



**HAL**  
open science

# Unsupervised Deep Generative Models for Anomaly Detection in Neuroimaging: A Systematic Scoping Review

Youwan Mahé, Elise Bannier, Stéphanie Leplaideur, Elisa Fromont, Francesca Galassi

## ► To cite this version:

Youwan Mahé, Elise Bannier, Stéphanie Leplaideur, Elisa Fromont, Francesca Galassi. Unsupervised Deep Generative Models for Anomaly Detection in Neuroimaging: A Systematic Scoping Review. 2025. <hal-05304920v2>

**HAL Id: hal-05304920**

**<https://inria.hal.science/hal-05304920v2>**

Preprint submitted on 8 Mar 2026

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

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



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

# Unsupervised Deep Generative Models for Anomaly Detection in Neuroimaging: A Systematic Scoping Review

Youwan Mahé,<sup>1,2\*</sup> Elise Banner<sup>1,3</sup>, Stéphanie Leplaideur<sup>1,4,5</sup>,  
Elisa Fromont<sup>6§</sup> and Francesca Galassi,<sup>1\*§</sup>

<sup>1</sup>Univ Rennes, Inria, CNRS, Inserm, IRISA UMR 6074, Empenn, Rennes, France,

<sup>2</sup>Siemens Healthineers, Courbevoie, France

<sup>3</sup>CHU Rennes, Radiology Department, Rennes, France

<sup>4</sup>CHU Rennes, Physical Medicine and Rehabilitation Department, Rennes, France

<sup>5</sup>Centre de Kerpape, Ploemeur, France

<sup>6</sup>Univ Rennes, Inria, CNRS, IRISA, Rennes, France

§ These authors contributed equally as co-last authors.

\*Correspondence: [youwan.mahe, francesca.galassi]@inria.fr

March 7, 2026

**Abstract:** Unsupervised anomaly detection (UAD) based on deep generative modelling has been increasingly explored for identifying pathological brain abnormalities without requiring voxel-level annotations. By learning the distribution of healthy anatomy and generating pseudo-healthy reconstructions, these methods aim to localise deviations in a pathology-agnostic manner. Despite rapid methodological development - from autoencoders and variational autoencoders to generative adversarial networks and diffusion-based models - a structured synthesis of their application in structural neuroimaging is lacking.

We conducted a PRISMA-ScR-guided scoping review of studies published between January 2018 and December 2025 that applied unsupervised deep generative models to anomaly detection in brain MRI (and, less frequently, CT). Thirty-three studies met inclusion criteria. Methods were categorised by architectural family, and reported performance was synthesised across major pathology groups, with segmentation (Dice) and detection metrics (AUROC,

AUPRC) disaggregated by evaluation level (voxel, slice, subject). For transparency, we also summarised dataset characteristics, dimensionality (2D vs. 3D), and thresholding strategies.

Reported performance varied substantially across studies and pathologies. Dice scores generally ranged from approximately 0.30 to 0.75, while voxel-level AUROC values spanned roughly 0.58 to 0.97. Performance was typically higher for large focal lesions, such as brain tumours, than for small or diffuse abnormalities including white-matter hyperintensities. Across overlapping performance ranges, no architectural family consistently dominated across pathologies. Reported metrics were strongly influenced by evaluation level, thresholding strategy, and dataset composition, limiting direct cross-study comparability.

Overall, unsupervised generative approaches demonstrate potential for pathology-agnostic anomaly localisation, particularly in settings where annotated data are scarce. However, methodological heterogeneity, limited external validation, and sensitivity to dataset characteristics remain important challenges. Emerging paradigms - including anatomy-aware modelling, diffusion-based frameworks, and alternative normative evaluation metrics - seek to address these limitations and improve robustness and clinical relevance.

**Keywords:** Unsupervised Anomaly Detection (UAD), Deep Generative Modelling, Neuroimaging, Magnetic Resonance Imaging (MRI)

## 1 Introduction

Brain magnetic resonance imaging (MRI) plays a central role in the diagnosis, monitoring, and prognosis of neurological disease. Structural MRI enables high-resolution *in vivo* visualisation of abnormalities such as tumours, vascular infarcts, white-matter hyperintensities, demyelinating lesions, traumatic injury, and regional atrophy (Crimi & Bakas, 2021). Their appearance varies substantially across imaging sequences and pathology types (Vemuri et al., 2022; Villanueva-Meyer et al., 2017), complicating reliable delineation. Accurate segmentation is essential for deriving quantitative biomarkers - including lesion load, spatial distribution, and longitudinal volumetric change - that support clinical decision-making and research studies. However, manual detection and segmentation by expert radiologists remain time-consuming, require specialised expertise, and are subject to inter- and intra-rater variability (Darrault et al., 2025; Walsh et al., 2023). The development of robust automated tools is therefore a major priority.

Before the rise of deep learning, automated lesion detection relied on classical image-processing and sta-

tistical modelling approaches, including Gaussian mixture models, fuzzy c-means clustering, and Markov random fields (García-Lorenzo et al., 2013; Gonzalez & Woods, 2007). While effective in controlled settings, these techniques required expert-defined features and were sensitive to anatomical variability and acquisition differences (Xu et al., 2024). The advent of deep learning transformed medical image analysis by enabling multi-scale feature learning directly from data (Litjens et al., 2017). In neuroimaging, convolutional architectures such as U-Net (Ronneberger et al., 2015) and its derivatives, including nnU-Net (Isensee et al., 2024), have achieved state-of-the-art performance on curated benchmarks. Large public datasets such as BraTS for brain tumours (Menze et al., 2015) and MSSEG for multiple sclerosis (Commowick et al., 2021) have facilitated this progress.

Despite these advances, supervised models exhibit important limitations. Their performance depends critically on large voxel-level annotated datasets, which are costly to produce and often unavailable in rare or heterogeneous diseases (Lee et al., 2022). Moreover, supervised approaches operate under a closed-set assumption: they are trained to recognise predefined lesion categories and may struggle under distribution shift caused by variations in pathology subtype, patient population, or acquisition protocol (Ghafoorian et al., 2017). As a result, accuracy may deteriorate substantially outside curated benchmarks.

Unsupervised anomaly detection (UAD) offers a complementary paradigm. Rather than learning to segment predefined abnormalities, UAD methods aim to model the distribution of healthy anatomy and detect deviations from this normative representation. In neuroimaging, this is typically achieved by training generative models exclusively on healthy brain scans. At inference, a pathological image is mapped toward a pseudo-healthy reconstruction, and spatial discrepancies between the input and its reconstruction highlight candidate abnormalities. This reconstruction-based strategy enables pathology-agnostic detection and may be particularly valuable in settings where annotations are scarce, diseases are rare, or lesion phenotypes are heterogeneous. Recent open-set evaluations, such as the NOVA benchmark (Bercea, Li, et al., 2025), have demonstrated substantial performance drops for state-of-the-art models on rare and heterogeneous pathologies, underscoring the need for distribution-modelling approaches that do not rely on predefined lesion categories.

Over the past decade, deep generative modelling has evolved rapidly. Early work employed autoencoders and variational autoencoders (VAEs) to learn compact representations of healthy anatomy (Zimmerer et al., 2019). Generative adversarial networks (GANs) were subsequently introduced to improve reconstruction realism and boundary sharpness (Schlegl et al., 2019). More recently, denoising diffusion probabilistic models (DDPMs) (Ho et al., 2020) and related continuous-time generative frameworks have demonstrated strong capacity for modelling complex anatomical variability and producing high-fidelity pseudo-healthy reconstructions (Bercea et al., 2023; Pinaya, Tudosi, et al., 2022). These methodolog-

ical advances have expanded the landscape of unsupervised neuroimaging anomaly detection.

Several reviews have addressed subsets of this literature, including autoencoder and diffusion-based anomaly detection (Baur, Denner, et al., 2021; Bercea, Cattin, et al., 2025), GAN applications in medical imaging (Wang et al., 2023), and broader surveys of generative models in healthcare (Pang et al., 2021; Tschuchnig & Gadermayr, 2022). However, to our knowledge, no prior review has systematically synthesised unsupervised deep generative models for structural neuroimaging from a pathology-stratified perspective while explicitly disaggregating segmentation and detection metrics, evaluation levels (voxel, slice, subject), thresholding strategies, and dataset characteristics. Reported performance varies substantially depending on lesion type, evaluation protocol, and data composition, complicating cross-study comparison and interpretation.

In this PRISMA-ScR-guided scoping review, we systematically analyse unsupervised deep generative models for anomaly detection in structural neuroimaging published between 2018 and December 2025. We categorise methods into four principal architectural families - autoencoders, variational autoencoders, generative adversarial networks, and diffusion-based models - and synthesise their performance across major pathology groups, including brain tumours, stroke, multiple sclerosis, and white-matter hyperintensities. Beyond architectural comparison, we examine the influence of dimensionality (2D vs. 3D), dataset composition, thresholding strategies, and evaluation level on reported metrics. Finally, we discuss methodological limitations of reconstruction-based approaches and highlight emerging paradigms aimed at improving robustness, interpretability, and clinical relevance.

## 1.1 Review questions

This scoping review addresses the following questions:

- (i) Which unsupervised deep generative model families (AE, VAE, GAN, diffusion/flow) have been applied to anomaly detection and/or segmentation in structural neuroimaging since 2018?
- (ii) How are these methods evaluated (evaluation level, metrics, and thresholding strategy), and how do reported results vary across major pathology groups?
- (iii) What methodological design variations (e.g., 2D vs. 3D processing, patching or masking strategies, pretraining) are described in the literature, and how are they reported in relation to performance?
- (iv) What emerging paradigms aim to address current limitations of reconstruction-based anomaly detection?

## 2 Methods

### 2.1 Background: generative modelling for unsupervised anomaly detection

Generative models aim to learn the distribution of healthy brain images, thereby capturing normative anatomical variability. When an explicit probabilistic formulation is available, this can be written as approximating the data distribution  $p_{\text{data}}(\mathbf{x})$  with a distribution  $p_{\theta}(\mathbf{x})$  induced by a parameterised mapping. In neuroimaging UAD, models are trained on healthy brain images only, so the learned distribution represents normative anatomy rather than pathology.

Many generative approaches introduce a latent variable  $\mathbf{z} \in \mathbb{R}^d$ , drawn from a simple prior distribution  $p(\mathbf{z})$ , typically a standard multivariate Gaussian  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ , and a neural network  $G_{\theta}$  that maps latent codes to image space:

$$\mathbf{z} \sim p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad \mathbf{x} = G_{\theta}(\mathbf{z}), \quad \mathbf{x} \sim p_{\theta}(\mathbf{x}). \quad (1)$$

When an explicit encoder  $E_{\phi}$  is available (e.g., autoencoders and VAEs), a test image  $\mathbf{x}_{\text{test}}$  is mapped to latent space as

$$\mathbf{z} = E_{\phi}(\mathbf{x}_{\text{test}}), \quad \hat{\mathbf{x}} = G_{\theta}(\mathbf{z}). \quad (2)$$

For models without an encoder (e.g., classical GANs or some diffusion-based methods), a latent representation must instead be obtained through an inversion or optimisation procedure:

$$\mathbf{z}^* = \arg \min_{\mathbf{z}} \text{dist}(\mathbf{x}_{\text{test}}, G_{\theta}(\mathbf{z})), \quad \hat{\mathbf{x}} = G_{\theta}(\mathbf{z}^*), \quad (3)$$

where  $\text{dist}(\cdot, \cdot)$  denotes a similarity or dissimilarity measure (e.g.,  $\ell_1$ ,  $\ell_2$ , SSIM).

Because pathological patterns are absent from the healthy training distribution, they are typically not faithfully reconstructed in the pseudo-healthy image  $\hat{\mathbf{x}}$ . An anomaly map is therefore defined as

$$\mathbf{r} = \text{dist}(\mathbf{x}_{\text{test}}, \hat{\mathbf{x}}), \quad (4)$$

highlighting spatial deviations from normative anatomy without requiring voxel-level annotations.

Despite substantial architectural differences between autoencoders, variational autoencoders, generative adversarial networks, and diffusion models, this reconstruction-based principle underpins the majority of unsupervised generative approaches reviewed in this work. Subsequent sections compare how different

architectures, dimensionalities, and training strategies affect the quality of pseudo-healthy reconstructions and the reliability of the resulting anomaly maps.

## 2.2 Information sources and search strategy

This review was conducted in accordance with the PRISMA-ScR (2018) guidelines for scoping reviews (Tricco et al., 2018). We searched PubMed, Web of Science, ScienceDirect, Springer Link, IEEE Xplore, and ArXiv up to 17 December 2025, including early-access preprints. Boolean queries combined terms related to *unsupervised anomaly detection*, *neuroimaging* (MRI or CT), and *deep learning*, as summarised in Table 1. Reference lists of relevant articles and reviews were also screened. Searches were limited to articles published in English.

Table 1: Boolean queries used for database searching

Database	Query	Date
PubMed	(Anomaly AND Unsupervised) AND Brain) AND (MRI OR CT) AND (Machine Learning OR Deep Learning)	Dec 17 2025
Web Of Science	((TS=Anomaly) AND (TS=unsupervised) AND (TS=brain)) AND ((TS=MRI) OR (TS=CT)) AND ((TS=Machine Learning) OR (TS=Deep Learning))	Dec 17 2025
Science Direct	<i>Unsupervised Anomaly Brain MRI CT Machine Learning</i> Filter : Research Article	Dec 17 2025
ArXiv	Unsupervised AND Anomaly AND Brain AND "Deep learning", in Computer Science (cs)	Dec 17 2025
IEEE Xplore	Unsupervised AND Anomaly AND Brain AND "Deep Learning"	Dec 17 2025
Springer Nature Link	(Anomaly AND Unsupervised AND Brain) AND (MRI OR CT) AND (Machine Learning OR Deep Learning) AND ("Conference Paper" OR "Research Article") AND ("Computer Vision" OR "Machine Learning")	Dec 17 2025

## 2.3 Eligibility and screening

Studies were eligible if they:

- proposed or evaluated an unsupervised deep learning method for anomaly detection and/or lesion segmentation in structural neuroimaging (MRI or CT), trained exclusively on healthy data to model normative anatomy (i.e., without access to pathological labels);
- reported at least one quantitative evaluation metric relevant to detection (e.g., AUROC, AUPRC) or segmentation (e.g., Dice);
- used real human imaging data from public datasets or institutional cohorts.

Unsupervised learning can be characterised in multiple ways. A narrow definition restricts it to methods that lack any explicit label-driven training objective, thereby excluding self-supervised approaches (Dosovitskiy et al., 2014). In this review, we adopt a pragmatic definition aligned with anomaly detection in medical imaging: we include methods trained without using real pathological labels (neither voxel-level segmentations nor subject-/case-level diagnostic labels) in the training objective. This encompasses approaches that use proxy or self-supervised objectives to model the distribution of healthy data (e.g., masked autoencoding, lesion synthesis, or discriminative proxy tasks) (I. Goodfellow et al., 2016). In Table 3, methods employing such proxy or self-supervised objectives are denoted with a  $\mathcal{P}$ .

We excluded studies that employed rule-based or non-deep learning approaches; fully supervised or semi-supervised methods requiring pathological annotations; non-generative methods (except where discussed contextually); non-neuroimaging applications; animal-only or synthetic-only data; review papers; and non-research formats (e.g., abstracts, editorials). For transparency, self-supervised non-generative and hybrid approaches are discussed in a separate contextual subsection but are not included in the main quantitative synthesis because they do not satisfy the generative modelling inclusion criteria. Although some applications use deep generative models for functional or diffusion imaging (e.g., FDG-PET, diffusion tensor imaging, or functional MRI) (Hassanally et al., 2024; Muñoz-Ramírez et al., 2022; Sathyanarayanan & Muthukumaravel, 2025; Solal et al., 2023), we restricted the primary synthesis to structural imaging to ensure comparability across studies. Similarly, we excluded studies that did not report voxel-, slice-, or subject-level quantitative performance metrics.

