

GenTract: Generative Global Tractography

Alec Sargood¹ Lemuel Puglisi² Elinor Thompson¹ Mirco Musolesi³ Daniel C. Alexander¹
¹Hawkes Institute and Department of Computer Science, University College London, UK
²Department of Maths and Computer Science, University of Catania, Italy
³AI Centre and Department of Computer Science, University College London, UK

alec.sargood.23@ucl.ac.uk

Abstract

Tractography is the process of inferring the trajectories of white-matter pathways in the brain from diffusion magnetic resonance imaging (dMRI). Local tractography methods, which construct streamlines by following local fiber orientation estimates stepwise through an image, are prone to error accumulation and high false positive rates, particularly on noisy or low-resolution data. In contrast, global methods, which attempt to optimize a collection of streamlines to maximize compatibility with underlying fiber orientation estimates, are computationally expensive. To address these challenges, we introduce GenTract, the first generative model for global tractography. We frame tractography as a generative task, learning a direct mapping from dMRI to complete, anatomically plausible streamlines. We compare both diffusion-based and flow matching paradigms and evaluate GenTract’s performance against state-of-the-art baselines. Notably, GenTract achieves precision $1.8\times$ and $2.1\times$ higher than the next-best methods, DDTracking and TractOracle, respectively. This advantage becomes even more pronounced in challenging low-resolution and noisy settings, where it outperforms the closest competitor by a factor of 3.5. By producing tractograms with high precision on research-grade data while also maintaining reliability on imperfect, lower-resolution data, GenTract represents a promising solution for global tractography.

1. Introduction

The intricate network of white matter pathways forming the human connectome provides the structural architecture for all neural communication [9]. The ability to map these connections in vivo with diffusion-weighted Magnetic Resonance Imaging (dMRI) has offered insights into brain development [6], aging [20], and the pathophysiology of neurological disorders [7], while also serving as a tool for neurosurgical planning [15]. However, current methods to reconstruct these pathways, known as tractography, face fun-

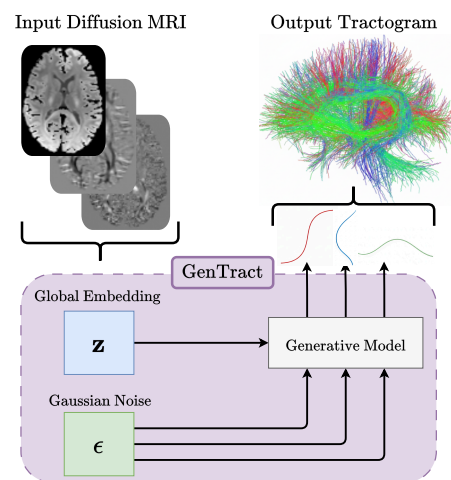


Figure 1. Overview of our proposed GenTract methodology. We provide a learned embedding of, \mathbf{z} , representing whole dMRI information. Starting from Gaussian noise (ϵ), all coordinates of individual streamlines are generated in parallel. Generated streamlines are collated to form the output tractogram.

damental challenges that have limited their accuracy and reliability [21, 23].

Tractography algorithms use information from a dMRI scan, which comprises a set of Diffusion Weighted Images (DWIs), each sensitive to water diffusion in a distinct direction. Because water diffuses more readily along brain fibers than across them, the full set of DWIs can be used to estimate the distribution of fiber orientations within each voxel. The dominant tractography paradigm, local tracking [5, 18, 24, 35, 36], uses these voxel-wise orientations to trace the trajectory of the white matter fibers step-by-step in the form of ‘streamlines’: lines in 3D space that represent the best estimates of the underlying fiber trajectories. However, global connectivity mapping using the local search strategy makes these methods prone to error accumulation, resulting in high false-positive connections [21, 30]. This

problem is especially pronounced in clinical dMRI scans, which are often limited to acquisition times of a few minutes leading to lower spatial resolution and signal-to-noise ratio compared to the longer, high-resolution protocols used in research. In addition, these methods depend on heterogeneous preprocessing pipelines, especially for constructing seeding masks that define streamline initiation points. This process, which may involve anything from manual segmentation to thresholding of tissue masks, introduces substantial operator- and method-dependent variability, undermining reproducibility [16, 35].

Global tractography [4] seeks to address these limitations by inferring the most probable streamline configurations across the entire brain simultaneously through a single optimization problem. Compared with local methods, global approaches reduce dependence on seeding masks and leverage broader anatomical context. However, existing global approaches depend on traditional optimization schemes that are computationally inefficient, susceptible to suboptimal solutions, and often fail to generate complete tractograms, thereby limiting their practical utility [14, 21, 29].

Contributions. In this paper, we introduce **GenTract** (**Generative Tractography**), reframing tractography as a generative task, conditioned on global information, as illustrated in Figure 1. GenTract is designed to resolve the limitations of both paradigms: it tackles the reliance of local tracking on step-wise information, while solving the computational inefficiency and optimization failures of existing global methods. Our primary contributions are:

1. We are the first to formulate tractography as a global, conditional generative task conditioned on the whole set of fiber orientation distribution functions (fODFs) for a given individual. GenTract treats tractography as a sampling process, where the model draws streamlines from a learned conditional distribution, producing all streamline coordinates simultaneously. This overcomes the sequential error propagation of step-wise methods, and also eliminates the need for a seeding mask, reducing method-dependent variability.
2. We propose a modular generative architecture for global tractography compatible with multiple generative paradigms. We provide the first implementation and analysis of both Diffusion and Flow Matching for the task of global tractography, establishing the effectiveness of both schemes.
3. We conduct a comparative analysis of GenTract against traditional and state-of-the-art (SOTA) tractography paradigms. This analysis demonstrates that GenTract achieves state-of-the-art level of precision on high-fidelity data while exhibiting greater robustness in noisy, low-resolution settings.

