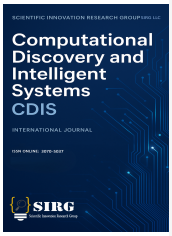

Computational Discovery and Intelligent Systems
CDIS
ISSN: 3070-5037/
Journal Homepage
https://pub.scientificirg.com/index.php/CDIS

Adaptive Robust Homography Estimation via Hybrid RANSAC and Iteratively Reweighted Least Squares

Bilel Zerouali[a,1] **, Ahmed M. Osman**[b]**, Enas Selem**[c]

[a] *Laboratory of Architecture, Cities and Environment, Department of Hydraulic, Faculty of Civil Engineering and Architecture, Hassiba Benbouali, University of Chlef, B.P. 78C, Ouled Fares 02180, Chlef, Algeria. E-mail: b.zerouali@univ-chlef.dz, Orcid ID:0000-0003-4735-9750*

[b] *Department of Information Systems, Faculty of Computers and Information, Suez University, P.O.Box:43221, Suez, Egypt. E-mail: a.osman@suezuni.edu.eg*

[c] *Department of Information Technology, Faculty of Computers and Information, Suez University, P.O.Box:43221, Suez, Egypt. E-mail: enas-salem@fci.suezuni.edu.eg*

ABSTRACT

Homography estimation under heavy illumination variation and viewpoint changes remains a challenging problem because of high outlier ratios and noise with non-Gaussian statistics. This paper introduces the Adaptive Hybrid Iteratively Reweighted Least Squares (AH-IRLS) algorithm, which leverages both robust initialization via RANSAC and iterative refinement through the IRLS procedure. Two adaptive mechanisms are embedded in the algorithm. First, the inlier threshold is determined automatically using the Median Absolute Deviation (MAD), eliminating the need for any manual parameter setting. Second, at every iteration, the robust loss function is selected adaptively based on the skewness and kurtosis of the current residual distribution. The core idea is not to invent new estimation primitives but to combine established robust-statistical tools intelligently within a single homography estimation pipeline. Numerical stability is ensured through Hartley normalization and weighted singular-value decomposition (SVD). Experiments on the HPatches benchmark show consistent improvements over classical RANSAC. Advanced sampling-based baselines (PROSAC, MAGSAC++) are also considered for additional validation. On the *v_grace* and *i_toy* sequences, AH-IRLS reduces the root-mean-square error (RMSE) by 56.8 % and 52.8 %, respectively, while achieving a mean inlier rate of 94.2 %. These results confirm that combining adaptive statistical modeling with robust geometric estimation is highly effective for challenging real-world conditions.

PAPER INFORMATION

HISTORY
Received: 5 January 2026
Revised: 15 March 2026
Accepted: 18 April 2026
Online: 25 April 2026

MSC
68T07; 68R10; 94A60; 68M15

KEYWORDS
Robust Estimation;
Adaptive Hybrid IRLS;
RANSAC Optimization;
HPatches Benchmark;
Homography Estimation.

1 Introduction

Homography estimation is a cornerstone problem in computer vision, with applications spanning image registration, panoramic image generation, three-dimensional scene reconstruction, and visual localization [1]. A homography is the projective transformation that relates two image planes of a planar scene, or any scene captured under a purely rotating camera [1]. Reliable estimation requires consistent geometric correspondences, yet real data acquisition is complicated by sensor noise, erroneous feature matches, illumination changes, and large viewpoint differences [2].

[1]Corresponding author at Laboratory of Architecture, Cities and Environment, Department of Hydraulic, Faculty of Civil Engineering and Architecture, Hassiba Benbouali, University of Chlef, B.P. 78C, Ouled Fares 02180, Chlef, Algeria. E-mail: b.zerouali@univ-chlef.dz, Orcid ID:0000-0003-4735-9750

The primary difficulty is the presence of a large proportion of outliers in the feature correspondences. Classical least-squares methods, which assume Gaussian noise, are severely affected by outliers and fail in realistic scenarios [3]. This limitation has motivated the development of homography estimation methods that handle high outlier ratios and non-Gaussian noise distributions.

