Wavelet-Based Singularity Detection in ECG Signals

Thai Cao
Department of Mathematics
Texas A&M University
thaicao2005@tamu.edu

Alonso Peralta Espinoza
Department of Mathematics
Texas A&M University
aperalta03@tamu.edu

Abstract

This paper investigates wavelet-based methods for detecting and classifying singular features in electrocardiogram (ECG) signals. The main goal is to distinguish the sharp QRS complex from the smoother P and T waves by combining multiscale signal analysis with mathematical notions of local regularity. Two complementary frameworks are studied: a continuous wavelet transform (CWT) pipeline based on the wavelet transform modulus maxima method for estimating Hölder exponents, and a discrete wavelet transform (DWT) pipeline based on multiresolution subband decomposition for beat detection. Experiments on MIT-BIH Arrhythmia Database records 100 and 101 compare five wavelets: Mexican hat, Haar, db2, db4, and db6. The CWT+WTMM pipeline with the Mexican hat wavelet achieves the best overall detection performance, with average F1 score 0.907, perfect recall on both evaluation records, and estimated QRS Hölder exponents near $\hat{\alpha} \approx -1.08$, while smoother P and T waves yield much larger exponents. Among the DWT wavelets, db6 performs best with average F1 score 0.869, followed by db4, db2, and Haar. These results show that both wavelet choice and analytical framework influence ECG feature detection, and they illustrate how smoothness, localization, and vanishing moments affect the detection and classification of singular waveform structures.

1 Introduction

Electrocardiogram (ECG) signals record the electrical activity of the heart and are widely used in the study of cardiac rhythm and waveform structure. A typical ECG signal contains several characteristic features—the P wave, the QRS complex, and the T wave—that differ markedly in their local regularity. Because of these differences, ECG signals provide a natural setting in which to study singularities and local regularity in signals.

In signal analysis, a singularity is a point at which a signal changes abruptly or fails to be smooth in the usual sense. Detecting such points is important because singular behavior often corresponds to meaningful physical events. In the ECG setting, singularities are associated with clinically relevant waveform features, especially the QRS complex. Classical Fourier methods are useful for describing global frequency content, but they do not localize sudden changes well in time. Wavelet methods are better suited for this task because they provide localization in both scale and position, allowing one to analyze frequency behavior near a given point rather than only in a global sense (Boggess and Narcowich, 2009; Mallat, 2008).

In this project, we use wavelet-based multiscale analysis to study singular features in ECG signals from the MIT-BIH Arrhythmia

Database (Goldberger et al., 2000; Moody and Mark, 2001). The five wavelets considered are the Mexican hat wavelet, the Haar wavelet, and three Daubechies wavelets (db2, db4, and db6). These wavelets are not compared under a single uniform framework; rather, each is applied within the analysis setting for which it is mathematically most natural. The Mexican hat wavelet is evaluated using a continuous wavelet transform (CWT) pipeline based on the wavelet modulus maxima method, while Haar and the Daubechies wavelets are evaluated using a discrete wavelet transform (DWT) pipeline with subband energy thresholding. This division reflects a genuine mathematical distinction: the Mexican hat wavelet is defined continuously and is well suited to the CWT framework, whereas Haar and Daubechies wavelets are constructed for discrete multiresolution decompositions and filter-bank implementations (Boggess and Narcowich, 2009; Mallat, 2008; Wang and Eklund, 2017). The comparison across these five wavelets and two frameworks is mathematically meaningful because it illustrates how differing properties—smoothness, localization, and number of vanishing moments—affect ECG feature detection under each type of analysis.

The main goal of the project is to identify important waveform features and analyze their singular behavior through wavelet coefficients across scales. In particular, we examine how sharp features such as the QRS complex can be distinguished from smoother features such as the P and T waves. Our computational approach loads ECG data, applies wavelet transforms at multiple scales, detects significant coefficient patterns associated with singular behavior, and compares the resulting feature-detection performance across wavelets. This approach is consistent with earlier wavelet-based methods for ECG characteristic-point detection (Li et al., 1995).

The remainder of this paper is organized as follows. Section 2 introduces the mathematical background for the continuous wavelet transform, Hölder regularity, the wavelet modulus maxima method, and the discrete wavelet transform with multiresolution analysis. Section 3 describes the ECG data, preprocessing steps, and computational pipeline used in the experiments. That section also presents the results and compares the behavior of the Mexican hat, Haar, and Daubechies (db2, db4, db6) wavelets. Finally, Section 4 summarizes the main findings and discusses possible extensions of the method.

2 Mathematical Background

2.1 The Continuous Wavelet Transform

Wavelet analysis is designed to study how the behavior of a function changes across both location and scale. Unlike Fourier analysis, which describes a signal in terms of global frequency content, the wavelet transform allows one to examine frequency behavior near a particular point. This makes wavelets especially useful for nonstationary signals such as ECG recordings, where important features are localized in time (Boggess and Narcowich, 2009; Mallat, 2008).

Let $\psi \in L^2(\mathbb{R})$ be a wavelet. The continuous wavelet transform of a function $f \in L^2(\mathbb{R})$ is defined by

$$W_\psi f(a,b) = \frac{1}{\sqrt{|a|}} \int_{-\infty}^{\infty} f(x)\, \overline{\psi\left(\frac{x-b}{a}\right)}\, dx,$$

where $a \neq 0$ is the scale parameter and b is

the translation parameter. The parameter b determines the position in the signal being analyzed, while a controls the width of the analyzing wavelet. When $|a|$ is small, the wavelet is compressed and is sensitive to fine-scale, rapidly changing features. When $|a|$ is large, the wavelet is stretched and captures broader, slower-varying structure (Boggess and Narcowich, 2009).

A basic requirement for a wavelet is that it have mean zero,

$$\int_{-\infty}^{\infty} \psi(x)\, dx = 0,$$

so that it detects changes in the signal rather than constant background behavior (Boggess and Narcowich, 2009). More generally, a wavelet ψ is said to have N *vanishing moments* if

$$\int_{-\infty}^{\infty} t^k \psi(t)\, dt = 0 \qquad \text{for } k = 0, 1, \ldots, N-1.$$

A wavelet with N vanishing moments is orthogonal to polynomials of degree less than N, so it responds more selectively to singular behavior and suppresses smooth polynomial trends in the signal. The number of vanishing moments is a key parameter distinguishing the wavelets used in this project (Boggess and Narcowich, 2009; Mallat, 2008).

