Scilight

Journal of Bio-optics

https://www.sciltp.com/journals/jbo

Review

Ophthalmic OCT-Based Image Generation Using GANs: A Scoping Review

Hadi Afsharan [1,2,3,*], Parmida Ghorbanian [1,2], Farzan Navaeipour [1], Najmeh Fayyazifar [1], Yiheng Lyu [1,4], Mohammed Bennamoun [4], Barry Cense [2] and Girish Dwivedi [1,3,5]

[1] Harry Perkins Institute of Medical Research, The University of Western Australia, Perth, WA 6009, Australia
[2] Optical + Biomedical Engineering Laboratory, Department of Electrical, Electronic and Computer Engineering, The University of Western Australia, Perth, WA 6009, Australia
[3] Victor Chang Cardiac Research Institute, Darlinghurst, NSW 2010, Australia
[4] Department of Computer Science and Software Engineering, The University of Western Australia, Perth, WA 6009, Australia
[5] Department of Cardiology, Fiona Stanley Hospital, Murdoch, WA 6150, Australia
* Correspondence: hadi.afsharan@UWA.edu.au

How To Cite: Afsharan, H.; Ghorbanian, P.; Navaeipour, F.; et al. Ophthalmic OCT-Based Image Generation Using GANs: A Scoping Review. *Journal of Bio-optics* **2026**, *2*(1), 1. https://doi.org/10.53941/jbiop.2026.100001

Received: 28 November 2025
Revised: 22 January 2026
Accepted: 9 March 2026
Published: 8 April 2026

Abstract: Recent advances in artificial intelligence (AI) and deep learning (DL) are reshaping ophthalmology, particularly in the domains of image analysis and image synthesis. Optical coherence tomography (OCT) has become indispensable for diagnosing and monitoring retinal diseases, and the application of generative models such as generative adversarial networks (GANs) now enables the creation of realistic synthetic OCT data. In this systematic review, we evaluate the current state of OCT image generation using DL techniques, with emphasis on retinal applications including common pathologies such as age-related macular degeneration, diabetic retinopathy, glaucoma, diabetic retinopathy and other retinal pathologies. We provide an overview of commonly employed architectures, including GAN variants, variational autoencoders, and emerging diffusion models, and highlight how they have been applied for data augmentation, cross-modality translation, noise reduction, and rare pathology synthesis. We discuss validation strategies, performance metrics, and limitations across existing studies, and emphasize the clinical opportunities these technologies present in improving diagnostic accuracy, education, and accessibility of advanced imaging. Finally, we identify gaps in dataset diversity, external validation, and regulatory considerations, and outline future directions to ensure responsible translation of synthetic OCT imaging into clinical practice.

Keywords: optical coherence tomography (OCT); OCT angiography (OCTA); polarization-sensitive OCT (PS-OCT); retinal imaging; deep learning; generative adversarial networks (GANs); synthetic image generation; medical image synthesis

1. Introduction

Optical coherence tomography (OCT) has transformed ophthalmic practice since the early 1990s, enabling non-invasive, high-resolution imaging of retinal architecture with unparalleled detail [1,2]. The transition from time-domain [3] to spectral-domain and swept-source OCT [4] has established the technology as the gold standard for diagnosing and monitoring diverse retinal diseases, including age-related macular degeneration [5], diabetic retinopathy [6], and glaucoma [7]. Beyond structural imaging, advanced modalities such as OCT angiography (OCTA) [8] and polarization-sensitive OCT (PS-OCT) [9] extend diagnostic capabilities by revealing vascular and tissue birefringence properties.

The integration of artificial intelligence (AI) with medical imaging has catalyzed a paradigm shift in retinal image analysis [10,11]. Deep learning (DL) now enables automated diagnosis, segmentation, and classification with remarkable accuracy [12,13]. Within this landscape, generative models and particularly generative adversarial networks (GANs), have emerged as powerful tools to address persistent challenges in medical imaging, such as data scarcity [14], class imbalance, and the generation of synthetic data for rare pathologies [15,16]. While OCT-based imaging encompasses a broader range of modalities, this review focuses specifically on retinal structural and functional imaging and does not cover other OCT-derived techniques such as OCE [17,18], which target tissue biomechanics and follow distinct acquisition and analytical paradigms. Building on this foundation, the next section outlines how research on OCT image generation has evolved, highlighting key milestones and methodological shifts.

1.1. Evolution of OCT Image Generation Research

DL-driven OCT image synthesis is a rapidly advancing field with important clinical and research implications [16,19]. Earlier data augmentation techniques, limited to geometric or intensity transformations, were insufficient to capture the complex anatomical variation and pathological diversity seen in practice [20,21]. The introduction of GANs in 2014 provided a framework to generate realistic medical images while preserving both structural integrity and disease-specific features [22].

In OCT, generative models support multiple applications: dataset expansion for robust model training [16], cross-modality translation [23], generation of synthetic pathologies for DL training [24], classification [25], and reconstruction of artifact-degraded images [26]. However, the intrinsic properties of OCT, including speckle noise, depth-dependent attenuation, and layered retinal architecture, pose unique challenges that demand specialized generative solutions.

1.2. Current State of Literature Reviews

Several reviews have addressed DL in ophthalmology, yet few have focused specifically on OCT image generation.

The seminal review by Yanagihara et al. [27] established a framework for DL in OCT, highlighting key limitations such as dataset size, annotation quality, and model interpretability. However, it did not address image synthesis directly. Kugelman et al. [28] provided the first targeted review of GANs in OCT, covering anomaly detection, image enhancement, and domain translation. You et al. [29] broadened this perspective, examining GANs across multiple ophthalmic imaging modalities, with insights into cross-modality synthesis and comparative model performance. Around the same time, Wang et al. [30] outlined one of the first structured evaluations of GAN-based ophthalmic image synthesis, emphasizing clinical implications of synthetic OCT and fundus images.

Subsequent work has addressed niche areas. Ahmed et al. [31] systematically reviewed DL-based OCT denoising, while Li et al. [32] conducted a gap analysis of OCT-focused DL, underscoring critical research needs.

Broader reviews have provided important context. Badar et al. [11] surveyed DL in retinal image analysis, covering OCT but without addressing synthesis. Hormel et al. [33] focused on AI in OCT angiography, underscoring the role of DL for vascular biomarker quantification. More recently, Masalkhi et al. [12] systematically reviewed GAN applications for ophthalmic image synthesis, including OCT, while Ali et al. [34] and Oulmalme et al. [35] highlighted cross-modality medical image synthesis and enhancement, discussing the rise of diffusion models and transformers alongside GANs.

Beyond GANs, other generative models have also been explored for image synthesis in ophthalmology and medical imaging more broadly. Variational autoencoders (VAEs) in particular have been reviewed in the context of synthetic medical image generation. Rais et al. [36] provided a comprehensive overview of VAE-based approaches, highlighting their ability to generate diverse, high-fidelity images in data-scarce clinical settings, while also contrasting their outputs with GANs and diffusion models. Kebaili et al. [37] extended this discussion by surveying DL-driven data augmentation strategies in medical imaging, with VAEs, GANs, and diffusion models all considered as core approaches for dataset expansion and synthesis.

Together, these contributions indicate that while GANs have dominated the literature, VAEs and diffusion models are increasingly recognized in both review and application studies as promising alternatives for OCT image synthesis.

While these reviews provide valuable context, they stop short of offering a comprehensive synthesis focused exclusively on OCT image generation, leaving an important gap that this work addresses.

1.3. Identified Research Gap and Study Rationale

Despite rapid progress, no comprehensive review has yet focused exclusively on DL-based OCT image generation and synthesis. This is a notable omission, given that OCT synthesis has matured into a distinct domain with specialized architectures, evaluation metrics, and clinical applications.

As mentioned before, recent developments, including diffusion models, vision transformers, and hybrid methods, have shown promise beyond traditional GANs [38–40]. Applications now extend from dataset augmentation to treatment outcome prediction [41], cross-modality synthesis [23], and the creation of rare pathologies for research [16] and training [24]. Meanwhile, evaluation methodologies have advanced from standard image metrics to clinical validation studies, perceptual similarity measures, and task-specific performance benchmarks. These developments warrant a dedicated review to consolidate knowledge, highlight best practices, and chart future directions.

1.4. Review Scope and Contributions

This work provides the first systematic review of DL-based OCT image generation, with emphasis on GAN-based approaches while also covering emerging paradigms such as diffusion models and transformers. We include conventional OCT, OCTA, and PS-OCT, reflecting the integrated workflows of modern retinal imaging.

The review makes five contributions:

1. A systematic analysis of state-of-the-art generative models applied to OCT synthesis.
2. Comprehensive evaluation of quality assessment and clinical validation methods.
3. Critical discussion of clinical applications, from augmentation to treatment prediction.
4. Identification of current technical challenges and limitations.
5. Projection of future research directions and emerging paradigms.

Our OCT focus is motivated by its ubiquity in ophthalmic practice, the ability to synthesize derived modalities (OCTA, PS-OCT) from structural data, and the unique imaging challenges, including coherent noise, depth resolution, and quantitative metrics that distinguish OCT from other medical imaging domains. While OCT is used in multiple anatomical sites, we focus exclusively on retinal/ophthalmic applications, where OCT has achieved the greatest clinical adoption and where synthetic image generation has the most immediate translational potential.

1.5. Clinical Relevance and Implications

High-quality OCT synthesis has wide-ranging implications for clinical practice, education, and research [42]. In resource-limited settings, synthetic data may enable algorithm development without extensive local datasets [43]. For medical education, generative models can produce diverse and rare pathologies to enrich training [16,42]. In research, synthetic OCT supports algorithm validation [41], accelerates multi-center studies [25], and facilitates exploration of uncommon disease presentations [15].

As ophthalmology advances toward personalized medicine and AI-driven care, the clinical value of synthetic imaging will only grow [14,44]. This review provides researchers and clinicians with a consolidated resource on current methods, validated evaluation strategies, and future opportunities in OCT image generation ultimately supporting the development of robust AI systems and improved patient outcomes.

The remainder of the article systematically reviews generative architectures, clinical applications, validation methodologies, and emerging research directions, serving as both a reference for researchers and a practical guide for clinical implementation.

2. OCT imaging Modalities

2.1. Optical Coherence Tomography

OCT is a non-invasive imaging technique that produces high-resolution cross-sectional views of retinal tissue through interferometric detection of backscattered light [45], Figure 1. Compared with ultrasound or MRI, OCT offers superior resolution, combining axial resolution on the order of ~5 μm (comparable to confocal microscopy) with lateral resolution of ~25 μm (similar to confocal scanning laser ophthalmoscopy) [46]. Unlike traditional optical modalities limited by pupil size, interferometric detection defines resolution by the light source itself, enabling consistently high-quality imaging [47].

The evolution from time-domain OCT (TD-OCT) to spectral-domain (SD-OCT) and swept-source OCT has dramatically increased imaging speed and quality by capturing depth information simultaneously rather than sequentially [45]. These advances have established OCT as the gold standard for retinal imaging, particularly in

macular diseases, diabetic retinopathy, and glaucoma [47]. Retinal layer segmentation enables quantitative thickness measurements, critical for detecting nerve fiber layer thinning in glaucoma progression [48,49].

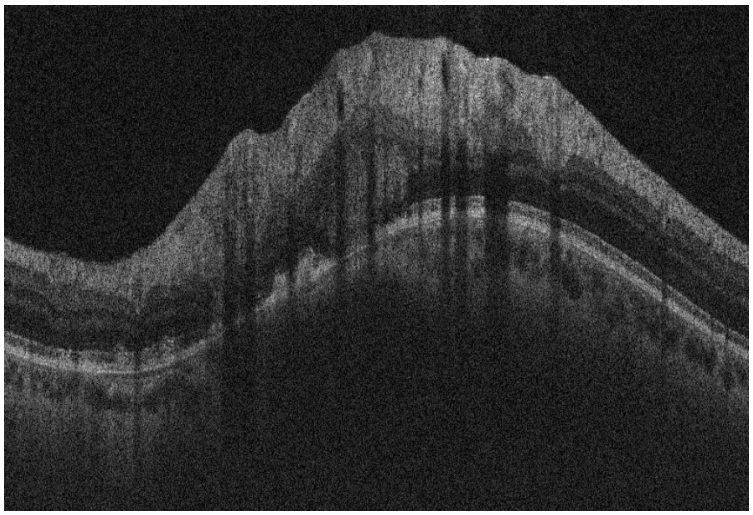

Figure 1. OCT circular B-scan (cross-section) of an area near the optic nerve of the right eye of a healthy volunteer. The image was acquired using a circular scan with a radius of 3.5 degrees from the center of the optic nerve.

2.2. OCT Angiography

OCTA extends conventional OCT by detecting blood flow without the need for intravenous dye injection. The technique acquires multiple consecutive OCT scans at identical locations and measures inter-scan differences to detect motion, distinguishing dynamic blood flow from static tissue. This enables detailed visualization of the superficial and deep retinal capillary plexuses, as well as the choriocapillaris [8].

As shown in Figure 2, OCTA provides quantitative vascular metrics including vessel density, flow void areas, and foveal avascular zone characteristics. Clinical advantages include rapid acquisition (3–6 s), absence of dye-related risks, and suitability for repeated follow-up examinations. OCTA has shown particular value in diabetic retinopathy staging and the evaluation of neovascularisation in age-related macular degeneration [50].

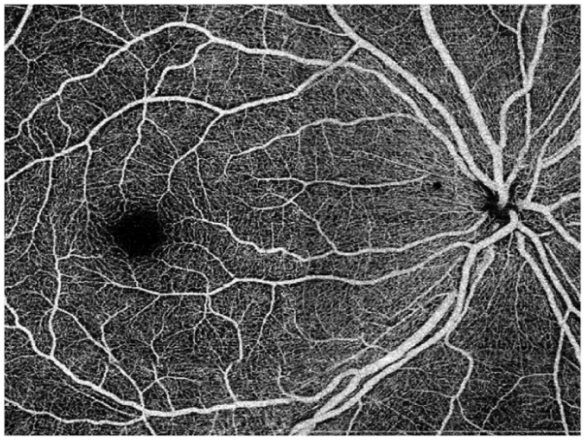

Figure 2. Swept source montage OCT-Angiogram showing fovea and OD of a normal eye, from [50].

