AIMS Mathematics, 11(5): 13530–13566.
DOI: 10.3934/math.2026557
Received: 23 January 2026
Revised: 04 May 2026
Accepted: 12 May 2026
Published: 15 May 2026

https://www.aimspress.com/journal/Math

Research article

A pointwise contextual multiscale texture operator and its application to 3D medical image segmentation

C. Coronado Gallardo[1], C. Fernández García[2], I. Suazo[1], J. A. Vega[1,3] and Z. Fernández Muñiz[2,*]

[1] Facultad de Ciencias de la Salud, Universidad Autónoma de Chile, Providencia-Santiago de Chile, Chile

[2] Departamento de Matemáticas, Universidad de Oviedo, Asturias, Spain

[3] Departamento de Morfología y Biología Celular, Universidad de Oviedo, Asturias, Spain

* **Correspondence:** Email: zulima@uniovi.es.

Abstract: We introduce a pointwise multiscale texture operator providing a mathematically grounded and interpretable description of how local image context evolves across observation scales. The operator is defined through the scale evolution of a contextual multiscale signature derived from Gaussian scale-space representations, thereby establishing direct connections with diffusion processes and multiresolution analysis. This formulation yields a stable descriptor of texture transitions that is robust to moderate noise, blurring, and acquisition variability, while remaining independent of any specific segmentation strategy. As a demonstration of its practical utility, we embed the proposed operator within a texture-coherence-driven region-growing framework for image segmentation. The resulting algorithm is dimension independent, computationally tractable, and does not rely on training data or complex architectural design. We illustrate its performance on 3D cone-beam computed tomography (CBCT) datasets, where it enables coherent segmentation and volumetric reconstruction of the dental pulp chamber under clinically relevant conditions involving low contrast, partial volume effects, and heterogeneous textures. Beyond this specific application, the proposed operator constitutes a general framework for multiscale texture analysis and structural characterization, particularly suited to imaging scenarios where interpretability, robustness to acquisition degradation, and limited annotated data are critical.

Keywords: pointwise multiscale texture operator; image segmentation; multi-resolution analysis; pulp cavity; volume estimation
Mathematics Subject Classification: 68U10, 65D18

1. Introduction

Image segmentation is a fundamental problem in image processing and computer vision. Broadly speaking, it consists of decomposing an image into meaningful regions that are internally coherent and distinct from their surroundings according to criteria such as intensity, texture, or geometric structure [1, 2]. From a mathematical standpoint, segmentation is often formulated as a partition problem, with well-established connections to variational models [3, 4], partial differential equation (PDE)-based methods [5], and multiscale decompositions [6].

A major difficulty in segmentation arises from the fact that real images rarely exhibit ideal, sharp discontinuities. In practice, transitions between regions are often blurred due to acquisition effects, noise, intrinsic structural smoothness, or anatomical variability. From a modeling perspective, the observed image can frequently be interpreted as a noisy and smoothed version of an underlying ideal signal, typically expressed as the convolution of a latent image with a point spread function plus additive noise [1, 7]. This smoothing erodes sharp boundaries and complicates the discrimination of regions based solely on intensity.

These challenges are further amplified in clinical 3D imaging, such as cone-beam computed tomography (CBCT), where low contrast, streaking artifacts, partial volume effects, and heterogeneous textures hinder reliable extraction of anatomical structures. In digital dentistry, for instance, accurate estimation of the pulp cavity volume is essential for diagnosis, personalized endodontic planning, minimally invasive treatments, and age estimation. However, CBCT-based segmentation remains particularly demanding under these conditions [8–11].

Classical segmentation techniques, including region growing, edge detection, thresholding, and texture-based descriptors, have played an important historical role due to their simplicity and computational efficiency [12]. Among them, Otsu's thresholding method remains widely used in contemporary applications [13–15], and modern variants further improve performance in noisy or textured scenarios [16]. Nevertheless, methods purely based on global or local thresholds typically struggle in low-contrast volumetric images such as CBCT, where structural smoothness and device-dependent variability may obscure intensity-based boundaries. Recent comprehensive reviews provide an integrated overview of current advances in medical image segmentation, covering classical intensity-based approaches, texture-driven methods, and modern deep learning architectures, while emphasizing persistent challenges related to robustness, generalization, and interpretability [17–19].

Beyond intensity-based segmentation, many works have incorporated entropy-based and texture-driven descriptors. These approaches exploit the fact that structural heterogeneity is often better captured by local spatial relationships than by intensity alone [20–22]. In parallel, deep learning has transformed medical image segmentation, with architectures such as U-Net, DeepLab, and transformer-based models achieving state-of-the-art performance across many modalities [23–25]. Before the widespread adoption of deep neural networks, dental CT segmentation had already been explored using variational methods and active contour formulations, particularly level-set-based approaches [26]. Building on these early variational approaches, U-Net-based architectures have subsequently been introduced for refined tooth and pulp segmentation in CBCT volumes, achieving substantial improvements in anatomical delineation and volumetric consistency [27,28]. More recently, convolutional neural network (CNN)-based systems have obtained excellent results in tooth and pulp cavity segmentation from CBCT scans [29–33]. Nevertheless, these approaches typically require

large annotated datasets and substantial computational resources and may suffer from issues related to generalization, interpretability, and clinical integration [8]. Beyond task-specific deep learning segmentation architectures, recent advances in representation learning have explored the extraction of structurally robust features through data-driven hierarchical models. In particular, deep architectures developed for complex scene parsing aim to capture multilevel contextual information and structural coherence under challenging imaging conditions, including noise, occlusions, and heterogeneous textures. Although these approaches are fundamentally learning-based and supervised in nature, they emphasize hierarchical contextual encoding under complex imaging conditions. In this broader perspective, such representation learning frameworks complement analytically grounded multiscale formulations by addressing structural stability from a data-driven standpoint. A representative example is the work of [34], which investigates robust feature representation learning for accurate structural parsing in complex visual environments. Unlike these supervised representation learning strategies, the framework proposed in this work is defined mathematically within Gaussian scale–space theory and does not rely on training data or large annotated corpora.

These limitations motivate the search for mathematically grounded descriptors that characterize local structure in a stable and interpretable manner across scales, independently of large training datasets.

Unlike classical multiscale texture descriptors computed at fixed resolutions—such as gray-level co-occurrence matrices, wavelet energy features, or entropy-based statistics—the operator introduced in this work is defined as the derivative with respect to scale of a contextual multiscale signature generated by Gaussian diffusion. This formulation embeds the descriptor within the axiomatic framework of scale–space theory and enables a dynamical characterization of local contextual evolution across scales, providing a mathematically principled alternative to purely statistical or transform-based descriptors.

Rather than proposing yet another segmentation algorithm, this work focuses on the definition of a pointwise multiscale texture operator that encodes the scale evolution of local image context in a mathematically explicit and structurally interpretable manner.

Multiscale and multidirectional texture analysis has proven highly effective in biomedical imaging, both from a theoretical perspective and through practical applications involving texture-driven classification and segmentation [35–37]. This perspective naturally connects to the mathematical foundations of scale-space theory, which provides a principled framework for multiscale image representations and contextual analysis. The scale-space theory pioneered by Lindeberg [38] establishes that Gaussian convolution is the unique linear, isotropic, and causality-preserving mechanism for generating representations of an image at multiple resolutions. Related multiscale stabilization and compensation strategies have also been explored in other signal-processing domains, where scale-dependent representations are used to mitigate environmental variability, acquisition drift, and noise effects [39].

The evolution of an image under Gaussian smoothing satisfies a diffusion equation in which the scale parameter plays the role of a pseudo-time variable, creating a natural bridge between segmentation, diffusion processes, and texture descriptors derived from derivatives with respect to scale. Building on these ideas, multiscale representations have been widely used to characterize texture in a manner robust to noise, blurring, and acquisition variability, and they are closely connected to classical multiresolution techniques such as wavelet analysis [6].

Motivated by this theoretical foundation, we introduce the operator defined as the derivative with respect to scale of a contextual multiscale signature obtained from localized Gaussian averages. This operator captures how local image context evolves as the observation scale increases, providing a descriptor of multiscale texture transitions that is independent of any particular segmentation paradigm.

As a concrete and clinically relevant application, we embed this operator within a texture-coherence-driven region-growing framework. The resulting algorithm is dimension-independent and computationally tractable, and it does not rely on training data or complex architectural design. Although the framework is broadly applicable, we illustrate its performance on 3D CBCT datasets. The proposed method consistently extracts structurally coherent anatomical regions even in the presence of blurred edges, artifacts, and nontrivial texture patterns.

As a representative application, we apply the approach to the segmentation and volumetric reconstruction of the dental pulp chamber, showing stable performance in noisy, low-contrast, and structurally intricate 3D data while maintaining interpretability and moderate computational complexity.

The proposed framework is particularly suitable in imaging scenarios where annotated datasets are scarce, acquisition conditions are heterogeneous, or interpretability is required for clinical integration. In such contexts—including CBCT imaging and low-dose tomography—training-dependent segmentation models may face generalization or transparency limitations. By contrast, the present operator provides a scale–space grounded operator operating independently of prior learning, making it appropriate for exploratory analysis and quantitative morphometric applications.

The primary contribution of this study is methodological rather than benchmark-oriented. The goal is not to introduce a competing segmentation architecture but to define a multiscale operator that provides an interpretable functional description of contextual texture evolution across observation scales.

Accordingly, the segmentation framework is used only as a proof-of-applicability illustrating how the operator can guide structural extraction under realistic imaging conditions. The work therefore falls within the scope of analytical imaging methodology rather than large-scale supervised segmentation benchmarking.

From a theoretical standpoint, the contribution of this work is twofold. First, it reformulates Gaussian scale-space derivatives into a functional contextual signature that characterizes, for each point, the evolution of its local neighborhood across a prescribed scale interval rather than relying on isolated filter responses at fixed resolutions. Second, it introduces an L^2-based multiscale coherence criterion between such signatures and integrates it into a region-growing framework so that segmentation decisions are governed by persistent contextual similarity across scales instead of single-scale responses or local intensity statistics. These elements depart from the standard use of fixed-scale texture descriptors and difference-of-Gaussians/Laplacian of Gaussian (DoG/LoG) constructions, in which scale derivatives are typically interpreted as pointwise feature detectors at individual scales, by instead treating them as components of a scale-dependent functional representation.

To formalize the setting in which the proposed operator is defined, we next describe the mathematical model of image formation and the acquisition effects that motivate a texture-based multiscale description.

2. Materials and methods

2.1. Preliminaries: image formation model and acquisition effects

In many imaging applications, and particularly in medical contexts, segmentation aims to identify regions of interest within a function $I : D \subset \mathbb{R}^n \to [0, 1]$, where the domain D is typically a rectangular box in $\mathbb{R}^n$ with $n = 2$ (2D images) or $n = 3$ (volumetric data). The value $I(\mathbf{p})$ represents the intensity of the underlying physical quantity at point $\mathbf{p}$. For visualization and processing purposes, the dynamic range is usually normalized to $[0, 1]$ [1].

However, the acquired image does not coincide with the ideal function I. Two main effects distort the measurement:

(1) convolutional blurring caused by the finite spatial resolution of the acquisition system; and
(2) additive noise arising from detector electronics, photon statistics, and reconstruction artifacts.

Consequently, an imaging device does not record the exact value $I(\mathbf{p})$ at each point, but rather a weighted average of intensities within a neighborhood determined by the PSF of the system. Denoting by ψ the PSF, the measured (blurred) signal satisfies $u(\mathbf{p}) \approx (I * \psi)(\mathbf{p})$, where $*$ denotes convolution.

In addition, the image is sampled on a discrete grid $\{(x_j, y_k)\} \subset D$ so that each recorded value represents a spatially averaged measurement rather than a pointwise evaluation. In discrete form, the recorded signal can be expressed as

$$u(x_j, y_k) \approx \int_{\mathbb{R}^2} I(x, y)\, \psi(x_j - x, y_k - y)\, dx\, dy = (I * \psi)(x_j, y_k).$$

In practice, the final captured data include additional noise, $I_c(\mathbf{p}) = (I * \psi)(\mathbf{p}) + \rho(\mathbf{p})$ with ρ representing an additive component often modeled as zero-mean white noise with small variance [7].

Convolutional averaging causes loss of high-frequency details and smooths transitions between anatomical structures, while the additive noise term further degrades boundary localization. This phenomenon is illustrated in Figures 1–3.

Figure 1. Ideal two-tone image.

Figure 2. Image after convolutional blurring.

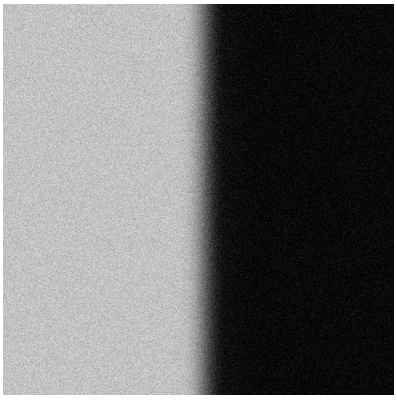

Figure 3. Blurred image with additive noise.

These effects are particularly relevant in CBCT imaging, where the data consist of a 3D function

$$I : D = [0, w] \times [0, h] \times [0, t] \subset \mathbb{R}^3 \to [0, 1],$$

sampled on a discrete voxel grid. The same acquisition model applies:

$$I_c(\mathbf{p}) = (I * \psi)(\mathbf{p}) + \rho(\mathbf{p}), \qquad \mathbf{p} \in D,$$

where $\psi : \mathbb{R}^3 \to \mathbb{R}$ is the 3D PSF and ρ the corresponding noise process. From the volumetric dataset, one may extract 2D slices or orthogonal sections, each of which inherits the same blurring and noise characteristics.

