Novelty Assessment Report

Paper: CardioComposer: Leveraging Differentiable Geometry for Compositional Control of Anatomical Diffusion Models

PDF URL: https://openreview.net/pdf?id=JyboUMeEUi

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-30

Abstract

Generative models of 3D cardiovascular anatomy can synthesize informative structures for clinical research and medical device evaluation, but face a trade-off between geometric controllability and realism. We propose CardioComposer: a programmable, inference-time framework for generating multi-class anatomical label maps based on interpretable ellipsoidal primitives. These primitives represent geometric attributes such as the size, shape, and position of discrete substructures. We specifically develop differentiable measurement functions based on voxel-wise geometric moments, enabling loss-based gradient guidance during diffusion model sampling. We demonstrate that these losses can constrain individual geometric attributes in a disentangled manner and provide compositional control over multiple substructures. Finally, we show that our method is compatible with a wide array of anatomical systems containing non-convex substructures, spanning cardiac, vascular, and skeletal organs.

Core Task Landscape

This paper addresses: **Compositional Geometric Control of Anatomical Diffusion Models**

A total of **13 papers** were analyzed and organized into a taxonomy with **10 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Geometric Constraint Mechanisms for Anatomical Generation**
- **Segmentation-Guided Anatomical Synthesis**
- **Multimodal Conditioning Architectures**
- **Semantic Augmentation and Domain-Agnostic Generation**
- **Generative Diffeomorphic Mapping for Neuroanatomical Registration**

Complete Taxonomy Tree

- Compositional Geometric Control of Anatomical Diffusion Models Survey Taxonomy
- Geometric Constraint Mechanisms for Anatomical Generation
  - Primitive-Based Compositional Control ★ (2 papers)
  - [0] CardioComposer: Leveraging Differentiable Geometry for Compositional Control of Anatomical Diffusion Models (Anon et al., 2026) View paper
  - [8] CardioComposer: Flexible and Compositional Anatomical Structure Generation with Disentangled Geometric Guidance (Kadry, 2025) View paper
  - Topological and Morphological Property Enforcement (1 papers)
  - [5] Anatomica: Localized Control over Geometric and Topological Properties for Anatomical Diffusion Models (Karim Kadry, 2025) View paper
  - Landmark and Skeletal Structure Guidance (1 papers)
  - [6] A Diffusion Model for Simulation Ready Coronary Anatomy with Morpho-Skeletal Control (Karim Kadry, 2024) View paper
- Segmentation-Guided Anatomical Synthesis
  - Medical Image Synthesis with Anatomical Masks (2 papers)
  - [3] Heartbeat: Towards controllable echocardiography video synthesis with multimodal conditions-guided diffusion models (Xinrui Zhou, 2024) View paper
  - [11] Anatomically-Controllable Medical Image Generation with Segmentation-Guided Diffusion Models (Konz, 2024) View paper
  - Surgical and Interventional Scene Generation (1 papers)
  - [1] Data augmentation for surgical scene segmentation with anatomy-aware diffusion models (Danush Kumar Venkatesh, 2025) View paper
  - Microstructure and Material Phase Segmentation Synthesis (1 papers)
  - [13] Phase-Fraction-Guided Denoising Diffusion Model for Multiphase Steel Microstructure Segmentation Augmentation Via Micrograph Image-Mask Pair Synthesis (Hoang Hai Nam Nguyen, n.d.) View paper
- Multimodal Conditioning Architectures
  - Text and Appearance-Guided Human Synthesis (2 papers)
  - [9] 3DMM-GAN: Multi-Modal Alignment With Adversarial Learning for Compositional 3D Human Image Synthesis (Tianyi Chen, 2025) View paper
  - [12] HairFree: Compositional 2D Head Prior for Text-Driven 360Â° Bald Texture Synthesis (M Ostrek, n.d.) View paper
  - Landmark-Tokenized Facial Editing (1 papers)
  - [10] LaTo: Landmark-tokenized Diffusion Transformer for Fine-grained Human Face Editing (Zhang Zheng-hao, 2025) View paper
- Semantic Augmentation and Domain-Agnostic Generation (2 papers)
  - [4] Semantic data augmentation with generative models (C Shivashankar, 2023) View paper