2. Related Works

Local Methods. Classical local tractography methods, such as SD Stream [5] and iFOD2 [36], track streamlines by using voxel-wise diffusion information to take the next tracking step. However, they break down in complex fiber configurations like crossing or kissing ‘bundles’ [29]: collections of fibers that form neuroanatomical pathways, connecting two regions of the brain. Machine learning (ML) based techniques have increasingly advanced the field of tractography, improving accuracy over classical methods [40]. For example, the authors in [24] propose the first ML-based approach, employing random forest classifiers to model the relationship between local image data and sequential tracking steps. More recently, TrackToLearn [34] reframes tractography as a reinforcement learning problem, where agents start from seed points and are guided by rewards that encourage geometrically plausible trajectories and alignment with the diffusion signal. Despite being able to produce more complete connectivity mappings, these methods still produce many false positives, especially in low-resolution data where blurred fiber orientations cause tracking drift. To address this issue, TractOracle [35] builds on TrackToLearn by incorporating anatomical priors into the reward function, providing agents with additional contextual information, aiding tracking and lowering false positives. Similarly, DDTracking [18], a supervised diffusion-based tracking algorithm, uses neighboring voxel information to generate more reliable next-step predictions. However, these methods still propagate streamlines in an autoregressive manner and are thus limited in their robustness to underlying data perturbations.

Global Methods. Existing global tractography methods [4, 13] address the error accumulation of local methods by treating fiber reconstruction as a single inverse problem, optimizing the entire whole-brain fiber configuration simultaneously to find the solution that best explains the underlying diffusion data. This concept is introduced in [22], which models short fiber segments as ‘spins’ and finds the lowest-energy state for the system. This is improved by [8], where these segments can move and rotate to better fit complex crossings. A major limitation is computation time, which is addressed by the authors in [28], who introduce an efficient algorithm that makes global reconstruction more practical by optimizing a global energy function using a faster particle-based approach. The current standard, tckglobal [2], further refines this by integrating a multi-tissue model. However, these classical optimization approaches still remain a significant computational bottleneck, limiting their application for generating dense whole-brain tractograms. Furthermore, they can exhibit optimization failures, and get stuck in sub-optimal solutions, causing segments to incor-

rectly connect or terminate prematurely [2, 29]. As such, although more robust to local perturbations, global methods have not been adopted as mainstream approaches for tractography. GenTract leverages deep generative modeling to learn a conditional distribution over streamlines based on the entire diffusion signal. This approach maintains the robustness of classical global methods while replacing their costly optimization with a generative process that produces all streamline coordinates in parallel. This process is guided by context from the entire dMRI volume, leading to improvements in both precision and efficiency.

3. Preliminaries

dMRI Representations. Most modern tractography algorithms rely on per-voxel estimates of the fODF for tracking. The fODF, $f(\theta, \phi)$, is a continuous function on the sphere, defined by polar angle θ and azimuthal angle ϕ , that represents the distribution of fiber orientations within the voxel. It is designed to capture complex fiber configurations, and is typically estimated from the raw DWIs. For computational tractability, the continuous fODF in each voxel is projected onto an orthogonal Spherical Harmonic (SH) basis, $Y_l^k(\theta, \phi)$:

$$f(\theta, \phi) \approx \sum_{l=0}^{L_{\max}} \sum_{k=-l}^l \alpha_{lk} Y_l^k(\theta, \phi). \quad (1)$$

Here, l and k are the SH degree and order, and L_{\max} is the maximum harmonic order (e.g., a common choice of $L_{\max} = 6$ yields 28 coefficients). This process yields a set of SH coefficients, α_{lk} , which provide a compact representation of the fODF for that voxel. The input to tractography is a 4D tensor of size $H \times W \times D \times m$, where each spatial location (H, W, D) contains a vector of m SH coefficients (e.g., $m = 28$ for $L_{\max} = 6$).

Diffusion and Flow Matching Models. Modern generative models learn to reverse a continuous-time process mapping between data and noise. Diffusion Models (DMs) [12] learn to reverse a forward noising process (often an SDE) that transforms data x_0 into noise $\epsilon \sim \mathcal{N}(0, I)$. Generation involves solving the learned reverse-time SDE. Training is done by learning to predict the added noise ϵ from a corrupted sample $x_t = \alpha_t x_0 + \sigma_t \epsilon$ using a network $\epsilon_\theta(x_t, t)$:

$$\mathcal{L}_D(\theta) = \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon_\theta(x_t, t) - \epsilon\|^2]. \quad (2)$$

In contrast, Flow Matching (FM) models [19] directly learn the deterministic vector field \mathbf{v}_θ of an ordinary differential equation (ODE), $\frac{dx_t}{dt} = \mathbf{v}_\theta(x_t, t)$, which transports mass from a noise distribution x_0 to the data distribution x_1 . For a simple linear interpolant $x_t = (1 - t)x_0 + tx_1$, the target

field is $\mathbf{v} = x_1 - x_0$. The model is trained via a direct regression loss:

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t, x_0, x_1} [\|\mathbf{v}_\theta(x_t, t) - \mathbf{v}\|^2], \quad (3)$$

and inference is performed by solving the learned ODE, starting from x_0 .

4. Methodology: GenTract

We now introduce the architecture of GenTract, which comprises two main components: a global fODF encoder and a conditional Transformer trained using either a Diffusion or FM objective. These components, summarized in Figure 2, jointly address the challenges outlined in the introduction. The encoder learns a global context embedding, which is used as a condition by the generator to sample all streamline coordinates simultaneously. Each streamline is thus produced as a whole rather than stepwise.

4.1. Anatomical Conditioning Encoder

The model’s first task is to convert the high-dimensional fODF data into a compact and informative conditioning tensor \mathbf{z} . The encoder processes these volumes in a two-stage process:

1. *Per-Coefficient Representation Learning.* We train m independent VAEs, each modeling a distinct SH coefficient volume (Figure 2A). The i -th VAE includes an encoder $\mathcal{E}^{(i)}$ that maps the volume (H, W, D) to a compact latent space (C_z, H_z, W_z, D_z) and a decoder $\mathcal{D}^{(i)}$ that reconstructs it. Each VAE is trained using a composite loss function comprising reconstruction, perceptual, and adversarial terms, along with a Kullback-Leibler regularization term, following the formulation in [11]. This process yields a set of m latent embeddings, $z^{(0)}, \dots, z^{(m-1)}$, that compactly represent the full fODF information.
2. *Fusion and Refinement.* Each $z^{(i)}$ is further refined by a 3D ResNet-like class-conditioned encoder \mathcal{E}^c , producing $\hat{z}^{(i)} = \mathcal{E}^c(z^{(i)}, i)$ of shape (C_c, H_c, W_c, D_c) . The weights of \mathcal{E}^c are shared across all m coefficients, while the index i is provided as a conditioning signal to preserve coefficient-specific context. Additionally, \mathcal{E}^c further reduces the spatial dimensionality of the latent representations to lower the overall computational cost. Unlike the VAEs, \mathcal{E}^c is trained jointly with the downstream generative model, enabling it to adaptively extract the most relevant information for the tractography task. The spatial dimensions of each $\hat{z}^{(i)}$ are flattened, and all m latents are concatenated along the channel dimension, resulting in the global conditioning tensor \mathbf{z} of shape $(H_c \times W_c \times D_c, m \times C_c)$.

