gao et al. *ApJ*, 988(1):84, July 2025. ISSN 0004-637X. doi: 10.3847/1538-4357/ade2dc. URL <https://doi.org/10.3847/1538-4357/ade2dc>.
- I. Goodfellow et al. *Advances in neural information processing systems*, 27, 2014. doi: 10.48550/arXiv.1406.2661. URL <https://doi.org/10.48550/arXiv.1406.2661>.
- I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- I. Goodfellow et al. *Communications of the ACM*, 63(11):139–144, 2020. ISSN 1557-7317. doi: 10.1145/3422622. URL <https://doi.org/10.1145/3422622>.
- J. Gui et al. *IEEE transactions on knowledge and data engineering*, 35(4):3313–3332, 2021. ISSN 2326-3865. doi: 10.1109/tkde.2021.3130191. URL <https://doi.org/10.1109/tkde.2021.3130191>.
- R. Guidotti et al. *ACM computing surveys (CSUR)*, 51(5):1–42, Aug. 2018. ISSN 1557-7341. doi: 10.1145/3236009. URL <https://doi.org/10.1145/3236009>.
- V. Gustafsson, M. Brüggen, and T. Enßlin. *A&A*, 692:A248, Dec. 2024. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202451265. URL <https://doi.org/10.1051/0004-6361/202451265>.
- C. Hale and F. Tabatabaei. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Hale01.
- P. J. Hancock et al. *Monthly Notices of the Royal Astronomical Society*, 422(2):1812–1824, 05

2012. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2012.20768.x. URL <https://doi.org/10.1111/j.1365-2966.2012.20768.x>.
- P. J. Hancock, C. M. Trott, and N. Hurley-Walker. *Publications of the Astronomical Society of Australia*, 35:e011, 2018. doi: 10.1017/pasa.2018.3. URL <https://doi.org/10.1017/pasa.2018.3>.
- M. J. Hardcastle et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Hardcastle01.
- K. He, X. Zhang, S. Ren, and J. Sun. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778. IEEE, 2016a. doi: 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- K. He, X. Zhang, S. Ren, and J. Sun. *Identity Mappings in Deep Residual Networks*, pages 630–645. Springer International Publishing, Cham, 2016b. ISBN 9783319464930. doi: 10.1007/978-3-319-46493-0_38. URL https://doi.org/10.1007/978-3-319-46493-0_38.
- K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn, Feb. 2018. ISSN 1939-3539. URL <https://doi.org/10.1109/tpami.2018.2844175>.
- G. Heald et al. *Galaxies*, 8(3):53, 2020. doi: 10.3390/galaxies8030053. URL <http://dx.doi.org/10.3390/galaxies8030053>.
- Y. D. Hezaveh, L. P. Levasseur, and P. J. Marshall. *Nature*, 548(7669):555–557, 2017. doi: 10.1038/nature23463. URL <http://dx.doi.org/10.1038/nature23463>.
- J. Ho, A. Jain, and P. Abbeel. *Advances in neural information processing systems*, 33:6840–6851, 2020. doi: 10.48550/arXiv.2006.11239. URL <https://doi.org/10.48550/arXiv.2006.11239>.
- P. Holderrieth and E. Erives. An introduction to flow matching and diffusion models, 2025. URL <https://arxiv.org/abs/2506.02070>.
- A. Hota et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Hota01.
- Y. Hu. *The Astrophysical Journal*, 990(1):76, 2025. doi: 10.3847/1538-4357/adf32e. URL <http://dx.doi.org/10.3847/1538-4357/adf32e>.
- Y. Hu and A. Lazarian. *The Astrophysical Journal*, 975(1):66, 2024. doi: 10.3847/1538-4357/ad7950. URL <http://dx.doi.org/10.3847/1538-4357/ad7950>.
- Y. Hu and A. Lazarian. *The Astrophysical Journal*, 981(1):58, 2025. doi: 10.3847/1538-4357/adaf97. URL <http://dx.doi.org/10.3847/1538-4357/adaf97>.
- C. Huyen. *Designing machine learning systems*. " O'Reilly Media, Inc.", [S.l.], 2022. ISBN 1098107969.
- P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134. IEEE, 2017. doi: 10.1109/cvpr.2017.632. URL <https://doi.org/10.1109/cvpr.2017.632>.