  - [7] Robust Polyp Detection and Diagnosis Through Compositional Prompt-Guided Diffusion Models. (Yu JIA, 2025) View paper
- Generative Diffeomorphic Mapping for Neuroanatomical Registration (1 papers)
  - [2] Solving the where problem and quantifying geometric variation in neuroanatomy using generative diffeomorphic mapping (Daniel J. Tward, 2025) View paper

Narrative

Core task: Compositional geometric control of anatomical diffusion models. The field centers on generating anatomically plausible medical images through diffusion models that respect geometric and structural constraints. The taxonomy reveals several complementary directions: geometric constraint mechanisms that enforce spatial and shape priors during generation, segmentation-guided approaches that leverage anatomical masks to steer synthesis, multimodal conditioning architectures that integrate diverse input modalities (text, sketches, landmarks), semantic augmentation methods that produce domain-agnostic or label-preserving synthetic data, and specialized diffeomorphic mapping techniques for neuroanatomical registration. Representative works span cardiac imaging (Heartbeat[3], Coronary Anatomy Diffusion[6]), broader anatomical synthesis (Anatomica[5]), and surgical or polyp detection scenarios (Surgical Scene Augmentation[1], Polyp Detection Diffusion[7]), illustrating how different branches address organ-specific versus general anatomical generation challenges.

A particularly active line of work explores primitive-based compositional control, where generation is decomposed into interpretable geometric elements such as landmarks, contours, or parametric shape descriptors. CardioComposer[0] exemplifies this approach by enabling fine-grained compositional manipulation of cardiac structures through geometric primitives, closely related to CardioComposer Flexible[8] which extends similar compositional strategies. This contrasts with segmentation-guided methods like Anatomically-Controllable Generation[11] that rely on dense mask conditioning, and with semantic augmentation frameworks such as Semantic Data Augmentation[4] that prioritize label consistency over explicit geometric control. Meanwhile, diffeomorphic techniques like Diffeomorphic Neuroanatomy[2] focus on smooth, topology-preserving transformations rather than direct synthesis. The original paper sits within the primitive-based cluster, emphasizing interpretable geometric handles for compositional control, distinguishing itself from mask-driven approaches by offering more modular, editable generation pathways that align with clinical workflows requiring precise anatomical adjustments.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. CardioComposer: Flexible and Compositional Anatomical Structure Generation with Disentangled Geometric Guidance

Authors: Kadry, Karim, Goraya, Shoaib, Karim Kadry, et al. (15 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Generative models of 3D anatomy, when integrated with biophysical simulators, enable the study of structure-function relationships for clinical research and medical device design. However, current models face a trade-off between controllability and anatomical realism. We propose a programmable and compositional framework for guiding unconditional diffusion models of human anatomy using interpretable ellipsoidal primitives embedded in 3D space. Our method involves the selection of certain tissues...

⚠ Similarity Notice

This paper appears to be a variant or shortened version of the original paper. Both papers share nearly identical titles (CardioComposer with similar subtitles), describe the same core technical approach using differentiable geometric guidance with ellipsoidal primitives for anatomical diffusion models, and present the same fundamental methodology of controlling size, shape, and position through geometric moment losses during inference.

Contributions Analysis

Overall novelty summary. CardioComposer introduces a programmable framework for generating multi-class anatomical label maps using interpretable ellipsoidal primitives and differentiable geometric moment functions to guide diffusion sampling. The paper resides in the 'Primitive-Based Compositional Control' leaf, which contains only two papers total (including the original work). This represents a relatively sparse research direction within the broader taxonomy of 13 papers across multiple branches, suggesting the primitive-based approach to anatomical generation is less explored compared to segmentation-mask conditioning methods.

The taxonomy reveals that CardioComposer's approach diverges from neighboring directions in meaningful ways. The sibling leaf 'Topological and Morphological Property Enforcement' focuses on persistent homology and topological features rather than primitive parameterization, while 'Landmark and Skeletal Structure Guidance' uses discrete point sets instead of continuous ellipsoidal representations. The broader 'Segmentation-Guided Anatomical Synthesis' branch (containing six papers across three leaves) represents a more crowded alternative paradigm that conditions on dense masks without explicit geometric parameterization, highlighting CardioComposer's distinct emphasis on interpretable, modular geometric handles.