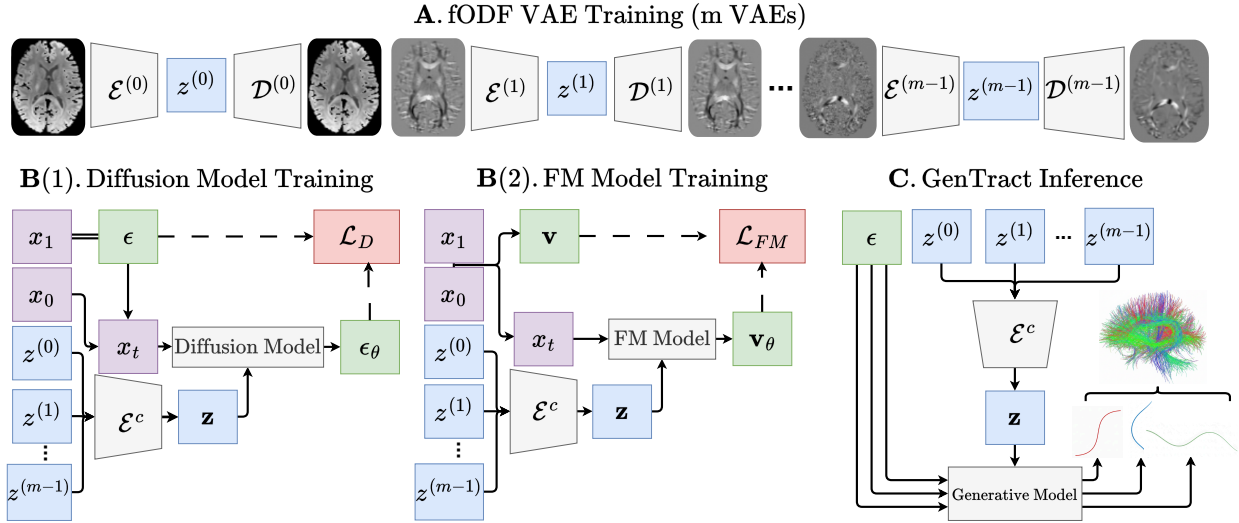


Figure 2. Overview of the GenTract framework. **A**: VAEs encode fODF coefficients into latent representations. **B(1)** and **B(2)**: The training protocol for Diffusion and FM models respectively, using the learned $z^{(i)}$ as input. Losses are back-propagated through both the generative model and the class-conditioned encoder \mathcal{E}^c . **C**: The inference process, where streamlines are generated by sampling from Gaussian noise conditioned on \mathbf{z} .

4.2. Conditional Transformer

The core generative model is a Transformer [39] designed to operate on batches of 3D coordinates representing streamlines. It is trained to learn the mapping from noise to a clean streamline, conditioned on the anatomical context tensor \mathbf{z} .

Architecture. The model receives three inputs: (i) a noisy streamline x_t of shape $(p, 3)$, where p is the number of sampled points on the streamline; (ii) the timestep t ; and (iii) the context tensor \mathbf{z} . To enable cross-attention between \mathbf{z} and the streamline representation, these inputs are linearly projected into a shared embedding space of dimension n . Following [25], we add a sinusoidal positional encoding and a learned temporal embedding to the streamline representation. The main body of the model is a stack of M Transformer layers containing two key attention mechanisms:

- *Self-Attention*: Models the geometric relationships and dependencies between points along the streamline itself, ensuring internal coherence.
- *Cross-Attention*: Integrates guidance from the projected n -dimensional anatomical context \mathbf{z} , allowing the global white matter structure to inform the streamline’s generation.

The final output from the Transformer stack is projected back to the original 3D coordinate space to compute the training loss.

Training Objectives. The Transformer backbone is trained to learn the dynamics of the continuous time-process from Gaussian noise to target streamlines, conditioned on

the learned \mathbf{z} , using either the Diffusion objective, \mathcal{L}_D , as described in (2) (Figure 2 B(1)), or the FM objective \mathcal{L}_{FM} , as in (3) (Figure 2 B(2)). In both cases, the loss is back-propagated through the Transformer network, as well as the class-conditioned encoder, \mathcal{E}^c . Given a global conditioning tensor \mathbf{z} , the training objective allows the model to learn the target distribution over streamlines.

4.3. Full Tractogram Inference

At inference for a new subject (Figure 2 (C)), we first compute the conditioning tensor \mathbf{z} from the input SH coefficients. The model then generates streamlines by applying the learned reverse process, starting from random noise and conditioned on the subject-specific \mathbf{z} . Repeating this process in batches allows us to sample from the full conditional distribution of streamlines learned by the model for a specific subject. An example of a generated tractogram from GenTract can be seen in Figure 1, along with further qualitative results in the Supplementary Material.

4.4. Implementation Details

We detail all the data dimensions and hyperparameters in the Supplementary Material. The VAEs were obtained by fine-tuning separate instances of the MAISI VAE [11]. The Transformer blocks, as well as training scripts and data processing, for both the Diffusion and FM models were implemented via MONAI [26]. All experiments are performed on an NVIDIA H100 GPU [1]. All training and inference code can be found at <https://github.com/alecsargood/GenTract>.

5. Experiments

This Section evaluates the GenTract framework. We first introduce the datasets and evaluation protocols. We then conduct an internal model comparison to determine GenTract’s optimal configuration, examining the choice of generative framework, model size, and number of inference steps. Next, we provide a comparative benchmark of this optimal model against SOTA methods on high-quality data. Finally, we perform a robustness analysis, assessing all methods under noisy and low-resolution conditions using both internal and external test sets. We provide details on all statistical tests performed in the experiments in the Supplementary Material.