Search results were imported into Rayyan (Ouzzani et al., 2016) for duplicate removal and blinded screening by two reviewers. Titles and abstracts were screened first, followed by full-text assessment of potentially relevant studies. Disagreements were resolved through discussion until consensus was reached.

Bibliographic records were retrieved using the export functions of each database (CSV or BibTeX) and imported into Rayyan for record matching and management. For arXiv, where no export function is available through the web interface, we used the API via custom Python scripts to obtain publication metadata. Full-text articles were accessed through institutional subscriptions where available.

Data extraction was performed independently by two reviewers using a structured extraction form developed for this review. Extracted variables included: publication year, model architecture, imaging modality, dimensionality (2D or 3D), training and test datasets, pathology type, evaluation level (voxel, slice, or subject), thresholding strategy, and reported quantitative performance metrics (Dice, AUROC, AUPRC). Extracted data were tabulated and used for qualitative and quantitative synthesis.

## 2.4 Risk of bias assessment

Risk of bias was assessed with a focus on methodological quality. Guided by PRISMA recommendations for systematic reviews (Page et al., 2021), two reviewers (YM, FG) independently screened and appraised all records in Rayyan using blind mode. Disagreements were resolved by discussion until consensus. As no validated risk-of-bias tool exists for anomaly detection in medical imaging, we applied a structured checklist covering dataset characteristics (e.g., public availability, diversity of pathologies), as well as reproducibility and transparency (e.g., code and data availability). For studies raising concerns about reporting integrity, we additionally screened them with the *Problematic Paper Screener* (Cabanac et al., 2022) and documented outcomes.

# 3 Results

## 3.1 Study selection

The initial search identified 574 records. After duplicate removal and title/abstract screening, 527 records were excluded according to predefined eligibility criteria (e.g., out of scope, non-neuroimaging, non-deep learning, rule-based methods, review articles). Forty-seven full-text articles were assessed for eligibility, of which 14 were excluded (e.g., absence of structural imaging, lack of quantitative metrics, non-research format). A total of 33 studies were included in the final synthesis (Fig. 1). The included studies were published between January 2018 and 17 December 2025 (including early-access articles), providing an eight-year overview of unsupervised deep generative models applied to structural neuroimaging.

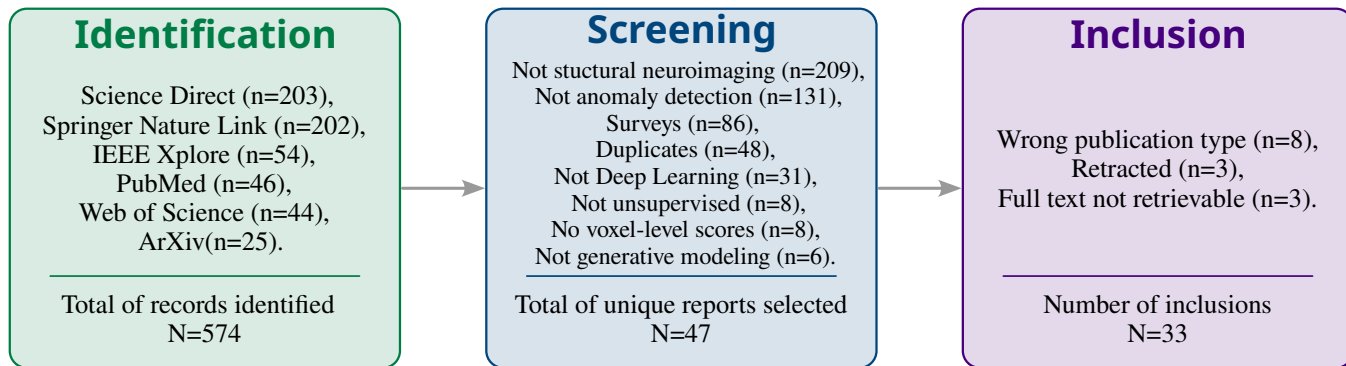


Figure 1: PRISMA flow diagram for scoping reviews, including database and register searches.

## 3.2 Study characteristics

Across the 33 included studies, four main architectural families were identified: autoencoders (AEs), variational autoencoders (VAEs), generative adversarial networks (GANs), and diffusion-based models. Although standard AEs do not explicitly define a probabilistic generative model, they are widely used as reconstruction-based normative baselines and are therefore included in this synthesis. Proxy or self-supervised pretext objectives were included when used to support pseudo-healthy reconstruction (e.g., masking/inpainting or structured corruption). In total, the 33 studies comprised 9 autoencoder-based approaches, 9 variational autoencoders, 4 GAN-based methods, and 11 diffusion-based models.

MRI constituted the primary imaging modality (T1-w, T1-w Gd, T2-w, and FLAIR), with occasional use of computed tomography. Most methods processed 2D slices, although an increasing subset employed volumetric 3D architectures, particularly among recent AE/VAE studies and selected GAN and diffusion approaches. Although some generative models have been applied to non-structural modalities such as FDG-PET for Alzheimer’s disease detection (Hassanly et al., 2024; Solal et al., 2023), this review focused on structural imaging to ensure methodological comparability across studies.

## 3.3 Evaluation metrics

Most studies reported either segmentation or detection performance. Segmentation accuracy was assessed using the Dice similarity coefficient (DSC or Dice score), defined as

$$\text{DSC} = \frac{2|X \cap Y|}{|X| + |Y|}, \quad (5)$$

where  $X$  is the predicted segmentation and  $Y$  the reference annotation. Dice scores depend not only on the quality of the anomaly map but also on the thresholding strategy used to binarise residual maps.

Threshold selection varied substantially across studies. In several cases, thresholds were retrospectively chosen to maximise Dice on the test set, yielding optimistic upper-bound estimates that may not reflect clinically deployable performance. Dice is also sensitive to lesion size and can disproportionately penalise small or sparse abnormalities. Detection was formulated as a binary classification problem (pathological vs. non-pathological) and quantified using AUROC and AUPRC. In highly imbalanced voxel-level tasks, AUPRC is often more informative than AUROC because it better reflects precision under class imbalance. Detection metrics were computed at different evaluation levels (voxel/pixel, slice, subject). Because these correspond to distinct classification tasks, values are not directly comparable across evaluation regimes.

*Comparability across studies.* Substantial heterogeneity exists in dataset composition, preprocessing pipelines, anomaly-map post-processing, thresholding strategies, and evaluation levels. Accordingly, reported performance values should be interpreted as descriptive indicators within each experimental setting rather than direct head-to-head comparisons across architectures.

### 3.4 Datasets and pathologies

Performance metrics were reported across a range of publicly available and institutional neuroimaging datasets. Three pathology groups predominated in the literature: brain tumours, multiple sclerosis (MS)/white-matter hyperintensities (WMH), and stroke. Fewer studies addressed neurodegenerative conditions (e.g., Alzheimer’s disease, Parkinson’s disease), traumatic brain injury, neonatal encephalopathy, or healthy ageing. Representative examples of central axial slices from the principal pathological datasets are shown in Figure 2. A structured overview of dataset characteristics - including cohort size, imaging modality, and documented biases - is provided in Table 2. Because dataset composition, acquisition protocols, and pathology prevalence vary substantially across studies, these factors should be considered when interpreting reported performance metrics.

**Brain tumours.** The majority of tumour studies relied on the *Brain Tumor Segmentation* (BraTS) challenge datasets (Menze et al., 2015), which provide multi-sequence MRI including T1-weighted (T1-w), contrast-enhanced T1-w (T1-w Gd), T2-weighted (T2-w), and FLAIR images. As detailed in Table 2, BraTS has expanded substantially across releases: BraTS 2012 included 50 cases, BraTS 2015 increased to 253 cases (with a strong high-grade glioma predominance), BraTS 2017-2020 ranged from 477 to 660 cases with expert-revised annotations across multiple centres, and the most recent BraTS 2021-2023 (GLI) editions comprise approximately 2,040 multi-institutional cases (Baid et al., 2021). A smaller number of studies used alternative cohorts, including the neuroimaging dataset of brain tumour patients from Pernet et al. (2016) (22 cases) and the Brain Tumor Dataset (BTD) from Cheng (2017) (233 cases),




Figure 2: Central axial slices from a healthy brain (IXI), a brain tumour case (BraTS), a chronic stroke case (ATLAS v2.0), and a multiple sclerosis case (MSSEG).

which are comparatively smaller and more homogeneous than BraTS. Gliomas, the primary pathology represented in these datasets, typically produce large focal lesions with substantial mass effect. Lesion morphology varies according to tumour grade, and progressive growth within the cranial cavity frequently distorts surrounding anatomical structures. Sequence selection is clinically motivated: T2-w and FLAIR emphasise water content and oedema, whereas T1-w Gd highlights blood-brain barrier disruption through contrast enhancement (Menze et al., 2015).

**Stroke.** Stroke lesions were primarily represented by the *Ischaemic Stroke Lesion Segmentation* (ISLES) 2015 dataset and the *Anatomical Tracings of Lesions After Stroke* (ATLAS) datasets (versions 1.1/1.2 and 2.0) (Hernandez Petzsche et al., 2022; Liew et al., 2021), as summarised in Table 2. ISLES 2015 comprises two subsets: the Sub-acute Ischemic Stroke Lesion Segmentation (SISS) cohort (64 subjects imaged with T1-w, T2-w, FLAIR, and DWI) and the Stroke Perfusion Estimation (SPES) cohort (50 subjects imaged with T1-w, T2-w, DWI, and perfusion modalities). These cohorts target acute and sub-acute ischaemic stroke, where diffusion-weighted imaging (DWI) is most sensitive to early infarction. As stroke progresses, diffusion signal characteristics evolve and lesion conspicuity may decrease, making delineation on structural sequences such as FLAIR more challenging (Hernandez Petzsche et al., 2022). ATLAS v1.1/1.2 includes 229-304 chronic stroke cases acquired as T1-weighted MRI, while ATLAS v2.0 comprises 955 manually segmented T1-w MRI cases collected from 33 cohorts worldwide (Liew et al., 2021). A key limitation of ATLAS is its reliance on T1-w imaging alone, where chronic infarcts often resemble cerebrospinal fluid and are accompanied by tissue loss and structural deformation. Overall, lesion stage (acute vs. chronic), available contrasts (multi-sequence vs. T1-w only), and lesion appearance vary

substantially across datasets; consequently, reported anomaly detection and segmentation performance is heterogeneous and not directly comparable across studies.

**Multiple sclerosis and white-matter hyperintensities.** Although multiple sclerosis (MS) and white-matter hyperintensities (WMH) arise from distinct pathological processes, they are described jointly here because they share similar imaging characteristics relevant to anomaly detection: small, sparse, FLAIR-hyperintense white-matter lesions with low lesion-to-brain volume ratios. As summarised in Table 2, MS was primarily evaluated using the *ISBI 2015 Longitudinal MS Lesion Segmentation* dataset (ISBI 2015) (Carass et al., 2017), the *Multiple Sclerosis Segmentation* dataset (MSSEG) (Commowick et al., 2018, 2021), and the *MSLUB* dataset (Lesjak et al., 2017). MSSEG (MICCAI 2016) comprises 53 subjects with four MRI contrasts (T1-w, T1-w Gd, proton density, and FLAIR), while ISBI 2015 includes 82 cases and is subject to mono-site bias. MSLUB is a smaller mono-site cohort of 30 subjects with T2-w, FLAIR, and proton density imaging. WMH were primarily evaluated using the *WMH Segmentation Challenge* dataset (WMH 2017; 170 subjects; T1-w and FLAIR), which is enriched for older participants and spans multiple scanners, introducing age- and scanner-related heterogeneity (Kuijf et al., 2022). Unlike tumours or many stroke lesions, MS plaques and WMH are typically small, sparse, and spatially distributed. They are most conspicuous on FLAIR and T2-w MRI, where they appear as focal white-matter hyperintensities. These properties contribute to substantial variability in segmentation and anomaly detection performance across studies and remain a major bottleneck for reconstruction-based methods.

**Traumatic brain injury.** Traumatic brain injury (TBI) was represented primarily by the *CENTER-TBI* cohort (Table 2), a large multicentre European dataset comprising  $\sim 4,500$  subjects, with imaging dominated by computed tomography (CT) (Steyerberg et al., 2019). On CT, acute TBI may present as focal hyperdense haemorrhagic lesions (e.g., epidural or subdural haematomas) as well as smaller haemorrhagic foci; however, diffuse axonal injury is often subtle or not directly visible on standard CT, which can limit sensitivity for detecting microscopic injury patterns (Currie et al., 2015).

**Neurodegeneration and other conditions.** A smaller subset of included studies leveraged neurodegenerative cohorts, most often through the *Alzheimer’s Disease Neuroimaging Initiative* (ADNI) (Beckett et al., 2015) and the *Open Access Series of Imaging Studies* (OASIS-3) (LaMontagne et al., 2019), as summarised in Table 2. ADNI is a large multi-phase public repository including structural MRI and PET imaging from cognitively normal individuals, subjects with mild cognitive impairment, and patients with Alzheimer’s disease ( $> 2,000$  participants across phases). OASIS-3 is a longitudinal, single-site cohort

comprising 1,098 participants (605 cognitively normal and 493 at-risk or demented individuals), with multi-contrast MRI including T1-w, T2-w, FLAIR, DWI, SWI, ASL, and resting-state fMRI.

**Healthy control datasets.** Across studies, normative training data were sourced either from dedicated healthy cohorts (e.g., IXI, UK Biobank, HCP) or from healthy subsets extracted from mixed pathology datasets (denoted H in Table 3). A smaller number of studies relied on in-house institutional cohorts. Explicit reporting of screening procedures and contamination control was inconsistent. Frequently used resources included the IXI dataset<sup>1</sup> (approximately 600 healthy volunteers scanned across three London hospitals; T1-w, T2-w, and proton density MRI), which was widely employed for training autoencoder, VAE, and diffusion models. Large population-based cohorts were also used, including the UK Biobank (UKB), which provides multi-modal MRI for over 50,000 participants, and the Human Connectome Project (HCP; Young Adult cohort,  $N \sim 1,200$ ; ages 22–35 years) (Van Essen et al., 2012). While UKB offers unprecedented scale, it is subject to healthy volunteer and geographical biases (Littlejohns et al., 2020). HCP, although high-quality, reflects a restricted young adult demographic. Smaller normative datasets included the Cambridge Centre for Ageing and Neuroscience (Cam-CAN;  $N \sim 650$ ) (Shafto et al., 2014), the Neurofeedback Skull-stripped (NFBS) dataset (125 T1-w scans) (Puccio et al., 2016), and institutional in-house cohorts. A single diffusion-based study used a healthy subset of the FastMRI+ dataset (1,001 axial brain MRI scans) (Zhao et al., 2022). The *Medical Out-of-Distribution* (MOOD) challenge (Zimmerer et al., 2020) dataset was used once. Notably, MOOD relies on synthetic anomalies injected into healthy HCP scans, which may introduce synthetic bias and inflate detection performance relative to real-world pathology. Although ADNI and OASIS-3 contain both healthy and pathological subjects, in the reviewed literature, they were primarily used as sources of cognitively normal control scans for training normative models rather than as primary anomaly detection benchmarks. It is important to note that ADNI controls exhibit a documented healthy volunteer bias (Gianattasio et al., 2021), and OASIS-3 reflects single-site demographic characteristics, which may influence normative modelling. A single study reported combining smaller healthy subsets of the *individual brain charting* (IBC), a high-resolution fMRI dataset for cognitive mapping (Pinho et al., 2018), and from a multi-parametric neuroimaging reproducibility study referred here to as *Kirby* (Landman et al., 2011)

### 3.5 Synthesis by architecture

To provide a structured overview, we synthesised findings by architectural family and pathology. For segmentation tasks, we report Dice similarity coefficient (Dice) scores; for detection tasks, we report

---

<sup>1</sup><http://brain-development.org/ixi-dataset/>

AUROC and AUPRC, always specifying the evaluation level (voxel/pixel, slice, or subject) to avoid misleading cross-study comparisons. Study characteristics and experimental setups are summarised in four complementary tables. Table 2 describes datasets (release, size, modality, population, and known biases). Table 3 summarises methods (architectural family, training/test datasets, imaging modality, and input dimensionality/processing strategy). Quantitative results are reported in Table 4 for detection and Table 5 for segmentation.