Among 29 candidates examined, the contribution-level analysis shows varied novelty profiles. The differentiable geometric measurement functions (10 candidates examined, 0 refutable) and inference-time guidance framework (9 candidates examined, 0 refutable) appear to have limited direct prior work within the search scope. However, the compositional control framework for multi-part anatomical constraints (10 candidates examined, 1 refutable) shows at least one overlapping candidate, suggesting some existing work addresses multi-structure compositional generation. The limited search scope means these findings reflect top-K semantic matches rather than exhaustive coverage.

Given the sparse primitive-based leaf and the limited 29-candidate search, CardioComposer appears to occupy a relatively underexplored niche within anatomical diffusion modeling. The single sibling paper and the refutation of one contribution among 29 candidates suggest moderate novelty, though the analysis cannot rule out relevant work outside the top-K semantic neighborhood. The framework's extension to cardiac, vascular, and skeletal systems may represent incremental breadth rather than fundamental methodological departure from the sibling work.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Differentiable geometric measurement functions for anatomical characterization

Description: The authors develop differentiable functions that measure voxel-wise geometric moments to characterize anatomical substructures. These functions compute size via zeroth-order moments, position via first-order moments, and shape via scale-normalized second-order moments, enabling gradient-based optimization during diffusion sampling.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Latent graph representations for critical view of safety assessment

URL: View paper

Brief Assessment

Critical View Graph[24] focuses on latent graph representations for surgical scene understanding using bounding boxes and visual features, not on differentiable geometric moment functions for measuring anatomical size, position, and shape during diffusion sampling.

2. A skeletonization algorithm for gradient-based optimization

URL: View paper

Brief Assessment

Skeletonization Optimization[28] focuses on extracting medial axes from binary images using geometric moments for skeletonization algorithms, not on measuring anatomical features from multi-class label maps for diffusion model guidance.

3. Diff-TRGN: Diffusion-based tooth root generation network with multimodal clinical guidance

URL: View paper

Brief Assessment

Diff-TRGN[25] uses a differentiable projection module for geometric consistency between point clouds and 2D masks in dental imaging, which is fundamentally different from the original paper's voxel-wise geometric moment functions (zeroth, first, and second-order) for characterizing 3D anatomical substructures in diffusion models.

4. Explicit differentiable slicing and global deformation for cardiac mesh reconstruction

URL: View paper

Brief Assessment

Cardiac Mesh Slicing[22] focuses on differentiable voxelization and slicing for mesh reconstruction from medical images, not on measuring anatomical features via geometric moments for diffusion model guidance.

5. An ensemble shape gradient features descriptor based nodule detection paradigm: a novel model to augment complex diagnostic decisions assistance

URL: View paper

Brief Assessment

Ensemble Shape Gradient[31] focuses on nodule detection using gradient-based features for diagnostic assistance, not on differentiable geometric measurement functions for anatomical characterization or diffusion model guidance.

6. Shape matters: detecting vertebral fractures using differentiable point-based shape decoding

URL: View paper

Brief Assessment

Vertebral Fracture Shape[27] focuses on shape auto-encoders for vertebrae fracture detection using pre-trained surface patches, not on differentiable geometric moment functions for gradient-based optimization during diffusion sampling.

7. Dual Consistency Enabled Weakly and Semi-Supervised Optic Disc and Cup Segmentation With Dual Adaptive Graph Convolutional Networks

URL: View paper

Brief Assessment

Optic Disc Segmentation[29] focuses on optic disc/cup segmentation using graph convolutional networks with signed distance functions, not on developing differentiable geometric moment functions (zeroth, first, second-order) for measuring anatomical features during diffusion model sampling.

8. Masks-to-skeleton: Multi-view mask-based tree skeleton extraction with 3d gaussian splatting

URL: View paper

Brief Assessment

Masks-to-Skeleton[26] focuses on tree skeleton reconstruction from multi-view images using graph optimization and 3D Gaussian splatting for silhouette rendering. It does not develop differentiable geometric measurement functions for computing anatomical features like size, position, and shape from label maps via voxel-wise geometric moments.

