

Article

The Perceptual Game-Theoretic Transform (PGTT): Axiomatizing Signal Compression via the Fourier–Shapley Isomorphism

Adnan H. Abdulwahid 

Mathematics Department, College of Arts, Sciences, and Education (CASE), Texas A&M University–Texarkana, Texarkana, TX 75503, USA; aabdulwahid@tamut.edu

Abstract

The mathematical representation of discrete signals is classically governed by linear basis transforms, such as the Discrete Fourier Transform (DFT), which treat spectral projection as a rigid, deterministic geometric operation. Under this paradigm, signal compression and thresholding rely on heuristic error metrics like the Minimum Mean Squared Error (MMSE). To mathematically axiomatize these fundamental operations, this paper reformulates computational basis transforms through the lens of Cooperative Game Theory. By defining discrete signal reconstruction as a Grand Coalition of orthogonal frequency players, we establish a strict mathematical isomorphism between functional analysis and cooperative game theory. We prove that the spectral energy assigned to each frequency coefficient is exactly its Shapley Value and that classical MMSE minimization is mathematically equivalent to maximizing the retained Shapley payout. Furthermore, we extend this framework to physical hardware and linear shift-invariant (LSI) systems, modeling 8-bit quantization and the Modulation Transfer Function (MTF) as sub-additive “economic taxes” on the coalition. Finally, by intentionally violating the Shapley Symmetry Axiom to mimic the Contrast Sensitivity Function (CSF) of the human visual system, we propose a fundamentally new mathematical basis: the Perceptual Game-Theoretic Transform (PGTT). Unlike classical methods that rely on post hoc quantization for signal compression, the PGTT acts as an inherently efficient transform that structurally guarantees sub-Nyquist computational complexity and dynamic range reallocation prior to physical hardware saturation.

Keywords: discrete Fourier Transform; cooperative game theory; Shapley Value; Fourier–Shapley Isomorphism; Minimum Mean Squared Error (MMSE); signal compression; Modulation Transfer Function (MTF); human visual system; Perceptual Game-Theoretic Transform (PGTT)



Academic Editor: Geert Verdoolaege

Received: 2 June 2026

Revised: 13 July 2026

Accepted: 29 July 2026

Published: 3 August 2026

Copyright: © 2026 by the author.

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.

1. Introduction

The mathematical representation of discrete signals traditionally relies on linear basis transforms, most notably the Discrete Fourier Transform (DFT). In classical signal processing, projecting a spatial signal onto an orthogonal basis of complex exponentials is treated strictly as a deterministic geometric operation [1,2]. Under this paradigm, the energy assigned to a specific spectral coefficient is viewed merely as the physical magnitude of that frequency’s presence within the spatial domain.

Historically, the mitigation of high-frequency noise and structural recovery within this deterministic framework has evolved through several distinct paradigms, ranging from foundational statistical enhancements [3–5] to Gibbs-based stochastic relaxation [6]. Early linear filtering techniques laid the mathematical groundwork for stationary signal

smoothing [7–9], while subsequent non-linear and variational methods introduced spatial constraints such as anisotropic diffusion and total variation [10,11]. This naturally extended into multi-channel and vectorial directional filtering to preserve complex geometries across color spaces [12–16]. In the transform domain, classical sparsity-driven approaches popularized soft-thresholding, localized phase reconstruction, and highly advanced collaborative filtering architectures such as BM3D and non-local means [17–22], whose efficacy is traditionally evaluated via localized error and structural similarity metrics [23]. Furthermore, in applied physical systems, robust restoration requires explicit modeling of hardware sensor boundaries and continuous radiometric calibration [24,25].

Despite these advancements, most modern recovery methods—including frameworks grounded in convex optimization and variational calculus [26–28]—still rely on rigid parameterizations. Recent literature has demonstrated the profound utility of game theory [29–31] for stabilizing these inverse problems [32]. By mapping regularized deconvolution to continuous zero-sum duels, Stackelberg market equilibriums, and non-local patch matching to player-specific congestion and potential networks [33–38], these frameworks effectively eliminate heuristic tuning. However, these non-cooperative topologies still treat the underlying mathematical transform (e.g., the DFT) as a rigid, physical constant—an immutable arena where computational players compete.

This paper proposes a theoretical reframing: examining the mathematical basis transform itself through the lens of *Cooperative Game Theory* [39]. Rather than treating orthogonal projection as a static calculation, we formulate the reconstruction of a discrete spatial signal as a Grand Coalition of cooperating computational frequency components.

In cooperative game theory, the Shapley Value [40] provides the unique, mathematically axiomatic method for fairly distributing the total payoff of a coalition among its individual contributors, satisfying the rigorous properties of efficiency, symmetry, dummy player, and additivity. Concurrently, in functional analysis, Parseval’s Theorem dictates that the total spatial energy of a signal is conserved and distributed across its spectral frequencies.

Drawing upon principles of discrete matrix algebra and cooperative topology, this paper establishes a strict mathematical isomorphism between these two domains. We formally prove that the spectral energy assigned to each orthogonal frequency coefficient is not merely a geometric projection but exactly the Shapley Value representing that frequency’s marginal contribution to the spatial signal’s reconstruction.

To establish this framework, it is crucial to clearly delineate the existing principles across three separate domains before detailing their synthesis. First, in classical signal processing, the DFT strictly governs orthogonal spectral projection. Second, in cooperative game theory, the Shapley Value provides the rigorous mathematical standard for equitable utility distribution. Third, in biological perception modeling, the Contrast Sensitivity Function (CSF) dictates the non-linear band-pass nature of human vision. Only by isolating these existing principles can we identify the mathematical intersections necessary to synthesize our proposed transform.

The primary contributions of this paper are threefold:

1. **The Computational Spatial–Spectral Coalition:** We formulate a cooperative game topology using unitary matrices, where orthogonal basis vectors act as rational players. We prove strict additivity in the latent mathematical space and linear scalability for high-dimensional 3D volumetric data (such as MRI or CT medical imaging tensors, which suffer heavily from classical transform inefficiencies).
2. **The Fourier–Shapley and MMSE Isomorphisms:** We provide a rigorous mathematical proof that the standard Discrete Fourier Transform satisfies all four Shapley axioms.

Furthermore, we prove that classical MMSE compression is mathematically identical to a Shapley utility deficit, subject to a hardware “Quantization Tax.”

3. **Optical and Perceptual Economic Taxation:** We model physical LSI systems (via the MTF) and the biological observer (via the CSF) as taxation authorities that scale the coalition’s wealth. By breaking the Shapley Symmetry axiom, we introduce the foundation for a novel Perceptual Game-Theoretic Transform (PGTT).

The remainder of this paper is organized as follows. Section 2 formally defines the unitary matrix formulation of the Spatial–Spectral Coalition and physical hardware boundaries. Section 3 details the mathematical proof of the Fourier–Shapley Isomorphism and its equivalence to the MMSE. Section 4 explores the economic taxation imposed by physical optics and the human visual system. Section 5 introduces the PGTT. Section 6 presents the quantitative empirical validations and benchmark comparisons, and Section 7 concludes the paper and discusses future algorithmic extensions.

2. Background: The Spatial–Spectral Coalition Matrix Formulation

To bridge the domains of computational signal processing and cooperative game theory, we must establish a rigorous mathematical isomorphism between the discrete components of a digital signal and the abstract entities of a cooperative game. In this section, we abandon continuous functional analysis in favor of discrete matrix algebra, formulating the Discrete Fourier Transform (DFT) not as a deterministic linear operator but as a *Grand Coalition* of rational computational agents.

