

000 001 002 003 004 005 FUNKAN: FUNCTIONAL KOLMOGOROV-ARNOLD 006 NETWORK FOR MEDICAL IMAGE SEGMENTATION 007 008 009

010 **Anonymous authors**
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

Paper under double-blind review

ABSTRACT

Medical image segmentation is pivotal for clinical diagnosis, however, remains challenging due to complex anatomies and imaging artifacts. While deep learning offers powerful solutions, prevailing architectures lack inherent interpretability and often rely on empirically designed components. Kolmogorov-Arnold networks provide a mathematically interpretable alternative but fail to preserve the spatial structure of visual data, as they process flattened feature vectors. To bridge this gap, we introduce Functional Kolmogorov-Arnold Network (FunKAN), a novel framework that generalizes the Kolmogorov-Arnold theorem to functional spaces. FunKAN parametrizes its inner functions via truncated spectral expansion over Hermite basis functions, enabling direct processing of 2D feature maps within a theoretically grounded, interpretable design. Leveraging this, we integrate FunKAN into the U-shaped architecture, yielding a new state-of-the-art segmentation model across diverse medical imaging modalities. Extensive benchmarks on BUSI (ultrasound), GlaS (histology), and CVC-ClinicDB (colonoscopy) datasets show that U-FunKAN outperforms strong baselines (U-Net, KAN, Mamba), achieving IoU and F1-score improvement and superior efficiency in terms of Gflops. Our work unites theoretical function approximation and practical medical image analysis, offering the novel state-of-the-art solution for clinical applications.

1 INTRODUCTION

Computer-aided diagnosis systems now constitute essential components of the contemporary medical imaging infrastructure, addressing critical issues such as rising diagnostic workloads and interpreter variability (Kadhim et al., 2022). Nevertheless, their clinical utility is often limited by segmentation accuracy. Although deep learning has revolutionized medical image analysis, prevailing neural architectures frequently rely on empirically derived components that lack theoretical justification and exhibit limited generalization across imaging modalities (Borys et al., 2023). To address these gaps, this research aims to develop a novel, theoretically-grounded deep learning architecture for cancer segmentation across diverse medical imaging contexts.

Extensive clinical evidence confirms that early-stage disease detection and subsequent diagnostic confirmation, whether through ultrasound imaging, histopathological analysis or colonoscopy, correlate strongly with enhanced long-term survival probabilities (Abhisheka et al., 2023). Diagnostic accuracy, however, is compromised by an exponential increase in imaging examinations coupled with a critical shortage of trained specialists, including radiologists and pathologists. Indeed, the World Health Organization (WHO) reported that in 2024 breast cancer affected 2.3 million women worldwide annually, resulting in 670000 deaths. This disease can develop at any age after puberty. Breast cancer outcomes exhibit a significant association with socioeconomic development levels: in very high-HDI nations women face a lifetime breast cancer incidence of 1 in 12 and a mortality rate of 1 in 71. In stark contrast, low-HDI countries demonstrate both a lower incidence rate of 1 in 27 and disproportionately higher mortality 1 in 48, highlighting the substantial deficiencies in early detection capabilities and therapeutic accessibility.

In light of the advancements of Kolmogorov-Arnold networks (KANs) (Liu et al., 2024b) for both medical image enhancement (Penkin & Krylov, 2025) and segmentation (Li et al., 2025), we propose a Functional Kolmogorov-Arnold Network (FunKAN) – a novel extension of the original KANs

054 to better address fundamental image processing requirements. While the original theorem (Kolmogorov, 1957) applies to continuous functionals $f(x_1, \dots, x_n)$ on \mathbb{R}^n , we hypothesize its generalization to continuous functionals $f(\chi_1, \dots, \chi_n)$ on H^n , where each χ_i states for an element from a Hilbert space H . The proposed functional extension enables representation-based feature extraction in a natural way for high-dimensional image latents by obviating feature flattening and considering each 2D feature map χ_i as an element of an underlying Hilbert space H , viewed on a spatial grid $h \times w$. Thus, the proposed approach preserves an intrinsic structure of imaging data and establishes a principled connection between the classical approximation theory and the contemporary deep learning approaches for image analysis. To ensure reproducible research, we release entire codebase on GitHub featuring PyTorch Lightning for modularity, Ruff for code quality enforcement, YAML-based configuration system for experiments management.

065 Our contributions can be summarized as follows:

- 067 **Theoretical contribution:** We propose an extension of the Kolmogorov-Arnold theorem onto functional spaces.
- 069 **Empirical validation:** We introduce Functional Kolmogorov-Arnold Network (FunKAN).
- 071 **State-of-the-art segmentation quality:** U-FunKAN achieves state-of-the-art segmentation accuracy across three distinct medical imaging modalities:
 - 072 • breast ultrasound, BUSI dataset (Al-Dhabyani et al., 2020),
 - 073 • histological gland structures, Glas dataset (Valanarasu et al., 2021),
 - 074 • colonoscopy polyp detection, CVC-ClinicDB dataset (Bernal et al., 2015).

076 2 RELATED WORK

078 2.1 MEDICAL IMAGE SEGMENTATION

080 The advent of deep learning has driven substantial progress in medical image analysis, particularly in the task of semantic segmentation. This technology has transitioned from a research curiosity to 081 an essential tool for the automated and precise delineation of anatomical structures and pathological 082 regions. Driven initially by convolutional neural networks (CNNs), like U-Net, and more recently 083 by vision transformers, these models have demonstrated remarkable capabilities in tumor detection 084 in oncology.

086 U-Net (Ronneberger et al., 2015) established a foundational encoder-decoder framework with skip 087 connections, facilitating an accurate localization through the integration of high-level semantic 088 information with low-level spatial details. Its widespread adoption in medical image analysis is largely 089 attributable to its robust performance with limited training data. However, the original U-Net archi- 090 tecture exhibits limitations in modeling long-range spatial dependencies and preserving fine struc- 091 tural details.

092 Early architectural innovations addressed these limitations through several key developments. At- 093 tention U-Net (Oktay et al., 2018) enhances feature selectivity through attention gates in skip 094 connections, dynamically emphasizing salient features, useful for a specific task, while suppressing 095 irrelevant ones. U-Net++ (Zhou et al., 2018) improves feature fusion through nested, dense skip 096 pathways, reducing the semantic gap between encoder and decoder features. By aggregating fea- 097 tures across multiple scales U-Net++ enhances segmentation quality for anatomically irregular tar- 098 gets (e.g., infiltrating tumor margins), while incurring a greater computational overhead.

