



FlowLet: Conditional 3D brain MRI synthesis using wavelet flow matching

Danilo Danese^a, Angela Lombardi^a, Matteo Attimonelli^{a,b}, Giuseppe Fasano^a,
Tommaso Di Noia^a

^a Politecnico di Bari, Bari, Italy

^b Sapienza University of Rome, Rome, Italy

ARTICLE INFO

Keywords:

Flow matching
MRI
3D brain synthesis
Generative models
Age-conditioned generation
Deep learning

ABSTRACT

Generative modeling for 3D brain MRI is challenged by a trade-off between anatomical fidelity, sample diversity, and computational efficiency. Diffusion-based approaches achieve strong visual quality but typically require hundreds to thousands of sampling steps, while latent-space compression can introduce reconstruction artifacts and degrade fine-grained anatomy. We introduce *FlowLet*, a conditional generative framework that performs Flow Matching in an invertible 3D wavelet domain. This representation enables multi-scale generation without learned latent compression, while deterministic ODE sampling allows fast inference. Age conditioning is modeled through complementary feature-wise modulation and spatially adaptive cross-attention, enabling explicit control over age-related morphological variation. Across multi-site neuroimaging datasets, FlowLet achieves competitive and, in several settings, superior global fidelity compared to diffusion-based baselines using as few as 10 sampling steps. Region-based evaluation across 95 cortical and subcortical brain regions demonstrates improved local anatomical plausibility beyond what is captured by global similarity metrics alone. In a downstream brain age prediction study, models augmented with FlowLet-generated data consistently reduce prediction error relative to real-only training and other generative baselines. Rather than focusing on a single dominant metric improvement, these results highlight a consistent trade-off between efficiency, controllability, and anatomically meaningful 3D brain MRI generation. The proposed framework is released as open-source to support reproducibility.

1. Introduction

Magnetic Resonance Imaging (MRI) provides a non-invasive, high-resolution representation of brain anatomy and is routinely used to study structural changes associated with development, aging, and neurological disorders (Bacon et al., 2024). In recent years, learning-based analysis of brain MRI has enabled quantitative biomarkers that capture subtle morphological variations beyond traditional volumetric measures (Rahman et al., 2025; Fu et al., 2025).

A prominent example is *Brain Age Prediction* (BAP), a data-driven framework that estimates an individual's biological brain age from structural MRI by learning population-level associations between anatomical features and chronological age (Cole et al., 2017; Baecker et al., 2021; Lombardi et al., 2021; Peng et al., 2021). The discrepancy between predicted and chronological age, commonly referred to as the brain-age gap, has been linked to cognitive decline, neurodegenerative disease, and altered aging trajectories, making BAP a clinically relevant biomarker (Gaser et al., 2024).

The performance and reliability of BAP models strongly depend on the availability of large, diverse, and age-balanced MRI datasets spanning the full lifespan. However, in practice, publicly available 3D brain MRI datasets exhibit substantial demographic imbalances. Certain age ranges, particularly young and middle-aged adults, are overrepresented, while pediatric and elderly populations are often under-sampled (Bashyam et al., 2020). Although large-scale cohorts exist (Sudlow et al., 2015), restricted access and acquisition costs limit their widespread use. These data limitations reduce generalizability, introduce age-related biases, and hinder the deployment of BAP models in clinical and epidemiological settings (Dinsdale et al., 2021).

Given the cost, logistical complexity, and ethical constraints associated with acquiring additional MRI data, synthetic data generation has emerged as a complementary strategy for dataset enrichment. In this context, *generative modeling* aims to learn the underlying distribution of brain MRIs and synthesize anatomically plausible volumes that can be used for augmentation, bias mitigation, and robustness

* Corresponding authors.

E-mail addresses: danilo.danese@poliba.it (D. Danese), angela.lombardi@poliba.it (A. Lombardi), matteo.attimonelli@poliba.it (M. Attimonelli), giuseppe.fasano@poliba.it (G. Fasano), tommaso.dinoia@poliba.it (T. Di Noia).

<https://doi.org/10.1016/j.media.2026.104161>

Received 8 February 2026; Received in revised form 15 April 2026; Accepted 6 June 2026

Available online 15 June 2026

1361-8415/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

improvement (Chintapalli et al., 2024). When properly designed, generative models offer a scalable mechanism to increase data diversity, particularly for underrepresented age groups.

However, generating realistic 3D brain MRIs remains technically challenging. Volumetric neuroimaging data are characterized by extremely high dimensionality and complex multi-scale anatomical structure (Fernandez et al., 2024; Yu et al., 2025). State-of-the-art generative approaches, including autoencoding architectures and diffusion models, often face trade-offs between computational efficiency and anatomical fidelity when applied to full-resolution 3D volumes (Dhariwal and Nichol, 2021). To alleviate computational burden, many methods rely on learned latent representations, which can introduce reconstruction artifacts and obscure fine-grained anatomical details that are critical for age-related analysis (Müller-Franzes et al., 2023).

Beyond realism, effective use of synthetic MRI in BAP requires *controlled generation*. In particular, models must be conditioned on age information such that synthesized volumes reflect meaningful and localized morphological changes associated with aging (Peng et al., 2024). Designing conditioning mechanisms that influence generation at multiple spatial scales while preserving anatomical coherence remains an open problem. Insufficient or overly coarse conditioning often results in weak age specificity or entangled anatomical effects, limiting downstream utility.

These challenges are closely related to a fundamental limitation in modern generative modeling, commonly referred to as the *generative modeling trilemma* (Xiao et al., 2022). The trilemma formalizes the trade-off between sample quality, diversity, and sampling efficiency, stating that improving one objective typically degrades at least one of the others. In the context of medical imaging, where high anatomical fidelity, population-level variability, and practical inference times are all required, mitigating this trade-off is particularly critical.

In this work, we introduce *FlowLet*, a conditional generative framework for 3D brain MRI synthesis designed to address these challenges. FlowLet combines Flow Matching with an invertible wavelet-based representation, enabling multi-scale modeling of volumetric brain anatomy without relying on learned latent compression. This design supports anatomically structured generation while remaining computationally tractable. Age information is incorporated through complementary feature-wise and spatial conditioning mechanisms, providing explicit control over age-related morphological variations. By leveraging Flow Matching, FlowLet further supports efficient sample generation and offers a practical alternative to the diffusion-based baseline implementations considered in this study.

We evaluate FlowLet on three neuroimaging datasets spanning over 12 multi-site sources to assess robustness to acquisition variability and demographic heterogeneity. Beyond standard global similarity metrics (FID, MMD, MS-SSIM), which are known to be limited in volumetric neuroimaging due to background dominance, we develop a comprehensive evaluation protocol. This includes (i) a downstream BAP task to assess functional utility as a clinically meaningful biomarker, and (ii) region-based anatomical metrics across 95 brain regions to quantify localized structural fidelity. This combined evaluation reveals cases where favorable global metrics fail to capture anatomically relevant errors. To support reproducibility and practical adoption, we release the complete open-source implementation of FlowLet.

Our main contributions are:

1. A publicly available wavelet-based Flow Matching framework for efficient and anatomically faithful 3D brain MRI synthesis;
2. A multi-level age-conditioning design that combines feature-wise and spatial mechanisms for localized morphological control;
3. A region- and task-aware evaluation protocol for generative brain MRI, integrating anatomical fidelity and downstream brain age prediction.

2. Related work

2.1. Augmentation

Despite an increase in general-purpose data availability, 3D neuroimaging remains constrained by small cohorts and high-dimensional voxel data, limiting population diversity and statistical reliability (Button et al., 2013). Deep learning models, which require large and diverse datasets to generalize well, are particularly affected. Data augmentation is thus essential for improving model robustness by artificially expanding training sets. Traditional methods such as adding noise, cropping, flipping or elastic transform, preserve labels but limit variability and, more importantly, they raise the risk of introducing distorted anatomical structures or amplify biases (Wu and Suk, 2017; Shorten and Khoshgoufar, 2019). These risks are especially pronounced in medical imaging, where anatomical fidelity is critical and validation often relies on expert assessment. Moreover, class imbalance remains a persistent challenge: oversampling strategies, though common, scale poorly in high-dimensional spaces and may fail to capture minority class variation (Nguyen et al., 2011; Blagus and Lusa, 2013). These limitations require specialized methods, i.e., synthetic data generation.

2.2. Generative models for 3D synthesis

Generative modeling has become a key approach for synthetic data generation in medical imaging, enabling the synthesis of anatomically plausible samples beyond simple geometric or intensity-based augmentations. Early deep generative models, such as Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) and Variational Autoencoders (VAEs) (Kingma and Welling, 2014), laid the foundation for data-driven image synthesis. GANs are known for producing sharp samples but often suffer from training instability and mode collapse (Salimans et al., 2016), while VAEs offer stable training and efficient sampling at the expense of overly smooth reconstructions (Kingma and Welling, 2014). These limitations become particularly pronounced when extending such models to high-resolution 3D neuroimaging data.

Denosing Diffusion Models (DDMs) (Ho et al., 2020) have recently achieved state-of-the-art performance across a wide range of imaging domains, including medical and neuroimaging applications (Wyatt et al., 2022; Pinaya et al., 2022; Wu et al., 2023; Durrer et al., 2023). These models formulate generation as the reversal of a gradual noising process, enabling high-quality and diverse sample synthesis. However, diffusion-based generation typically requires solving stochastic differential equations through hundreds or thousands of iterative steps (J. Song et al., 2021; Davtyan et al., 2023), resulting in substantial computational overhead. This iterative sampling becomes a major bottleneck when scaling to high-resolution 3D volumes.

To improve sampling efficiency, several approaches operate in compressed or discrete latent spaces. Vector-Quantized VAEs (VQ-VAEs) (van den Oord et al., 2017) and Latent Diffusion Models (LDMs) (Rombach et al., 2022) reduce computational costs by performing diffusion in a lower-dimensional representation. While effective in improving efficiency and global coherence, these methods introduce additional sources of complexity, such as learned compression artifacts or increased training overhead, which may degrade fine anatomical details. Extensions targeting ultra-high-resolution synthesis rely on specialized scheduling and hierarchical strategies (Zhang et al., 2025a), further increasing system complexity.

Flow Matching (FM) has recently emerged as an alternative paradigm for generative modeling that addresses the inefficiency of diffusion-based sampling (Lipman et al., 2023; Albergo and Vanden-Eijnden, 2023; Liu et al., 2023; Neklyudov et al., 2023). Instead of learning a stochastic denoising process, FM models a continuous-time velocity field that transports samples from a simple prior to the target data distribution by solving an ordinary differential equation. By encouraging straighter transport trajectories (Lee et al., 2023), FM

substantially reduces the number of inference steps required for sample generation, making it particularly attractive for high-resolution volumetric synthesis. Recent works have further explored conditional control and efficiency improvements within the FM framework (Gagneux et al., 2025a).

Wavelet Diffusion Models (WDMs) provide a learning-free alternative to latent diffusion for dimensionality reduction in generative modeling (Phung et al., 2023; Friedrich et al., 2024; Zhang et al., 2025b). By applying a fixed wavelet transform, WDMs decompose images into multi-resolution frequency components and perform diffusion directly in the wavelet domain. This approach significantly reduces memory requirements while preserving spatial structure and avoiding artifacts introduced by learned compression.

Despite these advantages, existing 3D WDM formulations (Friedrich et al., 2024) considered here rely on iterative diffusion sampling, in which inference requires hundreds to thousands of sequential denoising evaluations along the reverse trajectory. Each evaluation carries the full computational cost of a forward pass through a 3D U-Net, making the cumulative inference cost a significant bottleneck in high-resolution volumetric settings. The integration of more efficient generative paradigms, such as Flow Matching, within a wavelet-based representation remains largely unexplored. A conceptually related direction combining wavelets with normalizing flows has been investigated for high-resolution 2D image synthesis (Yu et al., 2020), but extensions to conditional 3D neuroimaging remain limited.

Building on the literature discussed above, FlowLet addresses several open gaps in generative modeling for 3D neuroimaging. While diffusion-based models achieve high sample quality, their iterative sampling limits scalability to high-resolution volumes. Latent and wavelet-based approaches mitigate computational costs but often remain tied to diffusion sampling or introduce compression artifacts. In parallel, Flow Matching offers efficient generation but has seen limited adoption in volumetric medical imaging and has not been explored in conjunction with wavelet representations. FlowLet bridges these directions by integrating Flow Matching within an invertible wavelet domain, enabling efficient, anatomically faithful, and controllable 3D brain MRI synthesis. In addition, FlowLet explicitly targets age-conditioned generation and introduces a task- and region-aware evaluation protocol, addressing limitations in both conditioning strategies and assessment practices in prior work.

3. Proposed method

3.1. System overview

FlowLet is a conditional generative framework for 3D brain MRI synthesis that operates entirely in the wavelet domain. The overall, single-stage architecture pipeline illustrated in Fig. 1, consists of three main blocks: (i) wavelet-based decomposition of volumetric MRI data, (ii) flow matching-based generative modeling in the multi-scale frequency domain, and (iii) reconstruction of the synthesized volume via inverse wavelet transform.

During training, each input MRI volume is first transformed using an invertible 3D Haar discrete wavelet transform (DWT), yielding one low-frequency subband capturing coarse anatomical structure and seven high-frequency subbands encoding localized fine details. A conditional neural network is then trained to predict continuous-time velocity fields that transport samples from Gaussian noise to the target data distribution in this wavelet space. Conditioning variables, such as age, are injected into the network to explicitly control age-dependent morphological characteristics during generation.

At inference time, synthesis begins from Gaussian noise in the wavelet domain and proceeds by solving an ordinary differential equation defined by the learned velocity field. The resulting wavelet coefficients are finally mapped back to the voxel domain through the inverse DWT, producing a full-resolution 3D brain MRI. This design enables efficient sampling, multi-scale anatomical modeling, and explicit control over clinically relevant attributes, while avoiding the artifacts and overhead associated with learned latent compression.

3.2. Preliminaries

To achieve tractable learning while maintaining anatomical fidelity, the orthonormal Haar Discrete Wavelet Transform (DWT) (Mallat, 1999) is adopted. This perfectly invertible transform decomposes 3D volumes into frequency components, reducing dimensionality while preserving structural information up to numerical precision (Bullmore et al., 2004).

Given a 3D volume $x \in \mathbb{R}^{D \times H \times W}$, where D , H , and W denote the three spatial dimensions of the volume, the Haar DWT applies 1D low-pass and high-pass filters sequentially along each axis. Following the preprocessing convention adopted in this work, which aligns all volumes to a common reference space, these axes correspond to the anterior-posterior, superior-inferior, and left-right directions, respectively. Specifically, the Haar analysis filters are

$$\mathbf{L} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \end{bmatrix}, \quad \mathbf{H} = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 & 1 \end{bmatrix}. \quad (1)$$

This produces eight frequency subbands,

$$\mathcal{B} = \{\text{LLL}, \text{LLH}, \text{LHL}, \text{LHH}, \text{HLL}, \text{HLH}, \text{HHL}, \text{HHH}\}, \quad (2)$$

where each letter indicates whether a Low- or High-pass filter is applied along the corresponding axis. The subbands are then concatenated into an 8-channel tensor,

$$x_w = \mathcal{W}(x) \in \mathbb{R}^{8 \times \frac{D}{2} \times \frac{H}{2} \times \frac{W}{2}}, \quad (3)$$

where $\mathcal{W}$ denotes the forward DWT. The transform is perfectly invertible via the inverse DWT,

$$\mathcal{I}DWT = \mathcal{W}^{-1}(x_w) = x. \quad (4)$$

Due to the orthonormality of the Haar basis (Parseval's theorem for orthonormal wavelets (Mallat, 1989; Daubechies, 1992)), energy preservation and perfect reconstruction are guaranteed across subbands.

3.3. FlowLet framework

FlowLet builds upon the wavelet representation by explicitly modeling generation in the multi-scale frequency domain. As shown in Fig. 1 (left), the transformed 3D MRI volume is factorized into a low-frequency approximation subband (LLL) and seven high-frequency detail subbands (LLH, LHL, ..., HHH). The LLL subband captures the dominant anatomical structure, while the remaining components isolate fine-grained, spatially localized details. As the signal energy is predominantly concentrated in the LLL subband and total energy is preserved across subbands by Parseval's theorem, interpolation trajectories in this coarse subspace are inherently smoother, supporting stable learning dynamics and reconstruction (Lipman et al., 2023). At the same time, neural networks can specialize across subbands, enhancing stability and convergence in generative frameworks (Ho et al., 2020; Vahdat et al., 2021). This frequency-based factorization also reduces global variability and overfitting by promoting spatial locality (Mallat, 1999). In practice, this representation allows the generative network to operate jointly over the eight wavelet subbands at reduced spatial resolution, yielding a compact multi-scale parameterization for volumetric synthesis. Although this reduces spatial resolution, the transform is invertible and no information is discarded, as the eight subbands jointly encode the original signal. In addition, the effect of this representation on low- and high-frequency components is explicitly analyzed in Appendices A.9 and A.10. Code is available at <https://github.com/sisinflab/FlowLet>.

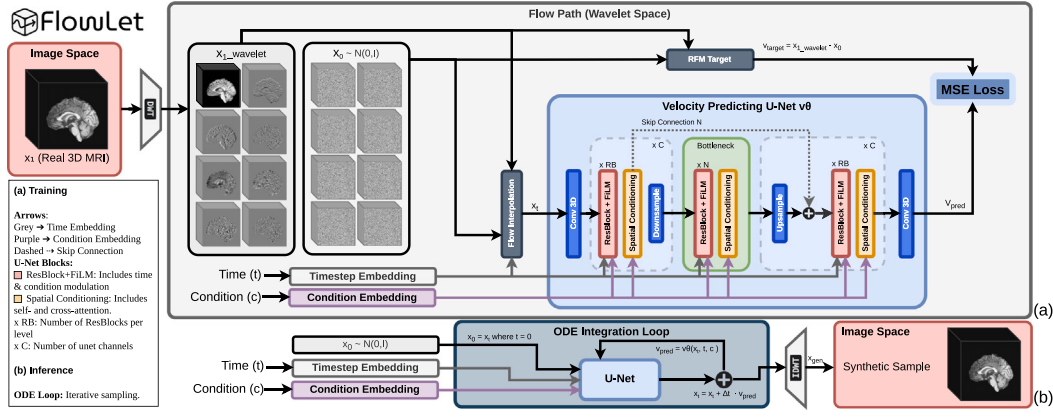


Fig. 1. FlowLet overview. **(a) Training:** a preprocessed 3D MRI volume x is mapped to the invertible 3D Haar wavelet domain via the DWT, producing eight subbands (LLL + seven detail bands). The conditional 3D U-Net $v_\theta(x_t, t, c)$ (condition c : age) is trained to predict the FM velocity field in wavelet space along the interpolation path x_t between noise x_0 and data x_1 . **(b) Inference:** starting from Gaussian noise x_0 in wavelet space and a target condition c , samples are generated by integrating the learned ODE $\partial_t x_t = v_\theta(x_t, t, c)$ with a small number of solver steps; the final wavelet coefficients are mapped back to voxel space with the inverse DWT (IDWT) to obtain the synthesized MRI.