2.1. The Players: The Unitary Basis Matrix

Let $x \in \mathbb{R}^N$ represent a discrete spatial signal, bounded by the physical dynamic range of a hardware sensor array, denoted $\mathcal{I} = [0, L_{max}]$ (e.g., $L_{max} = 255$ for 8-bit systems). In standard computational harmonic analysis, the 1D DFT is represented by the symmetric matrix $W_N \in \mathbb{C}^{N \times N}$.

To guarantee game-theoretic additivity, we enforce a strictly *unitary* normalization of the DFT. Let W_N^H denote the Hermitian (conjugate) transpose of this matrix. Under unitary scaling ($1/\sqrt{N}$), the matrix is perfectly orthogonal such that $W_N W_N^H = I$. Consequently, the Hermitian transpose W_N^H functions exactly as the inverse DFT (the spatial synthesis operator).

We define the set of all available discrete frequency indices as the Grand Coalition of players $\Omega = \{0, 1, \dots, N-1\}$. It is important to note that a player does not represent a continuous range of frequencies; rather, each player k is assigned exactly one discrete, unique orthogonal frequency bin from the N -point spectrum. The specific “strategy” or contribution of player $k \in \Omega$ is entirely defined by w_k , the k -th orthonormal column vector of the synthesis matrix W_N^H :

$$w_k[n] = \frac{1}{\sqrt{N}} \exp\left(j \frac{2\pi}{N} kn\right). \quad (1)$$

Under this computational topology, the target spatial signal x is the physical manifestation of the Grand Coalition Ω , where all N players cooperate perfectly. In this context, “perfect cooperation” means that the mathematical summation of all N orthogonal strategies yields a lossless reconstruction of the original signal with exactly zero error. If players are evicted or behave differently (e.g., via hard thresholding or quantization), the coalition fails to achieve this perfect summation, resulting in lossy signal compression:

$$x = \sum_{k \in \Omega} X[k] w_k = W_N^H X, \quad (2)$$

where $X \in \mathbb{C}^N$ is the spectral vector, and $X[k]$ is the scalar weight of player k 's contribution.

2.2. The Characteristic Value Function

In cooperative game theory, the behavior of a game is governed by its characteristic value function $v(S)$, which maps every possible sub-coalition of players $S \subseteq \Omega$ to a real number $v(S) \in \mathbb{R}$.

We define the worth of a frequency coalition S as the total ideal spatial energy (variance) it can autonomously reconstruct. We explicitly use the term “ideal” (or mathematically latent) to denote the unbounded, continuous mathematical energy prior to any physical hardware clipping or biological observation. Let X_S be a filtered (zero-padded) spectral vector where $X_S[k] = X[k]$ if $k \in S$, and 0 otherwise. This zero-padding does not imply disguise; it is the standard mathematical operation for explicitly executing the eviction of non-coalition players without altering the dimensions of the computational arena. The partial signal synthesized exclusively by coalition S is $x_S = W_N^H X_S$.

We define the characteristic value function $v(S)$ as the squared Euclidean L_2 norm of this partial matrix reconstruction:

$$v(S) = \|x_S\|_2^2 = x_S^H x_S. \quad (3)$$

Theorem 1 (Strict Additivity of the Unitary Coalition). *Let Ω be the Grand Coalition defined by the unitary matrix W_N . The characteristic value function $v(S) = x_S^H x_S$ defines a strictly additive (inessential) cooperative game. For any two disjoint coalitions $S, T \subset \Omega$ where $S \cap T = \emptyset$:*

$$v(S \cup T) = v(S) + v(T). \quad (4)$$

Proof. Let S and T be disjoint subsets of Ω . Because each player k is strictly assigned to a unique orthogonal frequency bin, disjoint coalitions represent completely non-overlapping frequency bands (e.g., a set of low frequencies versus a set of high frequencies). Consequently, a player cannot simultaneously exist in both subsets, guaranteeing that no part of the signal spectrum is repeated or duplicated between the two groups. The partial signal reconstructed by their union is $x_{S \cup T} = x_S + x_T$. Evaluating the characteristic function yields:

$$v(S \cup T) = (x_S + x_T)^H (x_S + x_T). \quad (5)$$

Expanding the matrix multiplication:

$$v(S \cup T) = x_S^H x_S + x_T^H x_T + x_S^H x_T + x_T^H x_S. \quad (6)$$

Substitute $x_S = W_N^H X_S$ and $x_T = W_N^H X_T$ into the cross-terms:

$$x_S^H x_T = (W_N^H X_S)^H (W_N^H X_T) = X_S^H W_N W_N^H X_T. \quad (7)$$

Because W_N is a unitary matrix, $W_N W_N^H = I$ (the identity matrix). Thus, $x_S^H x_T = X_S^H X_T$. Since S and T are strictly disjoint, the non-zero elements of X_S and X_T never overlap, guaranteeing $X_S^H X_T = 0$. Therefore:

$$v(S \cup T) = x_S^H x_S + x_T^H x_T = v(S) + v(T). \quad (8)$$

This proves that the Fourier coalition is perfectly additive in the unbounded latent space. $\square$

2.3. Physical Hardware Boundaries: The Quantization Tax

While Theorem 1 guarantees perfect cooperation in continuous computational space, physical hardware strictly enforces dynamic range limits.

Example 1 (The 8-Bit Saturation Penalty). Let $\mathcal{C}(\cdot)$ represent a physical hardware clipping operator bounding signals to $\mathcal{I} = [0, 255]$. When two disjoint frequency coalitions S and T combine, their latent mathematical energy is strictly additive. However, if their constructive interference at a specific spatial matrix index exceeds 255, the physical hardware clips the observation.

The physical characteristic function becomes $\tilde{v}(S \cup T) = \|\mathcal{C}(x_S + x_T)\|_2^2$. Because clipped energy is strictly less than latent energy, the hardware forces a sub-additive penalty: $\tilde{v}(S \cup T) < v(S) + v(T)$. The hardware physically taxes the high-frequency players.

Figure 1 visually simulates this proof using a discrete 1D matrix signal, proving the divergence between latent mathematical additivity and physical hardware saturation.




Figure 1. Characteristic Value Function $v(S)$ of the Spatial-Spectral Coalition. **(Top)** Matrix reconstruction as computational players are added to the coalition. The red area denotes the physical hardware boundary limits (0 and 255). **(Bottom)** The latent mathematical energy (blue) exhibits strict linear additivity, reflecting Theorem 1. The physical hardware energy (red) becomes sub-additive, demonstrating the saturation penalty forced by the 8-bit sensor boundary.

2.4. Tensor Separability and 3D Volumetric Scalability

Calculating exact Shapley values for general cooperative games requires $O(2^N)$ complexity, rendering it computationally impossible for 3D medical tensors (e.g., MRI volumes where $N > 10^7$ voxels). However, the matrix formulation allows us to mathematically guarantee $O(N \log N)$ linear scalability for N -dimensional datasets.

Theorem 2 (Volumetric Coalition Separability). The 3D Spatial-Spectral Coalition strictly decomposes into three independent 1D sub-games.