5.1. Datasets and Preprocessing

Training and Test Datasets. We use dMRI data from 1,042 subjects (age 22-35, 45.6% female) from the publicly available Human Connectome Project (HCP) Young Adult dataset [38]. For the supervised training target, we use the fODF and tractography data processed by [16], which employs the PyAFQ tractography pipeline. PyAFQ first applies constrained spherical deconvolution to the raw dMRI data to estimate the fODF by computing its SH coefficients, and then generates a tractogram using probabilistic tractography [3]. It then registers the subject’s data to a template and uses its own filtering pipeline to filter streamlines based on geometric plausibility and alignment with a predefined white-matter atlas containing 24 known bundles (collection of streamlines that form neuroanatomical pathways connecting two regions of the brain) [16]. Each subject is thus represented by a pair: a 4D SH coefficient image and its corresponding filtered tractogram. We use a strict subject-level split for training (75%), validation (10%), and test (15%) sets. The training set is augmented using deterministic rotations ($\pm 15^\circ$, $\pm 30^\circ$, $\pm 45^\circ$) of the tractograms and their corresponding SH volumes. We use the HCP test set to establish baseline performance and then evaluate model robustness under several challenging conditions. First, we assess performance on noisy data by adding Rician noise ($\sigma = 0.005$) to the input fODF signal of the HCP test set. Next, to simulate low-field data, we synthetically downsample the HCP test set to 3 mm^3 isotropic resolution and similarly add Rician noise ($\sigma = 0.005$), with resolution parameters aligned to previous studies of tractography on low-field MR scanners [10]. Finally, we assess GenTract’s out-of-domain performance in low-resolution (LR) and noisy conditions using the 28-subject test set from TractoInferno [27], a highly heterogeneous dataset compiled from multiple studies (ages 18–62, with varying acquisition protocols) and extensively preprocessed to ensure research-grade quality. Specifically, we synthetically downsample this external dataset and introduce noise in the same manner as the HCP test set, enabling evaluation of robustness

under LR, noisy, and cross-dataset conditions.

Data Preprocessing. To reduce computational load for training of the model, all SH coefficient volumes are downsampled to 1.875 mm^3 isotropic resolution. Each of the SH coefficient volumes are z-score normalized, and all streamline coordinates are min-max scaled to $[-1, 1]$, using statistics computed from the training set.

5.2. Evaluation

Evaluation Metrics. Evaluating tractograms is challenging as there is no absolute biological ‘ground truth’, thus we must rely on proxies to assess model performance. As done in previous studies [17, 34], we define a streamline as a True Positive (TP) if it is retained by a given filtering tool (the proxy for truth), and a False Positive (FP) if it is discarded. Given this, we evaluate all methods on three criteria:

1. *Precision (Percentage of Retained Streamlines):* This is defined as $\frac{TP}{TP+FP}$. A high-precision score indicates that a large proportion of the generated streamlines are valid (TP), and that the model minimizes the creation of false positive streamlines (FP).
2. *Number of Discovered Bundles:* This metric assesses the model’s ability to reconstruct a predefined set of bundles, using a white matter atlas as a reference. We count how many of these canonical bundles are successfully populated by plausible streamlines. This metric also serves as a proxy for streamline recall, where a missed bundle (and thus missed streamlines) represent false negatives (FNs).
3. *Inference Time:* We report the mean computation time required to generate a full tractogram for a test subject.

Internal Model Comparison Tools. For the internal model comparisons, we evaluate tractograms using the PyAFQ filtering mechanism [16]. This is the same filtering mechanism used to filter the supervised training data, and we use this tool to calculate both precision (AFQ % P) and number of recovered bundles (AFQ Bundles, up to 24). This provides a consistent metric to evaluate performance on the internal test set to determine the optimal configuration for GenTract.

SOTA Comparison Evaluation Tools. To avoid any potential bias from using PyAFQ (which was used in training), the main SOTA comparison relies on two independent and established tools. The first, BundleSeg (BS) [33], provides a measure of precision (BS % P) and number of recovered bundles (BS Bundles). BundleSeg operates by registering a subject’s streamlines into a standard template space. There, streamlines are segmented against a high-resolution, population-based atlas (51 bundles) that serves

as a comprehensive reference for ‘valid’ pathways. This subject-specific registration penalizes any streamlines not well-aligned with the subject’s unique anatomy. The second, TractOracle-Net (TO-Net) [35], serves as an independent measure of precision (TO-Net % P). TO-Net is a pre-trained transformer that classifies streamlines as plausible or implausible based on their geometric properties. Its diverse training provides an assessment of the model’s output, independent of a single atlas.

5.3. Internal Model Comparison Study

Choice of Architecture and Generative Framework. We perform a comparison study by: (1) training GenTract using either the Diffusion or FM frameworks, (2) varying the number of Transformer layers ($M \in \{4, 6, 8\}$), and (3) evaluating different model embedding dimensions for the best-performing configuration. We do not consider transformer layers beyond $M = 8$ because of computational expense. Inference for all architectures is conducted using 50 steps. The quantitative results are summarized in Table 1. All model configurations successfully identify all 24 AFQ bundles. The primary differences appear in the AFQ % P, where performance improves consistently with increasing numbers of layers for both architectures. Moreover, Diffusion models outperform FM models at each corresponding layer depth. The 8-layer Diffusion model ($M = 8$) achieves the highest precision at $85.0 \pm 1.5\%$, compared to $82.45 \pm 1.67\%$ for its FM counterpart. Finally, Table 1 identifies $n = 256$ as the optimal embedding dimension for a fixed architecture, as smaller dimensions ($n = 128$) underperform and larger dimensions ($n = 512$) cause overfitting. We therefore select the Diffusion model framework, with $M = 8$ and $n = 256$, as the optimal GenTract architecture.

Table 1. Performance summary of tractography generation models. The table shows the number of bundles found and the percentage of streamlines retained. Values are mean \pm std. Best performing results in each column in **bold** and second best underlined.