2.3. Polarization-Sensitive OCT

While conventional OCT provides excellent structural detail, it lacks intrinsic tissue specificity. PS-OCT addresses this limitation by analysing changes in the polarization state of light as it interacts with birefringent tissue [51]. Metrics such as the degree of polarization uniformity (DOPU) allow differentiation of tissue microstructure and integrity [52]. PS-OCT has demonstrated promise in automated segmentation of the retinal pigment epithelium and early detection of drusen, and is particularly sensitive to fibrotic, collagen-rich tissues [53]. Recently, it has been associated with systemic diseases such as diabetes [54] and hypertension [55]. However, clinical translation has been constrained by the complexity and cost of PS-OCT systems, along with reduced portability and the need for specialised expertise. Recent research has therefore focused on computational synthesis

of PS-OCT contrasts from standard OCT images using DL, as a strategy to expand accessibility. Figure 3 shows OCT and corresponding PS-OCT images of an area near the optic nerve head of a healthy eye.

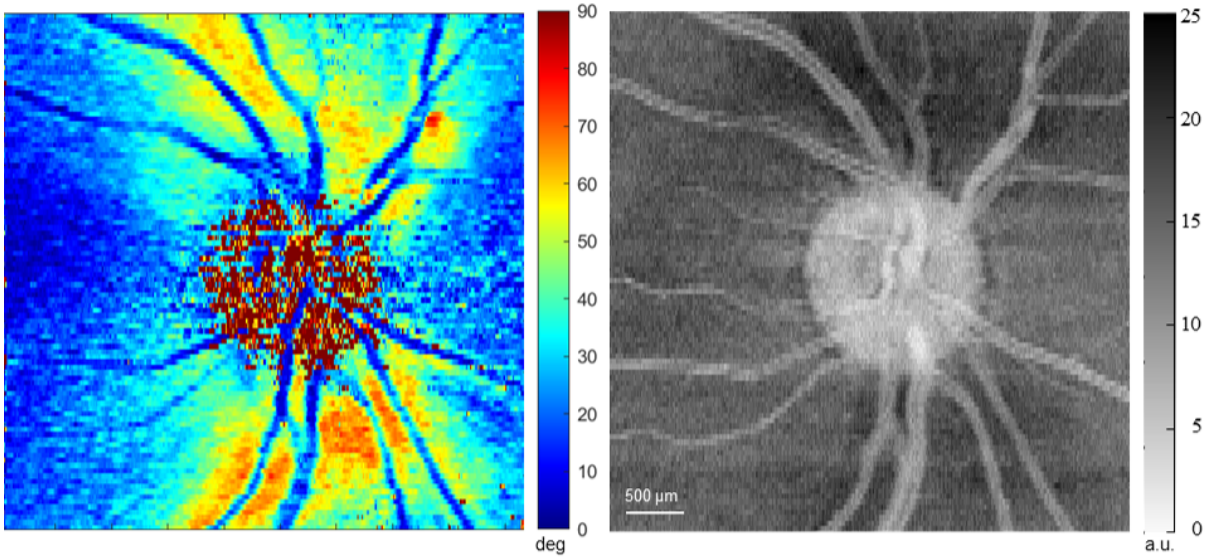

Figure 3. OCT (**right**) and PS-OCT (**left**) retardation en-face images of the optic nerve head of a normal eye. The right panel shows conventional OCT intensity, while the left panel displays PS-OCT-specific retardation measurements (in degrees).

2.4. Clinical Integration and Context for Image Generation

Modern OCT systems increasingly integrate multiple modalities: structural OCT, OCTA, and PS-OCT, to provide a comprehensive assessment of retinal health [1]. Each modality presents unique challenges for DL-based image generation: structural OCT requires faithful preservation of laminar architecture and speckle statistics [19]; OCTA demands accurate reproduction of vascular topology [56]; and PS-OCT necessitates preservation of birefringence-related contrasts [57].

The ability to synthetically generate these modalities offers multiple opportunities: dataset augmentation for training, cross-modality translation to reduce reliance on costly or invasive techniques, and the creation of rare pathological presentations for research and education. These applications form the foundation for the DL approaches explored in the subsequent sections of this review.

To rigorously assess this emerging field, we adopted a systematic review approach, as detailed in the following methodology.

3. Methods

3.1. Study Design and Protocol

This review was conducted in accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines [58]. The protocol was developed to systematically identify and critically evaluate studies on DL-based approaches for OCT image generation, with a particular focus on GANs and their clinical applications in ophthalmology.

3.1. Literature Search Strategy

A comprehensive search was conducted across PubMed/MEDLINE, IEEE Xplore, Web of Science, Embase, MDPI, The JAMA Network, and Investigative Ophthalmology & Visual Science (ARVO). Google Scholar was additionally consulted to capture gray literature. The search covered studies published between January 2020 and August 2024. Both Medical Subject Headings (MeSH) and free-text terms were used. The Boolean query applied (example in PubMed; adapted to each database syntax) was listed in Box 1.

Box 1. Search strategy and boolean queries used for database screening.

("DL" **OR** "artificial intelligence" **OR** "machine learning" **OR** "generative adversarial network" OR GAN **OR** "image generation" **OR** "synthetic image*" **OR** "image synthesis"**) AND (**"optical coherence tomography" **OR** OCT **OR** "OCT angiography" **OR** OCTA **OR** "polarization-sensitive OCT" **OR** PS-OCT **OR** "retinal imaging" **OR** "retinal image*"**) AND (**"age-related macular degeneration" **OR** AMD **OR** "diabetic retinopathy" **OR** DR **OR** glaucoma **OR** "retinal disease*"**)**

"image*" → retrieves image, images, imaging. "retinal image*" → retrieves retinal image, retinal images, retinal imaging. "synthetic image*" → retrieves synthetic image, synthetic images.

To ensure comprehensive coverage, reference lists of relevant articles were manually screened. Preprint servers (arXiv, bioRxiv, medRxiv) and institutional repositories were also searched for recent or unpublished work.

3.2. Study Selection Criteria

Eligible studies were required to report original research that applied DL methods, particularly GANs, for OCT image generation or synthesis in retinal or ophthalmic applications. Only studies published in English between January 2020 and August 2024 were considered. To be included, studies needed to describe their datasets, report performance metrics, and provide sufficient methodological detail to allow assessment of quality and reproducibility. Our search strategy was intentionally focused on retinal pathologies (AMD, diabetic retinopathy, glaucoma) as these represent the most clinically prevalent and well-studied OCT applications. This focus ensures clinical relevance while maintaining sufficient scope for comprehensive synthesis of generative modeling approaches.

We included studies using both public and private datasets, provided that sufficient detail was reported to assess data quality and enable partial reproducibility. Required information included: dataset size (number of images or volumes), disease categories, and imaging device/modality. Studies were excluded only when dataset information was entirely absent or too vague to assess methodological validity

We excluded reviews, editorials, and commentaries, as well as studies lacking adequate technical description, dataset information, or validation. Papers that focused exclusively on non-OCT modalities without cross-modality synthesis, investigated non-retinal OCT applications, or addressed only classification and segmentation without an image-generation component were also excluded. Duplicate and non-English publications were removed. The detailed characteristics of the included studies are provided in Table 1.

Table 1. Eligibility for inclusion in this review was determined according to the criteria summarized in Box 1.

Inclusion Criteria	Exclusion Criteria
– Studies with GAN output being synthesized OCT, OCTA, or PS-OCT images. – Studies that used GANs primarily for the purpose of image synthesis. – Studies that contained measurable performance metrics including accuracy, sensitivity, recall, specificity, AUC, F1-score, PSNR, SSIM, KID. – Papers that provide adequate description of training datasets, including source, size, disease distribution, and acquisition parameters (publicly available or institutional datasets with sufficient detail)	– Studies performed with DL but did not include image generation/synthesis. – Studies performed with eye images other than OCT, such as color fundus photography, ultra-wide fundus, and fluorescein angiography fundus, as the input to the GAN model. – Studies aimed to segment, analyze or classification of OCT images using GANs or DLs and are not related to synthesis or generation. – Studies that provide no description of training data source, size, or characteristics, preventing reproducibility assessment – Studies that did not use DL techniques in the synthesis stage. – Studies where images of skin, intravascular, lungs, etc. and not retinas were generated (excluding PS-OCT studies, because of lack of available datasets and studies on retinal PS-OCT synthesis). – Studies where DL model performance metrics were not reported, or where reported metrics were

substantially below baseline levels (defined as >20% lower than comparable studies using similar datasets and tasks) without adequate methodological justification.

Exception for PS-OCT: Given the limited availability of retinal PS-OCT datasets and the nascent state of PS-OCT clinical adoption, we included non-retinal PS-OCT synthesis studies to provide a more complete picture of PS-OCT generative modeling capabilities. This exception is justified by: (i) the scarcity of retinal PS-OCT work, (ii) the transferability of PS-OCT synthesis principles across anatomical sites, and (iii) the goal of highlighting emerging methodologies that may be adapted for retinal applications".

3.3. Study Selection Results

The initial search identified 1,847 records. After removing 1,046 duplicates (56% of the initial set, mainly from overlap between Web of Science and PubMed), 801 unique records remained for title and abstract screening. Of these, 734 were excluded according to the predefined criteria, leaving 67 studies for full-text review.

During full-text assessment, 46 studies were excluded: 18 did not employ DL for image generation, 12 provided insufficient methodological detail, 8 focused on non-retinal OCT applications, 5 reported inadequate validation, and 3 lacked dataset description. Ultimately, 21 studies met all inclusion criteria and were retained for analysis. The complete selection process is illustrated in Figure 4.

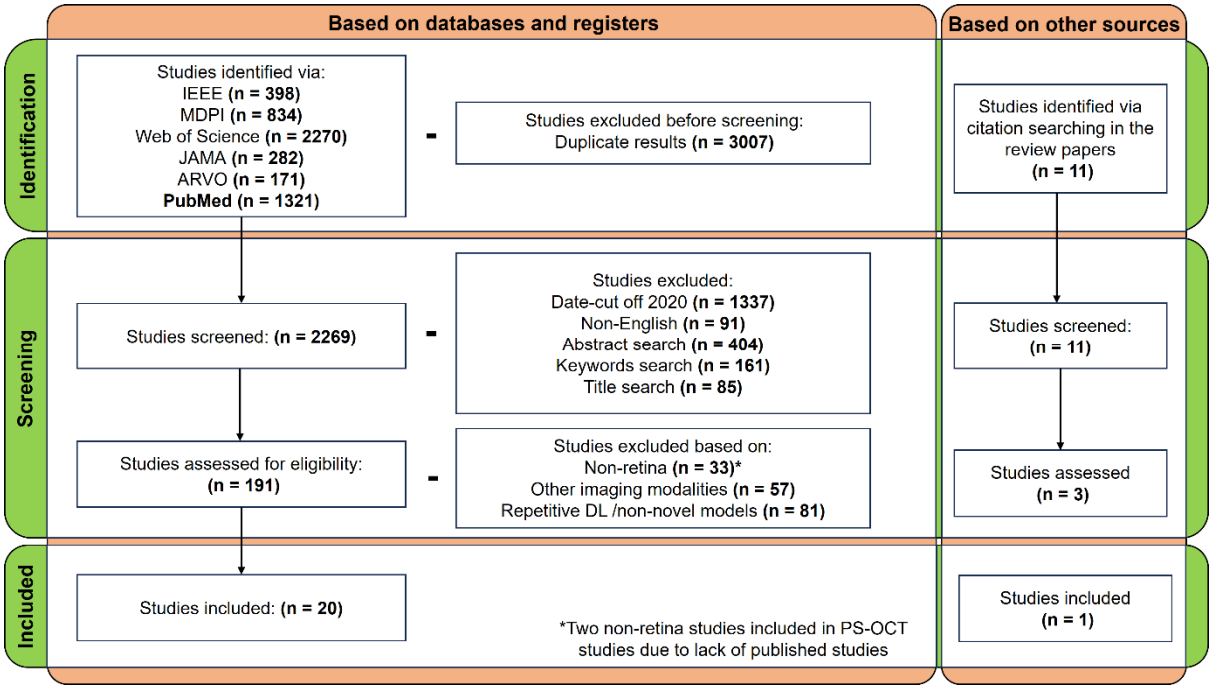

Figure 4. Flow diagram of the study selection process following PRISMA guidelines, showing the number of records identified, screened, excluded, and included in the review.

3.4. Screening and Data Extraction

All search results were imported into EndNote X20, where duplicates were removed automatically and verified manually. Two reviewers (H.A. and F.N.) independently screened titles and abstracts. Full-text review was then performed using a standardized extraction form, which captured details on study characteristics, methodology, datasets, evaluation metrics, results, and limitations.

Risk of bias was assessed based of four different domains:

1. Dataset quality—(number of images used for training, private or publicity of the dataset)
2. Evaluation metric appropriateness and completeness
3. Clinical validation (expert grading, diagnostic task performance)
4. Reproducibility (code availability, parameter specification)

Studies were rated as low, moderate, or high risk in each domain. These assessments were also performed independently by two reviewers (H.A. and F.N.), with disagreements resolved through discussion.