3.3.1. Flow matching formulations

The framework supports multiple flow matching formulations to enable flexible and adaptable modeling, all operating in the wavelet domain, where samples x_t , noise x_0 , data x_1 , and velocity fields v (the instantaneous time-derivative $\partial_t x_t$ along the path) belong to $\mathbb{R}^{8 \times \frac{D}{2} \times \frac{H}{2} \times \frac{W}{2}}$. The data sample x_1 is the wavelet transform $\mathcal{W}(x_1^{\text{voxel}})$ of an original sample $x_1^{\text{voxel}} \sim p_{\text{data}}$, and noise x_0 is standard Gaussian in wavelet space. FlowLet supports a modular family of FM strategies that define different continuous-time interpolation paths x_t from noise to data in the wavelet domain. Each variant specifies a target velocity field $v_{\text{target}}(x_t, t)$. The model learns a parameterized velocity network $v_\theta(x_t, t, c)$, conditioned on signal c , by minimizing the MSE loss:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{x_t, t, v_{\text{target}}} \left[\left\| v_\theta(x_t, t, c) - v_{\text{target}}(x_t, t) \right\|^2 \right]. \quad (5)$$

The learned velocity field governs a deterministic ordinary differential equation (ODE) on $t \in [0, 1]$, solved via Euler integration.

The selected formulations, chosen for their foundational role in the FM literature and their progressively increasing trajectory curvature, are ordered to reflect differences in the geometry of the interpolation path, an aspect known to influence the expressiveness and stability of the learned velocity field (Lee et al., 2023). This setup enables a systematic evaluation of how curvature impacts training stability and synthesis quality. An overview of the implemented FM variants is provided.

Rectified Flow Matching (RFM). RFM (Liu et al., 2023) performs linear interpolation between a Gaussian noise sample and a data sample in the wavelet domain:

$$x_t = (1 - t)x_0 + tx_1, \quad v_{\text{target}} = x_1 - x_0. \quad (6)$$

The target velocity field is constant along the straight line path, resulting in zero path curvature, which promotes stable training and yields low-variance gradient estimates.

Conditional Flow Matching (CFM). CFM (Lipman et al., 2023; Albergo and Vanden-Eijnden, 2023) constructs a linear path between a data sample x_1 and a sampled noise x_0 in wavelet space for each training instance. The target velocity field is then defined as the instantaneous direction from x_t to x_1 , scaled by the remaining time:

$$x_t = (1 - t)x_0 + tx_1, \quad v_{\text{target}} = \frac{x_1 - x_t}{1 - t + \epsilon}, \quad (7)$$

where a small $\epsilon > 0$ is added to prevent divergence as $t \rightarrow 1$. While the underlying path is a straight line, the target velocity field is explicitly dependent on the current state x_t and time t , making it more dynamic than the constant velocity of RFM. This introduces non-zero curvature and increases sensitivity as t approaches 1.

Variance-Preserving diffusion matching (VP). Inspired by Variance-Preserving diffusion and the corresponding deterministic probability-flow formulation (Y. Song et al., 2021; J. Song et al., 2021), VP defines a nonlinear interpolation from data to noise governed by a linear variance schedule $\beta(t) = \beta_{\min} + t(\beta_{\max} - \beta_{\min})$, $t \in [0, 1]$. Signal and noise scaling coefficients are:

$$\bar{\alpha}(t) = \exp \left(-\frac{1}{2} \int_0^t \beta(s) ds \right), \quad \sigma(t) = \sqrt{1 - \bar{\alpha}(t)^2}. \quad (8)$$

The forward noising process generates intermediate samples x_t via interpolation between a data sample x_1 and standard Gaussian noise $\xi \sim \mathcal{N}(0, \mathbf{I})$, while the corresponding target velocity field $v_{\text{target}}(x_t, t)$, governing the reverse-time dynamics, is defined by the gradient (score) of the marginal distribution $\nabla_{x_t} \log p_t(x_t)$:

$$x_t = \bar{\alpha}(t)x_1 + \sigma(t)\xi, \quad v_{\text{target}}(x_t, t) = -\frac{1}{2}\beta(t)x_t - \beta(t)\nabla_{x_t} \log p_t(x_t). \quad (9)$$

This nonlinear velocity field leads to curved reverse trajectories characteristic of diffusion models. A small positive constant is typically added to denominators during training for numerical stability.

Trigonometric flow. Trigonometric (Nichol and Dhariwal, 2021) uses a circular interpolation on the unit half-circle in wavelet space:

$$x_t = \cos \left(\frac{\pi}{2} t \right) x_0 + \sin \left(\frac{\pi}{2} t \right) x_1, \quad (10)$$

with time derivative velocity

$$v_{\text{target}} = \frac{\pi}{2} \left[-\sin \left(\frac{\pi}{2} t \right) x_0 + \cos \left(\frac{\pi}{2} t \right) x_1 \right]. \quad (11)$$

This formulation maintains constant norm $\|x_t\|$ and has constant curvature $\frac{\pi^2}{4}$, introducing smooth curved trajectories with stable, non-straight, velocity fields.

Implementation-level details for each FM variant (including the exact code correspondence for x_t and v_{target}) are provided in Appendix A.1, while Appendix A.2 gives the full VP derivation linking the SDE to the deterministic probability-flow ODE and the exact conditional target velocity used in training.

3.3.2. Conditional U-net architecture

FlowLet employs a conditional 3D U-Net, v_θ , designed to predict the velocity field within the 8-channel wavelet domain, as illustrated in Fig. 1(a). The model input consists of interpolated wavelet coefficients $x_t \in \mathbb{R}^{8 \times \frac{D}{2} \times \frac{H}{2} \times \frac{W}{2}}$, a timestep t , and a conditioning vector c ; the output is the predicted velocity field v_{pred} in the same wavelet domain, which is integrated via Euler ODE and mapped back to the volume domain through the inverse discrete wavelet transform (IDWT). The U-Net

backbone follows encoder, bottleneck and decoder stages with skip connections, enabling hierarchical feature extraction for 3D data. The primary building blocks are residual blocks (*ResBlocks*) incorporating normalization layers and SiLU activations. Temporal conditioning is intrinsic to the flow matching formulation: the timestep t is embedded via sinusoidal positional encoding (Vaswani et al., 2017), followed by a multi-layer perceptron (MLP), yielding a time embedding e_{time} that conditions feature computation throughout the network.

3.3.3. Unified conditional embedding

Beyond temporal conditioning, FlowLet is explicitly designed to synthesize brain MRIs according to clinically relevant attributes, with age being the focus of this work. Scalar conditioning variables are first normalized to a consistent numerical range; in the case of age, values are scaled to $[0, 1]$ based on the minimum and maximum observed in the training data. Each normalized scalar is then projected into a high-dimensional latent space using a dedicated two-layer MLP with SiLU activations, transforming the 1-dimensional normalized input into a 512-dimensional vector. When multiple conditioning variables are present, their projected representations are combined via element-wise summation, yielding a unified conditioning embedding e_{cond} . This unified representation is injected into the U-Net through two complementary mechanisms: depth-wise feature modulation and spatially adaptive attention.

3.3.4. Depth-wise conditioning via FiLM

To ensure the conditioning information influences feature computation at every level of the network, we employ Feature-wise Linear Modulation (FiLM) (Perez et al., 2018) within every residual block of the U-Net. FiLM applies an affine transformation to intermediate feature maps, allowing the network to dynamically adjust the scale and bias of activations on a per-instance basis. In our model, both the flow matching timestep embedding (e_{time}) and the condition embedding (e_{cond}) generate independent scale (γ) and bias (b) parameters. Specifically, given intermediate features h , conditioning is applied sequentially as:

$$h' = \text{Norm}(h) \cdot (1 + \gamma_{\text{time}}) + b_{\text{time}}, \quad (12)$$

$$h_{\text{mod}} = h' \cdot (1 + \gamma_{\text{cond}}) + b_{\text{cond}}. \quad (13)$$

The first transformation (12) adapts features according to the current point along the flow trajectory, while the second (13) refines them based on the conditioning variable. By integrating FiLM into every residual block, the conditional signal is made pervasively available, influencing both low-level texture formation and high-level anatomical structure across encoder and decoder paths.

3.3.5. Spatially adaptive conditioning via cross-attention

While FiLM enables robust global modulation, it applies identical transformations across all spatial locations within a feature channel. To capture localized, region-specific anatomical changes associated with aging, FlowLet further incorporates spatial conditioning modules inspired by transformer architectures (Rombach et al., 2022). These modules are inserted at lower spatial resolutions of the U-Net, where feature maps encode abstract, semantically meaningful representations with large receptive fields. Each spatial conditioning block first applies self-attention to model long-range spatial dependencies within the feature map. This is followed by a cross-attention operation, in which spatial features act as queries and the unified conditioning embedding provides both keys and values:

$$h_{\text{attn}} = \text{Attention}(Q, K, V), \quad (14)$$

$$Q = W_Q h, \quad K = W_K e_{\text{cond}}, \quad V = W_V e_{\text{cond}}. \quad (15)$$

This mechanism enables the network to selectively modulate specific anatomical regions as a function of age, refining spatially localized

morphological patterns such as ventricular enlargement or cortical thinning.

While FiLM and cross-attention are individually well-established conditioning techniques, their roles in this framework are complementary and specifically motivated by the wavelet-domain formulation. Because the network input consists of eight subbands encoding distinct frequency components, FiLM provides channel-wise modulation that can differentially scale and shift each subband as a function of the conditioning variable, effectively controlling the relative emphasis on low-versus high-frequency content at each layer. Cross-attention complements this by introducing spatially selective conditioning: age-related morphological changes (e.g., ventricular enlargement, cortical thinning) are localized phenomena, and cross-attention allows the network to attend to different spatial regions depending on the conditioning input. The combination thus provides both frequency-selective and spatially-selective control, which we argue is particularly relevant when operating on a multi-band wavelet representation rather than on single-channel image data.

3.3.6. Inference

Sampling, illustrated in Fig. 1(b), is performed by solving an ordinary differential equation in the wavelet domain, starting from Gaussian noise $x_0 \sim \mathcal{N}(0, I)$. Wavelet coefficients are iteratively updated using the learned velocity field v_{pred} . After integration, the final MRI volume is reconstructed by applying the IDWT. Additional training and inference benchmarks are provided in Appendix A.4.

4. Datasets and preprocessing

This section describes the multi-site datasets and the unified preprocessing pipeline used to build a full range age-spanning cohort. This setup aims to provide consistent training inputs and to enable fair evaluation of both synthesis quality and downstream BAP performance.

The training cohort was constructed by integrating T1-weighted MRI scans from three publicly available datasets: OpenBHB,¹ ADNI,² and OASIS-3.³ Only cognitively normal subjects were retained. This integration was motivated by the pronounced age imbalance of OpenBHB, which predominantly represents younger populations, and by the need to ensure adequate coverage of older age ranges for age-conditioned synthesis and downstream BAP evaluation.

An additional external dataset (DLBS) is used for evaluation purposes and is described in Appendix A.12.

4.1. OpenBHB

OpenBHB aggregates T1-weighted MRIs from 10 publicly available datasets (e.g., IXI, ABIDE I/II, GSP, CoRR), spanning 62 imaging sites across North America, Europe, and Asia (Dufumier et al., 2022). The mean age of the subjects is 24.92 ± 14.29 years, with a strong concentration around early adulthood. OpenBHB provides a predefined training-validation split obtained via stratified sampling based on age, sex, and acquisition site; in this work, the predefined validation subset was reserved exclusively for testing downstream tasks.

4.2. ADNI

From the Alzheimer's Disease Neuroimaging Initiative (ADNI-1, ADNI-2, and ADNI-3), 769 T1-weighted scans were selected from cognitively normal individuals aged 60–91 years (mean: 76.97 ± 4.99). To avoid repeated measures, only one scan per subject was retained, and all individuals with any cognitive impairment were excluded.

¹ <https://baobablab.github.io/bhb/dataset>.

² <https://adni.loni.usc.edu/>.

³ <https://sites.wustl.edu/oasisbrains/>.




Fig. 2. Age distribution across datasets. Overlaid histograms show the number of samples per age for OpenBHB, OASIS-3, and ADNI. OpenBHB is concentrated in younger adults (strong peak around the early 20 s), whereas OASIS-3 and ADNI primarily cover older adults (50–90 years), with highest density around the 70–80 age range.

4.3. OASIS-3

The Open Access Series of Imaging Studies (OASIS-3) is a longitudinal neuroimaging dataset spanning the adult lifespan. From an initial pool of 1314 scans, 1041 T1-weighted MRIs of cognitively normal subjects were retained after filtering, covering ages 42–95 years (mean: 71.10 ± 8.93).

4.4. Age distribution and data splits

Fig. 2 illustrates the combined age distribution of the integrated training cohort. OpenBHB contributes the majority of samples, producing a pronounced peak in younger age groups (approximately 10–30 years). The inclusion of ADNI and OASIS-3 substantially enriches representation in older age ranges (60–95 years), resulting in a more balanced coverage across adulthood. This broader distribution is essential for learning age-dependent anatomical variability and for evaluating age-conditioned synthesis. For generative evaluation, test samples were drawn from a 20% split of the integrated training distribution, preserving the same age profile. In contrast, BAP evaluation was performed on the independent OpenBHB validation set, which follows the original OpenBHB age distribution and provides an unbiased benchmark for downstream prediction performance.

4.5. Preprocessing pipeline

All MRI volumes were processed using a standardized preprocessing pipeline applied uniformly across datasets. In particular, skull stripping plays a critical role in preventing the model from exploiting non-brain structures (e.g., skull, scalp, or neck regions) that may correlate with age but are unrelated to brain morphology. Such spurious correlations can induce a *Clever Hans* effect (Wallis and Buvat, 2022), where predictions rely on superficial cues rather than meaningful neuroanatomical patterns. By isolating brain tissue, the model is forced to focus on age-relevant neuroanatomical patterns only. The following preprocessing steps were applied uniformly using tools from ANTs (Tustison et al., 2021) and FSL:

- **Bias field correction:** N4ITK correction (Tustison et al., 2010) was applied using ANTs to remove smooth, low-frequency intensity variations that commonly affect MRI scans;
- **Spatial normalization:** Affine registration to the MNI152 template (Fonov et al., 2009) using FSL FLIRT (Jenkinson et al., 2002);
- **Skull stripping:** Brain extraction was conducted on registered images using FSL BET (Smith, 2002);

- **Resampling:** Volumes were resampled to a common isotropic resolution of $91 \times 109 \times 91$, selected for compatibility with the downstream BAP model;
- **Intensity normalization:** Final volumes were normalized using z-score normalization (zero mean, unit variance).

Scans that failed automated preprocessing were excluded from the final dataset. No additional manual curation was performed. A complete list of retained samples and associated metadata is provided in the codebase, along with the full preprocessing implementation.

5. Experimental setup

This section describes the experimental protocol, baseline models, and training configuration. Evaluation metrics and downstream validation tasks are described in the following sections.

5.1. Baselines

We compare FlowLet against the following seven state-of-the-art generative models for 3D brain MRI synthesis:

- **WDM:** unconditional Wavelet Diffusion Model operating in the 3D Haar wavelet domain (Friedrich et al., 2024);
- **MD:** Medical Diffusion, a DDPM trained on a VQ-GAN backbone (Khader et al., 2022);
- **MLDM:** MONAI Latent Diffusion Model (Pinaya et al., 2022);
- **BS:** BrainSynth, a VQ-VAE and Transformer-based model for brain MRI synthesis (Tudosiu et al., 2024);
- **MOTFM:** a recent Flow Matching-based framework (Yazdani et al., 2025).

Several of these baselines are unconditional in their original public implementations. To enable a fair comparison on the age-conditioned synthesis and downstream BAP task, we introduce age-conditioned variants of WDM and MOTFM, denoted as **WDMa** and **MOTFMa**, respectively. Conditioning is implemented using the same strategy adopted in FlowLet: the scalar age value is normalized, projected into a 512-dimensional embedding, and injected into the U-Net backbone via cross-attention. This ensures that all conditional models receive age information in a consistent manner, without introducing architectural advantages specific to FlowLet.

We exclude subject-specific brain aging models (Litrico et al., 2024; Pombo et al., 2023; Yeganeh et al., 2025), which aim to transform an existing MRI of a given subject to a different age. These approaches address a fundamentally different task, as they rely on identity preservation and do not perform synthesis from noise, making them not directly comparable to our setting.

5.2. Training protocol and implementation details

FlowLet is implemented in PyTorch, using the AdamW optimizer with cosine annealing and mixed precision. For more memory-efficient attention, `xformers`⁴ is optionally available. All experiments were run on an NVIDIA A6000 GPU (48 GB VRAM); notably, while baseline models required this large memory capacity, FlowLet can be trained in 24 GB, enhancing accessibility. To compare different flow formulations (RFM, CFM, VP, and Trigonometric), all variants were implemented and trained under identical architecture, hyperparameters, and optimization settings. All external baselines were likewise trained from scratch on the same dataset, using the provided default hyperparameters except for input channels and padding to match our volumes. Although faster diffusion samplers such as DDIM exist (J. Song et al., 2021), they are not part of the original released implementations of

⁴ <https://github.com/facebookresearch/xformers>.

the considered baselines. In this work, we retain the standard sampling procedures provided by each method to ensure a consistent and reproducible comparison across models. We note that FlowLet does not rely on accelerated sampling heuristics but directly benefits from the deterministic ODE formulation, which enables efficient generation without requiring additional sampling approximations. Complete training details and hyperparameters are provided in [Appendix A.3](#). Detailed efficiency benchmarks (VRAM, training time, inference throughput, and scaling with resolution) are reported in [Appendix A.4](#).

5.3. Ablation study design

To assess the contribution of individual design choices within FlowLet, we conduct a series of controlled ablation experiments. Specifically, we analyze the impact of: (i) the wavelet basis used for the invertible representation, (ii) the number of inference steps and the associated sampling efficiency, (iii) the conditioning mechanisms, including FILM and spatial conditioning via cross-attention, and (iv) the numerical solver employed for ODE integration.

In each ablation, a single component is modified or removed while keeping the remaining architecture, training procedure, and evaluation protocol unchanged. All ablation variants are trained and evaluated under the same experimental setup as the full model to ensure a fair and isolated assessment of each design choice. Quantitative and qualitative results of these ablations are reported in [Section 7.3](#).

6. Evaluation metrics

This section describes the quantitative metrics used to evaluate generative performance in terms of image fidelity, sample diversity, and anatomical plausibility, as well as their relevance to downstream clinical tasks.

6.1. Image fidelity and diversity

Image fidelity and distributional alignment between synthetic and real data were evaluated using the Fréchet Inception Distance (FID) ([Heusel et al., 2017](#)) and the Gaussian kernel-based Maximum Mean Discrepancy (MMD) ([Gretton et al., 2012](#)). Both metrics were computed on feature representations extracted from a ResNet-50 pretrained on medical images ([Chen et al., 2019](#)), following established evaluation practices for 3D medical image synthesis ([Friedrich et al., 2024](#)). Lower values of FID and MMD indicate closer alignment between the distributions of synthetic and real samples.

Sample diversity was assessed using an *intra-set* Multi-Scale Structural Similarity Index (MS-SSIM) ([Wang et al., 2003](#)), computed as the average pairwise similarity among generated samples. In this setting, higher intra-set MS-SSIM values indicate reduced inter-sample variability and may suggest mode collapse, whereas lower values correspond to increased diversity. Statistical significance of pairwise comparisons was assessed using two-sided Wilcoxon rank-sum tests with Bonferroni correction ($\alpha = 0.05$).