099 Emerging hybrid architectures have introduced novel computational paradigms, like U-Mamba (Ma 100 et al., 2024) and U-KAN (Li et al., 2025). U-Mamba integrates Mamba into the U-Net architec- 101 ture to capture long-range dependencies with linear computational complexity, making it partic- 102 ularly suitable for high-resolution medical imaging. U-KAN substitutes convolutional layers in 103 the backbone with Kolmogorov-Arnold-motivated adaptive activation functions. The spline-based 104 parameterization of these activations enables more accurate modeling of the complex biological 105 morphologies, particularly irregular tumor margins and vascular networks, while simultaneously 106 addressing the spectral bias (Rahaman et al., 2019) inherent in ReLU-based neural networks. How- 107 ever, the proposed KAN-based backbone processes spatial feature maps as unstructured coordinate 108 collections, thereby ignoring the locality priors essential for image representation. MedKAN (Yang

et al., 2025) and UKAGNet (Drokin, 2024) partially mitigate this limitation through a hybrid convolutional approach, combining adaptive spline-based nonlinearities with spatial inductive biases. However, MedKAN remains constrained by its classification design, lacking any proven generalizations onto medical image enhancement and segmentation pipelines. Whereas UKAGNet does not go beyond the original Kolmogorov-Arnold theorem to adapt the concept further for image-to-image processing pipelines.

2.2 KOLMOGOROV-ARNOLD NETWORKS

Foundational mathematical theorems provide a powerful starting point for constructing new deep learning architectures, enabling advanced modeling of complex data relationships (Li et al., 2023). A prime example is Kolmogorov-Arnold network (Liu et al., 2024b), which implements the theoretical framework of the Kolmogorov-Arnold theorem through adaptive B-spline embeddings.

The Kolmogorov-Arnold theorem (Kolmogorov, 1957) states that if $f: [0, 1]^n \rightarrow \mathbb{R}$ is a multivariate continuous function, then it can be written as a finite composition of continuous functions of a single variable and the binary operation of addition:

$$f(x_1, \dots, x_n) = \sum_{j=1}^{2n+1} \zeta_j \left(\sum_{i=1}^n \phi_{ji}(x_i) \right), \quad (1)$$

where $\phi_{ji}: [0, 1] \rightarrow \mathbb{R}$ and $\zeta_j: \mathbb{R} \rightarrow \mathbb{R}$ – continuous inner functions of a single variable. The theorem provides another justification that neural networks of sufficient depth and width are capable of forming dense subsets in the space of continuous functions defined over compact domains (Cybenko, 1989).

While the Kolmogorov-Arnold representation theorem offered a theoretically appealing reduction of high-dimensional function approximation to learning univariate functions, the pathological non-smoothness or even fractal character of its inner functions had severely constrained its applications. In 2024 the authors Liu et al. (2024b) presented Kolmogorov-Arnold network, relaxing the original theorem’s constraints while preserving its fundamental principles. Unlike the classical representation limited to two nonlinear layers with $(2n + 1)$ hidden terms, their architecture permits arbitrary width and depth, leading to the modern differentiable KAN definition:

$$KAN(\mathbf{x}) = (\Phi_L \circ \Phi_{L-1} \circ \dots \circ \Phi_1)(\mathbf{x}), \quad (2)$$

where $\{\Phi_l\}_{l=1}^L$ – the Kolmogorov-Arnold layers, defined as:

$$\begin{bmatrix} x_{l+1,1} \\ x_{l+1,2} \\ \vdots \\ x_{l+1,m} \end{bmatrix} = \begin{bmatrix} \phi_{l,11} & \phi_{l,12} & \dots & \phi_{l,1n} \\ \phi_{l,21} & \phi_{l,22} & \dots & \phi_{l,2n} \\ \vdots & \vdots & \vdots & \vdots \\ \phi_{l,m1} & \phi_{l,m2} & \dots & \phi_{l,mn} \end{bmatrix} \begin{bmatrix} x_{l,1} \\ x_{l,2} \\ \vdots \\ x_{l,n} \end{bmatrix}, \quad (3)$$

$$x_{l+1,j} = \sum_{i=1}^n \phi_{l,ji}(x_{l,i}), \quad (4)$$

where $\phi_{l,ji}: \mathbb{R} \rightarrow \mathbb{R}$ – continuous inner functions, parameterized in a smooth differentiable manner by B-splines.

Recent KAN architectures (Seydi, 2024) have improved efficiency by replacing B-splines with Gaussian radial basis functions (RBFs), resulting in FastKAN (Li, 2024). In ChebyKAN (SS et al., 2024) the authors employ Chebyshev polynomials as a complete orthogonal system to substitute B-splines, achieving enhanced training stability.

Although Kolmogorov-Arnold networks demonstrate powerful capabilities in multivariate function approximation, their naive application to image processing is fundamentally limited. The limitation stems from KANs’ treatment of inputs as scalars, thereby ignoring the essential two-dimensional geometric structure inherent in visual data. We resolve this limitation by formulating a functional-space generalization of the Kolmogorov-Arnold theorem, considering each feature map as an element of an underlying Hilbert space H . This theoretical advancement motivates our functional Kolmogorov-Arnold network, seamlessly suitable for image processing pipelines by incorporating spatial awareness.