4. DL Fundamentals for OCT Image Generation

4.1. Transition from Traditional to DL Approaches

Early approaches to retinal image analysis relied heavily on handcrafted features such as Speeded-up Robust Features (SURF), Scale Invariant Feature Transform (SIFT), and Histogram of Oriented Gradients (HOG) [59]. While effective for basic texture and edge detection, these methods lacked adaptability across diverse OCT datasets and struggled with speckle noise, depth attenuation, and the layered retinal structure [60]. DL has overcome these limitations by automatically learning hierarchical representations directly from raw images, enabling robust detection of anatomical boundaries and pathological features [13,14].

4.2. Convolutional Neural Networks (CNNs) in OCT Imaging

Convolutional Neural Networks (CNNs) remain central to OCT image analysis. Architectures such as U-Net, ResNet, and DenseNet have been widely adopted for tasks including retinal layer segmentation, fluid detection, and disease classification [61]. Their strength lies in automatically extracting spatially hierarchical features from retinal boundaries to subtle pathological structures without manual preprocessing [13]. Importantly, CNN-based models also serve as evaluators of synthetic data, allowing quantification of whether GAN-generated OCT images preserve clinically relevant features [62].

4.3. GANs for Image Synthesis

GANs have been the dominant framework for synthetic OCT generation. Their adversarial training paradigm, putting a generator against a discriminator, enables the production of realistic, high-resolution retinal images [16].

4.3.1. Core Architecture and Adversarial Training

A GAN consists of two primary components: Generator: A neural network that learns to create synthetic images from random noise vectors (latent space). The goal of the generator is to produce images that are indistinguishable from real OCT scans. Starting from random input, the generator produces synthetic images that should match the distribution of real training data. Discriminator: A neural network that learns to distinguish between real OCT images (from the training dataset) and synthetic images (produced by the generator). The discriminator outputs a probability score indicating whether an input image is real or synthetic.

4.3.2. The Adversarial Process

During training, these two networks engage in a minimax game. The generator tries to fool the discriminator by creating increasingly realistic synthetic images. The discriminator tries to correctly identify which images are real versus synthetic. As training progresses, the generator improves at creating realistic images (to fool the discriminator), while the discriminator becomes better at detecting subtle differences. Training reaches equilibrium when the discriminator can no longer reliably distinguish real from synthetic images (discriminator output ≈ 0.5 for synthetic images).

Several GAN variants have been developed to address specific OCT synthesis challenges: Variants such as Pix2Pix (paired translation) [63], CycleGAN (unpaired translation) [64], StyleGAN (progressive growing) [65], and PatchGAN (local coherence) [66] have been applied to tasks ranging from denoising and artifact reduction [26] to cross-modality synthesis (e.g., OCT ↔ FA) [23]. Yet, GANs still suffer from mode collapse, hallucinated features, and limited generalization, which hinder clinical adoption [14,19,67].

4.3.3. Advantages and Limitations

The adversarial training paradigm offers significant advantages: it produces sharper, more realistic images compared to alternatives like VAEs, and requires no explicit loss function for realism (the discriminator learns this automatically). However, GANs also suffer from important limitations:

- Mode collapse: The generator may learn to produce limited varieties of images, failing to capture the full diversity of the training distribution
- Training instability: The adversarial process can be difficult to balance, with one network dominating the other
- Hallucinated features: GANs may generate anatomically plausible but clinically incorrect features (discussed in detail in Section 6.5)
- Limited generalization: Models trained on one dataset/device may not generalize well to others

4.4. Beyond GANs: VAEs, Diffusion Models, and Hybrid Paradigms

While GANs have received most of the attention in ophthalmic image synthesis, other generative models are becoming increasingly relevant. VAEs provide stable training and efficient latent representations, although they often suffer from blurred outputs that limit their clinical utility [38]. Diffusion models, in contrast, iteratively denoise random noise to create highly realistic images [39]. Recent studies suggest they offer superior preservation of fine anatomical detail in retinal imaging compared with GANs [68], though their advantages come at the cost of long inference times and higher computational demands [69]. Hybrid and emerging models, including transformer-based generators [70] and GAN-diffusion or GAN-VAE hybrids [38], aim to combine stability with high fidelity. While these approaches remain underexplored in OCT, their success in other medical imaging domains indicates they are likely to play a larger role in the future.

4.5. Synergistic Roles of CNNs and GANs

CNNs and GANs frequently operate in tandem: CNNs are used to evaluate the fidelity of synthetic images, while GANs augment CNN training datasets with rare or balanced examples. Transfer learning from pre-trained CNNs can also accelerate GAN convergence and improve anatomical realism [70]. This synergy is particularly valuable in OCT applications where rare pathologies or longitudinal treatment-response data are limited.

4.6. Model Evaluation in OCT Image Synthesis

Assessing the validity of synthetic OCT requires both technical and clinical evaluation. Common quantitative metrics include Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and Fréchet Inception Distance (FID), while newer metrics such as Learned Perceptual Image Patch Similarity (LPIPS) capture perceptual realism [71]. Equally important is clinical validation: expert grading, segmentation accuracy, or diagnostic performance on tasks trained with synthetic data. Balanced reporting of both technical and clinical metrics is essential to ensure the real-world utility of generative methods.

5. Retinal Image Synthesis

5.1. Retinal Image Datasets

Progress in OCT image synthesis depends heavily on the availability and quality of datasets. Several public resources, including AROI, OPTIMA, SEED, RETOUCH, and AREDS2, provide annotated OCT scans for common conditions such as age-related macular degeneration (AMD), diabetic retinopathy, and glaucoma. OCT angiography datasets such as OCTA-500 and ROSE-1/2 have also facilitated vessel segmentation and cross-modality synthesis. Table 2 lists commonly used datasets in OCT image generation research.

Table 2. Specifications of publicly available retinal image datasets for retinal image generation via DL models.

Dataset	Imaging Modality	Dataset Characteristics	Annotation	Ref
AROI	OCT	25 image sets of neovascular-AMD patients	Manually annotated	[72]
OPTIMA	OCT	30 volumetric scans (1136 B-scans) from patients with diabetic retinopathy and AMD	Manually annotated	[73]
SEED	OCT	33,000 images from healthy and diseased eyes with a wide range of pathology	Manually annotated	[74]
RETOUCH	OCT	112 OCT volumes (11,334 B-scans)	Manually annotated for fluid detection	[75]
AREDS2	OCT	OCT volumetric images of 4203 participants with AMD	No annotation	[76]
A2A SD-OCT (Ancillary SD-OCT)	SD-OCT	OCT volumetric images (100 B-scans, 1000 A-scans per B-scan) of 312 AMD participants and 122 control group	No annotation	[77]
OCTA-500	OCTA	Over 360,000 images of 300 subjects with 6 mm × 6 mm FOV and 200 subjects with 3 mm × 3 mm FOV.	Pixel-level labels: Retinal vessel and foveal segmentation	[78]

Table 2. *Cont.*

Dataset	Imaging Modality	Dataset Characteristics	Annotation	Ref
ROSE-1	OCTA	117 OCTA images from 39 subjects (26 with Alzheimer's disease, 13 healthy controls)	Centerline-level and pixel-level annotations for vessels	[79]
ROSE-2	OCTA	112 OCTA images from 112 eyes (3x3 mm^2 centered on the fovea)	Centerline-level vessel annotations	[79]
DME-OCT (Duke dataset)	OCT	Data recorded from 16 patients with diabetic macular edema (DME) pathology	Manually annotated	[80]
UMN public database	OCT	725 images from 29 DME subjects	Manually segmented	[81]
ONH-OCT	OCT	416 ONH volumetric scans of 71 healthy and 341 patients (each volumetric scan contains 45 B-scans, 384 A-scans per B-scan), 15° × 15° scanning angle	40 scans manually annotated	[82]
Bascom Palmer Eye Institute	OCT	610 B-scans of 10 diabetic patients with image dimensions of 768 × 496 pixels	50 B-scans manually annotated	[83]

Despite these resources, important limitations remain. Most publicly available OCT datasets are small by contemporary deep learning standards, often comprising fewer than 30,000 images, which constrains the training and generalisability of high-capacity generative models. Annotation quality is heterogeneous, with many datasets providing labels for isolated features such as retinal fluid rather than comprehensive, layer-resolved retinal maps. Disease representation is also imbalanced, with a predominance of moderate to advanced pathology and relative underrepresentation of early-stage disease. For example, datasets frequently include advanced age-related macular degeneration with large drusen or established geographic atrophy, while lacking images of early drusen or nascent drusenoid pigment epithelium detachment. This imbalance limits the utility of synthetic augmentation for early detection algorithms and may contribute to diagnostic delay in clinical practice.

In addition, most public OCT datasets predominantly represent populations of European ancestry from developed healthcare systems. Underrepresentation of other ethnic and demographic groups, including Asian, African, and Indigenous populations, raises concerns regarding generalisability. Continued reliance on private datasets further limits reproducibility and external validation.

These limitations explain the growing reliance on synthetic data for augmentation, particularly when developing models for rare conditions or treatment-response prediction.

5.2. Synthetic Data as a Solution

The limitations of existing datasets have led to the development of synthetic data generation as a complementary approach in medical imaging. High-quality synthetic images can expand small datasets, thereby improving the robustness of model training. They also help to balance disease distributions, ensuring that algorithms are not biased toward more common conditions. Importantly, synthetic datasets offer privacy-preserving alternatives to real patient data, reducing concerns around confidentiality and data governance. Another advantage is the ability to generate rare pathological presentations that may otherwise be unavailable, providing valuable material for both research and educational purposes. Synthetic images can also facilitate multi-center studies by enabling data sharing without direct transfer of sensitive patient information.

Beyond these advantages, cross-modality synthesis provides a further opportunity to minimise costs and technical complexity. For example, generating OCTA or PS-OCT images directly from standard structural OCT can reduce the need for additional hardware components or dye-based procedures. Such approaches not only make advanced imaging more accessible in resource-limited settings but also accelerate the adoption of multimodal analysis pipelines without the associated infrastructure burden.

The integration of synthetic and real data represents a promising path forward in addressing current dataset limitations. By combining the scalability and flexibility of synthetic data with the authenticity of real-world imaging, researchers can maintain clinical relevance and diagnostic accuracy while overcoming many of the constraints associated with limited datasets. This capability forms the foundation for the methodological approaches described in the following sections.

5.3. Generative Models and Applications in OCT Synthesis

A wide range of generative models has been applied to OCT synthesis. While GAN-based approaches dominate, newer frameworks, which we will briefly describe in sections 5.5 and 5.6, are emerging. In this section, and for clarity, applications of GANs in OCT synthesis are grouped into four categories as shown in Figure 5. More information on the studies is provided in Table 3.

1. Image quality enhancement and structural preservation

Conditional GANs (cGANs), PatchGAN variants, and more recent super-resolution frameworks such as ESRGAN and SRFlow have been applied to correct data imbalance [67,84], improve segmentation fidelity [56], preserve layered retinal anatomy [19,84], and reduce artifacts [85]. Incorporating SSIM and Fourier-domain loss functions has further improved preservation of layered retinal anatomy [19,86].

2. Treatment outcome prediction and monitoring

Models such as cGAN, RegGAN, CycleGAN, UNIT and Pix2PixHD have been trained to predict post-treatment OCT scans in diabetic macular edema (DME) [41], AMD [87,88], and retinal vein occlusion (RVO) [89]. These methods demonstrate that synthetic OCT can approximate clinical outcomes with sufficient accuracy to assist in prognosis and patient counseling.

3. Dataset augmentation and rare disease simulation

Progressive GANs (PGGANs) and related architectures have been used to generate synthetic datasets for glaucoma [25], and retinal pathologies including choroidal neovascularisation (CNV) and DME [16]. Pix2Pix-based approaches have been employed to augment DME classifiers [24] and generate balanced training cohorts from limited datasets. Multi-GAN strategies have also been applied to AMD/DME/normal scans [20], improving classifier accuracy and robustness. These synthetic samples improve downstream CNN performance and provide rare pathological examples for education and training.

4. Cross-modality translation

Pix2Pix, CycleGAN, U-Net hybrids, and residual dense networks have enabled translation across modalities, including OCT↔FA [23], fundus→OCTA [90], OCT→OCTA with motion artifact suppression [26], OCT→PS-OCT [57,91], PS-OCT→PS-OCTA [92]. GAN-based TD→SD conversion has also been demonstrated [93]. Collectively, these cross-modality synthesis strategies reduce reliance on invasive dye-based imaging and lower costs of advanced modalities, while expanding access to functional information.

Figure 5 as a two-dimensional bubble map, was constructed to visually summarize the studies included in this review based on their relative clinical proximity and methodological novelty.

The clinical proximity and methodological novelty scores represent semi-quantitative assessments based on systematic review of study characteristics. While we applied consistent criteria across all studies (detailed in Supplementary Materials), these scores contain inherent subjectivity and have not been externally validated. This visualization should be interpreted as a conceptual framework for understanding the landscape of OCT synthesis research rather than as precise quantitative rankings.

Each study (S_i) was assigned a coordinate (C_i, N_i), where C_i reflects the degree of clinical translation (Clinical proximity) and N_i is the level of architectural innovation (Methodological novelty).

$$S_i = (C_i, N_i) \tag{1}$$

C_i was derived as a weighted composite score incorporating (i) clinical feasibility (U_i; interpretability, privacy preservation, or annotation efficiency), (ii) dataset origin and diversity (V_i; multi-center versus single-site or synthetic), and (iii) endpoint validation (E_i; e.g., expert grading, treatment prediction, or segmentation accuracy), as:

$$C_i = 0.4U_i + 0.3V_i + 0.3E_i \tag{2}$$

N_i was similarly quantified from (i) model-level innovation (M_i; e.g., baseline GAN vs. Fourier- or diffusion-augmented architectures), (ii) conditioning or architectural extensions (C_i'; e.g., multimodal inputs, spectral priors, semantic guidance), and (iii) research originality (R_i; novelty of application or first-in-domain implementation), as:

$$N_i = 0.4M_i + 0.3C_i' + 0.3R_i \tag{3}$$

The resulting coordinates were normalized to the 0–1 range to allow consistent comparison across studies. Table S1 lists definitions and scoring ranges for all quantitative variables used to derive the clinical proximity and methodological novelty indices. Bubbles were coloured according to the primary application domain (image quality enhancement: blue, treatment outcome prediction: pink, data augmentation: green, or cross-modality translation: yellow).