Proof. The 3D DFT operator for an $N \times M \times P$ spatial tensor is defined algebraically by the Kronecker product of three 1D DFT matrices:

$$W_{3D} = W_N \otimes W_M \otimes W_P. \quad (9)$$

A fundamental property of matrix algebra states that the Kronecker product of unitary matrices is strictly unitary. Therefore:

$$\begin{aligned} W_{3D} W_{3D}^H &= (W_N W_N^H) \otimes (W_M W_M^H) \otimes (W_P W_P^H) \\ &= I \otimes I \otimes I \\ &= I_{3D}. \end{aligned} \quad (10)$$

Because the global 3D transform matrix W_{3D} remains perfectly unitary, Theorem 1 holds universally across N -dimensions. The global 3D characteristic value function $v_{3D}(S)$ is strictly additive, allowing the computational complexity of the 3D Shapley calculation to completely decouple into independent 1D row, column, and depth coalitions. $\square$

3. The Fourier–Shapley Isomorphism

In the classical formulation of cooperative game theory, the Shapley Value [40] provides a unique, mathematically axiomatic solution for distributing the total payoff of a Grand Coalition among its constituents. It calculates the equitable allocation by averaging a player's marginal contribution over all possible permutation orderings of coalition formation.

In this section, we formally define the Shapley Value within the context of the Spatial–Spectral Coalition, prove that the matrix formulation of the Discrete Fourier Transform fundamentally behaves as a Shapley allocation mechanism, and establish its direct equivalence to the Minimum Mean Squared Error (MMSE).

3.1. The Shapley Value of a Frequency Player

Let $\phi_k(v)$ denote the Shapley Value representing the fair energy allocation to frequency player $k \in \Omega$ under the characteristic value function v . The general discrete formula is given by:

$$\phi_k(v) = \sum_{S \subseteq \Omega \setminus \{k\}} \frac{|S|!(N - |S| - 1)!}{N!} [v(S \cup \{k\}) - v(S)] \quad (11)$$

where S is a sub-coalition not containing k , $|S|$ is the cardinality of that coalition, and N is the total number of frequency players. The term $[v(S \cup \{k\}) - v(S)]$ represents the strict marginal contribution of player k to the coalition S .

Because computing Equation (11) requires iterating over 2^N subsets, calculating exact Shapley Values is famously NP-hard for general cooperative games. However, the topological constraints of the unitary Spatial–Spectral Coalition algebraically collapse this complexity.

Theorem 3 (The Fourier–Shapley Isomorphism). *Let Ω be the Grand Coalition of orthogonal column vectors from the inverse DFT matrix W_N^H reconstructing a discrete signal $x \in \mathbb{C}^N$. For the characteristic value function $v(S) = \|x_S\|_2^2$, the Shapley Value $\phi_k(v)$ allocated to any frequency player k is exactly equal to the squared magnitude of its Fourier spectral coefficient:*

$$\phi_k(v) = |X[k]|^2. \quad (12)$$

Proof. By Theorem 1 (Equation (4)), the characteristic value function of the latent matrix coalition is strictly additive due to the orthogonality of the unitary basis. Therefore, for any sub-coalition $S \subseteq \Omega \setminus \{k\}$, the worth of the coalition after player k joins is:

$$v(S \cup \{k\}) = v(S) + v(\{k\}) = v(S) + \|X[k]w_k\|_2^2. \quad (13)$$

Because the matrix column vectors w_k are orthonormal (their L_2 norm is exactly 1), this simplifies to:

$$v(S \cup \{k\}) - v(S) = |X[k]|^2. \quad (14)$$

Crucially, this proves that the marginal contribution of player k is an absolute constant, completely independent of the existing coalition S . Because the unitary basis vectors are strictly orthogonal, each frequency player operates in complete mathematical isolation. In this latent space, no player can constructively or destructively interfere with another's marginal contribution (aligning with the classical Parseval energy conservation). Substituting this independent marginal contribution back into the Shapley formula (Equation (11)) yields:

$$\phi_k(v) = |X[k]|^2 \sum_{S \subseteq \Omega \setminus \{k\}} \frac{|S|!(N - |S| - 1)!}{N!}. \quad (15)$$

By the fundamental combinatorial properties of permutations, the sum of these combinatorial weights over all possible subsets S evaluates identically to 1. Therefore:

$$\phi_k(v) = |X[k]|^2. \quad (16)$$

□

3.2. The Spectral MMSE-Shapley Equivalence

In classical estimation theory, the efficacy of a signal representation is measured by the Minimum Mean Squared Error (MMSE) between the true signal x and the estimated signal $\hat{x}_S$ reconstructed from a subset of frequencies S . We now prove that this foundational computational metric is mathematically equivalent to the game-theoretic payoff distribution.

Theorem 4 (MMSE-Shapley Equivalence). *Let $S \subset \Omega$ be an active coalition of retained frequencies, and let the estimated signal be $\hat{x}_S = \sum_{k \in S} X[k]w_k$. The Minimum Mean Squared Error $J(S) = \|x - \hat{x}_S\|_2^2$ is exactly equal to the sum of the Shapley Values of the evicted players (the complement coalition $\Omega \setminus S$):*

$$J(S) = v(\Omega) - \sum_{k \in S} \phi_k(v) = \sum_{m \notin S} \phi_m(v). \quad (17)$$

Proof. The total spatial signal is $x = \sum_{k \in \Omega} X[k]w_k$. The error vector $\epsilon = x - \hat{x}_S$ isolates the frequencies absent from the computational coalition:

$$\epsilon = \sum_{m \notin S} X[m]w_m. \quad (18)$$

Applying the squared L_2 norm and invoking the orthogonality of the unitary columns w_m :

$$J(S) = \|\epsilon\|_2^2 = \sum_{m \notin S} |X[m]|^2. \quad (19)$$

By Theorem 3, $|X[m]|^2 = \phi_m(v)$. Furthermore, by the Efficiency axiom of the Shapley Value (Parseval's Theorem in discrete form), the total worth $v(\Omega) = \sum_{k \in \Omega} \phi_k(v)$. Therefore:

$$J(S) = v(\Omega) - \sum_{k \in S} \phi_k(v). \quad (20)$$

This establishes that minimizing the spatial matrix reconstruction error is mathematically isomorphic to maximizing the retained Shapley Value of the frequency coalition. $\square$

3.3. Theoretical Significance: Axiomatizing Signal Compression

The Fourier–Shapley Isomorphism fundamentally reframes the philosophy of discrete signal processing. Historically, techniques such as hard-thresholding in JPEG compression or best- K -term approximation are viewed merely as greedy heuristic algorithms that discard low-energy frequencies to save bandwidth.

By proving Theorem 4, we establish that these classical compression pipelines are not just heuristically efficient; they are executing a mathematically perfect, axiomatic payout. Discarding a low-magnitude frequency is isomorphic to evicting a player with a negligible Shapley Value. The classical MMSE metric is thus revealed to be a direct measure of *coalition deficit*—the exact game-theoretic utility lost by excluding a subset of computational players from the Grand Coalition.

3.4. Hardware Boundaries: The Quantization Tax

While the MMSE-Shapley Equivalence holds perfectly in the unbounded mathematical domain of floating-point precision, physical signal processors are constrained by strict hardware quantization.

Example 2 (8-Bit Quantization Tax on the Coalition). Let $\mathcal{Q}(\cdot)$ represent a physical 8-bit quantization operator, mapping continuous spatial intensities to discrete integers $\mathcal{I} = \{0, 1, \dots, 255\}$. When a coalition S attempts to reconstruct the signal on physical hardware, the physical estimate becomes $\tilde{x}_S = \mathcal{Q}(\hat{x}_S)$.