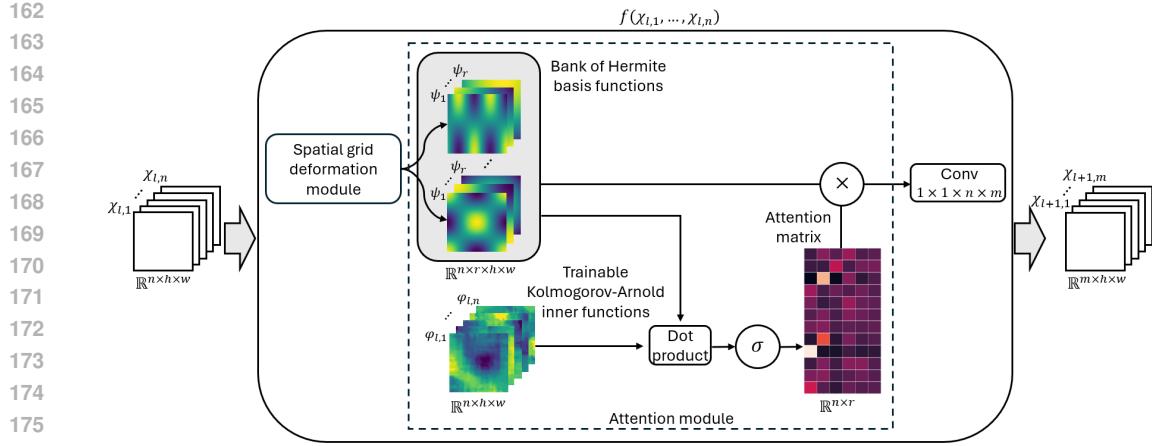


Figure 1: Architecture of Functional Kolmogorov-Arnold Network (FunKAN). The network implements trainable inner functions $\{\varphi_{l,i}\}_{i=1}^n$ through Fourier decompositions over the basis Hermite functions $\{\psi_k\}_{k=1}^r$, where each function is visualized on $h \times w$ spatial grid matching the input feature dimensions, and each decomposition is defined by normalized Fourier coefficients stored in the rows of the attention matrix.

3 METHOD

This work introduces a novel neural architecture termed Functional Kolmogorov-Arnold Network (FunKAN), illustrated in Fig. 1. The model is constructed upon a theoretical foundation – a proposed functional generalization of the classical Kolmogorov-Arnold theorem. While the original theorem provides representation for multivariate functions on real-valued inputs, our generalization extends this concept to continuous functionals operating on elements of a Hilbert space:

Statement 3.1 *If f is a continuous functional on H^n , then it can be represented as a composition of linear continuous functionals from the dual space H^* , continuous functions of a single variable and the binary operation of addition:*

$$f(\chi_1, \dots, \chi_n) \rightsquigarrow \sum_j \zeta_j \left(\sum_i \varphi_{ji}(\chi_i) \right),$$

where H is a Hilbert space, $\chi_i \in H$, $\varphi_{ji} \in H^*$ and $\zeta_j: \mathbb{R} \rightarrow \mathbb{R}$.

The proposed functional extension of the Kolmogorov-Arnold theorem hypothesizes that continuous operators on H^n , modeling activation mappings between feature spaces, may be approximated by functionals from the dual space H^* .

To construct a computationally tractable model, we ground our approach in the Riesz representation theorem. This theorem establishes an isomorphism between a Hilbert space H and its dual space H^* , allowing us to replace the dual element $\varphi_{l,ji}$ with its corresponding representative $\varphi_{l,j}(\chi_{l,i})$ in H . We then parameterize these inner functions in a differentiable manner by expressing them via a Fourier expansion truncated to the first r Hermite functions $\{\psi_k\}_{k=1}^r$. Using the fact that Hermite functions are the eigenfunctions of the integral Fourier transform in $L_2(\mathbb{R})$ (Titchmarsh, 1948), the spectral truncation strategy retains the most informative modes and mirrors the frequency-truncation principle in Fourier neural operators (Li et al., 2023). So, the next-layer feature map is computed as:

$$\chi_{l+1,j} = \sum_{i=1}^n \sum_{k=1}^r \langle \varphi_{l,j}(\chi_{l,i}), \psi_k \rangle \psi_k. \quad (5)$$

Thus, we conceptualize each 2D feature map $\chi_{l,i}$ as an element of a Hilbert space H , whose values are discretized over the spatial domain $h \times w$. This construction allows for a seamless integration into standard deep learning architectures for image processing, eliminating the necessity of feature flattening.

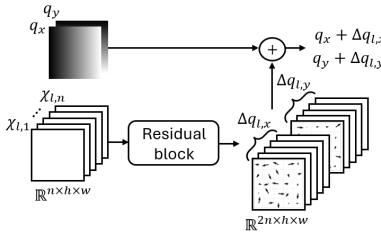


Figure 2: Architecture of the spatial grid deformation module, illustrating learned spatial deformation through residual network-generated offset tensors $\Delta q_{l,x}, \Delta q_{l,y}$. These predicted offsets are combined additively with a broadcasted uniform reference grid $\{q_x, q_y\}$ to produce the deformed sampling grid for the basis Hermite functions evaluation.

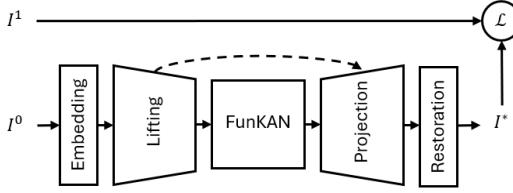


Figure 3: Overview of FunKAN as a multi-purpose backbone for medical image processing. The model processes an input image I^0 , being supervised by a target image I^1 via the loss function $\mathcal{L}(I^*, I^1)$.

The final computational form is derived in two steps. First, the j -index is factored out, which is equivalent to a 1×1 convolution with weights $W_l = \{\omega_{l,j}\}_{j=1}^m, \omega_{l,j} \in \mathbb{R}^{1 \times 1 \times n}$. Second, the Hermite basis functions are evaluated on a deformed grid, a strategy inspired by modern implicit architectures (Agro et al., 2024), inducing an additional learnable dependency upon input functions $\{\chi_{l,i}\}_{i=1}^n$. Consequently, the layer output is computed as:

$$\chi_{l+1,j} = \sum_{i=1}^n \omega_{l,j} \left(\sum_{k=1}^r \langle \varphi_l(\chi_{l,i}), \psi_{l,k}(\chi_{l,i}) \rangle \psi_{l,k}(\chi_{l,i}) \right). \quad (6)$$

Hence, a FunKAN layer is parameterized by three core components: a tensor of learnable inner functions $\varphi_l \in \mathbb{R}^{n \times h \times w}$, a set of deformed Hermite basis functions $\psi_l \in \mathbb{R}^{n \times r \times h \times w}$ and convolutional weights $W_l \in \mathbb{R}^{1 \times 1 \times n \times m}$.

The spatial coordinates for the basis Hermite functions evaluation are deformed by a learned vector field $\Delta q_l = \{\Delta q_{l,x}, \Delta q_{l,y}\}$, generated through a residual block (He et al., 2016) (see Fig. 2). Formally, the grid deformation is computed as $q + \Delta q_l$, where $q = \{q_x, q_y\}$ is a uniform grid.