The presence of blurred interfaces is clearly visible in dental CBCT images, where anatomical regions, such as enamel, dentin, pulp chamber, and surrounding bone, do not exhibit sharp boundaries. Instead, their separation is dominated by differences in local texture rather than by clean edges.

This acquisition model motivates the use of descriptors that integrate contextual information over increasing spatial supports rather than relying on pointwise intensity or sharp gradient-based indicators. Such descriptors are naturally grounded in scale–space diffusion theory, where progressive smoothing is interpreted as a diffusion process generating a one-parameter family of image representations across observation scales [40, 41]. This behavior is illustrated in Figure 4.

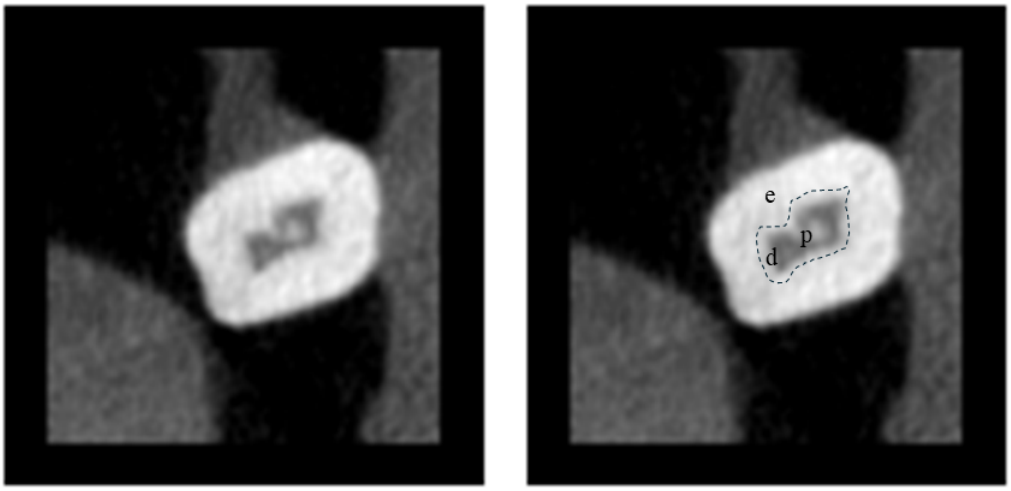

Figure 4. CBCT slice of a tooth: Anatomical regions are primarily distinguished by texture rather than sharp edges.

2.2. Overall description of the method

The proposed segmentation method is based on multiscale analysis and the pointwise multiscale texture operator. Many real images exhibit blurred boundaries, noise, and structural heterogeneity, which limits the effectiveness of segmentation strategies relying exclusively on intensity or gradient information [1, 2]. To overcome these limitations, the proposed method exploits texture information encoded through the scale evolution of local contextual structure across multiple spatial scales.

The algorithm integrates three core components:

(1) **Gaussian smoothing at multiple scales**, which provides a multiresolution decomposition in the sense of classical scale–space theory [38];

(2) The **pointwise multiscale texture operator**, defined through the scale evolution of the contextual multiscale signature of each point; and
(3) A **coherence criterion**, expressed as an L^2 similarity condition over a scale interval, determining whether a candidate point is consistent with the texture of a reference region within the current segmentation.

The procedure begins with a reference point known to lie within the region of interest. Starting from this initial set, the algorithm follows an iterative region-growing strategy: At each iteration, points adjacent to the current region are examined, and those whose pointwise texture is compatible (within a prescribed threshold) with the texture of the reference region are incorporated. This strategy results in a progressive expansion of the region guided by multiscale texture coherence.

The pointwise multiscale texture operator provides several advantages. At coarse scales, Gaussian smoothing attenuates high-frequency noise and reduces texture variability, facilitating the identification of homogeneous regions. The derivative with respect to scale emphasizes structural transitions in a manner less sensitive to local intensity fluctuations, capturing variations arising from spatial organization in addition to sharp intensity discontinuities [20]. Consequently, the method is particularly effective in images with diffuse boundaries, convolutional blurring [7], or low contrast.

Although the formulation is fully general and applies to both 2D and 3D images, we illustrate its performance by segmenting the dental pulp chamber in 3D CBCT scans. This also enables empirical validation through the computation of the segmented volume.

The mathematical formulation of the multiscale descriptors and their theoretical properties are detailed in the following subsections.

2.3. Contextual multiscale signature

Let $I : \mathbb{R}^n \to \mathbb{R}$ denote an image (with $n = 2$ or $n = 3$ in applications). For $\sigma > 0$, let

$$\gamma_\sigma(\mathbf{x}) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left(-\frac{\|\mathbf{x}\|^2}{2\sigma^2}\right)$$

be the Gaussian kernel with variance σ^2. The contextual multiscale signature of a point $\mathbf{p} \in \mathbb{R}^n$ is defined by the convolution

$$g_{\mathbf{p},I}(\sigma) = (I * \gamma_\sigma)(\mathbf{p}) = \int_{\mathbb{R}^n} I(\mathbf{x})\, \gamma_\sigma(\mathbf{p} - \mathbf{x})\, d\mathbf{x}. \tag{2.1}$$

This quantity represents the average value of the image around the point $\mathbf{p}$ at scale σ, encoding how the point relates to its surrounding structures as the observation window expands. As σ increases, the signature integrates progressively larger neighborhoods, capturing the participation of the point in textures or anatomical patterns across scales.

Importantly, the contextual multiscale signature is defined independently of any segmentation or classification procedure.

Unlike fixed-scale contextual averages, the scale-dependent evolution of $g_{\mathbf{p},I}(\sigma)$ enables the extraction of differential texture information, which forms the basis of the operator introduced in the next subsection.

2.4. Pointwise texture operator

The *pointwise texture operator* is defined as the derivative of the signature with respect to scale:

$$\tau_{\mathbf{p},I}(\sigma) = \frac{d}{d\sigma}\, g_{\mathbf{p},I}(\sigma). \tag{2.2}$$

This operator captures the scale evolution of the local contextual structure as the smoothing scale increases. Large absolute values are typically associated with contextual transitions or structural boundaries; small values correspond to locally homogeneous regions. Furthermore, if $I_2 = aI_1 + b$ with $a, b \in \mathbb{R}$, then

$$\tau_{\mathbf{p},I_2}(\sigma) = a\, \tau_{\mathbf{p},I_1}(\sigma),$$

hence the operator is invariant under additive intensity shifts.

As a scale-dependent quantity, $\tau_{\mathbf{p},I}(\sigma)$ provides a pointwise characterization of structural transitions that is naturally suited to multiscale analysis.

2.5. Scale–space diffusion identity

A fundamental property linking the contextual multiscale signature (2.1) to diffusion processes underlies its stability and physical interpretability.

Theorem 2.1. *Let $I \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$ and let γ_σ denote the Gaussian kernel. Then the contextual multiscale signature*

$$g_{p,I}(\sigma) = (I * \gamma_\sigma)(p)$$

satisfies the scale–space diffusion equation

$$\frac{\partial g_{p,I}}{\partial \sigma}(\sigma) = \sigma \Delta g_{p,I}(\sigma), \tag{2.3}$$

where Δ denotes the spatial Laplacian. In particular, for every $\sigma > 0$, the map $p \mapsto g_{p,I}(\sigma)$ is smooth.

Equation (2.3) shows that Gaussian smoothing generates the unique linear, isotropic, and causality-preserving scale–space [38]. The pointwise texture operator is therefore closely related to the scale derivative of a diffusion process.

Although the diffusion identity is classical within scale–space theory, its role here is not merely formal. It provides the theoretical foundation supporting the use of scale derivatives as stable, physically consistent multiscale texture operators. In this context, the novelty does not lie in the diffusion equation per se, but in its reinterpretation as a mechanism for pointwise texture characterization driving the coherence-driven region-growing framework.

Additional properties and interpretation. This identity provides a physically meaningful interpretation of the contextual multiscale signature. As the scale parameter σ increases, the Gaussian kernel progressively smooths the image, acting as a diffusion process in which σ plays the role of a pseudo-time variable. Coarser scales therefore correspond to increasingly smoothed neighborhoods of $\mathbf{p}$, and the evolution of $g_{\mathbf{p},I}(\sigma)$ reflects the gradual suppression of fine structural details.

A classical consequence of Gaussian scale–space theory is the conservation of mass:

$$\int_{\mathbb{R}^n} g_{\mathbf{p},I}(\sigma)\, d\mathbf{p} = \text{constant},$$

showing that diffusion redistributes intensity without altering its total amount. This reinforces the interpretation of $g_{\mathbf{p},I}$ as a physically consistent multiscale descriptor.

Moreover, the texture operator $\tau_{\mathbf{p},I}(\sigma)$ inherits several desirable properties from the underlying diffusion process. The smoothness of $g_{\mathbf{p},I}(\sigma)$ for every $\sigma > 0$ ensures that $\tau_{\mathbf{p},I}$ is well defined and continuous, while standard stability results for Gaussian convolution imply that small perturbations of the image produce small variations in both $g_{\mathbf{p},I}$ and its scale derivative. These properties are essential for robust texture characterization under noise and acquisition artifacts.

Finally, multiplying $\tau_{\mathbf{p},I}(\sigma)$ by σ connects the operator with well-established mechanisms in edge detection within scale–space theory, where scale-normalized derivatives provide consistent responses across resolutions. Thus, the pointwise texture operator not only quantifies local textural transitions but also aligns with the fundamental principles of multiscale analysis. This theoretical correspondence underpins the stability, interpretability, and effectiveness of the proposed segmentation framework when applied to images with blurred interfaces and multiscale texture patterns.

2.6. r-adic discretization and pseudo-wavelet representation

For computational purposes, we discretize the scale variable using a grid,

$$\sigma_s = r^s \sigma_0, \qquad s \in \mathbb{Z},$$

with $r > 1$, in accordance with logarithmic perceptual scaling (Weber–Fechner law [42]) and classical multiresolution analysis [6]. The scale derivative in (2.2) is approximated by a finite difference:

$$\tau_{\mathbf{p},I}(\sigma_s) \approx \frac{g_{\mathbf{p},I}(\sigma_s) - g_{\mathbf{p},I}(\sigma_{s-1})}{\sigma_s - \sigma_{s-1}} = (I * \Psi_s)(\mathbf{p}),$$

where Ψ_s is a DoG filter:

$$\Psi_s(\mathbf{x}) = \gamma_{\sigma_s}(\mathbf{x}) - \gamma_{\sigma_{s-1}}(\mathbf{x}).$$

The dyadic case corresponds to $r = 2$; we shall refer to the general case as r-adic.

The family $\{\Psi_s\}$ behaves as a set of band-pass pseudo-wavelet filters:

- they satisfy the dilation property $\Psi_{s+1}(\mathbf{x}) = \Psi_s(\mathbf{x}/r)$;
- their Fourier transform is bounded away from zero except near $\omega = 0$, where they fail the admissibility condition but retain finite energy over $[\varepsilon, \infty)$ for any $\varepsilon > 0$;
- the central frequency scales as $\omega_s \propto (r^s \sigma_0)^{-1}$, ensuring logarithmic coverage of scales.

Thus, the pointwise multiscale texture operator can be interpreted as analyzing the image through an r-adic family of band-pass filters, connecting the method with classical wavelet and multiresolution techniques [6].

Although the discretization yields DoG filters, the proposed framework does not use isolated filter responses as descriptors. Instead, each point is represented through the evolution of the operator across the scale axis, forming a multiscale contextual signature.

Interpretation of the r-adic discretization. The choice of an r-adic scale grid is not only computationally convenient but also avoids redundancy in the multiscale representation. Consecutive Gaussian kernels with arbitrarily close standard deviations would yield highly correlated values of $g_{\mathbf{p},I}(\sigma)$, providing little additional information. The r-adic sampling $\sigma_s = r^s \sigma_0$ ensures a logarithmic

separation between scales, analogous to the construction of classical wavelet families, and is consistent with the Weber–Fechner law, which models perceptual sensitivity as a logarithmic function of stimulus magnitude.

From a scale–space perspective, the family $\{\Psi_s\}$ provides a natural decomposition of the signal into frequency bands whose central frequencies are distributed logarithmically. Thus, the discrete texture operator

$$\tau_{\mathbf{p},I}(r^s\sigma_0) = (I * \Psi_s)(\mathbf{p})$$

quantifies how the contextual identity of the point $\mathbf{p}$ evolves when the effective observation window is scaled by a factor r. Small values of $\tau_{\mathbf{p},I}(r^s\sigma_0)$ indicate that the neighborhood of $\mathbf{p}$ is homogeneous at that resolution, while large values correspond to significant structural transitions. Strong variations across scales reveal that the point participates in multiple texture regimes, reflecting different levels of detail in its spatial context.

Viewed in the joint space–scale (or space–frequency) domain, the texture operator therefore provides a compact multiresolution characterization of the local structure around $\mathbf{p}$, capturing how the surrounding texture evolves as the scale varies. This perspective reinforces the connection between the proposed method and well-established multiscale representations such as wavelet and scale–space analysis.