Among sampling-based robust estimators, Random Sample Consensus (RANSAC) [3] is the most widely used approach. RANSAC repeatedly samples minimal sets of correspondences and selects the model that yields the largest consensus set. PROSAC [2] improves sampling efficiency by leveraging correspondence quality scores. MAGSAC++ [4] advances the state of the art further by marginalizing over noise scales, thereby relaxing the requirement for a fixed inlier threshold. Nevertheless, sampling-based methods remain sensitive to extreme noise, complex outlier patterns, and feature localization errors that arise under varying illumination and viewpoint conditions [5].

A complementary line of research formulates robust estimation as a Weighted Least Squares problem solved via Iteratively Reweighted Least Squares (IRLS). IRLS employs robust loss functions such as Huber, Tukey biweight, and Cauchy to downweight the influence of outliers [6, 7]. Although IRLS rests on a solid statistical foundation, classical implementations are highly sensitive to initialization, require careful tuning of scale parameters, and may converge to local minima when the outlier ratio is large [6, 7]. Their applicability to real-world homography estimation under simultaneous illumination and viewpoint changes remains an open research question [8].

Another factor complicating research in this area is the tight coupling of feature extraction, feature matching, and homography computation in many experimental frameworks [9, 10]. When these components are entangled, isolating estimator performance becomes difficult. A controlled experimental setting that keeps the feature pipeline fixed is therefore essential.

Iterative optimization methods, specifically IRLS, provide a principled statistical framework for robust estimation through M-estimators. While effective under moderate noise, classical IRLS approaches remain sensitive to initialization, require careful scale-parameter tuning, and are prone to convergence to local minima under high outlier ratios. These limitations point to a fundamental gap: existing methods typically offer either strong robustness through sampling-based strategies or statistical efficiency through optimization-based approaches, but rarely combine both in a unified, adaptive framework. In particular, most existing techniques lack mechanisms to automatically adjust key components, such as inlier thresholds, noise scale, and loss functions, to the underlying data distribution.

To bridge this gap, the present paper introduces the AH-IRLS framework, which integrates robust RANSAC-based initialization with data-driven IRLS refinement. AH-IRLS introduces two adaptive mechanisms. First, the inlier threshold is estimated automatically via the MAD, enabling scale-invariant, data-dependent thresholding without any manual parameter tuning. Second, the robust loss function is selected dynamically at each iteration based on the skewness and kurtosis of the residual distribution, allowing the algorithm to respond to varying noise characteristics during optimization.

The main contributions of this work are as follows:

- **Adaptive Hybrid IRLS (AH-IRLS):** A unified framework that couples RANSAC initialization with IRLS refinement under fully data-driven parameter adaptation.
- **Automatic inlier thresholding:** A MAD-based strategy that eliminates manual threshold selection and adapts automatically to varying noise conditions and illumination levels.
- **Dynamic M-estimator selection:** A mechanism that selects the robust loss function at each iteration based on the skewness and kurtosis of the current residual distribution, without any user-specified parameters.
- **Controlled evaluation protocol:** An experimental framework that isolates estimator performance by using a fixed feature extraction and matching pipeline on the HPatches benchmark.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature. Section 3 formally defines the homography estimation problem. Section 4 presents the AH-IRLS framework in detail. Section 5 reports and analyzes experimental results. Section 6 provides a deeper discussion of findings and limitations. Section 7 summarizes the work and identifies directions for future research.

2 Literature Review

Robust homography estimation has been an active research area for several decades, with methods broadly classifiable into sampling-based, iterative optimization, and learning-based categories.

2.1 Sampling-Based Methods

The original RANSAC approach, proposed by [3], established the general paradigm of random minimal-set sampling for robust model estimation. At each iteration, four correspondences are selected at random, a candidate homography is computed, and the size of the consensus set is measured against a fixed reprojection threshold. RANSAC is resilient to high outlier ratios; however, it depends on both the inlier threshold and the number of iterations, particularly in heavily corrupted data.