A detailed scoring matrix for all included studies, including the sub-component values (U_i, V_i, E_i, M_i, C_i' and R_i), weighting scheme, and rationale for each score, is provided in Supplementary Table S2. Taken together, these studies provide a foundation for identifying broader patterns and common challenges, which we discuss next.

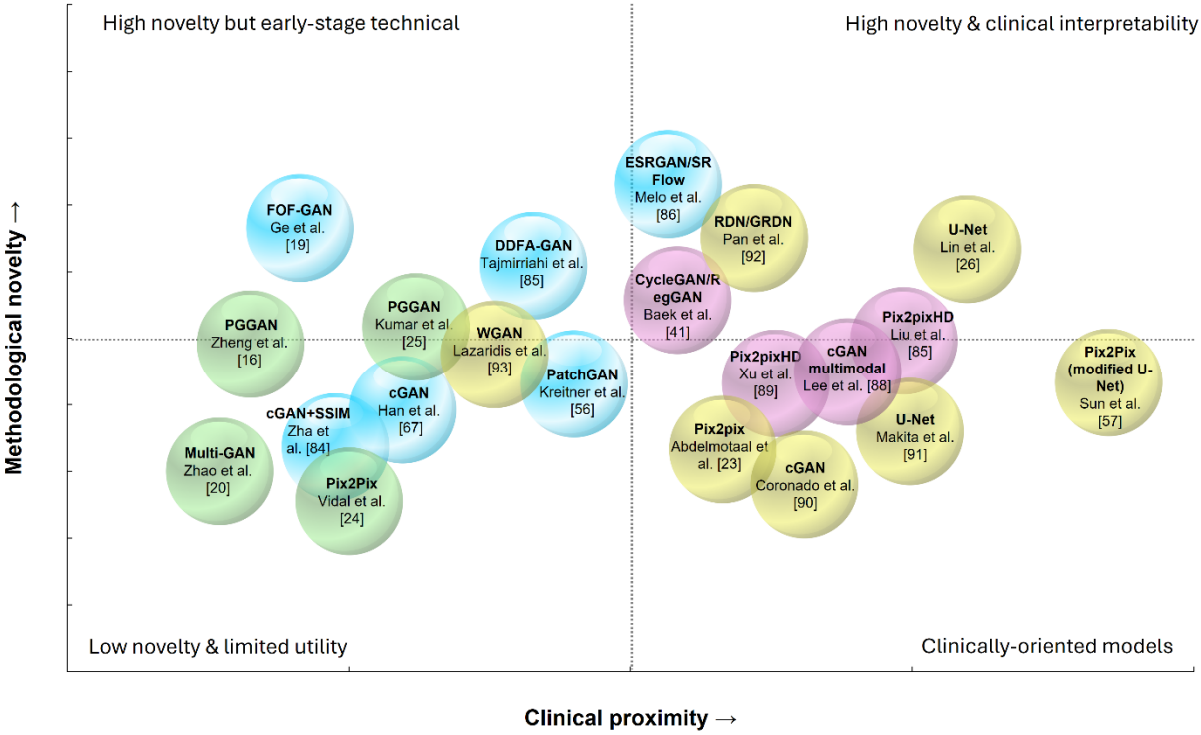

Figure 5. A bubble map of generative models used with modalities based on OCT. The vertical axis represents methodological novelty (from well-established frameworks at the bottom to more sophisticated approaches at the top), while the horizontal axis represents increasing clinical proximity (from technical demonstrations on the left to direct clinical utility on the right). Colours differentiate the models based on their application: pink: treatment outcome prediction & monitoring, blue: image quality enhancement, yellow: cross-modality translation, green: data augmentation. Scores details can be found in the supplementary information and Table S2. Note: Scoring represents semi-quantitative assessment but has not been externally validated. Bubble positions should be interpreted as approximate representations of study characteristics rather than precise measurements.

5.3.1. Trends in Using GAN Models for OCT

Across the studies summarized in Figure 5, several consistent patterns emerge. First, GAN-based models remain the dominant approach, particularly for image augmentation and cross-modality translation. In contrast, diffusion and hybrid models are almost absent despite their growing prominence in other imaging domains. Second, most studies rely on private datasets, which limits reproducibility and cross-institutional generalisability; only a handful (e.g., Kermany, Srinivasan, OCTA-500) use widely accessible datasets. Third, most validation relies on technical metrics (SSIM, PSNR, FID), with relatively few studies incorporating clinical endpoints such as diagnostic accuracy or expert grading. Finally, while applications span enhancement, prognosis, and rare disease simulation, there is little work on few-shot synthesis, multimodal integration (e.g., OCT + FA), or rigorous external testing, areas that represent important opportunities for advancing the field.

Based on Figure 5, architectures introducing spectral or Fourier-domain priors, such as FOF-GAN [19] and DDFA-GAN [85], cluster high on methodological novelty but remain earlier in clinical translation. Flow-based super-resolution alongside adversarial SR (SRFlow vs ESRGAN [86]) sit in the same innovation quadrant, emphasizing frequency-aware fidelity and perceptual realism over validated clinical utility. By contrast, application-proximal models occupy the high-proximity axis. Post-treatment prediction pipelines based on paired translation Pix2PixHD [87,89] and related adversarial translators such as CycleGAN/UNIT/Pix2PixHD/RegGAN [41], and cGAN [88] are positioned closest to clinical tasks through lesion-level grading, fluid detection, and thickness-error readouts on DME/AMD/RVO cohorts. Cross-modality U-Net/GAN hybrids such as Pix2Pix (OCT↔FA [23]), U-Net with motion-artifact suppression (OCT→OCTA [26]), cGAN (fundus→OCTA [90]), U-Net (OCT→PS-OCT [91]), modified U-Net/Pix2Pix (OCT→PS-OCT [57]) and RDN/GRDN (PS-OCT→PS-OCTA [92]), also map near this axis given expert-readability, vessel-density/Dice measures, or modality-specific

clinical surrogates. In a similar clinically leaning direction, Lazaridis et al. used standard/WGAN variants (TD-OCT→SD-OCT [93]) to bridge technology and trial design by improving signal quality and segmentation in the UKGTS dataset. A second cluster comprises augmentation-centric syntheses that are essential for training and benchmarking but are one step further from immediate clinical endpoints. These include cGAN/PatchGAN disease-aware synthesis for imbalance correction (Zha et al. [84]), therapeutically motivated augmentation on Kermany dataset (Han et al. [67]), Pix2Pix augmentation for DME with mixed public/curated sources (Vidal et al. [24]), and class-specific GANs trained on Srinivasan/Kermany variants (Zhao et al. [20]). High-fidelity PGGAN studies, such as Kumar et al. for circumpapillary/glaucoma [25] and Zheng et al. for CNV/DME/drusen/normal OCT [16], prioritize realism and educational/diagnostic training value; while not yet outcome-linked, they consistently improve downstream classifiers and specialist grading. On the analytic side of OCTA segmentation, Kreitner et al. combine ResNet and PatchGAN with public OCTA-500/ROSE-1/Giarratano resources to lift vessel-level performance [56].

Taken together, the field is moving beyond proof-of-concept realism towards outcome-aware, clinically interpretable tools that either (i) forecast therapy response [87–89], (ii) virtualize costly or invasive modalities [23,26,57,90–92], or (iii) deliver training-grade realism for robust model development [16,20,24,25,56,67,84].

These literature are also listed in Table 3 and represented in more details along the same categories mentioned earlier in Figure 5: (1) image quality enhancement and denoising (blue bubbles in Figure 5), (2) predictive modelling of treatment outcomes (pink bubbles in Figure 5), (3) data augmentation for diagnostic training (green bubbles in Figure 5) and (4) cross-modality Synthesis (yellow bubbles in Figure 5). Below, we outline these four themes in more depth.

Table 3. Advanced models for retinal image synthesis using deep neural networks. Links to the literature are provided under the authors' section (click on the author's name).

Authors	Method/Model	Imaging Modality	Application	Training Dataset
Zha et al. [84]	cGAN including a PatchGAN	OCT	Improving data imbalance issue	Private dataset
Han et al. [67]	cGAN	OCT	Image augmentation for therapeutic reasons	Kermany dataset [94]
Lazaridis et al. [93]	Standard GAN, and different WGANs	OCT (TD → SD)	Quality enhancement, glaucoma trials	UKGTS dataset
Baek et al. [41]	CycleGAN, UNIT, Pix2PixHD, and RegGAN	OCT	Prediction and visualization of long-term treatment	Private dataset (KINGFISHER study)
Lee et al. [88]	cGAN	OCT	Prediction of outcome after post-treatment	Private dataset
Liu et al. [87]	Pix2pixHD	OCT	Therapeutic tool	Private dataset
Xu et al. [89]	Pix2pixHD	OCT	Therapeutic tool	Private dataset
Vidal et al. [24]	Pix2pix	OCT	Image augmentation	Private and ImageNet dataset
Zhao et al. [20]	Three different GAN models	OCT	Image augmentation	Srinivasan dataset [95]
Kumar et al. [25]	PGGAN	OCT	Image augmentation	Private dataset
Zheng et al. [16]	PGGAN	OCT	Realistic image generation for educational and DL model training	Kermany et al. [96], and private SSH dataset
Lin et al. [26]	U-Net	OCT →OCTA	Suppression of motion artifact	Private dataset
Ge et al. [19]	FOF-GAN	OCT	Generation of higher quality images	Kermany et al. [96]
Tajmirriahi et al. [85]	DDFA-GAN	OCT	Image enhancement	Private dataset
Melo et al. [86]	ESRGAN and SRFlow network	OCT	Generation of higher quality images	DIV2K, Flickr2K and UCSD
Abdelmotaal et al. [23]	Pix2pix	OCT → FA and FA → OCT	Generation of clinically useful OCT and FA images from a single source	Private dataset
Kreitner et al. [56]	ResNet and PatchGAN	OCTA	Image augmentation for segmentation purposes	OCTA-500, ROSE-1, and Giarratano [97] dataset
Coronado et al. [90]	cGAN	Fundus → OCTA	Improve vessel segmentation to minimize manual procedures	Private dataset
Makita et al. [91]	U-Net	OCT → PS-OCT	Reduction of operational costs	Private dataset
Pan et al. [92]	RDN and GRDN	PS-OCT → PS-OCTA	Reduction of operational costs, therapeutic tool	Private dataset
Sun et al. [57]	Pix2Pix (modified U-Net)	OCT → PS-OCT	Reduction of operational costs	Private dataset

5.3.2. Image Quality and Structural Fidelity

Early GAN frameworks often struggled to preserve laminar boundaries and disease-specific features in OCT synthesis. Zha et al. [84] addressed this limitation by incorporating a SSIM loss into a cGAN, demonstrating improved structural fidelity over conventional L1 loss. Han et al. [67] applied cGANs to synthesize OCT images of retinal disease, showing that synthetic scans achieved classification accuracies above 97% across multiple pretrained models and were nearly indistinguishable from real images in interpretable saliency analyses. Lazaridis et al. [93] advanced this line of work by using ensembles of GAN variants (WGAN, perceptual loss GANs, and cycle-consistent GANs) to convert low-quality TD-OCT into SD-OCT, improving signal-to-noise ratio, RNFL segmentation, and PSNR/SSIM compared to conventional approaches.

More recently, Fourier-domain modifications have been introduced to improve anatomical detail. Ge et al. [19] proposed Fourier-FastGAN (FOF-GAN), embedding Fourier attention into the discriminator to better preserve frequency content, achieving lower FID and improved classification accuracy for DME and drusen compared with baseline FastGAN. Tajmirriahi et al. [85] further extended this by introducing a dual-discriminator Fourier acquisitive GAN (DDFA-GAN), combining spectral and Fourier domain discriminators to preserve high-frequency retinal information, with superior FID and MS-SSIM relative to DCGAN, WGAN, and LS-GAN baselines. Together, these works show a clear trajectory: from intensity-driven realism toward clinically verifiable fidelity, where loss functions and spectral priors are explicitly tuned to OCT's layered structure.

Despite improvements in perceptual quality metrics, an important limitation of perceptual-loss–driven and GAN-based image synthesis methods is the risk of introducing artificial or misleading features. These features may appear anatomically plausible yet deviate from the true underlying tissue structure, potentially confounding clinical interpretation. Mehdizadeh et al. showed that while deep feature losses can enhance perceptual sharpness in OCT denoising, certain loss configurations introduce additional textures and alter fine retinal features compared with averaged reference images, despite favorable perceptual metrics [98]. This issue is particularly critical in retinal imaging, where subtle pathological changes (e.g., early drusen, microaneurysms, intraretinal cysts) must be distinguished from processing-induced artifacts. Consequently, quantitative metrics such as PSNR and SSIM, while useful for technical comparison, cannot substitute for clinical validation by expert graders capable of identifying implausible or spurious features."

5.3.3. Predictive Synthesis for Treatment Outcomes

A second trend is the use of generative models to forecast OCT appearances after anti-VEGF therapy in DME and neovascular AMD. Baek et al. [41] compared four GAN architectures (CycleGAN, UNIT, Pix2PixHD, and RegGAN) for predicting week-52 OCT outcomes in DME, finding that RegGAN produced the most gradable images and achieved superior lesion-wise agreement for residual fluid and exudates. Lee et al. [88] employed a cGAN conditioned on OCT alone or in combination with fluorescein and indocyanine green angiography, showing that multimodal input improved detection of intraretinal and subretinal fluid. Liu et al. [87] used Pix2PixHD on a swept-source OCT dataset to predict post-treatment images in neovascular AMD, reporting that 91% of generated images were clinically interpretable and that central macular thickness could be predicted with a mean absolute error of ~26 μm. Xu et al. [89] applied a similar Pix2PixHD pipeline for retinal vein occlusion, again demonstrating strong correlation between synthetic and real post-therapeutic scans. Collectively, these studies highlight a shift: synthetic imaging is no longer confined to static augmentation but is now being used for dynamic prognosis, with demonstrated utility for simulating clinical trial outcomes and stratifying patient response.