Although MS-SSIM is traditionally used to compare a generated image with a reference target, it has been widely adopted in 3D brain MRI synthesis as an intra-set metric to quantify structural similarity among generated samples ([Pinaya et al., 2022](#); [Friedrich et al., 2024](#)). When used in this manner, MS-SSIM should be interpreted as a relative measure under consistent evaluation conditions rather than against an absolute threshold. A well-performing generative model is therefore expected to preserve anatomically coherent structures while maintaining realistic inter-subject variability.

Because intra-set MS-SSIM values are influenced by the conditioning range used during generation, diversity metrics depend on the age interval over which samples are synthesized. To account for this effect, we perform an additional age-stratified evaluation in which fidelity and diversity metrics are computed independently within restricted age ranges (non-overlapping age bins 15–30, 40–55, and 65–80 years). The detailed age-stratified protocol and corresponding quantitative results are reported in [Appendix A.6](#).

6.2. Region-based anatomical plausibility

Global similarity metrics are effective at assessing overall image quality and distributional alignment, but they may overlook fine-grained anatomical inconsistencies that are critical in clinical neuroimaging ([Wu et al., 2025](#)). In volumetric MRI in particular, metrics such as FID and MS-SSIM can be dominated by non-informative voxels, potentially producing favorable scores even when clinically relevant brain structures differ ([Jafarsteh et al., 2025](#)). To complement the global evaluation, we therefore perform a region-based anatomical analysis that explicitly focuses on local realism and morphological consistency across anatomically defined brain regions. This evaluation leverages FastSurfer,⁵ a deep-learning pipeline for automated segmentation and parcellation of structural MRIs. Each volume is segmented into 95 cortical and subcortical regions of interest (ROIs). Let $\mathcal{R}$ denote the set of ROIs. For each region $r \in \mathcal{R}$, let $V_r^{(R)}$ and $V_r^{(S)}$ denote the sets of voxels belonging to region r in the segmentations of the real reference volume R and the synthetic volume S , respectively. Region-based metrics are computed independently for each ROI and subsequently averaged across $\mathcal{R}$, yielding summary scores that reflect anatomical plausibility across the full brain rather than a limited subset of structures. A key aspect of this evaluation is the pairing strategy between synthetic and real volumes. Each synthetic brain volume is compared against a *different age-matched* real subject, and metrics are computed independently for each ROI prior to averaging across $\mathcal{R}$. This one-to-one, age-matched evaluation prevents a mode-collapsed model from appearing artificially strong: a single repeated anatomy cannot simultaneously match the diverse set of real age-matched references. Consequently, favorable scores reflect genuine anatomical plausibility and alignment with natural inter-subject variability rather than artifacts stemming from reduced diversity.

Models and sample generation. The region-based evaluation includes all baseline models (WDM, WDMa, MD, MLDM, MOTFM, MOTFma, BS) as well as all FlowLet variants (RFM, CFM, VP, and Trigonometric). Unless otherwise specified, FlowLet variants are configured for 10-step generation. For each model, 500 synthetic brain volumes are generated with age conditions linearly spanning the full age range of the training set. To ensure comparability across FlowLet variants despite differences in the underlying flow formulation, the same random seed is used for all generations. All synthetic volumes are segmented into 95 anatomical classes using FastSurfer. A reference set of 500 real brain MRIs, spanning the same age range, is processed through the identical segmentation pipeline and used as anatomical ground truth.

Region-wise metric definitions. Due to natural anatomical variability and potential segmentation mismatches, ROI segmentations extracted from synthetic and real volumes may not perfectly coincide, resulting in different voxel supports for a given region. Using only one of the two supports could therefore bias region-wise comparisons. To ensure a robust and unbiased evaluation, we define the comparison domain for each region as the union of voxel sets derived from both the real and synthetic segmentations,

$$U_r = V_r^{(R)} \cup V_r^{(S)}.$$

All region-wise metrics are computed over this unified voxel set and subsequently averaged across all regions in $\mathcal{R}$, providing a stable summary of anatomical plausibility relative to the real reference.

We define three complementary summary metrics:

$$\text{iMAE} = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \left(\frac{1}{|U_r|} \sum_{v \in U_r} |R_v - S_v| \right), \quad (16)$$

$$\text{KLD} = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} D_{\text{KL}}(P_r^{(R)} \parallel P_r^{(S)}), \quad (17)$$

⁵ <https://deep-mi.org/research/fast-surfer/> (v2.4.2).




Fig. 3. Visual assessment of image fidelity and realism for different 3D brain MRI synthesis models. Each column displays standard Axial, Coronal, and Sagittal views for the real reference scan (Subject of 72 years old) and specified generative method (Ours is RFM 10 steps).

$$\text{DICE} = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \frac{2|V_r^{(R)} \cap V_r^{(S)}|}{|V_r^{(R)}| + |V_r^{(S)}|}. \quad (18)$$

The overall intensity Mean Absolute Error (iMAE) quantifies local intensity realism by measuring the average absolute voxel-wise difference between R and S samples within each ROI. The overall Kullback–Leibler divergence (KLD) evaluates the similarity of regional intensity distributions by comparing normalized histograms of real ($P_r^{(R)}$) and synthetic ($P_r^{(S)}$) intensities computed over U_r . Finally, the overall Dice Similarity Coefficient (DICE) assesses morphological consistency through the spatial overlap between real and synthetic ROI segmentations. Dice values range from 0 (no overlap) to 1 (perfect overlap), while lower values of iMAE and KLD indicate improved intensity and distributional alignment.

Interpretation alongside global metrics. Region-based anatomical fidelity metrics are intended to be interpreted in conjunction with global distributional and diversity measures, including FID, MMD, and intra-set MS-SSIM. A reliable generative model should achieve high anatomical accuracy (e.g., low iMAE and high DICE) while simultaneously preserving sufficient inter-sample variability (low MS-SSIM) and global distributional alignment (low FID and MMD). Importantly, a mode-collapsed model may result in a deceptively strong region-wise score for a single anatomy, but will be revealed by elevated MS-SSIM and degraded global metrics. Considering these measures jointly therefore provides a more comprehensive and reliable assessment of generative quality.

6.3. Downstream brain age prediction

The clinical usefulness and functional fidelity of the synthetic data were evaluated through their impact on BAP. This task is particularly relevant in our setting because older adults are underrepresented in the training distribution, yet are the population most susceptible to cognitive decline. Establishing a normative trajectory of healthy aging therefore requires that BAP be defined and evaluated exclusively on cognitively normal subjects. Following the protocol of De Bonis et al. (2024), we assess whether synthetic data generated by each method improves age regression performance when training a 3D BAP predictor and evaluating it on a real, held-out test set.

BAP model and training protocol. We employ a 3D DenseNet-121 architecture configured for regression, with a linear output layer predicting chronological age from structural T1-weighted MRI volumes, following (De Bonis et al., 2024). Input intensities are normalized to the $[0, 1]$ range using the 5th and 95th percentile values computed from the training set to avoid test-set leakage. Models are trained using stochastic gradient descent with a cosine annealing warm restarts scheduler.

For each generative method, the downstream BAP evaluation is performed according to the following protocol:

1. **Synthetic data generation:** 3000 synthetic brain MRI volumes are generated for each method;
2. **Age labeling:** since BAP is a supervised task, each synthetic volume is associated with an age label. For conditional generative models, age is used directly as the conditioning variable during generation, with samples spanning the full training age range (5.9–95.5 years). For unconditional generators, age labels are assigned by sampling from the empirical age distribution of the training set to ensure a fair comparison;
3. **BAP model training:** a separate instance of the BAP network is trained for each generative method using the same architecture and optimization protocol;
4. **Evaluation dataset:** all trained BAP predictors are evaluated on the independent OpenBHB validation set, restricted to cognitively normal subjects aged 44 years and older;
5. **Performance metric:** prediction performance is quantified using Mean Absolute Error (MAE) in years, where lower values indicate more accurate estimation of chronological age.

7. Results

We evaluate FlowLet against baselines, combining qualitative inspection with quantitative evidence: global metrics, BAP performance, and ROI-based anatomical plausibility.

7.1. Qualitative evaluation

Fig. 3 presents representative samples generated by each model, providing a first qualitative assessment of image-level fidelity and anatomical realism. The FlowLet-generated volumes exhibit anatomically coherent structures and consistent global organization across views. However, visual differences in fine-scale anatomical details (e.g., cortical folding and cerebellar structures) remain subtle when compared to strong baselines such as WDM and MOTFMa. In particular, some degree of blurring is still observable in high-frequency regions across all methods, including FlowLet.

In addition, to isolate the effect of age conditioning, we generate a fixed-seed trajectory in which only the conditioning variable is varied while keeping the initialization constant (Appendix A.11, Fig. 12). This controlled setting highlights coherent age-dependent morphological changes.

A more detailed analysis of the wavelet-domain behavior is provided in Appendix A.10 (Figs. 11 and 10).

The segmentation outputs provide an additional qualitative perspective on anatomical consistency across models. Fig. 4 reports representative segmentation outputs for synthetic samples generated at the same target age (72 years), compared against a real reference subject. For unconditional baselines (MOTFM, MD, and WDM), the synthetic sample is selected randomly, whereas FlowLet samples are




Fig. 4. Qualitative comparison of FastSurfer segmentations of synthetic data from different models. The leftmost column displays the segmentation from a real 72-year-old subject, while each subsequent column shows axial, coronal, and sagittal views relative to a synthetic sample of the same age generated by a specific model. FlowLet samples are generated at 10 steps.

generated with explicit age conditioning. All FlowLet variants produce broadly consistent and anatomically plausible segmentations, with the Trigonometric formulation exhibiting slightly less regular boundaries.

Visual differences across methods remain subtle in this setting, particularly for models that do not exhibit clear failure modes. As a result, qualitative inspection alone is not sufficient to establish meaningful differences in anatomical fidelity. For this reason, we complement this analysis with region-based quantitative metrics (Section 7.2.3), which provide a more sensitive and reliable evaluation across anatomical regions.

A step-wise qualitative comparison across flow variants for 1–200 ODE steps is provided in A.5 for completeness.

7.2. Quantitative evaluation

7.2.1. Image fidelity and diversity

We first assess the generative performance of FlowLet in terms of image fidelity and sample diversity, with particular attention to sampling efficiency. Table 1a reports global similarity metrics for FlowLet variants and competing generative models. All FlowLet variants achieve highly competitive Fréchet Inception Distance (FID) and Maximum Mean Discrepancy (MMD) scores using only 10 ODE sampling steps, outperforming diffusion-based baselines such as WDM, MD, and MLDM, which require substantially more sampling iterations. Among the evaluated variants, the Trigonometric formulation attains the lowest FID and MMD values, indicating strong alignment with the real data distribution in feature space. To contextualize these values, we also report real-to-real baselines computed between disjoint subsets of the held-out data, as well as on the independent test set (Appendix A.13), providing an empirical reference for the lower bound of each metric.

Notably, improvements in global fidelity are not obtained at the expense of sample diversity. All FlowLet variants maintain intra-set MS-SSIM values that are comparable to, or lower than, those of baseline models, suggesting that enhanced similarity does not arise from mode collapse or reduced inter-sample variability. This indicates that FlowLet preserves a realistic level of diversity while improving global distributional alignment.

Fig. 5 further illustrates the trade-off between sampling efficiency and image fidelity for FlowLet variants. FID consistently improves as the number of sampling steps increases, with performance saturating at approximately 10 steps for RFM, CFM, and VP. Beyond this point, additional steps do not materially improve these formulations, while the Trigonometric path becomes unstable at 200 steps and degrades in both quantitative and downstream behavior.

7.2.2. Brain age prediction

We next evaluate whether synthetic brain MRIs generated by FlowLet improve performance on the downstream BAP task, a clinically relevant benchmark that is particularly sensitive to age-related

Table 1

Overall mean values for synthetic sample quality. **Bold** and underlined indicate the best and second-best models per metric. Metrics are computed over 100 random bootstrap resamples of the full generated sets. Standard deviations ($\leq 10^{-3}$) are omitted for conciseness.

(a) Ours (10 steps) vs. baselines					
	Method	Steps	FID ↓	MMD ↓	MS-SSIM ↓
Ours	RFM	10	<u>0.2981</u>	<u>0.0119</u>	0.9508
	CFM	10	0.3098	0.0124	0.9707
	VP	10	0.3079	0.0123	0.9706
	Trigon.	10	0.2854	0.0114	0.9660
Baselines	WDM	1000	0.3073	0.0123	0.9456
	WDMa	1000	0.3167	0.0123	<u>0.9430</u>
	MD	1000	0.3843	0.0153	0.9595
	MLDM	1000	0.3590	0.0144	0.9538
	MOTFM	10	0.3696	0.0147	0.9676
	MOTFMa	10	0.3747	0.0145	0.9528
	BS	–	0.3454	0.0138	0.9346
(b) Ablations of ours (steps + conditioning)					
	Method	Steps	FID ↓	MMD ↓	MS-SSIM ↓
RFM		1	0.3334	0.0133	0.9886
		2	0.3232	0.0129	0.9838
		5	0.3130	0.0125	0.9746
		200	0.2978*	<u>0.0119*</u>	<u>0.9487</u>
RFM DB4		10	0.3141	0.0125	0.9663
CFM		1	0.3361	0.0134	0.9899
		2	0.3258	0.0130	0.9858
		5	0.3146	0.0126	0.9771
		200	0.3044	0.0122	0.9508*
VP		1	0.3341	0.0133	0.9898
		2	0.3234	0.0129	0.9858
		5	0.3132	0.0125	0.9771
		200	0.3004	0.0120	0.9513*
Trigon.		1	0.3292	0.0131	0.9211
		2	0.2974*	<u>0.0119*</u>	0.9521
		5	0.2859	0.0114	0.9680
		200	0.3527	0.0141	0.9557
Trigon. RK4		1	<u>0.2974</u>	<u>0.0119</u>	0.9485
		2	0.3112	0.0124	0.9576
		5	0.3862	0.0154	0.9604
		10	0.3820	0.0152	0.9579
		200	0.3773	0.0151	0.9518
FiLM		10	0.3252	0.0130	0.9861
Spatial		10	0.3234	0.0129	0.9846
Uncond.		10	0.3181	0.0127	0.9803

* Marks results not significantly different from FlowLet-RFM 10 steps, based on pairwise Wilcoxon rank-sum tests with Bonferroni correction ($\alpha = 0.05$).

anatomical variation. Table 2a reports BAP results for cognitively normal subjects aged 44 years and older, comparing models trained on real data alone with models augmented using synthetic samples from




Fig. 5. FID vs. steps for FlowLet variants. The shaded bands indicate standard deviation.

different generative methods. To further assess generalization beyond the training distribution, we additionally evaluate the augmentation pipeline on a fully independent external dataset (DLBS), with results reported in [Appendix A.14](#).

BAP models trained with FlowLet-generated data consistently achieve lower test MAE than those trained on real data alone and those augmented with synthetic samples from baseline generators. Among FlowLet variants, the RFM and CFM formulations yield the best test performance, outperforming both unconditional baselines (e.g., WDM and MD) and the strong conditional MLDM baseline. This indicates that explicit age-conditioned synthesis from noise is beneficial for improving downstream age regression in the present setting.

The comparison between MOTFM and its age-conditioned counterpart MOTFMa further highlights the importance of conditioning. While the original MOTFM framework performs poorly on BAP, applying the same conditioning strategy used in FlowLet leads to a substantial reduction in test error. This confirms that the observed performance gains are driven by effective age conditioning rather than by architectural differences alone.

Notably, BAP models trained with FlowLet-RFM and FlowLet-CFM synthetic data achieve lower test error than the model trained exclusively on real data, suggesting that FlowLet-generated samples provide complementary age-relevant information that mitigates dataset imbalance and improves generalization.

7.2.3. Region-based anatomical plausibility

We further assess whether improvements in global image fidelity correspond to anatomically meaningful realism by evaluating local structural consistency across anatomically defined brain regions. [Table 2b](#) reports region-based metrics computed on FastSurfer segmentations, including intensity-based errors (iMAE, KLD) and morphological overlap (DICE), averaged over 95 cortical and subcortical regions.

FlowLet variants, particularly RFM and VP, rank among the strongest methods on the region-based metrics, with lower iMAE and KLD values and competitive DICE scores relative to most evaluated baselines. In contrast, several baseline models that exhibit competitive global fidelity metrics show degraded regional performance. For example, while the MD baseline achieves relatively low global FID values, its substantially lower DICE score (0.294) reveals deficiencies in capturing accurate anatomical shapes at the regional level.

These region-level results also help contextualize the downstream BAP findings. Although the Trigonometric formulation achieves strong global similarity metrics, it underperforms relative to RFM and CFM on both region-based anatomical measures and brain age prediction. This discrepancy indicates that alignment in global feature space does not necessarily guarantee preservation of fine-grained anatomical structure that is critical for age-related tasks.

To complement the overall metrics, we report additional age-stratified results across three non-overlapping age bins in [A.6](#), coupled with the corresponding p-value results and the complete significance-testing protocol for pairwise comparisons in [A.7](#).

Table 2

Downstream evaluations: (a) Brain age prediction (BAP); (b) Region-based.

(a) BAP for the Age ≥ 44 years group. Lower Mean Absolute Error (MAE, in years) = better accuracy. Ours: 10-step samples.			
	Model	Train MAE $\downarrow$	Test MAE $\downarrow$
Ours	Real data	1.15 $\pm$ 1.02	4.91 $\pm$ 3.92
	RFM	1.46 $\pm$ 0.59	4.01 $\pm$ 3.38
	CFM	1.39 $\pm$ 0.59	4.06 $\pm$ 3.37
	VP	1.02 $\pm$ 0.49	4.68 $\pm$ 3.78
	Trigon.	1.09 $\pm$ 0.48	4.27 $\pm$ 3.33
Ablat.	FiLM	0.57 $\pm$ 0.51	6.40 $\pm$ 4.70
	Spatial	0.87 $\pm$ 0.54	5.05 $\pm$ 3.84
	Uncond.	0.67 $\pm$ 0.64	5.65 $\pm$ 3.74
Baselines	WDM	1.63 $\pm$ 1.36	6.36 $\pm$ 5.22
	WDMa	0.33 $\pm$ 0.42	4.93 $\pm$ 4.09
	MD	2.54 $\pm$ 2.78	7.62 $\pm$ 6.40
	MLDM	0.98 $\pm$ 0.47	5.30 $\pm$ 3.86
	MOTFM	2.10 $\pm$ 2.82	10.88 $\pm$ 9.58
	MOTFMa	0.85 $\pm$ 0.53	4.90 $\pm$ 3.57
	BS	0.90 $\pm$ 0.40	4.16 $\pm$ 3.38

(b) Segmentation quality: lower is better for iMAE/KLD, higher for DICE.				
	Model	iMAE $\downarrow$	KLD $\downarrow$	DICE $\uparrow$
Ours	RFM	37.68 $\pm$ 10.22	0.855 $\pm$ 0.599	0.420 $\pm$ 0.169
	CFM	42.93 $\pm$ 11.19	1.395 $\pm$ 1.058	0.424 $\pm$ 0.171
	VP	37.61 $\pm$ 10.20	0.865 $\pm$ 0.615	0.423 $\pm$ 0.171
	Trigon.	43.35 $\pm$ 12.06	1.614 $\pm$ 1.277	0.379 $\pm$ 0.172
	Uncond.	39.95 $\pm$ 10.14	1.188 $\pm$ 1.094	0.409 $\pm$ 0.159
Baselines	WDM	47.52 $\pm$ 9.45	1.088 $\pm$ 0.781	0.368 $\pm$ 0.160
	WDMa	56.43 $\pm$ 14.44	2.112 $\pm$ 1.090	0.383 $\pm$ 0.172
	MD	38.44 $\pm$ 10.44	0.863 $\pm$ 0.593	0.294 $\pm$ 0.156
	MLDM	46.93 $\pm$ 11.62	1.040 $\pm$ 0.645	0.331 $\pm$ 0.154
	MOTFM	41.67 $\pm$ 11.12	0.915 $\pm$ 0.620	0.409 $\pm$ 0.163
	MOTFMa	42.93 $\pm$ 11.18	1.394 $\pm$ 1.058	0.391 $\pm$ 0.162
	BS	43.90 $\pm$ 11.53	0.862 $\pm$ 0.589	0.356 $\pm$ 0.158

7.3. Ablation studies

Ablation studies were conducted to assess how sensitive FlowLet is to key design choices. Specifically, we analyze the impact of the wavelet basis, the number of inference steps, the conditioning mechanisms, and the numerical solver. Furthermore, these experiments help disentangle which components drive sample quality and downstream utility; full results (including age-stratified evaluations and statistical tests) are reported in [Appendix](#).