Table 2: Comprehensive overview of neuroimaging datasets and their characteristics.

Dataset (Release)	Modality	Pathology	Population	Cases (N)	Known Biases & Characteristics
<b>Tumour</b>					
BraTS '12	MRI (4 seqs)	Glioma	Adult	50	Small sample size.
BraTS '15	MRI (4 seqs)	Glioma	Adult	253	Skewed to HGG.
BraTS '17-20	MRI (4 seqs)	Glioma	Adult	477-660	Expert-revised labels; pre-operative scan bias, n=19 centres.
BraTS '21-23(GLI)	MRI (4 seqs)	Glioma	Adult	2,040	pre-operative scan bias; multi-site diversity.
Pernet et al. (2016)	MRI (3 seqs)	Tumor	Adult	22	Small sample size, pre-operative.
Cheng (2017) (BTD)	MRI (T1w)	Tumor	Adult	233	Monocentric; private.
<b>Stroke</b>					
ISLES 2015 (SISS)	MRI (4 seqs)	Sub-acute isch. Stroke	Adult	64	Mostly monocentric; ischemic stroke bias.
ISLES 2015 (SPES)	MRI (4 seqs)	acute isch. Stroke	Adult	50	monocentric; ischemic stroke bias.
ATLAS v1.1/1.2	MRI (T1w)	Chronic Stroke	Adult	229-304	Modality bias (T1w only), multicentric.
ATLAS v2.0	MRI (T1w)	Chronic Stroke	Adult	955	Modality bias (T1w only), multicentric.
<b>Multiple Sclerosis (MS) and White Matter Hyperintensities (WMH)</b>					
ISBI 2015	MRI (4 seqs)	Multiple Sclerosis	Adult	82	Mono-site bias (JHU, USA).

*Continued on next page*

Dataset (Release)	Modality	Pathology	Population	Cases (N)	Known Biases & Characteristics
MSSEG	MRI (4 seqs)	Multiple Sclerosis	Adult	53	Small sample size.
MSLUB	MRI (3 seqs)	Multiple Sclerosis	Adult	30	Mono-site bias (UMCL, Slovenia).
WMH 2017	MRI (T1w, FLAIR)	White Matter Hyp.	Adult	170	Age bias (elderly); scanner bias (3 sites).
<b>Dementia</b>					
ADNI (1/2/3)	MRI, PET	Alzheimer	Adult	>2,000	Clinical trial participants (elderly).
OASIS-3	MRI, PET	Alzheimer	Adult	1,098	Mono-site bias; longitudinal study characteristics.
<b>Healthy Controls</b>					
UK Biobank	MRI (Various)	Healthy	Adult	~50,000	Volunteer bias (healthier/wealthier demographics).
HCP (Young)	MRI (T1w, T2w)	Healthy	Young Adult	~1,200	Age bias (22-35y); high-quality bias.
Cam-CAN	MRI (T1w, T2w)	Healthy	Adult	~650	Mono-site bias (Cambridge); lifestyle homogeneity.
IXI	MRI (T1w, T2w, PD)	Healthy	Adult	~600	Multi-site (3 London hospitals); demographic bias.
NFBS	MRI (T1w)	Healthy	Young Adult	125	Age bias (21-45), healthy/addict/phobic subjects.
Kirby	MRI (Various)	Healthy	Adult	21	Mono-site bias, Small sample size.
IBC	MRI (Various)	Healthy	Adult	12	Mono-site bias, Small sample size.
<b>Others</b>					

*Continued on next page*

Dataset (Release)	Modality	Pathology	Population	Cases (N)	Known Biases & Characteristics
Center-TBI	CT	Traumatic Brain Injury	Adult	~4,500	Diverse injury mechanisms, multicentric (Europe).
MOOD 2020	MRI (Brain)	Synthetic	Adult	~800	Synthetic bias (anomalies injected into healthy data from young HCP).
fastMRI+	MRI (T1w, T2w, FLAIR)	Not reported	Clinical Population	1,001	Pathology/Healthy mix, axial only.

Table 3: Comparison of unsupervised anomaly detection methods and their experimental setups. MS: Multiple Sclerosis; TBI: Traumatic Brain Injury; WMH: White Matter Hyperintensities. \* denotes the presence of a working link to a GitHub repository. <sup>H</sup> denotes the use of healthy samples drawn from the same dataset as the pathological test set. <sup>P</sup> denotes the use of proxy or self-supervised objectives during training (e.g., masking, synthetic corruption). Specific proxy mechanisms are described in the corresponding architecture subsections. Input Dimensionality refers to the sampling or patching strategy used during model training.

Method	Train Dataset	Test Dataset	Imaging Modality	Input Dimensionality
<b>Autoencoders</b>				
Baur et al., 2018	In-house	In-house (MS)	T1w, FLAIR	Full 2D axial
Baur, Wiestler, et al., 2021	In-house	In-house (Tumour)	FLAIR	Full 3D
Ghorbel et al., 2023	OASIS-3 <sup>H</sup>	BraTS'20 MSLUB	FLAIR FLAIR	Patch 2D axial
Kascenas et al., 2023 * <sup>P</sup>	BraTS'21 <sup>H</sup>	BraTS'21	T1w, T1w Gd, T2w & FLAIR	Full 2D axial
Luo et al., 2023 *	In-house <sup>H</sup>	In-house (Multiple)	CT	Full 3D
	In-house	BraTS'19 In-house (Stroke)	T2w T2w	

*Continued on next page*

Method	Train Dataset	Test Dataset	Imaging Modality	Input Dimensionality
Meissen et al., 2023 *	Cam-CAN <sup>2</sup> & MOOD <sup>3</sup>	In-house (MS) BraTS'20	T2w T1w	Partial 2D axial
Jiménez-García et al., 2024 <sup>P</sup>	IXI	BraTS'20	T1w, T1w Gd, T2w & FLAIR	Full 3D
Lu et al., 2024 <sup>P</sup>	BTD <sup>4H</sup>	BTD	T1w Gd	Full 2D patch
Qu et al., 2026	BraTS'21 <sup>H</sup>	BraTS'21	T1w, T1w Gd T2w & FLAIR	Full 2D axial
	In-house	In-house (Stroke)	CT	
<b>Variational Autoencoders</b>				
Sato et al., 2019	IXI	BraTS'17 ATLAS 1.1	T1w, T2w T1w	Full 2D axial
Uzunova et al., 2019 <sup>P</sup>	BraTS'15 <sup>H</sup>	BraTS'15	T1w Gd, T2w, FLAIR	Full 2D axial & 3D
Zimmerer et al., 2019 *	HCP	BraTS'17	not reported	Full 2D axial
Bengs et al., 2021 <sup>P</sup>	In-house	BraTS'19 ATLAS 1.1	T1w T1w	Full 3D
Lambert et al., 2021	ADNI <sup>H</sup> & IBC & Kirby	BraTS'18	FLAIR	Full 3D
Chatterjee et al., 2022 <sup>*P</sup>	IXI & MOOD <sup>H</sup>	MSSEG, WMH BraTS'17	FLAIR T1w Gd, T2w	Full 2D axial
Pinaya et al., 2022	UKB <sup>H</sup>	BraTS'18	FLAIR	Full 2D ax. & 3D

*Continued on next page*

<sup>2</sup>Shafiq et al., 2014

<sup>3</sup>Zimmerer et al., 2020

<sup>4</sup>Cheng, 2017

Method	Train Dataset	Test Dataset	Imaging Modality	Input Dimensionality
Lüth et al., 2023	HCP	MSLUB	FLAIR	Full 2D axial
		WMH	FLAIR	
		BraTS'17	T2w	
Wijanarko et al., 2024	IXI	ISLES 2015	T2w	Full 2D axial
		BraTS'20	T2w	
<b>Generative Adversarial Networks</b>				
Simarro et al., 2020	Center-TBI <sup>5H</sup>	Center-TBI	CT	Full 3D
Dey et al., 2021	BraTS'19 <sup>H</sup>	BraTS'19	not reported	Full 2D axial
		ISBI'15 <sup>H</sup>	T2w or FLAIR not reported	
Nguyen et al., 2021	NFBS	<i>Pernet et al.</i> <sup>6</sup>	T1w	Full 2D axial
Wu et al., 2021	BraTS'12/18 <sup>H</sup>	BraTS'12/18	T2w	Patch 3D
<b>Diffusion Models</b>				
Pinaya, Graham, et al., 2022	UKB <sup>H</sup>	BraTS'18	FLAIR	Full 2D axial
		MSLUB	FLAIR	
		WMH	FLAIR	
Wyatt et al., 2022 *	NFBS	<i>Pernet et al.</i> <sup>7</sup>	T1w	Partial 2D axial
Behrendt et al., 2023 *	IXI	BraTS'21	T2w	Patch 2D axial
		MSLUB	T2w	
Bercea et al., 2023 <sup>*P</sup>	IXI & FastMRI+	ATLAS 2.0	T1w	Partial 2D axial
Iqbal et al., 2023 <sup>*P</sup>	IXI	BraTS'21	T2w	Full 2D axial
				<i>Continued on next page</i>

<sup>5</sup>Steyerberg et al., 2019<sup>6</sup>Pernet et al., 2016<sup>7</sup>Pernet et al., 2016

Method	Train Dataset	Test Dataset	Imaging Modality	Input Dimensionality
Behrendt et al., 2024 *	IXI	MSLUB	T2w	Partial 3D
		BraTS'21	T2w	
		ATLAS 2.0	T1w	
		MSLUB	T2w	
Bercea et al., 2024 <sup>*P</sup>	IXI	WMH	T1w	Full 2D axial
		ATLAS 2.0	T1w	
Fontanella et al., 2024 <sup>P</sup>	BraTS'21 <sup>H</sup>	BraTS'21	T1, T1w Gd, T2w & FLAIR	Partial 2D axial
Kumar Trivedi et al., 2024 <sup>*P</sup>	IXI	WMH	T1w, FLAIR	Full 2D axial
		BraTS'20/21	T2w	
		MSLUB	T2w	
Bi et al., 2025 *	BraTS'23 <sup>H</sup>	BraTS'23	FLAIR	Partial 2D axial
Beizae et al., 2026 <sup>*P</sup>	BraTS'21 <sup>H</sup>	BraTS'21	T1w, T1w Gd, T2w & FLAIR	Partial 2D axial
	ATLAS 2.0 <sup>H</sup>	ATLAS 2.0	T1w	

Table 4: Performance comparison using detection metrics (AUROC/AUPRC). Methods are grouped by model family (Autoencoders, VAEs, GANs, Diffusion models, and related approaches), as defined in Table 3. Only studies that explicitly reported AUROC and/or AUPRC are included.

Method	Test Dataset	AUROC	AUPRC	Eval. level	Notes
<b>Autoencoders</b>					
<i>Continued on next page</i>					

Baur, Wiestler, et al., 2021	In-house	–	0.300	Voxel	Evaluated on MS
Ghorbel et al., 2023	BraTS'20	0.780	0.425	Pixel	
	MSLUB	0.886	0.203		
Kascenas et al., 2023 *	BraTS'21	–	0.833	Pixel	
	In-house (Multiple)	–	0.693	Pixel	
Luo et al., 2023 *	BraTS'19	0.844	0.741	Slice	
	In-house	0.807	0.705		Evaluated on Stroke
	In-house	0.858	0.731		Evaluated on MS
Meissen et al., 2023 *	BraTS'20	0.790	–	Pixel	
Jiménez-García et al., 2024	BraTS'20	0.838	–	Voxel	
Lu et al., 2024	BTD <sup>8</sup>	0.992	–	Slice	
Qu et al., 2026	BraTS'21	0.905	–	Slice	
	In-house	0.839	–	Slice	Evaluated on Stroke
<b>Variational Autoencoders</b>					
Sato et al., 2019	BraTS'17	0.582	–	Pixel	Evaluated on T1w
		0.788	–	Pixel	Evaluated on T2
		ATLAS 1.1	0.672	–	
Uzunova et al., 2019	BraTS'15	0.950	–	Pixel	Evaluated in 2D on T1w Gd
		0.950	–	Voxel	Evaluated in 3D on T1w Gd
		0.940	–	Voxel	Evaluated in 3D on fusion of T1w Gd, T2w & FLAIR
Zimmerer et al., 2019 *	BraTS'17	0.820	–	Pixel	
Bengs et al., 2021	BraTS'19	–	0.279	Slice	

*Continued on next page*

<sup>8</sup>Cheng, 2017

	ATLAS 1.1	–	0.256	
Pinaya et al., 2022	MSLUB	–	0.272	Pixel
		0.866	–	Slice
	BraTS'18	–	0.555	Pixel
	WMH	–	0.320	Pixel
Lüth et al., 2023	BraTS'17	0.942	0.380	Voxel
	ISLES 2015	0.898	0.186	
Wijanarko et al., 2024	BraTS'20	0.968	0.462	Voxel
<b>Generative Adversarial Networks</b>				
Simarro et al., 2020	Center-TBI	0.749	–	Voxel
<b>Diffusion Models</b>				
Wyatt et al., 2022 *	<i>Pernet et al.</i> <sup>9</sup>	0.863	–	Pixel
Behrendt et al., 2023 *	BraTS'21	–	0.541	Voxel
	MSLUB	–	0.106	
Bercea et al., 2023 *	ATLAS 2.0	–	0.145	Pixel
Iqbal et al., 2023 *	BraTS'21	–	0.590	Pixel
	MSLUB	–	0.106	
Kumar Trivedi et al., 2024 *	BraTS'20	–	0.417	Pixel
	BraTS'21	–	0.578	Pixel
	MSLUB	–	0.067	
Bi et al., 2025 *	BraTS'23	0.922	–	Slice

<sup>9</sup>Pernet et al., 2016

Table 5: Performance comparison using segmentation metrics (Dice similarity coefficient). Methods are grouped by model family (Autoencoders, VAEs, GANs, and Diffusion models), as defined in Table 3.. This summary details the reported Dice scores alongside the specific thresholding strategies employed to binarise residual maps (e.g., post-hoc optimisation, validation set, or percentile-based thresholding). Only studies that explicitly reported segmentation performance using the Dice metric are included.