Theorem 2.2. *Let*

$$\psi_s(\mathbf{x}) = \frac{1}{r^{2s-1}\sigma_0^2\sqrt{2\pi}} e^{-\frac{x^2}{r^{2s+1}\sigma_0^2}} - \frac{1}{r^{2s-2}\sigma_0^2\sqrt{2\pi}} e^{-\frac{x^2}{r^{2s-1}\sigma_0^2}}, \quad \forall \mathbf{x} \in \mathbb{R}$$

and define $\Psi_s(\mathbf{x}) = \prod_{k=1}^{n} \psi_s(x_k)$. *Then* ψ_s *satisfies the following properties:*

(1) ***Dilation.*** *For all* $x \in \mathbb{R}$,

$$\psi_s\left(\frac{x}{r}\right) = \psi_{s+1}(x).$$

(2) ***Pseudo-wavelet property.*** *Although* $\widehat{\psi_s}(0) \neq 0$, *the function* ψ_s *satisfies*

$$C_\varepsilon(\psi_s) = \int_\varepsilon^\infty \frac{|\widehat{\psi_s}(\omega)|^2}{\omega}\, d\omega < \infty$$

for every $\varepsilon > 0$, *and therefore behaves as a pseudo-wavelet.*

(3) ***Band-pass filtering.*** *The 1D function* ψ_s *acts as a band-pass filter whose effective bandwidth is*

$$\Delta\omega_s = \sqrt{\frac{\displaystyle\int_0^\infty \omega^2 |\widehat{\psi_s}(\omega)|^2\, d\omega}{\displaystyle\int_0^\infty |\widehat{\psi_s}(\omega)|^2\, d\omega}} \approx 1.0331\, \frac{1}{r^s\sigma_0}.$$

(4) ***Optimal logarithmic scale coverage.*** *The r-adic spacing provides constant relative resolution:*

$$\frac{\Delta\sigma_s}{\sigma_s} = \frac{1}{r}, \qquad \forall s \in \mathbb{Z}.$$

2.7. Interpretation and scope of the pointwise texture operator

From a theoretical standpoint, the proposed operator extends fixed-scale texture descriptors by encoding the scale derivative of a diffusion-generated contextual field rather than static local statistics. Its novelty does not reside in Gaussian filtering itself—which follows classical scale–space theory—but in the functional definition of a contextual multiscale signature, its scale-derivative interpretation, and its role as a driver of multiscale coherence within a coherence-driven region-growing framework.

In contrast to classical scale–space frameworks based on derivatives evaluated at individual scales, the present formulation characterizes each point through the evolution of its contextual signature over a scale interval.

To further clarify the technical contribution beyond classical multiscale filtering frameworks, the novelty of the proposed operator can be summarized along four complementary dimensions:

(1) Functional contextual representation. The operator is formulated through a continuous contextual multiscale signature rather than as a finite set of filter responses, enabling the characterization of contextual evolution across observation scales.
(2) Scale-derivative interpretation. The pointwise multiscale texture operator is interpreted as a differential measure of contextual identity evolution, rather than as a band-pass or blob-detection response as in DoG/LoG schemes.
(3) Multiscale coherence similarity. Segmentation decisions are based on the similarity between full operator signatures across a scale interval, measured through an L^2 functional coherence criterion.
(4) Segmentation integration. The pointwise multiscale texture operator is embedded within a coherence-driven region-growing framework in which region propagation is governed by multiscale structural compatibility rather than local intensity or gradient magnitude.

These elements collectively distinguish the proposed framework from conventional DoG/LoG-based texture descriptors, whose use typically remains restricted to localized feature detection or fixed-scale statistical characterization.

Within this framework, the operator $\tau_{\mathbf{p},I}(\sigma)$ provides a functional characterization of contextual identity across scales. By encoding how local contextual information evolves as the observation window expands, it captures transitions between homogeneous and heterogeneous regimes rather than acting as a classical edge or gradient-based feature.

While in this work the operator is used to guide a region-growing algorithm, its definition is independent of segmentation and can be naturally exploited in other contexts, including texture classification, anomaly detection, registration based on multiscale similarity, and the analysis of structural patterns in biological or industrial imaging. This highlights the operator as a general-purpose multiscale operator rather than a task-specific segmentation feature.

2.7.1. Functional texture signature

The pointwise multiscale texture operator introduced in this work is naturally formalized as a functional signature along the scale axis.

Definition 2.1 (Pointwise multiscale texture signature)**.** *Let $I \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$ and let $g(p, \sigma; I)$ denote the contextual multiscale signature defined in (2.1) for $\sigma \in [\sigma_{\min}, \sigma_{\max}]$. We define the associated pointwise*

multiscale texture signature as the function

$$\mathcal{T}_I(p) : [\sigma_{\min}, \sigma_{\max}] \to \mathbb{R}, \qquad \mathcal{T}_I(p)(\sigma) = \partial_\sigma g(p, \sigma; I),$$

whenever $\partial_\sigma g(p, \sigma; I)$ exists in the classical sense, and in the distributional sense otherwise. For fixed p, we view $\mathcal{T}_I(p)$ as an element of $L^2([\sigma_{\min}, \sigma_{\max}])$.

In contrast to classical DoG/LoG descriptors, which are typically based on isolated responses at individual scales, the proposed operator represents each point p by the full scale-dependent function $\mathcal{T}_I(p)$, regarded as a curve in the functional space $L^2([\sigma_{\min}, \sigma_{\max}])$ endowed with the usual L^2 norm in scale, which will later be used to quantify multiscale coherence between texture signatures. This formulation should be understood as a change of representation: Instead of encoding local structure through a finite collection of filter responses, each point is associated with a scale-dependent trajectory describing the evolution of its contextual neighborhood under Gaussian diffusion.

The construction of the contextual multiscale signature $g(p, \sigma; I)$ and its scale derivative is grounded in standard Gaussian scale–space theory. In particular, the diffusion identity, mass conservation, and the approximation of scale derivatives by DoG filters under logarithmic scale sampling are classical results that provide the analytical foundation of the present framework.

2.7.2. Relation to classical DoG/LoG constructions

While the discrete implementation of the scale derivative involves DoG filters, the proposed framework is not equivalent to classical DoG/LoG descriptors. In standard scale–space feature detection, DoG/LoG responses at individual scales are used as local indicators of edges, blobs, or keypoints. In contrast, the present work treats the scale derivative as a trajectory in scale, and segmentation is driven by the coherence between entire trajectories $\mathcal{T}_I(p)$ and $\mathcal{T}_I(q)$ over a scale interval. This functional and contextual viewpoint departs from conventional uses of DoG/LoG in the literature, where responses are typically interpreted independently at each scale, and is central to the methodological contribution of the operator.

2.8. Coherence-driven region-growing framework

The segmentation strategy proposed in this work relies on an iterative region-growing procedure guided by the multiscale coherence criterion. While the previous sections introduced the theoretical formulation of the texture operator and its scale–space properties, we now describe the practical algorithm used to select and expand the segmented region in a discrete image setting.

Definition 2.2 (Multiscale coherence distance)**.** *Given two points p, q and their texture signatures $\mathcal{T}_I(p), \mathcal{T}_I(q) \in L^2([\sigma_{\min}, \sigma_{\max}])$, we define their relative multiscale texture dissimilarity as*

$$d(p, q) = \frac{\left\|\mathcal{T}_I(p) - \mathcal{T}_I(q)\right\|_{L^2([\sigma_{\min}, \sigma_{\max}])}}{\max\left\{\|\mathcal{T}_I(p)\|_{L^2([\sigma_{\min}, \sigma_{\max}])}, \|\mathcal{T}_I(q)\|_{L^2([\sigma_{\min}, \sigma_{\max}])}\right\}}.$$

This functional distance is invariant under global intensity scaling and aggregates discrepancies between scale–space trajectories across the entire interval $[\sigma_{\min}, \sigma_{\max}]$. In the discrete implementation, the L^2-norm is approximated by trapezoidal numerical integration along the sampled scale grid, leading to the coherence functional used in Section 2.8.

Proposition 2.1 (Stability of the texture signature)**.** *Let $I, J \in L^2(\mathbb{R}^n)$ and let $[\sigma_{\min}, \sigma_{\max}] \subset (0, \infty)$ be a fixed scale interval. For any point $p \in \mathbb{R}^n$, define the pointwise multiscale texture signature as*

$$\mathcal{T}_I(p)(\sigma) = \frac{\partial}{\partial\sigma}(I * \gamma_\sigma)(p),$$

where $\gamma_\sigma(\mathbf{x}) = (2\pi\sigma^2)^{-n/2} \exp(-\|\mathbf{x}\|^2/(2\sigma^2))$ is the Gaussian kernel. Then there exists a constant $C > 0$, depending only on n, $\sigma_{\min}$, and $\sigma_{\max}$ such that

$$\|\mathcal{T}_I(p) - \mathcal{T}_J(p)\|_{L^2([\sigma_{\min},\sigma_{\max}])} \leq C \|I - J\|_{L^2(\mathbb{R}^n)}.$$

In particular, the estimate is uniform in p.

The full proof is provided in Appendix.

2.8.1. Initialization and parameter selection

The algorithm starts from a reference set K_0, chosen so as to be fully contained within the target region. In practice, K_0 is defined as a small neighborhood (typically a square or ball) around an expert-selected seed point.

The scale variable is discretized according to an r-adic rule,

$$\sigma_k = r^{-k} w,$$

where w denotes the image width and k takes a finite set of integer values determining the resolution levels used in the analysis. This parametrization is consistent with the r-adic discretization introduced in Section 2.6, with the image width w providing a natural upper bound for the largest meaningful scale.

The base $r > 1$ is a parameter to be chosen by the expert, depending on the importance to be given to small variations in the texture signature. Larger values of r produce a coarser scale sampling, while values close to 1 yield a finer sampling of scales, emphasizing subtle texture changes. A convenient parametrization is $r = 2^\alpha$ with $\alpha \in (0, 1]$, where small α gives fine sampling and $\alpha = 1$ recovers the standard dyadic scheme.

This choice is constrained by two practical considerations: For very large negative k, the effective radius becomes smaller than a pixel and loses meaning, whereas for values of σ_k comparable to w/r or larger, the resulting averages become nearly constant across the image. The selected range of scales therefore reflects both the image resolution and the expected size of relevant structures.

A texture coherence threshold $\nu > 0$ is also fixed. This parameter controls the admissible relative deviation between the multiscale texture of a candidate point and that of the reference region K_0. In practice, ν has a clear interpretation as a relative tolerance in multiscale texture space and does not require extensive fine tuning across similar images.

Finally, a discrete growth pattern is defined on the image grid. Given a binary image A that takes value 1 on the current region and 0 elsewhere, the associated candidate set $\widetilde{A}$ is defined by standard 4-neighbor connectivity:

$$(x_i, y_j) \in \widetilde{A} \iff (x_{i\pm1}, y_j) \in A \text{ or } (x_i, y_{j\pm1}) \in A.$$

This choice ensures a simple and robust growth mechanism consistent with the underlying pixel lattice.

2.8.2. Implementation details and parameter configuration

Scale discretization and numerical derivative. The scale variable is discretized according to an r-adic sampling strategy. In synthetic validation experiments (e.g., the mandrill study), dyadic sampling ($r = 2$) was adopted, with scale indices

$$s \in \{-4, -5, -6, -7\}$$

yielding four analysis scales. The Gaussian radii are defined as

$$\sigma_s = 2^s W,$$

where W denotes the image width.

For CBCT datasets, analogous logarithmic sampling grids were defined according to the spatial resolution of the volumes and the expected anatomical scale of the structures of interest.

The scale derivative defining the texture operator is numerically approximated through first-order finite differences:

$$\tau_{\mathbf{p},I}(\sigma_s) \approx \frac{g_{\mathbf{p},I}(\sigma_s) - g_{\mathbf{p},I}(\sigma_{s-1})}{\sigma_s - \sigma_{s-1}}.$$

Gaussian convolution. Gaussian smoothing is computed through discrete convolution using standard separable implementations. In practice, filtering is performed using Octave's `imsmooth` function, which implements isotropic Gaussian kernels with variance σ^2.

Seed region and neighborhood definition. The initial reference set K_0 is defined as a dilated neighborhood around an expert-selected seed point. In the synthetic experiments, the dilation radius corresponds to 10% of the image width, implemented via morphological dilation.

Texture coherence is evaluated over a local neighborhood defined by a square (2D) or cubic (3D) structuring element of size determined by the parameter *entorno*.

Connectivity and region propagation. Region growing is driven by lattice adjacency. Neighborhood expansion is defined through standard 4-connectivity in 2D grids, implemented via morphological dilation, with its natural extension to volumetric adjacency in 3D datasets.

Texture coherence threshold. Segmentation decisions are governed by the relative coherence threshold ν. Unless otherwise specified, a value $\nu = 0.5$ is used in synthetic experiments. In CBCT segmentations, the operational threshold was set to $\nu = 0.3438$, empirically determined to ensure stable region propagation while preventing leakage across anatomical boundaries.

Numerical integration. The L^2 coherence functional is computed through trapezoidal numerical integration along the discretized scale axis, ensuring consistent approximation of the continuous formulation.

Stopping criterion. The region-growing process iterates until stabilization, defined as the first iteration for which no additional neighboring points satisfy the coherence condition. This corresponds to the first stationary segmented set.