As shown in Fig. 2, the residual block generates spatial offset tensors: $\Delta q_{l,x} \in \mathbb{R}^{n \times h \times w}$ and $\Delta q_{l,y} \in \mathbb{R}^{n \times h \times w}$. These offsets are then added to a uniform grid, yielding the deformed sampling coordinates used for the basis Hermite functions evaluation. Consistent with the pre-activation approach (Duta et al., 2021), the residual block processes activations through batch normalization (Balestriero & Baraniuk, 2022) and ReLU before the subsequent convolution:

$$\Delta q_l = \mathcal{W}_{l,0} * BN(\chi_l) + \mathcal{F}_l(\chi_l), \quad (7)$$

$$\mathcal{F}_l = \mathcal{W}_{l,2} * ReLU\{BN(\mathcal{W}_{l,1} * ReLU\{BN(\chi_l)\})\}, \quad (8)$$

where $\Delta q_l \in \mathbb{R}^{2n \times h \times w}$ – learned spatial offsets, BN – batch normalization and $\{\mathcal{W}_{l,i}\}_{i=0}^2$ – 3×3 convolutional layers, maintaining the spatial resolution.

While both pre- and post-activation residual architectures possess equivalent theoretical representational capacity, empirical evidence demonstrates superior gradient propagation in pre-activation architectures (Duta et al., 2021). Our implementation ensures stable optimization through batch normalization layers and skip connections.

270 We leverage FunKAN as a backbone architecture within the U-shaped segmentation framework (see
 271 Fig. 3), resulting in the proposed U-FunKAN architecture. The model is configured as follows:
 272

- 273 1. Embedding: 3×3 convolution, projecting an input image into 16-dimensional feature
 274 space.
- 275 2. Lifting: four consecutive U-Net-like encoder residual blocks with progressively increasing
 276 filter count: $32 (C_1) \rightarrow 64 (C_2) \rightarrow 128 (C_3) \rightarrow 128$, – each halving the spatial resolution
 277 through strided 3×3 convolution.
- 278 3. Backbone: sequence of three FunKAN blocks, interconnected with skip connections in
 279 128-dimensional feature space ($n = 128$), encapsulating the spectral encoding of each
 280 inner function over the first six Hermite basis functions ($r = 6$).
- 281 4. Projection: four consecutive U-Net-like decoder residual blocks with gradually decreasing
 282 filter count: $128 (C_3) \rightarrow 64 (C_2) \rightarrow 32 (C_1) \rightarrow 16$, – each performing $\times 2$ upsampling via
 283 nearest-neighbor interpolation, followed by 3×3 convolution for feature refinement and
 284 skip connection from the corresponding lifting module.
- 285 5. Restoration: 1×1 convolution, projecting ReLU pre-activated features to logits.

287 The model is trained in a supervised manner using the weighted combination of binary cross-entropy
 288 and dice loss:

$$289 \mathcal{L}_{segm} = \frac{1}{N} \sum_{i=1}^N 0.1 \cdot CE(I_i^*, I_i^1) + Dice(I_i^*, I_i^1), \quad (9)$$

292 where N – batch size, equals to 8.

294 The Hermite basis is utilized because of the inherent dual localization exhibited by Hermite func-
 295 tions, a property stemming from their role as integral Fourier transform eigenfunctions (Grünbaum,
 296 1982). The number of basis functions ($r = 6$) is determined following the methodology of Penkin
 297 & Krylov (2025), which performed a grid search over three candidate bases – B-splines, Chebyshev
 298 polynomials and Hermite functions, ultimately selecting six basis functions as optimal.

300 4 EXPERIMENTS

302 We conducted a comprehensive evaluation of FunKAN on three datasets, selected to validate our
 303 method’s robustness across anatomical diversity, encompassing oncological, histological and endo-
 304 scopic structures, and modality variations, including ultrasound, histopathology and colonoscopy.

305 The experimental framework is implemented in Python 3.12 using PyTorch 2.5, with all models
 306 trained and evaluated with full precision on NVIDIA RTX A6000 GPU. The software stack employs
 307 PyTorch Lightning 2.5.1, CUDA 11.8 and cuDNN 9. Computational reproducibility is ensured
 308 through the seeds setup and YAML-based experiments management. The models were trained from
 309 scratch till convergence using Adam (Diederik, 2014) stochastic optimization algorithm ($\beta_1 = 0.9$,
 310 $\beta_2 = 0.999$, $\epsilon = 10^{-8}$) with learning rate manual scheduling upon the scheme: 10^{-4} , $5 \cdot 10^{-5}$,
 311 10^{-5} . To enhance generalization, segmentation training datasets were augmented through random
 312 vertical/horizontal flips, rotations and transpositions, each applied with a probability of 0.5.

313 **BUSI dataset.** BUSI dataset (Al-Dhabyani et al., 2020) is a publicly available dataset for breast
 314 tumor segmentation in ultrasound imaging. It consists of 780 grayscale breast ultrasound images in
 315 PNG format, collected from 600 female patients (aged 25–75) and categorized into three classes:
 316 133 images with no visible tumors (normal), 437 images of non-cancerous lesions (benign) and 210
 317 images of confirmed cancerous tumors (malignant). We utilized 647 benign and malignant images,
 318 resized to 256×256 .

319 **GlaS dataset.** GlaS dataset (Valanarasu et al., 2021) is a widely used dataset, specifically designed
 320 for gland segmentation. It contains 165 Hematoxylin and Eosin (H&E) stained histology RGB
 321 images. Our study utilized 165 images, resized to 512×512 . Despite the predefined train-test
 322 division, we re-partitioned them into training (80%) and testing (20%) subsets using a randomized
 323 split with the seed 42. Such approach ensures a fair comparison with competitors and the way we
 split BUSI and CVC datasets.