- P. Kairouz et al. Advances and open problems in federated learning, 2021. URL <https://doi.org/10.1561/9781680837896>.
- G. E. Karniadakis et al. *Nature Reviews Physics*, 3(6):422–440, 2021. doi: 10.1038/s42254-021-00314-5. URL <https://doi.org/10.1038/s42254-021-00314-5>.
- T. Karras, M. Aittala, T. Aila, and S. Laine. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022. doi: 10.52202/068431-1926. URL <https://doi.org/10.52202/068431-1926>.

- 068431–1926.
- E. F. Keane et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Keane01.
- D. P. Kingma and M. Welling. *arXiv preprint arXiv:1312.6114*, 2013. doi: 10.48550/arXiv.1312.6114. URL <https://doi.org/10.48550/arXiv.1312.6114>.
- A. Kirillov et al. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4015–4026. IEEE, Oct. 2023. doi: 10.1109/iccv51070.2023.00371. URL <https://doi.org/10.1109/iccv51070.2023.00371>.
- B. M. Kirk, U. Rau, and R. Ramyaa. *The Astronomical Journal*, 169(1):43, 2024. doi: 10.3847/1538-3881/ad88f6. URL <http://dx.doi.org/10.3847/1538-3881/ad88f6>.
- I. Kobyzev, S. J. Prince, and M. A. Brubaker. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):3964–3979, Nov. 2020. ISSN 1939-3539. doi: 10.1109/tpami.2020.2992934. URL <https://doi.org/10.1109/tpami.2020.2992934>.
- D. Korber, M. Bianco, E. Tolley, and J.-P. Kneib. *Monthly Notices of the Royal Astronomical Society*, 521(1):902–915, 2023. doi: 10.1093/mnras/stad615. URL <http://dx.doi.org/10.1093/mnras/stad615>.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. *Commun. ACM*, 60(6):84–90, May 2017. ISSN 0001-0782. doi: 10.1145/3065386. URL <https://doi.org/10.1145/3065386>.
- D. Kuiper, G. Contardo, D. Huppenkothen, and J. W. T. Hessels. *Monthly Notices of the Royal Astronomical Society*, 538(1):408–425, 03 2025. ISSN 0035-8711. doi: 10.1093/mnras/staf306. URL <https://doi.org/10.1093/mnras/staf306>.
- P. La Plante and M. Ntampaka. *ApJ*, 880(2):110, July 2019. ISSN 0004-637X. doi: 10.3847/1538-4357/ab2983. URL <https://doi.org/10.3847/1538-4357/ab2983>.
- C. D. Lagos et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Lagos01.
- E. Lastufka et al. *A&A*, 690:A310, Oct. 2024. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202449964. URL <https://doi.org/10.1051/0004-6361/202449964>.
- E. Lastufka et al. *Astronomy & Astrophysics*, 703:A217, 2025. doi: 10.1051/0004-6361/202553691. URL <http://dx.doi.org/10.1051/0004-6361/202553691>.
- Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. ISSN 0018-9219. doi: 10.1109/5.726791. URL <https://doi.org/10.1109/5.726791>.
- Y. LeCun et al. *A tutorial on energy-based learning*. MIT Press, 2006. URL <https://nyuscholars.nyu.edu/en/publications/a-tutorial-on-energy-based-learning>.
- Y. LeCun, Y. Bengio, and G. Hinton. *nature*, 521(7553):436–444, May 2015. ISSN 1476-4687. doi: 10.1038/nature14539. URL <https://doi.org/10.1038/nature14539>.
- R. Legin, Y. Hezaveh, L. P. Levasseur, and B. Wandelt. *CoRR*, 2021. URL <http://arxiv.org/abs/2112.05278v2>.
- W. Li et al. *Mon Not R Astron Soc*, 485(2):2628–2637, May 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz582. URL <https://doi.org/10.1093/mnras/stz582>.
- Z. Li et al. *Astronomy and Computing*, 36:100482, 2021. ISSN 2213-1337. doi: <https://doi.org/10.1016/j.ascom.2021.100482>. URL <https://www.sciencedirect.com/science/article/pii/S2213133721000366>.