Model Type	M	AFQ Bundles (/ 24)	AFQ % P
Diffusion	4	24.0 \pm 0.0	81.8 \pm 1.5
	6	24.0 \pm 0.0	<u>83.4 \pm 1.7</u>
	8	24.0 \pm 0.0	85.0 \pm 1.5
Flow Matching	4	24.0 \pm 0.0	75.50 \pm 2.14
	6	24.0 \pm 0.0	80.56 \pm 1.78
	8	24.0 \pm 0.0	82.45 \pm 1.67
Diffusion	$n = 128$	24.0 \pm 0.0	80.6 \pm 1.1
	$n = 256$	24.0 \pm 0.0	85.0 \pm 1.5
	$n = 512$	24.0 \pm 0.0	79.3 \pm 1.5

Number of Inference Steps. Figure 3 illustrates the trade-off between precision (AFQ % P) and computational time across different numbers of Denoising Diffusion Implicit Models (DDIM) inference steps. We observe a substantial performance jump from 5 steps (69.7% AFQ P) to 10 steps (82.9% AFQ P), but we note that gains diminish thereafter. The 25-step (85.4%) and 50-step (85.2%) configurations offered only small precision improvements while more than doubling the inference time of the 10-step model (from 391s to 1976s). All variations recovered 24 bundles. We therefore select 10 DDIM inference steps for all subsequent experiments.

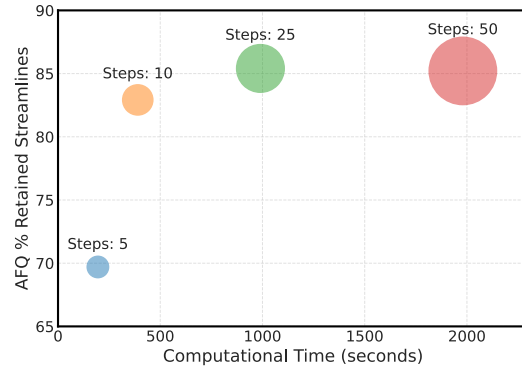


Figure 3. AFQ streamlines retained (AFQ % P) vs. Computational Time by Inference Step for the Diffusion ($M = 8$, $n = 256$) model. The size of each bubble is proportional to the number of DDIM inference steps used.

5.4. Comparison with State-of-the-Art

Baselines. We benchmark the optimal GenTract configuration against classical local tractography methods (iFOD2 [36] and SD Stream [5]), deep learning-based local methods (TractOracle [35] and DDTracking [18]), and a standard global method (tckglobal [2]). The iFOD2, SD Stream, and tckglobal baselines are implemented using MRtrix3 [37] and tractograms generated by iFOD2 and SD Stream are further processed with SIFT filtering [32], following standard practice. For TractOracle and DDTracking, we used the pretrained models provided in their open-source implementations, as this empirically yielded better performance than retraining. Further details of all inference pipelines settings are in the Supplementary Material.

Quantitative Comparison: Base Performance. Table 2 presents the quantitative comparison on the HCP test set, which highlights a clear trade-off between streamline precision (minimizing FPs) and bundle count (minimizing FNs). GenTract achieves a state-of-the-art precision: we observe that on the BS % P metric, GenTract (61.95%) outperforms the next-best deep learning methods, DDTracking

(35.20%) and TractOracle (28.93%), by a factor of 1.8 and 2.1, respectively. This advantage is consistent on the TO-Net % P metric, where GenTract (56.35%) again surpasses DDTracking (30.70%) and TractOracle (39.55%). This gain in precision comes at the cost of a higher false negative count, as the other methods identify more bundles than GenTract (36.62). While GenTract has a higher FN rate than the baselines, we hypothesize that this is a consequence of training on a constrained 24-bundle (PyAFQ) distribution.

Table 2. Comparison of tractography models on HCP test set. Best performance is **bolded** and second-best is underlined (mean \pm std).

Model	BS % P	BS Bundles (/51)	TO-Net % P
tckglobal	0.19 \pm 0.05	42.91 \pm 2.75	17.83 \pm 1.03
iFOD2	1.96 \pm 0.31	<u>48.88 \pm 1.49</u>	6.30 \pm 0.60
SD Stream	4.71 \pm 0.67	47.85 \pm 1.63	11.10 \pm 0.80
DDTracking	<u>35.20 \pm 3.49</u>	49.69 \pm 0.64	30.70 \pm 6.00
TractOracle	28.93 \pm 3.28	48.20 \pm 0.72	<u>39.55 \pm 1.99</u>
GenTract	61.95 \pm 4.50	36.62 \pm 1.21	56.35 \pm 4.25

Quantitative Comparison: Robustness to Rician Noise.

Table 3 summarizes the quantitative comparison under Rician noise, as described in Section 5.1. We observe that iFOD2, SD Stream, and DDTracking are severely impacted by noise, with precision rates falling substantially and the number of bundles recovered dropping. TractOracle exhibits moderate robustness, with its BS % P decreasing by 23.8%, though it retains the highest bundle count. Tckglobal remains relatively stable under noise, still reconstructing many bundles, but remains uncompetitive in terms of precision. In contrast, GenTract shows strong stability across all metrics, with only a 2.6% decrease in BS % P, 3.8% in BS Bundles, and 3.3% in TO-Net % P, and although it has a higher FN rate than tckglobal and TractOracle, it maintains the highest precision among all baseline methods. These findings indicate that GenTract preserves high streamline precision under noisy conditions, showing its potential as a promising choice for challenging diffusion MRI data present in clinical scenarios.

Table 3. Comparison of tractography models under Rician noise. Best performance is **bolded** and second-best is underlined (mean \pm std).

Model	BS % P	BS Bundles (/51)	TO-Net % P
tckglobal	0.20 \pm 0.05	43.11 \pm 2.84	17.85 \pm 1.05
iFOD2	0.08 \pm 0.04	13.17 \pm 2.22	5.03 \pm 0.73
SD Stream	0.11 \pm 0.04	11.40 \pm 2.60	9.53 \pm 0.92
DDTracking	5.41 \pm 1.84	39.43 \pm 2.60	12.24 \pm 1.95
TractOracle	<u>22.06 \pm 3.08</u>	48.29 \pm 0.59	34.90 \pm 2.35
GenTract	60.32 \pm 4.12	35.20 \pm 1.43	54.52 \pm 4.02

Quantitative Comparison: Robustness to Low-Resolution and Noise.