324
 325 Table 1: Performance comparison of state-of-the-art segmentation models across three clinically
 326 distinct medical imaging scenarios. Results report average intersection over union (IoU) and F1
 327 scores with standard deviation over three random runs

328 Methods	329 BUSI		330 GlaS		331 CVC	
	332 IoU \uparrow	333 F1 \uparrow	334 IoU \uparrow	335 F1 \uparrow	336 IoU \uparrow	337 F1 \uparrow
338 U-Net (Ronneberger et al., 2015)	339 57.22 \pm 4.74	340 71.91 \pm 3.54	341 86.66 \pm 0.91	342 92.79 \pm 0.56	343 83.79 \pm 0.77	344 91.06 \pm 0.47
345 Att-Unet (Oktay et al., 2018)	346 55.18 \pm 3.61	347 70.22 \pm 2.88	348 86.84 \pm 1.19	349 92.89 \pm 0.65	350 84.52 \pm 0.51	351 91.46 \pm 0.25
352 U-Net++ (Zhou et al., 2018)	353 57.41 \pm 4.77	354 72.11 \pm 3.90	355 87.07 \pm 0.76	356 92.96 \pm 0.44	357 84.61 \pm 1.47	358 91.53 \pm 0.88
359 U-NeXt (Valanarasu & Patel, 2022)	360 59.06 \pm 1.03	361 73.08 \pm 1.32	362 84.51 \pm 0.37	363 91.55 \pm 0.23	364 74.83 \pm 0.24	365 85.36 \pm 0.17
366 Rolling-UNet (Liu et al., 2024a)	367 61.00 \pm 0.64	368 74.67 \pm 1.24	369 86.42 \pm 0.96	370 92.63 \pm 0.62	371 82.87 \pm 1.42	372 90.48 \pm 0.83
373 U-Mamba (Ma et al., 2024)	374 61.81 \pm 3.24	375 75.55 \pm 3.01	376 87.01 \pm 0.39	377 93.02 \pm 0.24	378 84.79 \pm 0.58	379 91.63 \pm 0.39
380 UKAGNet (Drokin, 2024)	381 63.45	382 77.64	383 87.31	384 93.23	385 76.85	386 86.91
387 U-KAN (Li et al., 2025)	388 63.38 \pm 2.83	389 76.40 \pm 2.90	390 87.64 \pm 0.32	391 93.37 \pm 0.16	392 85.05 \pm 0.53	393 91.88\pm0.29
394 U-FunKAN (Ours)	395 68.49\pm0.62	396 77.37 \pm 0.58	397 88.02\pm0.24	398 93.50\pm0.12	399 85.93\pm0.72	400 91.42 \pm 0.61

338
 339 Table 2: Efficiency comparison of floating-point operations (Gflops) and trainable parameters
 340 (Params) across state-of-the-art segmentation algorithms. Results are measured for 256×256 input
 341 resolution using THOP Python tool for profiling PyTorch models

342 Methods	343 Gflops \downarrow	344 Params (M) \downarrow
345 U-Net (Ronneberger et al., 2015)	346 524.2	347 34.53
348 Att-Unet (Oktay et al., 2018)	349 533.1	350 34.9
351 U-Net++ (Zhou et al., 2018)	352 1109	353 36.6
354 U-NeXt (Valanarasu & Patel, 2022)	355 4.58	356 1.47
357 Rolling-UNet (Liu et al., 2024a)	358 16.82	359 1.78
360 U-Mamba (Ma et al., 2024)	361 2087	362 86.3
363 U-KAN (Li et al., 2025)	364 14.02	365 6.35
366 U-FunKAN (Ours)	367 4.35	368 3.6

352 **CVC-ClinicDB dataset.** CVC-ClinicDB dataset (Bernal et al., 2015) is a dataset for polyp segmen-
 353 tation in colonoscopy images. It contains 612 high-resolution colonoscopy RGB frames, extracted
 354 from 29 video sequences with varied lighting conditions, specular reflections and mucosal textures.
 355 All images were resized to 256×256 .

356 We measured anatomical structures segmentation accuracy on BUSI, GlaS and CVC-ClinicDB
 357 datasets in terms of IoU and F1 scores. To ensure full reproducibility, we released CSV files speci-
 358 fying the data splits for each dataset in our GitHub repository.

360 Table 1 presents a quantitative evaluation of our proposed U-FunKAN model against several leading
 361 segmentation architectures, including convolutional U-shaped networks, attention-based variants
 362 and recent hybrid models, incorporating Mamba and KAN layers. The reported metrics are averaged
 363 over the last fifty epochs, coupled with uncertainty estimations, derived from three independent
 364 training runs with the seeds: 50, 100 and 150. The analysis reveals that U-FunKAN sets a new state-
 365 of-the-art, achieving the highest IoU on all three datasets, also being the most efficient algorithm in
 366 terms of Gflops (see Table 2). U-FunKAN also attains the highest F1-score on GlaS datasets, while
 367 minor underperforming on BUSI and CVC datasets compared to UKAGNet (Drokin, 2024) and U-
 368 KAN (Li et al., 2025), respectively, which requires $\times 3$ greater computational complexity (in Gflops)
 369 and $\times 1.7$ more parameters than our approach. The observed performance profile, characterized by
 370 a high IoU and a slightly lower F1-score, suggests that the model produces fewer false positives
 371 (a characteristic highly prioritized by IoU) at the cost of a modest increase in false negatives (as
 372 reflected by the F1-score). Deeper U-FunKAN architectures mitigate this trade-off, achieving state-
 373 of-the-art quality in both metrics (see Table 3). Beyond accuracy, U-FunKAN exhibits a low variance
 374 across multiple runs, indicating training stability and reliability.

375 The comparative analysis of U-FunKAN model variants with different channel settings is reported
 376 in Table 3. The proposed channel setting: $32 (C_1) \rightarrow 64 (C_2) \rightarrow 128 (C_3)$, – achieves an optimal
 377 balance between computational efficiency and segmentation performance, yielding state-of-the-art
 378 results. Deeper configurations can further improve accuracy, while they incur additional computa-
 379 tional overhead.

378

379

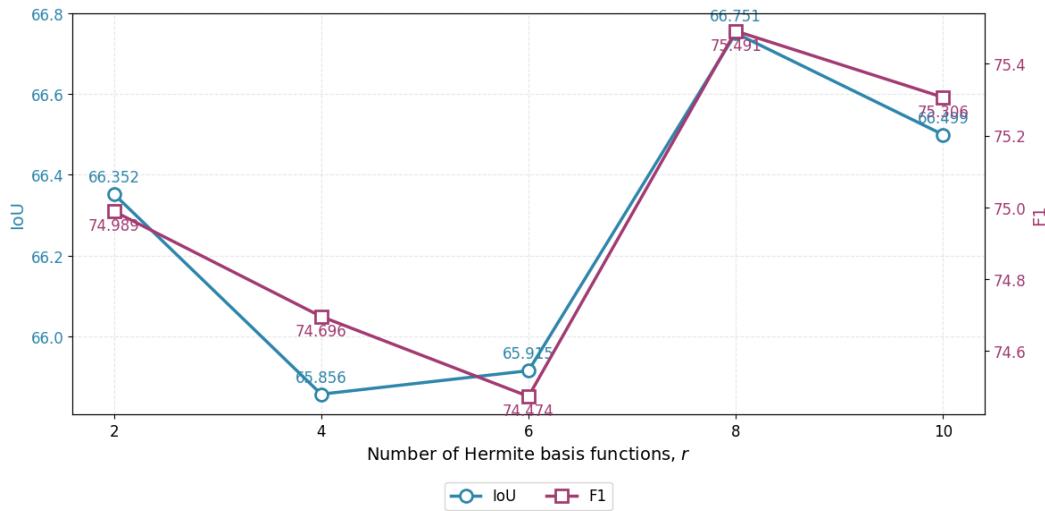
380 Table 3: Ablation study on impact of the channel scaling in U-FunKAN on segmentation performance and efficiency. The best intersection over union (IoU) and F1 scores are reported
381