Wavelet selection. We analyze the impact of the wavelet basis used in the invertible representation on reconstruction fidelity and generative performance. Among the evaluated wavelet families (Haar, Daubechies-4, Symlet-4, Coiflet-2, and Biorthogonal 3.3), Haar consistently yields the lowest round-trip reconstruction error (mean MAE 6.08×10^{-8}) and the most stable generative behavior. Replacing Haar with smoother bases, such as Daubechies-4, degrades global image quality and downstream performance, as reflected by higher FID and reduced BAP accuracy ([Tables 1b, 3, and 4](#)). Based on these results, Haar provides the most reliable trade-off between reconstruction fidelity, stability, and computational efficiency, and is adopted as the default wavelet basis. Detailed reconstruction analyses and extended quantitative comparisons are reported in [Appendix A.8](#).

Step count and sampling efficiency. We evaluate the impact of the number of ODE sampling steps on generative quality and computational cost. Across the FlowLet variants, most gains are achieved within the first 10 steps. For RFM, CFM, and VP, image fidelity improves rapidly at low step counts and then largely saturates, indicating that longer trajectories provide limited practical benefit in these formulations ([Fig.](#)

5 and Table 1b). A different behavior is observed for the Trigonometric formulation. While it attains favorable global metrics at low step counts, its performance does not continue to improve at high step counts and instead degrades at 200 steps. This deterioration is most clearly visible in the quantitative metrics and is also reflected in downstream and region-based evaluation. Thus, the step-count study indicates that the high-step regime is not uniformly beneficial across flow formulations and that the Trigonometric path exhibits a distinct instability at large integration steps. Consistent trends are observed in downstream evaluation. BAP performance improves rapidly from 1 to 10 steps and then saturates, with no meaningful benefit at higher step counts for the stable formulations (Table 1b). Notably, even low-step configurations outperform the real-only baseline, indicating that FlowLet captures age-relevant anatomical information without requiring long inference trajectories. These results identify 10 steps as an effective trade-off between sampling efficiency and generative quality for FlowLet. Detailed timing analyses, extended step-wise evaluations, and solver-dependent behaviors are reported in Appendices A.4 and A.5.

Effect of conditioning mechanisms. We analyze the contribution of the conditioning strategy used to inject age information into FlowLet. Specifically, we compare the full model, which combines FiLM and spatial conditioning via cross-attention, against ablated variants using only FiLM, only spatial conditioning, or no conditioning.

Removing either conditioning component degrades performance across multiple evaluation axes. While globally conditioned variants retain competitive FID and MMD scores, downstream BAP performance and region-based anatomical metrics are substantially affected (Tables 1b and 2a). In particular, models using only spatial conditioning or only FiLM exhibit higher BAP error and reduced anatomical consistency compared to the full model, whereas the unconditional variant performs worst overall.

Solver analysis. We further investigate whether the observed high-step degradation can be attributed primarily to the numerical solver used for ODE integration. In addition to the default Euler method, we therefore evaluate a fourth-order Runge–Kutta (RK4) solver for the Trigonometric flow formulation. Using RK4 does not resolve the Trigonometric failure mode. Instead, the degradation persists across evaluation axes: the overall FID worsens from 0.3112 at 2 steps to 0.3774 at 200 steps, the OpenBHB validation test MAE increases from 4.10 ± 3.17 to 4.43 ± 3.83 , and the region-based metrics deteriorate from 37.21 ± 9.96 to 44.28 ± 11.05 in iMAE and from 0.915 ± 0.668 to 1.664 ± 1.027 in KLD. These results show that increasing solver accuracy alone does not restore stable behavior for the curved Trigonometric path. Taken together, the solver comparison indicates that the limitations observed for the Trigonometric formulation are linked more strongly to the geometry of the learned transport path than to solver discretization alone (Tables 1b, 3, and 4).

Flow matching variants. We compare different flow matching formulations within the same architectural and training setup to assess the effect of interpolation geometry on generative performance. Specifically, we evaluate RFM, CFM, VP diffusion matching, and the Trigonometric formulation.

Across evaluations, RFM and CFM consistently provide the most stable and reliable performance, achieving competitive global fidelity while maintaining strong downstream BAP accuracy and regional anatomical plausibility. In contrast, the Trigonometric formulation attains favorable global metrics at low step counts but exhibits degraded regional consistency and downstream performance as the number of steps increases. The VP formulation does not yield a clear advantage over RFM or CFM to offset its additional complexity (Tables 1, 2a, and 2b).

These results support RFM as the preferred formulation for FlowLet, offering a robust balance between generative quality, stability, and computational efficiency.

Table 3

Brain age prediction performance for the age ≥ 44 years group on the merged dataset. Lower values indicate better prediction accuracy. Results are reported as Mean Absolute Error (MAE, in years) and Standard deviation.

Model	Steps	Train MAE ↓	Test MAE ↓
Real training data		1.15 ± 1.02	4.91 ± 3.92
RFM	1	1.32 ± 0.91	4.81 ± 4.43
	2	1.03 ± 0.61	4.74 ± 3.85
	5	1.42 ± 0.46	4.23 ± 3.52
	10	1.46 ± 0.59	4.01 ± 3.38
	200	1.83 ± 0.40	4.80 ± 3.91
Trigon. RK4	1	1.09 ± 0.45	4.26 ± 3.40
	2	0.33 ± 0.40	4.10 ± 3.17
	5	0.93 ± 0.56	4.12 ± 3.73
	10	0.76 ± 0.41	4.27 ± 3.49
	200	0.63 ± 0.42	4.43 ± 3.83
RFM db4	10	0.50 ± 0.35	4.61 ± 4.33

Table 4

Segmentation (ROI) quality metrics for the additional experiments. Lower values indicate better performance for iMAE and KLD; Higher values indicate better performance for DICE.

Model	Steps	iMAE ↓	KLD ↓	DICE ↑
Trigon. RK4	1	40.14 ± 10.53	1.186 ± 0.676	0.038 ± 0.099
	2	37.21 ± 9.96	0.915 ± 0.668	0.401 ± 0.182
	5	37.43 ± 10.27	0.868 ± 0.628	0.407 ± 0.174
	10	38.63 ± 10.50	0.969 ± 0.707	0.370 ± 0.166
	200	44.28 ± 11.05	1.664 ± 1.027	0.394 ± 0.172
RFM db4	10	40.48 ± 10.85	1.065 ± 0.751	0.427 ± 0.171

8. Discussion

8.1. Efficient 3D MRI synthesis without latent compression

FlowLet was designed to address a practical bottleneck in 3D neuroimaging generation: achieving high anatomical fidelity and controllable synthesis while keeping sampling costs tractable. Most state-of-the-art generative models for volumetric MRI rely on diffusion-based sampling (Ho et al., 2020; J. Song et al., 2021), which typically requires hundreds to thousands of iterations to reach high-quality outputs. To mitigate the associated computational burden, many approaches operate in learned latent spaces (Rombach et al., 2022; Pinaya et al., 2022), trading spatial resolution and anatomical detail for improved efficiency.

Our results show that this trade-off is not unavoidable. By performing Flow Matching directly in an invertible wavelet domain Lipman et al. (2023), FlowLet achieves competitive or superior global fidelity compared to latent diffusion and wavelet diffusion baselines (Friedrich et al., 2024), while requiring an order-of-magnitude fewer sampling steps (see Section 7.2.1). Importantly, this efficiency gain does not rely on learned compression: the wavelet representation preserves exact invertibility and fine-grained spatial detail, avoiding the reconstruction artifacts commonly observed in latent autoencoder-based pipelines (Müller-Franzes et al., 2023).

The ablation results further clarify the factors enabling this efficiency. Performance saturates at approximately ten ODE steps across all FlowLet variants, indicating that long stochastic trajectories are unnecessary when transport is learned in a multi-scale, invertible representation. In contrast to diffusion-based models, where increasing the number of steps is often essential to reduce noise-induced artifacts (J. Song et al., 2021), FlowLet benefits from smooth and stable trajectories already at low step counts.

Compared to wavelet diffusion models, which reduce memory usage but remain constrained by slow iterative sampling (Friedrich et al.,

2024), FlowLet demonstrates that wavelet representations can be effectively combined with deterministic flow-based generation.

8.2. Computational efficiency and accessibility

Beyond sampling efficiency, FlowLet substantially reduces the computational resources required for training and inference. As reported in Section 5.2 and Appendix A.4, operating in the wavelet domain results in an approximately 8× reduction in memory consumption compared to voxel-space diffusion models, enabling training with batch size 4 using approximately 22 GB of VRAM. In contrast, diffusion-based baselines such as WDM and MLDM require over 40 GB of VRAM under comparable settings. This reduced memory footprint allows FlowLet to be trained and deployed on widely available consumer-grade GPUs (e.g., RTX 3090/4090), lowering the barrier to entry for large-scale 3D MRI synthesis. Inference efficiency further reinforces this advantage: FlowLet generates a full-resolution 3D brain MRI in approximately 1.6 s using 10 ODE steps, representing a speedup of over one order of magnitude relative to the original released conditional diffusion baseline implementations used in this study.

These practical considerations are particularly relevant for research environments with limited computational resources and support the scalability of FlowLet to large neuroimaging cohorts.

8.3. Global metrics are not enough in volumetric neuroimaging

Global distributional metrics such as FID, MMD, and MS-SSIM are widely adopted to evaluate generative models, including in medical imaging applications (Pinaya et al., 2022; Friedrich et al., 2024). However, our results highlight that strong performance on these metrics does not necessarily translate into anatomically meaningful realism when synthesizing high-dimensional volumetric data such as 3D brain MRI (see Section 7.2.3).

Several baseline models achieve competitive global fidelity scores, yet exhibit degraded regional anatomical consistency and weaker downstream performance. This discrepancy is particularly evident in diffusion- and latent-based approaches, which may generate globally plausible intensity distributions while failing to preserve fine-grained structural details in anatomically relevant regions (Ho et al., 2020; Rombach et al., 2022).

In volumetric MRI, the predominance of background voxels further amplifies this effect, allowing global metrics to be dominated by non-informative regions and potentially masking localized anatomical errors (Friedrich et al., 2024). Our region-based evaluation explicitly addresses this limitation by quantifying anatomical fidelity across 95 cortical and subcortical regions.

This analysis reported in Section 7.2.3 reveals cases in which models with favorable global similarity scores nonetheless show reduced Dice overlap and increased regional intensity discrepancies. Importantly, these anatomical deficiencies are reflected in downstream brain age prediction performance, indicating that global similarity alone is insufficient to assess functional utility in clinically relevant tasks.

Overall, no single metric fully captures the behavior of the proposed model. While improvements are not uniformly large across all evaluation criteria, FlowLet consistently achieves competitive or improved performance across global, regional, and downstream evaluations. This pattern suggests that the contribution of the model lies in the combination of efficiency, controllability, and anatomically meaningful generation, rather than in a single dominant metric. Region-based and downstream evaluations provide essential insight into model behavior that would otherwise remain hidden when relying exclusively on global similarity measures, and should be considered standard components of future assessments in 3D medical image synthesis.

8.4. The role of explicit age conditioning

A central finding of this work is that explicit age conditioning is important for generating synthetic brain MRIs that remain coherent across age trajectories and useful for downstream age-related tasks. While several generative models can approximate the overall distribution of brain MRI intensities, our results show that unconditional synthesis is insufficient when the target application requires control over biologically meaningful attributes such as age.

In the absence of explicit conditioning, generative models tend to learn an average representation of the population, producing samples that appear globally realistic but lack consistent age-specific morphological patterns. This behavior is reflected in our experiments by the poor brain age prediction performance of unconditional models, including diffusion-based and flow-based baselines, despite their competitive global fidelity scores (as shown in Section 7.2.2).

This pattern is also reflected in the region-based evaluation, where the unconditional FlowLet variant remains broadly plausible overall but yields weaker anatomical scores than the age-conditioned variants.

The ablation analysis reported in Section 7.3, further highlights the importance of how conditioning information is injected into the model. Using either feature-wise modulation or spatial conditioning alone leads to degraded downstream and region-based performance, even when global metrics remain largely unaffected. Only the combination of feature-wise conditioning via FiLM and spatially adaptive conditioning via cross-attention consistently preserves age-relevant anatomical variations across the brain. This suggests that effective conditioning in volumetric neuroimaging requires both global control of feature statistics and localized modulation of spatial structures.

These observations are reinforced by the comparison between MOTFM and its age-conditioned variant MOTFMa. While the original MOTFM framework performs poorly on brain age prediction, applying the same age-conditioning strategy used in FlowLet yields a substantial improvement in downstream performance. This indicates that the observed gains are driven primarily by effective conditioning rather than by architectural differences alone.

Overall, our findings suggest that age-conditioned synthesis from noise remains underexplored in 3D neuroimaging and is particularly valuable for applications that require controlled generation along biologically interpretable dimensions. Explicit conditioning mechanisms therefore represent an important design direction for clinically meaningful MRI synthesis.

8.5. Flow geometry matters more than solver accuracy

Our ablation results indicate that the geometry of the learned transport path plays a more critical role in generative performance than the numerical accuracy of the ODE solver. While higher-order solvers are often assumed to improve sample quality in continuous-time generative models, our experiments show that increasing solver accuracy alone does not compensate for unfavorable flow geometry in high-dimensional anatomical settings.

In particular, replacing Euler integration with a fourth-order Runge-Kutta (RK4) solver for the Trigonometric formulation does not lead to improvements in either global fidelity or downstream performance. Instead, models based on curved interpolation paths exhibit increased instability and degraded anatomical consistency as the number of steps increases. This behavior suggests that the limitations observed for these formulations are primarily attributable to the structure of the velocity field rather than to discretization error.

In contrast, Rectified Flow Matching and Conditional Flow Matching, which rely on straight or near-straight interpolation paths, consistently provide stable and reliable performance across global, regional, and downstream evaluations (see Section 7.3). These findings are consistent with recent theoretical and empirical observations in the flow matching literature, which highlight the benefits of low-curvature

transport paths in terms of training stability and expressiveness (Lipman et al., 2023; Lee et al., 2023).

Taken together, these results suggest that, in the context of volumetric medical image synthesis, the choice of flow formulation and interpolation geometry has a greater impact on anatomical plausibility and downstream utility than the choice of numerical solver. Prioritizing simple and stable transport paths therefore appears more effective than increasing solver order when designing efficient generative models for high-dimensional anatomical data.

9. Limitations and future directions

Despite the encouraging results, several limitations should be acknowledged. First, while region-based anatomical metrics and downstream brain age prediction provide stronger proxies than global similarity scores, they do not replace expert-driven clinical assessment. Future work should therefore incorporate structured evaluations by neuroradiologists to further validate anatomical realism.

Second, although FlowLet supports explicit conditioning, this work focuses exclusively on age. Extending the framework to multi-attribute conditioning (e.g., sex, pathology, cognitive scores) introduces additional challenges related to disentanglement and robustness, which require systematic investigation. In addition, the current evaluation matches synthetic and real subjects based solely on age. Other covariates, such as sex, are not consistently available across all datasets and are therefore not included in the matching procedure. While the preprocessing pipeline mitigates part of the inter-subject variability, incorporating multi-attribute matching represents an important direction for future work.

Third, while the proposed preprocessing pipeline mitigates spurious correlations, comprehensive audits of fairness, robustness across demographic subgroups, and potential privacy leakage remain open challenges. Addressing these aspects will be essential for the deployment of generative models in sensitive clinical settings.

Finally, this study is limited to T1-weighted structural MRI. Future work will explore the application of FlowLet to other 3D imaging modalities and the integration of uncertainty estimation to better characterize the reliability of synthetic cohorts.

10. Conclusion

This work introduced FlowLet, a conditional generative framework for 3D brain MRI synthesis that combines flow matching with an invertible wavelet representation. By avoiding learned latent compression, FlowLet achieves efficient generation with few sampling steps while preserving anatomically meaningful structure. Extensive evaluation demonstrates that improvements in global fidelity translate into stronger regional anatomical plausibility and improved performance on a clinically relevant downstream task, brain age prediction.

Our results highlight the importance of combining efficient generative modeling with explicit conditioning and anatomy-aware evaluation when synthesizing volumetric medical images. While limitations remain, FlowLet provides a scalable, open-source, and controllable approach to 3D MRI synthesis and represents a step toward more practical and clinically meaningful generative models in neuroimaging.

CRedit authorship contribution statement

Danilo Danese: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Angela Lombardi:** Writing – review & editing, Writing – original draft, Validation, Supervision, Project administration, Methodology, Investigation, Formal analysis, Conceptualization. **Matteo Attimonelli:** Writing – review & editing, Writing – original draft, Methodology, Validation. **Giuseppe Fasano:** Writing – review & editing, Visualization, Methodology, Investigation, Data curation. **Tommaso Di Noia:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was partially supported by the following projects: “LIFE: the Italian system wide Frailty network” - local code: T2-AN-12 (CUP D93C22000640001); DEMETRA: “Development of an ensemble learning-based, multidimensional sensory impairment score to predict cognitive impairment in an elderly cohort of Southern Italy” (CUP D99J22001970006) Missione 6/componente 2/Investimento: 2.1 “Rafforzamento e potenziamento della ricerca biomedica del SSN”, funded by European Commission NextGenerationEU; We acknowledge the CINECA award under the ISCRA initiative (Project IsCc1 SynBrain), (Project IsCd1 FlowMRI), (Project IsCd3 EMBRAIN) for the availability of high performance computing resources and support”; This work has been carried out while Matteo Attimonelli was enrolled in the Italian National Doctorate on Artificial Intelligence run by Sapienza University of Rome in collaboration with Politecnico Di Bari. Data collection and sharing for the Alzheimer’s Disease Neuroimaging Initiative (ADNI) is funded by the National Institute on Aging (National Institutes of Health Grant U19 AG024904). The grantee organization is the Northern California Institute for Research and Education. In the past, ADNI has also received funding from the National Institute of Biomedical Imaging and Bioengineering, the Canadian Institutes of Health Research, and private sector contributions through the Foundation for the National Institutes of Health (FNIH), including generous contributions from the following: AbbVie, Alzheimer’s Association, Alzheimer’s Drug Discovery Foundation, Araclon Biotech, BioClinica Inc., Biogen, Bristol-Myers Squibb Company, CereSpir Inc., Cogstate, Eisai Inc., Elan Pharmaceuticals Inc., Eli Lilly and Company, EuroImmun, F. Hoffmann-La Roche Ltd and its affiliated company Genentech Inc., Fujirebio, GE Healthcare, IXICO Ltd., Janssen Alzheimer Immunotherapy Research & Development LLC., Johnson & Johnson Pharmaceutical Research & Development LLC., Lumosity, Lundbeck, Merck & Co. Inc., Meso Scale Diagnostics LLC., NeuroRx Research, Neurotrack Technologies, Novartis Pharmaceuticals Corporation, Pfizer Inc., Piramal Imaging, Servier, Takeda Pharmaceutical Company, and Transition Therapeutics.