Table 4 summarizes the results on the low-resolution and noisy HCP dataset, as described in Section 5.1. While an absolute drop is expected for all methods on degraded data, we observe that GenTract retains a higher level of relative performance, while baseline methods degrade more severely. Local methods (iFOD2, SD Stream) struggle in this setting, with BS % P dropping to 0%. This behavior is consistent with their autoregressive design, which is highly sensitive to ambiguous local information and prone to error propagation. The global method (tckglobal) also fails completely. This is likely due to its classical optimization which forms incorrect streamline connections under low-resolution ambiguity, leading to a collapse in precision. Although TractOracle continues to produce the highest number of bundles, its BS % P rate decreases to nearly 1%, and TO-Net % P to less than 6%. DDTracking similarly has low-precision scores. In contrast, GenTract still achieves a BS % P of 15.73%, an order of magnitude higher than the next-best result, and a TO-Net %-P of 43.43%, while also recovering the second-highest bundle count. These results suggest that GenTract’s performance advantage is maintained on this low-resolution and noisy data.

Table 4. Comparison of tractography models under low-resolution and noise. Best performance is **bolded** and second-best is underlined (mean \pm std).

Model	BS % P	BS Bundles (/51)	TO-Net % P
tckglobal	0.00 \pm 0.00	0.00 \pm 0.00	3.62 \pm 1.30
iFOD2	0.00 \pm 0.00	0.18 \pm 0.39	0.55 \pm 0.34
SD Stream	0.00 \pm 0.00	0.00 \pm 0.00	3.40 \pm 2.51
DDTracking	<u>4.44 \pm 1.45</u>	31.71 \pm 3.03	<u>8.42 \pm 0.60</u>
TractOracle	1.12 \pm 0.30	37.79 \pm 2.45	5.41 \pm 0.73
GenTract	15.73 \pm 2.82	<u>31.74 \pm 1.26</u>	43.43 \pm 3.68

Quantitative Comparison: External Robustness to Low-Resolution and Noise.

Table 5 presents the external robustness evaluation, assessing GenTract’s generalization on the synthetically-degraded TractoInferno dataset. We observe that most models perform better on this dataset than on the degraded HCP data, which we hypothesize is due to the better underlying image quality of the TractoInferno dataset as a result of its extensive preprocessing. Despite this general uplift, the performance hierarchy remains clear. Classical and global methods (iFOD2, SD Stream, tckglobal) still fail completely. DDTracking shows only modest TO-Net % P (9.79%). TractOracle, while recovering the most bundles (42.18), achieves a limited precision of 9.74% BS % P. In contrast, GenTract achieves the highest precision by a wide margin (24.94% BS % P and 40.88% TO-Net % P) and the second-highest bundle count (35.05). These re-

sults suggest that GenTract retains its performance advantage on this low-quality, out-of-domain data, indicating its design generalizes beyond the training domain.

Table 5. Performance on low-resolution and noisy TractoInferno dataset. Best performance is **bolded** and second-best is underlined (mean \pm std).

Model	BS % P	BS Bundles (/51)	TO-Net % P
tckglobal	0.00 \pm 0.00	0.00 \pm 0.00	2.24 \pm 0.68
iFOD2	0.01 \pm 0.02	0.36 \pm 0.49	1.34 \pm 0.89
SD Stream	0.01 \pm 0.04	0.18 \pm 0.39	4.24 \pm 2.75
DDTracking	6.35 \pm 7.73	17.96 \pm 9.89	9.79 \pm 1.44
TractOracle	<u>9.74 \pm 3.53</u>	42.18 \pm 3.38	<u>10.34 \pm 1.34</u>
GenTract	24.94 \pm 8.09	<u>35.04 \pm 1.54</u>	40.88 \pm 5.01

Qualitative Comparison. Figure 4 illustrates the Right Superior Longitudinal Fasciculus (SLFR) bundle for a single HCP test subject (segmented by BundleSeg), for all methods under different data conditions. GenTract has a higher streamline retention compared with baseline methods under the Noise and LR + Noise settings. This supports the quantitative performance differences we report in the previous tables. SD Stream, iFOD2, and tckglobal fail to produce any valid streamlines for this bundle in LR + Noise setting. This aligns with their low bundle counts and precision scores in the quantitative analysis. Further qualitative results can be found in the Supplementary Material.

Inference Times. Figure 5 shows that GenTract is 2-3 \times faster than classical methods and achieves inference times comparable to other deep learning methods.

6. Discussion

In this study, we present GenTract, a new paradigm that frames global tractography as a conditional generative task. GenTract addresses limitations of prior methods: it avoids step-wise error propagation of local approaches by generating all streamline coordinates simultaneously. Moreover, this generative formulation removes the need for initial seeding points, a source of operator-dependent variability in local tracking. We hypothesize that leveraging global conditioning from the entire dMRI contributes to its robustness against local noise and data perturbations. These properties show GenTract’s potential for clinical applications, where scans are often time-constrained and of lower quality, supporting the practical adoption of global tractography.

Future Work. Our study highlights important directions for future research. A primary challenge is the absence of ground truth in tractography, which gives rise to two related issues. First, there is potentially suboptimal supervision for

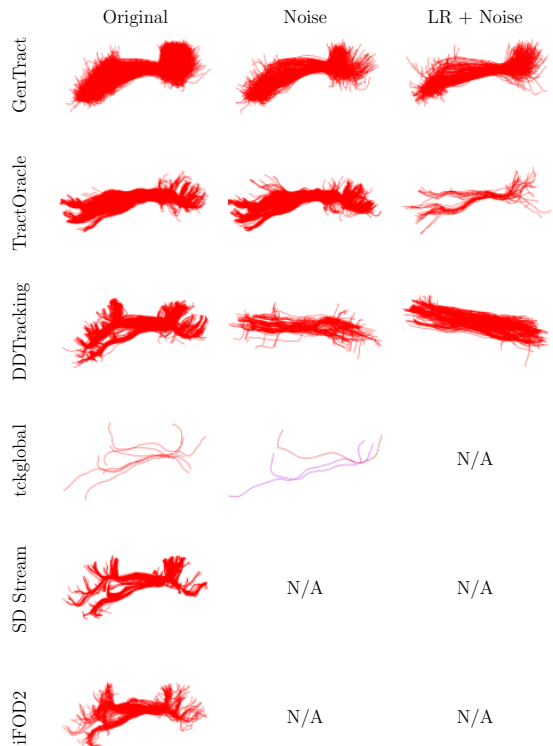


Figure 4. Qualitative result showing SLFR segmented by BundleSeg for each method, on original, noised, and LR+Noise for a HCP test subject.