U-FunKAN			BUSI			
C_1	C_2	C_3	IoU \uparrow	F1 \uparrow	Gflops \downarrow	Params (M) \downarrow
32	64	128	69.11	77.95	4.35	3.6
64	96	128	69.94	78.42	10.84	4.1
128	160	256	69.49	78.39	40.42	15.7
256	320	512	70.62	79.31	161.43	62.4

387

388

389



404

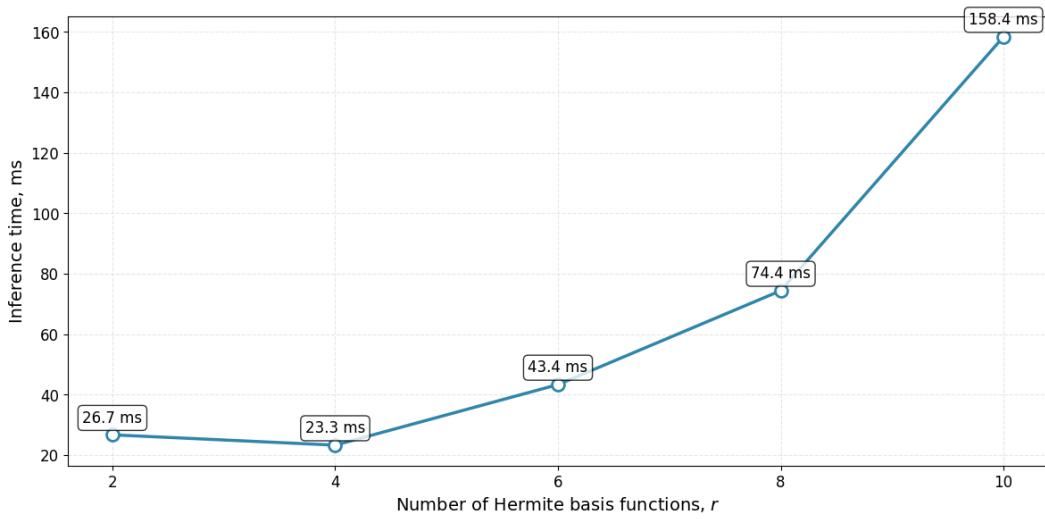
405

406

407 Figure 4: Ablation study on impact of the number of the basis Hermite functions in U-FunKAN on
408 breast cancer segmentation quality on BUSI ultrasound dataset. Averaged IoU and F1 scores are
409 reported over the last fifty epochs for each experiment.

410

411



427

428

429

430

431

432 Figure 5: Ablation study on impact of the number of the basis Hermite functions in U-FunKAN on
433 efficiency. Median inference times are reported over 100 runs for a 256×256 input on Intel Core
434 i7-14700HX.

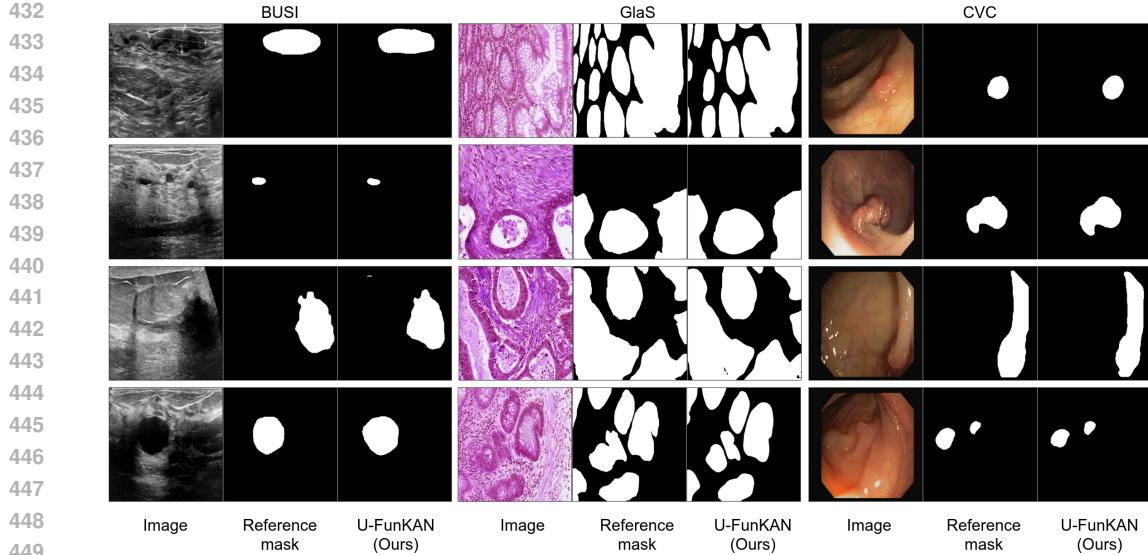


Figure 6: Qualitative results of U-FunKAN over three heterogeneous medical scenarios.

Fig. 4 shows an ablation study evaluating the sensitivity of U-FunKAN’s segmentation quality to the number of Hermite basis functions, r . Models were trained from scratch with learning rate 10^{-4} till convergence on BUSI dataset to ensure a fair comparison. The results demonstrate that both IoU and F1-score improve as r increases from the default value of 6 to 8 or 10. However, this extra gain in accuracy comes at the cost of increased computational latency, as detailed in Fig. 5.

Fig. 6 presents qualitative segmentation results of the proposed U-FunKAN model across three heterogeneous medical imaging scenarios: breast cancer detection in ultrasound images, glands segmentation in histopathology and polyps segmentation in colonoscopy images.