Appendix

This section provides additional details and extended experiments supporting the main work, including Flow Matching implementations, hyperparameters and efficiency benchmarks, step-wise qualitative results, age-stratified metrics, wavelet analyses, and significance testing.

A.1. Flow matching implementations

This subsection clarifies the implementation of the flow matching formulations. In all cases, the training objective is to minimize the MSE loss between a predicted velocity field v_θ and a target velocity field v_{target} . For each training step, a time value t is sampled uniformly from $[0, 1]$. The data sample x_1 corresponds to the variable `x1_wavelet` in the code, and the noise sample $x_0 \sim \mathcal{N}(0, I)$ is represented by variables named `x0_wavelet`.

Rectified Flow Matching (RFM). The implementation directly translates the linear interpolation path and constant velocity from the paper. The path $x_t = (1 - t)x_0 + tx_1$ is computed as `xt = (1 - t_broadcast) * x0_wavelet + t_broadcast * x1_wavelet`. The target velocity $v_{\text{target}} = x_1 - x_0$ corresponds to the variable `v_target = x1_wavelet - x0_wavelet`.

Conditional Flow Matching (CFM). The state-dependent target velocity $v_{\text{target}} = \frac{x_1 - x_t}{1 - t + \epsilon}$ is implemented as $v_{\text{target}} = (x_1_{\text{wavelet}} - x_t) / (1 - t_{\text{broadcast}} + 1e-8)$. The term x_t is computed via the same linear interpolation as in RFM, and a small $\epsilon = 10^{-8}$ is used for numerical stability.

Trigonometric flow. The circular interpolation path $x_t = \cos(\frac{\pi}{2}t)x_0 + \sin(\frac{\pi}{2}t)x_1$ is implemented as $x_t = \text{torch.cos(angle)} * x_0_{\text{wavelet}} + \text{torch.sin(angle)} * x_1_{\text{wavelet}}$, where angle represents $\frac{\pi}{2}t$. The corresponding velocity field v_{target} is computed as its time derivative: $v_{\text{target}} = -\text{torch.sin(angle)} * (\pi/2) * x_0_{\text{wavelet}} + \text{torch.cos(angle)} * (\pi/2) * x_1_{\text{wavelet}}$.

A.2. Variance-Preserving (VP) diffusion matching

This section presents the explanation of the VP Diffusion Matching implementation, establishing the connection between the high-level SDE formulation from the main text and the exact conditional velocity field used for training, as implemented in the codebase. The derivation summarizes and follows rigorous treatments found in the literature, particularly (Y. Song et al., 2021; Lipman et al., 2023; Tong et al., 2023; Gagneux et al., 2025b).

Core definitions and time conventions. The VP formulation is defined over a continuous data-to-noise time interval $t \in [0, 1]$, where $t = 0$ corresponds to clean data and $t = 1$ to pure noise. The process is governed by a linear variance schedule $\beta(t)$:

$$\beta(t) = \beta_{\min} + t(\beta_{\max} - \beta_{\min}), \quad (19)$$

where our implementation uses the standard hyperparameters $\beta_{\min} = 0.1$ and $\beta_{\max} = 20.0$ (J. Song et al., 2021). From this schedule, we define the signal and noise scaling coefficients, $\bar{\alpha}(t)$ and $\sigma(t)$, as follows:

$$\bar{\alpha}(t) = \exp\left(-\frac{1}{2} \int_0^t \beta(s) ds\right), \quad \text{and} \quad \sigma(t) = \sqrt{1 - \bar{\alpha}(t)^2}. \quad (20)$$

A noised sample x_t is generated from a real sample x_1 via the interpolation $x_t = \bar{\alpha}(t)x_1 + \sigma(t)\xi$, where $\xi \sim \mathcal{N}(0, \mathbf{I})$.

Crucially, FlowLet training loop samples a time $t_{\text{flow}} \in [0, 1]$ along the generative noise-to-data path. This is mapped to the theoretical data-to-noise time via the relation $t = 1 - t_{\text{flow}}$. All formulas below use the theoretical time t .

From stochastic SDE to deterministic ODE. To make explicit the link between the VP diffusion model and the velocity field used in FlowLet, we first recall the full reverse-time SDE associated with the VP forward process (Y. Song et al., 2021):

$$dx_t = \left[-\frac{1}{2}\beta(t)x_t - \beta(t)\nabla_{x_t} \log p_t(x_t)\right] dt + \sqrt{\beta(t)} d\bar{W}_t. \quad (21)$$

For a generic diffusion SDE of the form $dx_t = f(x_t, t) dt + g(t) dW_t$, the corresponding Probability Flow ODE that evolves the same marginals p_t is (J. Song et al., 2021):

$$\frac{dx_t}{dt} = f(x_t, t) - \frac{1}{2} g(t)g(t)^\top \nabla_{x_t} \log p_t(x_t). \quad (22)$$

Applying this identity to the VP reverse SDE, with $g(t) = \sqrt{\beta(t)}\mathbf{I}$, yields the deterministic VP flow:

$$v_{\text{ODE}}(x_t, t) = -\frac{1}{2}\beta(t)x_t - \frac{1}{2}\beta(t)\nabla_{x_t} \log p_t(x_t). \quad (23)$$

which explains the factor $1/2$ multiplying the score term. This is the velocity field that FlowLet is trained to approximate.

Deriving the computable target velocity. Directly using Eq. (23) for training is intractable because the true score, $\nabla_{x_t} \log p_t(x_t)$, is unknown. In the conditional flow matching framework, we circumvent this by computing the velocity conditioned on the target data sample x_1 . This

Table 5

Training hyperparameters for FlowLet models.

Parameter	Value
Optimizer & Scheduler	
Optimizer	AdamW
Learning rate	3e-6
Weight decay	1e-5
Adam β_1, β_2	0.9, 0.999
LR scheduler	CosineAnnealingLR
Scheduler eta_min	1e-7
Training	
Epochs	200
Batch size	4
Gradient clipping norm	1.0
Model & Architecture	
U-Net base channels	128
U-Net channel multipliers	(1, 2, 4, 8)
U-Net attention resolutions	(4, 8)
U-Net attention heads	8
Condition embedding dim	512
xformers attention	Enabled
Mixed precision	Enabled

is achieved by substituting the intractable score with the analytical conditional score, $\nabla_{x_t} \log p_t(x_t | x_1)$, whose closed form is given by Tweedie's formula (Efron, 2011):

$$\nabla_{x_t} \log p_t(x_t | x_1) = -\frac{x_t - \bar{\alpha}(t)x_1}{\sigma(t)^2}. \quad (24)$$

Substituting this into our ODE velocity field (Eq. (23)) gives the final, computable target velocity that our network learns to predict. This connection is rigorously established in recent flow matching literature (Lipman et al., 2023; Tong et al., 2023; Gagneux et al., 2025b):

$$\begin{aligned} v_{\text{target}}(x_t, t | x_1) &= -\frac{1}{2}\beta(t)x_t - \frac{1}{2}\beta(t)\left(-\frac{x_t - \bar{\alpha}(t)x_1}{\sigma(t)^2}\right) \\ &= -\frac{1}{2}\beta(t)\left(x_t - \frac{x_t - \bar{\alpha}(t)x_1}{1 - \bar{\alpha}(t)^2}\right) \\ &= -\frac{1}{2}\beta(t)\frac{\bar{\alpha}(t)^2x_t - \bar{\alpha}(t)x_1}{1 - \bar{\alpha}(t)^2}. \end{aligned} \quad (25)$$

Equivalence to the code implementation. The target velocity in Eq. (25) is mathematically equivalent to our implementation, which adopts a numerically stable reformulation of the same expression. We recall that $\bar{\alpha}(t) = e^{-\frac{1}{2}T(t)}$ and $\bar{\alpha}(t)^2 = e^{-T(t)}$. Additionally, in the generative training loop, the time variable is reversed as $t_{\text{flow}} = 1 - t$. Finally, Eq. (25) takes the following form, as implemented in FlowLet.

$$v_{\text{target}}(x_t, t | x_1) = -\frac{1}{2}\beta(1-t)\frac{e^{-T(1-t)}x_t - e^{-\frac{1}{2}T(1-t)}x_1}{1 - e^{-T(1-t)}}. \quad (26)$$

All the implemented formulations are available in `/flowlet/models/flow_matching.py`.

A.3. Training hyperparameters

All models were trained using PyTorch⁶ on a system running Ubuntu 24.04, equipped with a single NVIDIA A6000 GPU (48 GB). Reproducibility was ensured across the entire codebase by enforcing deterministic behavior across PyTorch, NumPy, and CuDNN. The training procedure was standardized across all Flow Matching formulations and baselines to ensure fair comparison. The training hyperparameters for FlowLet and the BAP model are reported in Tables 5 and 6, respectively.

⁶ version 2.6.0.




Fig. 6. Qualitative comparison of FlowLet flow matching formulation across different ODE step counts. Each column shows axial, coronal, and sagittal views for a given method and step count. All samples share the same noise seed and age condition, yielding anatomically consistent outputs with similar structural compartments across flow formulations.

Table 6

Training hyperparameters for the BAP model.

Parameter	Value
<i>Optimizer & Scheduler</i>	
Optimizer	SGD
Initial learning rate	0.01
LR Scheduler	CosineAnnealingWarmRestarts
Scheduler T_0	17
Scheduler T_mult	2
Scheduler eta_min	1e-5
<i>Training</i>	
Epochs	100
Batch size	16
Loss function	L1 Loss (MAE)
<i>Model & Architecture</i>	
Model architecture	DenseNet-121
Input channels	1
Output classes	1 (Age)
<i>Data preprocessing</i>	
Normalization method	Min-Max scaling
Scaling min value	5th percentile (of training set)
Scaling max value	95th percentile (of training set)

Data preparation and loading. The dataset was constructed from a centralized metadata file containing all MRI NIfTI file paths and associated subject metadata (e.g., age). Only subjects labeled as “Cognitively Normal” (CN) were retained. From this filtered dataset ($N = 5794$), a deterministic 80%–20% split was applied using a fixed random seed, resulting in 4635 training and 1159 validation subjects, ensuring no overlap across tasks. This split was used consistently across all training runs. The minimum and maximum age values were computed over the full dataset and used to normalize the age condition to the $[0, 1]$ range across all samples.

The preprocessing pipeline for generative modeling was applied to each 3D MRI volume and consisted of the following steps:

1. Load the NIfTI volume and normalize intensities by clipping to the 0.5th and 99.5th percentiles, then scaling to the range $[-1, 1]$;
2. Pad the volume to the required input size of $112 \times 112 \times 112$ using replication padding;
3. For training samples, apply minimal, non-invasive data augmentation using the MONAI⁷ library, including random 3D rotations, intensity scaling, and Gaussian noise. These augmentations are designed to preserve anatomical structures while increasing data variance. No augmentations were applied to the validation set;
4. Apply a 3D Haar DWT, implemented using the PyWavelets⁸ library, to convert the single-channel input into an 8-channel tensor representing approximation and detail subbands. This tensor is directly used as input to the U-Net.

Training dynamics. The model was trained for 200 epochs using the AdamW optimizer with an initial learning rate of 3×10^{-6} , decayed via cosine annealing to a minimum of 1×10^{-7} . Automatic mixed precision (AMP) training was enabled, using a *GradScaler* to ensure numerical stability. Memory-efficient attention layers from the *xformers*⁹ library were used to reduce memory consumption. Gradients were clipped to a maximum L2 norm of 1.0. The model was trained using MSE loss between the predicted and target velocity fields in the wavelet domain.

Final model selection was based on the lowest validation MSE over the complete training epochs.

The full implementation is provided in the FlowLet_CODE folder of the codebase.

⁷ <https://monai.io/> (1.4.0).

⁸ <https://pywavelets.readthedocs.io/> (1.8.0).

⁹ <https://github.com/facebookresearch/xformers> (0.0.29.post3).

Table 7

Computational resources analysis. Training times are for 200 epochs unless specified otherwise in parentheses. Inference time is measured for generating 3000 samples.

Method	Stage/Component	Batch size	VRAM (GB)	Training time	Inference time
FlowLet (Ours)	Single stage	4	22	80 h	1 h 18 min
MLDM	Stage 1	2	39	120 h	10 h
	Stage 2	2	14	70 h (350 epochs)	
MOTFM	Single stage	1	25	315 h	1 h 40 min
MOTFMa	Single stage	1	25	315 h	1 h 40 min
MD	Stage 1	4	16	74 h	13 h 40 min
	Stage 2	2	39	200 h	
	Stage 3	4	10	1 h	
	Stage 4	16	11	16 h 40 min	
WDM	Single stage	1	40	96 h	58 h 20 min
WDMa	Single stage	1	43	101 h	58 h 20 min
BS	Stage 1	16	7	205 h (500 epochs)	4 h 18 min
	Stage 2 (Extraction)	–	4	1 h	
	Stage 3	2	5	27 h (500 epochs)	

Table 8

Ablation study on spatial resolution, memory consumption, and speed for FlowLet model. Peak VRAM usage is reported for Batch size 1 and 4. Inference time is measured per sample. The entry “–” indicates required VRAM exceeds 48 GB.

Resolution	Peak VRAM (Batch size 1)	Peak VRAM (Batch size 4)	Inference time
112 ³	18 GB	22 GB	1.57 s
128 ³	23 GB	31 GB	2.12 s
256 ³	42 GB	–	6.80 s

A.4. Efficiency benchmarks

To evaluate the computational efficiency of FlowLet, we report peak VRAM usage, training time, and inference throughput for all major baselines in Table 7. FlowLet requires substantially less memory during training (approximately 22 GB with a batch size of 4), compared to over 40 GB for diffusion-based baselines such as WDM and MLDM. This reduced memory footprint allows FlowLet to be trained and deployed on consumer GPUs (e.g., RTX 3090/4090). Inference is also considerably faster: FlowLet produces a full 3D sample in 1.57 s using 10 ODE steps, representing roughly a 45× speedup over the conditional diffusion baseline WDMa (about 70 s for 1000 steps).

We further assess how FlowLet scales with increasing spatial resolution, as summarized in Table 8. FlowLet scales efficiently, requiring 42 GB of VRAM to train at 256³ with batch size 1. This behavior is largely due to operating in the wavelet domain. Sampling times remain stable across resolutions, increasing from roughly 1.6 s per volume at 112³ to 6.8 s at 256³ (10 ODE steps).

A.5. Additional qualitative assessment of flow variants

To assess the generative behavior of different flow matching formulations, synthetic brain MRIs were generated using FlowLet with RFM, CFM, VP, and Trigonometric flows. For each method, volumes were sampled using 1, 2, 5, 10, and 200 ODE solver steps, with a fixed random seed and constant age condition (normalized age value $0.5 \in [0, 1]$, corresponding to 51 years). This controlled setup isolates the effect of the flow formulation and integration efficiency on the resulting images. Fig. 6 displays representative axial, coronal, and sagittal mid-slices for qualitative comparison. At high step counts (e.g., 200 steps), all methods converge toward anatomically plausible structures, confirming consistency in the asymptotic regime. In contrast, notable differences arise in low-step regimes: RFM, CFM, and VP produce stable and coherent volumes with as few as 2–5 steps, while the Trigonometric flow shows instability and structural artifacts. These results illustrate the trade-off between integration curvature and sampling stability, emphasizing the advantages of lower-curvature flows for efficient and anatomically faithful synthesis.

The sample generation process with FlowLet is implemented in `scripts/generate_linear.py`.

A.6. Additional age-stratified quantitative evaluation

Benchmark. Three non-overlapping age bins were considered: 15–30, 40–55, and 65–80 years. For each generative model and each age bin, 200 synthetic brain MRI volumes were generated. Age conditions were uniformly sampled within each bin. All metrics (FID, MMD, and intra-set MS-SSIM) were computed independently for each age group using identical feature extractors and evaluation settings as in the main experiments.

Table 11 (overall, 15–30, 40–55 and 65–80 age groups) reports the complete quantitative results. Cells marked with a dash (“–”) under the unconditional (*Uncond.*) baselines (RFM-Uncond., WDM, MOTFM and MD) indicate that age-stratified metrics could not be computed, as these models lack explicit age conditioning and therefore cannot be evaluated within the defined age range.

The performance trend observed in the overall evaluation is consistently reflected across all age groups. Among the evaluated age ranges, the 65–80 group achieves the best performance in terms of both FID and intra-set MS-SSIM, indicating that FlowLet produces high-fidelity and diverse samples precisely in the demographic segment where data augmentation is most needed. This group exhibits the lowest MS-SSIM values across all age bins, suggesting increased anatomical variability in later adulthood, potentially reflecting diverse neuroanatomical aging trajectories (Bethlehem et al., 2022). In contrast, higher MS-SSIM scores observed in younger age groups likely reflect more homogeneous structural patterns.

To complement the results presented in the main text, Table 10 provides the complete quantitative results for the new ablation experiments, including the performance of FlowLet-RFM trained with the Daubechies-4 (db4) basis and the Trigonometric flow variant integrated using the 4th-Order Runge–Kutta (RK4) solver. The performance degradation observed in the overall metrics for the RFM-db4 variant (FID 0.3142 vs. 0.2981 for Haar) is consistent across all age bins (Table 10), confirming that the superior reconstruction fidelity of the Haar basis translates into better generalized generative performance, even in age-stratified contexts. Similarly, the RK4 integration of the Trigonometric path, while intended to improve stability, failed to yield superior results compared to the simpler Euler-integrated RFM, often showing inferior FID/MMD across multiple step counts and age groups.

These findings underscore the importance of contextualizing intra-set MS-SSIM values with respect to the demographic characteristics of the generated data. All comparisons are reported within age-matched intervals to ensure fair and meaningful evaluation of generative diversity.




Fig. 7. FID vs. number of sampling steps for FlowLet variants calculated on overall (5.9–95.5 years) and by age range. We report FID as a function of the sampling budget for RFM, CFM, VP, and Trigonometric formulations. Shaded regions denote standard deviation across samples. All methods improve as steps increase, reaching marginal improvements after 10 steps. The Trigonometric variant consistently becomes unstable at 200 steps.




Fig. 8. The leftmost column shows a real MRI volume in three views (axial, coronal, sagittal). Subsequent columns display the absolute reconstruction error for the same sample after a round-trip transform using different wavelet bases. The reported MAE values were computed over the entire set of voxel intensities in the 3D volume for this single instance. The color scale is adjusted to emphasize the narrow dynamic range of errors and facilitate visual comparison.