GenTract. The results suggest that the specific choice of the PyAFQ filtered dataset leads to learning a constrained bundle distribution. While we train a high-precision model, it comes at a cost of a higher false negative rate. Future work should explore strategies to mitigate this trade-off, perhaps by expanding the distribution of streamlines present in the training set, or by developing global training approaches that do not require supervised data. Second, since no biological ground truth exists for tractography, we must rely on proxy evaluation tools that may have inaccuracies or introduce their own biases. A further limitation is the reliance on synthetic degradation for the robustness analysis. While the results on this data suggests that GenTract is more robust to noisy and low-resolution settings, a next step is to validate this performance on real-world clinical scans. Assessing its practical utility will require evaluating performance on non-synthetically degraded clinical and low-resolution data. Additionally, as we transition toward clinical viability, future work can exploit the probabilistic sampling nature of GenTract to quantify uncertainty in the generated tractograms.

Acknowledgements

AS is supported by the EPSRC-funded UCL Centre for Doctoral Training i4health (EP/S021930/1) and by the Wellcome Trust CU-MONDAI project (WT 221915). LP is supported by the PNRR initiative (DM 118/2023). ET is supported by the Wellcome Trust CU-MONDAI project (WT 221915). DCA acknowledges funding from the E-DADS project (EU-JPND 2019; UKRI MR/T046422/1), the Wellcome Trust CU-MONDAI project (WT 221915), the Wellcome Trust Democratizing MRI project (WT 317797), and the NIHR UCLH Biomedical Research Centre.

References

- [1] Jack Choquette. Nvidia hopper h100 gpu: Scaling performance. *IEEE Micro*, 43(3), 2023. 4
- [2] Daan Christiaens, Marco Reisert, Thijs Dhollander, Stefan Sunaert, Paul Suetens, and Frederik Maes. Global tractography of multi-shell diffusion-weighted imaging data using a multi-tissue model. *NeuroImage*, 123, 2015. 2, 3, 6
- [3] Dipy Contributors, Eleftherios Garyfallidis, Matthew Brett, Bagrat Bagrat Amirbekian, Ariel Rokem, Stefan van der Walt, Maxime Descoteaux, and Ian Nimmo-Smith. Dipy, a library for the analysis of diffusion MRI data. *Frontiers in Neuroinformatics*, 8, 2014. 5
- [4] Alessandro Daducci, Simona Schiavi, Daan Christiaens, Robert Smith, and Daniel C Alexander. Global tractography. In *Handbook of Diffusion MR Tractography*. Elsevier, 2025. 2
- [5] Maxime Descoteaux, Rachid Deriche, Thomas R Knosche, and Alfred Anwander. Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE Transactions on Medical Imaging*, 28(2), 2008. 1, 2, 6
- [6] Jessica Dubois, Ghislaine Dehaene-Lambertz, Sofya Kulikova, Cyril Poupon, Petra S Hüppi, and Lucie Hertz-Pannier. The early development of brain white matter: a review of imaging studies in fetuses, newborns and infants. *Neuroscience*, 276, 2014. 1
- [7] R Douglas Fields. White matter in learning, cognition and psychiatric disorders. *Trends in Neurosciences*, 31(7), 2008. 1
- [8] Pierre Fillard, Cyril Poupon, and Jean-François Mangin. A novel global tractography algorithm based on an adaptive spin glass model. In *Proceedings of the 12th International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2009. 2
- [9] Christopher M Filley and R Douglas Fields. White matter and cognition: making the connection. *Journal of Neurophysiology*, 116(5), 2016. 1
- [10] James Gholam, Phil Schmid, Joshua Ametepe, Alix Plumley, Leandro Beltrachini, Francesco Padormo, Rui Teixeira, Rafael OHalloran, Kaloian Petkov, Klaus Engel, et al. Diffusion Tensor MRI and Spherical-Deconvolution-Based Tractography on an Ultra-Low Field Portable MRI System. *arXiv preprint arXiv:2506.04473*, 2025. 5
- [11] Pengfei Guo, Can Zhao, Dong Yang, Ziyue Xu, Vishwesh Nath, Yucheng Tang, Benjamin Simon, Mason Belue, Stephanie Harmon, Baris Turkbey, et al. Maisi: Medical ai for synthetic imaging. In *Proceedings of the 2025 IEEE/CVF Winter Conference on Applications of Computer Vision*. IEEE, 2025. 3, 4
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [13] Saad Jbabdi, Mark W Woolrich, Jesper LR Andersson, and TEJ Behrens. A bayesian framework for global tractography. *NeuroImage*, 37(1), 2007. 2
- [14] Ben Jeurissen, Maxime Descoteaux, Susumu Mori, and Alexander Leemans. Diffusion MRI Fiber Tractography of the Brain. *NMR in Biomedicine*, 32(4), 2019. 2
- [15] Koji Kamagata, Christina Andica, Wataru Uchida, Kaito Takabayashi, Yuya Saito, Matthew Lukies, Akifumi Hagiwara, Shohei Fujita, Toshiaki Akashi, Akihiko Wada, et al. Advancements in diffusion MRI tractography for neurosurgery. *Investigative Radiology*, 59(1), 2024. 1
- [16] John Kruper, Jason D Yeatman, Adam Richie-Halford, David Bloom, Mareike Grotheer, Sendy Caffarra, Gregory Kiar, Iliana I Karipidis, Ethan Roy, Bramsh Q Chandio, et al. Evaluating the reliability of human brain white matter tractometry. *Aperture Neuro*, 1(1), 2021. 2, 5
- [17] Jon Haitz Legarreta, Laurent Petit, François Rheault, Guillaume Theaud, Carl Lemaire, Maxime Descoteaux, and Pierre-Marc Jodoin. Filtering in tractography using autoencoders (FINTA). *Medical Image Analysis*, 72, 2021. 5
- [18] Yijie Li, Wei Zhang, Xi Zhu, Ye Wu, Yogesh Rath, Lauren J O'Donnell, and Fan Zhang. DDTracking: A Deep Generative Framework for Diffusion MRI Tractography with Streamline Local-Global Spatiotemporal Modeling. *arXiv preprint arXiv:2508.04568*, 2025. 1, 2, 6
- [19] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow Matching for Generative Modeling. *arXiv preprint arXiv:2210.02747*, 2022. 3
- [20] Huan Liu, Yuanyuan Yang, Yuguo Xia, Wen Zhu, Rehana K Leak, Zhishuo Wei, Jianyi Wang, and Xiaoming Hu. Aging of cerebral white matter. *Ageing Research Reviews*, 34, 2017. 1
- [21] Klaus H Maier-Hein, Peter F Neher, Jean-Christophe Houde, Marc-Alexandre Côté, Eleftherios Garyfallidis, Jidan Zhong, Maxime Chamberland, Fang-Cheng Yeh, Ying-Chia Lin, Qing Ji, et al. The challenge of mapping the human connectome based on diffusion tractography. *Nature Communications*, 8(1), 2017. 1, 2
- [22] Jean-François Mangin, Cyril Poupon, Yann Cointepas, Denis Rivière, Dimitri Papadopoulos-Orfanos, Chris A Clark, Jean Régis, and Denis Lrke Bihan. A framework based on spin glass models for the inference of anatomical connectivity from diffusion-weighted MR data—a technical review. *NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In Vivo*, 15(7-8), 2002. 2
- [23] Jean-François Mangin, Pierre Fillard, Yann Cointepas, Denis Le Bihan, Vincent Frouin, and Cyril Poupon. Toward global tractography. *NeuroImage*, 80, 2013. 1