2.8.3. Iterative region-growing procedure

Starting from K_0, the algorithm constructs a sequence of nested sets $\{K_j\}_{j\geq 0}$ as follows. At iteration j, we consider the set of candidate points adjacent to the current region K_j according to the discrete connectivity rule described above.

$$\widetilde{K}_j = \{\mathbf{x} \in D : \mathbf{x} \text{ is adjacent to } K_j\}.$$

The candidate set $\widetilde{K}_j$ is first formed by adding all grid points adjacent to K_j according to the connectivity rule described above. The next iterate K_{j+1} is then defined by retaining only those candidate points whose multiscale texture is sufficiently coherent with that of the initial reference set K_0.

Given an arbitrary point $\mathbf{p}$, we define the coherence criterion as the minimum relative $L^2([\sigma_{\min}, \sigma_{\max}])$ error between its contextual multiscale signature and the signatures of the points belonging to the initial set K_0:

$$\min_{\mathbf{q}\in K_0} \sqrt{\frac{\displaystyle\int_{\sigma_{\min}}^{\sigma_{\max}} \left|\tau_{\mathbf{p},I}(\sigma) - \tau_{\mathbf{q},I}(\sigma)\right|^2 d\sigma}{\displaystyle\int_{\sigma_{\min}}^{\sigma_{\max}} \left|\tau_{\mathbf{q},I}(\sigma)\right|^2 d\sigma}}.$$

To compute this value numerically, we approximate the integral by means of the trapezoidal rule:

$$e_{K_0}(\mathbf{p}) = \min_{\mathbf{q}\in K_0} \sqrt{\frac{\displaystyle\sum_k \Delta\sigma_k \left(\left|\tau_{\mathbf{p},I}(\sigma_k) - \tau_{\mathbf{q},I}(\sigma_k)\right|^2 + \left|\tau_{\mathbf{p},I}(\sigma_{k-1}) - \tau_{\mathbf{q},I}(\sigma_{k-1})\right|^2\right)/2}{\displaystyle\sum_k \Delta\sigma_k \left(\left|\tau_{\mathbf{q},I}(\sigma_k)\right|^2 + \left|\tau_{\mathbf{q},I}(\sigma_{k-1})\right|^2\right)/2}},$$

where $\Delta\sigma_k = \sigma_k - \sigma_{k-1}$.

A point $\mathbf{p}_i \in \widetilde{K}_j$ is accepted if the value of this function is below a threshold ν chosen by the expert:

$$\mathbf{p}_i \in K_{j+1} \quad \Longleftrightarrow \quad e_{K_0}(\mathbf{p}_i) < \nu.$$

Using the initial reference set K_0 rather than the evolving region avoids drift effects and anchors the segmentation to the intrinsic texture of the target structure.

By construction, the sequence $\{K_j\}$ is non-decreasing and bounded, and therefore stabilizes after a finite number of iterations. The final segmented region is defined as the first set K_j such that $K_{j+1} = K_j$. This stopping condition corresponds to a state in which no additional neighboring points satisfy the texture coherence criterion.

Basic properties of the discrete region-growing scheme. Let $\{K_j\}_{j\geq 0}$ be the sequence of regions generated by the iterative procedure described above. Then:

(1) **Monotonicity.** The sequence is non-decreasing:

$$K_j \subseteq K_{j+1}, \quad \forall j \geq 0.$$

(2) **Boundedness.** Since the image domain D is finite, the sequence $\{K_j\}$ is bounded by D.

(3) **Finite termination.** There exists $j^* \in \mathbb{N}$ such that

$$K_{j^*} = K_{j^*+1},$$

that is, the algorithm terminates after a finite number of iterations.

(4) **Stability with respect to rejection.** Points that do not satisfy the coherence criterion are permanently excluded from subsequent iterations, preventing oscillatory behavior.

These properties ensure that the discrete segmentation process is well defined and terminates in a finite number of steps. A full theoretical analysis of the sensitivity of the algorithm with respect to the initialization, the coherence threshold, and the scale discretization, as well as convergence rates, remains an interesting direction for future work.

2.8.4. Interpretation and convergence behavior

The use of the initial reference set K_0 as the reference operator baseline ensures that the growth of the region is driven by intrinsic multiscale properties rather than by mere spatial proximity. Points belonging to the same anatomical or structural region exhibit similar scale–space evolution of their texture signatures and are therefore progressively incorporated, whereas points across boundaries fail the coherence test.

The monotonicity of the construction and the exclusion of rejected points guarantee finite stabilization of the discrete iterative process. Since the sequence of segmented sets is non-decreasing and bounded within a finite image domain, the algorithm necessarily reaches a stationary configuration after a finite number of iterations. This discrete stabilization should be distinguished from a full analytical convergence proof in the continuous setting, which remains outside the scope of the present work.

Sensitivity to seed initialization and region of interest selection. Although the proposed segmentation framework follows a region-growing strategy, its evolution is not driven by local intensity similarity but by multiscale texture coherence. Consequently, the influence of the initial seed is partially mitigated by the functional nature of the operator signatures guiding the propagation.

Empirical tests conducted under moderate perturbations of the seed location indicate that, provided the initialization remains within the target anatomical structure, the final segmented region remains stable. This behavior reflects the intrinsic coherence of multiscale texture signatures within homogeneous regions, which naturally drives the growth toward a consistent structural support.

If the seed is placed outside the target structure, the iterative process stabilizes around the texture regime corresponding to that initialization, as expected in any region-growing framework. This behavior confirms that the method is structurally driven rather than intensity-driven.

Similarly, the bounding region used to localize the segmentation acts only as a computational constraint. Variations in its size within anatomically plausible margins do not significantly affect the outcome since candidate points are accepted exclusively on the basis of the multiscale coherence criterion. Points outside the structural regime of interest are therefore naturally rejected during propagation. These observations suggest that the segmentation outcome is primarily governed by multiscale structural coherence rather than by the precise initialization configuration.

Theoretical interpretation of the multiscale coherence criterion. The coherence functional introduced in Section 2.8.3 admits a natural interpretation within the functional setting defined by the

texture operator. Since the pointwise multiscale texture operator $\tau_{\mathbf{p},I}(\sigma)$ is defined as a scale-dependent function over the interval $[\sigma_{\min}, \sigma_{\max}]$, the comparison between two points cannot be reduced to a single-scale operator response but must be formulated as a similarity measure between functions.

Within this framework, the L^2 norm provides a mathematically consistent metric to evaluate the discrepancy between operator signatures. This choice is supported by several theoretical considerations. First, Gaussian scale–space representations define a diffusion-generated functional space in which energy norms are stable under convolution and small perturbations. Consequently, the L^2 distance between operator signatures is robust to moderate noise and acquisition blur.

Second, the normalization term in Section 2.8.3 yields a relative coherence measure, ensuring invariance under global intensity scaling and enabling consistent comparison across regions exhibiting different multiscale operator magnitudes.

Finally, integrating the discrepancy across a scale interval enforces structural compatibility over multiple observation scales. The segmentation decision is therefore based on the persistence of contextual similarity across scales rather than on isolated local features, constituting a stronger homogeneity condition consistent with the diffusion-based construction of the operator.

2.9. Example: characterization of different textures in an image

To illustrate how the pointwise multiscale texture operator $\tau_{\mathbf{p},I}(\sigma)$ provides meaningful information, we conducted a numerical experiment using a well-known test image: the mandrill (*Mandrillus sphinx*). The image has a resolution of 270×270 pixels. Five distinct points were selected in order to perform the comparison (see the points marked in Figure 5a).

The pointwise texture was computed at each selected location for the following scale values:

$$\sigma_s = 2^{-s} \cdot \text{width}, \quad \text{with } s = -11, -10, \ldots, -3 \quad \text{and} \quad \text{width} = 270.$$

This scale range is chosen to span from very fine resolutions, comparable to the pixel size, to scales of the order of the main structures present in the image. Small values of σ emphasize local microtextural variations and high-frequency details, whereas larger scales encode more global contextual information associated with macrostructures. The r-adic discretization ensures an efficient multiscale exploration while avoiding redundancy between consecutive scales.

This choice is consistent with the r-adic scale discretization introduced in Section 2.6, adapted here to the image resolution for visualization and analysis purposes.

As observed in Figure 5b, each resolution level σ corresponds to macrostructures (for large σ) or microstructures (for small σ). Intermediate scales, on the other hand, contain relevant information about specific anatomical details, such as the structure of the animal's snout.

This observation highlights the inherently multiscale nature of the proposed operator. In particular, intermediate scales tend to be the most discriminative from a textural perspective, as they capture persistent structural transitions without being dominated by noise or excessive smoothing. Consequently, in this example, $\tau_{\mathbf{p},I}(\sigma)$ acts as a contextual multiscale operator that integrates information across scales rather than as a purely local edge detector.

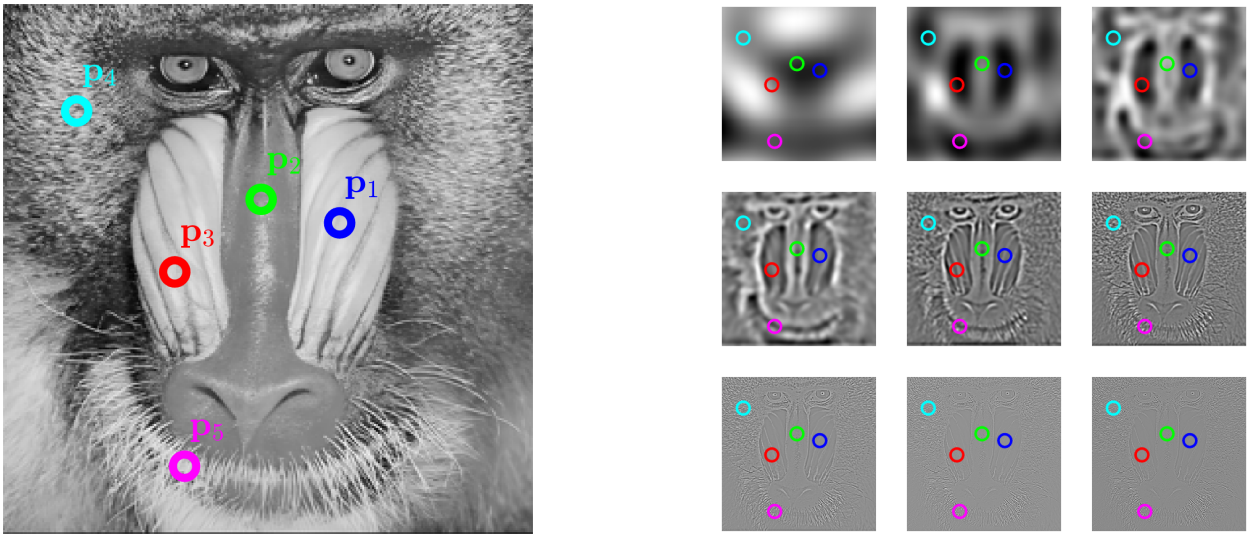

(a) Original image with the five selected analysis points.

(b) Visualization of extracted structures at different scales σ.

Figure 5. Experimental configuration and multiscale structural visualization.

This experiment is intended as an interpretability study illustrating the discriminative capability of the pointwise multiscale texture operator rather than as a performance evaluation of the segmentation algorithm.

By plotting the functions $\tau_{\mathbf{p},I}(\sigma)$ obtained for each point (see Figure 6), clear differences are observed between points $\mathbf{p}_4$ and $\mathbf{p}_5$ with respect to the others. Points $\mathbf{p}_1$, $\mathbf{p}_2$, and $\mathbf{p}_3$ exhibit similar behavior, with the exception of a slight divergence in the trend of point $\mathbf{p}_3$ at larger values of σ.

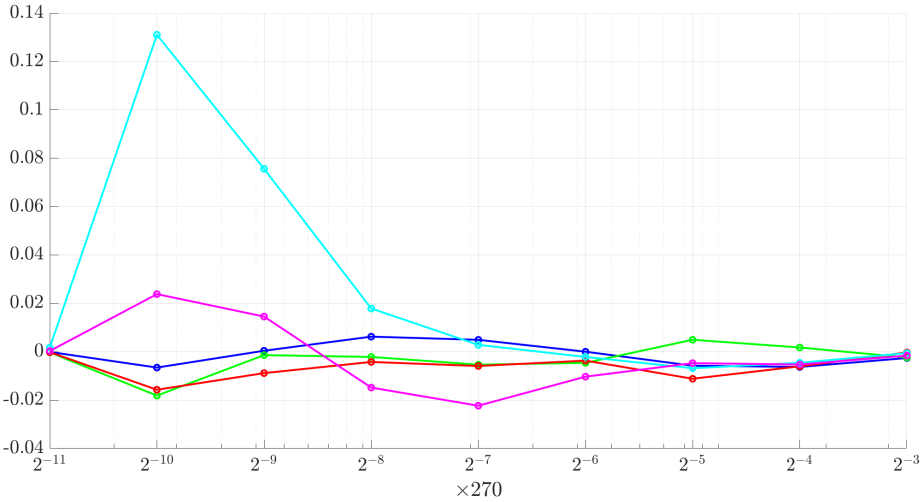

Figure 6. Pointwise texture functions $\tau_{\mathbf{p},I}(\sigma)$ computed for the five selected points.

To quantitatively compare the multiscale texture behavior at different points, we introduce a measure of relative similarity between the texture at a point $\mathbf{p}$ and that at another point $\mathbf{q}$ using the previous expression on $\mathbf{q}$ and the set $\{\mathbf{p}\}$ defined in Section 2.8.3:

$$d(\mathbf{p}, \mathbf{q}) = e_{\{\mathbf{p}\}}(\mathbf{q}). \tag{2.4}$$

The resulting relative texture dissimilarity matrix is reported in Table 1.