9. Lung nodule detection and classification based on geometric fit in parametric form and deep learning

URL: View paper

Brief Assessment

Lung Nodule Geometric[30] focuses on gradient-based clustering for lung nodule detection, not on differentiable geometric moment functions for anatomical diffusion models or gradient-based optimization during sampling.

10. Comparative Analysis of Feature Extraction Techniques for Facial Paralysis Classification

URL: View paper

Brief Assessment

Facial Paralysis Features[23] focuses on facial landmark-based geometric features (distances and angles between facial points) for paralysis classification, not on differentiable voxel-wise geometric moment functions for 3D anatomical diffusion models.

Contribution 2: Inference-time guidance framework for controlling substructure geometry

Description: The authors present an inference-time method that uses gradients from geometric loss functions to guide unconditional diffusion models. This approach enables independent or joint control of substructure attributes without retraining the model, where substructures can consist of one tissue class or unions of multiple classes.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Guiding Diffusion Models for Spatially Consistent Image Generation

URL: View paper

Brief Assessment

Spatially Consistent Diffusion[41] focuses on spatial consistency in general image generation using vision-language models and layout-based guidance, not on geometric control of anatomical substructures through differentiable geometric loss functions.

2. Inference-Time Alignment Control for Diffusion Models with Reinforcement Learning Guidance

URL: View paper

Brief Assessment

Inference-Time Alignment[39] focuses on controlling diffusion models using RL-derived rewards for tasks like human preference alignment and compositional control, not geometric loss gradients for anatomical substructure control. The candidate's guidance mechanism (RL-based reward conditioning) differs fundamentally from the original's geometric moment-based losses.

---

3. Ditto: Diffusion inference-time t-optimization for music generation

URL: View paper

Brief Assessment

Ditto[32] focuses on music generation through optimizing initial noise latents for audio spectrograms, not anatomical diffusion models or geometric control of 3D medical segmentations. The technical domains and applications are fundamentally different.

---

4. ADPro: a Test-time Adaptive Diffusion Policy via Manifold-constrained Denoising and Task-aware Initialization for Robotic Manipulation

URL: View paper

Brief Assessment

ADPro[38] focuses on test-time adaptation for robotic manipulation using manifold-constrained denoising and task-aware initialization. The original paper addresses anatomical diffusion models with geometric loss functions for medical segmentations, while ADPro[38] applies diffusion policies to robot control with pose-based guidance. These are fundamentally different application domains with distinct technical approaches.

---

5. Unified Control for Inference-Time Guidance of Denoising Diffusion Models

URL: View paper

Brief Assessment

Unified Inference Control[35] focuses on general image generation tasks using diffusion models with reward-based guidance, while the original paper specifically addresses anatomical segmentation with geometric moment-based losses for controlling discrete substructures in medical imaging. The candidate does not demonstrate prior work in the specific domain of anatomical diffusion model control.

---

6. GeoGuide: Geometric Guidance of Diffusion Models

URL: View paper

Brief Assessment

GeoGuide[36] focuses on improving classifier-free guidance for class-conditional image generation in diffusion models, not on controlling geometric attributes of anatomical substructures using geometric loss gradients.

---

7. Controllable Music Production with Diffusion Models and Guidance Gradients

URL: View paper

Brief Assessment

Controllable Music Production[40] applies guidance gradients to music generation tasks, not anatomical diffusion models. The technical domain (audio vs. 3D medical segmentation) and the specific geometric moment-based losses for anatomical substructures are fundamentally different.

---

8. Non-Differentiable Diffusion Guidance for Improved Molecular Geometry

URL: View paper

Brief Assessment

Molecular Geometry Guidance[34] focuses on molecular geometry optimization using quantum mechanical oracles (GFN2-XTB), not anatomical segmentation control. The technical domains and applications are fundamentally different.