Method	Test Dataset	Dice	Thresh. Strategy	Notes
<b>Autoencoders</b>				
Baur et al., 2018	In-house	0.605	Percentile-based	Evaluated on MS
Baur, Wiestler, et al., 2021	In-house	0.390	Validation set	Evaluated on Tumour
Ghorbel et al., 2023	BraTS'20	0.502	Post-hoc optimisation	
	MSLUB	0.173		
Kascenas et al., 2023 *	BraTS'21	0.773	Post-hoc optimisation	
	In-house (Multiple)	0.674		
Luo et al., 2023 *	BraTS'19	0.462	Validation set	
Meissen et al., 2023 *	BraTS'20	0.400	Percentile-based	
Jiménez-García et al., 2024	BraTS'20	0.471	Validation set	
<b>Variational Autoencoders</b>				
Uzunova et al., 2019	BraTS'15	0.550	Validation set	Evaluated in 2D on T1w Gd
		0.320		Evaluated in 3D on T1w Gd
		0.500		Evaluated in 3D on fusion of T1w Gd, T2w & FLAIR
Zimmerer et al., 2019 *	BraTS'17	0.440	Validation set	
Bengs et al., 2021	BraTS'19	0.301	Validation set	
	ATLAS 1.1	0.331		
Lambert et al., 2021	BraTS'18	0.650	Validation set	
<i>Continued on next page</i>				

Method	Test Dataset	Dice	Thresh.	Strategy	Notes
	WMH, MSSEG	0.463			
Chatterjee et al., 2022 *	BraTS'17	0.531		Histogram-based	Evaluated on T1w Gd
		0.642			Evaluated on T2w
Pinaya et al., 2022	MSLUB	0.378		Post-hoc optimisation	Evaluated in 2D
	BraTS'18	0.537			
	WMH	0.429			
	MSLUB	0.133			Evaluated in 3D
	BraTS'18	0.617			
	WMH	0.133			
Wijanarko et al., 2024	BraTS'20	0.606		Fixed	
<b>Generative Adversarial Networks</b>					
Dey et al., 2021	BraTS'19	0.680		Histogram-based	
	ISBI'15	0.482			
Nguyen et al., 2021	<i>Pernet et al.</i> <sup>10</sup>	0.770		Histogram-based	
Wu et al., 2021	BraTS'12/18	0.630		Percentile-based	
<b>Diffusion Models</b>					
Pinaya, Graham, et al., 2022	BraTS'18	0.398		Post-hoc optimisation	
	MSLUB	0.247			
	WMH	0.298			
Wyatt et al., 2022 *	<i>Pernet et al.</i> <sup>11</sup>	0.383		Fixed	
Behrendt et al., 2023 *	BraTS'21	0.490		Validation set	
	MSLUB	0.105			

*Continued on next page*

<sup>10</sup>Pernet et al., 2016

<sup>11</sup>Pernet et al., 2016

Method	Test Dataset	Dice	Thresh.	Strategy	Notes
Bercea et al., 2023 *	ATLAS 2.0	0.228		Post-hoc optimisation	
Iqbal et al., 2023 *	BraTS'21	0.530		Post-hoc optimisation	
	MSLUB	0.107			
Behrendt et al., 2024 *	BraTS'21	0.574		Validation set	
	ATLAS 2.0	0.148			
	MSLUB	0.061			
	WMH	0.132			
Bercea et al., 2024 *	ATLAS 2.0	0.297		Post-hoc optimisation	
Fontanella et al., 2024	BraTS'21	0.699		Validation set	
	WMH	0.569			
Kumar Trivedi et al., 2024 *	BraTS'20	0.371		Post-hoc optimisation	
	BraTS'21	0.506			
	MSLUB	0.055			
Bi et al., 2025 *	BraTS'23	0.738		Post-hoc optimisation	
Beizae et al., 2026 *	BraTS'21	0.697		Post-hoc optimisation	Evaluated on T1w
		0.730			Evaluated on T1w Gd
		0.803			Evaluated on T2w
		0.851			Evaluated on FLAIR
	ATLAS 2.0	0.416			

### 3.5.1 Autoencoders

Autoencoders do not define an explicit probabilistic density over the data but instead learn a deterministic encoder-decoder mapping that reconstructs input images through a compressed latent representation. In unsupervised anomaly detection, AEs are trained on healthy data to model normative anatomical variability; anomalies are inferred from residual differences between the input and its reconstruction. Although not strictly generative in a probabilistic sense, AEs established the reconstruction-based paradigm that underpins later probabilistic approaches such as anomaly detection with variational autoencoders and diffusion models.

**Study characteristics.** We identified nine studies employing AE-based frameworks for UAD in neuroimaging (Table 3). Most studies focused on brain tumours (Baur, Wiestler, et al., 2021; Ghorbel et al., 2023; Jiménez-García et al., 2024; Kascenas et al., 2023; Lu et al., 2024; Luo et al., 2023; Meissen et al., 2023; Qu et al., 2026). Applications to multiple sclerosis (MS) were reported by Baur et al. (2018), Ghorbel et al. (2023), and Luo et al. (2023). Stroke was addressed by Luo et al. (2023) and Qu et al. (2026). All studies used MRI (T1-w, T1-w Gd, T2-w and/or FLAIR), with one study additionally evaluating CT in a mixed tumour–stroke setting (Kascenas et al., 2023). Regarding input dimensionality, six studies predominantly used 2D slice-wise processing (full axial slices or patches), while three implemented full 3D convolutional AEs (Baur, Wiestler, et al., 2021; Jiménez-García et al., 2024; Luo et al., 2023). The AE studies were published between 2018 and 2026, spanning early dense-bottleneck implementations (Baur et al., 2018) to more recent 3D convolutional and attention-augmented architectures (Jiménez-García et al., 2024; Lu et al., 2024; Qu et al., 2026).

**Architecture recap.** Autoencoders were among the first deep learning architectures applied to unsupervised anomaly detection in neuroimaging. They consist of an encoder–decoder structure trained to reconstruct inputs through a compressed latent representation. In UAD settings, AEs are trained on healthy data, and anomalies are inferred from residual differences between input and reconstruction. Early implementations relied on fully connected bottlenecks, which imposed strong compression and limited spatial correspondence between latent representations and anatomical structures, often resulting in oversmoothed reconstructions (Baur et al., 2018). Subsequent work adopted fully convolutional architectures with skip connections and residual blocks to better preserve spatial information. Residual connections were introduced to facilitate gradient flow in deeper networks (He et al., 2016), and some studies additionally incorporated attention mechanisms inspired by vision transformers (Dosovitskiy et al., 2021; Ghorbel et al., 2023; Qu et al., 2026).

**Architectural variations.** Several extensions were introduced within the AE framework. Denoising autoencoders treated pathological regions as structured corruption and reconstructed plausible healthy patterns (Kascenas et al., 2023). Training strategies included elastic deformation-based augmentations (Jiménez-García et al., 2024) and composite reconstruction losses integrating contrastive or discriminative components (Lu et al., 2024). Some approaches incorporated pretrained convolutional feature extractors to guide latent representations (Meissen et al., 2023). Both 2D slice-wise and full 3D volumetric implementations were reported; volumetric models increased spatial context but required careful tuning of latent dimensionality to avoid oversmoothing or partial reconstruction of anomalous structures (Luo et al., 2023).

**Quantitative synthesis.** Performance varied across pathologies and evaluation levels (Tables 4 and 5).

*Brain tumours.* Dice ranged from 0.390 (validation-set threshold) to 0.773 (best-possible/post-hoc threshold) (Baur, Wiestler, et al., 2021; Kascenas et al., 2023). Voxel/pixel-level AUROC on BraTS datasets ranged from 0.780 to 0.838 (Ghorbel et al., 2023; Jiménez-García et al., 2024), while slice-level AUROC ranged from 0.844 to 0.992 (Lu et al., 2024; Luo et al., 2023). Voxel/pixel-level AUPRC ranged from 0.425 to 0.833 (Ghorbel et al., 2023; Kascenas et al., 2023), and slice-level AUPRC reached 0.741 (Luo et al., 2023).

*Stroke.* Dice was 0.674 in CT-based evaluation (Kascenas et al., 2023). Slice-level AUROC ranged from 0.807 to 0.839 (Luo et al., 2023; Qu et al., 2026), with AUPRC up to 0.705 (Luo et al., 2023).

*Multiple sclerosis.* Dice ranged from 0.173 (Ghorbel et al., 2023) to 0.605 (Baur et al., 2018). Voxel/pixel-level AUROC reached 0.886 (Ghorbel et al., 2023), with pixel-level AUPRC ranging from 0.203 (Ghorbel et al., 2023) to 0.300 (Baur, Wiestler, et al., 2021). Slice-level AUROC reached 0.858 and AUPRC reached 0.731 (Luo et al., 2023).

**Closing.** Across AE studies, tumour segmentation achieved the highest Dice values (up to 0.773 under best-threshold optimisation), whereas MS segmentation was lower and more variable (down to 0.173). Detection metrics were consistently higher at the slice level (up to 0.992 AUROC) than at the voxel/pixel level (up to 0.886 AUROC), reflecting differences in task formulation rather than superior localisation performance. Although some studies employed volumetric 3D autoencoders (Table 3), no consistent performance advantage over 2D slice-wise approaches was observed. Reported outcomes varied substantially with thresholding strategy (Table 5) and dataset composition, limiting direct numerical comparability across studies. Within AE studies, architectural variations such as denoising objectives, attention mechanisms, or composite reconstruction losses were explored in tumour segmentation settings; how-

ever, reported Dice values remained within the overall AE tumour interval. In addition, analyses of latent dimensionality showed that overly compressed representations led to oversmoothing and reduced lesion contrast, whereas excessively large latent spaces risked partial reconstruction of anomalous structures (Luo et al., 2023). These observations indicate that latent capacity tuning influenced reconstruction fidelity within AE frameworks.

### 3.5.2 Variational autoencoders

**Study characteristics.** We included nine records using variational autoencoders (VAEs) as the main UAD method (Table 3). Of these, all focused on brain tumours (Bengs et al., 2021; Chatterjee et al., 2022; Lambert et al., 2021; Lüth et al., 2023; Pinaya et al., 2022; Sato et al., 2019; Uzunova et al., 2019; Wijanarko et al., 2024; Zimmerer et al., 2019), with three also addressing stroke (Bengs et al., 2021; Lüth et al., 2023; Sato et al., 2019). WMH were studied by Pinaya et al. (2022) and Lambert et al. (2021) (the latter combining WMH with MS due to visual similarities). A single study focused on MS exclusively (Pinaya et al., 2022). There was a 7:4 split between 2D and 3D VAEs.

**Architecture recap.** Variational autoencoders extend the deterministic autoencoder framework by introducing a probabilistic latent space. Instead of mapping each input to a single latent code, the encoder outputs the parameters of a distribution - typically a Gaussian defined by mean and variance - regularised toward a prior, most often  $\mathcal{N}(0, I)$  (Kingma & Welling, 2013). During training, latent samples are drawn from this parameterised posterior distribution and passed through the decoder to reconstruct the input. The learning objective combines a reconstruction term with a Kullback–Leibler (KL) divergence that encourages alignment between the learned latent distribution and the prior. In unsupervised anomaly detection, anomalies are inferred from reconstruction residuals and/or latent-space deviations.

**Architectural variations.** Several extensions of the canonical VAE were proposed. Spatial VAEs preserved the latent representation as a low-resolution feature map rather than collapsing it into a dense vector, maintaining spatial correspondence between latent units and anatomical regions (Bengs et al., 2021; Lambert et al., 2021). Vector-quantised VAEs (VQ-VAEs) replaced continuous latent variables with discrete codebooks of embeddings to capture more structured latent representations (Oord et al., 2017). Building on this formulation, transformer-based autoregressive models were trained on healthy data to modify pathological latent codes toward healthy counterparts (Marimont & Tarroni, 2021; Pinaya et al., 2022).

Loss formulations also evolved. Some studies modified reconstruction terms or latent regularisation to alter sensitivity to anatomical variation (Sato et al., 2019). Structural similarity measures and composite reconstruction objectives integrating  $\ell_1$ ,  $\ell_2$ , and perceptual terms were explored (Meissen et al., 2023; Wijanarko et al., 2024). Context-encoding VAEs (ceVAEs) combined probabilistic reconstruction with masked inpainting branches sharing encoder–decoder weights to encourage contextual feature learning (Chatterjee et al., 2022). Additional variations included contrastive pretraining stages and alternative decoder distributions such as Gaussian mixture models and normalising flows (Koller & Friedman, 2009; Lüth et al., 2023; Rezende & Mohamed, 2015).

Conditioning mechanisms and positional encodings were also investigated in both 2D and 3D settings (Uzunova et al., 2019). Anomaly scoring strategies ranged from pure reconstruction residuals to ELBO-informed metrics incorporating the KL divergence term (Zimmerer et al., 2019).

**Quantitative synthesis.** Performance varied across pathologies and evaluation levels (Tables 4 and 5).

*Brain tumours.* Dice ranged from 0.301 (Bengs et al., 2021) to 0.650 (Lambert et al., 2021), depending on dimensionality and thresholding strategy. Voxel-level AUROC ranged from 0.582 to 0.968 (Sato et al., 2019; Wijanarko et al., 2024), while pixel-level AUPRC achieved 0.555 (Pinaya et al., 2022). Slice-level AUROC reached 0.866 (Pinaya et al., 2022), with slice-level AUPRC up to 0.279 (Bengs et al., 2021).

*Stroke.* Dice reached 0.331 on ATLAS 1.1 (Bengs et al., 2021). Voxel-level AUROC ranged from 0.672 to 0.898 (Lüth et al., 2023; Sato et al., 2019), with voxel-level AUPRC ranging from 0.186 to 0.256 (Bengs et al., 2021; Lüth et al., 2023). Slice-level AUPRC reached 0.256 (Bengs et al., 2021).

*Multiple sclerosis.* Dice ranged from 0.133 (Pinaya et al., 2022) to 0.463 (Lambert et al., 2021), depending on dimensionality and dataset. Voxel-level AUPRC reached 0.272 on MSLUB (Pinaya et al., 2022). Slice-level AUROC reached 0.866 (Pinaya et al., 2022).

*White matter hyperintensities.* Dice ranged from 0.133 (Pinaya et al., 2022) to 0.429 (Lambert et al., 2021). Pixel-level AUPRC reached 0.320 (Pinaya et al., 2022).

**Closing.** Across VAE studies, tumour segmentation achieved the highest Dice values (up to 0.650), whereas performance on smaller and more diffuse lesions (MS and WMH) remained substantially lower. Detection metrics showed wide AUROC ranges (up to 0.968), but AUPRC values were generally moderate, consistent with strong class imbalance at the voxel level. No consistent performance advantage was observed between 2D and 3D implementations across studies. Reported Dice values were highly sensitive to thresholding strategy, with several studies reporting post-hoc best-achievable Dice that should be

interpreted as upper-bound estimates rather than fixed operating points. However, across the broader VAE literature, tumour Dice remained bounded within approximately 0.30-0.65, with substantial overlap with AE and GAN ranges.

### 3.5.3 Generative adversarial networks

**Study characteristics.** We identified four studies employing generative adversarial networks (GANs) for unsupervised anomaly detection in neuroimaging (Table 3). Most focused exclusively on brain tumours using BraTS MRI data (Nguyen et al., 2021; Wu et al., 2021), with another one addressing multiple sclerosis lesions on MSSEG (Dey et al., 2021) or traumatic brain injury only with CT images from the Center-TBI cohort (Simarro et al., 2020). Two approaches used 2D slice-wise GANs, and the two others applied volumetric 3D architectures to TBI and tumours (Simarro et al., 2020; Wu et al., 2021).

**Architecture recap.** Generative adversarial networks, introduced by I. J. Goodfellow et al. (2014), consist of a generator trained to synthesise samples and a discriminator trained to distinguish generated from real data. Through adversarial optimisation, the generator approximates the training distribution. In neuroimaging UAD, GANs are typically trained on healthy brain images and used to generate pseudo-healthy reconstructions of pathological inputs, with anomalies inferred from residual differences between input and reconstruction. The most commonly adopted medical UAD framework was f-AnoGAN (Schlegl et al., 2019), an extension of AnoGAN (Schlegl et al., 2017) incorporating an encoder for faster inference and a Wasserstein-based objective to improve training stability. As with other adversarial models, GANs remain sensitive to training instability and mode collapse.

**Architectural variations.** Several task-specific adaptations were introduced within the GAN framework. Symmetry-driven designs incorporated contralateral hemisphere information as a structural prior for anomaly localisation (Wu et al., 2021). Partition-based strategies separated candidate anomalous regions prior to adversarial evaluation (Dey et al., 2021). Multi-stage pipelines introduced additional refinement modules to modify coarse reconstructions (Nguyen et al., 2021). Both 2D slice-wise and 3D implementations were represented among the included GAN studies (Table 3).

**Quantitative synthesis.** Performance varied across pathologies and evaluation levels (Tables 4 and 5).

*Brain tumours.* Dice ranged from 0.630 to 0.770 across BraTS and in-house datasets (Dey et al., 2021; Nguyen et al., 2021; Wu et al., 2021). No voxel/pixel- or slice-level AUROC/AUPRC values were reported

for tumour localisation in the included GAN studies.

*Multiple sclerosis.* Dice reached 0.482 on ISBI 2015 (Dey et al., 2021). No voxel/pixel- or slice-level detection metrics were reported for MS using GAN-based methods.