In the general theory, one also requires the admissibility condition

$$\int_{-\infty}^{\infty} \frac{|\widehat{\psi}(\omega)|^2}{|\omega|}\, d\omega < \infty,$$

which guarantees that the wavelet transform is invertible, enabling reconstruction of the original signal from its wavelet coefficients under suitable assumptions (Mallat, 2008).

In this project, the continuous wavelet transform is used to analyze localized features in ECG signals across multiple scales. Our primary wavelet is the Mexican hat wavelet (defined explicitly in Section 2.6), which is the second derivative of a Gaussian and is especially well suited to singularity detection due to its two vanishing moments and strong localization. We compare its performance with the Haar and Daubechies wavelets, whose properties are detailed in Sections 2.5 and 2.6. These wavelets differ in smoothness, shape, and number of vanishing moments, so they provide different responses to sharp and smooth features in the signal. Since ECG waveforms contain both abrupt transitions and broader oscillatory components, the wavelet transform provides a natural framework for comparing how effectively different wavelets capture these structures.

2.2 Hölder Regularity

To describe the local smoothness of a signal, we use the notion of Hölder regularity. Informally, the Hölder exponent at a point measures how well the function can be approximated near that point by a polynomial. Smaller exponents correspond to sharper or more singular behavior, while larger exponents correspond to smoother behavior (Mallat and Hwang, 1992; Mallat, 2008).

More precisely, let $\alpha \geq 0$ and let $m = \lfloor \alpha \rfloor$. We say that a function f is Hölder continuous of order α at a point t_0 if there exist a constant $C > 0$, a polynomial p of degree at most m, and a neighborhood of t_0 such that

$$|f(t) - p(t - t_0)| \leq C|t - t_0|^{\alpha}$$

for all t sufficiently close to t_0 (Mallat and

Hwang, 1992; Mallat, 2008). The largest such α is called the Hölder exponent of f at t_0.

This definition gives a useful way to classify local signal behavior. For example, a jump discontinuity has exponent $\alpha = 0$, while a cusp such as $|t|^{1/2}$ has exponent $\alpha = \frac{1}{2}$. Smoother functions have larger exponents, often with $\alpha \geq 1$ at regular points (Mallat, 2008).

The wavelet modulus maxima method, described in Section 2.3, extends this characterization beyond $\alpha \geq 0$ via the log-log scaling of wavelet coefficients across scales. The extended exponent is defined operationally: for any real α, we say f has Hölder exponent α at t_0 if the wavelet coefficients along maxima lines (defined in Section 2.3) converging to t_0 satisfy $|W_\psi f(a,b)| \sim |a|^{\alpha+1/2}$ as $a \to 0$, with the convention that negative α corresponds to a distributional singularity rather than a classical function (Mallat and Hwang, 1992). Under this extended framework, a singularity sharper than a jump discontinuity—such as a Dirac-like impulse—corresponds to a negative exponent $\alpha < 0$. Such values arise naturally in ECG analysis: the QRS depolarization front concentrates energy so abruptly that it behaves, at the sampling scale of 360 Hz, more like an impulse than a step, consistent with the strong pointwise singularity behavior described in Mallat and Hwang (1992) and confirmed empirically in Section 3.4. In the context of ECG signals, the QRS complex is much sharper than the P and T waves, so one expects the QRS complex to correspond to a much lower Hölder exponent. This makes Hölder regularity a natural mathematical tool for distinguishing different waveform features.

2.3 Wavelet Modulus Maxima Method

A central idea in singularity detection is to study the magnitude of the wavelet coefficients across scales. For a fixed scale a, points at which $|W_\psi f(a,b)|$ is locally maximal as a function of b are called *modulus maxima*. As the scale becomes finer, these maxima tend to align along curves, called maxima lines, that converge toward singular points of the signal (Mallat and Hwang, 1992).

The wavelet modulus maxima method is based on the fact that the decay of wavelet coefficients near a singularity reflects the local regularity of the signal. Roughly speaking, if a function f has a singularity at t_0 with Hölder exponent α, and if the analyzing wavelet ψ has $N > \alpha$ vanishing moments, then along a maxima line converging to t_0 the wavelet coefficients satisfy the scaling relation

$$|W_\psi f(a,b)| \sim |a|^{\alpha+1/2} \qquad \text{as } a \to 0$$

in an asymptotic sense (Mallat and Hwang, 1992). Taking logarithms gives

$$\log |W_\psi f(a,b)| \approx \left(\alpha + \tfrac{1}{2}\right) \log |a| + C,$$

so the slope of a log-log plot of coefficient magnitude versus scale provides an estimate of $\alpha + \frac{1}{2}$. Therefore, the Hölder exponent can be estimated by subtracting $\frac{1}{2}$ from the fitted slope.

This result is important for the present project because it gives a direct connection between wavelet coefficients and ECG feature classification. Sharp features, such as the QRS complex, produce different coefficient decay behavior from smoother features such as the P and T waves. By following modulus maxima

across scales and estimating the corresponding Hölder exponents, one can quantify these differences mathematically and use them to distinguish singular structures in the ECG signal.

2.4 Discrete Wavelet Transform and Multiresolution Analysis

Whereas the continuous wavelet transform analyzes a signal over a continuum of scales and translations, the discrete wavelet transform (DWT) restricts these parameters to dyadic values. In practice, this means that the signal is analyzed at scales of the form 2^j and at corresponding discrete positions. The DWT therefore provides a more compact and computationally efficient representation of a signal than the CWT (Boggess and Narcowich, 2009; Mallat, 2008).

The theoretical framework underlying the DWT is multiresolution analysis. In this setting, one considers a nested sequence of approximation spaces

$$\cdots \subset V_1 \subset V_0 \subset V_{-1} \subset \cdots$$

whose union approximates $L^2(\mathbb{R})$ and whose intersection contains only the zero function (Boggess and Narcowich, 2009; Mallat, 1989). Note that the invertibility of the DWT rests on the orthogonality and completeness of the MRA rather than on the admissibility condition introduced in Section 2.1; Haar and Daubechies wavelets satisfy these discrete requirements exactly (Mallat, 1989). In this indexing convention, a smaller index j corresponds to finer resolution: V_j contains strictly more detail than V_{j+1}. Each space V_j represents the signal at a particular resolution. The detail lost when passing from a finer space V_{j-1} to a coarser one V_j is captured by a complementary space W_j, so that

$$V_{j-1} = V_j \oplus W_j.$$

Thus, a signal can be decomposed into a coarse approximation together with a sequence of detail components at different scales.