- T. I. Liaudat et al. *CoRR*, 2023. URL <http://arxiv.org/abs/2312.00125v3>.

- Y. Lipman et al. Flow matching for generative modeling, 2023. URL <https://doi.org/10.48550/arXiv.2210.02747>.
- Y. Lipman et al. Flow matching guide and code, 2024. URL <https://doi.org/10.48550/arXiv.2412.06264>.
- X. Liu, C. Gong, and Q. Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow, 2022. URL <https://doi.org/10.48550/arXiv.2209.03003>.
- Y. Liu, E. d. L. Acedo, and P. Sims. Bayesian model comparison and validation with Gaussian Process Regression for interferometric 21-cm signal recovery, Nov. 2025a. URL <https://doi.org/10.48550/arXiv.2511.10499>.
- Y. Liu et al. *The Astrophysical Journal*, 978(2):151, 2025b. doi: 10.3847/1538-4357/ad97ba. URL <http://dx.doi.org/10.3847/1538-4357/ad97ba>.
- F. Loi et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Loi01.
- Y. K. Ma et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Ma01.
- D. J. C. MacKay. *Neural Computation*, 4(3):448–472, May 1992. ISSN 1530-888X. doi: 10.1162/neco.1992.4.3.448. URL <https://doi.org/10.1162/neco.1992.4.3.448>.
- Y. Mahida et al. *Journal of Cosmology and Astroparticle Physics*, 2025(12):055, Dec. 2025. ISSN 1475-7516. doi: 10.1088/1475-7516/2025/12/055. URL <https://doi.org/10.1088/1475-7516/2025/12/055>.
- T. Mangena, S. Hassan, and M. G. Santos. *Mon Not R Astron Soc*, 494(1):600–606, May 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa750. URL <https://doi.org/10.1093/mnras/staa750>.
- B. Margalef-Bentabol et al. *Monthly Notices of the Royal Astronomical Society*, 496(2):2346–2361, 08 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1647. URL <https://doi.org/10.1093/mnras/staa1647>.
- T. V. Martínez, N. B. Pérez, and M. Brüggen. *A&A*, 691:A360, Nov. 2024. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202451429. URL <https://doi.org/10.1051/0004-6361/202451429>.
- J. McKean et al. In *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, page 84, Apr. 2015. doi: 10.22323/1.215.0084. URL <https://doi.org/10.22323/1.215.0084>.
- H. B. McMahan et al. Communication-efficient learning of deep networks from decentralized data, 2023. URL <https://doi.org/10.48550/arXiv.1602.05629>.
- F. G. Mertens et al. *A&A*, 698:A186, June 2025. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202554158. URL <https://doi.org/10.1051/0004-6361/202554158>.
- A. Mesinger, S. Furlanetto, and R. Cen. *Monthly Notices of the Royal Astronomical Society*, 411(2):955–972, 2010. doi: 10.1111/j.1365-2966.2010.17731.x. URL <http://dx.doi.org/10.1111/j.1365-2966.2010.17731.x>.
- J. Miller-Jones, K. Rajwade, P. A. Woudt, and J. W. Hessels. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Miller-Jones01.
- T. M. Mitchell. *Machine learning*. McGraw-Hill, New York, 1997. URL <https://www.cs.cmu.edu/~tom/mlbook.html>.
- V. Mnih et al. *Nature*, 518(7540):529–533, Feb. 2015. ISSN 1476-4687. doi: 10.1038/nature14236. URL <https://doi.org/10.1038/nature14236>.

- N. Mohan and D. Rafferty. PyBDSF: Python Blob Detection and Source Finder. *Astrophysics Source Code Library*, record ascl:1502.007, Feb. 2015.
- C. Molnar. *Interpretable Machine Learning*. 3 edition, 2025. ISBN 978-3-911578-03-5. URL <https://christophm.github.io/interpretable-ml-book>.
- W. R. Morningstar et al. *CoRR*, 2018. URL <http://arxiv.org/abs/1808.00011v1>.
- W. R. Morningstar et al. *The Astrophysical Journal*, 883(1):14, 2019. doi: 10.3847/1538-4357/ab35d7. URL <http://dx.doi.org/10.3847/1538-4357/ab35d7>.