*Traumatic brain injury.* Voxel-level AUROC reached 0.749 on CT images from the CENTER-TBI cohort (Simarro et al., 2020). No AUPRC values were reported.

**Closing.** GAN-based approaches reported tumour Dice values ranging from 0.630 to 0.770, numerically overlapping with the AE tumour interval (Dey et al., 2021; Nguyen et al., 2021; Wu et al., 2021). Evidence for smaller or diffuse lesions remains limited; for MS, Dice reached 0.482 in a single evaluation (Dey et al., 2021). Quantitative detection metrics were sparsely reported, with voxel-level AUROC available only for CT-based traumatic brain injury (0.749) (Simarro et al., 2020). Although adversarial training aims to enhance reconstruction sharpness and boundary delineation, tumour Dice values remained within the upper range observed for autoencoder-based models. Limited reporting of voxel-level detection metrics restricts assessment of whether adversarial objectives confer consistent localisation advantages.

### 3.5.4 Diffusion models

**Study characteristics.** We identified 11 studies applying diffusion models to UAD in neuroimaging (Table 3). Nine addressed brain tumour detection and segmentation (Behrendt et al., 2023, 2024; Beizae et al., 2026; Bi et al., 2025; Fontanella et al., 2024; Iqbal et al., 2023; Kumar Trivedi et al., 2024; Pinaya, Graham, et al., 2022; Wyatt et al., 2022), with additional work on stroke (three) (Behrendt et al., 2024; Beizae et al., 2026; Bercea et al., 2023, 2024), MS (five) (Behrendt et al., 2023, 2024; Iqbal et al., 2023; Kumar Trivedi et al., 2024; Pinaya, Graham, et al., 2022), and WMH (Behrendt et al., 2024; Fontanella et al., 2024; Pinaya, Graham, et al., 2022). All but one study (Behrendt et al., 2024) processed data in a 2D slice-wise manner.

**Architecture recap.** Most diffusion-based UAD methods rely on Denoising Diffusion Probabilistic Models (DDPMs) (Ho et al., 2020). In DDPMs, a forward process progressively corrupts a training image by adding Gaussian noise over multiple timesteps. A neural network - typically a U-Net (Ronneberger et al., 2015) - is trained to predict and remove this noise at each timestep, thereby learning to approximate the healthy data distribution. In neuroimaging UAD, diffusion models are trained exclusively on healthy images. At inference, pathological inputs are partially noised and then denoised toward the learned healthy manifold. Anomaly maps are derived from discrepancies between the input and the pseudo-

healthy reconstruction (Bercea et al., 2023; Pinaya, Graham, et al., 2022; Wyatt et al., 2022). Because voxel-space DDPMs are computationally demanding for high-resolution 3D MRI, several studies adopted latent diffusion strategies in which images are first compressed via an autoencoder and diffusion is performed in latent space (Kumar Trivedi et al., 2024; Pinaya, Graham, et al., 2022). This reduces memory and computational requirements while preserving structural content.

**Architectural variations.** Several methodological directions were explored within diffusion-based UAD.

*Noise design.* While early approaches used Gaussian corruption, structured noise patterns such as Perlin or Simplex noise were investigated to modify the corruption process (Bercea et al., 2024; Wyatt et al., 2022).

*Similarity metrics.* Beyond voxel-wise intensity residuals, structure-aware metrics such as SSIM (Behrendt et al., 2024) and perceptual feature-based measures including LPIPS (Chen et al., 2019; Zhang et al., 2018) were used to define anomaly maps.

*Guided restoration.* To mitigate excessive modification of healthy tissue during denoising, some methods reintroduced uncorrupted image regions after an initial pass (Bercea et al., 2023) or selectively preserved anatomical regions during reconstruction (Bercea et al., 2024), while others, like MCDDPM (Kumar Trivedi et al., 2024), utilise multichannel latent representations and cross-attention to integrate contextual information directly into the denoising process.

*Patch- and mask-based designs.* Patch-based diffusion models restricted corruption to selected spatial regions (Behrendt et al., 2023), while masked DDPM variants modified spatial or frequency components of the input (Iqbal et al., 2023).

*Hybrid and multi-stage inference.* Some pipelines combined diffusion with saliency-driven or counterfactual guidance mechanisms (Fontanella et al., 2024). Multi-stage inference cascades applied repeated denoising passes to progressively attenuate pathological structures (Bi et al., 2025).

*Emerging continuous-time frameworks.* Beyond stochastic DDPMs, deterministic alternatives based on conditional flow matching and rectified flows were introduced (Albergo & Vanden-Eijnden, 2023; Lipman et al., 2023; Liu et al., 2023). Early neuroimaging implementations such as REFLECT (Beizae et al., 2026) operated in latent space and relied on self-supervised proxy corruptions to construct training objectives.

**Quantitative synthesis.** Performance varied across pathologies and evaluation levels (Tables 4 and 5).

*Brain tumours.* Dice ranged from 0.371 (Kumar Trivedi et al., 2024) to 0.851 (Beizae et al., 2026), with the highest values reported under Dice-optimised threshold selection. Voxel/pixel-level AUROC reached 0.863 (Wyatt et al., 2022), while voxel/pixel-level AUPRC ranged from 0.417 to 0.590 (Iqbal et al., 2023; Kumar Trivedi et al., 2024). Slice-level AUROC reached 0.920 (Bi et al., 2025).

*Stroke.* Dice ranged from 0.147 to 0.416 under Dice-optimised thresholds on ATLAS 2.0 (Behrendt et al., 2024; Beizae et al., 2026). Voxel-level AUPRC reached 0.145 (Bercea et al., 2023). No AUROC values were reported for stroke localisation.

*Multiple sclerosis.* Dice ranged from 0.061 to 0.247 (Behrendt et al., 2024; Pinaya, Graham, et al., 2022). Voxel/pixel-level AUPRC ranged from 0.067 (Kumar Trivedi et al., 2024) to 0.106 (Behrendt et al., 2023). No slice-level detection metrics were reported.

*White-matter hyperintensities.* Dice ranged from 0.132 to 0.569 (Behrendt et al., 2024; Fontanella et al., 2024). No AUROC/AUPRC values were reported for WMH in the included diffusion studies.

**Closing.** Diffusion-based approaches achieved their highest segmentation performance for tumours, with Dice values reaching 0.851 under Dice-optimised thresholding. Performance was substantially lower for stroke (up to 0.416) and particularly for multiple sclerosis (typically below 0.250), while WMH yielded intermediate values (up to 0.569). Detection metrics followed a similar pattern, with tumour localisation AUPRC exceeding 0.500 but remaining markedly lower for MS and stroke. Tumour Dice values therefore extended to a higher maximum than those reported for AE and VAE frameworks, but lower tumour Dice values still overlapped with reconstruction-based approaches. For smaller or sparse lesions, diffusion models did not shift the performance range substantially, with Dice remaining low. Patch- and mask-based diffusion variants, as well as multi-stage inference strategies, were associated with some of the highest reported tumour Dice values. In contrast, several other diffusion configurations reported tumour Dice in the mid-range (approximately 0.37–0.57). However, performance intervals overlapped substantially across diffusion variants, limiting attribution of gains to a single architectural component.

### 3.6 Comparison across studies

Unsupervised brain tumour segmentation has been investigated more extensively than other pathologies, likely reflecting the relatively high availability of curated datasets such as BraTS. Across model families,

tumours generally achieved the highest Dice scores. MRI remained the predominant imaging modality across all pathologies, with T1-w, T2-w, and FLAIR being the most frequently utilised sequences (Figure 3). T1-w Gd was commonly used in glioma datasets, while CT was sporadically employed for stroke or traumatic brain injury.

Input dimensionality varied considerably. In total, 10 studies adopted full volumetric 3D processing, whereas 16 processed complete 2D axial stacks sequentially. A further 7 studies used partial 2D processing restricted to selected axial slices, typically centred around mid-brain regions. Such design choices may influence anatomical context, class imbalance, and computational cost, and may therefore affect reported accuracy. Overall, 2D processing remained predominant, particularly among diffusion-based approaches.

Dice thresholding standardisation represents a major obstacle for cross-study comparison. Ten studies reported an upper-bound metric obtained through retrospective selection of the threshold yielding the best possible Dice, whereas 10 employed validation-set thresholding. The remaining eight used fixed or algorithmic thresholding strategies (e.g., Otsu, histogram-based, Felzenszwalb). Post-hoc best achievable Dice reporting became more frequent after 2022, complicating temporal comparisons and potentially inflating reported segmentation performance.

Across pathologies, smaller or sparser lesions such as MS and WMH consistently yielded lower Dice values across all model families. Sub-acute and chronic stroke segmentation was less frequently studied (six studies), and reported scores remained below those typically observed for tumours, reflecting lesion heterogeneity and less sharply defined boundaries (Hernandez Petzsche et al., 2022).

Taken together, pathology type exerted a stronger influence on performance than architectural family. While Dice ranges overlapped substantially across AE, VAE, GAN, and diffusion models, lesion size, contrast, and dataset characteristics appeared to drive performance differences more than model class. Importantly, heterogeneous datasets, preprocessing pipelines, evaluation levels, and thresholding strategies limit direct numerical comparability of absolute values across studies.

## 4 Discussion

**Principal findings.** This scoping review synthesised eight years (2018-2025) of research on unsupervised deep generative models for anomaly detection in structural neuroimaging. Across 33 studies, four principal architectural families were identified: autoencoders (AEs), variational autoencoders (VAEs), generative adversarial networks (GANs), and diffusion-based models. Across all model families, pathology type exerted a stronger influence on performance than architecture. Brain tumours - typically large




Figure 3: Distribution of MRI sequences and imaging modalities across major pathology categories.

and hyperintense - consistently achieved the highest segmentation and detection metrics. In contrast, smaller, sparse, or heterogeneous abnormalities such as multiple sclerosis (MS), white-matter hyperintensities (WMH), and stroke remained substantially more challenging. Dice scores for MS and WMH frequently fell below 0.50, and sometimes below 0.30, regardless of architecture. Diffusion models aim to approximate complex anatomical distributions through iterative denoising (Ho et al., 2020), and several neuroimaging studies qualitatively described anatomically plausible pseudo-healthy reconstructions. Nevertheless, segmentation performance did not consistently exceed that of AE-, VAE-, or GAN-based approaches across pathologies. GANs reported competitive tumour Dice values in selected studies, while AEs and VAEs achieved overlapping performance ranges. No unsupervised family approached the performance of contemporary supervised baselines on curated benchmarks such as BraTS (Bonato et al., 2025), where Dice scores above 0.90 are now routinely reported<sup>12</sup>. Despite these limitations, reconstruction-based generative models offer a distinct conceptual advantage: the ability to produce subject-specific pseudo-healthy images that explicitly visualise deviation from normative anatomy. This paradigm resembles clinical image interpretation, where abnormalities are assessed relative to expected anatomical structure. Finally, part of the observed performance variability likely reflects dataset effects rather than architecture alone. Reconstruction-based models are sensitive to the integrity and representativeness of the healthy training distribution: even 3% contamination by pathological cases can significantly reduce tumour-detection AUROC (Behrendt et al., 2022). In addition, normative cohorts such as UK Biobank and ADNI are affected by selection and demographic biases (Gianattasio et al., 2021; Littlejohns et al., 2020), which may exacerbate domain shift and limit generalisation.

<sup>12</sup>See <https://www.synapse.org/Synapse:syn53708249>

**Results interpretation.** The observed performance patterns appear to be driven primarily by interactions between lesion characteristics and model inductive biases. Large, well-contrasted tumours in curated benchmarks such as BraTS provide favourable conditions for residual-based detection, owing to their relatively high lesion-to-background contrast and volumetric prominence. In contrast, MS and WMH lesions are typically small and sparse, creating severe class imbalance and making Dice highly sensitive to minor localisation errors. Stroke lesions introduce additional complexity through heterogeneous intensity profiles and ill-defined boundaries (Hernandez Petzsche et al., 2022), which reduce residual conspicuity even when lesion volumes are comparable to tumours. Architectural refinements modified - but did not fundamentally overcome - these pathology-driven constraints. Differences between 2D and 3D implementations were inconsistent across studies (Pinaya et al., 2022; Uzunova et al., 2019), likely reflecting trade-offs between receptive field size, optimisation stability, volumetric class imbalance, and dataset scale rather than intrinsic superiority of a given dimensional paradigm. Patch-based diffusion variants (Behrendt et al., 2023; Kumar Trivedi et al., 2024) may further limit anatomical context when applied to small or diffuse abnormalities. Across model families, design modifications such as masking, tailored losses, pretraining, and hybrid latent strategies improved performance incrementally, but generally within pathology-dependent ceilings. These findings suggest that dataset composition and lesion morphology constrain achievable performance to a greater extent than architectural family alone.

**Methodological constraints of reconstruction-based UAD.** A central limitation of most reviewed approaches lies in the reconstruction-based anomaly paradigm itself. These models assume that pathological patterns are poorly represented in the learned healthy distribution and will therefore be attenuated during reconstruction. However, when lesion intensities overlap with normal anatomical variation, or when the network learns an identity shortcut, anomalous structures may be partially or fully reconstructed, leading to false negatives. Recent work has shown that improving reconstruction fidelity alone does not guarantee improved anomaly detection performance (Meissen et al., 2022). In such cases, voxel-wise residual maps become insufficiently discriminative, particularly for low-contrast, diffuse, or texture-preserving abnormalities. Several mitigation strategies have been proposed. Latent-space masking techniques aim to prevent trivial identity mappings by forcing inpainting from normative representations (Beizae et al., 2025). Decoder-level perturbation and attention-based masking mechanisms similarly constrain shortcut learning (Tang et al., 2025). Other approaches improve residual contrast through structured post-processing (Muñoz-Ramírez et al., 2021). Nevertheless, these solutions expose a broader trade-off: increasing reconstruction fidelity can reduce anomaly sensitivity when pathological patterns are also reconstructed, thereby diminishing residual-based contrast and increasing false negatives.

**Future directions.** Our synthesis highlights several priorities for advancing unsupervised anomaly detection (UAD) in neuroimaging. The most persistent challenge remains performance on small, sparse, or ill-defined abnormalities, where Dice scores consistently lag behind those achieved for large, hyper-intense tumours. Addressing this gap will require progress in both architectural design and evaluation methodology. *Architectural advances.* More efficient continuous-time generative models - such as flow matching (Lipman et al., 2023) or score-based diffusion variants (Song et al., 2021) - offer promising alternatives to classical DDPMs, with improved sampling efficiency and potentially greater stability. Latent-space hybrids combining strong encoders (e.g., VQ-VAEs or spatial VAEs) with diffusion- or flow-based decoders may enable scalable volumetric pseudo-healthy reconstruction (Bengs et al., 2021; Lambert et al., 2021; Pinaya et al., 2022). Incorporating anatomy-aware priors, including symmetry constraints or atlas guidance, may improve localisation of subtle or unilateral abnormalities. Perceptually informed residual measures (e.g., SSIM, LPIPS) and hybrid generative–discriminative strategies may further enhance anomaly contrast while limiting false positives (Behrendt et al., 2024; Chen et al., 2019; Huijben et al., 2024; Zhang et al., 2018). Nevertheless, architectural refinement alone is unlikely to eliminate pathology-dependent performance ceilings.

*Self-supervised and anomaly-aware paradigms.* A subset of recent studies adopts self-supervised strategies in which models are exposed to synthetic anomalies during training (e.g., corrupted latent representations in Beizae et al. (2026)). These approaches often report higher Dice scores by explicitly learning deviation patterns. However, synthetic perturbations may not faithfully capture the appearance of real-world clinical anomalies, particularly low-contrast or diffuse lesions. This introduces a conceptual distinction between purely normative generative models and anomaly-aware self-supervised frameworks, with uncertain implications for generalisability across heterogeneous pathologies.