- [24] Peter F Neher, Marc-Alexandre Côté, Jean-Christophe Houde, Maxime Descoteaux, and Klaus H Maier-Hein. Fiber tractography using machine learning. *NeuroImage*, 158, 2017. 1, 2
- [25] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 4
- [26] Walter HL Pinaya, Mark S Graham, Eric Kerfoot, Petru-Daniel Tudosiu, Jessica Dafflon, Virginia Fernandez, Pedro Sanchez, Julia Wolleb, Pedro F Da Costa, Ashay Patel, et al. Generative AI for medical imaging: extending the MONAI framework. *arXiv preprint arXiv:2307.15208*, 2023. 4
- [27] Philippe Poulin, Guillaume Theaud, Francois Rheault, Etienne St-Onge, Arnaud Bore, Emmanuelle Renauld, Louis de Beaumont, Samuel Guay, Pierre-Marc Jodoin, and Maxime Descoteaux. TractoInferno-A large-scale, open-source, multi-site database for machine learning dMRI tractography. *Scientific Data*, 9(1), 2022. 5
- [28] Marco Reisert, Irina Mader, Constantin Anastasopoulos, Matthias Weigel, Susanne Schnell, and Valerij Kiselev. Global fiber reconstruction becomes practical. *NeuroImage*, 54(2), 2011. 2
- [29] Francois Rheault, Philippe Poulin, Alex Valcourt Caron, Etienne St-Onge, Kurt G Schilling, Laurent Petit, Flavio Dell'Acqua, Alexander Leemans, and Maxime Descoteaux. Current challenges and opportunities for tractography. *Handbook of Diffusion MR Tractography*, 2025. 2, 3
- [30] Kurt G Schilling, Chantal MW Tax, Francois Rheault, Bennett A Landman, Adam W Anderson, Maxime Descoteaux, and Laurent Petit. Prevalence of white matter pathways coming into a single white matter voxel orientation: The bottleneck issue in tractography. *Human Brain Mapping*, 43(4), 2022. 1
- [31] Robert E Smith, Jacques-Donald Tournier, Fernando Calamante, and Alan Connelly. Anatomically-constrained tractography: improved diffusion mri streamlines tractography through effective use of anatomical information. *NeuroImage*, 62(3), 2012. 1
- [32] Robert E Smith, Jacques-Donald Tournier, Fernando Calamante, and Alan Connelly. Sift: Spherical-deconvolution informed filtering of tractograms. *NeuroImage*, 67, 2013. 6
- [33] Etienne St-Onge, Kurt G Schilling, and Francois Rheault. Bundleseg: A versatile, reliable and reproducible approach to white matter bundle segmentation. In *International Workshop on Computational Diffusion MRI*. Springer, 2023. 5
- [34] Antoine Théberge, Christian Desrosiers, Maxime Descoteaux, and Pierre-Marc Jodoin. Track-to-learn: A general framework for tractography with deep reinforcement learning. *Medical Image Analysis*, 72, 2021. 2, 5
- [35] Antoine Théberge, Maxime Descoteaux, and Pierre-Marc Jodoin. TractOracle: towards an anatomically-informed reward function for RL-based tractography. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2024. 1, 2, 6
- [36] J Donald Tournier, Fernando Calamante, Alan Connelly, et al. Improved probabilistic streamlines tractography by 2nd order integration over fibre orientation distributions. In *Proceedings of the International Society for Magnetic Resonance in Medicine*. Stockholm, 2010. 1, 2, 6
- [37] J Donald Tournier, Robert Smith, David Raffelt, Rami Tabbara, Thijs Dhollander, Maximilian Pietsch, Daan Christiaens, Ben Jeurissen, Chun-Hung Yeh, and Alan Connelly. Mrtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *NeuroImage*, 202, 2019. 6, 1
- [38] David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, Wu-Minn HCP Consortium, et al. The WU-Minn Human Connectome Project: an Overview. *NeuroImage*, 80, 2013. 5
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017. 4
- [40] Fan Zhang, Antoine Théberge, Pierre-Marc Jodoin, Maxime Descoteaux, and Lauren J O'Donnell. Think deep in the tractography game: deep learning for tractography computing and analysis. *Brain Structure and Function*, 230(6), 2025. 2