Table 1. Relative texture dissimilarity matrix $d(\mathbf{p}_k, \mathbf{p}_\ell)$ between the selected analysis points, computed according to (2.4). Rows correspond to $\mathbf{p}_k$ and columns to $\mathbf{p}_\ell$. Diagonal entries are zero by definition.

	$\mathbf{p}_1$	$\mathbf{p}_2$	$\mathbf{p}_3$	$\mathbf{p}_4$	$\mathbf{p}_5$
$\mathbf{p}_1$	0.00	1.48	0.84	2.70	1.59
$\mathbf{p}_2$	2.10	0.00	2.42	4.51	2.36
$\mathbf{p}_3$	0.64	1.31	0.00	2.31	1.01
$\mathbf{p}_4$	0.98	1.15	1.10	0.00	0.95
$\mathbf{p}_5$	1.03	1.08	0.85	1.70	0.00

Smaller values of $d(\mathbf{p}_k, \mathbf{p}_\ell)$ indicate higher multiscale textural similarity between the corresponding points.

As expected given their symmetric location with respect to the snout, the point whose texture is most similar to that of $\mathbf{p}_1$ is $\mathbf{p}_3$, and vice versa. This result validates the ability of the descriptor $\tau_{\mathbf{p},I}(\sigma)$ to group regions with similar multiscale textural properties. It is worth noting that this grouping is driven by multiscale contextual coherence rather than by mere spatial proximity, confirming that the operator captures intrinsic textural similarities between image regions.

Stability with respect to blur and additive noise

The inclusion of degraded versions of the mandrill image is methodologically motivated. Dental CBCT data inherently exhibit blur, noise, and weakly separated textures due to acquisition physics and reconstruction processes. Evaluating the segmentation framework under controlled degradations therefore provides a proxy setting to assess robustness under conditions analogous to clinical imaging.

A key motivation for the proposed segmentation framework is its robustness under image degradation. Since the contextual multiscale signature is constructed through Gaussian averaging and its scale derivative, the resulting operator inherently suppresses high-frequency perturbations, consistent with the theoretical properties discussed in Sections 2.5 and 2.6.

To assess this stability empirically, we consider the 2D mandrill image and fix the same initial seed point $\mathbf{p}_3 = (x_3, y_3)$ in all experiments. The reference set K_0 is defined as a square neighborhood centred at $\mathbf{p}_3$ with side length equal to 10% of the image width, corresponding to 27 pixels in each direction:

$$K_0 = \left\{\mathbf{p} = (x, y) \in \mathbb{R}^2 : \max(|x - x_3|, |y - y_3|) < 27\right\}.$$

The scale parameters are chosen as $\sigma_s = 2^{-s} \times 270$ for $s = 4, 5, 6, 7$, and the texture coherence threshold is fixed to $\nu = 0.5$. All remaining parameters of the algorithm are kept unchanged throughout the experiments.

We first apply the segmentation algorithm to the original image. The resulting region, shown in Figure 7, serves as a baseline for comparison. Next, the image is degraded by applying a local averaging filter, simulating convolutional blur induced by the acquisition process. This operation is used exclusively to simulate acquisition-induced blur and does not correspond to any averaging step within the proposed segmentation method. Using exactly the same parameters, the segmentation obtained from the blurred image is shown in Figure 8. The segmented region remains visually and geometrically consistent with the original case, indicating that moderate smoothing does not disrupt the multiscale texture coherence criterion.

Finally, to evaluate robustness against stochastic perturbations, additive white Gaussian noise is superimposed on the blurred image. The noise level is selected such that the peak signal-to-noise ratio satisfies

$$\mathrm{PSNR} = 10\log_{10}\left(\frac{1}{0.05^2}\right)\ \mathrm{dB} \approx 26.02\ \mathrm{dB}.$$

The segmentation result obtained under these conditions is displayed in Figure 9. Despite the combined effects of convolutional blur and additive noise, the extracted region remains stable and closely matches those obtained in the previous experiments.

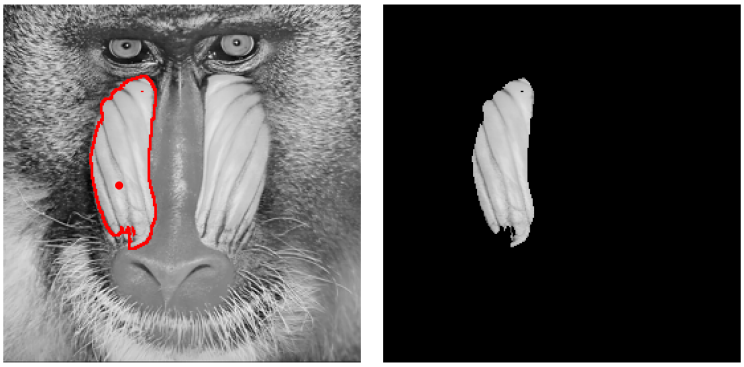

Figure 7. Segmentation obtained from the original mandrill image.

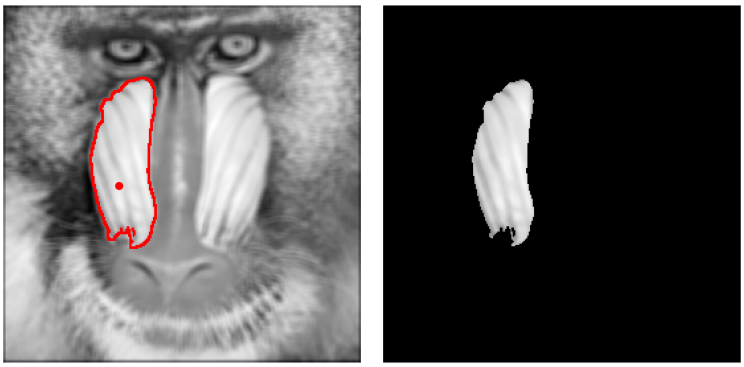

Figure 8. Segmentation obtained after local averaging blur.

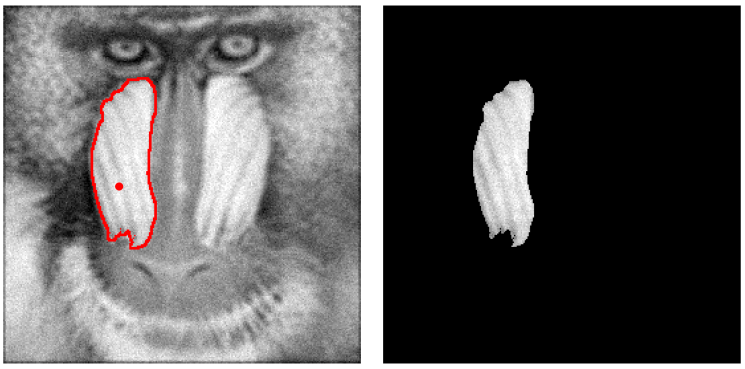

Figure 9. Segmentation obtained after blurring and additive white Gaussian noise (PSNR $\approx$ 26.02 dB).

This robustness can be attributed to the weighted averaging inherent in the construction of the contextual multiscale signature and to the integration of texture information across several scales. High-frequency noise is attenuated by Gaussian smoothing, while the ℓ^2-based coherence criterion enforces consistency of the entire scale-dependent texture profile rather than relying on isolated local features. As a consequence, spurious fluctuations do not propagate during the region-growing process, and the algorithm preserves structurally coherent anatomical morphology of the segmented region even under significant image degradation.

This experiment thus connects controlled synthetic validation with clinically relevant acquisition conditions, supporting the robustness of the proposed framework under degradation effects.

3. Results

This section presents the results obtained by applying the proposed coherence-driven region-growing framework to CBCT data. We report both quantitative volumetric estimates and qualitative segmentation outcomes, with the aim of assessing the geometric consistency and descriptive capacity of the method across heterogeneous subjects.

The experiments reported in this section illustrate the applicability of the proposed multiscale texture operator under clinically realistic CBCT acquisition conditions. The evaluation therefore focuses on structural coherence, geometric consistency, and robustness of the operator-driven segmentation process rather than on competitive benchmarking against supervised segmentation frameworks.

This evaluation philosophy is consistent with methodological studies introducing analytical operators, where validation focuses on stability, interpretability, and applicability rather than dataset-scale performance comparison.

3.1. Pulp chamber volumetric estimates

Table 2 summarizes the demographic characteristics of the studied subjects together with the final pulp chamber volumetric estimates. The dataset includes male and female individuals spanning a wide age range, as well as subjects with different anatomical characteristics and demographic attributes. The

reported volumes correspond to the 3D reconstructions obtained by aggregating the segmented regions across consecutive CBCT slices.

Table 2. Demographic information and final chamber volume of the studied subjects.

Patient ID	Age (years)	Sex	Indigenous	Volume (mm^3)
1	18	F	No	27.942
2	17	F	No	30.699
3	63	F	No	6.013
4	58	F	Yes	7.935
5	48	F	Yes	9.537
6	45	F	No	21.312
7	40	M	No	15.080
8	29	F	No	12.119
9	58	F	No	5.664
10	70	M	No	10.610
11	76	M	No	8.966
12	49	M	No	7.906
13	54	F	No	7.479
14	55	M	No	13.394
15	62	F	No	8.935
16	41	M	No	15.822
17	60	F	No	6.642
18	31	F	No	9.693
19	67	F	No	15.205
20	50	M	No	15.574
21	33	M	No	10.394

Figure 10 displays the scatter plot of the final pulp chamber volume as a function of age. An overall decreasing trend of the estimated pulp chamber volume with increasing age is observed across the studied cohort. Although no inferential statistical testing is conducted due to the limited cohort size, this behavior aligns with established dental aging processes, in which progressive secondary dentin deposition leads to gradual reduction of pulp cavity volume. Larger volumes are predominantly associated with younger individuals, whereas older subjects tend to exhibit smaller pulp chamber volumes, although some dispersion is present.

To further explore inter-subject variability, Figure 11 shows the distribution of the final chamber volume stratified by sex using boxplots. Male subjects exhibit a higher median pulp chamber volume, whereas female subjects show greater variability, indicating a sex-related difference in the observed volume distribution. This observation is reported descriptively and should not be interpreted as statistically conclusive.

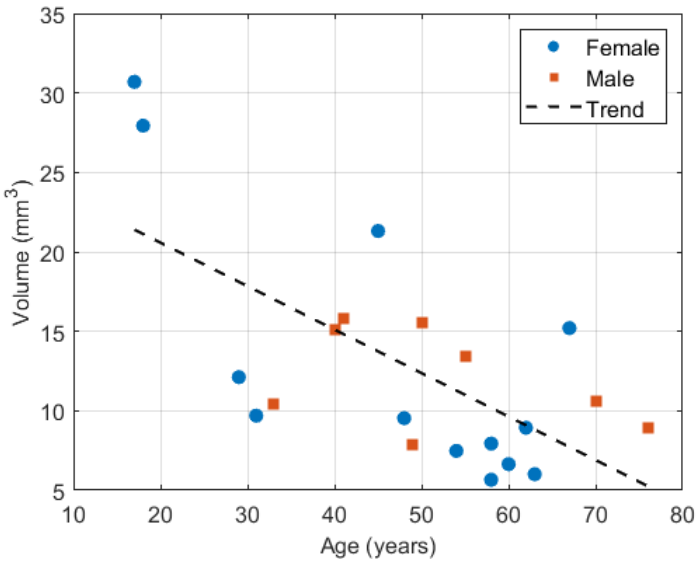

Figure 10. Scatter plot of the final chamber volume as a function of age. Male and female subjects are indicated with different markers. A decreasing trend of volume with increasing age is observed.

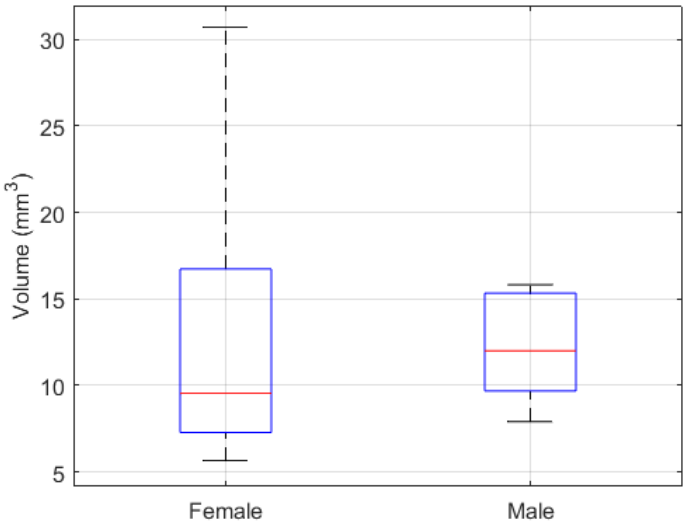

Figure 11. Boxplot of the final chamber volume stratified by sex. Male subjects exhibit a higher median pulp chamber volume, whereas female subjects show greater variability. Given the limited cohort size, these observations should be interpreted as descriptive rather than statistically conclusive.

3.2. Qualitative assessment of segmentation and 3D reconstruction

It is important to note that dental CBCT images inherently present low contrast, partial volume effects, and blurred interfaces between anatomical tissues. These acquisition characteristics impose intrinsic segmentation challenges and reflect realistic clinical imaging conditions.

