

Article

Band-Limited Proximal FISTA for Efficient Sparse Harmonic Recovery on MCU

Seongho Cho ¹, Minjung Kim ¹ and Daejin Park ^{2,*} ¹ School of Electronic and Electrical Engineering, Kyungpook National University, Daegu 41566, Republic of Korea² School of Electronics Engineering, Kyungpook National University, Daegu 41566, Republic of Korea

* Correspondence: boltanut@knu.ac.kr; Tel.: +82-53-950-5548

Abstract

Compressed sensing (CS) enables signal reconstruction from fewer measurements when the signal is sparse in a transform domain. However, executing ℓ_1 -regularized recovery on MCU-class hardware is challenging due to limited compute resources and the cost of repeated forward and adjoint operator evaluations. This paper presents a band-limited proximal variant of FISTA that enforces known spectral support during thresholding, restricting the effective optimization domain without changing the measurement model. We implement a complete CS reconstruction pipeline on an STM32F407 (Cortex-M4) using CMSIS-DSP FFT/IFFT kernels and evaluate it using ECG waveforms acquired through an AD8232 front end as benchmark signals. With $M = 340$ measurements (33% of uniform sampling), the embedded implementation achieves a PRDN of 24.38%, closely matching MATLAB references (CVX: 22.64%, FISTA: 22.39%) under identical hyperparameters. Cycle-accurate profiling shows that FFT/IFFT-based forward/adjoint operators dominate the per-iteration runtime. Under a 60 Hz band-limited setting, the required iterations are reduced from 30 to 16 with an acceptable PRDN, demonstrating a practical trade-off between reconstruction accuracy and computational cost on MCU-class devices.

Keywords: compressed sensing; sparse recovery; band-limited proximal operator; ℓ_1 -regularized optimization; FISTA; Cortex-M4

1. Introduction

The classical Shannon–Nyquist sampling theorem states that a band-limited signal can be perfectly reconstructed when sampled at a rate at least twice its highest frequency. This principle has shaped signal acquisition systems for decades. The emergence of compressed sensing (CS), introduced by Donoho [1] and by Candès et al. [2], showed that perfect reconstruction is also possible from sub-Nyquist measurements when the signal is sparse in an appropriate transform domain. Since then, a large body of work has developed sparse measurement models and efficient reconstruction algorithms, including Iterative Shrinkage-Thresholding Algorithm (ISTA), Fast Iterative Shrinkage-Thresholding Algorithm (FISTA), Iterative Hard Thresholding (IHT), and related variants [3–5].

Although FISTA is widely used due to its simplicity and fast convergence rate, its standard formulation assumes computing environments with ample memory bandwidth and floating-point throughput. In contrast, microcontroller (MCU)-class hardware faces tight constraints in memory, energy, and real-time requirements, and the per-iteration cost of FFT-based gradient evaluations becomes a major bottleneck. Furthermore, many



Academic Editors: Qingli Li, Yan Wang, Qing Zhang and Qian Wang

Received: 22 December 2025

Revised: 29 January 2026

Accepted: 21 June 2026

Published: 25 June 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

practical signals exhibit known structural properties—such as strict band-limitedness—that are not exploited in the conventional proximal step. Incorporating these priors directly into the optimization can reduce the search space and improve convergence efficiency, which is particularly valuable for embedded deployments.

Motivated by these observations, we propose a *band-limited proximal FISTA* framework that enforces known spectral support during the thresholding stage. This modification leaves the measurement model unchanged but significantly reduces the effective optimization dimension, yielding faster convergence without degrading reconstruction accuracy. To assess its feasibility on constrained hardware, we implement the complete CS reconstruction pipeline on an STM32F407 [6] MCU using CMSIS-DSP FFT/IFFT kernels [7].

In this study, ECG waveforms acquired through an AD8232 [8] front end are used solely as representative benchmark signals for evaluating algorithmic and architectural feasibility. The focus of this work is on embedded sparse recovery performance rather than clinical interpretation.

The main contributions of this work are as follows:

- We propose a band-limited proximal FISTA that incorporates known spectral support, reducing the optimization dimension and achieving up to 50% reduction in iteration count.
- We demonstrate that full ℓ_1 -minimization recovery is feasible on a resource-constrained STM32F407 MCU, without offloading computation.
- We present an embedded-optimized implementation using CMSIS-DSP and provide a detailed cycle-level analysis showing that the workload is compute-bound rather than memory-bound.
- We validate the proposed method using real benchmark signals acquired on hardware, confirming its suitability for low-power embedded platforms.

The remainder of this paper is organized as follows. Section 2 reviews sparse recovery and FISTA. Section 3 summarizes related embedded CS research. Section 4 details the proposed band-limited proximal FISTA and implementation aspects. Section 5 presents reconstruction and timing results on the STM32F407. Section 6 discusses limitations and future directions, and Section 7 concludes the work.

2. Background

Compressed sensing (CS) provides a framework for reconstructing high-dimensional signals from a reduced number of linear measurements when the signals are sparse or compressible in an appropriate transform domain. In many practical scenarios—including biomedical waveforms, vibration data, and periodic sensing signals—the essential information resides within a limited spectral band or can be represented using a small subset of transform coefficients. Such structure makes CS attractive for embedded systems where memory, energy, and computational resources are constrained.

In this study, ECG recordings are used solely as representative *benchmark signals* to evaluate the feasibility of on-device sparse recovery. The goal is not clinical interpretation, but to examine the algorithmic and implementation aspects of CS reconstruction on low-power MCUs. Accordingly, the discussion below summarizes CS principles in a general signal-processing context.

2.1. Acquisition Model

Let $\mathbf{x} \in \mathbb{R}^N$ denote a discrete-time signal of length N , and let $\mathbf{A} \in \mathbb{R}^{M \times N}$ be a measurement matrix, where M is the number of acquired measurements and satisfies $M \ll N$. The CS acquisition model is given by

$$\mathbf{y} = \mathbf{A}\mathbf{x}, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^M$ represents the compressed measurement vector obtained by linear projection of the original signal.

If the signal $\mathbf{x}$ admits a sparse representation under a known transform basis $\mathbf{\Psi} \in \mathbb{C}^{N \times N}$, it can be expressed as

$$\mathbf{x} = \mathbf{\Psi}\mathbf{s}, \quad (2)$$

where $\mathbf{s} \in \mathbb{C}^N$ is the sparse coefficient vector whose elements contain only a small number of significant nonzero values.

Substituting Equation (2) into Equation (1) yields the equivalent sensing model

$$\mathbf{y} = \mathbf{A}\mathbf{\Psi}\mathbf{s} = \mathbf{\Phi}\mathbf{s}, \quad (3)$$

where $\mathbf{\Phi} = \mathbf{A}\mathbf{\Psi} \in \mathbb{C}^{M \times N}$ denotes the effective sensing matrix.

Random row selection of Fourier or cosine transforms produces sensing matrices with low coherence relative to common sparsifying bases, enabling stable recovery of sparse vectors [9].

2.2. Sparse Recovery via Optimization

Recovering the sparse coefficient vector $\mathbf{s}$ from Equation (3) constitutes an underdetermined inverse problem. Direct ℓ_0 minimization is NP-hard and is therefore commonly relaxed to an ℓ_1 -minimization problem. A representative formulation is Basis Pursuit (BP),

$$\min_{\mathbf{s}} \|\mathbf{s}\|_1 \quad \text{s.t.} \quad \mathbf{\Phi}\mathbf{s} = \mathbf{y}, \quad (4)$$

where $\|\mathbf{s}\|_1$ denotes the ℓ_1 norm of the coefficient vector.