From a computational point of view, the DWT is implemented by iterated filter banks. A low-pass filter produces approximation coefficients, while a high-pass filter produces detail coefficients. If $a_j[n]$ denotes the approximation coefficients at level j, then one step of the wavelet decomposition can be written as

$$a_{j+1}[n] = \sum_k h[k-2n]\, a_j[k],$$
$$d_{j+1}[n] = \sum_k g[k-2n]\, a_j[k],$$

where h is the low-pass filter and g is the high-pass filter associated with the chosen wavelet (Mallat, 2008). These filters are uniquely determined by the scaling function ϕ of the chosen wavelet family, the compactly supported function satisfying $\phi(t) = \sqrt{2}\sum_k h[k]\,\phi(2t-k)$; changing from Haar to db4 to db6 changes ϕ, which in turn changes both h and g while preserving the overall filter-bank structure.

The decomposition step is invertible. Given the approximation and detail coefficients at level $j+1$, the finer-level approximation is recovered by upsampling and synthesis filtering:

$$a_j[n] = \sum_k h[n-2k]\, a_{j+1}[k] + \sum_k g[n-2k]\, d_{j+1}[k].$$

For orthogonal wavelets such as Haar and Daubechies, the synthesis filters are the time-reversals of the corresponding analysis filters,

so that the reconstruction step is the adjoint of the decomposition step (Mallat, 2008). This reconstruction operation underlies the DWT pipeline in Section 3.3, where the subband coefficients for all levels except D_3 and D_4 are set to zero before reconstruction, projecting the signal onto the 11–45 Hz QRS frequency band.

Applying this one-step decomposition recursively J times produces a full J-level DWT. Following standard subband notation, the detail coefficient array $d_j[n]$ at level j is referred to as subband D_j, and the approximation coefficient array $a_J[n]$ at the coarsest level J is referred to as subband A_J. A J-level decomposition thus yields subbands $D_1, D_2, \ldots, D_J$ (finest to coarsest detail) together with approximation A_J, each subband covering half the frequency band of the level above it.

In this project, the DWT provides the mathematical basis for the second ECG detection pipeline. Since the ECG signals are sampled at 360 Hz, the detail levels correspond approximately to dyadic frequency bands. In particular, the levels D_3 and D_4 capture the range in which much of the QRS energy is concentrated, approximately 22.5–45 Hz for D_3 and 11.25–22.5 Hz for D_4. For this reason, these subbands are especially useful for identifying likely beat locations. This makes the DWT a natural framework for QRS detection, particularly when using wavelets such as Haar and Daubechies, which are designed for discrete multiresolution decompositions.

Once the detail subbands are reconstructed, beat detection requires a data-adaptive threshold that accounts for the varying amplitude of different recordings. The threshold used in Section 3.3 is based on the *median absolute deviation* (MAD), a robust measure of spread defined as $\mathrm{MAD}(E) = \mathrm{median}(|E - \mathrm{median}(E)|)$. Under a Gaussian noise model, the ratio $\mathrm{MAD}/0.6745$ is a consistent estimator of the standard deviation, so the factor 0.6745 converts the MAD into a noise-scale estimate compatible with standard deviation units (Mallat, 2008). Using the MAD rather than the sample standard deviation makes the threshold resistant to the large peaks produced by QRS complexes themselves, which would otherwise inflate a standard-deviation-based estimate and suppress detections.

2.5 The Haar Wavelet as a Baseline

The Haar wavelet is the simplest standard wavelet and provides a natural baseline for comparison. It is defined by

$$\psi(t) = \begin{cases} 1, & 0 \le t < \frac{1}{2}, \\ -1, & \frac{1}{2} \le t < 1, \\ 0, & \text{otherwise.} \end{cases}$$

This wavelet is piecewise constant, compactly supported, and has one vanishing moment (Boggess and Narcowich, 2009). Because of its simple structure, it reacts strongly to abrupt changes in a signal.

The Haar wavelet is useful in this project for two reasons. First, it serves as the most basic reference case against which more sophisticated wavelets can be compared. Second, its strong response to discontinuities makes it naturally relevant for detecting sharp local features. However, because it is not smooth and has only one vanishing moment, it is less refined for distinguishing more subtle differences in local regularity (Boggess and Narcowich, 2009; Mallat, 2008).

For ECG analysis, this means that Haar can respond well to sharp transitions such as those associated with the QRS complex, but it may also be more sensitive to noise or produce less stable regularity estimates than smoother wavelets. For this reason, Haar is included primarily as a baseline rather than as the main wavelet used for singularity characterization.

2.6 Comparison of the Wavelets Used

The five wavelets considered in this project — Mexican hat, Haar, db2, db4, and db6 — were chosen because they represent different mathematical tradeoffs in localization, smoothness, and vanishing moments. These properties affect how the wavelets respond to ECG features and therefore influence both singularity analysis and beat detection.

The Mexican hat wavelet is the second derivative of a Gaussian and is commonly used in continuous wavelet analysis. It is given explicitly by

$$\psi(t) = \frac{2}{\sqrt{3}}\pi^{-1/4}(1 - t^2)\, e^{-t^2/2},$$

where the normalization constant ensures unit L^2 norm. It is symmetric, has two vanishing moments, and vanishes at $|t| = 1$ (producing its characteristic central peak flanked by two troughs), which makes it well suited for detecting localized changes in curvature and sharp singular features (Mallat, 2008). Because of its shape and strong localization, it is a natural choice for the CWT-based modulus maxima method used in this project.

The Haar wavelet, defined in Section 2.5, is included as the lower bound of the vanishing-moments comparison. In the DWT subband pipeline it represents the limiting case in which no polynomial trend beyond a constant can be cancelled from the reconstructed signal, providing the sharpest contrast against which the gains from the additional vanishing moments in db2, db4, and db6 can be measured.