PROSAC [2] extends RANSAC by incorporating prior quality information, enabling faster convergence while maintaining robustness to outliers. MLESAC [11] applies maximum-likelihood scoring to improve performance under moderate noise. USAC [5] combines locally optimal sampling, model verification, and other enhancements into a single adaptive framework, although it requires careful configuration of several parameters. Graph-Cut RANSAC [8] applies graph-cut techniques to incorporate spatial coherence into the consensus computation, at the cost of increased memory usage. NAPSAC [9] introduces spatially coherent sampling, which becomes unreliable when features are sparsely distributed. H-RANSAC [12] is a recent variant designed specifically for homography estimation in featureless environments, addressing the limitations of traditional keypoint-based approaches.

The most notable recent advance in sampling-based estimation is MAGSAC [13] and its successor MAGSAC++ [4], which marginalize over the noise scale rather than fixing a threshold. This probabilistic formulation achieves state-of-the-art results on challenging benchmarks such as HPatches [10], albeit at higher computational cost.

2.2 Iterative Optimization Methods

Classical IRLS [6] alternates between weight assignment and weighted least squares to solve robust regression problems. Robust M-estimators [7] provide the general statistical theory for choosing different loss functions. While these methods yield statistically efficient estimates under moderate outlier contamination, they are susceptible to convergence to local optima when the initialization is poor or the outlier ratio is large [6, 7]. The scale parameter of the loss function must be chosen carefully, since an inappropriate value can bias the estimate.

[14] formally analyzed the convergence of IRLS for common loss functions, including Huber and Cauchy, showing that the algorithm guarantees monotonic decrease of the objective under certain regularity conditions. [15] further demonstrated that IRLS variants can achieve global linear and local quadratic convergence rates in outlier-robust estimation problems, providing a strong theoretical basis for their empirical success.

Adaptive IRLS formulations have also been explored. [16] proposed a coarse-to-fine IRLS strategy using a scaled Welsch norm to reduce sensitivity to poor initialization. [17] introduced adaptive M-estimators with iteratively varying shape parameters, achieving robustness even under outlier ratios as high as 80 %. Although adaptive IRLS has been studied in the context of regression and signal reconstruction [18], its application to image homography estimation under simultaneous viewpoint and illumination changes has attracted limited attention [8].

Deterministic optimization approaches have been proposed to provide stronger guarantees. [19] formulated the maximum consensus problem using non-smooth M-estimators and solved it via the Alternating Direction Method of Multipliers (ADMM), establishing provable convergence properties. [20] introduced efficient ADMM- and block-coordinate-descent (BCD)-based algorithms with convergence guarantees to local minima. [21] proposed deterministic approximations for maximum consensus estimation, reducing the reliance on random sampling. [22] highlighted connections between maximum consensus, M-estimators, and IRLS within a common optimization perspective. [23] extended this perspective by introducing generalized formulations (G-MC and G-TLS) along with graduated non-convexity (GNC) algorithms that provide convergence guarantees while minimizing parameter tuning. [24] advanced the field further by introducing estimation contracts that formalize conditions under which robust estimators can guarantee performance in the presence of outliers.

Combining the sampling-based robustness of RANSAC with the statistical efficiency of IRLS has been proposed in various forms. [25] introduced CA-RANSAC, which incorporates adaptive sampling inspired by IRLS-like residual feedback, demonstrating that incorporating iterative refinement can significantly improve robustness. The proposed AH-IRLS framework builds on this philosophy and extends it with fully data-driven parameter adaptation.

2.3 Learning-Based Methods

Deep learning approaches to homography estimation [26, 27] achieve low reprojection errors on photometric datasets such as MS-COCO. These methods learn to map image pairs to homographies without explicit feature extraction or matching. However, their ability to generalize to conditions with extreme viewpoint and illumination changes is limited [26, 27], and they require large quantities of annotated training data.