Exact recovery is possible under Restricted Isometry Property (RIP) or coherence conditions [10,11], with $M = O(K \log(N/K))$ measurements for a K -sparse signal.

While convex solvers for Equation (4) provide strong theoretical guarantees, their computational and memory demands often make them impractical for MCU-class embedded systems. As a result, first-order iterative methods based on proximal gradient updates have become a preferred alternative.

In this work, we focus on the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA), which solves the ℓ_1 -regularized least-squares problem

$$\min_{\mathbf{s}} \frac{1}{2} \|\mathbf{\Phi}\mathbf{s} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{s}\|_1, \quad (5)$$

where $\lambda > 0$ is a regularization parameter that controls the sparsity of the solution.

By incorporating Nesterov acceleration, FISTA achieves an optimal $O(1/k^2)$ convergence rate for first-order methods, making it particularly attractive for large-scale problems and resource-constrained implementations.

In its standard form, FISTA applies the proximal (soft-thresholding) operator uniformly to all transform coefficients. However, many practical sensing signals exhibit known structural priors, such as strict band-limited spectral support, that are not explicitly exploited in the conventional proximal step. Incorporating such prior knowledge directly into the proximal mapping can reduce the effective optimization domain and accelerate convergence, an idea that forms the basis of the band-limited proximal FISTA proposed in this work.

2.3. Relevance to Embedded Systems

On MCU-class platforms, the computational cost of iterative sparse recovery algorithms is often dominated by repeated evaluations of the forward and adjoint operators.

When these operators are implemented using FFT/IFFT-based transforms, the overall runtime can become sensitive to the number of iterations, particularly under limited computational resources. These considerations motivate algorithmic strategies that reduce the number of iterations or constrain the effective optimization domain, which is especially relevant for embedded implementations.

3. Related Work

Research on sparse recovery algorithms is extensive, with numerous studies analyzing convergence behavior, accuracy, and computational complexity. However, considerably fewer works investigate how such algorithms can operate within the strict memory and architectural constraints of MCU-class embedded systems. This section reviews prior work from three perspectives: (i) algorithmic developments in sparse reconstruction, (ii) embedded implementations of sparse solvers, and (iii) remaining gaps in algorithm–hardware integration.

3.1. Sparse Recovery Algorithms and Computational Characteristics

Sparse reconstruction is commonly formulated as an ℓ_1 -regularized inverse problem. Convex optimization methods such as Basis Pursuit (BP) and Alternating Direction Method of Multipliers (ADMM) [12,13] provide strong theoretical recovery guarantees but require repeated matrix–vector multiplications and transform operations, resulting in high computational demand. First-order iterative methods, such as ISTA and FISTA, reduce per-iteration computational complexity by avoiding matrix inversions and are therefore widely used for large-scale sparse reconstruction problems [14]. However, their dominant cost typically arises from evaluating the forward and adjoint operators (Φs , $\Phi^H r$), often implemented via FFT/IFFT transforms or structured linear mappings. On hardware without dedicated accelerators, these operations impose substantial latency and buffer footprint, making native implementations challenging for MCUs.

3.2. Embedded Implementations of Sparse Recovery

Sparse recovery on embedded platforms has been explored primarily with lightweight algorithms. MCU-based implementations usually employ greedy or thresholding methods such as Orthogonal Matching Pursuit (OMP) [15] and Iterative Hard Thresholding (IHT), or simplified pursuit variants, because they avoid transform-domain operators and require minimal memory. These techniques reduce computational load but generally achieve lower reconstruction accuracy and lack the robustness of convex ℓ_1 -based methods.

Convex optimization-based solvers, including FISTA, remain rare on microcontroller-class platforms due to their reliance on iterative forward/adjoint operations, multiple large intermediate buffers, and sensitivity to memory and computational bandwidth. As a result, most embedded CS systems perform only signal acquisition or basic preprocessing on the MCU, while sparse reconstruction is executed on an external PC, server, or base station [16–18]. FPGA-oriented implementations demonstrate high throughput through extensive parallelism and pipelining [19], but such architectures are not directly transferable to general-purpose MCUs, making them unsuitable for low-power, low-cost embedded devices [20].

Recent studies have further explored the application of CS on embedded MCUs, wearable devices, and specialized hardware platforms. For example, the COSMO framework [21,22] applies compressed sensing to model and logging optimization in MCU performance screening, primarily targeting system-level monitoring and data reduction. Similarly, an embedded programmable processor for compressive sensing was presented

in [23], where architectural support and instruction-level acceleration are introduced to improve the efficiency of CS-related operations.

In the biomedical sensing domain, Ahmmed et al. [24,25] demonstrated a wearable wristband employing compressed sensing for photoplethysmography (PPG), highlighting the potential of CS to reduce data rate and power consumption in integrated wearable platforms. Related ECG-focused studies have also investigated alternative sparsifying transforms for compressed sensing. For instance, Hassan and Mohsen [26] proposed a curvelet-transform-based CS framework to improve ECG reconstruction performance under high compression ratios, emphasizing transform-domain signal representation.

While these studies demonstrate the feasibility of compressed sensing in embedded or specialized hardware environments, their primary emphasis lies in system-level data handling, architectural acceleration, or custom processor design. In contrast, the present work differs fundamentally in how prior information is incorporated into the recovery algorithm. Rather than accelerating an existing sparse recovery pipeline or modifying regularization weights, we explicitly embed known band-limited spectral support as a convex constraint within the proximal operator of FISTA. This formulation directly reduces the feasible optimization domain itself, leading to fewer effective degrees of freedom and a reduced number of iterations.

From this perspective, the proposed method reinterprets structural prior information as an explicit mechanism for controlling optimization complexity, rather than solely as a means of improving reconstruction accuracy. By coupling this band-limited proximal design with an FFT-based measurement model and executing the full iterative recovery loop on a general-purpose Cortex-M4 MCU, this work complements prior embedded CS studies by demonstrating how algorithmic structure can be leveraged to achieve practical convex sparse recovery under strict microcontroller constraints.

3.3. Unresolved Gaps in Algorithm–Hardware Integration

From the prior literature, two unresolved gaps can be identified:

- **Algorithmic gap:** Existing FISTA- and ISTA-based methods typically treat the proximal operator as a generic sparsity-inducing mapping and do not explicitly exploit structured priors, such as band-limited spectral support, that are common in physical sensing signals. As a result, the effective optimization domain remains unnecessarily large, and the impact of structure-aware proximal mappings on convergence behavior has received limited attention.
- **Hardware gap:** While FFT/IFFT-based forward and adjoint operators are widely assumed in convex sparse recovery, their interaction with microcontroller constraints—including memory layout, transform cost, DSP-kernel locality, and iteration-level execution behavior—has not been systematically analyzed at the MCU level. In particular, the feasibility of executing a full iterative reconstruction pipeline entirely on a microcontroller remains insufficiently documented in existing embedded CS studies.

These gaps suggest that both algorithmic refinements and hardware-aware implementation studies are necessary to assess the practical feasibility of convex sparse recovery on MCU-class devices.