5.3.4. Dataset Augmentation

Generative models are also widely applied for data augmentation and the creation of rare pathological presentations. Vidal et al. [24] used Pix2Pix with DenseNet-161 and VGG-19 feature extractors to generate synthetic OCT of DME, demonstrating improved classification performance when synthetic data were added to training. Zhao et al. [20] trained separate GANs on AMD, DME, and normal images, reporting high FID scores (~1) and improved classifier accuracy, sensitivity, and specificity when synthetic images were included; they also showed superiority over a VAE-based method. Kumar et al. [25] used progressive growing GANs (PGGANs) to synthesize circumpapillary OCT scans in glaucoma, finding that expert graders could not reliably distinguish synthetic from real images and that classifiers trained solely on synthetic data achieved comparable AUCs to those trained on real data. Zheng et al. [16] similarly demonstrated that PGGAN-generated OCT images of CNV, DME, drusen, and normals were nearly indistinguishable to retinal specialists, with synthetic-trained classifiers achieving only marginally lower AUC than real-trained models. Together, these studies confirm the translational and

educational value of synthetic corpora, while also highlighting differences across architectures: GANs yield sharper pathological features, whereas alternative models such as VAEs offer stability but with blurrier outputs.

5.3.5. Cross-Modality Synthesis and Functional OCT Extensions

Another frontier is cross-modality synthesis, where structural OCT is used to emulate functional modalities. Abdelmotaal et al. [23] used Pix2Pix GANs to translate between OCT thickness maps and fluorescein angiography, reporting FID-guided optimization and subjective quality ratings by experts. Coronado et al. [90] generated OCTA images from fundus photographs using a cGAN, showing superior representation of small vessels compared with vessel segmentation baselines. Lin et al. [26] proposed an attention-U-Net with a convolutional block attention module to generate OCTA from repeated structural OCT scans, achieving motion-artifact suppression and SSIM up to 0.93. Complementing these, Kreitner et al. [56] synthesised OCTA for segmentation augmentation (ResNet + PatchGAN), yielding better vessel-level performance while reducing reliance on manual labels.

PS-OCT synthesis represents a special case in our review. While most included studies focus on retinal imaging, we include non-retinal PS-OCT work due to limited retinal-specific publications and the high transferability of techniques. Makita et al. [91] synthesized DOPU maps from standard OCT, achieving F1 scores above 0.89 for RPE abnormalities, while Sun et al. [57] used Pix2Pix to generate PS-OCT retardation and DOPU images, with classifiers trained on synthetic data reaching AUCs of 0.97–0.99. Extending this, Pan et al. [92] used RDN/GRDN to translate PS-OCT→PS-OCTA, demonstrating feasibility and therapeutic utility through preservation of vascular features relevant to microvascular assessment. Melo et al. [86] further compared ESRGAN with flow-based SRFlow, with SRFlow achieving higher PSNR/SSIM and better downstream classification despite ESRGAN's perceptual sharpness.

These works collectively illustrate how "virtualized modalities" can lower cost and invasiveness, while retaining much of the clinical value of hardware-based OCTA and PS-OCT.

5.3.6. Methodological Considerations and Validation Practices

Based on Table 3, evaluation metrics for assessing the performance of the methods have gradually shifted beyond perceptual similarity indices like FID, PSNR, and SSIM toward clinically grounded measures. For example, Lazaridis et al. [93] assessed RNFL segmentation accuracy on GAN-enhanced TD-OCT, while Liu et al. [87] and Xu et al. [89] validated synthetic post-treatment OCT by comparing predicted versus measured central macular thickness, and Baek et al. [41] and Lee et al. [88] reported lesion-level sensitivity, specificity, and κ-agreement for residual fluid and exudates. Kumar et al. [25] and Zheng et al. [16] added human expert grading of synthetic versus real images, highlighting the importance of subjective validation.

Dataset limitations remain a critical bottleneck. Many studies, including Han et al. [67], Vidal et al. [24], Zhao et al. [20], and Zheng et al. [16], rely heavily on the Kermany dataset [94] or other single-center collections, raising concerns about representativeness and generalizability.

With respect to clinical translation, only a handful of models have been validated on external datasets or across devices. Kumar et al. [25] uniquely reported strong performance of synthetic-trained models on an external cohort, while most other works (e.g., Zha et al. [84], Han et al. [67]) limited testing to internal splits. This underscores a translational gap between algorithmic success and real-world deployment.

Furthermore, and as highlighted by Mehdizadeh et al. [98], GANs can generate images that score well on perceptual metrics while containing hallucinated features that would be immediately recognizable to clinicians as implausible. This discordance between technical and clinical fidelity underscores the need for systematic expert review as a mandatory component of validation, particularly for clinical translation.

Finally, there is increasing evidence for the educational potential of synthetic OCT. Zheng et al. [16] showed that specialists found synthetic OCT almost indistinguishable from real images, while Kumar et al. [25] and Vidal et al. [24] demonstrated that large synthetic datasets can enhance training for rare conditions such as glaucoma or DME. These findings suggest that synthetic OCT could become a valuable teaching tool, providing curated exposure to otherwise rare or underrepresented pathologies.

5.4. Emerging Directions in Retinal Image Synthesis

Although GAN-based models dominate current OCT synthesis research, several emerging approaches show potential to address their limitations and expand clinical utility.

- Diffusion models: By iteratively denoising random noise, diffusion frameworks achieve superior fidelity and structural preservation compared to GANs. Early studies in retinal imaging suggest that diffusion-generated

fundus and OCT images maintain fine vascular and laminar detail with greater stability, albeit at the cost of slower inference and higher computational demands. Their integration into OCT pipelines remains limited but represents a promising future direction.

- Hybrid architectures: Combining GANs with diffusion principles or transformer-based encoders offers the potential to balance stability with realism. Transformer-driven attention mechanisms can capture long-range retinal dependencies, while GAN components provide efficient sampling. Preliminary work in other medical imaging domains suggests these hybrids could mitigate mode collapse and improve anatomical consistency in OCT synthesis.
- Few-shot and semi-supervised learning: Rare retinal conditions are poorly represented in existing datasets. Few-shot GANs and semi-supervised generative frameworks provide strategies to synthesize realistic images from limited examples, enabling augmentation of underrepresented classes. These approaches may be critical for conditions such as inherited retinal dystrophies, where annotated scans are scarce.
- Federated and multi-center training: To overcome the reproducibility challenges of single-institution datasets, federated learning could allow generative models to be trained collaboratively across multiple centers without sharing raw data. This approach may reduce bias, enhance generalisability, and accelerate clinical translation.

Together, these emerging paradigms suggest that the future of OCT image synthesis will likely extend beyond traditional GANs, towards more robust, multimodal, and clinically validated hybrid models and generative systems. Building on these conceptual directions, the following section examines in detail the emerging non-GAN paradigms, including diffusion, autoencoder-based, and few-shot models, through representative studies and datasets.

5.5. Non-GAN Models: Current Implementations and Results

In the last few years, a growing body of work demonstrates the value of non-GAN paradigms and few-shot generative approaches. These methods address some of the limitations of GANs, such as unstable training, hallucination artifacts, and large data demands, while introducing complementary strengths.

5.5.1. Variational Autoencoder (VAE) Family and Diffusion-Based Models

Beyond adversarial frameworks, diffusion and autoencoder-based generative models are increasingly deployed for OCT and OCTA. A list of these studies is available in Table 4.

Table 4. Representative non-GAN and few-shot generative approaches for OCT image synthesis.

Representative Study	Model Type	Dataset Scale	Primary Clinical/Technical Endpoint
Wu Y. et al., 2024 [99]	Diffusion (DDPM; sketch-conditioned OCT)	100 OCT images	Synthetic-only training yields segmentation performance comparable to real-trained models
Huang K. et al. [100]	Diffusion (Volumetric CA-LDM)	OCTA-500	High-resolution volumetric synthesis under reduced GPU memory load
Morís D. I. et al. [101]	Diffusion (Semantic-guided LDM; OCT-A)	470 OCTA images (private)	Semantics-guided augmentation improves lesion-level discriminability for NV diagnosis
Rashid D. et al. [102]	Diffusion (LDM; OCT-A)	698 OCT-A (private)	Privacy-preserving synthetic OCT-A without rule-based simulation
Jebril H. et al. [103]	Autoencoder (VQ-VAE + AR)	DRAC; OCTA-500	Unsupervised OCTA anomaly detection (reported AUROC up to ~0.9+)
Hu D. et al. [104]	Diffusion (DDPM; denoising)	12 OCT scans (private)	Speckle suppression without paired clean targets
Aresta G. et al. [105]	Knowledge distillation (non-GAN)	3247 OCT volumes (private)	High AUC for OCT anomaly detection with interpretable maps
Yoo T. K. et al. [106]	Few-shot learning (classification)	≤20cases/class	Improved recognition of rare OCT pathologies under scarce labels
Tai C-Y. et al. [107]	Few-shot (incl. generative augmentation & attention)	10,000 per class (rare classes: 8000 generated + 2000 augmented) (from OCT2017)	Class-balancing + attention + generative augmentation improves rare-class performance
Zhou Y. et al. [108]	Few-shot (adversarial meta-learning)	492 raw OCT images (private)	Enhanced despeckling/feature preservation under limited labels
Mehta D. et al. [109]	Few-shot (vision-language prompting)	77 fundus (private) + RFMiD dataset [110]	Interpretable few-/zero-shot diagnosis; improved generalization via concept prompts

Denoising diffusion probabilistic models (DDPMs) have been used to synthesize circumpapillary OCT from coarse layer sketches, with segmentation networks trained purely on synthetics performing comparably to those trained on real data, thereby highlighting their potential for annotation-efficient workflows [99]. Cascaded amortized latent diffusion models (CA-LDMs) extend this capability to volumetric OCT, enabling high-resolution synthesis while mitigating GPU memory constraints [100]. In addition, semantic-guided generative latent diffusion augmentation has recently been proposed for improving neovascularization diagnosis in OCTA, demonstrating that clinically guided conditioning can enhance lesion-level discriminability while preserving anatomical plausibility [101]. Latent diffusion has also been applied to OCTA, removing the reliance on rule-based simulation and supporting privacy-preserving dataset curation for neurovascular research [102]. Kaplan and Lensu [111] introduced a VAE enhanced with contrastive learning to generate high-quality, disease-specific OCT images, showing that non-GAN methods can achieve clinically relevant synthesis.

In parallel, the VAE family continues to mature. Vector-quantized VAE (VQ-VAE) representations combined with autoregressive modeling have achieved unsupervised anomaly detection on en-face OCTA, reaching AUROCs of up to 0.92 [103]. Recent extensions include DDPM-based denoising of OCT B-scans under fully unsupervised settings, improving speckle suppression without requiring paired clean data [104]. Furthermore, a teacher/student knowledge distillation system trained exclusively on normal OCT volumes has shown high sensitivity for anomaly detection (volume-level AUC ≈ 0.94), while also generating interpretable anomaly maps, as an important step toward clinical transparency [105]. Collectively, these diffusion- and VAE-based approaches emphasize controllability, stability, interpretability, and data efficiency, positioning them as key complements to GAN-based synthesis in ophthalmic imaging.

5.5.2. Few-Shot GANs for Rare Ophthalmic Conditions

The Few-shot generative augmentation remains an active area of research to address data scarcity in rare ophthalmic diseases. Early frameworks demonstrated that CycleGAN-based synthesis, coupled with transfer learning, improved recognition of rare OCT pathologies such as macular telangiectasia and Stargardt disease [106]. More recently, updated pipelines have incorporated U-GAT-IT in place of CycleGAN, alongside class-balancing strategies and attention modules such as CBAM, achieving rare-class OCT classification accuracies approaching 98% [107].

Additional strategies highlight the diversity of this research direction. AMeta-FD, a meta-learning–augmented GAN framework, has been proposed for few-shot OCT despeckling, demonstrating improved restoration of image quality under limited training data [108]. Beyond purely image-based approaches, vision-language prompting has been introduced to guide interpretable few-shot retinal disease diagnosis, integrating concept-based signals from large language models to improve generalisation and transparency [109]. Together, these developments underscore that few-shot GAN strategies remain a vital complement to other generative approaches, expanding model robustness and classifier generalisability for conditions where annotated datasets remain scarce.

5.5.3. Synthesis

Taken together, these advances demonstrate that diffusion models emphasize stability, anatomical fidelity, and interpretability, autoencoder-based models provide latent control and cross-modality synthesis, and few-shot GANs directly address data scarcity in rare conditions. Importantly, evaluation should extend beyond image-centric metrics (FID, SSIM) toward clinically grounded endpoints such as retinal thickness prediction, RNFL segmentation accuracy, and expert grading concordance. By integrating these paradigms into discussions of regulatory approval, reproducibility standards, and deployment challenges, the field can move closer to real-world clinical translation.

This comparative overview of GAN and non-GAN paradigms leads naturally into a broader discussion of clinical opportunities, limitations, and priorities for future research.

6. Discussion

This review highlights the rapid evolution of DL–based methods for OCT image generation, with GANs driving much of the recent progress. Applications now extend beyond simple data augmentation to include treatment outcome prediction, cross-modality translation, and the synthesis of rare pathologies for education and research. Despite these advances, clinical translation remains constrained by methodological, technical, and regulatory challenges.