---

9. Flexible Geometric Guidance for Probabilistic Human Pose Estimation with Diffusion Models

URL: View paper

Brief Assessment

Geometric Pose Guidance[37] applies guidance to 2D-to-3D human pose estimation using 2D keypoint heatmap gradients, whereas the original paper controls 3D anatomical segmentation geometry using differentiable geometric moment losses. The domains and guidance mechanisms differ fundamentally.

Contribution 3: Compositional control framework for multi-part anatomical constraints

Description: The authors demonstrate that their framework supports compositional generation by combining multiple substructure-specific geometric losses. This enables complex anatomical constraints across arbitrary numbers of substructures, including non-convex geometries with branching or curved features.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. CardioComposer: Flexible and Compositional Anatomical Structure Generation with Disentangled Geometric Guidance

URL: View paper

Brief Assessment

CardioComposer Flexible[8] describes the same framework as the original paper. The candidate text appears to be an abstract or summary of the original work itself, not a distinct prior work that could refute novelty claims.

---

2. Significance of Anatomical Constraints in Virtual Try-On

URL: View paper

Brief Assessment

Virtual Try-On Constraints[15] addresses virtual clothing try-on using part-based warping for garments, not anatomical diffusion models or multi-substructure geometric control for medical simulations.

---

3. Relational Anatomical Supervision for Accurate 3D Multi-Chamber Cardiac Mesh Reconstruction

URL: View paper

Brief Assessment

Cardiac Mesh Supervision[19] focuses on enforcing spatial relationships between cardiac chambers (e.g., LV enclosed by myocardium, RV separated from LV) through occupancy-based mesh supervision, not on compositional generation by combining multiple substructure-specific geometric losses for arbitrary anatomical structures as in the original paper.

---

4. TreeNet: multi-loss deep learning network to predict branch direction for extracting 3D anatomical trees

URL: View paper

Brief Assessment

TreeNet[21] focuses on predicting branch directions for extracting 3D anatomical trees (vascular/neural structures), not on compositional control using multiple geometric losses for complex anatomical constraints across arbitrary substructures.

---

5. Anatomy-guided convolutional neural network for motion correction in fetal brain MRI

URL: View paper

Brief Assessment

Fetal Brain CNN[17] focuses on motion correction in fetal brain MRI using anatomy-guided convolutional networks with multi-loss frameworks, not on compositional generation of multi-part anatomical structures using geometric losses for diffusion models.

---

6. CMF-Net: craniomaxillofacial landmark localization on CBCT images using geometric constraint and transformer

URL: View paper

Brief Assessment

CMF-Net[14] focuses on landmark localization in craniomaxillofacial CBCT images using geometric constraints for spatial relationships between landmarks, not on compositional generation of multi-part anatomical structures using multiple geometric losses.

---

7. Neural-symbolic emotion-pose graph reasoning in AI-based human synthesis: A multimodal model integrating cognitive priors--Digital Restoration of the Aesthetics of â¦

URL: View paper

Brief Assessment

Neural-Symbolic Emotion[20] focuses on emotion-pose graph reasoning for human synthesis with symbolic constraints, not on compositional geometric control of anatomical diffusion models using differentiable losses.

---

8. Geometric Constrained Deep Learning for Motion Correction of Fetal Brain Mr Images

URL: View paper

Brief Assessment

Fetal Brain Motion[16] addresses motion correction in fetal brain MRI using geometric constraints between adjacent slices, not compositional generation with multiple substructure-specific geometric losses for arbitrary anatomical constraints.

---

9. RegioMorph-GAN: PPA Morphology-Driven Fundus Image Synthesis with Region-Focused Constraint

URL: View paper

Brief Assessment

RegioMorph-GAN[18] focuses on conditional GAN-based synthesis of fundus images with PPA lesion morphology control using mask inputs and region-focused losses. It does not address compositional control across multiple anatomical substructures using geometric losses as in the original paper.

---

10. Anatomica: Localized Control over Geometric and Topological Properties for Anatomical Diffusion Models

URL: View paper

Prior Art Analysis

Anatomica[5] demonstrates that compositional control using multiple geometric losses for complex anatomical constraints was implemented prior to the original paper's submission. Both papers use differentiable geometric moment-based losses (mass, centroid, covariance) to guide unconditional diffusion models, enabling compositional generation across multiple substructures. Anatomica[5] extends this to localized control domains with varying dimensionality and coordinate systems, while also incorporating topological constraints. The core mechanism of composing multiple substructure-specific geometric losses to enable complex anatomical constraints is present in both works, with Anatomica[5] providing evidence of this approach being developed independently or earlier.