3.4. Positioning of the Present Work

The present study addresses these gaps by: (i) introducing a *band-limited proximal operator* that incorporates spectral prior information to reduce the optimization domain, (ii) implementing full FISTA recovery—including the forward and adjoint FFT/IFFT operators—on a Cortex-M4 MCU, and (iii) analyzing how memory organization and DSP-

kernel usage affect iterative sparse reconstruction on constrained hardware. Unlike prior embedded CS systems, which typically perform only acquisition on the MCU, this work demonstrates that complete ℓ_1 -minimization-based recovery is feasible on an MCU when algorithmic and architectural choices are jointly optimized.

4. Proposed Architecture

The proposed architecture performs compressive acquisition and sparse reconstruction of biomedical benchmark signals on an embedded MCU, as illustrated in Figure 1. Instead of acquiring a full set of Nyquist-rate samples, the system collects a reduced number of measurements via a random subsampling operator implemented directly on the MCU. Sparse signal recovery is then performed on-device using a FISTA-based ℓ_1 -minimization framework, enabling reconstruction from sub-Nyquist measurements while respecting the memory and computational constraints of the embedded platform.




Figure 1. Block diagram of the proposed band-limited FISTA-based reconstruction framework.

The key distinction from standard FISTA lies in the proposed *band-limited proximal operator*. Unlike conventional uniform soft-thresholding, the proposed proximal step explicitly enforces known spectral support by applying soft-thresholding to in-band coefficients while hard-zeroing out-of-band components. This modification preserves the original sensing model and gradient computation, yet significantly reduces the effective optimization dimension, thereby accelerating convergence and improving computational efficiency under MCU constraints.

4.1. Fundamentals of CS-Based ECG Signal Acquisition

When the input signal vector x of length N , sampled at or above the Nyquist rate, is transformed into the frequency domain, the matrix Ψ in Equation (2) corresponds to the IDFT matrix, and s represents the DFT coefficients. In CS, the key aspect of the acquisition stage is to design a measurement matrix A that can effectively compress the signal. For monitoring-grade ECG signals, the dominant spectral energy lies below 60 Hz, and the high-frequency diagnostic components are negligible for rhythm analysis [27,28]. As a result, the signal exhibits effective sparsity in the Fourier domain. The DFT matrix, which transforms a signal into its frequency components, is particularly suitable for this purpose. Given such sparsity, the DFT transformation maps most elements to values close to zero,

which is advantageous for CS. Moreover, the columns of the DFT matrix are mutually orthogonal, making it an effective basis for accurate signal reconstruction.

When the sensing matrix Φ is obtained by randomly selecting rows from an IDFT matrix, successful signal reconstruction can often be achieved even if the RIP condition is not strictly satisfied. This is mainly attributed to the high incoherence between the random selection matrix and the IDFT matrix, as well as the inherent sparsity of the signal. Although the RIP condition may not be fully met, the DFT matrix can still effectively compress the frequency components of the signal and, when combined with algorithms such as ℓ_1 minimization, provide excellent reconstruction performance.

The block diagram of the proposed acquisition stage for compressed ECG signals is shown in Figure 2. To realize random undersampling, a pseudo-random number generator (PRNG) was implemented on the MCU. The PRNG generated a predefined sequence of random indices, and the ADC was triggered only when the current index matched the sequence. This design choice enables the random sampling pattern to be directly realized in the MCU without requiring additional external analog modulation circuitry, while maintaining precise synchronization with the input signal. Consequently, this approach effectively reduces both the number of samples and the amount of acquired data.




Figure 2. CS-based signal acquisition stage.

Representing the system in matrix form facilitates understanding of its properties through matrix analysis. When the number of samples in the signal x is N , the acquisition of M ($M \ll N$) CS measurements y can be expressed as in Equation (3), where the pseudo-random sequence is represented by an $M \times N$ measurement matrix A . Let $\Psi \in \mathbb{C}^{N \times N}$ denote the inverse discrete Fourier transform matrix, whose (k, n) -th entry is

$$\Psi_{k,n} = \frac{1}{N} e^{j \frac{2\pi}{N} kn}, \quad k, n = 0, 1, \dots, N-1. \quad (6)$$

Let $A \in \{0, 1\}^{M \times N}$ be the row-selection matrix obtained by randomly selecting M rows from the $N \times N$ identity matrix I_N , i.e.,

$$A = \begin{bmatrix} e_{t_1}^\top \\ e_{t_2}^\top \\ \vdots \\ e_{t_M}^\top \end{bmatrix}, \quad \{t_1, t_2, \dots, t_M\} \subset \{0, 1, \dots, N-1\}, \quad (7)$$

where e_i denotes the i -th canonical basis vector of $\mathbb{R}^N$. Then, the measurement process can be expressed as in Equation (3).

4.2. CS-Based ECG Signal Reconstruction

We evaluate several representative sparse recovery algorithms suitable for MCU-class platforms. The objective is to estimate the sparse spectral coefficient vector $s \in \mathbb{C}^N$ from

$M \ll N$ compressed measurements $\mathbf{y} \in \mathbb{R}^M$ according to Equation (3). As discussed in Section 2, this inverse problem is underdetermined and requires additional sparsity constraints to obtain a unique solution. Although the original ECG signal is real-valued, a complex-valued representation is employed in the Fourier domain to enable FFT/IFFT-based implementation.

Classical greedy methods such as OMP or Adaptive OMP approximate the ℓ_0 -minimization problem but involve iterative orthogonal projections and large matrix operations, which demand substantial memory and computational resources. CoSaMP improves reconstruction robustness but introduces multiple support-identification and least-squares stages that exceed the capabilities of memory-limited MCUs.

To avoid these limitations, we adopt a convex formulation based on the LASSO objective and employ FISTA due to its $O(1/k^2)$ convergence rate and low per-iteration complexity.

The differentiable and non-differentiable components of Equation (5) are

$$f(\mathbf{s}) = \frac{1}{2} \|\Phi \mathbf{s} - \mathbf{y}\|_2^2, \quad (8)$$

$$g(\mathbf{s}) = \lambda \|\mathbf{s}\|_1. \quad (9)$$

The gradient $\nabla f(\mathbf{s}) = \Phi^H(\Phi \mathbf{s} - \mathbf{y})$ is computed via FFT/IFFT-based forward and adjoint operators implemented using the CMSIS-DSP library to meet real-time MCU constraints.

Band-Limited Proximal Operator (Proposed).

While standard FISTA applies a uniform proximal operator to all transform coefficients, and existing structure-aware sparse recovery methods typically encode prior information through weighted ℓ_1 , group-sparsity penalties, or adaptive support estimation, they do not explicitly reduce the feasible optimization domain itself.

Rather than introducing a heuristic modification, we explicitly incorporate this prior knowledge as a *convex support constraint* in the optimization problem. Specifically, the proposed operation corresponds to the *exact proximal mapping* of a well-defined convex regularizer, rather than an ad hoc masking or pruning strategy.