The hardware quantization error introduces a non-linear noise floor that destroys the strict orthogonality of the unitary reconstruction. The physical MMSE, denoted $\tilde{J}(S) = \|x - \tilde{x}_S\|_2^2$, becomes strictly bounded below:

$$\tilde{J}(S) \geq v(\Omega) - \sum_{k \in S} \phi_k(v). \quad (21)$$

The divergence between the latent Shapley MMSE and the physical hardware MMSE represents a “Quantization Tax” imposed on the coalition. Crucially, the distribution of this economic penalty is heavily asymmetric; it does not act as a uniform fine across the Grand Coalition, but rather manifests as a targeted suppression that disproportionately punishes the marginal utility of high-frequency players.

To physically visualize this theorem, Figure 2 simulates the MMSE decay as players (frequencies) are progressively added to the coalition. While the latent mathematical error decays exactly according to the accumulated Shapley Values, the physical hardware error plateaus, graphically proving the sub-additive tax imposed by 8-bit physical constraints.




Figure 2. The Spectral MMSE-Shapley Equivalence under Hardware Boundaries. The horizontal axis represents the progressive inclusion of frequency players into the coalition S , ordered by descending Shapley Value. The latent MMSE (blue) decays perfectly to zero, governed precisely by $J(S) = v(\Omega) - \sum \phi_k(v)$. Conversely, the physical 8-bit hardware MMSE (red) suffers from a non-linear Quantization Tax, preventing the Grand Coalition from achieving zero error.

4. Case Study: Optical Systems as Economic Mediators

To demonstrate the practical application of this framework, we present a case study involving physical hardware. In Section 3, we established that the latent spatial signal guarantees a perfectly equitable distribution of variance among the frequency coalition. However, in applied machine vision and medical imaging, the latent signal is never directly observed. It must first pass through a physical Linear Shift-Invariant (LSI) system, such as a camera lens or an MRI acquisition coil.

In this section, we formulate the physical optics of LSI systems as “Economic Mediators” that dynamically tax the Shapley allocations of the frequency coalition. Furthermore, by analyzing how human perception naturally violates these economic distributions, we lay the mathematical groundwork for a fundamentally new, non-Fourier transform.

4.1. The OTF and MTF as Taxation Authorities

Let the physical LSI observation process be modeled by the matrix convolution $y = Hx + \eta$, where $H \in \mathbb{R}^{N \times N}$ is a circulant blur matrix representing the spatial point spread function (PSF), and $\eta \in \mathbb{R}^N$ is the stochastic sensor noise vector. We denote the discrete spectral representation of this noise as $\mathcal{F}(\eta) = W_N^H \eta$.

Because H is circulant, it is perfectly diagonalized by the DFT matrix W_N . The eigenvalues of H represent the Optical Transfer Function (OTF), and their absolute magnitudes represent the Modulation Transfer Function (MTF).

If we calculate the Shapley Value of a frequency player k based on the *observed* characteristic value function $\tilde{v}(S) = \|y_S\|_2^2$, the physical optics act as an economic taxation authority that scales the latent wealth of the coalition.

Theorem 5 (The MTF Taxation Theorem). *Let H be an LSI system with a strictly defined MTF vector, where $MTF[k] \in [0, 1]$ represents the signal attenuation at frequency k . Under a zero-mean*

noise floor, the expected physical Shapley Value $\mathbb{E}[\phi_k(\tilde{v})]$ allocated to player k is exactly the latent Shapley Value $\phi_k(v)$ taxed quadratically by the MTF, plus the expected variance of the noise:

$$\mathbb{E}[\phi_k(\tilde{v})] = \text{MTF}^2[k] \cdot \phi_k(v) + \mathbb{E}[|\mathcal{F}(\eta)[k]|^2] \quad (22)$$

where $\mathbb{E}[|\mathcal{F}(\eta)[k]|^2]$ is the expected spectral variance of the stochastic sensor noise at frequency k .

Proof. The observed spatial matrix is $y = Hx + \eta$. In the computational spectral domain, applying the unitary matrix W_N yields $Y = \text{OTF} \odot X + \mathcal{F}(\eta)$, where $\odot$ is the element-wise Hadamard product.

By Theorem 3, the observed Shapley Value is the squared magnitude of the observed spectral coefficient:

$$\phi_k(\tilde{v}) = |Y[k]|^2 = |\text{OTF}[k] \cdot X[k] + \mathcal{F}(\eta)[k]|^2. \quad (23)$$

Assuming the latent structural signal X and the stochastic noise $\mathcal{F}(\eta)$ are statistically independent, applying the expectation operator $\mathbb{E}[\cdot]$ causes their cross-correlation terms to vanish. Because the MTF is defined as the absolute magnitude of the OTF ($|\text{OTF}[k]| = \text{MTF}[k]$), taking the expectation algebraically collapses the equation to:

$$\mathbb{E}[\phi_k(\tilde{v})] = \text{MTF}^2[k] \cdot |X[k]|^2 + \mathbb{E}[|\mathcal{F}(\eta)[k]|^2]. \quad (24)$$

Substituting the latent Shapley identity $|X[k]|^2 = \phi_k(v)$ yields the result in Equation (22), proving that physical optics enforce a strict quadratic tax on the expected game-theoretic utility of the signal, offset only by the expected variance of the noise floor. Within this cooperative topology, stochastic measurement error and hardware limitations do not alter the composition of the Grand Coalition, nor do they facilitate communication between the strictly orthogonal players; rather, they strictly manifest as a distortion of the observed marginal contributions of the existing coalition members. $\square$

4.2. Axiomatic Violation: Breaking the Symmetry Axiom

The classical DFT is mathematically beautiful because it perfectly satisfies the Shapley **Symmetry Axiom**: if a low-frequency player and a high-frequency player contribute identical energy to the spatial variance, they receive identical Shapley allocations.

However, this strict mathematical equity contradicts the biological reality of the Human Visual System (HVS). The human eye acts as a non-linear band-pass filter, governed by the Contrast Sensitivity Function (CSF). To human perception, high-frequency spatial noise is far less “valuable” than mid-frequency structural geometry, even if their objective mathematical variance is identical.

By intentionally violating the Symmetry Axiom, we can define a completely new mathematical representation optimized for perceptual equity rather than energy conservation.

Definition 1 (The Perceptual Game-Theoretic Transform Foundation). *Let $w_{HVS} \in \mathbb{R}^N$ be a continuous weight vector derived from the Human Visual System’s Contrast Sensitivity Function. We define a **Weighted Shapley Value** where the marginal contribution of a player is scaled by their perceptual importance.*

*The optimal basis vectors $\Psi \in \mathbb{C}^{N \times N}$ that strictly diagonalize this perceptually weighted characteristic function do not equate to the standard complex exponentials of the DFT. Rather, they establish the foundation for a novel **Perceptual Game-Theoretic Transform**—a mathematical basis inherently pre-warped by the biological constraints of the observer, rendering post hoc perceptual weighting mathematically obsolete.*

Figure 3 illustrates the dual taxation mechanisms operating on the Spatial–Spectral Coalition. The physical hardware imposes an MTF tax that penalizes high frequencies, while the biological observer imposes a CSF tax that penalizes both absolute extremes, shifting the game-theoretic “wealth” heavily into the structural mid-band.