- R. I. J. Mostert et al. *Astronomy & Astrophysics*, 668:A28, 2022. doi: 10.1051/0004-6361/202243478. URL <http://dx.doi.org/10.1051/0004-6361/202243478>.
- K. Murphy. Reinforcement learning: An overview, 2025. URL <https://arxiv.org/abs/2412.05265>.
- S. W. Ndiritu et al. *Monthly Notices of the Royal Astronomical Society*, 502(4):5839–5853, 2021. doi: 10.1093/mnras/stab379. URL <http://dx.doi.org/10.1093/mnras/stab379>.
- R. M. Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- S. Neutsch, C. Heneka, and M. Brüggem. *Mon Not R Astron Soc*, 511(3):3446–3462, Apr. 2022. ISSN 0035-8711. doi: 10.1093/mnras/stac218. URL <https://doi.org/10.1093/mnras/stac218>.
- H. Nguyen et al. *Mon Not R Astron Soc*, 536(1):962–987, Jan. 2025. ISSN 0035-8711. doi: 10.1093/mnras/stae2631. URL <https://doi.org/10.1093/mnras/stae2631>.
- J. W. Nightingale, S. Dye, and R. J. Massey. *Monthly Notices of the Royal Astronomical Society*, 478(4):4738–4784, 2018. doi: 10.1093/mnras/sty1264. URL <http://dx.doi.org/10.1093/mnras/sty1264>.
- J. Oppliger et al. *Nature Machine Intelligence*, 6(2):180–186, 2024. doi: 10.1038/s42256-024-00790-1. URL <http://dx.doi.org/10.1038/s42256-024-00790-1>.
- S. P. O’Sullivan et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/OSullivan01.
- S. Pal et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/SabyasachiPal01.
- G. Pang, C. Shen, L. Cao, and A. V. D. Hengel. *ACM Comput. Surv.*, 54(2), Mar. 2021. ISSN 0360-0300. doi: 10.1145/3439950. URL <https://doi.org/10.1145/3439950>.
- G. Papamakarios et al. *The Journal of Machine Learning Research*, 22(1):2617–2680, 2021. URL <http://jmlr.org/papers/v22/19-1028.html>.
- S.-W. Park, J.-S. Ko, J.-H. Huh, and J.-C. Kim. *Electronics*, 10(10):1216, 2021. doi: 10.3390/electronics10101216. URL <https://doi.org/10.3390/electronics10101216>.
- L. Parker et al. *Monthly Notices of the Royal Astronomical Society*, 531(4):4990–5011, 07 2024. ISSN 0035-8711. doi: 10.1093/mnras/stae1450. URL <https://doi.org/10.1093/mnras/stae1450>.
- K. Pattle et al. *CoRR*, 2022. URL <http://arxiv.org/abs/2203.11179v1>.
- J. E. G. Peek and B. Burkhart. *The Astrophysical Journal Letters*, 882(1):L12, 2019. doi: 10.3847/2041-8213/ab3a9e. URL <http://dx.doi.org/10.3847/2041-8213/ab3a9e>.
- Y. Pietschke et al. *J. Cosmol. Astropart. Phys.*, 2025(10):039, Oct. 2025. ISSN 1475-7516. doi: 10.1088/1475-7516/2025/10/039. URL <https://doi.org/10.1088/1475-7516/2025/10/039>.

- D. Piras et al. *The Open Journal of Astrophysics*, 7, 2024. doi: 10.33232/001c.123368. URL <http://dx.doi.org/10.33232/001c.123368>.
- I. Prandoni and N. Seymour. In *Proceedings of Advancing Astrophysics with the Square Kilometre Array — PoS(AASKA14)*, AASKA14, page 67. Sissa Medialab, Apr. 2015. doi: 10.22323/1.215.0067. URL <https://doi.org/10.22323/1.215.0067>.
- I. Prandoni et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Prandoni01.
- D. Prelogović and A. Mesinger. *Monthly Notices of the Royal Astronomical Society*, 524(3):4239–4255, 09 2023. ISSN 0035-8711. doi: 10.1093/mnras/stad2027. URL <https://doi.org/10.1093/mnras/stad2027>.
- N. J. Pritchard, A. Wicenec, M. Bennamoun, and R. Dodson. *CoRR*, 2025. URL <http://arxiv.org/abs/2504.09796v1>.