Evidence

Evidence 1 - **Rationale**: Both papers claim compositional control through combining multiple geometric constraints. Anatomica[5] explicitly demonstrates 'compositional anatomical control' by combining different control domains and potentials, which directly parallels the original paper's claim of composing 'multiple substructure-specific geometric losses'. - **Original**: we validate that multiple substructure-specific geometric losses can be composed to enable more complex anatomical constraints. further, we show that this control extends to non-convex substructures with branching or curved geometry. - **Candidate**: we demonstrate that applying differentiable measurement and potential functions over anatomical substructures allows us to constrain localized properties through diffusion guidance. this includes geometric properties such as size, shape, position, and orientation, as well as topological properties s...

Evidence 2 - **Rationale**: The 'key insight' claimed by the original paper—using gradients from geometric loss functions applied to individual substructures—is implemented in Anatomica[5] through their geometric potential function that combines mass, position, and covariance losses for each substructure. - **Original**: ourkey insightis that unconditional diffusion models of multi-class anatomy can be constrained in a compositional manner by simple gradients derived from geometric loss functions applied individually to each substructure. - **Candidate**: we aim to penalize the deviations from the target geometric features ˉgk through a geometric potential functionl geo k . we formulate this potential function as a weighted combination of mean squared error lossesl mse: lgeo k =λ 0lmse(mk,ˉmk) +λ 1lmse(pk, ˉpk) +λ2lmse(σn k , ˉσ n k).

Evidence 3 - **Rationale**: Both papers use identical geometric moment formulations: zeroth-order moments for size/mass, first-order for position/centroid, and second-order for shape/covariance. This demonstrates that the differentiable geometric measurement approach is not novel to the original paper. - **Original**: we introduce a set of differentiable geometric measurement functions that compute physiologically relevant anatomical features from a substructure label map. we specifically measure voxel-wise geometric moments compute size via zeroth-order moments, position via first-order moments, and shape via sc... - **Candidate**: given a single parsed substructures k, we aim to differentiably extract the geometric moments by numerically integrating the zeroth, first, and second-order moments of the substructure. here, the zeroth moment represents the massm k ∈r, the first moment represents the centroidp k ∈r 3, and the secon...

Evidence 4 - **Rationale**: Anatomica[5] implements gradient-based guidance of unconditional diffusion models using geometric loss functions, enabling control without retraining—the same approach claimed as novel by the original paper. - **Original**: we demonstrate that simple gradients derived from differentiable geometric loss functions can guide unconditional latent diffusion models of discretized multi-class label maps. this enablesindependentorjoint control of substructure attributes without retraining - **Candidate**: following kadry et al. [29], we guide the diffusion process using the gradient derived from anatomical potential functions. at each sampling step, we first denoise the intermediately noised latentz σ to obtain a clean latent prediction ˆz0 =d θ(zσ;σ), which is thenparsedintoksubstructuress k

Appendix: Text Similarity Detection

Textual similarity detection checked 30 papers and found 5 similarity segment(s) across 2 paper(s).

The following **2 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. Anatomica: Localized Control over Geometric and Topological Properties for Anatomical Diffusion Models