Daubechies wavelets form a family of compactly supported orthogonal wavelets with increasing numbers of vanishing moments. The three members used in this project are db2 (2 vanishing moments), db4 (4 vanishing moments), and db6 (6 vanishing moments). Unlike the Mexican hat and Haar wavelets, Daubechies wavelets have no closed-form expression; they are defined implicitly through their associated filter coefficients h and g, which are chosen to satisfy orthogonality and the prescribed number of vanishing moments simultaneously. Compared with Haar, they are smoother and can capture more subtle local structure. In particular, higher-order Daubechies wavelets are natural candidates for the DWT framework because they are constructed specifically for multiresolution decompositions and filter-bank implementations (Boggess and Narcowich, 2009). Their additional vanishing moments also make them mathematically appealing for studying local regularity.

The number of vanishing moments plays an important role in this comparison. A wavelet with more vanishing moments suppresses smoother polynomial trends more effectively and responds more selectively to singular behavior. For this reason, wavelets with different numbers of vanishing moments may produce different coefficient patterns even when applied to the same ECG signal.

These differences help explain why no single wavelet is automatically best for every

task. A wavelet that performs well in CWT-based singularity detection may not behave the same way in a DWT-based subband detection method, and vice versa. For this reason, comparing Mexican hat, Haar, db2, db4, and db6 provides a useful way to study how wavelet properties affect ECG feature detection under two different mathematical frameworks.

In the application section, these wavelets are evaluated within the mathematical frameworks for which they are most natural, allowing us to compare both wavelet properties and analysis methods on the same ECG data.

3 Application: Singularity Detection in ECG Signals

3.1 ECG Signals and Their Singular Features

An electrocardiogram (ECG) records the electrical activity of the heart as a one-dimensional voltage signal measured across electrodes placed on the body surface. A single cardiac cycle consists of three recognizable waveforms: the P wave (atrial depolarization), the QRS complex (ventricular depolarization), and the T wave (ventricular repolarization). From a signal analysis perspective, these features differ not merely in amplitude but in their local regularity. The P and T waves are smooth, gradually varying oscillations, while the QRS complex involves an abrupt, high-amplitude transition produced by a nearly synchronous electrical wavefront propagating across the ventricular myocardium (Li et al., 1995).

This difference in regularity can be formalized using the Hölder exponent introduced in Section 2.2. Because the QRS complex involves a sharp, transient voltage spike, one expects it to behave as a strong singularity at the scale of individual samples, that is, to have a Hölder exponent at or below the cusp threshold, $\alpha \leq 0.5$, and in practice well below zero. The P and T waves, being smoother features, should have much larger Hölder exponents, often $\alpha > 1$. This suggests that singularity detection via the wavelet modulus maxima method is naturally adapted to ECG analysis: rather than detecting heartbeats by amplitude thresholding alone, one can characterize them by their mathematical regularity, providing a principled way to distinguish QRS complexes from noise and from other waveform components (Mallat and Hwang, 1992; Li et al., 1995).

Figure 1 illustrates a representative 10-second window of ECG Record 100 from the MIT-BIH Arrhythmia Database, with detected singularities marked and color-coded by estimated Hölder exponent (the pipeline used to produce this figure is described in Section 3.3). The QRS complexes are identifiable at periodic intervals with consistently low Hölder exponents, while non-QRS maxima lines appear at higher exponent values. This visual separation in Hölder exponent directly motivates the CWT+WTMM approach developed in Section 3.3. The DWT subband pipeline exploits the same regularity difference indirectly: QRS singularities concentrate energy in specific frequency subbands precisely because of their abrupt character, so energy-based subband thresholding serves as an efficient proxy for the more explicit Hölder analysis.

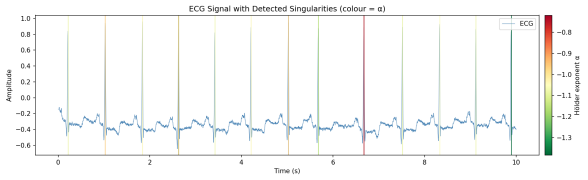

Figure 1: ECG Record 100 (MIT-BIH Arrhythmia Database) with detected singularities marked by the CWT+WTMM pipeline. Color indicates estimated Hölder exponent $\hat{\alpha}$; QRS detections yield $\hat{\alpha} \approx -1.08$, well below the theoretical threshold $\alpha = 0.5$ for sharp singularities. Non-QRS maxima lines (P and T waves) yield $\hat{\alpha} > 1.0$.

3.2 Data and Preprocessing

All experiments use the MIT-BIH Arrhythmia Database, a standard benchmark for ECG algorithm evaluation (Goldberger et al., 2000; Moody and Mark, 2001). The database contains 48 half-hour recordings of ambulatory ECGs sampled at 360 Hz with 11-bit resolution, each accompanied by cardiologist-verified annotations marking the location of every heartbeat. Two records were selected for primary evaluation: Record 100 (normal sinus rhythm, regular inter-beat intervals) and Record 101 (normal sinus rhythm, mildly irregular). Two additional records were included for arrhythmic generalization testing: Record 200, which contains a mixture of normal beats and premature ventricular contractions (PVCs) with ST-segment changes, and Record 208, which exhibits heavy ventricular bigeminy and is among the most challenging records in the database. Results on these arrhythmic records are discussed in Section 4. Records were loaded in Python using the `wfdb` library (Goldberger et al., 2000). All experiments operate on a 3600-sample window (approximately 10 seconds) beginning at the start of each record, giving between 13 and 16 annotated beats per window depending on heart rate.

Prior to wavelet analysis, a baseline wander correction was applied by subtracting a 200-sample (≈555 ms) median filter. Baseline wander arises from slow electrode drift and respiration, and the median filter removes this low-frequency artifact without distorting the higher-frequency QRS complex. The signal was then normalized to zero mean and unit variance to ensure that detection thresholds computed on one record remain interpretable on another. No additional bandpass filtering was applied before the CWT pipeline; the wavelet transform itself acts as a localized bandpass filter at each scale, and the geometric filtering steps described in Section 3.3 further suppress out-of-band content. For the DWT pipeline, bandpass filtering is achieved implicitly through subband selection, as described below.