- H. Qiu et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Qiu01.
- A. Radford, L. Metz, and S. Chintala. 2015. doi: 10.48550/arXiv.1511.06434. URL <https://doi.org/10.48550/arXiv.1511.06434>.
- A. Radford et al. Learning transferable visual models from natural language supervision, 2021. URL <https://arxiv.org/abs/2103.00020>.
- M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations, 2017a. URL <https://arxiv.org/abs/1711.10561>.
- M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics informed deep learning (part ii): Data-driven discovery of nonlinear partial differential equations, 2017b. URL <https://arxiv.org/abs/1711.10566>.
- M. Raissi, P. Perdikaris, and G. Karniadakis. *Journal of Computational Physics*, 378:686–707, 2019. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2018.10.045>. URL <https://www.sciencedirect.com/science/article/pii/S0021999118307125>.
- J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. doi: 10.1109/cvpr.2016.91. URL <https://doi.org/10.1109/cvpr.2016.91>.
- B. Remy et al. *Astronomy & Astrophysics*, 672:A51, 2023. doi: 10.1051/0004-6361/202243054. URL <http://dx.doi.org/10.1051/0004-6361/202243054>.
- S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks, 2016. URL <https://arxiv.org/abs/1506.01497>.
- S. Rezaei et al. *Monthly Notices of the Royal Astronomical Society*, 517(1):1156–1170, Oct. 2022. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stac2078. URL <https://doi.org/10.1093/mnras/stac2078>.
- D. J. Rezende and S. Mohamed. In *International conference on machine learning*, pages 1530–1538. PMLR, arXiv, 2015. doi: 10.48550/arXiv.1505.05770. URL <https://doi.org/10.48550/arXiv.1505.05770>.
- S. Riggi et al. *Astronomy and Computing*, 42:100682, Jan. 2023. ISSN 2213-1337. doi: 10.1016/j.ascom.2022.100682. URL <https://doi.org/10.1016/j.ascom.2022.100682>.
- S. Riggi, T. Cecconello, U. Becciani, and F. Vitello. *CoRR*, 2024a. URL <http://arxiv.org/>

- [abs/2411.08519v1](#).
- S. Riggi et al. *Publications of the Astronomical Society of Australia*, 41:e085, 2024b. doi: 10.1017/pasa.2024.84. URL <https://doi.org/10.1017/pasa.2024.84>.
- S. Riggi et al. *Publications of the Astronomical Society of Australia*, 42:e121, 2025. ISSN 1448-6083. doi: 10.1017/pasa.2025.10082. URL <http://dx.doi.org/10.1017/pasa.2025.10082>.
- C. Robert and G. Casella. *Statistical Science*, 26(1), Feb. 2011. ISSN 0883-4237. doi: 10.1214/10-sts351. URL <http://dx.doi.org/10.1214/10-ST351>.
- T. Robishaw et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Robishaw01.
- O. Ronneberger, P. Fischer, and T. Brox. In N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. ISBN 9783319245744. doi: 10.1007/978-3-319-24574-4_28. URL https://doi.org/10.1007/978-3-319-24574-4_28.
- J. Rüstig et al. *A&A*, 682:A146, Feb. 2024. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202348256. URL <https://doi.org/10.1051/0004-6361/202348256>.
- A. V. Sadr et al. *Monthly Notices of the Royal Astronomical Society*, 499(1):379–390, 2020. doi: 10.1093/mnras/staa2724. URL <http://dx.doi.org/10.1093/mnras/staa2724>.
- T. Sasmal et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Sasmal01.
- A. M. M. Scaife. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 378(2166):20190060, 2020. doi: 10.1098/rsta.2019.0060. URL <http://dx.doi.org/10.1098/rsta.2019.0060>.
- T. Schaeffer, S. K. Giri, and A. Schneider. *Monthly Notices of the Royal Astronomical Society*, 526(2):2942–2959, 2023. doi: 10.1093/mnras/stad2937. URL <http://dx.doi.org/10.1093/mnras/stad2937>.
- K. Schmidt et al. *Astronomy & Astrophysics*, 664:A134, 2022. doi: 10.1051/0004-6361/202142113. URL <http://dx.doi.org/10.1051/0004-6361/202142113>.