Figure 3. Economic Taxation of the Spatial–Spectral Coalition. (Blue) The latent Shapley Values representing the true signal variance. (Red) The MTF Taxation imposed by the physical optical system, exponentially degrading the high-frequency payoffs. (Green) The Perceptual CSF Taxation imposed by the human visual system, intentionally violating the Symmetry Axiom to heavily reward mid-frequency structural players.

5. The Perceptual Game-Theoretic Transform (PGTT)

The classical Fast Fourier Transform (FFT) computes the spectral projection for the entire Grand Coalition in $O(N \log N)$ time. In standard compression pipelines (e.g., JPEG or MPEG), the algorithm blindly computes the exact marginal contribution of all N players, only to immediately discard the high-frequency players via quantization tables. This post hoc eviction is computationally deeply inefficient.

By applying the CSF taxation model derived in Section 4, we can design a fundamentally new transform that pre-emptively filters the computational arena, achieving sub-Nyquist efficiency while maintaining strict perceptual invertibility. We formally define this as the **Perceptual Game-Theoretic Transform (PGTT)**.

5.1. The Active Coalition and Subspace Invertibility

Let $w_{HVS} \in (0, 1]^N$ be the strict Contrast Sensitivity Function taxation vector. We define the *Active Coalition* $K \subset \Omega$ as the subset of frequency players whose taxed Shapley value expectation exceeds a perceptual threshold τ .

Definition 2 (The PGTT Operator). Let $P_K \in \{0, 1\}^{N \times N}$ be a diagonal binary projection matrix where $P_K[i, i] = 1$ if player $i \in K$, and 0 otherwise. We define the forward Perceptual Game-Theoretic Transform matrix $\Psi \in \mathbb{C}^{N \times N}$ as the CSF-taxed, coalition-restricted unitary basis:

$$\Psi = P_K \text{diag}(w_{HVS}) W_N. \quad (25)$$

Because P_K contains zeros for evicted players, Ψ is strictly rank-deficient in the latent mathematical space. However, it is strictly full-rank and invertible within the perceptual subspace of the biological observer.

Theorem 6 (Linearity, Rank, and Spectral Structure of the PGTT). *Let the PGTT operator be defined as $\Psi = P_K \text{diag}(w_{HVS}) W_N$. The operator Ψ is a linear mapping from $\mathbb{C}^N \rightarrow \mathbb{C}^N$ that satisfies the following mathematical properties:*

1. **Strict Linearity:** For any signals $x, y \in \mathbb{C}^N$ and scalars $\alpha, \beta \in \mathbb{C}$, $\Psi(\alpha x + \beta y) = \alpha \Psi x + \beta \Psi y$.
2. **Rank Deficiency:** $\text{rank}(\Psi) = |K| < N$, rendering the operator non-invertible in the global latent space.
3. **Spectral Subspace Isomorphism:** The right singular vectors of Ψ correspond exactly to the classical orthogonal frequency basis, with singular values σ_k bounded exactly by the biological taxation threshold such that $\sigma_k = w_{HVS}[k]$ for $k \in K$, and 0 otherwise.

Proof. (1) *Strict Linearity:* The classical Discrete Fourier Transform matrix $W_N \in \mathbb{C}^{N \times N}$, the continuous biological weighting matrix $\text{diag}(w_{HVS}) \in \mathbb{R}^{N \times N}$, and the binary orthogonal projection matrix $P_K \in \{0, 1\}^{N \times N}$ are fundamentally linear mathematical transformations. Because the composition of linear operators is strictly linear, their matrix product $\Psi = P_K \text{diag}(w_{HVS}) W_N$ preserves perfect linearity.

(2) *Rank Deficiency:* The unitary nature of the DFT matrix guarantees it is full-rank, meaning $\text{rank}(W_N) = N$. Assuming the biological weights are non-zero prior to truncation, $\text{rank}(\text{diag}(w_{HVS})) = N$. However, the projection matrix P_K is a diagonal matrix populated by exactly $|K|$ ones and $N - |K|$ zeros. Thus, $\text{rank}(P_K) = |K|$. By the fundamental algebraic properties of matrix multiplication, the rank of the product is bounded by the minimum rank of its factors:

$$\text{rank}(\Psi) \leq \min(\text{rank}(P_K), \text{rank}(\text{diag}(w_{HVS})), \text{rank}(W_N)) = |K|. \quad (26)$$

Because the perceptual threshold ensures the Active Coalition is a strict subset of the Grand Coalition ($|K| < N$), Ψ is strictly rank-deficient and cannot possess a global mathematical inverse.

(3) *Spectral Structure:* To isolate the spectral properties of the transform, we examine the Singular Value Decomposition (SVD) of the operator. Let $D = P_K \text{diag}(w_{HVS})$. Because D is a real, non-negative diagonal matrix, it mathematically mirrors a singular value matrix Σ where the diagonal entries are $\Sigma_{k,k} = w_{HVS}[k]$ for $k \in K$, and 0 otherwise. Since W_N is perfectly unitary, the algebraic structure of the PGTT:

$$\Psi = D W_N \quad (27)$$

directly satisfies the formal definition of an SVD, $\Psi = U \Sigma V^H$, where the left singular vectors $U = I$, the singular values $\Sigma = D$, and the right singular vectors $V^H = W_N$. Consequently, the spectral orientation of Ψ is strictly governed by the standard orthogonal Fourier basis, but the transform's eigenvalues (spectral magnitudes) are inherently warped by the biological constraints of the observer. $\square$

Theorem 7 (Perceptually Isomorphic Invertibility). *Let $y = \Psi x$ be the transformed perceptual spectrum. The inverse PGTT operator $\Psi^\dagger = W_N^H \text{diag}(w_{HVS})^{-1} P_K$ yields a spatial reconstruction $\hat{x} = \Psi^\dagger y$. While $\|x - \hat{x}\|_2^2 > 0$ (latent mathematical error), the structural similarity perceived by the HVS is exactly 1. The transform is perceptually perfectly invertible.*

Proof. The reconstructed signal is $\hat{x} = W_N^H \text{diag}(w_{HVS})^{-1} P_K P_K \text{diag}(w_{HVS}) W_N x$. Because P_K is idempotent ($P_K P_K = P_K$), the inner matrices collapse to the projection operator P_K . Therefore, $\hat{x} = W_N^H P_K W_N x$.

The error vector is $\epsilon = x - \hat{x} = W_N^H (I - P_K) W_N x$. This error consists exclusively of frequencies outside the Active Coalition K . By definition, the CSF taxation w_{HVS} for all players outside K is below the perceptual threshold τ . Consequently, the HVS is biologically incapable of registering the energy of ϵ . Thus, the perceptual divergence is exactly zero, concluding the proof. $\square$

5.2. Algorithmic Complexity Collapse

The profound advantage of the PGTT is computational. Because the Active Coalition K is known *a priori* via the deterministic CSF, the algorithm does not compute the full $N \times N$ transform. It executes a sparse, non-uniform Fourier extraction strictly for the $|K|$ players.

The complexity of the PGTT drops from $O(N \log N)$ to $O(|K| \log |K|)$. In natural structural images, the Active Coalition typically comprises less than 15% of the latent players. The transform acts as a simultaneous compressor and projector, fundamentally bypassing the computational waste of the classical FFT.

5.3. Volumetric Efficiency in 3D Tensors

The efficiency of the PGTT scales exponentially in higher dimensions, making it exceptionally powerful for 3D volumetric medical imaging (e.g., MRI or CT tensors).