Furthermore, to demonstrate that our method has general applicability beyond binary segmentation, we apply FunKAN to the task of MRI reconstruction, outperforming KAN learnable backbones (Seydi, 2024) and classical state-of-the-art Kellner method (Kellner et al., 2016) for Gibbs-ringing suppression. In particular, we leverage FunKAN as a backbone architecture for MRI enhancement with the following setup:

1. Embedding: 5×5 convolution, projecting an input image into 16-dimensional feature space.
2. Lifting: 3×3 convolution, projecting ReLU pre-activated features into 32-dimensional feature space, maintaining the spatial resolution.
3. Backbone: sequence of three FunKAN blocks, interconnected with skip connections in 32-dimensional feature space ($n = 32$), encapsulating the spectral encoding of each inner function over the first six Hermite basis functions ($r = 6$).
4. Projection: 3×3 convolution, projecting ReLU pre-activated features into 16-dimensional feature space, maintaining the spatial resolution.
5. Restoration: 1×1 convolution, projecting ReLU pre-activated features to the target color space.

This model is trained in a supervised manner using mean squared error loss function on IXI dataset, mitigating Gibbs-ringing after discarding 75% of high-frequency information:

$$\mathcal{L}_{enh} = \frac{1}{N} \sum_{i=1}^N \|I_i^* - I_i^1\|_2^2, \quad (10)$$

where N – batch size, equals to 8.

Methods	IXI	
	PSNR \uparrow	TV
I^0	31.33	1476.55
I^1	–	1255.40
Kellner (Kellner et al., 2016)	31.09	1120.05
MLP (Cybenko, 1989)	37.96	1145.57
KAN (Liu et al., 2024b)	38.10	1161.63
ChebyKAN (SS et al., 2024)	38.01	1156.56
HermiteKAN (Seydi, 2024)	38.04	1161.31
FunKAN (Ours)	39.05	1174.86

Table 4: Application of FunKAN as MRI enhancement backbone within the convolutional architecture. Results report average peak signal-to-noise ratio (PSNR) and total variation (TV) across 2617 test images from IXI dataset. I^0 , I^1 denote artifact-corrupted and artifact-free images, respectively.

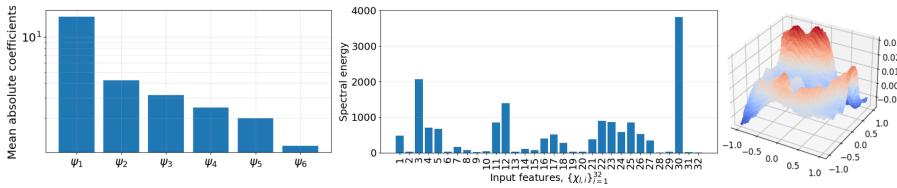


Figure 7: Interpretability of FunKAN as MRI enhancement backbone. Aggregated contribution of Hermite functions $\{\psi_k\}_{k=1}^r$ across all learned inner functions. Feature ranking intrinsics, computed as the total spectral energy, towards feature selection. A learned inner function $\varphi_{l,i} \in \mathbb{R}^{h \times w}$.

IXI dataset comprises approximately 600 MRI scans (255×255) from healthy subjects, collected using 1.5T and 3T scanners from Philips and GE Healthcare. It is composed of T1, T2 and PD 3D volumes, encoded in NIFTI format. Firstly, the intersection of these volumes was taken, producing 577 volumes that have all three modalities: T1, T2 and PD. Then, the first 400 volumes were utilized to synthesize the training set, the next 100 volumes to create the testing set and the rest of the data to generate the validating set, forming 10473, 2617 and 2016 image pairs, respectively.

A quantitative comparison (see Table 4) of backbones within the same convolutional architecture for MRI enhancement on IXI dataset demonstrates that substituting MLP with KAN and subsequently retraining the entire model results in a modest PSNR improvement of 0.1 dB. In contrast, FunKAN backbone achieves a markedly higher gain, outperforming all considering here KANs (B-spline, Chebyshev, Hermite) by 1 dB, owing to its inherent capacity to incorporate the geometric relations of visual data by design. Performance is benchmarked with superior restoration quality indicated by higher PSNR values and TV measures approaching the reference I^1 .

From an interpretability standpoint (see Fig. 7), the Hermite decomposition of a feature provides a spectral characterization of its complexity by construction. A concentration of spectral energy in the low-order Hermite coefficients indicates a smooth inner function, which is empirically associated with robust, generalizable features. Conversely, significant energy in higher-order coefficients suggests a propensity for overfitting and heightened adversarial vulnerability, as high-frequency components often correspond to non-robust, data-specific artifacts (Harder et al., 2021).

5 CONCLUSION

This paper introduced FunKAN, a novel neural framework that advances the state-of-the-art in medical image segmentation. The core of our contribution is a generalization of the Kolmogorov-Arnold representation theorem to functional spaces, which we have empirically validated. Through extensive evaluations on diverse tasks, including breast tumor detection in ultrasound, gland segmentation in histology and polyp identification in colonoscopy, the proposed U-FunKAN architecture demonstrated superior performance over contemporary models, such as convolutional networks, attention-based mechanisms and recent architectures based on Mamba and KAN. To support reproducibility and future research, our code will be made publicly available upon acceptance of this paper.

540 REFERENCES
541

542 Barsha Abhisheka, Saroj Kumar Biswas, and Biswajit Purkayastha. A Comprehensive Review on
543 Breast Cancer Detection, Classification and Segmentation using Deep Learning. *Archives of Com-
544 putational Methods in Engineering*, 30(8):5023–5052, 2023.

545 Ben Agro, Quinlan Sykora, Sergio Casas, Thomas Gilles, and Raquel Urtasun. UnO: Unsupervised
546 Occupancy Fields for Perception and Forecasting. In *Proceedings of the IEEE/CVF Conference
547 on Computer Vision and Pattern Recognition*, pp. 14487–14496, 2024.

548 Walid Al-Dhabayani, Mohammed Gomaa, Hussien Khaled, and Aly Fahmy. Dataset of Breast Ultra-
549 sound Images. *Data in brief*, 28:104863, 2020.

550 Randall Balestrieri and Richard G Baraniuk. Batch Normalization Explained. *arXiv preprint
551 arXiv:2209.14778*, 2022.

552 Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Debora Gil, Cristina Rodríguez, and
553 Fernando Vilariño. WM-DOVA Maps for Accurate Polyp Highlighting in Colonoscopy: Val-
554 idation vs. Saliency Maps from Physicians. *Computerized medical imaging and graphics*, 43:
555 99–111, 2015.