- B. Schosser, C. Heneka, and T. Plehn. *SciPost Phys. Core*, 8:037, 2025. doi: 10.21468/SciPostPhysCore.8.2.037. URL <https://scipost.org/10.21468/SciPostPhysCore.8.2.037>.
- D. Sculley et al. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’15, page 2503–2511, Cambridge, MA, USA, 2015. MIT Press. URL <https://proceedings.neurips.cc/paper/2015/file/86df7dcfd896fcaf2674f757a2463eba-Paper.pdf>.
- B. Semelin, R. Mériot, A. Mishra, and D. Cornu. *A&A*, 698:A35, 2025. doi: 10.1051/0004-6361/202453115. URL <https://doi.org/10.1051/0004-6361/202453115>.
- R. M. Shannon et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Shannon01.
- R. Sharma et al. *Astronomy and Computing*, 54:101004, 2026. doi: 10.1016/j.ascom.2025.101004. URL <http://dx.doi.org/10.1016/j.ascom.2025.101004>.
- SKA Observatory. SRCNet documents. <https://confluence.skatelescope.org/display/SNC/SRCNet+Documents>, 2025.

- I. V. Slijepcevic et al. 2023. doi: 10.48550/ARXIV.2305.16127. URL <https://arxiv.org/abs/2305.16127>.
- J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli. In *International conference on machine learning*, pages 2256–2265. PMLR, arXiv, 2015. doi: 10.48550/arXiv.1503.03585. URL <https://doi.org/10.48550/arXiv.1503.03585>.
- Y. Song and S. Ermon. *Advances in neural information processing systems*, 32, 2019. doi: 10.48550/arXiv.1907.05600. URL <https://doi.org/10.48550/arXiv.1907.05600>.
- Y. Song and S. Ermon. *Advances in neural information processing systems*, 33:12438–12448, 2020. doi: 10.48550/arXiv.2006.09011. URL <https://doi.org/10.48550/arxiv.2006.09011>.
- Y. Song and D. P. Kingma. *arXiv preprint arXiv:2101.03288*, 2021. doi: 10.48550/arXiv.2101.03288. URL <https://doi.org/10.48550/arXiv.2101.03288>.
- Y. Song et al. *arXiv preprint arXiv:2011.13456*, 2020. doi: 10.48550/arXiv.2011.13456. URL <https://doi.org/10.48550/arxiv.2011.13456>.
- Y. Song, C. Durkan, I. Murray, and S. Ermon. *Advances in Neural Information Processing Systems*, 34:1415–1428, 2021. URL <https://openreview.net/forum?id=AklttWFnxS9>.
- R. Sortino et al. *Experimental Astronomy*, 56(1):293–331, 2023. doi: 10.1007/s10686-023-09893-w. URL <http://dx.doi.org/10.1007/s10686-023-09893-w>.
- M. Spinelli et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Spinelli01.
- N. Sravan, M. J. Graham, C. Fremling, and M. W. Coughlin. 2022. URL <https://arxiv.org/abs/2112.05897>.
- SRCNet. SRCNet v0.1 implementation plan. Technical Report SRC-0000009-01, Feb 2025. URL <https://docs.google.com/document/d/11DnUf52hDejtUr76DCaCss60P9K0IEzxeR-tUx0kGyM/edit?usp=sharing>.
- X. Sun et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Sun01.
- R. S. Sutton, A. G. Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998. doi: 10.1109/tnn.1998.712192. URL <https://doi.org/10.1109/tnn.1998.712192>.
- P. Swierc, M. Tamargo-Arizmendi, A. Ćiprijanović, and B. D. Nord. *CoRR*, 2024. URL <http://arxiv.org/abs/2410.16347v1>.
- M. Tahani et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Tahani01.
- A. Tajja et al. *Monthly Notices of the Royal Astronomical Society*, page staf1082, 2025. doi: 10.1093/mnras/staf1082. URL <http://dx.doi.org/10.1093/mnras/staf1082>.
- K. Takahashi. *Publications of the Astronomical Society of Japan*, 75:S50–S84, 2023. doi: 10.1093/pasj/psac111. URL <http://dx.doi.org/10.1093/pasj/psac111>.