FID vs. sampling steps. Fig. 7 plots the FID as a function of ODE steps for all FlowLet variants, stratified by age group. All curves exhibit a consistent monotonic improvement with increasing step counts, in agreement with trends observed in the aggregate metrics. A performance plateau is observed between 10 and 200 steps, supporting the selection of 10 steps as an optimal balance between sampling efficiency and image fidelity across age groups. Notably, the Trigonometric flow, despite achieving competitive FID scores at low step counts, shows greater variability and degraded performance in high-step regimes.

A.7. Significance testing and P-value analysis

To assess the reliability of the performance differences observed in our study, we computed statistical significance relative to the reference configuration (FlowLet-RFM at 10 steps). These p-values correspond to

Table 9

Mean Absolute Error (MAE) and Standard deviation (Std) computed over the entire training set ($N = 5794$) for each wavelet. The error reflects voxel-wise intensity differences between original and reconstructed volumes.

Wavelet type	MAE $\pm$ Std $\downarrow$
Haar	$6.08 \times 10^{-8} \pm 1.60 \times 10^{-8}$
Daubechies (db4)	$1.03 \times 10^{-7} \pm 2.30 \times 10^{-8}$
Symlet (sym4)	$1.11 \times 10^{-7} \pm 2.81 \times 10^{-8}$
Coiflet (coif2)	$9.98 \times 10^{-8} \pm 2.24 \times 10^{-8}$
Biorthogonal (bior3.3)	$1.32 \times 10^{-7} \pm 3.70 \times 10^{-8}$

the global quality metrics (FID, MMD, and MS-SSIM) reported in Tables 10 and 11 and are shown in Fig. 9. In the heatmap, cells colored in red indicate results that are non-statistically significant ($p > 0.05$). Notably, the prevalence of red cells in the 200-step rows for the RFM, CFM, and VP solvers confirms that extending the inference trajectory beyond 10 steps yields statistically negligible improvements in sample quality for these methods. Conversely, the remaining colored regions (green to yellow) represent statistically significant differences, validating the distinct performance gains achieved by the Trigonometric solver and the variations observed in the baseline comparisons.

A.8. Additional wavelet evaluation

Five wavelet families were compared, Haar, Daubechies-4 (db4), Symlet-4 (sym4), Coiflet-2 (coif2), and Biorthogonal 3.3 (bior3.3), to assess their suitability for 3D MRI volume reconstruction. The analysis focused on filter structure, computational efficiency, reconstruction accuracy, and the presence of voxel-domain artifacts. Detailed properties, including filter definitions and theoretical foundations, are available in Mallat (1989), Daubechies (1992). Table 9 reports quantitative reconstruction errors, and Fig. 8 visualizes error distributions for a representative sample.

Table 10

Comparison of synthetic sample quality metrics for RFM DB4 and Trigon. RK4 variants. The metrics are computed over 100 bootstrap resamples.

Method	Type	Steps	Overall			Age 15–30		
			FID ↓	MMD ↓	MS-SSIM ↓	FID ↓	MMD ↓	MS-SSIM ↓
Ours	RFM DB4	10	0.3142 ± 0.0018	0.0125 ± 0.0001	0.9663 ± 0.0153	0.3376 ± 0.0031	0.0136 ± 0.0001	0.9822 ± 0.0015
Ours	Trigon. RK4	1	0.2974 ± 0.0017	0.0119 ± 0.0001	0.9485 ± 0.0149	0.3202 ± 0.0029	0.0129 ± 0.0001	0.9632 ± 0.0008
		2	0.3112 ± 0.0022	0.0124 ± 0.0001	0.9576 ± 0.0132	0.3359 ± 0.0031	0.0135 ± 0.0001	0.9721 ± 0.0009
		5	0.3862 ± 0.0021	0.0154 ± 0.0001	0.9604 ± 0.0114	0.4145 ± 0.0031	0.0166 ± 0.0001	0.9743 ± 0.0016
		10	0.3820 ± 0.0022	0.0152 ± 0.0001	0.9579 ± 0.0129	0.4061 ± 0.0035	0.0163 ± 0.0001	0.9742 ± 0.0028
		200	0.3774 ± 0.0023	0.0151 ± 0.0001	0.9518 ± 0.0172	0.4061 ± 0.0040	0.0163 ± 0.0002	0.9728 ± 0.0042
Method	Type	Steps	Age 40–55			Age 65–80		
			FID ↓	MMD ↓	MS-SSIM ↓	FID ↓	MMD ↓	MS-SSIM ↓
Ours	RFM DB4	10	0.3338 ± 0.0043	0.0134 ± 0.0002	0.9804 ± 0.0027	0.2739 ± 0.0039	0.0109 ± 0.0002	0.9780 ± 0.0035
Ours	Trigon. RK4	1	0.3175 ± 0.0041	0.0127 ± 0.0002	0.9642 ± 0.0010	0.2563 ± 0.0037	0.0102 ± 0.0001	0.9611 ± 0.0018
		2	0.3307 ± 0.0040	0.0132 ± 0.0002	0.9716 ± 0.0012	0.2685 ± 0.0040	0.0107 ± 0.0002	0.9676 ± 0.0021
		5	0.4080 ± 0.0049	0.0163 ± 0.0002	0.9717 ± 0.0021	0.3405 ± 0.0048	0.0136 ± 0.0002	0.9665 ± 0.0029
		10	0.4003 ± 0.0057	0.0160 ± 0.0002	0.9697 ± 0.0041	0.3394 ± 0.0054	0.0136 ± 0.0002	0.9614 ± 0.0056
		200	0.3974 ± 0.0050	0.0159 ± 0.0002	0.9627 ± 0.0095	0.3310 ± 0.0045	0.0132 ± 0.0002	0.9483 ± 0.0129

Haar (db1). The Haar wavelet (Daubechies-1) is the simplest wavelet, defined by a step-function basis. It is discontinuous and uses a length-2 filter, providing the shortest support and highest computational efficiency. Haar achieved the lowest reconstruction error (mean MAE: 6.08×10^{-8}), with minimal boundary artifacts and numerically exact reconstruction. Although its constant basis functions can cause blockiness under compression, no such artifacts were observed under lossless reconstruction.

Daubechies (db4). Daubechies-4 is an orthonormal wavelet with 4 vanishing moments and an 8-tap filter. Its smoother basis functions improve energy compaction and reduce blockiness compared to Haar. However, its longer support can introduce ringing near sharp transitions. In the experiments, db4 showed reconstruction errors roughly an order of magnitude higher than Haar (MAE $\sim 10^{-7}$), likely due to boundary effects.

Symlets (sym4). Sym4 is a near-symmetric variant of db4, preserving orthonormality and using an 8-tap filter. Its symmetry helps reduce phase distortion and shift sensitivity. Like db4, its reconstruction errors remained in the 10^{-7} range and were primarily localized in background voids.

Coiflets (coif2). Coiflet-2 uses a 12-tap, near-symmetric filter with vanishing moments in both wavelet and scaling functions. It enables smooth intensity transitions at higher computational cost and boundary sensitivity. It achieved the second-lowest reconstruction error (MAE $\sim 10^{-8}$), though some oscillatory artifacts appeared near high-contrast edges.

Biorthogonal (bior3.3). Bior3.3 employs separate analysis and synthesis filters with three vanishing moments each and an 8-tap symmetric analysis filter. This biorthogonal design ensures linear phase and supports shift-invariant, edge-aligned reconstruction. The MAE was on par with other 8-tap wavelets ($\sim 10^{-7}$), with mild structured artifacts near edges, likely due to non-orthogonality.

Qualitative reconstruction analysis. Fig. 8 presents the absolute voxel-wise reconstruction error for a representative MRI volume. This qualitative view highlights spatial error patterns following a single round-trip transform. Haar shows minimal, localized error within the brain, while the other bases introduce structured artifacts in background voids, reflecting longer filter supports or non-orthogonal behavior. These visual differences are consistent with the quantitative findings in Table 9. Among all evaluated wavelets, Haar consistently achieved the best trade-off between reconstruction fidelity, computational cost, and artifact suppression, especially in low-signal regions, supporting its choice as the default basis for FlowLet.

A.9. Low-frequency ablations

To assess how much of the current evaluation signal can already be explained by coarse low-frequency structure, we considered two complementary RFM-based *LLL* ablations. In the first setting, FlowLet was trained and sampled using only the approximation band *LLL*, with all seven high-frequency subbands removed throughout both training and inference. In the second setting, the standard full-wavelet FlowLet checkpoints were used, but the generated high-frequency bands were set to zero only at inference time before reconstruction. These two experiments separate the effect of learning exclusively from coarse structure from the effect of suppressing high-frequency content only at reconstruction time.

Table 12 summarizes the results across the full evaluation setup, including global metrics, brain age prediction, and region-based anatomical measures. Both ablations remain competitive under the present evaluation protocol, confirming that coarse anatomical structure carries a substantial fraction of the measurable signal. At the same time, the model trained on full wavelets and evaluated after high-frequency suppression remains consistently stronger than the model trained exclusively on *LLL*. Relative to the dedicated *LLL*-only model, it improves the global metrics (*FID*: 0.3145 vs. 0.3195; *MMD*: 0.0126 vs. 0.0127), the external DLBS age-restricted test MAE (5.32 ± 4.12 vs. 5.62 ± 4.44), the held-out OpenBHB validation test MAE (4.54 ± 3.01 vs. 4.85 ± 3.14), and the region-based overlap score (*DICE*: 0.428 ± 0.165 vs. 0.424 ± 0.172). These results indicate that coarse structure explains a substantial part of the current metric behavior, but they do not imply that the high-frequency bands are irrelevant. Rather, they show that full-wavelet training still improves the retained coarse-scale reconstruction even when the generated high-frequency coefficients are removed before reconstruction.

A.10. High-frequency wavelet coefficient analysis

We additionally evaluated FlowLet directly in wavelet-coefficient space to characterize how closely the generated subbands match the corresponding real-cohort distributions. The analysis used the same 3D Haar transform employed by the model and compared the combined 5793 real MRIs against 3000 generated FlowLet samples in model space. Statistics were computed for all eight bands separately.

The Fig. 10 shows that the seven detail bands remain populated and do not collapse toward zero. The one-high-pass bands retain the strongest overlap with the real cohort, whereas the two-high-pass and fully diagonal bands become progressively less well matched. The *HHH* is the sparsest subband, making it the most challenging band in the visual comparison. This pattern indicates that FlowLet retains

Table 11

Synthetic sample quality metrics across different age groups. Bold and underlined values are best and second-best models. The metrics are computed over 100 bootstrap resamples. Values reported as “–” indicate configurations where age-specific metrics could not be computed, such as unconditional models lacking explicit age control.

Method	Type	Steps	Overall			Age 15–30		
			FID ↓	MMD ↓	MS-SSIM ↓	FID ↓	MMD ↓	MS-SSIM ↓
Ours	RFM	1	0.3334 ± 0.0018	0.0133 ± 0.0001	0.9886 ± 0.0102	0.3525 ± 0.0033	0.0142 ± 0.0001	0.9981 ± 0.0002
		2	0.3232 ± 0.0018	0.0129 ± 0.0001	0.9838 ± 0.0112	0.3425 ± 0.0031	0.0138 ± 0.0001	0.9942 ± 0.0004
		5	0.3130 ± 0.0022	0.0125 ± 0.0001	0.9746 ± 0.0121	0.3333 ± 0.0033	0.0134 ± 0.0001	0.9863 ± 0.0008
		10	0.2981 ± 0.0017	<u>0.0119</u> ± 0.0001	0.9508 ± 0.0195	0.3223 ± 0.0030	<u>0.0130</u> ± 0.0001	0.9698 ± 0.0043
		200	<u>0.2978*</u> ± 0.0020	<u>0.0119*</u> ± 0.0001	0.9487 ± 0.0206	<u>0.3218*</u> ± 0.0031	<u>0.0130*</u> ± 0.0001	0.9687 ± 0.0047
Ours	CFM	1	0.3361 ± 0.0021	0.0134 ± 0.0001	0.9899 ± 0.0093	0.3556 ± 0.0030	0.0143 ± 0.0001	0.9978 ± 0.0005
		2	0.3258 ± 0.0019	0.0130 ± 0.0001	0.9858 ± 0.0104	0.3464 ± 0.0027	0.0139 ± 0.0001	0.9945 ± 0.0006
		5	0.3146 ± 0.0020	0.0126 ± 0.0001	0.9771 ± 0.0111	0.3363 ± 0.0027	0.0135 ± 0.0001	0.9870 ± 0.0009
		10	0.3098 ± 0.0019	0.0124 ± 0.0001	0.9707 ± 0.0117	0.3321 ± 0.0027	0.0134 ± 0.0001	0.9815 ± 0.0014
		200	0.3044 ± 0.0021	0.0122 ± 0.0001	0.9508* ± 0.0182	0.3281 ± 0.0034	0.0132 ± 0.0001	0.9675 ± 0.0055
Ours	VP	1	0.3341 ± 0.0021	0.0133 ± 0.0001	0.9898 ± 0.0092	0.3533 ± 0.0030	0.0142 ± 0.0001	0.9979 ± 0.0003
		2	0.3234 ± 0.0018	0.0129 ± 0.0001	0.9858 ± 0.0101	0.3445 ± 0.0029	0.0138 ± 0.0001	0.9948 ± 0.0004
		5	0.3132 ± 0.0019	0.0125 ± 0.0001	0.9771 ± 0.0109	0.3349 ± 0.0032	0.0135 ± 0.0001	0.9871 ± 0.0008
		10	0.3079 ± 0.0018	0.0123 ± 0.0001	0.9706 ± 0.0115	0.3301 ± 0.0026	0.0133 ± 0.0001	0.9817 ± 0.0013
		200	0.3004 ± 0.0019	0.0120 ± 0.0001	0.9513* ± 0.0183	0.3248 ± 0.0032	0.0131 ± 0.0001	0.9685* ± 0.0057
Ours	Trigon.	1	0.3292 ± 0.0020	0.0131 ± 0.0001	0.9211 ± 0.0084	0.3486 ± 0.0033	0.0140 ± 0.0001	0.9292 ± 0.0004
		2	0.2974* ± 0.0017	0.0119* ± 0.0001	0.9521 ± 0.0112	0.3176 ± 0.0029	0.0128 ± 0.0001	0.9624 ± 0.0004
		5	0.2859 ± 0.0019	0.0114 ± 0.0001	0.9680 ± 0.0130	0.3086 ± 0.0030	0.0124 ± 0.0001	0.9799 ± 0.0008
		10	0.2854 ± 0.0016	0.0114 ± 0.0001	0.9660 ± 0.0119	0.3084 ± 0.0029	0.0124 ± 0.0001	0.9775 ± 0.0011
		200	0.3527 ± 0.0027	0.0141 ± 0.0001	0.9557 ± 0.0165	0.3824 ± 0.0038	0.0154 ± 0.0002	0.9723* ± 0.0039
Ours	FILM	10	0.3252 ± 0.0020	0.0130 ± 0.0001	0.9861 ± 0.0008	0.3469 ± 0.0029	0.0139 ± 0.0001	0.9862 ± 0.0007
	Spatial	10	0.3234 ± 0.0017	0.0129 ± 0.0001	0.9846 ± 0.0022	0.3489 ± 0.0030	0.0140 ± 0.0001	0.9849 ± 0.0011
	Uncond.	10	0.3181 ± 0.0021	0.0127 ± 0.0001	0.9803 ± 0.0014	0.3384 ± 0.0011	0.0136 ± 0.0000	–
WDM	Uncond.	1000	0.3073 ± 0.0018	0.0123 ± 0.0001	0.9456 ± 0.0248	0.3284 ± 0.0013	0.0132 ± 0.0001	–
WDMa	Cond.	1000	0.3166 ± 0.0018	0.0127 ± 0.0001	0.9431 ± 0.0253	0.3315 ± 0.0030	0.0133 ± 0.0001	0.9694* ± 0.0005
MD	Uncond.	1000	0.3843 ± 0.0026	0.0153 ± 0.0001	0.9595 ± 0.0289	0.4072 ± 0.0024	0.0163 ± 0.0001	–
MLDM	Cond.	1000	0.3590 ± 0.0021	0.0144 ± 0.0001	0.9538 ± 0.0259	0.3733 ± 0.0032	0.0150 ± 0.0001	0.9784 ± 0.0024
MOTFM	Uncond.	10	0.3692 ± 0.0024	0.0147 ± 0.0001	0.9677 ± 0.0105	0.3926 ± 0.0021	0.0158 ± 0.0001	–
MOTFMa	Cond.	10	0.3747 ± 0.0027	0.0150 ± 0.0001	0.9529 ± 0.0203	0.3539 ± 0.0033	0.0142 ± 0.0001	0.9775 ± 0.0028
BS	Cond.	–	0.3454 ± 0.0020	0.0138 ± 0.0001	<u>0.9346</u> ± 0.0281	0.3495 ± 0.0035	0.0140 ± 0.0001	<u>0.9600</u> ± 0.0097
Method	Type	Steps	Age 40–55			Age 65–80		
			FID ↓	MMD ↓	MS-SSIM ↓	FID ↓	MMD ↓	MS-SSIM ↓
Ours	RFM	1	0.3511 ± 0.0045	0.0140 ± 0.0002	0.9966 ± 0.0018	0.2955 ± 0.0043	0.0118 ± 0.0002	0.9958 ± 0.0019
		2	0.3409 ± 0.0043	0.0136 ± 0.0002	0.9927 ± 0.0020	0.2844 ± 0.0045	0.0113 ± 0.0002	0.9911 ± 0.0026
		5	0.3303 ± 0.0041	0.0132 ± 0.0002	0.9838 ± 0.0031	0.2743 ± 0.0043	0.0109 ± 0.0002	0.9812 ± 0.0037
		10	0.3169 ± 0.0042	<u>0.0127</u> ± 0.0002	0.9580 ± 0.0129	0.2578 ± 0.0040	0.0103 ± 0.0002	0.9465 ± 0.0145
		200	<u>0.3160*</u> ± 0.0044	<u>0.0127*</u> ± 0.0002	0.9560* ± 0.0138	<u>0.2560*</u> ± 0.0041	<u>0.0102*</u> ± 0.0002	0.9433 ± 0.0155
Ours	CFM	1	0.3547 ± 0.0039	0.0142 ± 0.0002	0.9973 ± 0.0009	0.2971 ± 0.0042	0.0118 ± 0.0002	0.9956 ± 0.0021
		2	0.3445 ± 0.0043	0.0138 ± 0.0002	0.9940 ± 0.0010	0.2851 ± 0.0046	0.0114 ± 0.0002	0.9915 ± 0.0027
		5	0.3335 ± 0.0044	0.0133 ± 0.0002	0.9858 ± 0.0018	0.2741 ± 0.0036	0.0109 ± 0.0002	0.9827 ± 0.0033
		10	0.3283 ± 0.0038	0.0131 ± 0.0002	0.9793 ± 0.0030	0.2692 ± 0.0035	0.0107 ± 0.0001	0.9754 ± 0.0043
		200	0.3226 ± 0.0041	0.0129 ± 0.0002	0.9579* ± 0.0114	0.2631 ± 0.0038	0.0105 ± 0.0002	0.9474* ± 0.0149
Ours	VP	1	0.3524 ± 0.0044	0.0141 ± 0.0002	0.9969 ± 0.0017	0.2953 ± 0.0038	0.0118 ± 0.0002	0.9959 ± 0.0019
		2	0.3416 ± 0.0037	0.0137 ± 0.0002	0.9937 ± 0.0020	0.2842 ± 0.0040	0.0113 ± 0.0002	0.9919 ± 0.0024
		5	0.3320 ± 0.0041	0.0133 ± 0.0002	0.9854 ± 0.0028	0.2731 ± 0.0037	0.0109 ± 0.0002	0.9829 ± 0.0031
		10	0.3258 ± 0.0037	0.0130 ± 0.0002	0.9788 ± 0.0041	0.2685 ± 0.0041	0.0107 ± 0.0002	0.9752 ± 0.0041
		200	0.3184* ± 0.0040	<u>0.0127*</u> ± 0.0002	0.9584* ± 0.0134	0.2596* ± 0.0036	0.0104* ± 0.0001	0.9472* ± 0.0151
Ours	Trigon.	1	0.3484 ± 0.0041	0.0139 ± 0.0002	0.9281 ± 0.0011	0.2910 ± 0.0041	0.0116 ± 0.0002	0.9239 ± 0.0011
		2	0.3168* ± 0.0041	0.0127* ± 0.0002	0.9611* ± 0.0015	0.2599* ± 0.0039	0.0104* ± 0.0002	0.9584 ± 0.0016
		5	0.3062 ± 0.0036	0.0123 ± 0.0002	0.9789 ± 0.0023	0.2473 ± 0.0042	0.0099 ± 0.0002	0.9756 ± 0.0023
		10	0.3042 ± 0.0041	0.0122 ± 0.0002	0.9758 ± 0.0028	0.2451 ± 0.0036	0.0098 ± 0.0001	0.9715 ± 0.0031
		200	0.3731 ± 0.0049	0.0149 ± 0.0002	0.9629* ± 0.0115	0.3046 ± 0.0049	0.0122 ± 0.0002	0.9506* ± 0.0131
Ours	FILM	10	0.3455 ± 0.0040	0.0138 ± 0.0002	0.9860 ± 0.0009	0.2855 ± 0.0049	0.0114 ± 0.0002	0.9862 ± 0.0008
	Spatial	10	0.3450 ± 0.0042	0.0138 ± 0.0002	0.9861 ± 0.0008	0.2809 ± 0.0039	0.0112 ± 0.0002	0.9874 ± 0.0007
	Uncond.	10	0.3383 ± 0.0026	0.0135 ± 0.0001	–	0.2787 ± 0.0017	0.0111 ± 0.0001	–
WDM	Uncond.	1000	0.3277 ± 0.0020	0.0131 ± 0.0001	–	0.2694 ± 0.0017	0.0108 ± 0.0001	–
WDMa	Cond.	1000	0.3335 ± 0.0044	0.0134 ± 0.0002	0.9566* ± 0.0160	0.2825 ± 0.0039	0.0113 ± 0.0002	0.9401 ± 0.0208
MD	Uncond.	1000	0.4074 ± 0.0038	0.0163 ± 0.0002	–	0.3426 ± 0.0024	0.0137 ± 0.0001	–
MLDM	Cond.	1000	0.3681 ± 0.0043	0.0147 ± 0.0002	0.9776 ± 0.0034	0.3311 ± 0.0054	0.0132 ± 0.0002	0.9478 ± 0.0242
MOTFM	Uncond.	10	0.3914 ± 0.0027	0.0157 ± 0.0001	–	0.3269 ± 0.0021	0.0131 ± 0.0001	–
MOTFMa	Cond.	10	0.4001 ± 0.0048	0.0160 ± 0.0002	0.9697 ± 0.0068	0.3685 ± 0.0049	0.0147 ± 0.0002	0.9630 ± 0.0101
BS	Cond.	–	0.3601 ± 0.0055	0.0144 ± 0.0002	<u>0.9463</u> ± 0.0286	0.3188 ± 0.0050	0.0127 ± 0.0002	<u>0.9382</u> ± 0.0234