Beyond quantitative measurements, the segmentations were visually inspected to evaluate spatial coherence and anatomical plausibility. Figure 12 presents representative examples corresponding to a young adult subject, an older adult subject, and an indigenous subject. For each case, consecutive CBCT slices are shown with the segmented pulp chamber superimposed.

Across all cases, the segmented regions remain spatially consistent from slice to slice, forming connected structures without evident discontinuities or anatomically implausible leakage into surrounding tissues, within the spatial resolution limits of CBCT imaging. The corresponding 3D reconstructions exhibit coherent geometries that align with the expected anatomical morphology of the pulp chamber.

These qualitative results indicate that the coherence-driven region-growing framework produces stable and geometrically meaningful segmentations across subjects with different anatomical characteristics and image quality conditions. The consistency observed in both 2D slices and 3D reconstructions supports the volumetric measurements reported above.

To perform the segmentations shown in Figure 12, images were upscaled by a factor of 2. This interpolation step improves boundary localization and is used solely to refine volumetric discretization without modifying the segmentation criterion itself.

The upscaling process resulted in 3D datasets of size $512 \times 512 \times 16$ for Patient 21, $512 \times 512 \times 18$ for Patient 15, and $512 \times 512 \times 14$ for Patient 5. In all cases, the voxel dimensions were $0.0749 \times 0.0749 \times 0.25$ mm^3. The scale parameters were instantiated according to the r-adic discretization strategy described in Section 2.6, with $r = 2^{1/4}$ and $s = -11, -10, \ldots, -2$.

For visualization and volume estimation purposes, the pulp chamber region was localized using a bounding box centered at expert-provided reference points, given in pixel coordinates along the x, y, and z directions: $(134, 133, 11)$ for Patient 21, $(103, 111, 12)$ for Patient 15, and $(157, 139, 11)$ for Patient 5. The box had physical dimensions of $0.3745 \times 0.3745 \times 1.25$ mm^3, corresponding to a size of 5 voxels along each axis from the center point. This localization step restricts the computational domain but does not alter the segmentation criterion or the volumetric estimation within the region of interest.

A constant texture coherence threshold (as defined in Section 2.8) was applied uniformly across all datasets within this region of interest, with value $\nu = 0.3438$, in order to ensure comparability of the resulting volume estimates. The resulting approximate pulp chamber volumes are summarized in the preceding table.

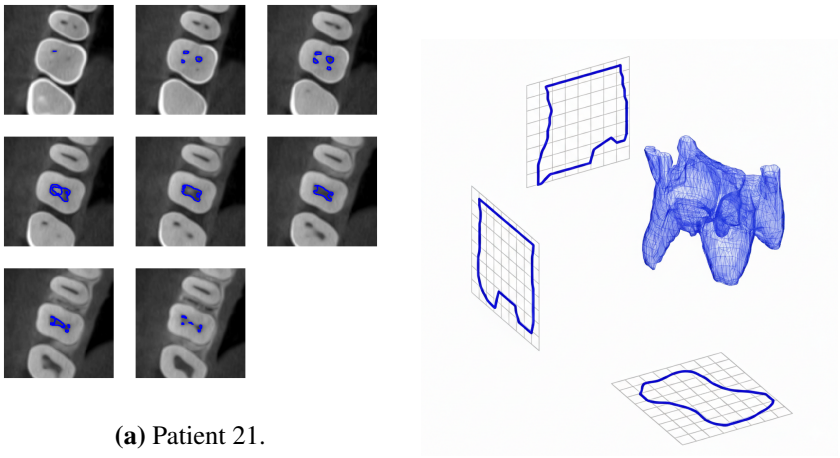

(a) Patient 21.

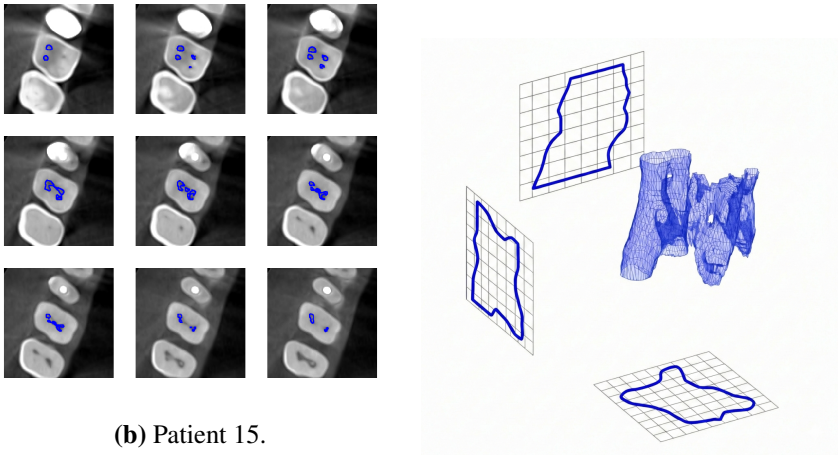

(b) Patient 15.

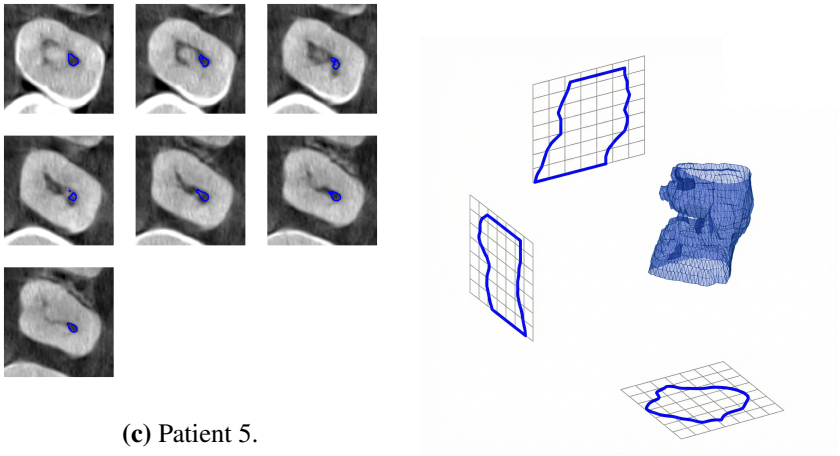

(c) Patient 5.

Figure 12. Representative examples of pulp chamber segmentation obtained with the proposed algorithm. From top to bottom: young adult subject (Patient 21), older adult subject (Patient 15), and indigenous subject (Patient 5). For each case, consecutive CBCT slices with the segmented pulp chamber superimposed and the corresponding 3D reconstruction are shown.

Expert-guided validation protocol

All segmentations presented in this study were performed under clinical expert supervision. Specialists in dental imaging evaluated the anatomical plausibility of the extracted pulp chambers and contributed to the calibration of the texture-coherence threshold, ensuring structural consistency across slices.

In several cases, the algorithm revealed subtle anatomical structures that were subsequently confirmed by expert evaluation but were not immediately apparent during the initial visual inspection of the CBCT sections.

Rather than relying exclusively on voxelwise overlap metrics, the validation protocol combined: visual inspection of segmentation boundaries; continuity assessment across consecutive CBCT slices; geometric consistency of the reconstructed 3D pulp chamber; and morphometric plausibility of the resulting volumes.

This expert-guided evaluation ensured that the segmented regions remained confined to the anatomical pulp cavity and did not exhibit leakage into dentin or surrounding bone structures.

Although exhaustive voxelwise annotations were not initially available for the full dataset, expert-annotated reference masks were subsequently generated for the subset used in the baseline comparison presented in Section 3.3, allowing the computation of overlap-based metrics such as the Dice coefficient, Intersection over Union (IoU), Precision, and Recall. However, large-scale fully annotated CBCT datasets suitable for extensive benchmark studies remain limited, and the construction of broader annotation corpora constitutes an important direction for future work.

3.3. Positioning with respect to classical segmentation baselines

The coherence-driven region-growing framework differs conceptually from classical intensity-driven or gradient-based approaches. Thresholding methods, such as Otsu-type strategies, rely on histogram separability and are therefore sensitive to low contrast and partial volume effects, both of which are intrinsic to CBCT imaging.

Similarly, conventional region-growing algorithms driven by intensity or gradient magnitude may exhibit leakage across blurred anatomical interfaces. Variational methods, including active contours and level-set formulations, require carefully tuned regularization terms to prevent boundary instability under acquisition noise.

Graph-based formulations offer global optimality properties but remain dependent on intensity or edge-weight definitions that degrade under diffuse texture transitions. Deep learning approaches, such as 3D U-Net architectures, provide powerful data-driven segmentation capabilities but require large fully annotated volumetric datasets for training, which remain limited for the anatomical structures considered in this study.

For qualitative baseline comparison, representative intensity-based segmentation methods were applied to the same CBCT slices and regions of interest used for the proposed framework.

In particular, global Otsu thresholding was employed as a fully automatic reference method, since the segmentation threshold is determined directly from the image histogram without manual parameter adjustment. After thresholding, the connected component containing the same initialization seed used for the proposed method was selected to ensure anatomical consistency and fair comparison across methods.

Classical intensity-based region growing was also considered using the same initialization seed employed by the proposed method. To avoid arbitrary parameter tuning, the tolerance parameter was automatically determined from the local intensity variability around the seed point by computing the standard deviation of a fixed 49×49 local neighborhood centered at the initialization point. This window size was chosen to provide sufficient local contextual representation while avoiding influence from adjacent anatomical structures.

No manual post-processing or case-specific parameter refinement beyond seed initialization was performed for the baseline methods. This strategy ensures that the comparative analysis is reproducible and not influenced by manual optimization of the reference methods. The qualitative results of this baseline comparison are shown in Figure 13.

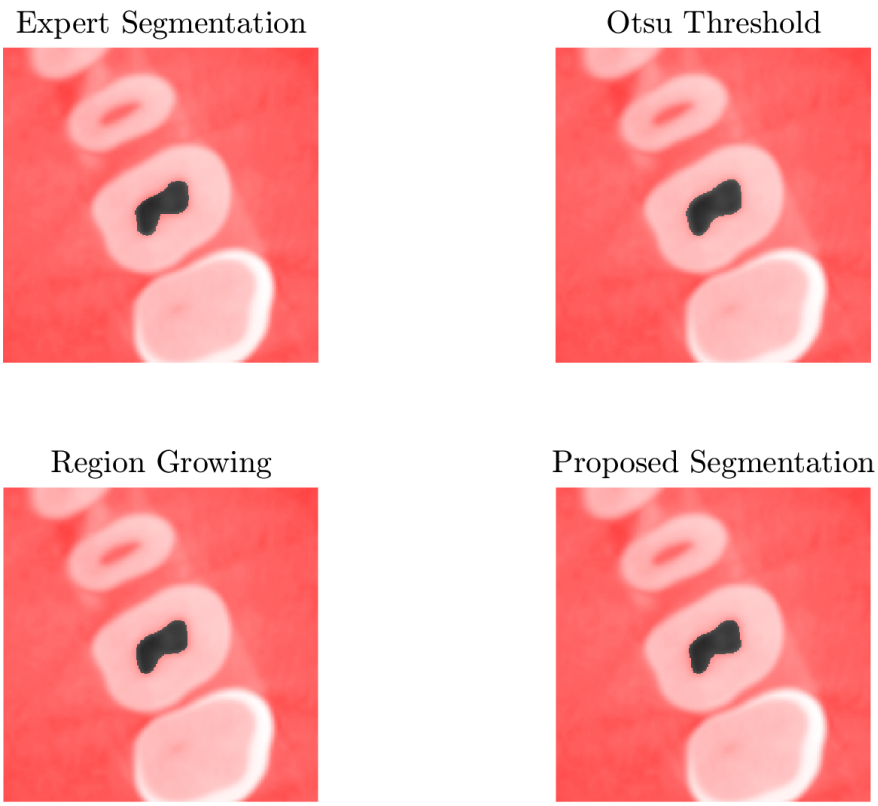

Figure 13. Qualitative comparison between expert reference segmentation, global Otsu thresholding, classical intensity-based region growing, and the proposed multiscale coherence-driven framework on a representative CBCT slice of the dental pulp chamber. All methods provide comparable delineation of the target structure, while the proposed method shows slightly improved boundary consistency with the expert segmentation and reduced sensitivity to local intensity variations near blurred anatomical interfaces.

To complement the qualitative comparison shown in Figure 13, a quantitative evaluation was performed using expert-annotated reference masks generated for all CBCT cases included in this study. Table 3 summarizes the performance of global Otsu thresholding, classical intensity-based region growing, and the proposed multiscale coherence-driven framework in terms of Dice coefficient, IoU, Precision, and Recall.

The proposed method achieved the highest mean performance in all principal overlap-based metrics, with a Dice coefficient of 0.887 ± 0.064 and an IoU of 0.801 ± 0.098, showing consistently better agreement with the expert reference masks than both Otsu thresholding and classical region growing across the full dataset. The improvement is particularly relevant in Precision, where the proposed framework reached 0.917 ± 0.066, indicating reduced false-positive propagation toward surrounding dentin and adjacent low-contrast structures. Recall values remained comparable across methods, with the proposed approach preserving similar sensitivity while improving boundary precision and reducing over-segmentation effects. In addition, the lower standard deviation observed across all evaluated cases suggests improved segmentation stability and reduced sensitivity to local intensity fluctuations and blurred anatomical boundaries.