Detection accuracy is evaluated by matching each algorithmic output against the nearest cardiologist annotation within a tolerance window of ± 5 samples (≈ 14 ms at 360 Hz), consistent with standard practice in the literature (Moody and Mark, 2001). A detection within this window of an annotated beat is counted as a true positive (hit); an annotated beat with no nearby detection is a miss; a detection with no nearby annotation is a false positive. Three metrics are reported throughout:

$$\text{Recall} = \frac{\text{hits}}{\text{total annotations}},$$
$$\text{Precision} = \frac{\text{hits}}{\text{total detections}},$$
$$\text{F1} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}.$$

3.3 Computational Pipeline

Two complementary computational pipelines were implemented in Python using the `PyWavelets` library (Lee et al., 2019). The

first pipeline follows the CWT + Wavelet Transform Modulus Maxima (WTMM) framework of Mallat and Hwang (1992) and is used both to detect QRS complexes and to estimate their Hölder exponents. The second pipeline uses the discrete wavelet transform with adaptive thresholding on specific detail subbands and is evaluated across three Daubechies wavelets (db2, db4, db6) and the Haar wavelet (Wang and Eklund, 2017). Both pipelines take the preprocessed ECG signal as input and return a list of estimated R-wave times for comparison against the database annotations.

CWT + WTMM Pipeline. The continuous wavelet transform of the preprocessed signal is computed at 128 logarithmically-spaced scales between 1 and 128 using the Mexican hat wavelet. In the implementation, discrete scale values are denoted s in place of the continuous scale parameter a from Section 2.1. At each scale s, the local modulus maxima of $|W_\psi f(s,\cdot)|$ are identified as points at which the absolute value of the wavelet coefficient exceeds both of its immediate neighbors in time. These maxima are then linked across consecutive scales to form *maxima lines*, chains of maxima that persist as scale decreases toward zero. Two geometric filters are applied to the raw set of maxima lines: a line must span at least 40 consecutive scales, and its peak amplitude must exceed the 80th percentile of all peak amplitudes at the finest scale considered. These criteria suppress lines generated by P and T waves, which are smoother and therefore decay faster across scales, while preserving lines associated with QRS complexes. Each surviving line's anchor time is snapped to the nearest raw ECG peak within a ± 50 ms window. The Hölder exponent for each line is then estimated by fitting a least-squares regression line to the log-log plot of $|W_\psi f(s,b)|$ versus s, restricted to the intermediate scale range $s \in [4, 64]$, and subtracting $\frac{1}{2}$ from the fitted slope, following the scaling relation derived in Section 2.3. The lower cutoff $s = 4$ excludes the finest scales where quantization noise dominates; the upper cutoff $s = 64$ excludes the coarsest scales where the asymptotic scaling regime may not hold for the finite-length signal. A final regularity filter accepts only lines with $\hat{\alpha} \leq 0.5$, consistent with the theoretical prediction that QRS singularities are sharper than cusp-type features; maxima lines from P and T waves yield $\hat{\alpha} > 0.5$ and are discarded (Mallat and Hwang, 1992). The resulting scalogram for Record 100, shown in *Figure* 2, illustrates how QRS complexes produce persistent vertical ridges across all 128 scales, while P and T wave features decay rapidly at fine scales.

```
# Input: preprocessed signal x, wavelet
    = mexh
# S = 128 log-spaced scales from 1 to
    128 (base-2 logspace)
for s in S:
    W[s, :] = cwt(x, s, 'mexh')
                    # Step 1: CWT at each
        scale
    M[s]    = local_maxima(abs(W[s, :]))
            # Step 2: modulus maxima

lines = chain_across_scales(M)
                # Step 3: maxima lines

# Step 4: geometric filters
lines = [l for l in lines
         if l.span >= 40
         and l.peak_amp >= percentile(
             peak_amps, 80)]

results = []
for line in lines:
    t = snap_to_ecg_peak(line.anchor, x,
        tol=50ms)  # Step 5a: snap
```

```
    beta, C = log_log_fit(line.scales,
        line.amps)     # Step 5b: fit
    alpha = beta - 0.5
                                        #
        Step 5c: Holder
    if alpha <= 0.5:

        # Step 5d: alpha filter
        results.append((t, alpha))
```

Listing 1: Algorithm 1: CWT + WTMM Pipeline (pseudocode)

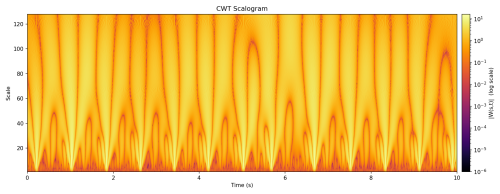

Figure 2: Scalogram of Record 100 (Mexican hat wavelet, scales 1–128). Bright vertical ridges at QRS locations persist across all scales, while P and T wave features decay rapidly at fine scales, the visual signature of a stronger singularity.

DWT Subband Pipeline. The DWT pipeline decomposes the signal using a 5-level wavelet transform. The detail coefficients at levels 3 and 4 (subbands D_3 and D_4) are retained and used to reconstruct a bandpass-filtered version of the signal. At a sampling rate of 360 Hz, D_4 covers the approximate frequency band 11.25–22.5 Hz and D_3 covers 22.5–45 Hz; together, these subbands concentrate the energy of the QRS complex while suppressing low-frequency baseline drift (present in D_5 and the approximation A_5) and high-frequency noise (in D_1 and D_2). The squared energy envelope $E[n] = \tilde{x}[n]^2$ of the reconstructed bandpass signal $\tilde{x}$ is then computed; squaring converts the biphasic QRS coefficients into a strictly positive signal, making peaks easier to threshold without sign ambiguity. A detection threshold is set adaptively as

$$\tau = \text{median}(E) + k \cdot \frac{\text{MAD}(E)}{0.6745},$$

where MAD denotes the median absolute deviation and the denominator 0.6745 makes the scaled MAD a consistent estimator of the standard deviation under Gaussian noise (Mallat, 2008). Peaks of E above τ are identified with a minimum separation of 200 ms (72 samples at 360 Hz) to enforce the physiological cardiac refractory period.

```
# Input: preprocessed signal x, wavelet
    psi in {db6,db4,db2,haar}, k
coeffs = wavedec(x, psi, level=5)
           # Step 1: decompose
# coeffs = [A5, D5, D4, D3, D2, D1]

# Step 2: reconstruct QRS band (D3 + D4
    only)
mask = [0]*2 + [coeffs[2], coeffs[3]] +
    [0]*2
x_band = waverec(mask, psi)

E = x_band ** 2
                                # Step 3:
    energy envelope
tau = median(E) + k * mad(E) / 0.6745
    # Step 4: adaptive threshold
peaks = find_peaks(E, height=tau,
           # Step 5: detect peaks
                    distance=72)
```