Detected in: Contribution: contribution_3

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

2. CardioComposer: Flexible and Compositional Anatomical Structure Generation with Disentangled Geometric Guidance

Detected in: Core Task (sibling), Contribution: contribution_3

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] CardioComposer: Leveraging Differentiable Geometry for Compositional Control of Anatomical Diffusion Models View paper
- [1] Data augmentation for surgical scene segmentation with anatomy-aware diffusion models View paper
- [2] Solving the where problem and quantifying geometric variation in neuroanatomy using generative diffeomorphic mapping View paper
- [3] Heartbeat: Towards controllable echocardiography video synthesis with multimodal conditions-guided diffusion models View paper
- [4] Semantic data augmentation with generative models View paper
- [5] Anatomica: Localized Control over Geometric and Topological Properties for Anatomical Diffusion Models View paper
- [6] A Diffusion Model for Simulation Ready Coronary Anatomy with Morpho-Skeletal Control View paper
- [7] Robust Polyp Detection and Diagnosis Through Compositional Prompt-Guided Diffusion Models. View paper
- [8] CardioComposer: Flexible and Compositional Anatomical Structure Generation with Disentangled Geometric Guidance View paper
- [9] 3DMM-GAN: Multi-Modal Alignment With Adversarial Learning for Compositional 3D Human Image Synthesis View paper
- [10] LaTo: Landmark-tokenized Diffusion Transformer for Fine-grained Human Face Editing View paper
- [11] Anatomically-Controllable Medical Image Generation with Segmentation-Guided Diffusion Models View paper
- [12] HairFree: Compositional 2D Head Prior for Text-Driven 360Â° Bald Texture Synthesis View paper
- [13] Phase-Fraction-Guided Denoising Diffusion Model for Multiphase Steel Microstructure Segmentation Augmentation Via Micrograph Image-Mask Pair Synthesis View paper
- [14] CMF-Net: craniomaxillofacial landmark localization on CBCT images using geometric constraint and transformer View paper
- [15] Significance of Anatomical Constraints in Virtual Try-On View paper
- [16] Geometric Constrained Deep Learning for Motion Correction of Fetal Brain Mr Images View paper
- [17] Anatomy-guided convolutional neural network for motion correction in fetal brain MRI View paper
- [18] RegioMorph-GAN: PPA Morphology-Driven Fundus Image Synthesis with Region-Focused Constraint View paper
- [19] Relational Anatomical Supervision for Accurate 3D Multi-Chamber Cardiac Mesh Reconstruction View paper
- [20] Neural-symbolic emotion-pose graph reasoning in AI-based human synthesis: A multimodal model integrating cognitive priors--Digital Restoration of the Aesthetics of â¦ View paper
- [21] TreeNet: multi-loss deep learning network to predict branch direction for extracting 3D anatomical trees View paper
- [22] Explicit differentiable slicing and global deformation for cardiac mesh reconstruction View paper
- [23] Comparative Analysis of Feature Extraction Techniques for Facial Paralysis Classification View paper
- [24] Latent graph representations for critical view of safety assessment View paper
- [25] Diff-TRGN: Diffusion-based tooth root generation network with multimodal clinical guidance View paper
- [26] Masks-to-skeleton: Multi-view mask-based tree skeleton extraction with 3d gaussian splatting View paper
- [27] Shape matters: detecting vertebral fractures using differentiable point-based shape decoding View paper
- [28] A skeletonization algorithm for gradient-based optimization View paper
- [29] Dual Consistency Enabled Weakly and Semi-Supervised Optic Disc and Cup Segmentation With Dual Adaptive Graph Convolutional Networks View paper
- [30] Lung nodule detection and classification based on geometric fit in parametric form and deep learning View paper
- [31] An ensemble shape gradient features descriptor based nodule detection paradigm: a novel model to augment complex diagnostic decisions assistance View paper
- [32] Ditto: Diffusion inference-time t-optimization for music generation View paper
- [33] DragonDiffusion: Enabling Drag-style Manipulation on Diffusion Models View paper
- [34] Non-Differentiable Diffusion Guidance for Improved Molecular Geometry View paper
- [35] Unified Control for Inference-Time Guidance of Denoising Diffusion Models View paper
- [36] GeoGuide: Geometric Guidance of Diffusion Models View paper
- [37] Flexible Geometric Guidance for Probabilistic Human Pose Estimation with Diffusion Models View paper
- [38] ADPro: a Test-time Adaptive Diffusion Policy via Manifold-constrained Denoising and Task-aware Initialization for Robotic Manipulation View paper
- [39] Inference-Time Alignment Control for Diffusion Models with Reinforcement Learning Guidance View paper
- [40] Controllable Music Production with Diffusion Models and Guidance Gradients View paper
- [41] Guiding Diffusion Models for Spatially Consistent Image Generation View paper