These quantitative results support the qualitative observations of Figure 13 and confirm that the multiscale coherence criterion provides more structurally consistent delineation than purely intensity-driven segmentation strategies under typical CBCT acquisition conditions. Despite achieving the highest mean values across all overlap-based metrics, pairwise Wilcoxon signed-rank tests between the proposed method and the baseline methods did not show statistical significance at the 5% significance level ($p > 0.05$). This is consistent with the limited dataset size and the relatively small performance differences observed across clinically challenging CBCT cases.

Table 3. Quantitative comparison of representative baseline segmentation methods using expert-annotated reference masks across all CBCT cases included in this study. Mean and standard deviation are reported for Dice, IoU, Precision, and Recall.

Segmentation method	Dice	IoU	Precision	Recall
Otsu thresholding	0.852 ± 0.082	0.750 ± 0.117	0.896 ± 0.083	0.844 ± 0.177
Classical region growing	0.863 ± 0.080	0.766 ± 0.117	0.890 ± 0.094	0.863 ± 0.159
Proposed method	$\mathbf{0.887 \pm 0.064}$	$\mathbf{0.801 \pm 0.098}$	$\mathbf{0.917 \pm 0.066}$	$\mathbf{0.864 \pm 0.095}$

While the present work is not intended as a comprehensive benchmark against all segmentation paradigms, it is useful to summarize the main operational differences among representative segmentation families under CBCT imaging conditions. In particular, aspects such as training requirements, dependence on annotated datasets, and dominant computational behavior strongly influence the practical applicability of each approach. Table 4 provides a qualitative comparison of these characteristics for the proposed framework and several commonly used segmentation families.

The objective of this baseline comparison is not to establish state-of-the-art competitive performance but to verify that the proposed operator provides structurally coherent segmentation under conditions where purely intensity-driven methods typically exhibit instability. As shown in Table 4, the proposed framework occupies an intermediate position between purely intensity-based methods and fully supervised deep learning architectures. While its computational cost is higher than simple thresholding strategies due to the multiscale Gaussian filtering stage, the overall complexity remains

dominated by the computation of Gaussian convolutions across a finite number of scales, leading to an asymptotic cost of order $O(N \cdot S)$, where N denotes the number of pixels or voxels and S the number of sampled scales. In practice, this remains computationally less demanding than training-dependent 3D deep learning approaches, while preserving full interpretability of the segmentation process and avoiding the need for large annotated datasets.

Table 4. Qualitative comparison of representative segmentation method families in terms of training requirements, annotation dependence, and dominant computational behavior under CBCT imaging conditions.

Method family	Training phase	Annotation requirement	Computational behavior
Proposed multiscale operator	No	Minimal (seed only)	$O(N \cdot S)$
Thresholding methods	No	None	Low ($O(N)$)
Classical region growing	No	Minimal	Moderate
Level-set / variational methods	No	Moderate	High iterative cost
Graph-based methods	No	Moderate	High depending on graph size
3D U-Net / deep learning	Yes	Extensive	Very high (training + inference)

In contrast, the proposed method operates on multiscale texture coherence rather than on single-scale intensity or gradient cues. This functional characterization enables robust delineation of anatomical regions exhibiting weak boundaries and contextual structural transitions, as demonstrated in the presented CBCT experiments.

4. Discussion

The experimental results obtained on CBCT data provide empirical support for the theoretical properties of the proposed coherence-driven region-growing framework. In particular, the spatial coherence and geometric consistency observed across consecutive slices (see Figure 12) suggest that the region-growing procedure driven by multiscale texture coherence stabilizes into connected 3D structures with stable morphology. Within this framework, Gaussian convolution acts as an implicit regularization mechanism, while the r-adic discretization of the scale variable ensures a logarithmic separation between successive scale levels of the pointwise multiscale texture operator.

As established by scale–space theory, Gaussian filtering generates a family of representations that are linear, isotropic, and causality preserving. Within this framework, the contextual multiscale signature evolves smoothly with respect to the scale parameter, and its derivative with respect to scale provides a robust indicator of structural transitions. The observed spatial coherence of the segmented regions suggests that the L^2-based texture coherence criterion effectively constrains the evolution of the region-growing process, preventing spurious oscillations and promoting stable region evolution. Together, these mechanisms provide the mathematical basis for the numerical stability observed in the segmentation results.

The r-adic discretization of the scale variable plays a complementary role by enforcing a logarithmic separation between successive resolution levels. This choice reduces redundancy in the operator signatures while preserving sensitivity to relevant structural features, in accordance with classical multiresolution analysis and pseudo-wavelet representations. Combined with Gaussian regularization, the r-adic sampling enables the algorithm to balance local responsiveness and global coherence,

yielding numerically stable segmentations under noise and acquisition blur without reliance on sharp intensity discontinuities.

From a conceptual standpoint, it is important to distinguish between the pointwise multiscale texture operator introduced in this work and the specific segmentation strategy used to exploit it. The region-growing algorithm constitutes one possible operational realization, chosen for its interpretability and computational simplicity, but the operator itself remains independent of any particular segmentation paradigm. Its mathematical grounding in scale–space diffusion theory ensures stability and interpretability, making it suitable for integration into alternative segmentation frameworks, including variational formulations, PDE-based models, texture classification, anomaly detection, registration, or hybrid analytical–learning pipelines.

To contextualize the contribution within the current segmentation landscape, it is useful to relate the proposed framework to recent developments in learning-based segmentation. While recent advances have been largely driven by deep learning architectures requiring large annotated datasets and extensive training procedures, the operator proposed here follows a complementary philosophy: It provides an analytically defined descriptor grounded in scale–space theory, particularly suited to scenarios where annotated data are limited or interpretability is required. Rather than replacing learning-based approaches, the proposed framework may serve either as a mathematically interpretable alternative or as a complementary structural prior in hybrid analytical–learning pipelines.

The reconstructed pulp chamber volumes exhibit consistent variations with respect to age and sex. From a methodological perspective, these trends indicate that the segmentation maintains structurally meaningful information consistent with anatomical expectations rather than collapsing anatomical structures through excessive smoothing. Although the present study is limited to descriptive analysis, the observed trends suggest that the underlying operator and region-growing mechanism are sensitive to genuine structural differences in the data, which is essential for subsequent quantitative analyses based on the segmented geometry.

From this perspective, the observed stability of the segmentations highlights the interpretability of the proposed framework, whose behavior can be directly traced to its mathematical formulation. The absence of training phases and the limited dependence on manually tuned parameters reduce the risk of overfitting and enhance reproducibility, while preserving direct control over the segmentation process. Rather than positioning itself as a benchmark-driven alternative to deep learning methods, the present framework provides a complementary paradigm in which segmentation behavior follows directly from scale–space diffusion principles and multiscale texture coherence.

Beyond the specific application considered here, the proposed framework opens several natural directions for further development. First, although the present work focuses on the extraction of a single region of interest, the texture coherence criterion could be extended to a multiclass setting by considering multiple reference signatures and allowing the simultaneous growth of competing regions with distinct multiscale behaviours. This would enable the segmentation of complex anatomical scenes involving several tissue types within a unified mathematical formulation.

Second, the dependence of the algorithm on a small number of interpretable parameters suggests the possibility of automatic or adaptive parameter selection strategies. For instance, the scale interval could be estimated directly from the image by analyzing the stability of the texture operator across scales, while seed initialization could be guided by extrema or coherence maxima of the multiscale texture field, reducing the need for expert intervention.

These methodological extensions should be interpreted within the constraints of real clinical imaging, particularly in CBCT data, where acquisition physics naturally produce blurred interfaces, noise, and heterogeneous textures. Given the methodological focus of the present work, extensive quantitative benchmarking against supervised learning segmentation frameworks was not considered within the scope of this study. Such comparisons, together with the integration of the operator into hybrid analytical–learning pipelines, constitute a natural continuation of the present framework and will be addressed in future studies using publicly available CBCT datasets and standard segmentation metrics.

Beyond the lack of benchmark comparisons, several operational limitations inherent to the proposed framework must be acknowledged. As a region-growing method driven by multiscale texture coherence, the algorithm may encounter difficulties in structures exhibiting extremely low contrast combined with weak textural differentiation. Similarly, highly fragmented anatomical regions or tissues with overlapping multiscale signatures may challenge the coherence criterion, which may occur in advanced mineralization stages or severe imaging artifacts.

Although the diffusion-based construction attenuates moderate noise, severe acquisition artifacts or strongly non-Gaussian noise perturbations may degrade the stability of the texture operator. Moreover, as with most seed-based methods, inappropriate initialization may bias the growth process, although the multiscale coherence constraint mitigates uncontrolled leakage. These aspects delineate the operational envelope of the method and motivate future work on adaptive seeding and hybrid segmentation strategies.

From a mathematical standpoint, several open questions remain. The current analysis does not include convergence proofs for the discrete region-growing process under general conditions, nor a formal sensitivity analysis with respect to parameter choices such as the scale interval or the coherence threshold. Addressing these aspects constitutes a natural direction for future work and would further strengthen the theoretical foundation of the proposed framework.

5. Conclusions

The primary contribution of this work is the introduction of a pointwise multiscale texture operator that provides an interpretable characterization of local contextual evolution across scales, grounded in scale–space diffusion theory and distinct from fixed-resolution statistical texture descriptors.

Its application to CBCT data demonstrates that the coherence-driven region-growing framework produces spatially coherent segmentations and connected 3D reconstructions of the dental pulp chamber. The resulting volumetric estimates exhibit variations consistent with known anatomical aging processes, suggesting that the reconstructed geometries preserve structurally meaningful information suitable for quantitative morphometric analysis.

From a methodological standpoint, the framework offers a transparent segmentation strategy whose behavior follows directly from its mathematical formulation, enabling explicit interpretation of the roles of scale, diffusion regularization, and texture coherence.

Future work will address formal convergence analysis of the discrete growth process, sensitivity studies with respect to scale sampling and coherence thresholds, and the extension of the framework to additional anatomical structures and imaging modalities.

Beyond the specific application considered, the proposed operator provides a principled basis

for segmentation methodologies designed to remain stable and explainable in imaging environments dominated by blur, noise, and structural heterogeneity.

Within this perspective, the proposed framework contributes a mathematically grounded approach to multiscale image analysis, highlighting the potential of analytically defined operators for robust and structurally consistent segmentation under challenging imaging conditions.

Author contributions

Data curation: C.C.G., C.F.G., Z.F.M.; Formal analysis: C.F.G., Z.F.M.; Investigation: C.C.G., C.F.G., J.A.V., Z.F.M.; Methodology: C.C.G., C.F.G., J.A.V., Z.F.M.; Resources: C.C.G., I.S., J.A.V.; Software: C.F.G.; Conceptualization, Supervision, Validation: C.C.G., C.F.G., I.S., J.A.V., Z.F.M.; Writing-original draft: C.C.G., C.F.G., J.A.V., Z.F.M.; Writing-review & editing: C.C.G., C.F.G., I.S., J.A.V., Z.F.M. All authors have read and approved the final version of the manuscript for publication.

Use of Generative-AI tools declaration

The authors declare they have no used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of interest

Authors declare no conflicts of interest.

References

1. R. C. Gonzalez, R. E. Woods, *Digital image processing*, New York: Pearson, 2018.
2. P. Soille, *Morphological image analysis: principles and applications*, Berlin: Springer, 1999.
3. D. Mumford, J. Shah, Optimal approximations by piecewise smooth functions and associated variational problems, *Commun. Pure Appl. Math.*, **42** (1989), 577–685. https://doi.org/10.1002/cpa.3160420503
4. A. Chambolle, Image segmentation by variational methods: Mumford and Shah functional and the discrete approximations, *SIAM J. Appl. Math.*, **55** (1995), 827–863. https://doi.org/10.1137/S0036139993257132
5. L. Alvarez, P. L. Lions, J. M Morel, Image selective smoothing and edge detection by nonlinear diffusion (II), *SIAM J. Numer. Anal.*, **29** (1992), 845–866. https://doi.org/10.1137/0729052
6. S. Mallat, *A wavelet tour of signal processing*, San Diego: Academic Press, 1999.
7. G. Aubert, P. Kornprobst, *Mathematical problems in image processing: partial differential equations and the calculus of variations*, Applied Mathematical Sciences, Vol. 147, New York: Springer, 2006. https://doi.org/10.1007/978-0-387-44588-5
8. Y. Yu, C. Wang, Q. Fu, R. Kou, F. Huang, B. Yang, et al., Techniques and challenges of image segmentation: a review, *Electronics*, **12** (2023), 1199. https://doi.org/10.3390/electronics12051199