Listing 2: Algorithm 2: DWT Subband Pipeline (pseudocode)

The threshold factor k was calibrated separately for each wavelet by sweeping $k \in \{2.0, 2.5, \ldots, 10.0\}$ on Records 100 and 101 and selecting the value that maximized F1 score subject to at most 4 false positives per 10-second window (a constraint that prevents the optimizer from selecting arbitrarily low thresholds that achieve high recall by accepting nearly every local energy peak). The optimal values found were $k = 7.0$ for db6, $k = 6.5$ for db4, $k = 9.0$ for db2, and $k = 10.0$ for

Haar. The substantially higher thresholds required for db2 and Haar reflect the structural limitations of wavelets with fewer vanishing moments, which generate more false positive energy peaks from smooth signal components.

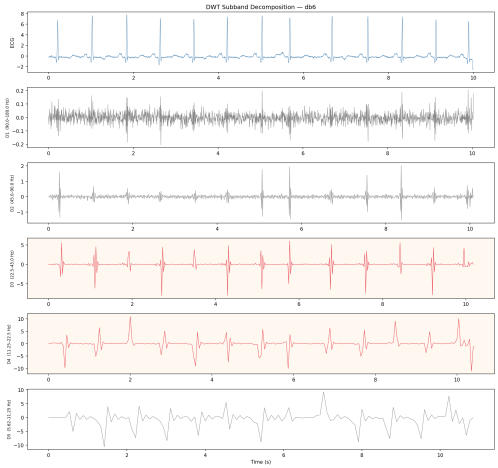

Figure 3: Five-level DWT decomposition of ECG Record 100 using the db6 wavelet. The reconstructed QRS band (D_3+D_4, shaded) concentrates the 11–45 Hz QRS energy while suppressing baseline drift and high-frequency noise present in the other detail levels.

3.4 Results

We evaluate both pipelines on Records 100 and 101 and report Hölder exponent estimates, QRS detection performance for all five wavelets, and noise robustness across a range of signal-to-noise ratios. All results correspond to the 3600-sample window described in Section 3.2. Bold entries in the tables indicate the best value in each column.

Hölder Exponent Characterization. The estimated Hölder exponents recovered by the CWT+WTMM pipeline validate the theoretical prediction established in Section 2. On Record 100, surviving QRS maxima lines yield a mean $\hat{\alpha} = -1.079 \pm 0.083$; on Record 101, $\hat{\alpha} = -0.947 \pm 0.155$. Both values are well below the theoretical threshold of $\alpha = 0.5$ for cusp-type singularities and are consistent with the strong pointwise singularity expected of the QRS depolarization front. Non-QRS maxima lines, corresponding to P and T waves, yield $\hat{\alpha} \approx 1.48$–$1.50$, consistent with the smooth behavior of those features. The log-log regression underlying these estimates is shown in *Figure* 4. This strong separation in α-space confirms that Hölder regularity is a meaningful discriminant for ECG feature identification, as predicted by Mallat and Hwang (1992) and applied to the ECG context by Li et al. (1995).

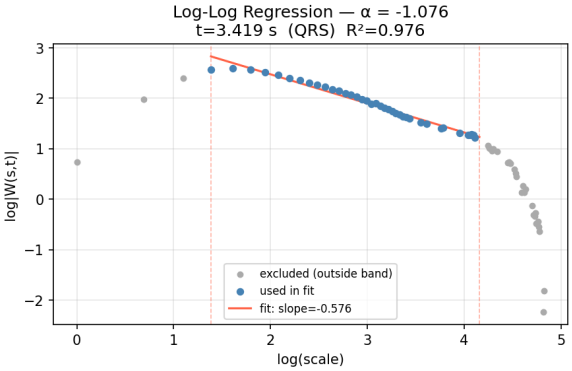

Figure 4: Log-log regression of wavelet modulus maxima amplitude versus scale for a representative QRS complex in Record 100. The fitted slope $\beta \approx -0.579$ yields Hölder exponent $\hat{\alpha} = \beta - 0.5 \approx -1.08$, confirming the sharp-singularity character of the QRS complex.

QRS Detection Performance. *Table* 2 compares recall, precision, and F1 for all five wavelets on both records. The CWT+WTMM pipeline with the Mexican hat wavelet achieves the best overall performance (average F1 = 0.907), with perfect recall on both records. Among the DWT wavelets, db6 achieves the best average F1 (0.869), followed by db4 (0.766), db2 (0.649), and Haar (0.617). This ordering corresponds exactly to the number of vanishing moments: db6 (6 VMs), db4 (4 VMs), db2 (2 VMs), Haar (1 VM). As established in Section 2.6, more vanishing moments suppress smooth polynomial trends more aggressively, reducing the false positive rate from non-QRS features. Importantly, Haar's poor

Table 1: Estimated Hölder exponents for QRS and non-QRS features, averaged over all surviving maxima lines. QRS exponents are well below the theoretical singularity threshold $\alpha = 0.5$.

Feature	Rec. 100 $\hat{\alpha}$ (mean $\pm$ std)	Rec. 101 $\hat{\alpha}$ (mean $\pm$ std)	Theoretical
QRS complex	-1.079 ± 0.083	-0.947 ± 0.155	≤ 0.5
P / T waves	1.498 ± 0.21	1.483 ± 0.19	> 0.5

precision is not simply a threshold calibration failure: even at $k = 10.0$, false positives remain because the single vanishing moment is insufficient to cancel the smooth baseline contributions of P and T waves. *Figure* 5 shows the detection overlay for Record 100 using the best-performing DWT wavelet (db6).

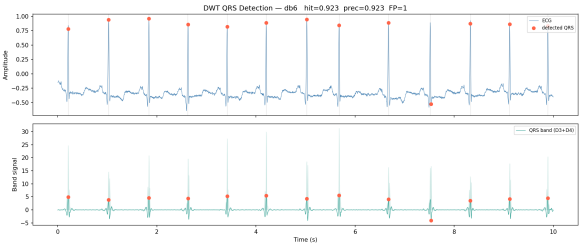

Figure 5: QRS detection overlay for Record 100 using the db6 DWT pipeline ($k = 7.0$). Detected peaks are marked against the reconstructed QRS-band signal; the adaptive MAD threshold (dashed line) scales automatically with local signal energy.