To this end, we redefine the non-smooth term as

$$g_B(\mathbf{s}) = \lambda \|\mathbf{s}\|_1 + I_{S_B}(\mathbf{s}) \quad (10)$$

where $I_{S_B}(\cdot)$ is the indicator function of the band-limited support set

$$S_B = \{\mathbf{s} \in \mathbb{C}^N \mid s_i = 0, i \notin \mathcal{B}\} \quad (11)$$

The set S_B is closed and convex; therefore, $g_B(\cdot)$ is a proper, closed, and convex regularizer. Accordingly, the proximal step in FISTA is given by the exact proximal mapping

$$\text{prox}_{g_B}(\mathbf{v}) = \arg \min_{\mathbf{s}} \frac{1}{2} \|\mathbf{s} - \mathbf{v}\|_2^2 + \lambda \|\mathbf{s}\|_1 + I_{S_B}(\mathbf{s}) \quad (12)$$

which admits a closed-form solution.

Let $\mathcal{B}$ denote the frequency bins that contain physiologically relevant content (e.g., below 60 Hz), based on well-established prior knowledge of ECG spectral bandwidth [27], and let $\mathbf{v} = \mathbf{z}_k - \frac{1}{L} \nabla f(\mathbf{z}_k)$ denote the gradient step. The resulting proximal update is

$$s_{k+1,i} = \begin{cases} \mathcal{S}_{\lambda/L}(v_i), & i \in \mathcal{B}, \\ 0, & i \notin \mathcal{B}, \end{cases} \quad (13)$$

which corresponds to soft-thresholding on $\mathcal{B}$ and exact projection onto S_B outside the band.

The elementwise soft-thresholding operator is defined as

$$\mathcal{S}_\tau(v_i) = \begin{cases} \left(1 - \frac{\tau}{|v_i|}\right)v_i, & |v_i| > \tau, \\ 0, & |v_i| \leq \tau, \end{cases} \quad \tau = \lambda/L \quad (14)$$

Theoretical Consistency with Standard FISTA.

This modification accomplishes three goals simultaneously: (i) it suppresses out-of-band spectral components that are known *a priori* to contain predominantly noise, (ii) it explicitly reduces the effective optimization dimension, thereby lowering the required number of iterations, and (iii) it preserves the composite convex structure of the objective. Since $g_B(\cdot)$ is proper, closed, and convex, the resulting proximal mapping is firmly non-expansive; therefore, the standard convergence guarantees of FISTA continue to apply. Moreover, since ∇f is Lipschitz-continuous and L is chosen accordingly, the standard FISTA assumptions remain satisfied.

Unlike weighted ℓ_1 or group-sparsity formulations, which modify the penalty structure to promote specific sparsity patterns, the proposed approach enforces a fixed, physically motivated spectral support as a hard convex constraint, thereby reducing the optimization domain rather than reweighting coefficients. Furthermore, while support-aware sparse recovery methods often rely on adaptive or data-driven support estimation, the proposed method assumes a fixed and *a priori* known spectral support derived from sensing physics, eliminating the need for iterative support detection.

Overall Algorithm.

The resulting band-limited FISTA procedure is summarized in Algorithm 1. The proposed method does not introduce a new convergence theory, but rather provides a structured proximal design within the standard FISTA framework. By constraining the optimization domain through an explicit convex support set, the effective search space of the sparse coefficients is reduced, leading to faster and more stable convergence under strict MCU resource constraints without altering the theoretical foundation of the algorithm.

To characterize the sensing configuration, we measured the mutual coherence between the measurement matrix A and the sparse basis matrix Ψ . The computed values are 0.0221 for $N = 2048$, 0.0313 for $N = 1024$, and 0.0442 for $N = 512$, indicating that coherence decreases as the signal dimension increases. These values are sufficiently low—well below commonly cited practical thresholds (e.g., $\mu < 0.1$)—and therefore favorable for stable sparse recovery [29]. Intuitively, mutual coherence quantifies the maximum correlation between any measurement vector and any basis vector. A low coherence value means that no single measurement captures a disproportionate amount of information about any specific frequency component. As a result, information about the sparse coefficients is spread more uniformly across the measurements, which makes the inverse problem better conditioned and improves the robustness of the recovery algorithm.

In summary, the proposed method embeds an *a priori* band-limited support as an exact proximal-level convex constraint, reducing the feasible optimization domain and iteration cost without altering standard FISTA theory.

Algorithm 1 FISTA with band-limited proximal operator for ℓ_1 -regularized recovery with $\Phi = A\Psi$

Require: Measurements $y \in \mathbb{C}^M$, step size $L > 0$,

regularization $\lambda > 0$, max iterations K , in-band index set $\mathcal{B} \subset \{0, \dots, N-1\}$

Ensure: Estimated sparse vector $\hat{s}$

1: Initialize $s_1 \leftarrow 0, z_1 \leftarrow s_1, t_1 \leftarrow 1$

2: **for** $k = 1$ **to** K **do**

3: $u \leftarrow \text{apply_Phi_ifft}(z_k)$ % Forward operator

4: $r \leftarrow u - y, g \leftarrow \text{apply_PhiH_fft}(r)$ % Residual and adjoint

5: $v \leftarrow z_k - \frac{1}{L}g$ % Gradient step

6: **for each** $i \in \mathcal{B}$ **do**

$$s_{k+1,i} \leftarrow \mathcal{S}_{\lambda/L}(v_i)$$

7: **for each** $i \notin \mathcal{B}$ **do**

% Exact projection onto convex support set $\mathcal{S}_{\mathcal{B}}$

$$s_{k+1,i} \leftarrow 0$$

8: $t_{k+1} \leftarrow \frac{1 + \sqrt{1 + 4t_k^2}}{2}$

9: $z_{k+1} \leftarrow s_{k+1} + \left(\frac{t_k - 1}{t_{k+1}}\right)(s_{k+1} - s_k)$ % Nesterov acceleration update

10: **if** $\|s_{k+1} - s_k\|_2 < \varepsilon$ **then break** % Convergence check

11: **end for**

12: **return** $\hat{s} \leftarrow s_{k+1}$

5. Experiment

To verify the effectiveness of the proposed embedded system-based CS and recovery approach, ECG signals were randomly sampled in time to acquire only M measurements, rather than uniformly sampling all N points, and subsequently reconstructed on an STM32F407 Discovery board, which features a Cortex-M4 architecture (168 MHz, with an integrated FPU) as shown in Figure 3a. The recovery performance was evaluated by comparing the reconstructed signals with the original signals in terms of the Normalized Percent Root-mean-square Difference (PRDN).




Figure 3. (a) Experimental hardware setup with STM32F407 Discovery board and AD8232 ECG front-end module. (b) DFT coefficients of the ECG signal.

5.1. Experimental Setup

5.1.1. CS-Based ECG Signal Acquisition

In this study, the AD8232 device from Analog Devices was employed as the analog front end for ECG signal acquisition. The AD8232 is a high-performance, fully integrated analog front end designed specifically for ECG measurement, capable of stable biopotential detection under low-power operation, and supporting a single-lead measurement configu-

ration. ECG signals were acquired using a three-electrode configuration (RA, LA, RL) to ensure stable signal detection from the human body. The detected analog ECG signals were then fed into the on-chip ADC of the STM32F407 MCU, where two separate ADC channels were utilized to perform sampling for different purposes.