6.1. Key Benefits and Opportunities

DL–driven OCT synthesis offers multiple opportunities. Synthetic datasets can augment small collections of images, addressing class imbalance and enhancing the robustness of diagnostic models. Cross-modality translation (e.g., OCT to OCTA or FA) may reduce the need for invasive dye-based imaging and lower the costs of advanced modalities such as PS-OCT. Prognostic models that predict post-treatment OCT images could support personalised care by helping clinicians counsel patients and tailor therapeutic strategies. In education and training, synthetic images provide unlimited access to rare conditions that are seldom encountered in routine practice.

6.2. Challenges and Limitations

Despite these benefits, several obstacles still hinder widespread clinical adoption. A major concern is the risk of artifacts and hallucinations, where GANs might produce anatomically plausible but clinically misleading features. Limited datasets and imbalance further slow progress, as most studies depend on small, single-institution datasets with little diversity, which challenges reproducibility and generalizability across different populations and imaging devices. Evaluation methods also vary while technical metrics like SSIM or FID are often reported, they do not reliably indicate clinical usefulness. Few studies incorporate clinician-based grading or directly assess diagnostic outcomes. Lastly, the high computational power needed to train large models creates barriers, especially for groups working in resource-limited settings.

The predominance of private, single-institution datasets in this field (evident in Table 3) limits reproducibility and cross-institutional validation. While we included such studies when adequately described, this highlights the critical need for more public benchmark datasets to enable fair comparisons and external validation.

Population and disease stage representation: The underrepresentation of early disease stages and non-European populations in training datasets also poses significant equity and clinical utility concerns. Synthetic augmentation trained on biased datasets will inherit and potentially amplify these biases. Future work must prioritize diverse, population-representative data collection and develop techniques to assess generalization across demographic groups.”

6.3. Mitigation Strategies

Several strategies are being developed to overcome these barriers. Federated learning allows collaborative training across institutions while protecting patient privacy, thereby increasing dataset diversity and model robustness. Explainable AI frameworks can enhance transparency by visualizing how synthetic features are generated or validated, which may build clinician trust. Advances in model design, including GAN–diffusion hybrids and transformer-based architectures, are starting to improve stability and preserve fine anatomical details. At the same time, benchmarking and reproducibility standards are essential. The availability of open datasets, transparent reporting of model parameters, and public release of source code will be vital for fair comparisons, reproducibility, and regulatory approval.

6.4. Beyond GANs: Diffusion Models and Hybrid Architectures

Recent advances in generative modelling have introduced diffusion models, which iteratively refine noisy data into high-fidelity synthetic images. Unlike GANs, which rely on adversarial training and are prone to mode collapse, diffusion models learn the full data distribution and demonstrate greater training stability. In retinal and OCT imaging, they show superior preservation of fine anatomical structures, such as retinal layers and microvasculature, while reducing the risk of hallucinated features. These properties make them particularly attractive for clinical applications where anatomical fidelity is critical.

Hybrid frameworks are also emerging, combining the realism and efficiency of GANs with the stability of diffusion models, or integrating transformer-based architectures to capture long-range dependencies in OCT data. Early results in other imaging domains (MRI, CT, fundus photography) suggest such approaches can yield more reliable outputs, with strong potential for OCT synthesis. Future work should benchmark GANs, diffusion models, and hybrids systematically to determine optimal trade-offs between realism, efficiency, and clinical usability.

6.5. Critical Limitations of GANs and Clinical Implications

While GANs have enabled substantial advances in OCT image synthesis, several intrinsic limitations constrain their safe and reliable translation into clinical practice. A clear understanding of these failure modes is essential, as synthetic images that appear visually plausible may nevertheless encode systematic errors with direct implications for diagnosis, prognostication, and treatment decision-making.

Mode collapse represents a fundamental limitation of adversarial training, arising when the generator converges on a narrow subset of the true data distribution rather than modelling its full phenotypic variability. In practical terms, this results in synthetic OCT images that are individually realistic yet insufficiently diverse. For example, a GAN trained on diabetic macular oedema may repeatedly generate cysts confined to stereotyped central macular locations, while underrepresenting perifoveal or peripheral patterns that are well recognised in clinical practice. Such restricted diversity risks propagating bias into downstream diagnostic or prognostic models, which may become overly confident in common disease presentations while failing to recognise rarer but clinically important variants. From an educational perspective, exposure to synthetic datasets affected by mode collapse may also limit trainee familiarity with the full spectrum of pathological morphology encountered in real-world practice.

A second and potentially more consequential limitation is the generation of hallucinated features, structures that are anatomically plausible in isolation but clinically incorrect or physiologically impossible. Because GANs are optimised to satisfy a discriminator rather than explicit anatomical constraints, they may introduce artefactual features that evade detection by conventional image quality metrics. In OCT, this may manifest as synthetic drusen appearing in incorrect retinal layers, fluid compartments violating known anatomical boundaries, or spurious hyperreflective foci lacking any biological correlate. Such hallucinations pose a direct risk to clinical validity: algorithms trained on these images may internalise artefactual patterns as true disease markers, while clinicians exposed to synthetic datasets may encounter diagnostically misleading representations. Quantitative analyses derived from these hallucinated images, such as measurements of retinal thickness or fluid volume, may also be systematically biased. This will undermine the use of synthetic images in disease monitoring or treatment response prediction.

Limited generalisation further constrains the clinical applicability of GAN-based OCT synthesis. Models trained on data from specific devices, imaging protocols, or patient populations frequently exhibit degraded performance when applied to unseen domains. For instance, a GAN trained on spectral-domain OCT acquired using a single vendor platform may generate artefact-laden or anatomically inconsistent images when tasked with synthesising swept-source OCT data from a different system. Similarly, underrepresentation of certain age groups or ethnic populations during training may result in synthetic images that poorly reflect anatomical variation across diverse patient cohorts, thereby neglecting existing inequities in ophthalmic care. Sensitivity to imaging parameters such as scan density, field of view, and wavelength, further limits the robustness of these models in multicentre or longitudinal settings.

These challenges are compounded by the inherent instability of adversarial training. The optimisation process requires a delicate balance between generator and discriminator, and failure to achieve equilibrium may result in vanishing gradients, oscillatory behaviour, or high sensitivity to hyperparameter selection. Such instability undermines reproducibility across research groups and raises concerns regarding the reliability of synthetic image quality in safety-critical applications. In the context of clinical translation, where consistency and auditability are important, these properties represent a substantial barrier to regulatory acceptance.

A range of mitigation strategies has been proposed, including architectural refinements such as Wasserstein GANs, spectral normalisation, and progressive training schemes, as well as hybrid approaches that integrate diffusion-based generative models to improve stability. Rigorous validation frameworks incorporating expert clinician grading, external dataset testing, and performance evaluation on clinically meaningful tasks are increasingly recognised as essential safeguards. Multi-institutional and federated learning strategies may also improve generalisability by exposing models to broader population and device heterogeneity. Nevertheless, none of these approaches fully eliminate the risks inherent to GAN-based synthesis.

Taken together, these limitations underscore the need for caution when deploying GAN-generated OCT images beyond exploratory or educational contexts. Without stringent validation and transparent reporting of failure modes, synthetic OCT data risk introducing subtle but consequential errors into clinical decision-making pipelines. Future work should therefore prioritise interpretability, anatomical constraint enforcement, and clinically grounded evaluation frameworks to ensure that generative models augment, rather than compromise, patient care.

6.6. Limitation of the Current Metrics

Current evaluation practices over-rely on technical metrics that correlate imperfectly with clinical utility. As demonstrated in denoising studies [98], images with superior PSNR/SSIM may introduce artifacts that compromise diagnostic accuracy. This discrepancy between technical and clinical validity demands more rigorous expert-based validation before synthetic images can be safely deployed in clinical or educational contexts.

6.7. Future Directions

Future research must move beyond proof-of-concept demonstrations and tackle several pressing gaps. Rigorous external validation across diverse patient cohorts and imaging devices will be essential to ensure generalisability. Multimodal synthesis, integrating OCT with fundus photography, fluorescein angiography, or genetic and clinical metadata, could enable more comprehensive models of retinal disease and reduce reliance on invasive diagnostics. Few-shot and semi-supervised learning approaches will be critical for synthesising rare or under-represented conditions, such as inherited retinal dystrophies. Minimising mode collapse and hallucination risk should remain a priority, with architectural innovations such as diffusion–GAN hybrids offering promising solutions. Finally, clinical translation will depend on prospective validation studies, integration into clinical workflows, and the establishment of clear regulatory pathways to ensure safety, reproducibility, and ethical deployment. These priorities are summarised in Table 5, which outlines key areas for advancing OCT image synthesis toward safe and clinically meaningful translation.

Table 5. Key research priorities for advancing OCT image synthesis toward clinical translation.

Priority area	Rationale	Example Strategies
External validation	Ensure generalisability across populations and devices	Multi-center trials; cross-platform testing
Multimodal synthesis	Leverage complementary imaging and clinical data	OCT + fundus + FA; integration with genetic/clinical metadata
Few-shot & semi-supervised learning	Address the rarity of inherited and under-represented diseases	Few-shot GANs; semi-supervised generative frameworks
Mitigation of mode collapse	Improve stability and reduce hallucinations	Hybrid GAN–diffusion models; regularisation losses
Clinical integration & regulation	Translate into safe, trusted clinical use	Prospective studies; explainable AI; adherence to regulatory frameworks

7. Conclusions

DL–based OCT synthesis represents a rapidly advancing frontier in ophthalmic imaging. While GANs remain the primary workhorse, the emergence of diffusion models, transformer architectures, and few-shot approaches signals an important shift toward more robust and clinically relevant generative frameworks. Real-world adoption will depend not only on technical advances but also on validation in diverse populations, integration into multimodal workflows, and development of transparent regulatory pathways. With continued innovation and collaborative effort, synthetic OCT has the potential to transform ophthalmic diagnostics, research, and education, ultimately improving patient outcomes.

Supplementary Materials

The additional data and information can be downloaded at: https://media.sciltp.com/articles/others/26040211 15082876/JBiOp-25110159-SM.pdf. Table S1: Definitions and scoring ranges for all quantitative variables used to derive clinical proximity and methodological novelty indices. Table S2: Quantitative mapping of clinical proximity (C_i) and methodological novelty (N_i) for all reviewed studies. Each row lists the study reference, model type, computed indices, and contributing sub-scores ($U_i,\ V_i,\ E_i,\ M_i,\ C_i',\ R_i$). Scores were assigned on a 0–1 scale by consensus review and normalised across all included publications. Bubble positions in Figure 5—in the main manuscript - correspond directly to the ($C_i,\ N_i$) coordinates.

Author Contributions

H.A.: Conceptualization, investigation, methodology, writing—original draft preparation, visualization, funding acquisition. P.G: writing—original draft preparation. F.N.: investigation, writing—review and editing, N.F: writing—review and editing. Y.L.: writing—review and editing, M.B.: writing—review and editing, validation. B.C.: writing—review and editing, supervision. G.D.: writing—review and editing, supervision, funding acquisition. All authors have read and agreed to the published version of the manuscript.

Funding

Hadi Afsharan received funding through the Western Australia Department of Health Future Health Research and Innovation Fund (FHRI) innovation fellowship 2024 (IF2024/1) and FHRI Translational Fellowship—burden

of disease and genomics (TFBG2024/4). Parmida Ghorbanian acknowledges the University of Western Australia RTP Fees Offset and HPI scholarships.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

No new datasets were generated or analyzed during the current study.

Acknowledgments

P.G. acknowledges the financial support of the Australian Government Research Training Program (RTP) Fee Offset and Novel Heart Disease Imaging Through AI HDR Scholarships. Y.L. acknowledges the support provided through the University of Western Australia's High Degree by Research scholarships. F.N. acknowledges the financial support from the Australian Government RTP Fee Offset and RTP stipend scholarships.

Conflicts of Interest

The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies

During the preparation of this work, the authors used Anthropic Claude to review and improve the grammar of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

1. Chopra, R.; Wagner, S.K.; Keane, P.A.; et al. Optical Coherence Tomography in the 2020s—Outside the Eye Clinic. *Eye* **2021**, *35*, 236–243.
2. Ge, X. High Resolution Optical Coherence Tomography. *J. Light. Technol.* **2021**, *39*, 3824–3835.
3. Vizzeri, G. Agreement between Spectral-Domain and Time-Domain OCT for Measuring RNFL Thickness. *Br. J. Ophthalmol.* **2009**, *93*, 775–781.
4. Chen, T.C. Spectral Domain Optical Coherence Tomography: Ultra-High Speed, Ultra-High Resolution Ophthalmic Imaging. *Arch. Ophthalmol.* **2005**, *123*, 1715–1720.
5. Schmidt-Erfurth, U.; Klimscha, S.; Waldstein, S.M.; et al. A View of the Current and Future Role of Optical Coherence Tomography in the Management of Age-Related Macular Degeneration. *Eye* **2017**, *31*, 26–44. https://doi.org/10.1038/eye.2016.227.
6. Sun, Z.; Yang, D.; Tang, Z.; et al. Optical Coherence Tomography Angiography in Diabetic Retinopathy: An Updated Review. *Eye* **2021**, *35*, 149–161. https://doi.org/10.1038/s41433-020-01233-y.
7. Geevarghese, A.; Wollstein, G.; Ishikawa, H.; et al. Optical Coherence Tomography and Glaucoma. *Annu. Rev. Vis. Sci.* **2021**, *7*, 693–726. https://doi.org/10.1146/annurev-vision-100419-111350.
8. Matsunaga, D.; Yi, J.; Puliafito, C.A.; et al. OCT Angiography in Healthy Human Subjects. *Ophthalmic Surg. Lasers Imaging Retin.* **2014**, *45*, 510–515. https://doi.org/10.3928/23258160-20141118-04.
9. Cense, B.; Chen, T.C.; Park, B.H.; et al. In Vivo Birefringence and Thickness Measurements of the Human Retinal Nerve Fiber Layer Using Polarization-Sensitive Optical Coherence Tomography. *JBO* **2004**, *9*, 121–125. https://doi.org/10.1117/1.1627774.
10. Gonçalves, M.B. Image Quality Assessment of Retinal Fundus Photographs for Diabetic Retinopathy in the Machine Learning Era: A Review. *Eye* **2024**, *38*, 426–433.
11. Badar, M.; Haris, M.; Fatima, A.; et al. Application of Deep Learning for Retinal Image Analysis: A Review. *Comput. Sci. Rev.* **2020**, *35*, 100203. https://doi.org/10.1016/j.cosrev.2019.100203.
12. Masalkhi, M.; Sporn, K.; Kumar, R.; et al. Ophthalmic Image Synthesis and Analysis with Generative Adversarial Network Artificial Intelligence. *J. Digit. Imaging. Inform. Med.* **2025**, *39*, 732–75. https://doi.org/10.1007/s10278-025-01519-1.
13. Al-Absi, H.R.H.; Pai, A.; Naeem, U.; et al. DiaNet v2 Deep Learning Based Method for Diabetes Diagnosis Using Retinal Images. *Sci. Rep.* **2024**, *14*, 1595. https://doi.org/10.1038/s41598-023-49677-y.