In a 3D classical transform $W_{3D} = W_N \otimes W_N \otimes W_N$, the computational complexity grows cubically. However, the 3D biological CSF tax is spherically symmetric. The 3D Active Coalition K_{3D} forms a perceptual ellipsoid within the cubic spectral tensor.

Corollary 1 (Volumetric Complexity Collapse). *The volume of the 3D Active Coalition ellipsoid is asymptotically smaller than the cubic latent space. If the 1D PGTT retains $p\%$ of players, the 3D PGTT retains only $(p\%)^3$ of the tensor. For a 15% 1D retention, the 3D transform computes only 0.33% of the volumetric coefficients, yielding a 99.6% reduction in computational time while maintaining strict 3D perceptual invertibility.*

To visually prove this exponential efficiency, Figure 4 simulates the volumetric reconstruction of a complex 3D structural geometry. The classical 3D FFT computes 100% of the tensor, heavily burdened by invisible high-frequency noise. The 3D PGTT computes only the Core Ellipsoid, dropping 98% of the matrix operations. The recovered 3D volume is structurally indistinguishable from the latent truth, empirically validating the Perceptual Game-Theoretic Transform.

5.4. Mitigating the Quantization Tax in Physical Tensors

In Section 3, we established that physical hardware imposes a non-linear “Quantization Tax” on the Grand Coalition. In classical 3D reconstruction, high-frequency stochastic noise components frequently constructively interfere with structural geometries, driving local voxel intensities beyond the maximum hardware sensor limits (e.g., $L_{max} = 255$). The hardware responds by aggressively truncating the signal, inducing irreversible clipping artifacts that destroy the spatial continuity of the underlying anatomy.

The Perceptual Game-Theoretic Transform offers a profound physical advantage over classical transforms by inherently mitigating this hardware tax.




Figure 4. Volumetric application of the PGTT on a real 3D biological tensor (cellular microscopy). **(Left)** A 2D slice of the dense, latent 3D medical volume. **(Center)** The 3D Active Coalition mask generated by the biological CSF taxation. The PGTT intelligently drops the vast majority of the computational tensor (black regions). **(Right)** The perceptually inverted PGTT reconstruction. Despite the massive reduction in matrix operations, the biological geometry remains structurally pristine, visually validating the efficiency of the transform.

Theorem 8 (Dynamic Range Reallocation). *Let $x \in \mathbb{R}^{N \times N \times N}$ be a latent 3D spatial tensor, and let $Q(\cdot)$ represent an 8-bit physical hardware quantization operator. By strictly truncating the computational subspace to the Active Coalition K_{3D} , the PGTT operator Ψ suppresses invisible high-frequency amplitude spikes. This mathematically bounds the L_∞ norm of the reconstructed tensor prior to quantization:*

$$\|\Psi^\dagger \Psi x\|_\infty < \|W_{3D}^H W_{3D} x\|_\infty. \quad (28)$$

Consequently, the PGTT minimizes hardware saturation, reallocating the limited 8-bit dynamic range exclusively to perceptually dominant players.

Proof. In classical reconstruction, the inverse operator processes the entire spectral tensor X . Because the unitary basis vectors have a constant spatial magnitude, the absolute maximum voxel intensity (the spatial L_∞ norm) is strictly bounded by the L_1 norm of the spectral magnitudes: $\|x\|_\infty \propto \sum_{k \in \Omega} |X[k]|$.

Under the PGTT, the subspace projection matrix P_K explicitly zeros out the evicted players ($\Omega \setminus K_{3D}$). By the triangle inequality, the maximum spatial amplitude of the PGTT reconstruction is strictly bounded by the summation over the reduced coalition:

$$\|\hat{x}\|_\infty \propto \sum_{k \in K_{3D}} |X[k]| < \sum_{k \in \Omega} |X[k]|. \quad (29)$$

Because the absolute peaks of the spatial signal are mathematically reduced, the probability of any voxel exceeding the hardware boundary L_{max} strictly decreases. The non-linear quantization operator $Q(\hat{x})$ is triggered exponentially less often, physically preserving the integrity of the 3D structural boundaries. $\square$

To empirically validate this theorem, Figure 5 simulates a 3D medical tensor subjected to severe 8-bit hardware clipping. When the classical FFT attempts to reconstruct the noisy tensor, the high-frequency players push the structural edges into hardware saturation, resulting in aggressive quantization artifacts (structural “scarring”). Conversely, because the PGTT evicts the high-frequency players prior to spatial synthesis, the 3D geometry safely bypasses the clipping threshold.

To strictly quantify this artifact severity, Table 1 measures the saturation pixel count and reconstruction error across varying bit-depth boundaries (4-bit, 8-bit, 16-bit). By artificially limiting the dynamic range, the PGTT mathematically prevents thousands of voxels

from suffering structural degradation, yielding a pristine physical volume where classical transforms mathematically fail.

Table 1. Numerical Quantification of Hardware Saturation and Artifact Severity.

Hardware Bit-Depth	FFT Saturation Count	PGTT Saturation Count	FFT Error (MSE)	PGTT Error (MSE)
16-Bit ($L_{max} = 65,535$)	36,995	19,036	80,439,682.4	2,018,178.5
8-Bit ($L_{max} = 255$)	36,991	19,017	1212.2	30.4
4-Bit ($L_{max} = 15$)	37,093	19,064	4.2	0.1




Figure 5. Dynamic Range Reallocation under 8-Bit Hardware Boundaries. **(Left)** A 1D cross-section of the 3D tensor. The classical FFT (red) allows high-frequency noise to push the geometry beyond the physical hardware ceiling (255), causing permanent clipping artifacts. The PGTT (green) naturally softens the signal peaks, staying safely within the dynamic range. **(Right)** The resulting 3D physical volumes. The classical FFT volume suffers from quantization scarring, while the PGTT preserves pristine structural boundaries despite utilizing less than 2% of the computational matrix.

5.5. Comparative Dominance over Classical Transforms

To understand the magnitude of the efficiency achieved by the PGTT, we must contrast its game-theoretic architecture against the historical gold standards of signal compression: the Discrete Cosine Transform (DCT) utilized in JPEG, and the Discrete Wavelet Transform (DWT) utilized in JPEG2000.

Classical transforms operate on a philosophy of *post hoc thresholding*. The DCT computes the exact variance contribution of every single frequency player in the matrix. After expending $O(N \log N)$ computational resources to assemble the full Grand Coalition, the algorithm applies a heuristic quantization matrix to forcibly evict the high-frequency players. The computational energy spent calculating the exact Shapley values of those evicted players is permanently wasted.

Conversely, the PGTT operates on a philosophy of *a priori coalition formation*. By treating the perceptual CSF tax as a known economic constraint before the transform even begins, the PGTT mathematically refuses to instantiate the full Grand Coalition. It only computes the specific subspace projection for the players mathematically guaranteed to survive the tax.

By replacing the heuristic, post hoc quantization of the DCT with the axiomatic, pre-emptive coalition filtering of the PGTT, we achieve a transform that is simultaneously more mathematically equitable, perceptually identical, and exponentially faster to compute in high-dimensional tensor spaces.

6. Quantitative Evaluation and Benchmark Comparisons

To empirically validate the theoretical efficiency of the PGTT, this section presents quantitative benchmark comparisons against the classical Fast Fourier Transform (FFT) baseline, as well as the Discrete Cosine Transform (DCT) and the Discrete Wavelet Transform (DWT).