556 Katarzyna Borys, Yasmin Alyssa Schmitt, Meike Nauta, Christin Seifert, Nicole Krämer,
557 Christoph M Friedrich, and Felix Nensa. Explainable AI in Medical Imaging: An Overview for
558 Clinical Practitioners – Beyond Saliency-based XAI Approaches. *European journal of radiology*,
559 162:110786, 2023.

560 George Cybenko. Approximation by Superpositions of a Sigmoidal Function. *Mathematics of
561 control, signals and systems*, 2(4):303–314, 1989.

562 P Kingma Diederik. Adam: A Method for Stochastic Optimization. *(No Title)*, 2014.

563 Ivan Drokin. Kolmogorov-Arnold Convolutions: Design Principles and Empirical Studies. *arXiv
564 preprint arXiv:2407.01092*, 2024.

565 Ionut Cosmin Duta, Li Liu, Fan Zhu, and Ling Shao. Improved Residual Networks for Image and
566 Video Recognition. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pp.
567 9415–9422. IEEE, 2021.

568 F Alberto Grünbaum. The Eigenvectors of the Discrete Fourier Transform: A Version of the Hermite
569 Functions. *Journal of Mathematical Analysis and Applications*, 88(2):355–363, 1982.

570 Paula Harder, Franz-Josef Pfreundt, Margret Keuper, and Janis Keuper. SpectralDefense: Detecting
571 Adversarial Attacks on CNNs in the Fourier Domain. In *2021 International Joint Conference on
572 Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2021.

573 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image
574 Recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
575 pp. 770–778, 2016.

576 Yezi Ali Kadhim, Muhammad Umer Khan, and Alok Mishra. Deep learning-based Computer-Aided
577 Diagnosis (CAD): Applications for Medical Image Datasets. *Sensors*, 22(22):8999, 2022.

578 Elias Kellner, Bibek Dhital, Valerij G Kiselev, and Marco Reisert. Gibbs-ringing Artifact Removal
579 based on Local Subvoxel-shifts. *Magnetic resonance in medicine*, 76(5):1574–1581, 2016.

580 Andrei Nikolaevich Kolmogorov. On the Representations of Continuous Functions of Many Vari-
581 ables by Superposition of Continuous Functions of One Variable and Addition. In *Dokl. Akad.
582 Nauk USSR*, volume 114, pp. 953–956, 1957.

583 Chenxin Li, Xinyu Liu, Wuyang Li, Cheng Wang, Hengyu Liu, Yifan Liu, Zhen Chen, and Yix-
584 uan Yuan. U-KAN Makes Strong Backbone for Medical Image Segmentation and Generation.
585 *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(5):4652–4660, 2025.

586 Ziyao Li. Kolmogorov-Arnold Networks are Radial Basis Function Networks. *arXiv preprint
587 arXiv:2405.06721*, 2024.

594 Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier Neural Operator
 595 with Learned Deformations for PDEs on General Geometries. *Journal of Machine Learning
 596 Research*, 24(388):1–26, 2023.

597

598 Yutong Liu, Haijiang Zhu, Mengting Liu, Huaiyuan Yu, Zihan Chen, and Jie Gao. Rolling-UNet:
 599 Revitalizing MLP’s Ability to Efficiently Extract Long-distance Dependencies for Medical Image
 600 Segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pp.
 601 3819–3827, 2024a.

602 Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić,
 603 Thomas Y Hou, and Max Tegmark. KAN: Kolmogorov-Arnold Networks. *arXiv preprint
 604 arXiv:2404.19756*, 2024b.

605 Jun Ma, Feifei Li, and Bo Wang. U-Mamba: Enhancing Long-range Dependency for Biomedical
 606 Image Segmentation. *arXiv preprint arXiv:2401.04722*, 2024.

607

608 Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Matthias Heinrich, Kazunari Misawa,
 609 Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention U-Net:
 610 Learning Where to Look for the Pancreas. *arXiv preprint arXiv:1804.03999*, 2018.

611 MA Penkin and AS Krylov. Adaptive Method for Selecting Basis Functions in Kolmogorov-Arnold
 612 Networks for Magnetic Resonance Image Enhancement. *Programming and Computer Software*,
 613 51(3):167–172, 2025.

614

615 Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua
 616 Bengio, and Aaron Courville. On the Spectral Bias of Neural Networks. In *International confer-
 617 ence on machine learning*, pp. 5301–5310. PMLR, 2019.

618 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomed-
 619 ical Image Segmentation. In *International Conference on Medical image computing and computer-
 620 assisted intervention*, pp. 234–241. Springer, 2015.

621

622 Seyd Teymoor Seydi. Exploring the Potential of Polynomial Basis Functions in Kolmogorov-
 623 Arnold Networks: A Comparative Study of Different Groups of Polynomials. *arXiv preprint
 624 arXiv:2406.02583*, 2024.

625 Sidharth SS, Keerthana AR, Anas KP, et al. Chebyshev Polynomial-based Kolmogorov-Arnold
 626 Networks: An Efficient Architecture for Nonlinear Function Approximation. *arXiv preprint
 627 arXiv:2405.07200*, 2024.

628

629 Edward C Titchmarsh. *Introduction to the theory of Fourier integrals*. Oxford: Clarendon Press;
 630 London: Oxford University Press, 1948.

631

632 Jeya Maria Jose Valanarasu and Vishal M Patel. UNeXt: MLP-based Rapid Medical Image Segmen-
 633 tation Network. In *International conference on medical image computing and computer-assisted
 634 intervention*, pp. 23–33. Springer, 2022.

635

636 Jeya Maria Jose Valanarasu, Poojan Oza, Ilker Hacihaliloglu, and Vishal M Patel. Medical Trans-
 637 former: Gated Axial-attention for Medical Image Segmentation. In *International conference on
 638 medical image computing and computer-assisted intervention*, pp. 36–46. Springer, 2021.

639

640 Zhuoqin Yang, Jiansong Zhang, Xiaoling Luo, Zheng Lu, and Linlin Shen. MedKAN: An
 641 Advanced Kolmogorov-Arnold Network for Medical Image Classification. *arXiv preprint
 642 arXiv:2502.18416*, 2025.

643

644 Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. UNet++:
 645 A Nested U-Net Architecture for Medical Image Segmentation. In *International workshop on
 646 deep learning in medical image analysis*, pp. 3–11. Springer, 2018.

647