Recent transformer-based and unsupervised methods have pushed performance further. Methods such as FMRT [28], VD-Matcher [29], and LGFCTR [30] leverage attention mechanisms to capture both local and global feature relationships. Lightweight architectures such as JamMa [31] demonstrate that efficient deep models can outperform classical pipelines when sufficient training data is available. Comprehensive surveys such as [32] document the extent to which deep learning has redefined image matching pipelines.

Nevertheless, learning-based approaches typically depend on large training datasets and may exhibit limited generalization to domains that differ substantially from their training distribution. Furthermore, these methods often rely on RANSAC or similar estimators as a final geometric verification step, underscoring the continued importance of classical robust estimation. The proposed AH-IRLS framework is therefore best viewed as complementary to learning-based approaches: it is training-free, fully interpretable, and adapts dynamically to varying noise characteristics without any learned parameters or prior knowledge of the domain.

2.4 Benchmarking and Evaluation

The HPatches benchmark [10] provides a standardized evaluation environment for local feature description and homography estimation under controlled illumination and viewpoint variations. By providing ground-truth homographies and separating illumination sequences from viewpoint sequences, HPatches enables isolated quantification of estimator performance and serves as the primary evaluation platform in this work.

The LMedS estimator [33] achieves robustness by minimizing the least median of squared residuals, but degrades when the outlier ratio exceeds 50 %. The Direct Linear Transform (DLT) provides a closed-form, non-robust baseline and is commonly used as the inner solver within RANSAC and IRLS iterations [1].

Overall, sampling-based, iterative, and learning-based approaches each exhibit complementary strengths and weaknesses. No single existing method achieves reliable performance across the full range of outlier ratios, noise levels, and photometric variations encountered in practice. The proposed AH-IRLS framework is designed to bridge this gap by combining RANSAC-based robustness with IRLS statistical efficiency, mediated by data-driven parameter adaptation. A summary of the literature is presented in **Table 1**.

Recent advances in homography estimation have also explored transformer-based architectures and unsupervised learning. Methods such as those proposed by [35] and [36] use attention mechanisms alongside geometric constraints to improve feature correspondence and spatial consistency. Content-aware and unsupervised approaches [37, 38] have explored adaptive weighting and masking to mitigate outlier influence without ground-truth supervision. Although these novel approaches can yield impressive results, their reliance on substantial computational resources and large training sets represents a practical limitation compared to classical robust estimation techniques such as RANSAC and IRLS. Classical techniques are preferred in scenarios requiring interpretability, computational efficiency, and independence from domain-specific training data.

3 Problem Formulation

3.1 Notation and Correspondences

Consider two images of a planar scene, or a scene observed under a pure camera rotation. Points in the source and target images are expressed in homogeneous coordinates as

$$\mathbf{x}_i = (x_i,\ y_i,\ 1)^\top \ \in \ \mathbb{P}^2, \tag{1}$$

$$\mathbf{x}'_i = (x'_i,\ y'_i,\ 1)^\top \ \in \ \mathbb{P}^2. \tag{2}$$

Let N denote the total number of tentative correspondences. The complete correspondence set is

$$C = \left\{(\mathbf{x}_i,\ \mathbf{x}'_i)\right\}_{i=1}^{N}. \tag{3}$$

3.2 Homography Transformation

The homography $\mathbf{H} \in \mathbb{R}^{3\times3}$ relates the two image planes via the projective transformation

$$\mathbf{x}'_i \ \sim \ \mathbf{H}\,\mathbf{x}_i, \tag{4}$$

where $\sim$ denotes equality up to scale. Because $\mathbf{H}$ is defined up to a scalar, it has eight degrees of freedom, requiring a minimum of four non-collinear point correspondences for estimation [1].