14. Sonmez, S.C. Generative Artificial Intelligence in Ophthalmology: Current Innovations, Future Applications and Challenges. *Br. J. Ophthalmol.* **2024**, *108*, 1335–1340.
15. Veturi, Y.A. SynthEye: Investigating the Impact of Synthetic Data on Artificial Intelligence-Assisted Gene Diagnosis of Inherited Retinal Disease. *Ophthalmol. Sci.* **2023**, *3*, 100258.
16. Zheng, C.; Xie, X.; Zhou, K.; et al. Assessment of Generative Adversarial Networks Model for Synthetic Optical Coherence Tomography Images of Retinal Disorders. *Trans. Vis. Sci. Tech.* **2020**, *9*, 29. https://doi.org/10.1167/tvst.9.2.29.
17. Navaeipour, F.; Hepburn, M.S.; Li, J.; et al. In Situ Stress Estimation in Quantitative Micro-Elastography. *Biomed. Opt. Express BOE* **2024**, *15*, 3609–3626. https://doi.org/10.1364/BOE.522002.
18. Navaeipour, F.; Sanderson, R.W.; Li, J.; et al. Development of Breast-Mimicking Phantoms for Use in Optical Coherence Elastography. *JBO* **2025**, *30*, 124504. https://doi.org/10.1117/1.JBO.30.12.124504.
19. Ge, N.; Liu, Y.; Xu, X.; et al. A Fast Generative Adversarial Network for High-Fidelity Optical Coherence Tomography Image Synthesis. *Photonics* **2022**, *9*, 944. https://doi.org/10.3390/photonics9120944.
20. Zhao, M.; Lu, Z.; Zhu, S.; et al. Automatic Generation of Retinal Optical Coherence Tomography Images Based on Generative Adversarial Networks. *Med. Phys.* **2022**, *49*, 7357–7367. https://doi.org/10.1002/mp.15988.
21. Perez, L.; Wang, J. The Effectiveness of Data Augmentation in Image Classification Using Deep Learning. *arXiv* **2017**, arXiv:1712.04621.
22. Goodfellow, I. Generative Adversarial Nets. *Adv. Neural Inf. Process. Syst.* **2014**, *27*.
23. Abdelmotaal, H.; Sharaf, M.; Soliman, W.; et al. Bridging the Resources Gap: Deep Learning for Fluorescein Angiography and Optical Coherence Tomography Macular Thickness Map Image Translation. *BMC Ophthalmol.* **2022**, *22*, 355. https://doi.org/10.1186/s12886-022-02577-7.
24. Vidal, P.L.; de Moura, J.; Novo, J.; et al. Image-to-Image Translation with Generative Adversarial Networks via Retinal Masks for Realistic Optical Coherence Tomography Imaging of Diabetic Macular Edema Disorders. *Biomed. Signal Process. Control* **2023**, *79*, 104098. https://doi.org/10.1016/j.bspc.2022.104098.
25. Sreejith Kumar, A.J.; Chong, R.S.; Crowston, J.G.; et al. Evaluation of Generative Adversarial Networks for High-Resolution Synthetic Image Generation of Circumpapillary Optical Coherence Tomography Images for Glaucoma. *JAMA Ophthalmol.* **2022**, *140*, 974–981. https://doi.org/10.1001/jamaophthalmol.2022.3375.
26. Lin, Z.; Zhang, Q.; Lan, G.; et al. Deep Learning for Motion Artifact-Suppressed OCTA Image Generation from Both Repeated and Adjacent OCT Scans. *Mathematics* **2024**, *12*, 446. https://doi.org/10.3390/math12030446.
27. Yanagihara, R.T.; Lee, C.S.; Ting, D.S.W.; et al. Methodological Challenges of Deep Learning in Optical Coherence Tomography for Retinal Diseases: A Review. *Trans. Vis. Sci. Tech.* **2020**, *9*, 11. https://doi.org/10.1167/tvst.9.2.11.
28. Kugelman, J.; Alonso-Caneiro, D.; Read, S.A.; et al. A Review of Generative Adversarial Network Applications in Optical Coherence Tomography Image Analysis. *J. Optom.* **2022**, *15*, S1–S11. https://doi.org/10.1016/j.optom.2022.09.004.
29. You, A.; Kim, J.K.; Ryu, I.H.; et al. Application of Generative Adversarial Networks (GAN) for Ophthalmology Image Domains: A Survey. *Eye Vis.* **2022**, *9*, 6. https://doi.org/10.1186/s40662-022-00277-3.
30. Wang, Z.; Lim, G.; Ng, W.Y.; et al. Generative Adversarial Networks in Ophthalmology: What Are These and How Can They Be Used? *Curr. Opin. Ophthalmol.* **2021**, *32*, 459. https://doi.org/10.1097/ICU.0000000000000794.
31. Ahmed, H.; Zhang, Q.; Donnan, R.; et al. Denoising of Optical Coherence Tomography Images in Ophthalmology Using Deep Learning: A Systematic Review. *J. Imaging* **2024**, *10*, 86. https://doi.org/10.3390/jimaging10040086.
32. Li, D.; Ran, A.R.; Cheung, C.Y.; et al. Deep Learning in Optical Coherence Tomography: Where Are the Gaps? *Clin. Exp. Ophthalmol.* **2023**, *51*, 853–863. https://doi.org/10.1111/ceo.14258.
33. Hormel, T.T.; Hwang, T.S.; Bailey, S.T.; et al. Artificial Intelligence in OCT Angiography. *Prog. Retin. Eye Res.* **2021**, *85*, 100965. https://doi.org/10.1016/j.preteyeres.2021.100965.
34. Ali, M.; Ali, M.; Hussain, M.; et al. Generative Adversarial Networks (GANs) for Medical Image Processing: Recent Advancements. *Arch. Comput. Methods Eng.* **2025**, *32*, 1185–1198. https://doi.org/10.1007/s11831-024-10174-8.
35. Oulmalme, C.; Nakouri, H.; Jaafar, F.; et al. A Systematic Review of Generative AI Approaches for Medical Image Enhancement: Comparing GANs, Transformers, and Diffusion Models. *Int. J. Med. Inform.* **2025**, *199*, 105903. https://doi.org/10.1016/j.ijmedinf.2025.105903.
36. Rais, K.; Amroune, M.; Benmachiche, A.; et al. Exploring Variational Autoencoders for Medical Image Generation: A Comprehensive Study. *arXiv* **2024**, arXiv:2411.07348.
37. Kebaili, A.; Lapuyade-Lahorgue, J.; Ruan, S.; et al. Deep Learning Approaches for Data Augmentation in Medical Imaging: A Review. *J. Imaging* **2023**, *9*, 81. https://doi.org/10.3390/jimaging9040081.
38. Saeed, A.Q. Synthesizing Retinal Images Using End-To-End VAEs-GAN Pipeline-Based Sharpening and Varying Layer. *Multimed. Tools Appl.* **2024**, *83*, 1283–1307.
39. Go, S. Generation of Structurally Realistic Retinal Fundus Images with Diffusion Models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024.

40. Hatamizadeh, A.; Song, J.; Liu, G.; et al. DiffiT: Diffusion Vision Transformers for Image Generation. In *Computer Vision–ECCV 2024*; Leonardis, A., Ricci, E., Roth, S., et al., Eds.; Springer Nature: Cham, Switzerland, 2025; pp. 37–55.
41. Baek, J.; He, Y.; Emamverdi, M.; et al. Prediction of Long-Term Treatment Outcomes for Diabetic Macular Edema Using a Generative Adversarial Network. *Trans. Vis. Sci. Tech.* **2024**, *13*, 4. https://doi.org/10.1167/tvst.13.7.4.
42. Peng, J.; Xie, X.; Lu, Z.; et al. Generative Adversarial Networks Synthetic Optical Coherence Tomography Images as an Education Tool for Image Diagnosis of Macular Diseases: A Randomized Trial. *Front. Med.* **2024**, 11, 1424749. https://doi.org/10.3389/fmed.2024.1424749.
43. Mares, V.; Nehemy, M.B.; Bogunovic, H.; et al. AI-Based Support for Optical Coherence Tomography in Age-Related Macular Degeneration. *Int. J. Retin. Vitr.* **2024**, *10*, 31. https://doi.org/10.1186/s40942-024-00549-1.
44. Mokhtari, A.; Maris, B.M.; Fiorini, P.; et al. A Survey on Optical Coherence Tomography—Technology and Application. *Bioengineering* **2025**, *12*, 65. https://doi.org/10.3390/bioengineering12010065.
45. Yasuno, Y. Optical Coherence Tomography—Principles, Implementation, and Applications in Ophthalmology. *arXiv* **2022**, arXiv:2212.04380.
46. Aumann, S.; Donner, S.; Fischer, J.; et al. Optical Coherence Tomography (OCT): Principle and Technical Realization. In *High Resolution Imaging in Microscopy and Ophthalmology: New Frontiers in Biomedical Optics*; Bille, J.F., Ed.; Springer International Publishing: Cham, Switzerland, 2019; pp. 59–85. ISBN 978-3-030-16638-0.
47. Afsharan, H.; Hackmann, M.J.; Wang, Q.; et al. Polarization Properties of Retinal Blood Vessel Walls Measured with Polarization Sensitive Optical Coherence Tomography. *Biomed. Opt. Express BOE* **2021**, *12*, 4340–4362. https://doi.org/10.1364/BOE.426079.
48. Eladawi, N.; Elmogy, M.; Ghazal, M.; et al. Classification of Retinal Diseases Based on OCT Images. *FBL* **2018**, *23*, 247–264. https://doi.org/10.2741/4589.
49. Stalmans, P.; Spileers, W.; Dralands, L.; et al. The Use of Optical Coherence Tomography in Macular Diseases. *Bull. Soc. Belg. Ophtalmol.* **1999**, *272*, 15–30.
50. Spaide, R.F. Optical Coherence Tomography Angiography. *Prog. Retin. Eye Res.* **2018**, *64*, 1–55.
51. Boer, J.F.; Hitzenberger, C.K.; Yasuno, Y.; et al. Polarization Sensitive Optical Coherence Tomography—A Review. *Biomed. Opt. Express* **2017**, *8*, 1838–1873.
52. Zhou, X.; Maloufi, S.; Louie, D.C.; et al. Investigating the Depolarization Property of Skin Tissue by Degree of Polarization Uniformity Contrast Using Polarization-Sensitive Optical Coherence Tomography. *Biomed. Opt. Express BOE* **2021**, *12*, 5073–5088. https://doi.org/10.1364/BOE.424709.
53. Cense, B.; Miller, D.T.; King, B.J.; et al. Measuring Polarization Changes in the Human Outer Retina with Polarization-Sensitive Optical Coherence Tomography. *J. Biophotonics* **2018**, *11*, e201700134. https://doi.org/10.1002/jbio.201700134.
54. Afsharan, H.; Silva, D.; Joo, C.; et al. Non-Invasive Retinal Blood Vessel Wall Measurements with Polarization-Sensitive Optical Coherence Tomography for Diabetes Assessment: A Quantitative Study. *Biomolecules* **2023**, *13*, 1230. https://doi.org/10.3390/biom13081230.
55. Afsharan, H.; Anilkumar, V.; Silva, D.; et al. Hypertension-Associated Changes in Retinal Blood Vessel Walls Measured in Vivo with Polarization-Sensitive Optical Coherence Tomography. *Opt. Lasers Eng.* **2024**, *172*, 107838. https://doi.org/10.1016/j.optlaseng.2023.107838.
56. Kreitner, L.; Paetzold, J.C.; Rauch, N.; et al. Synthetic Optical Coherence Tomography Angiographs for Detailed Retinal Vessel Segmentation Without Human Annotations. *IEEE Trans. Med. Imaging* **2024**, *43*, 2061–2073. https://doi.org/10.1109/TMI.2024.3354408.
57. Sun, Y.; Wang, J.; Shi, J.; et al. Synthetic Polarization-Sensitive Optical Coherence Tomography by Deep Learning. *NPJ Digit. Med.* **2021**, *4*, 105. https://doi.org/10.1038/s41746-021-00475-8.
58. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; et al. The PRISMA 2020 Statement: An Updated Guideline for Reporting Systematic Reviews. *BMJ* **2021**, *372*, n71. https://doi.org/10.1136/bmj.n71.
59. Routray, S.; Ray, A.K.; Mishra, C.; et al. Analysis of Various Image Feature Extraction Methods against Noisy Image: SIFT, SURF and HOG. In Proceedings of the 2017 Second International Conference on Electrical, Computer and Communication Technologies (ICECCT), Coimbatore, India, 22–24 February 2017; pp. 1–5.
60. Makandar, A.; Wangi, K. Comparison and Analysis of Different Feature Extraction Methods versus Noisy Images. *Int. J. Comput. Appl.* **2022**, *184*, 45–49.
61. Herath, H.M.S.S.; Yasakethu, S.L.P.; Madusanka, N.; et al. Comparative Analysis of Deep Learning Architectures for Macular Hole Segmentation in OCT Images: A Performance Evaluation of U-Net Variants. *J. Imaging* **2025**, *11*, 53. https://doi.org/10.3390/jimaging11020053.
62. Zhang, L.; Xu, Z.; Barnes, C.; et al. Perceptual Artifacts Localization for Image Synthesis Tasks. In Proceedings of the 2023 IEEE/CVF International Conference on Computer Vision (ICCV), Paris, France, 1–6 October 2023; pp. 7545–7556.
63. Wang, T.-C. High-Resolution Image Synthesis and Semantic Manipulation with Conditional Gans. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–22 June 2018.