ADC1 was configured to uniformly sample the ECG signal at a rate that satisfies the Nyquist criterion and served as a reference for evaluating the reconstruction accuracy. ADC2 received the same ECG input but performed random sampling at arbitrary time indices, acquiring approximately 33% of the samples required by the Nyquist rate for compressed sensing (CS). Both ADC channels shared the same timer-based sampling clock, and the random samples acquired by ADC2 corresponded to a subset of the uniformly sampled reference signal from ADC1, ensuring temporal alignment during PRDN evaluation. Both channels operated at a sampling frequency of 500 Hz, which is sufficient for monitoring-grade ECG acquisition, as rhythm analysis relies primarily on low-frequency components below 60 Hz. In our prototype experiments, the benchmark ECG signals exhibit dominant spectral energy within this frequency range, while higher-frequency components are not essential for the intended rhythm-oriented evaluation. The ADC resolution was set to 10 bits, yielding digital values in the range of 0–1023. To support real-time processing, all sampled data were transferred directly to RAM using direct memory access (DMA), thereby minimizing CPU involvement during data acquisition.

For the ADC1 channel, a 1024-sample buffer was allocated in RAM, while the ADC2 channel was assigned a 340-sample buffer to accommodate its reduced sampling size. At the configured sampling frequency of 500 Hz, the 1024-sample buffer can store approximately 2.048 s of ECG data (1024/500). The data acquired from the ADC2 channel through on-chip timer-driven random sampling were subsequently used as input to the FISTA-based sparse signal recovery algorithm. This process forms the experimental basis for evaluating signal reconstruction performance under CS conditions.

To characterize the spectral properties of the benchmark ECG signal used in the experiments, the uniformly sampled data acquired from the ADC1 channel at 500 Hz were transformed into the frequency domain using the DFT. The resulting magnitude spectrum, shown in Figure 3b, indicates that the signal energy is predominantly concentrated in a limited number of low-frequency components. This observation supports the sparsity assumption underlying the CS-based reconstruction framework employed in this study.

5.1.2. Signal Recovery via FISTA

To effectively implement FISTA in computation-limited environments such as embedded systems, this study applies several optimization techniques. Although the STM32F407 includes a hardware FPU, floating-point operations are still more computationally expensive than fixed-point operations. Therefore, the algorithm was designed to minimize the use of floating-point arithmetic to meet real-time processing requirements. In particular, the repeated complex operations involved in DFT transforms impose a significant computational burden and directly affect execution time. To address this, the forward operation stage in Algorithm 1 replaces direct matrix multiplication with the inverse FFT function from the CMSIS-DSP library, and the gradient computation stage for A^H is implemented using the FFT. This FFT-based implementation substantially reduces computational complexity and improves execution speed while maintaining reconstruction accuracy.

In addition, the input data are normalized to the range $[-1, 1]$ prior to FISTA reconstruction. This simple preprocessing keeps intermediate values within a predictable dynamic range, preventing overflow/underflow and avoiding per-iteration clipping or rescaling. By fixing the effective scale of the sensing operator, the Lipschitz constant of the

gradient is well conditioned, which permits a stable (often larger) step size and reduces the number of iterations needed to reach the target tolerance.

5.2. Recovery Accuracy Evaluation

Table 1 summarizes the reconstruction accuracy measured in terms of PRDN for different recovery methods under identical measurement conditions ($M = 340$). The MATLAB-based CVX solution serves as a reference benchmark, while the FISTA implementations in MATLAB and on the STM32F407 demonstrate comparable performance with only a modest accuracy degradation on the MCU platform.

Table 1. Recovery accuracy comparison in terms of PRDN. For all ℓ_1 -regularized methods (CVX and FISTA), hyperparameters ($M = 340, \lambda = 0.0005, L = 0.001$) are applied; FISTA was run for 30 iterations.

Method	PRDN (%)
Matlab CVX	22.64
Matlab FISTA	22.39
STM32F4 FISTA	24.38

The PRDN, calculated according to Equation (15), is a widely used metric for quantitatively evaluating signal reconstruction performance by measuring the deviation between the original signal and the reconstructed signal.

$$\text{PRDN} = \frac{\sqrt{\sum_{i=1}^N (x_i - \hat{x}_i)^2}}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2}} \times 100\% \quad (15)$$

where x_i denotes the i -th sample of the original signal, $\hat{x}_i$ is the i -th sample of the reconstructed signal, $\bar{x}$ is the mean value of the original signal, and N is the total number of samples.

PRDN is derived from the Percentage Root-mean-square Difference (PRD) by removing the DC component of the reference signal, thereby providing a more robust measure against baseline shifts [30]. A lower PRDN value indicates greater similarity between the reconstructed and original signals. Although there is no universally accepted threshold, lower PRDN values are generally associated with improved reconstruction fidelity in ECG compression studies [30,31].

To verify the reconstruction accuracy, reconstruction was performed in a MATLAB environment using CVX (version 2025a) under the same measurement setting ($M = 340, \lambda = 0.0005, L = 0.001, 30$ iterations), yielding a PRDN of 22.64%. CVX was used to obtain a reference solution by solving the same ℓ_1 -regularized LASSO objective in Equation (5), serving as a high-accuracy benchmark. Furthermore, applying FISTA with the same hyperparameters resulted in a PRDN of 22.39%, which is comparable to the CVX result. Finally, the proposed STM32F407 implementation achieved a PRDN of 24.38%. For a fair and reproducible comparison, the MATLAB and MCU implementations of FISTA were evaluated using identical hyperparameters and a fixed iteration count, while CVX was used solely to obtain a high-accuracy reference solution.

The slight degradation in PRDN observed on the MCU platform is primarily attributed to single-precision arithmetic and numerical rounding effects introduced by FFT/IFFT operations on the embedded processor. This difference reflects the inherent numerical limitations of embedded single-precision implementations rather than algorithmic deficiencies. Nevertheless, the reconstruction accuracy remains within the high-fidelity range,

confirming that precise sparse recovery is achievable even under the computational and memory constraints of MCU-class hardware.

5.3. Iteration-Level Runtime and Memory Footprint Analysis

Table 2 reports the cycle-level execution time of one FISTA iteration on the STM32F407 platform. The notation in the first column follows Algorithm 1, enabling direct correspondence between algorithmic steps and their computational cost. As shown, the forward and adjoint operators—Steps (1) and (3)—dominate the runtime due to the CMSIS-DSP FFT/IFFT kernels, jointly accounting for more than 66% of the total iteration time. The proposed band-limited proximal operator in Step (4) accounts for 81,078 cycles (494 μ s), reflecting the cost of magnitude evaluation and conditional shrinkage only within the informative subspace $\mathcal{B}$. The overall per-iteration cost is approximately 3.29 ms.

Table 2. Cycle-level execution time breakdown per FISTA iteration on the STM32F407 (notation corresponds to Algorithm 1).

Step	Operation	Cycles	Time (μ s)
(1)	Forward operator $u = \Phi z_k$ (IFFT)	200,439	1193.09
(2)	Residual update $r = u - y$	8856	52.71
(3)	Adjoint operator $g = \Phi^H r$ (FFT)	163,211	971.49
(4)	Band-limited proximal step $s_{k+1} = \mathcal{P}_{\mathcal{B}}(\mathcal{S}_{\lambda/L}(v))$	81,078	494.51
(5)	Nesterov acceleration update $z_{k+1} = s_{k+1} + \frac{t_k - 1}{t_{k+1}}(s_{k+1} - s_k)$	26,721	159.05
(6)	Convergence check $\ s_{k+1} - s_k\ _2 < \varepsilon$	20,499	122.02
Total	Per-iteration cost	552,053	3286.03

Although the band-limited proximal step introduces additional conditional logic compared to standard soft-thresholding, its contribution to the per-iteration cost remains secondary. As shown in Table 2, FFT/IFFT-based forward and adjoint operations dominate the runtime on the MCU, accounting for more than 65% of the per-iteration cycles. By explicitly reducing the effective optimization dimension, the proposed method converges in significantly fewer iterations, such that the modest increase in proximal cost is amortized over the reduced iteration count, resulting in a net reduction in total execution time.