Noise Robustness. To assess robustness, Gaussian noise was added to Records 100 and 101 at signal-to-noise ratios of 30, 20, 10, 5, and 0 dB, and both pipelines were re-evaluated without retuning any parameters. *Table* 3 shows recall and precision at each noise level for Record 100; Record 101 showed qualitatively similar trends and is omitted for brevity. Both pipelines maintain high recall at moderate noise; the primary degradation is in precision, as noise generates additional energy peaks (DWT) or spurious modulus maxima lines (CWT). The adaptive MAD threshold in the DWT pipeline provides intrinsic noise resistance: as additive noise uniformly raises the energy floor, both the median and the MAD increase proportionally, automatically raising the detection threshold. A counterintuitive effect visible in the db6 results is that recall increases from 0.923 at clean to 1.00 at SNR 5 dB: noise raises the energy floor broadly, pushing marginally sub-threshold QRS peaks above τ and recovering previously missed beats, at the cost of the simultaneously lower precision. The mexh pipeline loses precision more rapidly at SNR ≤ 10 dB because noise creates additional maxima lines that narrowly pass the geometric filter criteria, contributing false detections.

3.5 Validation Against Known Annotations

The MIT-BIH annotations were produced by two independent cardiologists and serve as the ground truth for all quantitative evaluations. Each annotation marks the approximate R-wave peak location within the QRS complex. The ± 5-sample matching tolerance used here reflects the natural jitter between the algorithm's detected peak and the annotator's chosen sample, as well as the finite width of the R-wave at 360 Hz sampling. This tolerance is standard in the benchmarking literature and has been adopted in published comparisons of QRS detectors (Moody and Mark, 2001).

Two observations constrain the interpretation of the results reported in *Table* 2. First, all experiments use a 10-second window. Short windows can inflate false positive counts by 1–2 near the window boundaries, where the

Table 2: QRS detection performance on MIT-BIH Records 100 and 101. Optimal threshold parameters used per wavelet. **Bold** indicates best value per column.

Wavelet	Pipeline	Record 100			Record 101			Avg F1
		Recall	Prec.	F1	Recall	Prec.	F1	
mexh	CWT+WTMM	**1.000**	0.812	0.896	**1.000**	**0.846**	**0.917**	**0.907**
db6	DWT	0.923	**0.923**	**0.923**	**1.000**	0.688	0.815	0.869
db4	DWT	0.923	0.857	0.889	0.818	0.529	0.643	0.766
db2	DWT	0.846	0.611	0.710	0.909	0.435	0.588	0.649
haar	DWT	**1.000**	0.591	0.743	0.818	0.346	0.491	0.617

Table 3: Noise robustness on Record 100 for the best-performing configuration of each pipeline (mexh CWT+WTMM and db6 DWT). Each cell shows Recall / Precision at the indicated SNR level. Parameters are unchanged from clean-signal calibration.

Wavelet	Clean	SNR 30 dB	SNR 20 dB	SNR 10 dB	SNR 5 dB	SNR 0 dB
db6	.923/.923	.923/.923	.923/.923	.923/.750	1.00/.650	1.00/.684
mexh	1.00/.812	1.00/.929	1.00/.929	1.00/.619	1.00/.520	1.00/.333

detection algorithm has incomplete signal context for the first and last beats. The reported precision figures are therefore slightly conservative; longer evaluation windows would likely yield equal or higher precision for both pipelines. Second, and more significantly, both pipelines were calibrated entirely on Records 100 and 101, which represent normal sinus rhythm. Generalization to arrhythmic records is not guaranteed and constitutes an important limitation of the current study, examined in Section 4.

Figure 6 illustrates the contrast in detection quality between the best-performing pipeline (mexh CWT+WTMM) and the simplest baseline (Haar DWT) on the same 10-second window of Record 100. The mexh pipeline achieves perfect recall by following QRS maxima lines across all 128 scales before committing to a detection; the Haar pipeline's false positives arise from insufficient suppression of smooth signal components by a single-vanishing-moment wavelet. The intermediate Daubechies wavelets (db4, db6) occupy the performance space between these extremes, with db6 approaching mexh performance within the DWT framework.

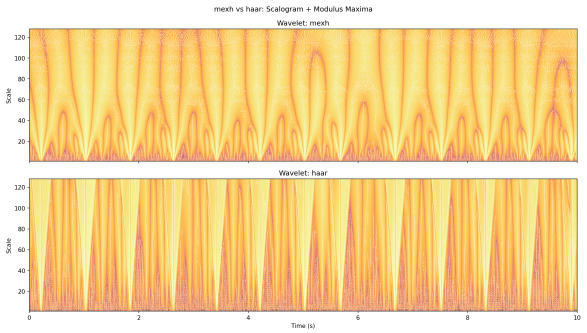

Figure 6: Side-by-side comparison of mexh (CWT+WTMM) and Haar (DWT) detections on Record 100. The mexh pipeline achieves perfect recall; Haar generates additional false positives due to its single vanishing moment, which cannot fully suppress P and T wave energy contributions.

4 Conclusions

This project applied wavelet-based singularity detection to ECG signals from the MIT-BIH Arrhythmia Database, using two distinct pipelines to compare the mathematical properties of five wavelets (Mexican hat, Haar, db2, db4, and db6). The analysis yielded four main findings.

Finding 1: Hölder regularity distinguishes QRS complexes from other ECG

features. The CWT+WTMM pipeline recovered Hölder exponents of $\hat{\alpha} \approx -1.08$ for QRS complexes across both evaluation records (see *Table* 1 and *Figure* 4). These values are well below the theoretical threshold $\alpha = 0.5$ for cusp-type singularities, confirming that the QRS depolarization front behaves as a strong pointwise singularity in the sense of Mallat and Hwang (1992). By contrast, P and T wave maxima lines yield $\hat{\alpha} \approx 1.5$, reflecting their smooth, oscillatory character. The separation between these two groups in α-space is large and consistent across records, validating Hölder regularity as a mathematically grounded tool for ECG feature discrimination. The scalogram in *Figure* 2 provides a visual complement to this result: QRS ridges persist across all 128 scales of the CWT, while P and T wave features decay rapidly, the expected signature of a sharper singularity.