64. Zhu, J.-Y. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017.
65. Karras, T.; Laine, S.; Aila, T. A Style-Based Generator Architecture for Generative Adversarial Networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019.
66. Isola, P. Image-to-Image Translation with Conditional Adversarial Networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017.
67. Han, K.; Yu, Y.; Lu, T.; et al. Transfer Learning and Interpretable Analysis-Based Quality Assessment of Synthetic Optical Coherence Tomography Images by CGAN Model for Retinal Diseases. *Processes* **2024**, *12*, 182.
68. Sordo, Z.; Chagnon, E.; Hu, Z.; et al. Synthetic Scientific Image Generation with VAE, GAN, and Diffusion Model Architectures. *J. Imaging* **2025**, *11*, 252. https://doi.org/10.3390/jimaging11080252.
69. Vivekananthan, S. Comparative Analysis of Generative Models: Enhancing Image Synthesis with VAEs, GANs, and Stable Diffusion. Available online: https://arxiv.org/html/2408.08751v1 (accessed on 19 September 2025).
70. Kamran, S.A. Vtgan: Semi-Supervised Retinal Image Synthesis and Disease Prediction Using Vision Transformers. In Proceedings of the IEEE/CVF international Conference on Computer Vision, Montreal, QC, Canada, 10–17 October 2021.
71. Dohmen, M.; Truong, T.; Baltruschat, I.M.; et al. Five Pitfalls When Assessing Synthetic Medical Images with Reference Metrics. In *Deep Generative Models*; Mukhopadhyay, A., Oksuz, I., Engelhardt, S., et al., Eds.; Springer Nature: Cham, Switzerland, 2025; pp. 150–159.
72. Melinščak, M. Annotated retinal optical coherence tomography images (AROI) database for joint retinal layer and fluid segmentation. *Autom. Časopis Za Autom. Mjer. Elektron. Računarstvo I Komun.* **2021**, *62*, 375–385.
73. Oguz, I. Graph-Based Retinal Fluid Segmentation from OCT Images. *Proceeding Optima Chall. MICCAI* **2015**.
74. Nawaz, M. Unraveling the Complexity of Optical Coherence Tomography Image Segmentation Using Machine and Deep Learning Techniques: A Review. *Comput. Med. Imaging Graph.* **2023**, *108*, 102269.
75. Bogunović, H.; Venhuizen, F.; Klimscha, S.; et al. RETOUCH: The Retinal OCT Fluid Detection and Segmentation Benchmark and Challenge. *IEEE Trans. Med. Imaging* **2019**, *38*, 1858–1874. https://doi.org/10.1109/TMI.2019.2901398.
76. Chew, E.Y.; The Age-Related Eye Disease Study 2 (AREDS2) Research Group; et al. The Age-Related Eye Disease Study 2 (AREDS2): Study Design and Baseline Characteristics (AREDS2 Report Number 1). *Ophthalmology* **2012**, *119*, 2282–2289.
77. Leuschen, J.N. Spectral-Domain Optical Coherence Tomography Characteristics of Intermediate Age-Related Macular Degeneration. *Ophthalmology* **2013**, *120*, 140–150.
78. Li, M. Ipn-v2 and octa-500: Methodology and dataset for retinal image segmentation. *arXiv* **2020**, *5*, 16. arXiv:2012.07261.
79. Ma, Y. ROSE: A Retinal OCT-Angiography Vessel Segmentation Dataset and New Model. *IEEE Trans. Med. Imaging* **2020**, *40*, 928–939.
80. Chiu, S.J. Kernel Regression Based Segmentation of Optical Coherence Tomography Images with Diabetic Macular Edema. *Biomed. Opt. Express* **2015**, *6*, 1172–1194.
81. Rashno, A. Fully Automated Segmentation of Fluid/Cyst Regions in Optical Coherence Tomography Images with Diabetic Macular Edema Using Neutrosophic Sets and Graph Algorithms. *IEEE Trans. Biomed. Eng.* **2018**, *65*, 989–1001.
82. Gawlik, K. Active Contour Method for ILM Segmentation in ONH Volume Scans in Retinal OCT. *Biomed. Opt. Express* **2018**, *9*, 6497–6518.
83. Tian, J. Performance Evaluation of Automated Segmentation Software on Optical Coherence Tomography Volume Data. *J. Biophotonics* **2016**, *9*, 478–489.
84. Zha, X.; Shi, F.; Ma, Y.; et al. Generation of Retinal OCT Images with Diseases Based on cGAN. In Proceedings of the Medical Imaging 2019: Image Processing, SPIE, San Diego, CA, USA, 15 March 2019; Volume 10949, pp. 544–549.
85. Tajmirriahi, M.; Kafieh, R.; Amini, Z.; et al. A Dual-Discriminator Fourier Acquisitive GAN for Generating Retinal Optical Coherence Tomography Images. *IEEE Trans. Instrum. Meas.* **2022**, *71*, 1–8. https://doi.org/10.1109/TIM.2022.3189735.
86. Melo, T.; Cardoso, J.; Carneiro, Â.; et al. OCT Image Synthesis through Deep Generative Models. In Proceedings of the 2023 IEEE 36th International Symposium on Computer-Based Medical Systems (CBMS), L'Aquila, Italy, 22–24 June 2023; pp. 561–566.
87. Liu, Y.; Yang, J.; Zhou, Y.; et al. Prediction of OCT Images of Short-Term Response to Anti-VEGF Treatment for Neovascular Age-Related Macular Degeneration Using Generative Adversarial Network. *Br. J. Ophthalmol.* **2020**, *104*, 1735–1740. https://doi.org/10.1136/bjophthalmol-2019-315338.
88. Lee, H.; Kim, S.; Kim, M.A.; et al. Post-treatment prediction of optical coherence tomography using a conditional generative adversarial network in age-related macular degeneration. *Retina* **2021**, *41*, 572. https://doi.org/10.1097/IAE.0000000000002898.
89. Xu, F.; Yu, X.; Gao, Y.; et al. Predicting OCT Images of Short-Term Response to Anti-VEGF Treatment for Retinal Vein Occlusion Using Generative Adversarial Network. *Front. Bioeng. Biotechnol.* **2022**, *10*, 914964. https://doi.org/10.3389/fbioe.2022.914964.

90. Coronado, I.; Pachade, S.; Trucco, E.; et al. Synthetic OCT-A Blood Vessel Maps Using Fundus Images and Generative Adversarial Networks. *Sci. Rep.* **2023**, *13*, 15325. https://doi.org/10.1038/s41598-023-42062-9.
91. Makita, S.; Miura, M.; Azuma, S.; et al. Synthesizing the Degree of Polarization Uniformity from Non-Polarization-Sensitive Optical Coherence Tomography Signals Using a Neural Network. *Biomed. Opt. Express BOE* **2023**, *14*, 1522–1543. https://doi.org/10.1364/BOE.482199.
92. Pan, M.; Wang, Y.; Gong, P.; et al. Feasibility of Deep Learning-Based Polarization-Sensitive Optical Coherence Tomography Angiography for Imaging Cutaneous Microvasculature. *Biomed. Opt. Express BOE* **2023**, *14*, 3856–3870. https://doi.org/10.1364/BOE.488822.
93. Lazaridis, G.; Lorenzi, M.; Ourselin, S.; et al. Improving Statistical Power of Glaucoma Clinical Trials Using an Ensemble of Cyclical Generative Adversarial Networks. *Med. Image Anal.* **2021**, *68*, 101906. https://doi.org/10.1016/j.media.2020.101906.
94. Kermany, D.; Zhang, K.; Goldbaum, M.; et al. Large Dataset of Labeled Optical Coherence Tomography (Oct) and Chest x-Ray Images. Available online: https://www.kaggle.com/datasets/anirudhcv/labeled-optical-coherence-tomography-oct (accessed on 19 September 2025).
95. Srinivasan, P.P.; Kim, L.A.; Mettu, P.S.; et al. Fully Automated Detection of Diabetic Macular Edema and Dry Age-Related Macular Degeneration from Optical Coherence Tomography Images. *Biomed. Opt. Express BOE* **2014**, *5*, 3568–3577. https://doi.org/10.1364/BOE.5.003568.
96. Kermany, D.S.; Goldbaum, M.; Cai, W.; et al. Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning. *Cell* **2018**, *172*, 1122–1131.e9. https://doi.org/10.1016/j.cell.2018.02.010.
97. Giarratano, Y.; Bianchi, E.; Gray, C.; et al. Automated Segmentation of Optical Coherence Tomography Angiography Images: Benchmark Data and Clinically Relevant Metrics. *Trans. Vis. Sci. Tech.* **2020**, *9*, 5. https://doi.org/10.1167/tvst.9.13.5.
98. Mehdizadeh, M.; MacNish, C.; Xiao, D.; et al. Deep Feature Loss to Denoise OCT Images Using Deep Neural Networks. *JBO* **2021**, *26*, 046003. https://doi.org/10.1117/JBO.26.4.046003.
99. Wu, Y.; He, W.; Eschweiler, D.; et al. Retinal OCT Synthesis with Denoising Diffusion Probabilistic Models for Layer Segmentation. In Proceedings of the 2024 IEEE International Symposium on Biomedical Imaging (ISBI), Athens, Greece, 27–30 May 2024.
100. Huang, K.; Ma, X.; Zhang, Y.; et al. Memory-Efficient High-Resolution OCT Volume Synthesis with Cascaded Amortized Latent Diffusion Models. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2024*; Linguraru, M.G., Dou, Q., Feragen, A., et al., Eds.; Springer Nature: Cham, Switzerland, 2024; pp. 478–487.
101. Morís, D.I.; de Moura, J.; Carmona, E.J.; et al. Semantic-Guided Generative Latent Diffusion Augmentation Approaches for Improving the Neovascularization Diagnosis in OCT-A Imaging. *Pattern Recognit. Lett.* **2025**, *189*, 31–37. https://doi.org/10.1016/j.patrec.2025.01.003.
102. Rashid, D.; Giarratano, Y.; MacGillivray, T.; et al. Using Latent Diffusion Models to Generate Synthetic OCTA Images. *Investig. Ophthalmol. Vis. Sci.* **2024**, *65*, 3741.
103. Jebril, H.; Esengönül, M.; Bogunović, H.; et al. Anomaly Detection in Optical Coherence Tomography Angiography (OCTA) with a Vector-Quantized Variational Auto-Encoder (VQ-VAE). *Bioengineering* **2024**, *11*, 682. https://doi.org/10.3390/bioengineering11070682.
104. Hu, D.; Tao, Y.K.; Oguz, I.; et al. Unsupervised Denoising of Retinal OCT with Diffusion Probabilistic Model. *SPIE* **2022**, *12032*, 25–34.
105. Aresta, G.; Araújo, T.; Schmidt-Erfurth, U.; et al. Anomaly Detection in Retinal OCT Images With Deep Learning-Based Knowledge Distillation. *Transl. Vis. Sci. Technol.* **2025**, *14*, 26. https://doi.org/10.1167/tvst.14.3.26.
106. Yoo, T.K.; Choi, J.Y.; Kim, H.K.; et al. Feasibility Study to Improve Deep Learning in OCT Diagnosis of Rare Retinal Diseases with Few-Shot Classification. *Med. Biol. Eng. Comput.* **2021**, *59*, 401–415. https://doi.org/10.1007/s11517-021-02321-1.
107. Tai, C.-Y.; Chen, C.-W.; Wu, C.-C.; et al. Improvement Strategies for Few-Shot Learning in OCT Image Classification of Rare Retinal Diseases. In Proceedings of the 2025 IEEE 14th Global Conference on Consumer Electronics (GCCE), Osaka, Japan, 23–26 September 2025.
108. Zhou, Y.; Peng, T.; Ahmed, T.S.; et al. AMeta-FD: Adversarial Meta-Learning for Few-Shot Retinal OCT Image Despeckling. *Comput. Med. Imaging Graph.* **2025**, *124*, 102597. https://doi.org/10.1016/j.compmedimag.2025.102597.
109. Mehta, D.; Jiang, Y.; Jan, C.; et al. Interpretable Few-Shot Retinal Disease Diagnosis with Concept-Guided Prompting of Vision-Language Models. In *Information Processing in Medical Imaging*; Oguz, I., Zhang, S., Metaxas, D.N., Eds.; Springer Nature: Cham, Switzerland, 2026; pp. 263–277.
110. Pachade, S.; Porwal, P.; Thulkar, D.; et al. Retinal Fundus Multi-Disease Image Dataset (RFMiD): A Dataset for Multi-Disease Detection Research. *Data* **2021**, *6*, 14. https://doi.org/10.3390/data6020014.
111. Kaplan, S.; Lensu, L. Contrastive Learning for Generating Optical Coherence Tomography Images of the Retina. In *Simulation and Synthesis in Medical Imaging*; Zhao, C., Svoboda, D., Wolterink, J.M., et al., Eds.; Springer International Publishing: Cham, Switzerland, 2022; pp. 112–121.