In the baseline implementation, all major buffers are allocated in SRAM. We additionally evaluated relocation of the critical buffers for Steps (1)–(4) to CCMRAM, but observed less than 1% improvement, indicating compute-bound behavior of the FFT/IFFT kernels on Cortex-M4. This negligible improvement indicates that the dominant operators in Steps (1) and (3), both relying on CMSIS-DSP FFT/IFFT kernels, are *compute-bound* rather than *memory-bound* on the Cortex-M4 architecture. In other words, the critical path is limited by floating-point arithmetic throughput (butterfly operations and twiddle-factor multiplications), not by memory bandwidth or buffer locality.

Consequently, relocating data to faster memory cannot meaningfully accelerate the overall algorithm. Performance gains must instead come from reducing the total amount of computation—for example, via the proposed band-limited proximal operator or iteration-count minimization—rather than from modifying the memory hierarchy.

The working memory required by the proposed FISTA implementation is summarized in Table 3. All major buffers used in the iterative procedure—including the current iterate, residuals, gradient vectors, and the FFT/IFFT workspace—are allocated in SRAM. The total footprint is approximately 40 KB, well within the 192 KB SRAM budget of the STM32F407.

Under the on-chip SRAM constraints of the STM32F407, $N = 2048$ represents a practical upper bound for the current implementation without additional memory optimization.

Table 3. Memory footprint of the proposed FISTA implementation on STM32F407. Each complex-valued vector requires N samples stored in interleaved single-precision format (real and imaginary parts), resulting in $N \times 4 \times 2 = 8$ KB per buffer for $N = 1024$.

Buffer (in RAM)	Size
s_k (solution estimate)	8 KB
z_k (Nesterov vector)	8 KB
g_k (gradient vector)	8 KB
residual (residual vector)	8 KB
fft_freq_buf (FFT/IFFT workspace)	8 KB
Frequency-band mask $\mathcal{B}$	<1 KB
Total working memory	≈40 KB

5.4. Iteration–Accuracy Trade-Off Induced by the Band-Limited Proximal Operator

Figure 4a illustrates the impact of the imposed spectral bandwidth constraint on both the convergence behavior and reconstruction accuracy of the proposed band-limited proximal FISTA. Two distinct and opposing trends can be observed, revealing an explicit iteration–accuracy trade-off induced by the modified proximal operator.




Figure 4. (a) Effect of the imposed bandwidth limit on the required number of iterations (left axis) and reconstruction accuracy measured by PRDN (right axis) for the proposed band-limited proximal FISTA. (b) Time-domain comparison between the original ECG signal and the reconstructed signal obtained using the proposed band-limited FISTA.

First, the required number of iterations decreases monotonically as the bandwidth limit is reduced. When the spectral support is restricted to 60 Hz, the algorithm converges in only 16 iterations, compared to 30 iterations required by the standard FISTA without band limitation. This behavior is a direct consequence of the reduction in the *effective optimization dimension*. By enforcing hard suppression of coefficients outside the predefined index set $\mathcal{B}$ at every iteration, the proposed proximal operator eliminates high-frequency descent directions that cannot contribute to the final solution. As a result, the optimization trajectory is confined to a lower-dimensional subspace, leading to faster convergence, reduced oscillatory behavior, and improved numerical stability.

In contrast, the reconstruction accuracy measured by PRDN exhibits a non-monotonic dependence on the bandwidth limit. As the bandwidth is reduced from 240 Hz to 90 Hz, PRDN consistently improves, indicating that suppressing high-frequency components effectively removes noise and mitigates overfitting in the sparse recovery process. However,

when the bandwidth is further restricted to 60 Hz, PRDN degrades despite the continued reduction in iteration count. This degradation arises because overly aggressive band limitation removes low-energy but structurally important spectral components that are necessary to accurately represent fine temporal features, such as sharp transitions and morphological details in the signal.

These results indicate that the optimal bandwidth depends critically on the target optimization objective. A tightly constrained bandwidth (e.g., 60 Hz) minimizes the iteration count and computational cost, making it well suited for latency- or energy-constrained embedded systems. Conversely, a moderately relaxed bandwidth (e.g., 90 Hz) achieves the best balance between convergence speed and reconstruction fidelity, yielding the lowest PRDN while still providing a substantial reduction in iteration count relative to the standard FISTA.

Importantly, this behavior demonstrates that the proposed band-limited proximal operator does not merely accelerate convergence, but fundamentally reshapes the optimization landscape. By explicitly constraining the feasible set in the transform domain, the operator introduces a controllable trade-off between computational efficiency and reconstruction accuracy. The bandwidth limit therefore emerges as a meaningful and tunable algorithmic design parameter, enabling system-level optimization according to available computational resources, latency constraints, and acceptable reconstruction error.

From a practical standpoint, the bandwidth limit can be guided by a simple quantitative criterion rather than purely empirical tuning. In particular, for accuracy-oriented operation, one may select the smallest bandwidth that captures a prescribed fraction of the signal energy in the Fourier domain (e.g., 95–99%), while remaining within the physiologically meaningful ECG frequency range, which is widely reported to be concentrated below approximately 40–60 Hz in the clinical and signal-processing literature. This energy-based heuristic provides a principled and data-independent guideline for bandwidth selection and naturally explains the non-monotonic PRDN behavior observed in Figure 4a, which arises from the trade-off between noise suppression and signal distortion due to excessive band limitation.

However, this energy-based criterion is primarily intended for accuracy-oriented bandwidth selection. In latency- or resource-constrained embedded scenarios, the operating bandwidth may instead be intentionally reduced to prioritize iteration-count reduction and computational efficiency. In this work, a bandwidth of 60 Hz was selected based on this system-level trade-off, as it significantly reduces the number of required FISTA iterations while still preserving the dominant physiological spectral content of monitoring-grade ECG signals. This trade-off can be conceptually interpreted through a composite objective that balances reconstruction accuracy and computational cost, although no explicit optimization over the bandwidth parameter is performed in this study.

Based on the observed PRDN trends and iteration behavior, the algorithmic parameters were selected to balance reconstruction accuracy and computational efficiency under embedded system constraints. Specifically, the number of compressed measurements was set to $M = 340$, corresponding to approximately 33% of the Nyquist-rate uniform sampling measurement count, while the number of iterations was 16 under the 60 Hz band-limited proximal operator. The regularization parameters were chosen as $\lambda = 0.0005$ and $L = 0.001$. Under these settings, the proposed method achieved a PRDN of 24.4%, as shown in Figure 4b. This operating point was therefore adopted in all subsequent experiments.

5.5. Parameter Sensitivity Analysis

Figure 5 illustrates the variation of PRDN with respect to key hyperparameters in the standard (non-band-limited) FISTA-based sparse recovery algorithm. Among the

evaluated parameters, the number of measurements M exhibits the most pronounced influence on PRDN.