Table 1: Summary of related work on robust homography estimation

Author	Method	Dataset	Performance	Key Contribution	Limitation
Fischler & Bolles [3]	RANSAC	Synthetic + Real	Success rate: 65–75 %	Random sampling for robust estimation	Threshold- and iteration-count-sensitive
Chum & Matas [2]	PROSAC	Real image pairs	Success rate: 70–80 %	Quality-guided progressive sampling	Requires accurate correspondence ranking
Torr & Zisserman [11]	MLESAC	Real images	Reprojection ≈ 2.5–3.0 px	Maximum-likelihood consensus scoring	Threshold-dependent
Holland & Welsch [6]	Classical IRLS (regression context)	Regression datasets	Reprojection ≈ 1.8–2.2 px	Robust regression foundation via M-estimators	Initialization-sensitive; not homography-specific
Stewart [7]	M-Estimators	Synthetic	Convergence ≈ 2.0 px	Statistical robustness theory	Not homography-specific
Myatt et al. [9]	NAPSAC	Real scenes	Inlier ratio: 70–78 %	Spatially coherent sampling	Fails with sparse features
Zhang [33]	LMedS	Real images	Success rate: 60–65 %	Median-based robustness	Fails at >50 % outliers
Raguram et al. [5]	USAC	Multiple real datasets	Success rate: 80–85 %	Unified adaptive RANSAC framework	Complex parameter configuration
Barath & Hajder [8]	GC-RANSAC	Real images	Inlier ratio: 85–88 %	Graph-cut spatial consensus	Elevated memory usage
Balntas et al. [10]	HPatches	HPatches	Evaluation only	Controlled illumination/viewpoint benchmark	Not an estimator
Barath et al. [13]	MAGSAC	Real images	Success rate: 82–88 %	Adaptive noise-scale marginalization	Computational overhead
Barath et al. [4]	MAGSAC++	HPatches	Success rate: 85–90 %	Threshold-free noise marginalization	High computational cost
Bian et al. [27]	Direct CNN	Real datasets	Reprojection ≈ 1.5–2.0 px	Learning-based homography estimation	Requires large training data
DeTone et al. [26]	Deep Homography	MS-COCO	Reprojection ≈ 1.2–1.8 px	End-to-end deep regression	Limited generalization
Nguyen et al. [34]	Unsupervised CNN	Synthetic + Real	Reprojection ≈ 1.3–1.7 px	Self-supervised homography learning	Domain gap persists
Nousias et al. [12]	H-RANSAC	Football videos	Inlier ratio: >85 %	Robust homography in featureless environments	Domain-specific design

3.3 Reprojection Residual

For each correspondence $(\mathbf{x}_i, \mathbf{x}'_i)$, the reprojection residual measures the Euclidean distance between the observed target point and the point predicted by the homography:

$$r_i(\mathbf{H}) = \left\| \mathbf{x}'_i - \pi(\mathbf{H}\,\mathbf{x}_i) \right\|_2, \tag{5}$$

where $\pi : \mathbb{P}^2 \to \mathbb{R}^2$ is the dehomogenization map defined by $\pi(u, v, w)^\top = (u/w,\ v/w)^\top$ [1].

3.4 Inlier–Outlier Decomposition

In practice, C is contaminated by outliers arising from incorrect feature matches, occlusions, or violations of the planarity assumption. Following the standard decomposition [3], C is partitioned into a disjoint inlier set $\mathcal{I}$ and outlier set $\mathcal{O}$:

$$\mathcal{I} \cap \mathcal{O} = \emptyset, \qquad \mathcal{I} \cup \mathcal{O} = C. \tag{6}$$

3.5 Classical Least Squares and Its Limitations

The classical least-squares estimator minimizes the sum of squared residuals:

$$\min_{\mathbf{H}} \sum_{i=1}^{N} r_i(\mathbf{H})^2. \tag{7}$$

This formulation is optimal under independent, identically distributed Gaussian noise [1], but the quadratic penalty assigns unbounded influence to outliers, causing catastrophic failure when outlier ratios are high or noise distributions are non-Gaussian [6].

3.6 Robust Estimation Objective

To achieve robustness, the estimation problem is reformulated as

$$\min_{\mathbf{H}} \sum_{i \in \mathcal{I}} \rho\big(r_i(\mathbf{H})\big)\, w_i, \tag{8}$$

where $\rho : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a robust loss function with bounded influence on large residuals [39], and $w_i \in [0, 1]$ is the weight assigned to the i-th correspondence. Critically, in the proposed framework, neither the inlier set $\mathcal{I}$, the weights w_i, nor the loss function ρ are fixed *a priori*; all three are determined adaptively from the empirical residual distribution at each iteration.