**Related foundation-model-based anomaly pipelines (outside normative modelling).** Recent work has also explored foundation-model-based anomaly detection pipelines that leverage large pretrained encoders or segmentation models. For example, Ma et al. (2025) used CLIP-generated pseudo-labels and a foundation segmentation model to train a downstream 3D U-Net in a self-supervised manner. Similarly, Rahmaniar and Suzuki (2026) proposed a teacher–student distillation framework built on an ImageNet-pretrained backbone for tumour detection. Unlike normative generative models, these approaches rely on pretrained semantic representations or pseudo-supervision and do not learn a healthy anatomical distribution through reconstruction. They were therefore considered adjacent but methodologically distinct from reconstruction-based generative UAD.

**Evaluation reform and benchmarking.** A major bottleneck for comparative analysis remains the absence of standardised evaluation protocols. Metrics are inconsistently reported across voxel-, slice-, and subject-level tasks, which are not directly comparable. Furthermore, post-hoc threshold optimisation frequently yields upper-bound “best achievable” Dice scores that may overestimate clinically realistic performance relative to validation-based thresholding.

Conventional metrics such as Dice and AUROC assess localisation accuracy on predefined lesions but do not capture the core objective of reconstruction-based generative models: robust normative representation learning. Importantly, reconstruction fidelity is not equivalent to anomaly sensitivity; improving reconstruction quality does not guarantee improved detection, particularly when pathological patterns overlap with normal anatomical variation (Meissen et al., 2022). Conversely, competitive Dice scores do not ensure preservation of healthy tissue or plausible correction of anomalies.

To address this gap, Bercea, Wiestler, et al. (2025) reframed UAD as a problem of *normative representation learning* and proposed task-specific indices tailored to pseudo-healthy reconstructions:

- *Restoration Quality Index (RQI)* — evaluates perceptual reconstruction behaviour using LPIPS (Zhang et al., 2018).
- *Anomaly-to-Healthy Index (AHI)* — measures how plausibly a pathological image is transformed toward a healthy distribution using Fréchet Inception Distance.
- *Conservation and Correction Index (CACI)* — quantifies whether healthy regions are preserved while anomalous regions are selectively corrected using structural similarity measures.

These indices complement Dice and AUROC by characterising pseudo-healthy reconstruction behaviour rather than lesion localisation alone. In a multi-reader study involving radiologists, stronger normative reconstruction metrics were associated with improved perceived plausibility and encouraging generalisation across unseen pathologies (Bercea, Wiestler, et al., 2025). Together, these findings suggest that evaluation of UAD models should consider both detection accuracy and the quality of normative modelling.

Benchmarks that incorporate rare or heterogeneous pathologies (e.g., NOVA) may further stress-test generalisation beyond widely used datasets such as BraTS or MSSEG.

**Clinical implications.** Unsupervised generative models are best positioned as broad anomaly detectors or triage tools, particularly where voxel-level annotations are unavailable. Their pseudo-healthy reconstructions provide interpretable contrastive visualisations that may complement supervised segmentation

models. However, for routine clinical segmentation - especially in small or diffuse pathologies - current accuracy remains insufficient. Controlled reader studies and prospective validation are required to determine whether pseudo-healthy reconstructions improve detection confidence, workflow efficiency, or longitudinal consistency under real-world conditions.

**Limitations.** This review has several limitations. First, although five major databases were searched (PubMed, Web of Science, ScienceDirect, Springer Nature, IEEE Xplore, and ArXiv), some relevant studies may have been missed. Restricting ArXiv searches to the *Computer Science (cs)* category may have excluded biomedical preprints. Second, the cut-off date (17 December 2025) introduces temporal bias in a rapidly evolving field. Third, quantitative synthesis relied primarily on Dice, AUROC, and AUPRC. Dice disproportionately penalises small lesions such as MS or WMH, AUROC can obscure class imbalance, and AUPRC introduces its own biases and is not inherently superior (McDermott et al., 2024). Other measures were reported inconsistently and could not be systematically compared. Fourth, evaluation strategies varied substantially across studies, particularly regarding residual threshold selection. Finally, conventional metrics do not fully capture the normative modelling objective of reconstruction-based UAD (Bercea, Wiestler, et al., 2025), underscoring the need for task-specific evaluation frameworks.

## 5 Conclusion

In this systematic scoping review, we compared generative AI-based methods for anomaly detection and segmentation in brain MRI, focusing on their ability to model healthy anatomy and detect deviations. None of the included studies solved the challenge across all pathologies. In detection, some methods achieved high AUROC values ( $> 0.9$ ), but performance was typically pathology- or dataset-specific. No generalisable detection framework has yet emerged across pathologies and datasets. Segmentation remains particularly challenging: Dice scores for large lesions were moderate (often between 0.6–0.8), whereas performance for small or sparse lesions frequently fell below 0.3 and occasionally below 0.1. We categorised studies by architecture (AE, VAE, GAN, diffusion) and summarised their main contributions (Table 3). Following PRISMA guidelines, we provided a transparent and reproducible synthesis, identifying consistent performance patterns across pathologies and highlighting emerging innovation. In summary, unsupervised generative models provide a valuable, annotation-free strategy for detecting and visualising neuroimaging anomalies. However, performance remains limited for small or sparse lesions, and these methods do not yet match supervised baselines. Future work should prioritise anatomy-aware architectures, standardised multi-pathology benchmarks, and prospective reader studies to establish whether

pseudo-healthy reconstructions can translate into clinically meaningful support, such as in research and follow-up exams.

## Data and materials availability

All extracted data, screening records, and analysis code (including scripts for figure generation) are openly available in the supporting GitHub repository<sup>13</sup>. The full PRISMA-ScR checklist is provided in the Supplementary Materials.

## Competing interests

The authors declare no conflicts of interest. No review protocol was prepared, and this review was not preregistered in PROSPERO or any other registry. Large language models (LLMs) were used exclusively for writing assistance, editing, and formatting. They did not contribute to study design, methodology, data analysis, or interpretation, and therefore did not affect the originality or scientific rigour of this work.

**CRedit authorship contribution statement** **YM**: Conceptualisation, Methodology, Investigation, Data curation, Writing – original draft. **EB**: Methodology, Supervision, Writing – review & editing. **SL**: Methodology, Supervision, Writing – review & editing. **EF**: Supervision, Writing – review & editing. **FG**: Methodology, Validation, Supervision, Writing – review & editing. All authors read and approved the final manuscript.

---

<sup>13</sup><https://github.com/youwanM/Unsupervised-Deep-Generative-Models-for-Anomaly-Detection-in-Neuroimaging>

# Supplementary Material

## Preferred Reporting Items for Systematic reviews and Meta-Analyses extension for Scoping Reviews (PRISMA-ScR) Checklist

SECTION	ITEM	PRISMA-ScR CHECKLIST ITEM	REPORTED ON PAGE #
<b>TITLE</b>			
Title	1	Identify the report as a scoping review.	1
<b>ABSTRACT</b>			
Structured summary	2	Provide a structured summary that includes (as applicable): background, objectives, eligibility criteria, sources of evidence, charting methods, results, and conclusions that relate to the review questions and objectives.	1
<b>INTRODUCTION</b>			
Rationale	3	Describe the rationale for the review in the context of what is already known. Explain why the review questions/objectives lend themselves to a scoping review approach.	1-4
Objectives	4	Provide an explicit statement of the questions and objectives being addressed with reference to their key elements (e.g., population or participants, concepts, and context) or other relevant key elements used to conceptualize the review questions and/or objectives.	4
<b>METHODS</b>			
Protocol and registration	5	Indicate whether a review protocol exists; state if and where it can be accessed (e.g., a Web address); and if available, provide registration information, including the registration number.	40
Eligibility criteria	6	Specify characteristics of the sources of evidence used as eligibility criteria (e.g., years considered, language, and publication status), and provide a rationale.	6-7
Information sources*	7	Describe all information sources in the search (e.g., databases with dates of coverage and contact with authors to identify additional sources), as well as the date the most recent search was executed.	6
Search	8	Present the full electronic search strategy for at least 1 database, including any limits used, such that it could be repeated.	6-8
Selection of sources of evidence†	9	State the process for selecting sources of evidence (i.e., screening and eligibility) included in the scoping review.	6-7
Data charting process‡	10	Describe the methods of charting data from the included sources of evidence (e.g., calibrated forms or forms that have been tested by the team before their use, and whether data charting was done independently or in duplicate) and any processes for obtaining and confirming data from investigators.	7
Data items	11	List and define all variables for which data were sought and any assumptions and simplifications made.	39
Critical appraisal of individual sources of evidence§	12	If done, provide a rationale for conducting a critical appraisal of included sources of evidence; describe the methods used and how this information was used in any data synthesis (if appropriate).	N/A
Synthesis of results	13	Describe the methods of handling and summarizing the data that were charted.	13



## Supplementary Material

SECTION	ITEM	PRISMA-ScR CHECKLIST ITEM	REPORTED ON PAGE #
<b>RESULTS</b>			
Selection of sources of evidence	14	Give numbers of sources of evidence screened, assessed for eligibility, and included in the review, with reasons for exclusions at each stage, ideally using a flow diagram.	9
Characteristics of sources of evidence	15	For each source of evidence, present characteristics for which data were charted and provide the citations.	14-24
Critical appraisal within sources of evidence	16	If done, present data on critical appraisal of included sources of evidence (see item 12).	N/A
Results of individual sources of evidence	17	For each included source of evidence, present the relevant data that were charted that relate to the review questions and objectives.	26-33
Synthesis of results	18	Summarize and/or present the charting results as they relate to the review questions and objectives.	33-34
<b>DISCUSSION</b>			
Summary of evidence	19	Summarize the main results (including an overview of concepts, themes, and types of evidence available), link to the review questions and objectives, and consider the relevance to key groups.	34-39
Limitations	20	Discuss the limitations of the scoping review process.	39
Conclusions	21	Provide a general interpretation of the results with respect to the review questions and objectives, as well as potential implications and/or next steps.	39-40
<b>FUNDING</b>			
Funding	22	Describe sources of funding for the included sources of evidence, as well as sources of funding for the scoping review. Describe the role of the funders of the scoping review.	40

JBIG = Joanna Briggs Institute; PRISMA-ScR = Preferred Reporting Items for Systematic reviews and Meta-Analyses extension for Scoping Reviews.

\* Where *sources of evidence* (see second footnote) are compiled from, such as bibliographic databases, social media platforms, and Web sites.

† A more inclusive/heterogeneous term used to account for the different types of evidence or data sources (e.g., quantitative and/or qualitative research, expert opinion, and policy documents) that may be eligible in a scoping review as opposed to only studies. This is not to be confused with *information sources* (see first footnote).

‡ The frameworks by Arksey and O'Malley (6) and Levac and colleagues (7) and the JBI guidance (4, 5) refer to the process of data extraction in a scoping review as data charting.

§ The process of systematically examining research evidence to assess its validity, results, and relevance before using it to inform a decision. This term is used for items 12 and 19 instead of "risk of bias" (which is more applicable to systematic reviews of interventions) to include and acknowledge the various sources of evidence that may be used in a scoping review (e.g., quantitative and/or qualitative research, expert opinion, and policy document).

From: Tricco AC, Lillie E, Zarin W, O'Brien KK, Colquhoun H, Levac D, et al. PRISMA Extension for Scoping Reviews (PRISMA-ScR): Checklist and Explanation. *Ann Intern Med.* 2018;169:467-473. doi: 10.7326/M18-0850.



**St. Michael's**  
Inspired Care.  
Inspiring Science.

## References

- Albergo, M. S., & Vanden-Eijnden, E. (2023). Building normalizing flows with stochastic interpolants. *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=li7qeBbCR1t>
- Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F. C., Pati, S., Prevedello, L. M., Rudie, J. D., Sako, C., Shinohara, R. T., Bergquist, T., Chai, R., Eddy, J., Elliott, J., Reade, W., . . . Bakas, S. (2021). The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. <https://arxiv.org/abs/2107.02314>
- Baur, C., Denner, S., Wiestler, B., Navab, N., & Albarqouni, S. (2021). Autoencoders for unsupervised anomaly segmentation in brain MR images: A comparative study. *Medical Image Analysis*, *69*, 101952. <https://doi.org/10.1016/j.media.2020.101952>
- Baur, C., Wiestler, B., Albarqouni, S., & Navab, N. (2018). Deep Autoencoding Models for Unsupervised Anomaly Segmentation in Brain MR Images. <https://doi.org/10.48550/ARXIV.1804.04488>
- Baur, C., Wiestler, B., Muehlau, M., Zimmer, C., Navab, N., & Albarqouni, S. (2021). Modeling Healthy Anatomy with Artificial Intelligence for Unsupervised Anomaly Detection in Brain MRI. *Radiology: Artificial Intelligence*, *3*(3), e190169. <https://doi.org/10.1148/ryai.2021190169>
- Beckett, L. A., Donohue, M. C., Wang, C., Aisen, P., Harvey, D. J., Saito, N., & Initiative, A. D. N. (2015). The alzheimer's disease neuroimaging initiative phase 2: Increasing the length, breadth, and depth of our understanding. *Alzheimer's & Dementia*, *11*(7), 823–831. <https://doi.org/https://doi.org/10.1016/j.jalz.2015.05.004>
- Behrendt, F., Bengs, M., Rogge, F., Kruger, J., Opfer, R., & Schlaefer, A. (2022). Unsupervised Anomaly Detection in 3D Brain MRI Using Deep Learning with Impured Training Data - 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI), 1–4. <https://doi.org/10.1109/ISBI52829.2022.9761443>
- Behrendt, F., Bhattacharya, D., Krüger, J., Opfer, R., & Schlaefer, A. (2023). Patched Diffusion Models for Unsupervised Anomaly Detection in Brain MRI. <https://doi.org/10.48550/ARXIV.2303.03758>
- Behrendt, F., Bhattacharya, D., Maack, L., Krüger, J., Opfer, R., Mieling, R., & Schlaefer, A. (2024). Diffusion Models with Ensembled Structure-Based Anomaly Scoring for Unsupervised Anomaly Detection - 2024 IEEE International Symposium on Biomedical Imaging (ISBI), 1–4. <https://doi.org/10.1109/ISBI56570.2024.10635828>
- Beizae, F., Hajimiri, S., Ben Ayed, I., Lodygensky, G., Desrosiers, C., & Dolz, J. (2026). Reflect: Rectified flows for efficient brain anomaly correction transport. In J. C. Gee, D. C. Alexander, J. Hong, J. E. Iglesias, C. H. Sudre, A. Venkataraman, P. Golland, J. H. Kim, & J. Park (Eds.), *Medical image*