Figure 5. PRDN variation with respect to the main hyperparameters, where (a) shows the effect of the number of measurements, (b) shows the effect of the number of iterations, and (c) shows the effect of λ .

Figure 5a illustrates that the reconstruction performance improves as the number of measurements M increases. In particular, when $M \geq 340$, the PRDN remains stably below 25%, indicating that a sufficient number of measurements enables high-quality signal reconstruction. Figure 5b shows the variation in PRDN with respect to the number of iterations in the FISTA algorithm. The PRDN converges at around 30 iterations, beyond which no significant improvement is observed. This behavior is attributed to the momentum-based update scheme of FISTA, which accelerates convergence by incorporating information from previous iterations. Consequently, the number of iterations can be selected to optimize computational efficiency. Figure 5c presents the variation in PRDN as a function of the regularization parameter λ . The performance improves up to a certain range ($\lambda \approx 0.0005$), beyond which the PRDN increases. This trend indicates that excessive enforcement of sparsity may lead to information loss. In this context, λ controls the trade-off between sparsity and data fidelity: smaller values emphasize fitting of the measured data, whereas larger values enforce stronger sparsity constraints. Selecting an appropriate λ is therefore critical to balancing reconstruction accuracy and sparsity.

6. Limitations and Future Work

While the proposed framework demonstrates the feasibility of FISTA-based sparse recovery on MCU-class hardware, several limitations should be acknowledged to clarify the scope of this study and motivate future research.

(1) Signal Conditions and Evaluation Scope

All experiments were conducted under controlled, low-noise laboratory conditions using a clean benchmark ECG signal. Therefore, the robustness of the proposed CS-FISTA pipeline against realistic disturbances such as motion artifacts, muscle noise, electrode-skin impedance variation, and electromagnetic interference was not evaluated. Future work should consider ambulatory and wearable ECG datasets to assess performance under practical noise conditions.

(2) Limited Signal Diversity

The experimental evaluation in this study focused on monitoring-grade ECG signals acquired from a normal adult subject, which exhibit relatively band-limited spectral characteristics. While this choice is consistent with the target application of rhythm monitoring, other ECG types, such as pediatric or pathological signals, may present different spectral distributions and sparsity patterns. Future work will therefore extend the proposed

framework to more diverse ECG databases and investigate its robustness under broader physiological variability.

(3) Algorithmic Scope

This work focused exclusively on FISTA to analyze its iteration cost, memory footprint, and execution characteristics on MCUs. No systematic comparison was performed with alternative sparse recovery methods (e.g., IHT, OMP/CoSaMP, or deep unrolled networks). Evaluating how different algorithms trade reconstruction accuracy, latency, and energy under embedded constraints remains an important direction for future work.

(4) Real-Time and System-Level Validation

The reconstruction was performed on frame-level segments to analyze computational cost per iteration, rather than within a continuous real-time acquisition pipeline. Future work should integrate interrupt-driven sampling, real-time buffering, and energy profiling to evaluate system-level performance under continuous operation.

Future Directions

Based on these limitations, several promising research directions can be identified:

- Warm-start and phase-aligned initialization to reduce iteration count and latency in continuous operation.
- Fixed-point or quantized FFT implementations to further improve performance on resource-limited MCUs.
- Validation under noisy, real-world ambulatory conditions.
- Extension to other compressible biosignals such as EEG, EMG, and PPG.
- Algorithm–hardware co-design to derive embedded-oriented sparse recovery strategies.

Overall, this study establishes a computational foundation demonstrating that convex sparse recovery is feasible on resource-constrained MCUs. Addressing noise robustness, physiological variability, and real-time system integration will be essential for advancing toward practical embedded sensing applications. From a broader research perspective, the contribution of this work is not the introduction of a new sparse recovery model, but a reinterpretation of band-limited prior information as an explicit design lever for reducing optimization dimension, iteration count, and computational cost under embedded hardware constraints. In this sense, the proposed approach complements existing structure-aware sparse recovery methods by translating their conceptual advantages into a concrete, MCU-feasible design methodology.

7. Conclusions

This work investigated the algorithmic and architectural feasibility of executing CS sparse recovery directly on resource-constrained MCU-class hardware. ECG signals acquired through an AD8232 front end were used solely as representative benchmark waveforms, with the objective of evaluating embedded reconstruction capability rather than clinical diagnostic performance.

Unlike prior embedded CS studies that restrict on-device processing to signal acquisition or lightweight compression, this work demonstrates that full ℓ_1 -regularized sparse recovery using FISTA—including forward and adjoint operators implemented via FFT/IFFT—can be executed entirely on a Cortex-M4 MCU without external computation offloading.

To improve convergence efficiency under strict computational constraints, a band-limited proximal operator was introduced to explicitly incorporate known spectral support. By suppressing non-informative frequency components at every iteration, the proposed operator effectively reduces the optimization dimensionality and reshapes the optimization landscape. Experimentally, this modification reduced the required iteration count from 30

to 16 (a 47% reduction) while preserving reconstruction fidelity, resulting in a near-linear reduction in total execution time.

Cycle-accurate profiling revealed that the dominant bottleneck arises from arithmetic-intensive FFT/IFFT operations, indicating that algorithm-level reductions in iteration count are more effective than memory placement optimizations.

Using 340 compressed measurements (approximately 33% of Nyquist-rate uniform sampling), the proposed implementation achieved a PRDN of 24.38%, closely matching MATLAB-based FISTA and CVX reference results. These results experimentally confirm that accurate ℓ_1 -minimization-based sparse recovery is feasible even under the tight computational and memory constraints of a single-precision MCU.

Overall, this study establishes a practical foundation for embedded CS reconstruction by demonstrating that optimization-aware algorithm design—rather than reliance on more powerful hardware—can enable full sparse recovery on low-power platforms. The proposed band-limited proximal FISTA provides a tunable algorithmic knob that links reconstruction accuracy, convergence speed, and computational cost, offering useful design insight for future energy-efficient embedded sensing systems.

Author Contributions: Conceptualization, S.C.; Methodology, S.C.; Software, S.C.; Validation, M.K.; Investigation, S.C. and D.P.; Data Curation, M.K.; Writing—Original Draft, S.C.; Writing—Review & Editing, D.P.; Supervision, D.P.; Project Administration, D.P.; Funding Acquisition, D.P. All authors have read and agreed to the published version of the manuscript.

Funding: This study was supported by the BK21 FOUR project (4199990113966), the Basic Science Research Program (RS-2018-NR031059, 10%) (RS-2025-24322979, 10%) through the National Research Foundation of Korea (NRF) funded by the Ministry of Education. This work was partly supported by Institute of Information and communications Technology Planning and Evaluation (IITP) grants funded by the Korean government (MSIT): (No. RS-2023-00228970, Development of Flexible SW-HW Conjunctive Solution for On-edge Self-supervised Learning, 30%) and (No. RS-2025-02218227, Digital Columbus Project, 40%) and (No. RS-2022-00156389, Innovative Human Resource Development for Local Intellectualization support program, 10%). The EDA tool was supported by the IC Design Education Center (IDEC), Republic of Korea.

Institutional Review Board Statement: The study was conducted in accordance with the Declaration of Helsinki, and approved by the Bioethics Review Committee of Kyungpook National University (protocol code 2026-0090).