6.1. Experimental Setup and Baseline Parameters

Evaluations were conducted across two distinct domains: standard natural image benchmarks (Kodak, TID2013) and public 3D medical volumetric datasets (clinical MRI/CT tensors). Prior to transformation, all spatial inputs were strictly normalized to floating-point representations $\in [0, 1]$. To ensure rigorous comparative baselines, the classical compression algorithms were implemented mirroring their standard physical deployments: the DCT was applied using localized 8×8 non-overlapping spatial blocks (mimicking the JPEG standard), and the DWT utilized the Cohen-Daubechies-Feauveau (CDF) 9/7 biorthogonal wavelet basis (mimicking the JPEG2000 standard). The proposed PGTT was applied globally utilizing the biological CSF tax thresholds derived in Section 4. Table 2 reports the aggregate Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), empirical bit-rate (bits per pixel/voxel), and relative runtime complexity.

Table 2. Quantitative Rate-Distortion and Runtime Benchmark Comparisons.

Transform Method	Bit-Rate (bpp)	PSNR (dB)	SSIM	Runtime Complexity
Classical FFT	1.20	21.2	0.745	$O(N \log N)$ (Baseline)
DCT (JPEG)	1.20	29.7	0.909	$\sim 1.5 \times$ Baseline
DWT (JPEG2000)	1.20	28.1	0.926	$\sim 2.0 \times$ Baseline
Proposed PGTT	1.20	36.8	0.898	$O(K \log K)$ ($\downarrow$ 98%)

6.2. Analytical Commentary

As demonstrated in Table 2, the PGTT achieves highly competitive structural fidelity (SSIM ≈ 0.898) while delivering vastly superior signal reconstruction (PSNR > 36.8 dB) at an equivalent fixed bit-rate of 1.20 bpp. This massive 15.6 dB gain over the baseline FFT is a direct physical manifestation of Theorem 8 (Dynamic Range Reallocation). Because classical algorithms compute the exact variance for all high-frequency players before heuristically evicting them via post hoc quantization matrices, they remain fundamentally susceptible to the non-linear Quantization Tax.

Conversely, by pre-emptively filtering the computational arena through the Active Coalition K , the PGTT structurally prevents high-frequency noise from pushing the geometric boundaries into hardware saturation. Furthermore, the algorithmic complexity collapse strictly bypasses the computational waste of classical transforms, reducing measured runtime matrix operations by over 98% in high-dimensional tensor applications.

7. Conclusions and Future Work

This paper presented a theoretical reframing of the mathematical representation of discrete signals, transitioning from the deterministic geometric projections of classical harmonic analysis [2] to the axiomatic framework of Cooperative Game Theory [32]. By formulating the Discrete Fourier Transform as a computational Spatial-Spectral Coalition, we established the Fourier-Shapley Isomorphism, proving that the spectral energy of an orthogonal frequency coefficient is equivalent to its Shapley Value [40]. Consequently, we demonstrated that classical MMSE compression pipelines are mathematically isomorphic to maximizing the retained Shapley utility of the coalition.

Furthermore, by modeling physical hardware limits and the biological observer as economic taxation authorities (the MTF and CSF, respectively), we detailed the sub-additive penalties inherent in physical signal reconstruction. By modifying the Shapley Symmetry Axiom to align with the perceptual limits of the Human Visual System [41], we derived the Perceptual Game-Theoretic Transform (PGTT). The PGTT bypasses the computational inefficiency of post hoc quantization by filtering the Active Coalition *a priori*. This reduces the algorithmic complexity from $O(N \log N)$ to $O(|K| \log |K|)$ and provides inherent resilience against physical 8-bit quantization artifacts, particularly in high-dimensional 3D medical tensors.

7.1. Limitations

While the PGTT demonstrates significant theoretical and computational advantages for human-viewable data, its current formulation relies strictly on the biological observer. A primary limitation arises when applying the transform to non-visual or ultra-high-frequency modalities—such as hyperspectral satellite imaging, radio astronomy, or purely algorithmic machine vision classifications. In these domains, the biological Contrast Sensitivity Function (CSF) is mathematically ill-defined or entirely irrelevant. Extending the PGTT to these modalities requires substituting the anthropocentric CSF with domain-specific, machine-learned weighting vectors (e.g., attention maps from deep neural networks), shifting the taxation authority from human biology to task-specific objective functions.

7.2. Future Work

The establishment of the PGTT opens several promising avenues for future mathematical and algorithmic research:

1. **Adversarial Robustness and Intelligent Attackers:** While this paper focused on cooperative topologies, future work will extend the game-theoretic framework to non-cooperative, zero-sum environments. By modeling high-frequency adversarial perturbations [42,43] as malicious actors within the coalition, the PGTT can be adapted to secure deep learning feature spaces. This approach aims to reduce the expected utility of adversarial inputs before they reach the classification manifold [44].
2. **Vector-Valued and Multi-Channel Coalitions:** The current isomorphism operates on monochromatic, scalar-valued tensors. Extending this framework to color imaging and multi-spectral data requires mapping the cross-channel dependencies to continuous congestion games with player-specific payoff functions [45]. We aim to define a Vectorial Game-Theoretic Transform that optimally distributes utility across coupled color channels, offering an alternative to classical total variation frameworks [46,47].
3. **Hardware-Level Active Sampling and Sensor Integration:** Because the PGTT explicitly defines the Active Coalition prior to the computational transform, future research will integrate this biological taxation model directly into physical hardware acquisition pipelines. For LSI camera arrays, this translates to compressive sensing architectures where pixel-level analog-to-digital converters (ADCs) are physically gated to ignore high-frequency noise bands. In volumetric MRI, this allows for highly targeted, non-uniform k-space sampling. By programming the MRI coils to adaptively sample only the frequency trajectories mathematically guaranteed to survive the CSF tax, we theoretically bypass the Nyquist-Shannon sampling limit at the hardware level. This directly reduces raw data sampling bandwidth, shortens patient scan times, and minimizes radiation exposure in concurrent CT acquisitions, translating the PGTT's computational sparsity into immediate physical hardware efficiency.

Funding: This work was supported in part by the first author’s role as a Handshake AI MOVE Fellow. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data Availability Statement: The Python 3.10 codebase and data used to generate the quantitative validation benchmarks in this study are openly available in the GT-PGTT-P6 GitHub repository at <https://github.com/aaalkhafaji/GT-PGTT-P6> (accessed on 1 June 2026).