- M. Terris, A. Dabbech, C. Tang, and Y. Wiaux. *Monthly Notices of the Royal Astronomical Society*, 518(1):604–622, 2022. doi: 10.1093/mnras/stac2672. URL <http://dx.doi.org/10.1093/mnras/stac2672>.
- H. Tiwari et al. *J. Cosmol. Astropart. Phys.*, 2022(04):045, Apr. 2022. ISSN 1475-7516. doi: 10.1088/1475-7516/2022/04/045. URL <https://doi.org/10.1088/1475-7516/2022/04/045>.

- K. D. Toennies. *An Introduction to Image Classification: From Designed Models to End-to-End Learning*. Springer Nature Singapore, 2024. ISBN 9789819978823. doi: 10.1007/978-981-99-7882-3. URL <https://doi.org/10.1007/978-981-99-7882-3>.
- E. Tolley. *RSL*, 5, 2024. doi: 10.46620/23-0028. URL <https://doi.org/10.46620/23-0028>.
- E. Tolley et al. *Astronomy and Computing*, 51:100920, 2025. doi: 10.1016/j.ascom.2024.100920. URL <http://dx.doi.org/10.1016/j.ascom.2024.100920>.
- A. Tripathi, A. Datta, A. Mazumder, and S. Majumdar. *J. Cosmol. Astropart. Phys.*, 2025(10):035, Oct. 2025. ISSN 1475-7516. doi: 10.1088/1475-7516/2025/10/035. URL <https://doi.org/10.1088/1475-7516/2025/10/035>.
- C. Trott et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Trott01.
- A. Van Den Oord, N. Kalchbrenner, and K. Kavukcuoglu. In *International conference on machine learning*, pages 1747–1756. PMLR, arXiv, 2016. doi: 10.48550/arXiv.1601.06759. URL <https://doi.org/10.48550/arxiv.1601.06759>.
- A. Vaswani et al. In I. Guyon et al., editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- T. Vernstrom, J. L. West, and C. Horellou. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Vernstrom01.
- J. Wang et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/JingWang01.
- R. Wang et al. *CoRR*, 2023. URL <http://arxiv.org/abs/2308.14610v2>.
- R. Wang et al. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(1):852–860, 2025a. doi: 10.1609/aaai.v39i1.32069. URL <http://dx.doi.org/10.1609/aaai.v39i1.32069>.
- S. Wang et al. *The Astronomical Journal*, 169(6):289, 2025b. doi: 10.3847/1538-3881/adc57b. URL <http://dx.doi.org/10.3847/1538-3881/adc57b>.
- D. Wright, K. Adámek, and W. Armour. *CoRR*, 2024. URL <http://arxiv.org/abs/2403.03091v1>.
- Z. Xiao, K. Kreis, and A. Vahdat. *arXiv preprint arXiv:2112.07804*, 2021. doi: 10.48550/arXiv.2112.07804. URL <https://doi.org/10.48550/arXiv.2112.07804>.
- D. Xu et al. *The Astrophysical Journal*, 980(1):52, 2025. doi: 10.3847/1538-4357/ada8a0. URL <http://dx.doi.org/10.3847/1538-4357/ada8a0>.
- L. Yang et al. *ACM Computing Surveys*, 56(4):1–39, Nov. 2023. ISSN 1557-7341. doi: 10.1145/3626235. URL <https://doi.org/10.1145/3626235>.
- S. Yatawatta. *CoRR*, 2023. URL <http://arxiv.org/abs/2301.03933v1>.
- L. Zhang et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/LeZhang01.
- X. Zhang, I. Cognard, and N. Dobigeon. *Astronomy and Computing*, 47:100822, 2024. ISSN 2213-1337. doi: <https://doi.org/10.1016/j.ascom.2024.100822>. URL <https://www.sciencedirect.com/science/article/pii/S2213133724000374>.
- Z. Zhang, B. Jiang, and Y. Zhang. *Publications of the Astronomical Society of the Pacific*, 134 (1036):064503, 2022. doi: 10.1088/1538-3873/ac67b1. URL <http://dx.doi.org/10.1088/1538-3873/ac67b1>.

J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232. IEEE, Oct. 2017. doi: 10.1109/iccv.2017.244. URL <https://doi.org/10.1109/iccv.2017.244>.