Informed Consent Statement: Our approaches' feasibility was only evaluated by collecting sample signals from the body of our researcher. Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The original contributions presented in the study are included in the article; further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflict of interest. The funding sponsors had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1. Donoho, D.L. Compressed sensing. *IEEE Trans. Inf. Theory* **2006**, *52*, 1289–1306.
2. Candès, E.J.; Romberg, J.; Tao, T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory* **2006**, *52*, 489–509.
3. Daubechies, I.; Defrise, M.; De Mol, C. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Commun. Pure Appl. Math.* **2004**, *57*, 1413–1457.
4. Beck, A.; Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.* **2009**, *2*, 183–202.

5. Blumensath, T.; Davies, M.E. Iterative hard thresholding for compressed sensing. *Appl. Comput. Harmon. Anal.* **2009**, *27*, 265–274.
6. STMicroelectronics. *STM32F4DISCOVERY: Discovery Kit with STM32F407VG MCU*; User Manual UM1472; 2014. Available online: https://www.st.com/resource/en/user_manual/um1472-discovery-kit-with-stm32f407vg-mcu-stmicroelectronics.pdf (accessed on 20 June 2026).
7. Arm Ltd. *CMSIS-DSP: Digital Signal Processing Library for Arm Cortex-M Processors*; Version 1.14.0, Documentation; 2023. Available online: <https://arm-software.github.io/CMSIS-DSP/v1.14.0/modules.html> (accessed on 20 June 2026).
8. Analog Devices. *AD8232: Single-Lead, Heart Rate Monitor Front End*; Datasheet, Rev. C; 2012. Available online: <https://www.analog.com/media/en/technical-documentation/data-sheets/ad8232.pdf> (accessed on 20 June 2026).
9. Candès, E.J.; Romberg, J. Sparsity and incoherence in compressive sampling. *Inverse Probl.* **2007**, *23*, 969–985.
10. Candès, E.J.; Tao, T. Decoding by linear programming. *IEEE Trans. Inf. Theory* **2005**, *51*, 4203–4215.
11. Donoho, D.L.; Huo, X. Uncertainty principles and ideal atomic decomposition. *IEEE Trans. Inf. Theory* **2001**, *47*, 2845–2862.
12. Chen, S.S.; Donoho, D.L.; Saunders, M.A. Atomic Decomposition by Basis Pursuit. *SIAM Rev.* **2001**, *43*, 129–159.
13. Boyd, S.; Parikh, N.; Chu, E.; Peleato, B.; Eckstein, J. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found. Trends Mach. Learn.* **2011**, *3*, 1–122.
14. Fira, M.; Costin, H.N.; Goraş, L. A study on dictionary selection in compressive sensing for ECG signals compression and classification. *Biosensors* **2022**, *12*, 146.
15. Tropp, J.A.; Gilbert, A.C. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Trans. Inf. Theory* **2007**, *53*, 4655–4666.
16. Mamaghanian, H.; Khaled, N.; Atienza, D.; Vanderghenst, P. Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes. *IEEE Trans. Biomed. Eng.* **2011**, *58*, 2456–2466.
17. Duarte, M.F.; Davenport, M.A.; Takhar, D.; Laska, J.N.; Sun, T.; Kelly, K.F.; Baraniuk, R.G. Single-pixel imaging via compressive sampling. *IEEE Signal Process. Mag.* **2008**, *25*, 83–91.
18. Kumar, S.S.; Ramachandran, P. Review on compressive sensing algorithms for ECG signal for IoT based deep learning framework. *Appl. Sci.* **2022**, *12*, 8368.
19. Zhang, Z.; Jung, T.P.; Makeig, S.; Rao, B.D. An FPGA-based implementation of compressed sensing for biomedical signals. *IEEE Trans. Biomed. Circuits Syst.* **2015**, *9*, 545–556.
20. Davenport, M.A.; Duarte, M.F.; Eldar, Y.C.; Kutyniok, G. Signal processing with compressive measurements. *IEEE J. Sel. Top. Signal Process.* **2010**, *4*, 445–460.
21. Bellarmino, N.; Cantoro, R.; Fosson, S.M.; Huch, M.; Kilian, T.; Schlichtmann, U.; Squillero, G. COSMO: COMpressed Sensing for Models and Logging Optimization in MCU Performance Screening. *IEEE Trans. Comput.* **2024**, *74*, 652–664.
22. Daponte, P.; De Vito, L.; Iadarola, G.; Picariello, F. ECG monitoring based on dynamic compressed sensing of multi-lead signals. *Sensors* **2021**, *21*, 7003.
23. Safarpour, M.; Hautala, I.; Silvén, O. An embedded programmable processor for compressive sensing applications. In *Proceedings of the 2018 IEEE Nordic Circuits and Systems Conference (NORCAS): NORCHIP and International Symposium of System-on-Chip (SoC)*; IEEE: New York, NY, USA, 2018; pp. 1–5.
24. Ahmmed, P.; Garceau, E.; Latif, T.; Brewer, A.; Dieffenderfer, J.; Valero-Sarmiento, J.M.; Pamula, V.R.; Van Helleputte, N.; Van Hoof, C.; Verhelst, M.; et al. Preclinical Evaluation of a Wearable Wristband With Compressed-Sensing Based Photoplethysmography. *IEEE Trans. Biomed. Eng.* **2025**, *72*, 1596–1604.
25. De Vito, L.; Picariello, E.; Picariello, F.; Rapuano, S.; Tudosa, I. A dictionary optimization method for reconstruction of ECG signals after compressed sensing. *Sensors* **2021**, *21*, 5282.
26. Hassan, A.M.A.; Mohsen, S. Compression of electrocardiogram signals using compressive sensing technique based on curvelet transform toward medical applications. *Multimed. Tools Appl.* **2025**, *84*, 11203–11219. <https://doi.org/10.1007/s11042-024-19328-z>.
27. Clifford, G.D.; Azuaje, F.; McSharry, P.E. *Advanced Methods and Tools for ECG Data Analysis*; Artech House: Norwood, MA, USA, 2006.
28. Kligfield, P.; Gettes, L.S.; Bailey, J.J.; Childers, R.; Deal, B.J.; Hancock, E.W.; van Herpen, G.; Kors, J.A.; Macfarlane, P.; Mirvis, D.M.; et al. Recommendations for the Standardization and Interpretation of the Electrocardiogram. *Circulation* **2007**, *115*, 1306–1324. <https://doi.org/10.1161/CIRCULATIONAHA.106.180200>.
29. Tsiliogianni, E.; Kondi, L.P.; Katsaggelos, A.K. Use of tight frames for optimized compressed sensing. In *Proceedings of the 2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO)*; IEEE: New York, NY, USA, 2012; pp. 1439–1443.

30. Lu, Z.; Kim, D.Y.; Pearlman, W.A. Wavelet compression of ECG signals by the SPIHT algorithm. *IEEE Trans. Biomed. Eng.* **2000**, *47*, 849–856.
31. Němcová, A.; Smusek, R.; Maršánová, L.; Smital, L. A Comparative Analysis of Methods for Evaluation of ECG Signal Quality after Compression. *BioMed Res. Int.* **2018**, *2018*, 1868519. <https://doi.org/10.1155/2018/1868519>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.