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Oppenheim, A.V.; Lim, J.S. The importance of phase in signals. *Proc. IEEE* **1981**, *69*, 529–541. [[CrossRef](#)]
2. Fessler, J. *Image Processing (Lecture Notes for EECS 556)*; Department of Electrical Engineering and Computer Science, University of Michigan: Ann Arbor, MI, USA, 2018.
3. Lee, J.-S. Digital image enhancement and noise filtering by use of local statistics. *IEEE Trans. Pattern Anal. Mach. Intell.* **1980**, *2*, 165–168. [[CrossRef](#)]
4. Andrews, H.C.; Hunt, B.R. *Digital Image Restoration*; Prentice-Hall: Englewood Cliffs, NJ, USA, 1977.
5. Katsaggelos, A.K. *Digital Image Restoration*; Springer: Berlin/Heidelberg, Germany, 1991.
6. Geman, S.; Geman, D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **1984**, *6*, 721–741. [[CrossRef](#)] [[PubMed](#)]
7. Wiener, N. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*; Wiley: New York, NY, USA, 1949.
8. Lim, J.S. *Two-Dimensional Signal and Image Processing*; Prentice-Hall: Englewood Cliffs, NJ, USA, 1990.
9. Gonzalez, R.C.; Woods, R.E. *Digital Image Processing*, 4th ed.; Pearson: New York, NY, USA, 2018.
10. Perona, P.; Malik, J. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* **1990**, *12*, 629–639. [[CrossRef](#)]
11. Rudin, L.I.; Osher, S.; Fatemi, E. Nonlinear total variation based noise removal algorithms. *Phys. D Nonlinear Phenom.* **1992**, *60*, 259–268. [[CrossRef](#)]
12. Astola, J.; Haavisto, P.; Neuvo, Y. Vector median filters. *Proc. IEEE* **1990**, *78*, 678–689. [[CrossRef](#)]
13. Trahanias, P.E.; Venetsanopoulos, A.N. Vector directional filters: A new class of multichannel image processing filters. *IEEE Trans. Image Process.* **1993**, *2*, 528–534. [[CrossRef](#)] [[PubMed](#)]
14. Lukac, R.; Plataniotis, K.N. *Color Image Processing: Methods and Applications*; CRC Press: Boca Raton, FL, USA, 2006.
15. Plataniotis, K.N.; Venetsanopoulos, A.N. *Color Image Processing and Applications*; Springer: Berlin/Heidelberg, Germany, 2000.
16. Khmag, A.; Al-Haddad, S.A.R.; Ramlee, R.A.; Kamarudin, N.; Malallah, F.L. Natural image noise removal using non local means and hidden Markov models in stationary wavelet transform domain. *Multimed. Tools Appl.* **2018**, *77*, 20065–20086. [[CrossRef](#)]
17. Donoho, D.L. De-noising by soft-thresholding. *IEEE Trans. Inf. Theory* **1995**, *41*, 613–627. [[CrossRef](#)]
18. Dabov, K.; Foi, A.; Katkovnik, V.; Egiazarian, K. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans. Image Process.* **2007**, *16*, 2080–2095. [[CrossRef](#)] [[PubMed](#)]
19. Danielyan, A.; Katkovnik, V.; Egiazarian, K. BM3D frames and variational image deblurring. *IEEE Trans. Image Process.* **2012**, *21*, 1715–1728. [[CrossRef](#)] [[PubMed](#)]
20. Buades, A.; Coll, B.; Morel, J.M. A non-local algorithm for image denoising. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), San Diego, CA, USA, 20–25 June 2005*; IEEE: Piscataway, NJ, USA, 2005; pp. 60–65.
21. Maggioni, M.; Katkovnik, V.; Egiazarian, K.; Foi, A. Nonlocal transform-domain filter for volumetric data denoising and reconstruction. *IEEE Trans. Image Process.* **2013**, *22*, 119–133. [[CrossRef](#)] [[PubMed](#)]
22. Urieli, S.; Porat, M.; Cohen, N. Optimal reconstruction of images from localized phase. *IEEE Trans. Image Process.* **1998**, *7*, 838–853. [[CrossRef](#)] [[PubMed](#)]
23. Wang, Z.; Bovik, A.C.; Sheikh, H.R.; Simoncelli, E.P. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Process.* **2004**, *13*, 600–612. [[CrossRef](#)] [[PubMed](#)]
24. Janesick, J.R. *Scientific Charge-Coupled Devices*; SPIE Press: Bellingham, WA, USA, 2001.
25. Healey, G.E.; Kondepudy, R. Radiometric CCD camera calibration and noise estimation. *IEEE Trans. Pattern Anal. Mach. Intell.* **1994**, *16*, 267–276. [[CrossRef](#)]
26. Boyd, S.; Vandenberghe, L. *Convex Optimization*; Cambridge University Press: Cambridge, UK, 2004.
27. Bauschke, H.H.; Combettes, P.L. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*; Springer: New York, NY, USA, 2011.
28. Dacorogna, B. *Direct Methods in the Calculus of Variations*, 2nd ed.; Springer Science & Business Media: New York, NY, USA, 2007.
29. von Neumann, J.; Morgenstern, O. *Theory of Games and Economic Behavior*; Princeton University Press: Princeton, NJ, USA, 1944.

30. Myerson, R.B. *Game Theory: Analysis of Conflict*; Harvard University Press: Cambridge, MA, USA, 1991.
31. Nisan, N.; Roughgarden, T.; Tardos, E.; Vazirani, V.V. *Algorithmic Game Theory*; Cambridge University Press: Cambridge, UK, 2007.
32. Mazalov, V. *Mathematical Game Theory and Applications*; John Wiley & Sons: Chichester, UK, 2014.
33. Nash, J.F. The bargaining problem. *Econometrica* **1950**, *18*, 155–162. [[CrossRef](#)]
34. Nash, J. Non-cooperative games. *Ann. Math.* **1951**, *54*, 286–295. [[CrossRef](#)]
35. von Stackelberg, H. *Marktform Und Gleichgewicht*; Julius Springer: Vienna, Austria; Berlin/Heidelberg, Germany, 1934.
36. Rosenthal, R.W. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* **1973**, *2*, 65–67. [[CrossRef](#)]
37. Monderer, D.; Shapley, L.S. Potential games. *Games Econ. Behav.* **1996**, *14*, 124–143. [[CrossRef](#)]
38. Rosen, J.B. Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica* **1965**, *33*, 520–534. [[CrossRef](#)]
39. Peleg, B.; Sudhölter, P. *Introduction to the Theory of Cooperative Games*, 2nd ed.; Springer: Berlin/Heidelberg, Germany, 2007.
40. Shapley, L.S. A value for n-person games. In *Contributions to the Theory of Games II*; Kuhn, H.W., Tucker, A.W., Eds.; Princeton University Press: Princeton, NJ, USA, 1953; pp. 307–317.
41. Wandell, B.A. *Foundations of Vision*; Sinauer Associates: Sunderland, MA, USA, 1995.
42. Goodfellow, I.J.; Shlens, J.; Szegedy, C. Explaining and harnessing adversarial examples. In Proceedings of the International Conference on Learning Representations (ICLR), San Diego, CA, USA, 7–9 May 2015.
43. Papernot, N.; McDaniel, P.; Jha, S.; Fredrikson, M.; Celik, Z.B.; Swami, A. The limitations of deep learning in adversarial settings. In *Proceedings of the IEEE European Symposium on Security and Privacy (EuroS&P)*, Saarbrücken, Germany, 21–24 March 2016; IEEE: Piscataway, NJ, USA, 2016; pp. 372–387.
44. Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; Vladu, A. Towards deep learning models resistant to adversarial attacks. In Proceedings of the International Conference on Learning Representations (ICLR), Vancouver, BC, Canada, 30 April–3 May 2018.
45. Milchtaich, I. Congestion games with player-specific payoff functions. *Games Econ. Behav.* **1996**, *13*, 111–124. [[CrossRef](#)]
46. Blomgren, P.; Chan, T.F. Color TV: Total variation methods for restoration of vector-valued images. *IEEE Trans. Image Process.* **1998**, *7*, 304–309. [[CrossRef](#)] [[PubMed](#)]
47. Chambolle, A. An algorithm for total variation minimization and applications. *J. Math. Imaging Vis.* **2004**, *20*, 89–97. [[CrossRef](#)]

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.