- computing and computer assisted intervention – miccai 2025* (pp. 456–466). Springer Nature Switzerland.
- Beizae, F., Lodygensky, G., Desrosiers, C., & Dolz, J. (2025). Mad-ad: Masked diffusion for unsupervised brain anomaly detection. *International Conference on Information Processing in Medical Imaging*, 139–153.
- Bengs, M., Behrendt, F., Krüger, J., Opfer, R., & Schlaefer, A. (2021). Three-dimensional deep learning with spatial erasing for unsupervised anomaly segmentation in brain MRI. *International Journal of Computer Assisted Radiology and Surgery*, 16(9), 1413–1423. <https://doi.org/10.1007/s11548-021-02451-9>
- Bercea, C. I., Cattin, P. C., Schnabel, J. A., & Wolleb, J. (2025). Denoising diffusion models for anomaly localization in medical images. *Machine Learning for Biomedical Imaging*, 3, 687–711. <https://doi.org/https://doi.org/10.59275/j.melba.2025-c586>
- Bercea, C. I., Li, J., Raffler, P., Riedel, E. O., Schmitzer, L., Kurz, A., Bitzer, F., Roßmüller, P., Canisius, J., Beyrle, M. L., Liu, C., Bai, W., Kainz, B., Schnabel, J. A., & Wiestler, B. (2025). Nova: A benchmark for anomaly localization and clinical reasoning in brain mri. <https://arxiv.org/abs/2505.14064>
- Bercea, C. I., Neumayr, M., Rueckert, D., & Schnabel, J. A. (2023). Mask, stitch, and re-sample: Enhancing robustness and generalizability in anomaly detection through automatic diffusion models. *ICML 3rd Workshop on Interpretable Machine Learning in Healthcare (IMLH)*. <https://openreview.net/forum?id=kTpafpXrqa>
- Bercea, C. I., Wiestler, B., Rueckert, D., & Schnabel, J. A. (2024). Diffusion models with implicit guidance for medical anomaly detection. In M. G. Linguraru, Q. Dou, A. Feragen, S. Giannarou, B. Glocker, K. Lekadir, & J. A. Schnabel (Eds.), *Medical image computing and computer assisted intervention – miccai 2024* (pp. 211–220). Springer Nature Switzerland.
- Bercea, C. I., Wiestler, B., Rueckert, D., & Schnabel, J. A. (2025). Evaluating normative representation learning in generative ai for robust anomaly detection in brain imaging. *Nature Communications*, 16(1), 1624. <https://doi.org/10.1038/s41467-025-56321-y>
- Bi, Y., Huang, L., Clarenbach, R., Ghotbi, R., Karlas, A., Navab, N., & Jiang, Z. (2025). Synomaly noise and multi-stage diffusion: A novel approach for unsupervised anomaly detection in medical images. *Medical Image Analysis*, 105, 103737. <https://doi.org/10.1016/j.media.2025.103737>
- Bonato, B., Nanni, L., & Bertoldo, A. (2025). Advancing precision: A comprehensive review of mri segmentation datasets from brats challenges (2012–2025). *Sensors*, 25(6). <https://doi.org/10.3390/s25061838>

- Cabanac, G., Labbé, C., & Magazinov, A. (2022). The 'Problematic Paper Screener' automatically selects suspect publications for post-publication (re)assessment [The theme of the conference is 'Fostering Research Integrity in an Unequal World']. <https://doi.org/10.48550/arXiv.2210.04895>
- Carass, A., Roy, S., Jog, A., Cuzzocreo, J. L., Magrath, E., Gherman, A., Button, J., Nguyen, J., Prados, F., Sudre, C. H., Jorge Cardoso, M., Cawley, N., Ciccarelli, O., Wheeler-Kingshott, C. A., Ourselin, S., Catanese, L., Deshpande, H., Maurel, P., Commowick, O., . . . Pham, D. L. (2017). Longitudinal multiple sclerosis lesion segmentation: Resource and challenge. *NeuroImage*, *148*, 77–102. <https://doi.org/https://doi.org/10.1016/j.neuroimage.2016.12.064>
- Chatterjee, S., Sciarra, A., Dünnwald, M., Tummala, P., Agrawal, S. K., Jauhari, A., Kalra, A., Oeltze-Jafra, S., Speck, O., & Nürnberger, A. (2022). StRegA: Unsupervised anomaly detection in brain MRIs using a compact context-encoding variational autoencoder. *Computers in Biology and Medicine*, *149*, 106093. <https://doi.org/10.1016/j.compbiomed.2022.106093>
- Chen, S., Ma, K., & Zheng, Y. (2019). Med3d: Transfer learning for 3d medical image analysis. <https://arxiv.org/abs/1904.00625>
- Cheng, J. (2017). brain tumor dataset. <https://doi.org/10.6084/m9.figshare.1512427.v8>
- Commowick, O., Istace, A., Kain, M., Laurent, B., Leray, F., Simon, M., Pop, S. C., Girard, P., Améli, R., Ferré, J.-C., Kerbrat, A., Tourdias, T., Cervenansky, F., Glatard, T., Beaumont, J., Doyle, S., Forbes, F., Knight, J., Khademi, A., . . . Barillot, C. (2018). Objective evaluation of multiple sclerosis lesion segmentation using a data management and processing infrastructure. *Scientific Reports*, *8*(1). <https://doi.org/10.1038/s41598-018-31911-7>
- Commowick, O., Kain, M., Casey, R., Ameli, R., Ferré, J.-C., Kerbrat, A., Tourdias, T., Cervenansky, F., Camarasu-Pop, S., Glatard, T., Vukusic, S., Edan, G., Barillot, C., Dojat, M., & Cotton, F. (2021). Multiple sclerosis lesions segmentation from multiple experts: The miccai 2016 challenge dataset. *NeuroImage*, *244*, 118589. <https://doi.org/https://doi.org/10.1016/j.neuroimage.2021.118589>
- Crimi, A., & Bakas, S. (2021). *Brainlesion: Glioma, multiple sclerosis, stroke and traumatic brain injuries: 6th international workshop, brainles 2020, held in conjunction with miccai 2020, lima, peru, october 4, 2020, revised selected papers, part i* (Vol. 12658). Springer Nature.
- Currie, S., Saleem, N., Straiton, J. A., Macmullen-Price, J., Warren, D. J., & Craven, I. J. (2015). Imaging assessment of traumatic brain injury. *Postgraduate Medical Journal*, *92*(1083), 41–50. <https://doi.org/10.1136/postgradmedj-2014-133211>
- Darrault, F., Dannhoff, G., Chauvel, M., Delmaire, T., Louchez, S., Poupon, C., Uszynski, I., Destrieux, C., Maldonado, I. L., & Andersson, F. (2025). A road map to manual segmentation of cerebral structures. *Journal of Anatomy*, *246*(5), 819–828. <https://doi.org/https://doi.org/10.1111/joa.14167>

- Dey, R., Sun, W., Xu, H., & Hong, Y. (2021). ASC-Net: Unsupervised Medical Anomaly Segmentation Using an Adversarial-based Selective Cutting Network. <https://doi.org/10.48550/ARXIV.2112.09135>
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations*. <https://openreview.net/forum?id=YicbFdNTTy>
- Dosovitskiy, A., Springenberg, J. T., Riedmiller, M., & Brox, T. (2014). Discriminative unsupervised feature learning with convolutional neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 27). Curran Associates, Inc.
- Fontanella, A., G. M., J. W., E. T., & A. S. (2024). Diffusion models for counterfactual generation and anomaly detection in brain images. *IEEE transactions on medical imaging*. <https://doi.org/10.1109/TMI.2024.3460391>
- García-Lorenzo, D., Francis, S., Narayanan, S., Arnold, D. L., & Collins, D. L. (2013). Review of automatic segmentation methods of multiple sclerosis white matter lesions on conventional magnetic resonance imaging. *Medical Image Analysis*, *17*(1), 1–18. <https://doi.org/https://doi.org/10.1016/j.media.2012.09.004>
- Ghafoorian, M., Mehrtash, A., Kapur, T., Karssemeijer, N., Marchiori, E., Pesteie, M., Guttman, C. R. G., de Leeuw, F.-E., Tempany, C. M., van Ginneken, B., Fedorov, A., Abolmaesumi, P., Platel, B., & Wells, W. M. (2017). Transfer learning for domain adaptation in mri: Application in brain lesion segmentation. In M. Descoteaux, L. Maier-Hein, A. Franz, P. Jannin, D. L. Collins, & S. Duchesne (Eds.), *Medical image computing and computer assisted intervention - miccai 2017* (pp. 516–524). Springer International Publishing.
- Ghorbel, A., Aldahdooh, A., Albarqouni, S., & Hamidouche, W. (2023). Transformer Based Models for Unsupervised Anomaly Segmentation in Brain MR Images - brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries (S. Bakas, A. Crimi, U. Baid, S. Malec, M. Pytlarz, B. Baheti, M. Zenk, & R. Dorent, Eds.). *13769*, 25–44. [https://doi.org/10.1007/978-3-031-33842-7\\_3](https://doi.org/10.1007/978-3-031-33842-7_3)
- Gianattasio, K. Z., Bennett, E. E., Wei, J., Mehrotra, M. L., Mosley, T., Gottesman, R. F., Wong, D. F., Stuart, E. A., Griswold, M. E., Couper, D., Glymour, M. M., Power, M. C., & for the Alzheimer's Disease Neuroimaging Initiative. (2021). Generalizability of findings from a clinical sample to a community-based sample: A comparison of adni and aric. *Alzheimer's & Dementia*, *17*(8), 1265–1276. <https://doi.org/https://doi.org/10.1002/alz.12293>
- Gonzalez, R. C., & Woods, R. E. (2007, August). *Digital image processing* (3rd ed.). Pearson.

- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning* [<http://www.deeplearningbook.org>]. MIT Press.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Hassanally, R., Brianceau, C., Solal, M., Colliot, O., & Burgos, N. (2024). Evaluation of pseudo-healthy image reconstruction for anomaly detection with deep generative models: Application to brain FDG PET. <https://doi.org/10.48550/ARXIV.2401.16363>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- Hernandez Petzsche, M. R., de la Rosa, E., Hanning, U., Wiest, R., Valenzuela, W., Reyes, M., Meyer, M., Liew, S.-L., Kofler, F., Ezhov, I., Robben, D., Hutton, A., Friedrich, T., Zarth, T., Bürkle, J., Baran, T. A., Menze, B., Broocks, G., Meyer, L., ... Kirschke, J. S. (2022). Isles 2022: A multi-center magnetic resonance imaging stroke lesion segmentation dataset. *Scientific Data*, 9(1). <https://doi.org/10.1038/s41597-022-01875-5>
- Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33, 6840–6851.
- Huijben, E. M. C., Amirrajab, S., & Pluim, J. P. W. (2024). Enhancing Reconstruction-Based Out-of-Distribution Detection in Brain MRI with Model and Metric Ensembles. <https://doi.org/10.48550/ARXIV.2412.17586>
- Iqbal, H., Khalid, U., Hua, J., & Chen, C. (2023). Unsupervised Anomaly Detection in Medical Images Using Masked Diffusion Model. <https://doi.org/10.48550/ARXIV.2305.19867>
- Isensee, F., Wald, T., Ulrich, C., Baumgartner, M., Roy, S., Maier-Hein, K., & Jäger, P. F. (2024). Nnu-net revisited: A call for rigorous validation in 3d medical image segmentation. In M. G. Linguraru, Q. Dou, A. Feragen, S. Giannarou, B. Glocker, K. Lekadir, & J. A. Schnabel (Eds.), *Medical image computing and computer assisted intervention – miccai 2024* (pp. 488–498). Springer Nature Switzerland.
- Jiménez-García, A., García, H. F., Cárdenas-Peña, D. A., Cárdenas-Bedoya, W., Porrás-Hurtado, G. L., & Orozco-Gutiérrez, Á. A. (2024). Unsupervised Anomaly Detection by Learning Elastic Transformations Within an Autoencoder Approach - 2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), 1–4. <https://doi.org/10.1109/EMBC53108.2024.10781622>
- Kascenas, A., Sanchez, P., Schrepf, P., Wang, C., Clackett, W., Mikhael, S. S., Voisey, J. P., Goatman, K., Weir, A., Pugeault, N., Tsiftaris, S. A., & O'Neil, A. Q. (2023). The role of noise in denoising

- models for anomaly detection in medical images. *Medical Image Analysis*, 90, 102963. <https://doi.org/10.1016/j.media.2023.102963>
- Kingma, D. P., & Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Koller, D., & Friedman, N. (2009, July). *Probabilistic graphical models: Principles and techniques*. MIT Press.
- Kuijf, H., Biesbroek, M., de Bresser, J., Heinen, R., Chen, C., van der Flier, W., Barkhof, Viergever, M., & Biessels, G. J. (2022). Data of the White Matter Hyperintensity (WMH) Segmentation Challenge. <https://doi.org/10.34894/AECRSD>
- Kumar Trivedi, V., Sharma, B., & Balamurugan, P. (2024). MCDDPM: Multichannel Conditional Denoising Diffusion Model for Unsupervised Anomaly Detection in Brain MRI - 2024 17th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI), 1–6. <https://doi.org/10.1109/CISP-BMEI64163.2024.10906217>
- Lambert, B., Louis, M., Doyle, S., Forbes, F., Dojat, M., & Tucholka, A. (2021). Leveraging 3d Information In Unsupervised Brain Mri Segmentation - 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), 187–190. <https://doi.org/10.1109/ISBI48211.2021.9433894>
- LaMontagne, P. J., Benzinger, T. L., Morris, J. C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A. G., Raichle, M. E., Cruchaga, C., & Marcus, D. (2019). Oasis-3: Longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. *medRxiv*. <https://doi.org/10.1101/2019.12.13.19014902>
- Landman, B. A., Huang, A. J., Gifford, A., Vikram, D. S., Lim, I. A. L., Farrell, J. A., Bogovic, J. A., Hua, J., Chen, M., Jarso, S., Smith, S. A., Joel, S., Mori, S., Pekar, J. J., Barker, P. B., Prince, J. L., & van Zijl, P. C. (2011). Multi-parametric neuroimaging reproducibility: A 3-t resource study. *NeuroImage*, 54(4), 2854–2866. <https://doi.org/https://doi.org/10.1016/j.neuroimage.2010.11.047>
- Lee, J., Liu, C., Kim, J., Chen, Z., Sun, Y., Rogers, J. R., Chung, W. K., & Weng, C. (2022). Deep learning for rare disease: A scoping review. *Journal of Biomedical Informatics*, 135, 104227. <https://doi.org/https://doi.org/10.1016/j.jbi.2022.104227>
- Lesjak, Ž., Galimzianova, A., Koren, A., Lukin, M., Pernuš, F., Likar, B., & Špiclin, Ž. (2017). A novel public mr image dataset of multiple sclerosis patients with lesion segmentations based on multi-rater consensus. *Neuroinformatics*, 16(1), 51–63. <https://doi.org/10.1007/s12021-017-9348-7>
- Liew, S.-L., Lo, B., Donnelly, M. R., Zavaliangos-Petropulu, A., Jeong, J. N., Barisano, G., Hutton, A., Simon, J. P., Juliano, J. M., Suri, A., Ard, T., Banaj, N., Borich, M. R., Boyd, L. A., Brodtmann, A., Buetefisch, C. M., Cao, L., Cassidy, J. M., Ciullo, V., . . . Yu, C. (2021). A large, curated, open-source stroke neuroimaging dataset to improve lesion segmentation algorithms. *medRxiv*. <https://doi.org/10.1101/2021.12.09.21267554>

- Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., & Le, M. (2023). Flow matching for generative modeling. *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=PqvMRDCJT9t>
- Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., van der Laak, J. A., van Ginneken, B., & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis, 42*, 60–88. <https://doi.org/10.1016/j.media.2017.07.005>
- Littlejohns, T. J., Holliday, J., Gibson, L. M., Garratt, S., Oesingmann, N., Alfaro-Almagro, F., Bell, J. D., Boulton, C., Collins, R., Conroy, M. C., Crabbtree, N., Doherty, N., Frangi, A. F., Harvey, N. C., Leeson, P., Miller, K. L., Neubauer, S., Petersen, S. E., Sellors, J., . . . Allen, N. E. (2020). The uk biobank imaging enhancement of 100,000 participants: rationale, data collection, management and future directions. *Nature Communications, 11*(1), 2624. <https://doi.org/10.1038/s41467-020-15948-9>
- Liu, X., Gong, C., & qiang liu. (2023). Flow straight and fast: Learning to generate and transfer data with rectified flow. *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=XVjTT1nw5z>
- Lu, S., Zhang, W., Guo, J., Liu, H., Li, H., & Wang, N. (2024). PatchCL-AE: Anomaly detection for medical images using patch-wise contrastive learning-based auto-encoder. *Computerized Medical Imaging and Graphics, 114*, 102366. <https://doi.org/10.1016/j.compmedimag.2024.102366>
- Luo, G., Xie, W., Gao, R., Zheng, T., Chen, L., & Sun, H. (2023). Unsupervised anomaly detection in brain MRI: Learning abstract distribution from massive healthy brains. *Computers in Biology and Medicine, 154*, 106610. <https://doi.org/10.1016/j.combiomed.2023.106610>
- Lüth, C. T., Zimmerer, D., Koehler, G., Jaeger, P. F., Isenensee, F., & Maier-Hein, K. H. (2023). Contrastive Representations for Unsupervised Anomaly Detection and Localization - bildverarbeitung für die Medizin 2023 (T. M. Deserno, H. Handels, A. Maier, K. Maier-Hein, C. Palm, & T. Tolxdorff, Eds.), 246–252. [https://doi.org/10.1007/978-3-658-41657-7\\_54](https://doi.org/10.1007/978-3-658-41657-7_54)
- Ma, X., Fu, J., Liao, W., Zhang, S., & Wang, G. (2025). Clisc: Bridging clip and sam by enhanced cam for unsupervised brain tumor segmentation. *2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*, 1–5. <https://doi.org/10.1109/ISBI60581.2025.10980784>
- Marimont, S. N., & Tarroni, G. (2021). Anomaly detection through latent space restoration using vector quantized variational autoencoders. *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 1764–1767. <https://doi.org/10.1109/ISBI48211.2021.9433778>
- McDermott, M., Zhang, H., Hansen, L., Angelotti, G., & Gallifant, J. (2024). A closer look at auroc and auprc under class imbalance. *Advances in Neural Information Processing Systems, 37*, 44102–44163. <https://arxiv.org/abs/2401.06091>