Finding 2: The number of vanishing moments predicts DWT detection quality. Across the three Daubechies wavelets and Haar, the ordering of average F1 scores on Records 100 and 101 is db6 (0.869) $>$ db4 (0.766) $>$ db2 (0.649) $>$ haar (0.617), matching exactly the ordering by vanishing moments (6, 4, 2, 1). This is not coincidental. A wavelet with N vanishing moments annihilates polynomials up to degree $N - 1$, suppressing smooth background trends and reducing false positives from P and T waves. Haar, with a single vanishing moment, cannot cancel even a constant baseline offset in the reconstructed subband signal, leading to persistent false detections that no threshold setting can fully eliminate. The detection overlay in *Figure* 5 illustrates the db6 result concretely, while the wavelet comparison in *Figure* 6 makes the gap between mexh and Haar visually clear.

Finding 3: The CWT+WTMM pipeline with the Mexican hat wavelet achieves the best overall detection performance. With perfect recall on both evaluation records and an average F1 of 0.907 (*Table* 2), the mexh pipeline outperforms all five DWT configurations. This advantage comes from the multiscale nature of the detection criterion: a QRS must produce a modulus maxima line spanning at least 40 consecutive scales and a peak amplitude above the 80th percentile before it is accepted as a detection. This geometric criterion is inherently more robust than a single-scale energy threshold, because it requires evidence of singular behavior across a range of resolutions rather than at one scale alone. The cost is a more complex implementation and, as discussed below, a loss of generalization to non-standard beat morphologies.

Finding 4: The DWT adaptive threshold provides greater noise robustness at very low SNR. As shown in *Table* 3, both pipelines maintain recall at or above 0.923 across all tested SNR levels, down to 0 dB; the primary degradation is in precision. mexh precision degrades faster than db6 (0.333 vs. 0.684 at 0 dB on Record 100) because additive noise introduces spurious modulus maxima lines that narrowly pass the geometric filters. The MAD-based threshold in the DWT pipeline scales proportionally with the noise floor, providing automatic adaptation that partly compensates for signal degradation. This robustness property makes the DWT approach practically attractive for real-world ambulatory recordings, where noise levels can vary substantially.

4.1 Limitations and Future Work

Both pipelines were calibrated exclusively on Records 100 and 101, which represent normal sinus rhythm. When evaluated on arrhythmic records, performance degrades substantially. On Record 200 (PVCs and ST-segment changes), the CWT+WTMM pipeline collapses entirely (F1 = 0.000), and on Record 208 (ventricular bigeminy), it recovers only a single beat (recall = 0.062). The db6 DWT pipeline degrades more gracefully, achieving recall ≈ 0.45 and F1 ≈ 0.44 on Record 200 and recall ≈ 0.44 and F1 ≈ 0.42 on Record 208, but still falls well short of its normal-sinus performance. These results are summarized in Table 4.

Table 4: Performance on arrhythmic records without recalibration. Parameters are fixed from the normal-sinus calibration on Records 100 and 101.

Wavelet	Pipeline	Record	Recall	F1
mexh	CWT+WTMM	200	0.000	0.000
mexh	CWT+WTMM	208	0.062	0.117
db6	DWT	200	0.450	0.440
db6	DWT	208	0.438	0.424

The CWT+WTMM collapse on arrhythmic records arises from a specific interaction between the $\alpha \leq 0.5$ filter and the morphology of PVC beats. Premature ventricular contractions have a broader, more heterogeneous waveform than normal QRS complexes, and their corresponding maxima lines yield Hölder exponents $\hat{\alpha} > 0.5$, above the filter threshold. This means the α gate, designed to accept only sharp singularities, inadvertently rejects the very beats it should detect. This is not a flaw in the theoretical framework; rather, it reveals that the Hölder exponent encodes beat *morphology*, not merely beat *presence*. PVC beats have a genuinely different local regularity from normal QRS complexes, and the wavelet modulus maxima method captures this distinction.

This observation suggests a natural direction for future work. Rather than applying a single threshold $\alpha \leq 0.5$ for all beats, one could use the estimated Hölder exponent as a *classification* feature: beats with $\hat{\alpha} \approx -1.08$ are classified as normal QRS; beats with $\hat{\alpha} > 0.5$ are flagged as morphologically abnormal (e.g., PVC candidates). Such an approach would transform the pipeline from a binary detector into an arrhythmia-aware classifier, exploiting the mathematical structure of Hölder regularity directly for clinical interpretation. Additionally, extending both pipelines to full 30-minute records rather than 10-second windows, and recalibrating threshold parameters across a broader set of rhythms, would be necessary steps toward a clinically deployable system.

References

Albert Boggess and Francis J. Narcowich. 2009. *A First Course in Wavelets with Fourier Analysis*, 2 edition. John Wiley & Sons.

Ary L. Goldberger, Luis A. N. Amaral, Leon Glass, Jeffrey M. Hausdorff, Plamen Ch. Ivanov, Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and H. Eugene Stanley. 2000. Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220.

Gregory Lee, Ralf Gommers, Filip Wasilewski, Kai Wohlfahrt, and Aaron O'Leary. 2019. PyWavelets: A python package for wavelet analysis. *Journal of Open Source Software*, 4(36):1237.

Cuiwei Li, Chongxun Zheng, and Changfeng Tai. 1995. Detection of ecg characteristic points using wavelet transforms. *IEEE Transactions on Biomedical Engineering*, 42(1):21–28.

Stephane Mallat. 2008. *A Wavelet Tour of Signal Processing: The Sparse Way*, 3 edition. Elsevier/Academic Press.

Stephane Mallat and Wen Liang Hwang. 1992. Singularity detection and processing with wavelets. *IEEE Transactions on Information Theory*, 38(2):617–643.

Stephane G. Mallat. 1989. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693.

George B. Moody and Roger G. Mark. 2001. The impact of the MIT-BIH arrhythmia database. *IEEE Engineering in Medicine and Biology Magazine*, 20(3):45–50.

Zifeng Wang and Johan M. Eklund. 2017. Real-time qrs complex detection using daubechies wavelets. In *2017 IEEE 30th Canadian Conference on Electrical and Computer Engineering (CCECE)*, pages 1–4.