* Marks results not significantly different from FlowLet-RFM at 10 steps ($p > 0.05$).

structured high-frequency content across all detail bands, while the degree of alignment varies naturally with band complexity.

Table 13 reports the exact band-wise statistics as mean ± standard deviation across volumes, while Fig. 11 provides a compact graphical

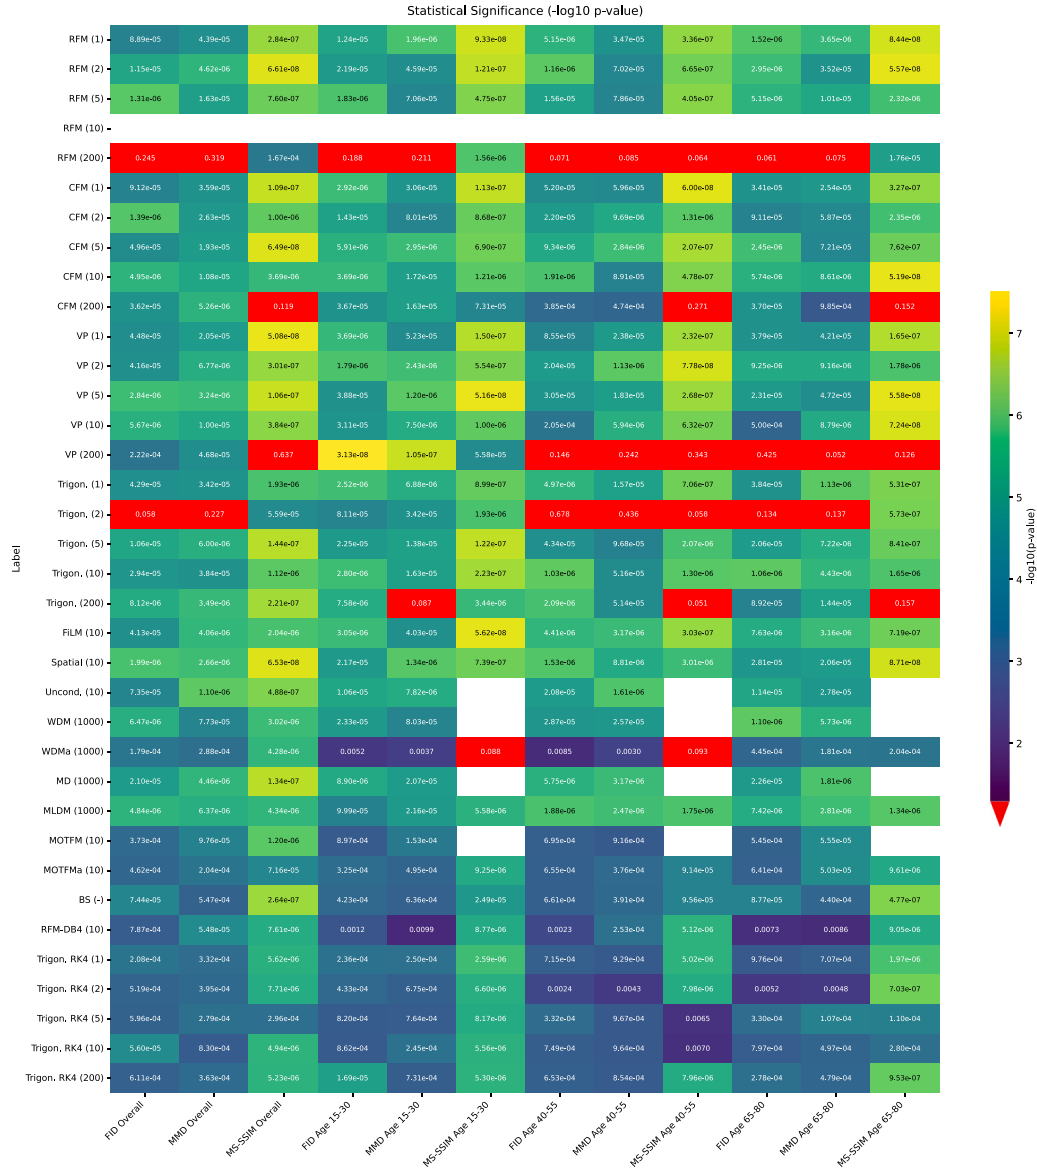


Fig. 9. Statistical Significance ($-\log_{10}$ p-value) of Pairwise Comparisons against FlowLet-RFM (10 Steps). The heatmap displays the results of the Bonferroni-corrected p-values obtained from two-sided Wilcoxon rank-sum tests for pairwise comparisons between the indicated models and the FlowLet-RFM (10 steps) baseline across various fidelity, diversity, and age-stratified metrics. Cells colored red indicate a non-statistically significant difference ($p > 0.05$) from the FlowLet-RFM (10 steps) baseline.

Table 12

Comparison between the two RFM-based low-frequency ablations. The first row corresponds to a FlowLet model trained and sampled using only the approximation band. The second row corresponds to a full-wavelet model reconstructed after zeroing the generated high-frequency bands at inference time.

Setting	Global metrics			BAP Test MAE ↓		ROI		
	FID ↓	MMD ↓	MS-SSIM ↓	DLBS	Held-out OpenBHB	iMAE ↓	KLD ↓	DICE ↑
LLL-only	0.3195 ± 0.0019	0.0127 ± 0.0001	0.9656 ± 0.0134	5.62 ± 4.44	4.85 ± 3.14	37.95 ± 10.14	1.189 ± 1.094	0.424 ± 0.172
HF-zeroed	0.3145 ± 0.0022	0.0126 ± 0.0001	0.9582 ± 0.0205	5.32 ± 4.12	4.54 ± 3.01	37.38 ± 10.68	1.190 ± 1.113	0.428 ± 0.165

summary of the same quantities on a logarithmic scale to facilitate inspection of small relative differences.

All reported quantities are computed from the wavelet coefficients c in each subband, where $|c|$ denotes their absolute value. Mean $|c|$ describes the coefficient magnitude within a band and summarizes how strongly that band is populated on average. Energy measures the overall signal content by giving greater weight to larger coefficients, making it sensitive to the presence of stronger responses. Kurtosis char-

acterizes the sharpness and tail-heaviness of the distribution, helping to distinguish more concentrated coefficient distributions from more diffuse ones. The statistics $|c|$ p95 and $|c|$ p99 denote the 95th and 99th percentiles of the absolute coefficient magnitudes, respectively, and are included to describe the upper tail of the distribution without relying on extreme maxima. Histogram entropy measures how spread the coefficients are across bins, providing a complementary view of distributional diversity. Together, these descriptors allow us to assess




Fig. 10. Absolute-value wavelet coefficient histograms for the real cohort and FlowLet, pooled over all subjects and shown for all eight subbands in model space. The detail bands remain populated and do not collapse toward zero. The strongest overlap with the real cohort is observed in the one-high-pass bands (*LLH*, *LHL*, and *HLL*), whereas the mismatch grows in the two-high-pass bands and is largest in the fully diagonal *HHH* band.

both the average strength of each band and the overall shape of its coefficient distribution relative to the real cohort.

These results show that FlowLet preserves the overall wavelet structure, while the higher-frequency statistics shift in a band-dependent way. The approximation band remains close to the real cohort, and the detail bands stay populated throughout the hierarchy. Across the seven detail bands, mean $|c|$ decreases by 9.78%, energy by 27.98%, $|c|$ p95 by 12.32%, and $|c|$ p99 by 15.53%, while histogram entropy increases by 14.96%. At the band level, the one-high-pass bands remain closest to the real cohort, whereas the most finely structured band shows the largest change, with mean $|c|$ decreasing by 20.19%, energy by 49.82%, $|c|$ p95 by 24.33%, and $|c|$ p99 by 28.65%.

These results show that the high-frequency mismatch is neither a simple collapse nor a uniform global shift. The analysis shows that nonzero detail coefficients are retained across all subbands. Fig. 11 and Table 13 show that the coefficient typical amplitude, tail strength, and kurtosis remain closer to the real cohort in the easier bands and

become progressively more challenging in the more diagonal ones. The higher-order scatter further shows that this is not just a rescaling effect: the generated high-frequency distributions shift toward lower energy, lower kurtosis, and higher histogram entropy, which is consistent with slightly weaker but more diffuse coefficient distributions.

A.11. Fixed-seed age-conditioning trajectory

Fig. 12 shows the effect of isolated age conditioning. We generated a controlled age trajectory with a single fixed seed (42) and a fixed target-age grid of 6, 15, 25, 35, 45, 55, 65, 75, 85, and 95 years. Only the conditioning variable was changed across columns, while the initialization and sampling steps (10) were held constant. The resulting sequence shows coherent age-dependent morphological variation, including progressive ventricular enlargement, sulcal widening, and cortical thinning in later decades.




Fig. 11. Band-wise comparison of magnitude, tail, and higher-order summary metrics in model space. FlowLet remains close to the real cohort in the one-high-pass bands, while the attenuation becomes stronger in the two-high-pass bands and is largest in *HHH*, where mean $|c|$, energy, upper-tail statistics, and kurtosis are reduced, while histogram entropy is increased.

Table 13

Summary of overall high-frequency wavelet statistics for real data and FlowLet. Each value is reported as mean $\pm$ standard deviation across volumes.

Band	Mean $ c $		Energy		Kurtosis	
	Real	FlowLet	Real	FlowLet	Real	FlowLet
LLH	0.0192 ± 0.0017	0.0180 ± 0.0015	0.0042 ± 0.0010	0.0034 ± 0.0007	27.598 ± 9.454	26.298 ± 3.474
LHL	0.0188 ± 0.0016	0.0180 ± 0.0015	0.0040 ± 0.0009	0.0034 ± 0.0007	26.995 ± 6.435	25.058 ± 3.243
LHH	0.0070 ± 0.0017	0.0061 ± 0.0015	0.0006 ± 0.0004	0.0004 ± 0.0002	28.902 ± 13.936	25.631 ± 3.902
HLL	0.0202 ± 0.0016	0.0196 ± 0.0015	0.0045 ± 0.0008	0.0040 ± 0.0007	25.203 ± 6.642	23.179 ± 3.205
HLH	0.0069 ± 0.0016	0.0061 ± 0.0015	0.0006 ± 0.0003	0.0004 ± 0.0002	29.905 ± 17.790	26.409 ± 3.239
HHL	0.0069 ± 0.0015	0.0063 ± 0.0015	0.0006 ± 0.0003	0.0004 ± 0.0002	29.704 ± 9.402	25.652 ± 3.532
HHH	0.0038 ± 0.0018	0.0031 ± 0.0014	0.0002 ± 0.0002	0.0001 ± 0.0001	32.007 ± 21.929	26.859 ± 5.031
Band	$ c $ p95		$ c $ p99		Hist. entropy	
	Real	FlowLet	Real	FlowLet	Real	FlowLet
LLH	0.1399 ± 0.0155	0.1289 ± 0.0121	0.3071 ± 0.0414	0.2784 ± 0.0319	1.293 ± 0.168	1.476 ± 0.151
LHL	0.1363 ± 0.0137	0.1273 ± 0.0111	0.3070 ± 0.0411	0.2805 ± 0.0329	1.195 ± 0.107	1.320 ± 0.143
LHH	0.0496 ± 0.0119	0.0415 ± 0.0099	0.1129 ± 0.0353	0.0905 ± 0.0250	1.301 ± 0.168	1.453 ± 0.154
HLL	0.1478 ± 0.0150	0.1409 ± 0.0125	0.3256 ± 0.0336	0.3041 ± 0.0298	1.224 ± 0.109	1.459 ± 0.148
HLH	0.0484 ± 0.0115	0.0417 ± 0.0102	0.1108 ± 0.0335	0.0906 ± 0.0249	1.293 ± 0.168	1.479 ± 0.153
HHL	0.0486 ± 0.0108	0.0425 ± 0.0099	0.1122 ± 0.0317	0.0927 ± 0.0243	1.185 ± 0.118	1.426 ± 0.144
HHH	0.0269 ± 0.0125	0.0203 ± 0.0096	0.0605 ± 0.0327	0.0432 ± 0.0219	1.288 ± 0.182	1.477 ± 0.174




Fig. 12. Fixed-seed age-conditioning trajectory generated with FlowLet-RFM. Each column corresponds to a different target age, while the seed is kept fixed across the full sequence. Axial, coronal, and sagittal views are shown for multiple ages.




Fig. 13. Age distribution of the external evaluation cohorts. The held-out OpenBHB validation subset is concentrated among younger adults, while the DLBS subset spans the adult lifespan, with greater coverage in middle to late adulthood.

A.12. Dallas lifespan brain study

We further evaluated FlowLet on an external cohort derived from the Dallas Lifespan Brain Study (DLBS), a longitudinal adult-lifespan study designed to jointly characterize brain and cognition across healthy adulthood.¹⁰ For this analysis work, we extracted the T1-weighted structural MRI subset, retained only healthy subjects, and applied the same minimal preprocessing pipeline adopted for the main experiments. The resulting processed subset comprises 956 scans, with mean age 57 ± 17 years, minimum age 21 years, and maximum age 89 years. In this form, DLBS serves as an independent adult-lifespan validation cohort with stronger coverage of later adulthood than OpenBHB, making it a useful complement to the merged training distribution used in the main study. This contrast is illustrated in Fig. 13, where the held-out OpenBHB subset remains concentrated in younger ages, whereas DLBS provides broader support across middle and older adulthood.

A.13. Real-to-real global metric references

To better contextualize the global similarity metrics, we computed real-to-real references using the same evaluation pipeline adopted for the synthetic comparisons. Specifically, we compared two independent real cohorts against the merged real-cohort reference: (i) the held-out OpenBHB validation subset reserved for downstream testing, and (ii) the processed DLBS subset introduced above. The held-out OpenBHB comparison remains close to the merged real cohort overall, while still showing non-zero differences, particularly in the older age bin, which reflects the expected shift between real subsets drawn from a heterogeneous lifespan distribution (Table 14). The DLBS comparison is also close to the merged real cohort (Table 15), indicating that after the same preprocessing and healthy-subject filtering, DLBS occupies a compatible but not identical region of the real-data distribution. Taken together, these real-to-real references provide a practical dataset-level scale for interpreting the synthetic global metrics, showing that non-zero FID, MMD, and MS-SSIM values also arise across genuine real cohorts due to natural biological variability and residual inter-cohort differences. While the synthetic-to-real values remain higher than the real-to-real references, these baselines confirm that non-zero metric values naturally arise even between genuine real cohorts, providing a practical scale for interpreting the reported results.

¹⁰ Park et al. *The Dallas Lifespan Brain Study: A Comprehensive Adult Lifespan Data Set of Brain and Cognitive Aging*, Scientific Data, 2025. DOI: 10.1038/s41597-025-04847-7.

Table 14

Real-to-real metric reference obtained by comparing the held-out OpenBHB validation subset with the merged real cohort.

Age group	Samples	FID ↓	MMD ↓	MS-SSIM ↓
Overall	757	0.0069	0.0104	0.9233
15–30	497	0.0037	0.0074	0.9497
40–55	35	0.0030	0.0066	0.9467
65–80	17	0.0186	0.0220	0.8883

Table 15

Real-to-real global metrics for the DLBS subset versus the merged real cohort.