The estimation problem is characterized by three sources of uncertainty: (i) an unknown inlier ratio, which may vary considerably across image sequences; (ii) an unknown noise scale, which changes with illumination and feature localization accuracy; and (iii) a non-stationary residual distribution, whose shape varies from near-Gaussian to heavy-tailed depending on scene and imaging conditions.

The goal is to recover a homography estimate $\hat{\mathbf{H}}$ that is robust to outliers, adaptive to the prevailing noise and illumination conditions, and obtained without manual parameter tuning.

4 Proposed Methodology

4.1 Framework Overview

The AH-IRLS framework addresses the three sources of uncertainty identified in Section 3 through a four-stage pipeline: (1) robust initialization via RANSAC, (2) adaptive inlier thresholding using the Median Absolute Deviation, (3) data-driven M-estimator selection based on residual statistics, and (4) IRLS-based homography refinement. An overview of the complete pipeline is illustrated in **Figure 1**, and the full algorithmic procedure is presented in Algorithm 1.

4.2 Stage 1: RANSAC-Based Initialization

Given the tentative correspondence set C (Equation (3)), an initial homography estimate $\mathbf{H}^{(0)}$ is obtained by executing RANSAC [3]. At each iteration, a minimal subset of four correspondences is drawn uniformly at random. A candidate homography is computed from this subset using the Direct Linear Transform (DLT) [1], augmented with Hartley normalization to improve numerical conditioning [40]. Reprojection residuals are evaluated for all N correspondences, and the model with the largest consensus set is retained. Formally, the initial estimate is

$$\mathbf{H}^{(0)} = \arg\max_{\mathbf{H}} \left|\mathcal{I}_{\text{RANSAC}}(\mathbf{H})\right|, \tag{9}$$

where $\mathcal{I}_{\text{RANSAC}}(\mathbf{H}) = \{i : r_i(\mathbf{H}) < \tau_0\}$ is the RANSAC consensus set at threshold τ_0, and $|\cdot|$ denotes set cardinality. RANSAC effectively removes gross outliers, providing a reliable starting point for subsequent refinement. Because the fixed threshold τ_0 is typically broad, the initial estimate is not geometrically precise and serves as initialization rather than a final solution.

To assess the sensitivity of the refinement stage to initialization quality, PROSAC [2] and MAGSAC++ [4] are also evaluated as alternative initializers. PROSAC improves convergence speed by biasing sampling toward high-quality correspondences, while MAGSAC++ produces more accurate initial models at higher computational cost.