- Meissen, F., Paetzold, J., Kaissis, G., & Rueckert, D. (2023). Unsupervised Anomaly Localization with Structural Feature-Autoencoders - brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries (S. Bakas, A. Crimi, U. Baid, S. Malec, M. Pytlarz, B. Baheti, M. Zenk, & R. Dorent, Eds.). *13769*, 14–24. [https://doi.org/10.1007/978-3-031-33842-7\\_2](https://doi.org/10.1007/978-3-031-33842-7_2)
- Meissen, F., Wiestler, B., Kaissis, G., & Rueckert, D. (2022). On the pitfalls of using the residual error as anomaly score. In E. Konukoglu, B. Menze, A. Venkataraman, C. Baumgartner, Q. Dou, & S. Albarqouni (Eds.), *Proceedings of the 5th international conference on medical imaging with deep learning* (pp. 914–928, Vol. 172). PMLR. <https://proceedings.mlr.press/v172/meissen22a.html>
- Menze, B. H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber, M.-A., Arbel, T., Avants, B. B., Ayache, N., Buendia, P., Collins, D. L., Cordier, N., . . . Van Leemput, K. (2015). The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging*, *34*(10), 1993–2024. <https://doi.org/10.1109/TMI.2014.2377694>
- Muñoz-Ramírez, V., Kmetzsch, V., Forbes, F., Meoni, S., Moro, E., & Dojat, M. (2022). Subtle anomaly detection: Application to brain MRI analysis of de novo Parkinsonian patients. *Artificial Intelligence in Medicine*, *125*, 102251. <https://doi.org/10.1016/j.artmed.2022.102251>
- Muñoz-Ramírez, V., Pinon, N., Forbes, F., Lartizen, C., & Dojat, M. (2021). Patch vs. global image-based unsupervised anomaly detection in mr brain scans of early parkinsonian patients. In A. Abdulkadir, S. M. Kia, M. Habes, V. Kumar, J. M. Rondina, C. Tax, & T. Wolfers (Eds.), *Machine learning in clinical neuroimaging* (pp. 34–43). Springer International Publishing.
- Nguyen, B., Feldman, A., Bethapudi, S., Jennings, A., & Willcocks, C. G. (2021). Unsupervised Region-Based Anomaly Detection In Brain MRI With Adversarial Image Inpainting - 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), 1127–1131. <https://doi.org/10.1109/ISBI48211.2021.9434115>
- Oord, A. v. d., Vinyals, O., & Kavukcuoglu, K. (2017). Neural discrete representation learning. <https://doi.org/10.48550/ARXIV.1711.00937>
- Ouzzani, M., Hammady, H., Fedorowicz, Z., & Elmagarmid, A. (2016). Rayyan—a web and mobile app for systematic reviews. *Systematic Reviews*, *5*(1). <https://doi.org/10.1186/s13643-016-0384-4>
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., . . . Moher, D. (2021). The prisma 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, *372*. <https://doi.org/10.1136/bmj.n71>
- Pang, G., Shen, C., Cao, L., & Hengel, A. V. D. (2021). Deep learning for anomaly detection: A review. *ACM Comput. Surv.*, *54*(2). <https://doi.org/10.1145/3439950>

- Pernet, C., Gorgolewski, K., & Ian, W. (2016). A neuroimaging dataset of brain tumour patients. <https://doi.org/10.5255/UKDA-SN-851861>
- Pinaya, W. H. L., Graham, M. S., Gray, R., Da Costa, P. F., Tudosiu, P.-D., Wright, P., Mah, Y. H., MacKinnon, A. D., Teo, J. T., Jager, R., Werring, D., Rees, G., Nachev, P., Ourselin, S., & Cardoso, M. J. (2022). Fast Unsupervised Brain Anomaly Detection and Segmentation with Diffusion Models - medical Image Computing and Computer Assisted Intervention – MICCAI 2022 (L. Wang, Q. Dou, P. T. Fletcher, S. Speidel, & S. Li, Eds.). *13438*, 705–714. [https://doi.org/10.1007/978-3-031-16452-1\\_67](https://doi.org/10.1007/978-3-031-16452-1_67)
- Pinaya, W. H. L., Tudosiu, P.-D., Dafflon, J., Da Costa, P. F., Fernandez, V., Nachev, P., Ourselin, S., & Cardoso, M. J. (2022). Brain imaging generation with latent diffusion models. In A. Mukhopadhyay, I. Oksuz, S. Engelhardt, D. Zhu, & Y. Yuan (Eds.), *Deep generative models* (pp. 117–126). Springer Nature Switzerland. [https://doi.org/10.1007/978-3-031-18576-2\\_12](https://doi.org/10.1007/978-3-031-18576-2_12)
- Pinaya, W. H., Tudosiu, P.-D., Gray, R., Rees, G., Nachev, P., Ourselin, S., & Cardoso, M. J. (2022). Unsupervised brain imaging 3d anomaly detection and segmentation with transformers. *Medical Image Analysis*, *79*, 102475. <https://doi.org/https://doi.org/10.1016/j.media.2022.102475>
- Pinho, A. L. G., Hertz-Pannier, L., & Thirion, B. (2018). "individual brain charting". OpenNeuro. <https://doi.org/10.18112/openneuro.ds000244.v1.0.0>
- Puccio, B., Pooley, J. P., Pellman, J. S., Taverna, E. C., & Craddock, R. C. (2016). The preprocessed connectomes project repository of manually corrected skull-stripped t1-weighted anatomical mri data. *GigaScience*, *5*(1), s13742-016-0150–5. <https://doi.org/10.1186/s13742-016-0150-5>
- Qu, G., Pan, H., Yang, C., & Lin, Y. (2026). Unsupervised anomaly detection for medical image classification using masked denoising autoencoder. *Applied Soft Computing*, *188*, 114388. <https://doi.org/https://doi.org/10.1016/j.asoc.2025.114388>
- Rahmaniar, W., & Suzuki, K. (2026). Multi-ad: Cross-domain unsupervised anomaly detection for medical and industrial applications. *Pattern Recognition*, *172*, 112486. <https://doi.org/https://doi.org/10.1016/j.patcog.2025.112486>
- Rezende, D., & Mohamed, S. (2015). Variational inference with normalizing flows. *International conference on machine learning*, 1530–1538. <https://proceedings.mlr.press/v37/rezende15.html>
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. M. Wells, & A. F. Frangi (Eds.), *Medical image computing and computer-assisted intervention – miccai 2015* (pp. 234–241). Springer International Publishing. [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
- Sathyanarayanan, S., & Muthukumaravel, A. (2025). Schizonet: A self-supervised learning architecture for unsupervised anomaly detection in schizophrenia fmri scans. *2025 International Conference*

- on *Smart & Sustainable Technology (INCSST)*, 1–6. <https://doi.org/10.1109/INCSST64791.2025.11210391>
- Sato, K., Hama, K., Matsubara, T., & Uehara, K. (2019). Predictable Uncertainty-Aware Unsupervised Deep Anomaly Segmentation - 2019 International Joint Conference on Neural Networks (IJCNN), 1–7. <https://doi.org/10.1109/IJCNN.2019.8852144>
- Schlegl, T., P, S., SM, W., G, L., & U, S.-E. (2019). F-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54, 30–44. <https://doi.org/10.1016/j.media.2019.01.010>
- Schlegl, T., Seeböck, P., Waldstein, S. M., Schmidt-Erfurth, U., & Langs, G. (2017). Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. <https://arxiv.org/abs/1703.05921>
- Shafto, M. A., Tyler, L. K., Dixon, M., Taylor, J. R., Rowe, J. B., Cusack, R., Calder, A. J., Marslen-Wilson, W. D., Duncan, J., Dalgleish, T., Henson, R. N., Brayne, C., Matthews, F. E., & Cam-CAN. (2014). The cambridge centre for ageing and neuroscience (cam-can) study protocol: A cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. *BMC Neurology*, 14(1), 204. <https://doi.org/10.1186/s12883-014-0204-1>
- Simarro, J., de la Rosa, E., Vyvere, T. V., Robben, D., & Sima, D. M. (2020). Unsupervised 3D Brain Anomaly Detection. <https://doi.org/10.48550/ARXIV.2010.04717>
- Solal, M., Hassanaly, R., & Burgos, N. (2023). Leveraging healthy population variability in deep learning unsupervised anomaly detection in brain FDG PET. <https://doi.org/10.48550/ARXIV.2311.12081>
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., & Poole, B. (2021). Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations*. <https://openreview.net/forum?id=PXTIG12RRHS>
- Steyerberg, E. W., Wiegers, E., Sewalt, C., Buki, A., Citerio, G., De Keyser, V., Ercole, A., Kunzmann, K., Lanyon, L., Lecky, F., Lingsma, H., Manley, G., Nelson, D., Peul, W., Stocchetti, N., von Steinbüchel, N., Vande Vyvere, T., Verheyden, J., Wilson, L., ... Zoerle, T. (2019). Case-mix, care pathways, and outcomes in patients with traumatic brain injury in center-tbi: A european prospective, multicentre, longitudinal, cohort study. *The Lancet Neurology*, 18(10), 923–934. [https://doi.org/10.1016/S1474-4422\(19\)30232-7](https://doi.org/10.1016/S1474-4422(19)30232-7)
- Tang, P., Yan, X., Hu, X., Cui, Y., Luo, D., Zhang, J., Xu, P., Peng, J., He, Q., Huang, F., Xue, S., & Lasser, T. (2025). Shortcutbreaker: Low-rank noisy bottleneck with global perturbation attention for multi-class unsupervised anomaly detection. <https://arxiv.org/abs/2510.18342>
- Tricco, A. C., Lillie, E., Zarin, W., O'Brien, K. K., Colquhoun, H., Levac, D., Moher, D., Peters, M. D., Horsley, T., Weeks, L., Hempel, S., Akl, E. A., Chang, C., McGowan, J., Stewart, L., Hartling, L.,

- Aldcroft, A., Wilson, M. G., Garritty, C., . . . Straus, S. E. (2018). Prisma extension for scoping reviews (prisma-scr): Checklist and explanation. *Annals of Internal Medicine*, *169*(7), 467–473. <https://doi.org/10.7326/m18-0850>
- Tschuchnig, M. E., & Gadermayr, M. (2022). Anomaly detection in medical imaging - a mini review. In P. Haber, T. J. Lampoltshammer, H. Leopold, & M. Mayr (Eds.), *Data science – analytics and applications* (pp. 33–38). Springer Fachmedien Wiesbaden. [https://doi.org/10.1007/978-3-658-36295-9\\_5](https://doi.org/10.1007/978-3-658-36295-9_5)
- Uzunova, H., Schultz, S., Handels, H., & Ehrhardt, J. (2019). Unsupervised pathology detection in medical images using conditional variational autoencoders. *International Journal of Computer Assisted Radiology and Surgery*, *14*(3), 451–461. <https://doi.org/10.1007/s11548-018-1898-0>
- Van Essen, D., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S., Della Penna, S., Feinberg, D., Glasser, M., Harel, N., Heath, A., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., . . . Yacoub, E. (2012). The human connectome project: A data acquisition perspective [Connectivity]. *NeuroImage*, *62*(4), 2222–2231. <https://doi.org/https://doi.org/10.1016/j.neuroimage.2012.02.018>
- Vemuri, P., Decarli, C., & Duering, M. (2022). Imaging markers of vascular brain health: Quantification, clinical implications, and future directions. *Stroke*, *53*(2), 416–426. <https://doi.org/10.1161/STROKEAHA.120.032611>
- Villanueva-Meyer, J. E., Mabray, M. C., & Cha, S. (2017). Current clinical brain tumor imaging. *Neurosurgery*, *81*(3), 397–415. <https://doi.org/10.1093/neuros/nyx103>
- Walsh, R., Meurée, C., Kerbrat, A., Masson, A., Hussein, B. R., Gaubert, M., Galassi, F., & Combés, B. (2023). Expert variability and deep learning performance in spinal cord lesion segmentation for multiple sclerosis patients. *2023 IEEE 36th International Symposium on Computer-Based Medical Systems (CBMS)*, 463–470.
- Wang, R., Bashyam, V., Yang, Z., Yu, F., Tassopoulou, V., Chintapalli, S. S., Skampardon, I., Sreepada, L. P., Sahoo, D., Nikita, K., Abdulkadir, A., Wen, J., & Davatzikos, C. (2023). Applications of generative adversarial networks in neuroimaging and clinical neuroscience. *NeuroImage*, *269*, 119898. <https://doi.org/https://doi.org/10.1016/j.neuroimage.2023.119898>
- Wijanarko, H., Calista, E., Chen, L.-F., & Chen, Y.-S. (2024). Tri-VAE: Triplet Variational Autoencoder for Unsupervised Anomaly Detection in Brain Tumor MRI - 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 3930–3939. <https://doi.org/10.1109/CVPRW63382.2024.00397>
- Wu, X., Bi, L., Fulham, M., Feng, D. D., Zhou, L., & Kim, J. (2021). Unsupervised brain tumor segmentation using a symmetric-driven adversarial network. *Neurocomputing*, *455*, 242–254. <https://doi.org/10.1016/j.neucom.2021.05.073>

- Wyatt, J., Leach, A., Schmon, S. M., & Willcocks, C. G. (2022). Anoddpm: Anomaly detection with denoising diffusion probabilistic models using simplex noise - 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 649–655. <https://doi.org/10.1109/CVPRW56347.2022.00080>
- Xu, Y., Quan, R., Xu, W., Huang, Y., Chen, X., & Liu, F. (2024). Advances in medical image segmentation: A comprehensive review of traditional, deep learning and hybrid approaches. *Bioengineering*, *11*(10), 1034. <https://doi.org/10.3390/bioengineering11101034>
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., & Wang, O. (2018). The unreasonable effectiveness of deep features as a perceptual metric. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 586–595. <https://doi.org/10.1109/CVPR.2018.00068>
- Zhao, R., Yaman, B., Zhang, Y., Stewart, R., Dixon, A., Knoll, F., Huang, Z., Lui, Y. W., Hansen, M. S., & Lungren, M. P. (2022). Fastmri+, clinical pathology annotations for knee and brain fully sampled magnetic resonance imaging data. *Scientific Data*, *9*(1), 152. <https://doi.org/10.1038/s41597-022-01255-z>
- Zimmerer, D., Isensee, F., Petersen, J., Kohl, S., & Maier-Hein, K. (2019). Unsupervised Anomaly Localization using Variational Auto-Encoders. <https://doi.org/10.48550/ARXIV.1907.02796>
- Zimmerer, D., Petersen, J., Köhler, G., Jäger, P., Full, P., Roß, T., Adler, T., Reinke, A., Maier-Hein, L., & Maier-Hein, K. (2020, March). Medical out-of-distribution analysis challenge. <https://doi.org/10.5281/zenodo.3961376>