Age group	Samples	FID ↓	MMD ↓	MS-SSIM ↓
Overall	956	0.0018	0.0014	0.9194
15–30	100	0.0011	0.0037	0.9348
40–55	200	0.0009	0.0027	0.9294
65–80	200	0.0003	0.0031	0.9212

Table 16

External DLBS brain age prediction performance for the Age ≥ 44 years group. Lower Mean Absolute Error (MAE, in years) = better accuracy.

	Model	Train MAE ↓	Test MAE ↓
	Real data	1.15 ± 1.02	6.38 ± 4.97
Ours	RFM	1.46 ± 0.59	5.24 ± 4.38
	CFM	1.39 ± 0.59	5.29 ± 4.26
	VP	1.02 ± 0.49	5.27 ± 4.49
	Trigon.	1.09 ± 0.48	5.26 ± 4.48
Ablat.	FiLM	0.57 ± 0.51	6.37 ± 5.06
	Spatial	0.87 ± 0.54	5.79 ± 4.53
	Uncond.	1.46 ± 1.06	5.97 ± 4.83
Baselines	WDM	1.63 ± 1.36	7.36 ± 6.20
	WDMa	0.33 ± 0.42	5.60 ± 4.58
	MD	2.54 ± 2.78	5.81 ± 4.37
	MLDM	0.98 ± 0.47	5.97 ± 4.75
	MOTFMa	0.85 ± 0.53	5.53 ± 4.48
	BS	0.90 ± 0.40	5.55 ± 4.33

A.14. External DLBS validation

We next evaluated the downstream augmentation pipelines on DLBS using the same age-restricted BAP setting adopted in the main paper for subjects aged 44 years and older. Results are shown in Table 16. On this external cohort, age-conditioned augmentation remained beneficial relative to the real-only baseline. The real-only model reached test MAE 6.38 ± 4.97 , whereas the FlowLet variants improved this value to 5.24 ± 4.38 (RFM), 5.29 ± 4.26 (CFM), 5.27 ± 4.49 (VP), and 5.26 ± 4.48 (Trigonometric). The strongest non-FlowLet baseline among the reported methods was MOTFMa with test MAE 5.53 ± 4.48 , followed by BrainSynth with 5.55 ± 4.33 and WDMa with 5.60 ± 4.58 . These findings indicate that the benefit of age-conditioned synthetic augmentation relative to the real-only training baseline extends to an independent external adult-lifespan cohort.

References

- Albergo, M.S., Vanden-Eijnden, E., 2023. Building normalizing flows with stochastic interpolants. In: ICLR. OpenReview.net.
- Bacon, E.J., He, D., Achi, N.A.D., Wang, L., Li, H., Yao-Digba, P.D.Z., Monkam, P., Qi, S., 2024. Neuroimage analysis using artificial intelligence approaches: a systematic review. *Med. Biol. Eng. Comput.* 62 (9), 2599–2627.
- Baecker, L., Garcia-Dias, R., Vieira, S., Scarpazza, C., Mechelli, A., 2021. Machine learning for brain age prediction: Introduction to methods and clinical applications. *EBioMedicine* 72.
- Bashyam, V.M., Erus, G., Doshi, J., Habes, M., Nasrallah, I.M., Truelove-Hill, M., Srinivasan, D., Mamourian, L., Pomponio, R., Fan, Y., et al., 2020. MRI signatures of brain age and disease over the lifespan based on a deep brain network and 14 468 individuals worldwide. *Brain* 143 (7), 2312–2324.

- Bethlehem, R.A.I., et al., 2022. Brain charts for the human lifespan. *Nature* 604 (7906), 525–533.
- Blagus, R., Lusa, L., 2013. SMOTE for high-dimensional class-imbalanced data. *BMC Bioinform.* 14, 106.
- Bullmore, E., Fadili, J., Maxim, V., Şendur, L., Whitcher, B., Suckling, J., Brammer, M., Breakspear, M., 2004. Wavelets and functional magnetic resonance imaging of the human brain. *NeuroImage* 23, <http://dx.doi.org/10.1016/j.neuroimage.2004.07.012>, Mathematics in Brain Imaging.
- Button, K., Ioannidis, J., Mokrysz, C., Nosek, B., Flint, J., Robinson, E., Munafò, M., 2013. Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Rev. Neurosci.* 14, <http://dx.doi.org/10.1038/nrn3475>.
- Chen, S., Ma, K., Zheng, Y., 2019. Med3D: Transfer learning for 3D medical image analysis. *CoRR abs/1904.00625*.
- Chintapalli, S.S., Wang, R., Yang, Z., Tassopoulou, V., Yu, F., Bashyam, V., Erus, G., Chaudhari, P., Shou, H., Davatzikos, C., 2024. Generative models of MRI-derived neuroimaging features and associated dataset of 18,000 samples. *Sci. Data* 11 (1), 1330.
- Cole, J.H., Poudel, R.P., Tsagkrasoulis, D., Caan, M.W., Steves, C., Spector, T.D., Montana, G., 2017. Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *NeuroImage* 163, 115–124.
- Daubechies, I., 1992. Ten Lectures on Wavelets. SIAM.
- Davtyan, A., Sameni, S., Favaro, P., 2023. Efficient video prediction via sparsely conditioned flow matching. In: ICCV. IEEE.
- De Bonis, M.L.N., Fasano, G., Lombardi, A., Ardito, C., Ferrara, A., Di Sciascio, E., Di Noia, T., 2024. Explainable brain age prediction: a comparative evaluation of morphometric and deep learning pipelines. *Brain Informatics* 11 (1), 33.
- Dhariwal, P., Nichol, A.Q., 2021. Diffusion models beat GANs on image synthesis. In: *NeurIPS*. pp. 8780–8794.
- Dinsdale, N.K., Bluemke, E., Smith, S.M., Arya, Z., Vidaurre, D., Jenkinson, M., Namburete, A.I., 2021. Learning patterns of the ageing brain in MRI using deep convolutional networks. *NeuroImage* 224, 117401.
- Dufumier, B., Grigis, A., Victor, J., Ambroise, C., Frouin, V., Duchesnay, E., 2022. Openbhb: a large-scale multi-site brain mri data-set for age prediction and debiasing [dataset]. *NeuroImage* 263.
- Durrer, A., Wolleb, J., Bieder, F., Sinnecker, T., Weigel, M., Sandkühler, R., Granziera, C., Yaldizli, Ö., Cattin, P.C., 2023. Diffusion models for contrast harmonization of magnetic resonance images. In: *MIDL*. In: *Proceedings of Machine Learning Research*, vol. 227, PMLR, pp. 526–551.
- Efron, B., 2011. Tweedie's formula and selection bias. *J. Am. Stat. Assoc.* 106 (496), 1602–1614.
- Fernandez, V., Pinaya, W.H.L., Borges, P., Graham, M.S., Tudosiu, P.-D., Vercauteren, T., Cardoso, M.J., 2024. Generating multi-pathological and multi-modal images and labels for brain MRI. *Med. Image Anal.* 97, 103278.
- Fonov, V., Evans, A., McKinstry, R., Almlí, C., Collins, D., 2009. Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage* 47, S102. [http://dx.doi.org/10.1016/S1053-8119\(09\)70884-5](http://dx.doi.org/10.1016/S1053-8119(09)70884-5), URL: <https://www.sciencedirect.com/science/article/pii/S1053811909708845>. Organization for Human Brain Mapping 2009 Annual Meeting.
- Friedrich, P., Wolleb, J., Bieder, F., Durrer, A., Cattin, P.C., 2024. WDM: 3D wavelet diffusion models for high-resolution medical image synthesis. In: *DGM4MICCAI@MICCAI*. In: *Lecture Notes in Computer Science*, vol. 15224, Springer.
- Fu, J., Zheng, Y., Dey, N., Ferreira, D., Moreno, R., 2025. Synthesizing individualized aging brains in health and disease with generative models and parallel transport. *Med. Image Anal.* 103669.
- Gagneux, A., Martin, S.T., Emonet, R., Bertrand, Q., Massias, M., 2025a. A visual dive into conditional flow matching. In: *The Fourth Blogpost Track At ICLR 2025*.
- Gagneux, A., Martin, S.T., Emonet, R., Bertrand, Q., Massias, M., 2025b. A visual dive into conditional flow matching. In: *The Fourth Blogpost Track At ICLR 2025*. URL: <https://openreview.net/forum?id=M6MTUQc4um>.
- Gaser, C., Kalc, P., Cole, J.H., 2024. A perspective on brain-age estimation and its clinical promise. *Nat. Comput. Sci.* 4 (10), 744–751.
- Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y., 2014. Generative adversarial nets. In: *NIPS*. pp. 2672–2680.
- Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A.J., 2012. A kernel two-sample test. *J. Mach. Learn. Res.* 13, 723–773.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S., 2017. GANs trained by a two-time-scale update rule converge to a local Nash equilibrium. In: *NIPS*. pp. 6626–6637.
- Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models. In: *NeurIPS*.
- Jafrasteh, B., Peng, W., Wan, C., Luo, Y., Adeli, E., Zhao, Q., 2025. WASABI: A metric for evaluating morphometric plausibility of synthetic brain MRIs. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 684–694.
- Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* 17 (2), 825–841. <http://dx.doi.org/10.1006/nimg.2002.1132>.
- Khader, F., Mueller-Franzes, G., Arasteh, S.T., Han, T., Haarbuerger, C., Schulze-Hagen, M., Schad, P., Engelhardt, S., Baeßler, B., Foersch, S., Stegmaier, J., Kuhl, C., Nebelung, S., Kather, J.N., Truhn, D., 2022. Medical diffusion - denoising diffusion probabilistic models for 3D medical image generation. *CoRR abs/2211.03364*.
- Kingma, D.P., Welling, M., 2014. Auto-encoding variational Bayes. In: *ICLR*.
- Lee, S., Kim, B., Ye, J.C., 2023. Minimizing trajectory curvature of ODE-based generative models. In: *ICML*. In: *Proceedings of Machine Learning Research*, vol. 202, PMLR, pp. 18957–18973.
- Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M., 2023. Flow matching for generative modeling. In: *ICLR*. *OpenReview.net*.
- Litrico, M., Guarnera, F., Giuffrida, M.V., Ravì, D., Battiato, S., 2024. TADM: Temporally-aware diffusion model for neurodegenerative progression on brain MRI. In: *MICCAI* (2). Springer, pp. 444–453.
- Liu, X., Gong, C., Liu, Q., 2023. Flow straight and fast: Learning to generate and transfer data with rectified flow. In: *ICLR*. *OpenReview.net*.
- Lombardi, A., Diacono, D., Amoroso, N., Monaco, A., Tavares, J.M.R., Bellotti, R., Tangaro, S., 2021. Explainable deep learning for personalized age prediction with brain morphology. *Front. Neurosci.* 15, 674055.
- Mallat, S., 1989. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.* 11 (7), 674–693.
- Mallat, S., 1999. A Wavelet Tour of Signal Processing, second ed. Academic Press.
- Müller-Franzes, G., Niehues, J.M., Khader, F., Arasteh, S.T., Haarbuerger, C., Kuhl, C., Wang, T., Han, T., Nolte, T., Nebelung, S., Kather, J.N., Truhn, D., 2023. A multimodal comparison of latent denoising diffusion probabilistic models and generative adversarial networks for medical image synthesis. *Sci. Reports* 13 (1), 12098.
- Neklyudov, K., Brekelmans, R., Severo, D., Makhzani, A., 2023. Action matching: Learning stochastic dynamics from samples. In: *ICML*. In: *Proceedings of Machine Learning Research*, vol. 202, PMLR.
- Nguyen, H.M., Cooper, E.W., Kamei, K., 2011. Borderline over-sampling for imbalanced data classification. *Int. J. Knowl. Eng. Soft Data Paradigms* 3 (1), 4–21.
- Nichol, A.Q., Dhariwal, P., 2021. Improved denoising diffusion probabilistic models. In: *ICML*. In: *Proceedings of Machine Learning Research*, vol. 139, PMLR, pp. 8162–8171.
- Peng, W., Bosschieter, T., Ouyang, J., Paul, R., Sullivan, E.V., Pfefferbaum, A., Adeli, E., Zhao, Q., Pohl, K.M., 2024. Metadata-conditioned generative models to synthesize anatomically-plausible 3D brain MRIs. *Med. Image Anal.* 98, 103325.
- Peng, H., Gong, W., Beckmann, C.F., Vedaldi, A., Smith, S.M., 2021. Accurate brain age prediction with lightweight deep neural networks. *Med. Image Anal.* 68, 101871.
- Perez, E., Strub, F., de Vries, H., Dumoulin, V., Courville, A.C., 2018. FiLM: Visual reasoning with a general conditioning layer. In: *AAAI*. AAAI Press, pp. 3942–3951.
- Phung, H., Dao, Q., Tran, A., 2023. Wavelet diffusion models are fast and scalable image generators. In: *CVPR*. IEEE.
- Pinaya, W.H., 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: *MICCAI Workshop on Deep Generative Models*. Springer, pp. 117–126.
- Pombo, G., Gray, R.J., Cardoso, M.J., Ourselin, S., Rees, G., Ashburner, J., Nachev, P., 2023. Equitable modelling of brain imaging by counterfactual augmentation with morphologically constrained 3D deep generative models. *Med. Image Anal.* 84, 102723.
- Rahman, M.T., Orka, N.A., Khan, A., Liò, P., Moni, M.A., 2025. Understanding neurocognition with deep learning and MRI: A systematic review. *IEEE Trans. Cogn. Dev. Syst.*
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B., 2022. High-resolution image synthesis with latent diffusion models. In: *CVPR*. IEEE, pp. 10674–10685.
- Salimans, T., Goodfellow, I.J., Zaremba, W., Cheung, V., Radford, A., Chen, X., 2016. Improved techniques for training GANs. In: *NIPS*. pp. 2226–2234.
- Shorten, C., Khoshgoftaar, T.M., 2019. A survey on image data augmentation for deep learning. *J. Big Data* 6, 60.
- Smith, S.M., 2002. Fast robust automated brain extraction. *Hum. Brain Mapp.* 17 (3), 143–155.
- Song, J., Meng, C., Ermon, S., 2021. Denoising diffusion implicit models. In: *ICLR*. *OpenReview.net*.
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B., 2021. Score-based generative modeling through stochastic differential equations. In: *ICLR*. *OpenReview.net*.
- Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., Liu, B., Matthews, P., Ong, G., Pell, J., Silman, A., Young, A., Sprosen, T., Peakman, T., Collins, R., 2015. UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med* 12 (3), e1001779.
- Tong, A., Malkin, N., Huguët, G., Zhang, Y., Rector-Brooks, J., Fatras, K., Wolf, G., Bengio, Y., 2023. Conditional flow matching: Simulation-free dynamic optimal transport. *CoRR abs/2302.00482*.
- Tudosiu, P., Pinaya, W.H.L., Costa, P.F.D., Dafflon, J., Patel, A., Borges, P., Fernandez, V., Graham, M.S., Gray, R.J., Nachev, P., Ourselin, S., Cardoso, M.J., 2024. Realistic morphology-preserving generative modelling of the brain. *Nat. Mac. Intell.* 6 (7), 811–819.
- Tustison, N.J., Avants, B.B., Cook, P.A., Zheng, Y., Egan, A., Yushkevich, P.A., Gee, J.C., 2010. N4ITK: improved N3 bias correction. *IEEE Trans. Med. Imaging*.

- Tustison, N.J., Cook, P.A., Holbrook, A.J., Johnson, H.J., Muschelli, J., Devenyi, G.A., Duda, J.T., Das, S.R., Cullen, N.C., Gillen, D.L., Yassa, M.A., Stone, J.R., Gee, J.C., Avants, B.B., 2021. The ANTsX ecosystem for quantitative biological and medical imaging. *Sci. Reports*.
- Vahdat, A., Kreis, K., Kautz, J., 2021. Score-based generative modeling in latent space. In: *NeurIPS*. pp. 11287–11302.
- van den Oord, A., Vinyals, O., Kavukcuoglu, K., 2017. Neural discrete representation learning. In: *NIPS*. pp. 6306–6315.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need. In: *NIPS*. pp. 5998–6008.
- Wallis, D., Buvat, I., 2022. Clever Hans effect found in a widely used brain tumour MRI dataset. *Med. Image Anal.* 77, 102368. <http://dx.doi.org/10.1016/j.media.2022.102368>.
- Wang, Z., et al., 2003. Multiscale structural similarity for image quality assessment. In: *The Thirty-Seventh Asilomar Conference on Signals, Systems and Computers*, 2003. 2, <http://dx.doi.org/10.1109/ACSSC.2003.1292216>.
- Wu, J., Fu, R., Fang, H., Zhang, Y., Yang, Y., Xiong, H., Liu, H., Xu, Y., 2023. Med-SegDiff: Medical image segmentation with diffusion probabilistic model. In: *MIDL*. In: *Proceedings of Machine Learning Research*, vol. 227, PMLR, pp. 1623–1639.
- Wu, G., Suk, H.-I., 2017. Deep learning in medical image analysis. *Annu. Rev. Biomed. Eng.* 19, <http://dx.doi.org/10.1146/annurev-bioeng-071516-044442>.
- Wu, J., Wang, Z., Hong, M., Ji, W., Fu, H., Xu, Y., Xu, M., Jin, Y., 2025. Medical sam adapter: Adapting segment anything model for medical image segmentation. *Med. Image Anal.* 102, 103547.
- Wyatt, J., Leach, A., Schmon, S.M., Willcocks, C.G., 2022. AnoDDPM: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In: *CVPR Workshops*. IEEE, pp. 649–655.
- Xiao, Z., Kreis, K., Vahdat, A., 2022. Tackling the generative learning trilemma with denoising diffusion GANs. In: *ICLR*. [OpenReview.net](https://openreview.net).
- Yazdani, M., Medghalchi, Y., Ashrafian, P., Hacihaliloglu, I., Shahriari, D., 2025. Flow matching for medical image synthesis: Bridging the gap between speed and quality. *arXiv preprint arXiv:2503.00266*.
- Yeganeh, Y., Farshad, A., Charisiadis, I., Hasny, M., Hartenberger, M., Ommer, B., Navab, N., Adeli, E., 2025. Latent drifting in diffusion models for counterfactual medical image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. CVPR.
- Yu, J.J., Derpanis, K.G., Brubaker, M.A., 2020. Wavelet flow: Fast training of high resolution normalizing flows. In: *NeurIPS*.
- Yu, Z., Zhao, B., Zhang, S., Chen, X., Yan, F., Feng, J., Peng, T., Zhang, X.-Y., 2025. HiFi-Syn: Hierarchical granularity discrimination for high-fidelity synthesis of MR images with structure preservation. *Med. Image Anal.* 100, 103390.
- Zhang, J., Huang, Q., Liu, J., Guo, X., Huang, D., 2025a. Diffusion-4K: Ultra-high-resolution image synthesis with latent diffusion models. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025*, Nashville, TN, USA, June 11–15, 2025. Computer Vision Foundation / IEEE, pp. 23464–23473. <http://dx.doi.org/10.1109/CVPR52734.2025.02185>, URL: https://openaccess.thecvf.com/content/CVPR2025/html/Zhang_Diffusion-4K_Ultra-High-Resolution_Image_Synthesis_with_Latent_Diffusion_Models_CVPR_2025_paper.html.
- Zhang, J., Huang, Q., Liu, J., Guo, X., Huang, D., 2025b. Diffusion-4K: Ultra-high-resolution image synthesis with latent diffusion models. *CoRR* abs/2503.18352.