9. M. L. Slim, R. Jacobs, R. M. de Souza Leal, R. C. Fontenele, AI-driven segmentation of the pulp cavity system in mandibular molars on CBCT images using convolutional neural networks, *Clin. Oral Invest.*, **28** (2024), 650. https://doi.org/10.1007/s00784-024-06009-2
10. L. Hernández-Alvarez, I. Vila-García, Z. Fernández-Muñiz, A. Cernea, L. C. Hernández-González, T. Cobo, et al., Measuring dental chamber volume with DICOM images from cone-beam computed tomography can be improved with a simple algorithm, *Appl. Sci.*, **14** (2024), 5365. https://doi.org/10.3390/app14135365
11. F. Akbarizadeh, B. M. Tabrizi, K. Fereidouni, M. M. K. Mazidi, Evaluation of the accuracy of age estimation based on ratio of pulp chamber to crown volume of the anterior teeth: a CBCT study, *Forensic Sci. Int.*, **377** (2025), 112610. https://doi.org/10.1016/j.forsciint.2025.112610
12. R. M. Haralick, L. G. Shapiro, Image segmentation techniques, *Comput. Vision Graphics Image Process.*, **29** (1985), 100–132. https://doi.org/10.1016/S0734-189X(85)90153-7
13. N. Otsu, A threshold selection method from gray-level histograms, *IEEE Trans. Syst. Man Cybern.*, **9** (1979), 62–66. https://doi.org/10.1109/TSMC.1979.4310076
14. S. Li, L. Ye, Multi-level thresholding image segmentation for rubber tree secant using improved Otsu's method and snake optimizer, *Math. Biosci. Eng.*, **20** (2023), 9645–9669. https://doi.org/10.3934/mbe.2023423
15. O. Vite-Chávez, J. Flores-Troncoso, R. Olivera-Reyna, J. U. Muñoz-Minjares, Improvement procedure for image segmentation of fruits and vegetables based on the Otsu method, *Image Anal. Stereol.*, **42** (2023), 185–196. https://doi.org/10.5566/ias.2939
16. X. Liu, Z. Zhang, Y. Hao, H. Zhao, Y. Yang, Optimized Otsu segmentation algorithm-based temperature feature extraction method for infrared images of electrical equipment, *Sensors*, **24** (2024), 1126. https://doi.org/10.3390/s24041126
17. Y. Xu, R. Quan, W. Xu, Y. Huang, X. Chen, F. Liu, Advances in medical image segmentation: A comprehensive review of traditional, deep learning and hybrid approaches, *Bioengineering*, **11** (2024), 1034. https://doi.org/10.3390/bioengineering11101034
18. M. E. Rayed, S. M. S. Islam, S. I. Niha, J. R. Jim, M. M. Kabir, M. F. Mridha, Deep learning for medical image segmentation: state-of-the-art advancements and challenges, *Inf. Med. Unlocked*, **47** (2024), 101504. https://doi.org/10.1016/j.imu.2024.101504
19. Y. Gao, Y. Jiang, Y. Peng, F. Yuan, X. Zhang, J. Wan, Medical image segmentation: a comprehensive review of deep learning-based methods, *Diagnostics*, **11** (2025), 52. https://doi.org/10.3390/tomography11050052
20. R. M. Haralick, K. Shanmugam, I. Dinstein, Textural features for image classification, *IEEE Trans. Syst. Man Cybern.*, **SMC-3** (1973), 610–621. https://doi.org/10.1109/TSMC.1973.4309314
21. C. E. Shannon, A mathematical theory of communication, *Bell Syst. Tech. J.*, **27** (1948), 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x
22. H. Ibrar, J. Muhammad, Efficient convex region-based segmentation for noising and inhomogeneous patterns, *Inverse Probl. Imag.*, **17** (2023), 708–725. https://doi.org/10.3934/ipi.2022074

23. X. Li, H. Ding, H. Yuan, W. Zhang, J. Pang, G. Cheng, et al., Transformer-based visual segmentation: a survey, *arXiv*, 2023. https://doi.org/10.48550/arXiv.2304.09854

24. G. Csurka, R. Volpi, B. Chidlovskii, Semantic image segmentation: two decades of research, *arXiv*, 2023. https://doi.org/10.48550/arXiv.2302.06378

25. Y. Wang, U. Ahsan, H. Li, M. Hagen, A comprehensive review of modern object segmentation approaches, *arXiv*, 2023. https://doi.org/10.48550/arXiv.2301.07499

26. H. Gao, O. Chae, Individual tooth segmentation from CT images using level set methods, *Pattern Recognit.*, **43** (2010), 2406–2417. https://doi.org/10.1016/j.patcog.2010.01.010

27. W. Duan, Y. Chen, Q. Zhang, X. Lin, X. Yang, Refined tooth and pulp segmentation using U-Net in CBCT image, *Dentomaxillofacial Rad.*, **50** (2021), 20200251. https://doi.org/10.1259/dmfr.20200251

28. Z. Chen, Q. Liu, J. Wang, N. Ji, Y. Gong, B. Gao, Tooth image segmentation and root canal measurement based on deep learning, *Front. Bioeng. Biotechnol.*, **13** (2025), 1565403. https://doi.org/10.3389/fbioe.2025.1565403

29. D. Sun, J. Wang, Z. Zuo, Y. Jia, Y. Wang, STS-transunet: semi-supervised tooth segmentation transformer U-net for dental panoramic image, *Math. Biosci. Eng.*, **21** (2024), 2366–2384. https://doi.org/10.3934/mbe.2024104

30. A. O. Santos-Junior, R. C. Fontenele, F. S. Neves, S. Ali, R. Jacobs, M. Tanomaru-Filho, A unique AI-based tool for automated segmentation of pulp cavity structures in maxillary premolars on CBCT, *Sci. Rep.*, **15** (2025), 5509. https://doi.org/10.1038/s41598-025-86203-8

31. M. Gamal, M. Baraka, M. Torki, Automatic mandibular semantic segmentation of teeth pulp cavity and root canals, and inferior alveolar nerve on Pulpy3D dataset, In: M. G. Linguraru, Q. Dou, A. Feragen, S. Giannarou, B. Glocker, K. Lekadir, et al., *Medical image computing and computer assisted intervention – MICCAI 2024*, Springer, **15008** (2024), 14–23. https://doi.org/10.1007/978-3-031-72111-3_2

32. W. Cui, Y. Wang, Q. Zhang, H. Zhou, D. Song, X. Zuo, et al., CTooth: a fully annotated 3D dataset and benchmark for tooth volume segmentation on cone beam computed tomography images, *arXiv*, 2022. https://doi.org/10.48550/arXiv.2206.08778

33. O. Boudraa, Segmentation of 3D dental images using deep learning, *arXiv*, 2022. https://doi.org/10.48550/arXiv.2207.09582

34. Y. Liu, C. Wang, M. Lu, J. Yang, J. Gui, S. Zhang, From simple to complex scenes: learning robust feature representations for accurate human parsing, *IEEE Trans. Pattern Anal. Mach. Intell.*, **46** (2024), 5449–5462. https://doi.org/10.1109/TPAMI.2024.3366769

35. A. Depeursinge, Multiscale and multidirectional biomedical texture analysis: finding the needle in the haystack, *Biomed. Texture Anal.*, 2017, 1–36. https://doi.org/10.1016/B978-0-12-812133-7.00002-8

36. A. Fanizzi, T. M. A. Basile, L. Losurdo, R. Bellotti, U. Bottigli, R. Dentamaro, et al., A machine learning approach on multiscale texture analysis for breast microcalcification diagnosis, *BMC Bioinf.*, **21** (2020), 91. https://doi.org/10.1186/s12859-020-3358-4

37. S. T. M. Ataky, D. Saqui, J. de Matos, A. de Souza Britto Junior, A. Lameiras Koerich, Multiscale analysis for improving texture classification, *Appl. Sci.*, **13** (2023), 1291. https://doi.org/10.3390/app13031291

38. T. Lindeberg, Scale-space theory: a basic tool for analyzing structures at different scales, *J. Appl. Stat.*, **21** (1994), 225–270.

39. W. Liu, J. Hu, F. Lv, Z. Tang, A new method for long-term temperature compensation of structural health monitoring by ultrasonic guided wave, *Measurement*, **252** (2025), 117310. https://doi.org/10.1016/j.measurement.2025.117310

40. J. J. Koenderink, The structure of images, *Biol. Cybern.*, **50** (1984), 363–370. https://doi.org/10.1007/BF00336961

41. T. Lindeberg, *Scale-space theory in computer vision*, The Springer International Series in Engineering and Computer Science, Vol. 256, Springer, 1994. https://doi.org/10.1007/978-1-4757-6465-9

42. G. T. Fechner, *Elemente der psychophysik*, Breitkopf und Härtel, Leipzig, 1860.

Appendix

Proof of Theorem 2.2. **(1) Dilation.** The identity

$$\psi_s\left(\frac{x}{r}\right) = \psi_{s+1}(x)$$

follows directly from substituting $x/2$ into the definition of ψ_s and using the relation $\sigma_{s+1} = r\sigma_s$.

(2) Pseudo-wavelet property. If $f(x) = e^{-x^2/a^2}$, then

$$\widehat{f}(\omega) = a\sqrt{\pi}\, e^{-\pi^2\omega^2/a^2}.$$

Applying this to the two Gaussian terms in ψ_s yields

$$\widehat{\psi_s}(\omega) = \frac{1}{r^{s-1}\sigma_0} e^{-r^{2s+1}\pi^2\sigma_0^2\omega^2} - \frac{1}{r^{s-1}\sigma_0} e^{-r^{2s-1}\pi^2\sigma_0^2\omega^2}.$$

Since both terms decay super-exponentially,

$$\int_\varepsilon^\infty \frac{|\widehat{\psi_s}(\omega)|^2}{\omega}\, d\omega < \infty, \qquad \forall \varepsilon > 0,$$

showing that ψ_s satisfies the pseudo-wavelet admissibility condition.

(3) Band-pass behaviour. Direct computation gives

$$|\widehat{\psi_s}(\omega)|^2 = \frac{1}{r^{2s-2}\sigma_0^2}\left(e^{-r^{2s+2}\pi^2\sigma_0^2\omega^2} - 2e^{-r^{2s-1}\pi^2\sigma_0^2\omega^2} + e^{-r^{2s}\pi^2\sigma_0^2\omega^2}\right).$$

Using

$$\int_0^\infty e^{-a^2\omega^2}\,d\omega = \frac{\sqrt{\pi}}{2a}, \qquad \int_0^\infty \omega^2 e^{-a^2\omega^2}\,d\omega = \frac{\sqrt{\pi}}{4a^3},$$

one obtains

$$\Delta\omega_s = \frac{1}{\sqrt{5}\,\pi}\sqrt{\frac{225-32\sqrt{10}}{15-4\sqrt{10}}}\,\frac{1}{r^s\sigma_0} \approx 1.0331\,\frac{1}{r^s\sigma_0}.$$

(4) Logarithmic scale coverage. Since $\sigma_s = 2^s\sigma_0$, one has

$$\Delta\sigma_s = \sigma_s - \sigma_{s-1} = r^s\sigma_0 - r^{s-1}\sigma_0 = \frac{1}{r}\sigma_s,$$

so the relative resolution is constant:

$$\frac{\Delta\sigma_s}{\sigma_s} = \frac{1}{r}.$$

□

Proof of Proposition 2.1. The operator $\mathcal{T}$ is linear with respect to I. Differentiating the convolution under the integral sign gives

$$\mathcal{T}_I(p)(\sigma) = (I * K_\sigma)(p), \qquad \text{where} \qquad K_\sigma(\mathbf{x}) = \frac{\partial}{\partial\sigma}\gamma_\sigma(\mathbf{x}).$$

A direct computation yields

$$K_\sigma(\mathbf{x}) = \frac{\sigma}{(2\pi\sigma^2)^{n/2}}\left(\frac{\|\mathbf{x}\|^2}{\sigma^2} - n\right)\exp\left(-\frac{\|\mathbf{x}\|^2}{2\sigma^2}\right).$$

For each fixed σ, the kernel K_σ belongs to $L^2(\mathbb{R}^n)$. By translation invariance and the Cauchy–Schwarz inequality,

$$\begin{aligned}|\mathcal{T}_I(p)(\sigma) - \mathcal{T}_J(p)(\sigma)| &\le \int_{\mathbb{R}^n} |I(\mathbf{y}) - J(\mathbf{y})|\,|K_\sigma(p-\mathbf{y})|\,d\mathbf{y}\\ &\le \|I-J\|_{L^2(\mathbb{R}^n)}\,\|K_\sigma\|_{L^2(\mathbb{R}^n)}.\end{aligned}$$

Squaring the absolute value and integrating with respect to σ, we obtain

$$\begin{aligned}\|\mathcal{T}_I(p) - \mathcal{T}_J(p)\|^2_{L^2([\sigma_{\min},\sigma_{\max}])} &= \int_{\sigma_{\min}}^{\sigma_{\max}} |\mathcal{T}_I(p)(\sigma) - \mathcal{T}_J(p)(\sigma)|^2\,d\sigma\\ &\le \|I-J\|^2_{L^2(\mathbb{R}^n)}\int_{\sigma_{\min}}^{\sigma_{\max}} \|K_\sigma\|^2_{L^2(\mathbb{R}^n)}\,d\sigma.\end{aligned}$$

It remains to control the integral of $\|K_\sigma\|^2_{L^2(\mathbb{R}^n)}$. Since the map

$$\sigma \mapsto \|K_\sigma\|_{L^2(\mathbb{R}^n)}$$

is continuous on the compact interval $[\sigma_{\min},\sigma_{\max}]$, there exists a constant $M > 0$ such that

$$\|K_\sigma\|_{L^2(\mathbb{R}^n)} \le M \qquad \text{for all } \sigma \in [\sigma_{\min},\sigma_{\max}].$$

Therefore,

$$\int_{\sigma_{\min}}^{\sigma_{\max}} \|K_\sigma\|_{L^2(\mathbb{R}^n)}^2 \, d\sigma \leq M^2(\sigma_{\max} - \sigma_{\min}) < \infty.$$

Taking

$$C = M\sqrt{\sigma_{\max} - \sigma_{\min}}$$

yields

$$\|\mathcal{T}_I(p) - \mathcal{T}_J(p)\|_{L^2([\sigma_{\min},\sigma_{\max}])} \leq C\|I - J\|_{L^2(\mathbb{R}^n)},$$

which proves the desired estimate, uniformly in p. □