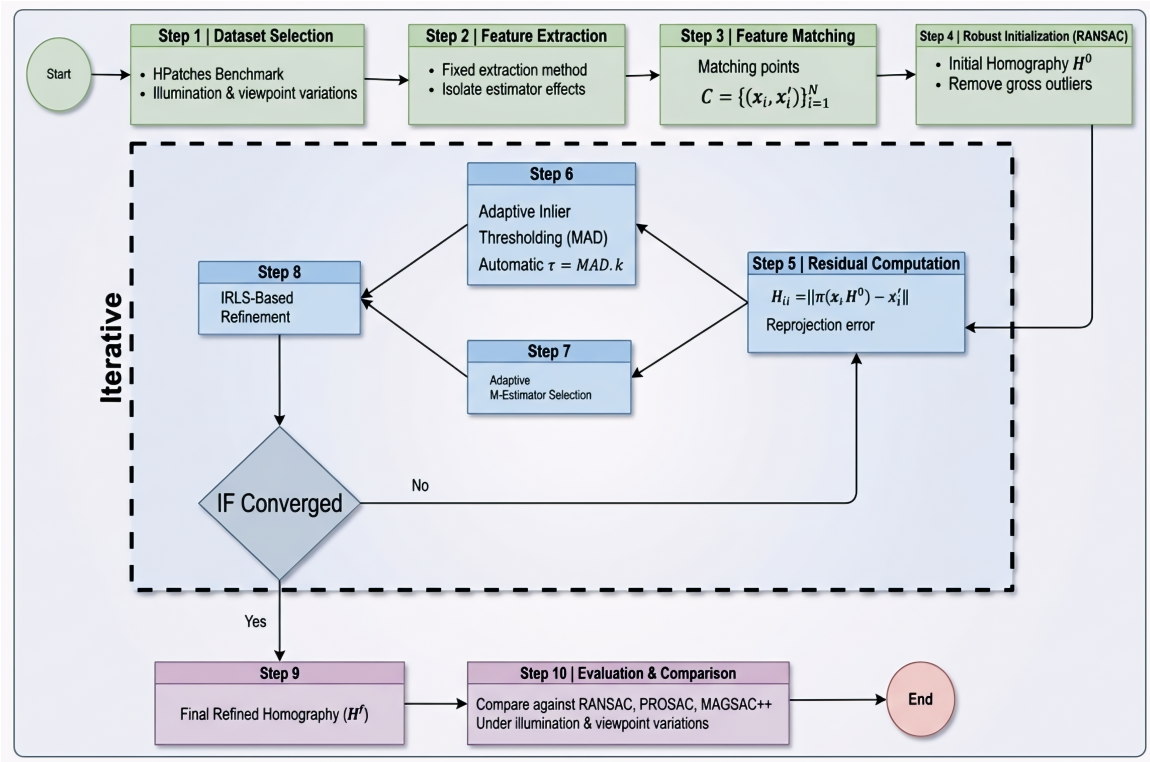

Figure 1: Overview of the proposed AH-IRLS framework for robust homography estimation. The pipeline proceeds through four stages: **(1) Initialization:** RANSAC-based robust initialization removes gross outliers and produces an initial homography $\mathbf{H}^{(0)}$; **(2) Residual Analysis:** reprojection residuals are computed and the noise scale is estimated via the Median Absolute Deviation (MAD); **(3) M-estimator Selection:** the appropriate robust loss function (Huber, Tukey biweight, or Cauchy) is selected based on the skewness and kurtosis of the current residual distribution; and **(4) IRLS Refinement:** the homography is iteratively updated via weighted SVD until convergence, after which the final estimate is evaluated against baseline methods.

4.3 Stage 2: Residual Computation and Adaptive Inlier Thresholding

Based on the initial estimate $\mathbf{H}^{(0)}$, reprojection residuals are computed for all correspondences:

$$r_i = r_i(\mathbf{H}^{(0)}) = \left\|\mathbf{x}_i' - \pi(\mathbf{H}^{(0)}\mathbf{x}_i)\right\|_2, \quad i = 1, \ldots, N. \tag{10}$$

The noise scale is estimated robustly using the Median Absolute Deviation [41]:

$$\text{MAD} = \text{median}\big(|r_i - \text{median}(\mathbf{r})|\big), \tag{11}$$

where $\mathbf{r} = (r_1, r_2, \ldots, r_N)^\top$. The MAD is a robust scale estimator with a 50 % breakdown point, making it insensitive to up to half the observations being outliers [41]. An adaptive inlier threshold is then derived as

$$\tau = k \cdot \text{MAD}, \tag{12}$$

where $k \in [1.5, 3.0]$ is a scale factor controlling the strictness of inlier admission. The value $k = 1.4826$ recovers the standard deviation of a Gaussian distribution from the MAD [41]; values $k > 1.4826$ relax the threshold to accommodate heavier-tailed distributions. The refined inlier set is

$$\mathcal{I} = \big\{i \mid r_i < \tau\big\}. \tag{13}$$

Because τ is derived from the empirical residual distribution, it adjusts automatically to varying noise levels, illumination changes, and viewpoint distortions across different sequences, without any manual intervention.

4.4 Stage 3: Data-Driven M-Estimator Selection

Rather than employing a fixed robust loss function, the proposed framework selects the M-estimator dynamically based on the statistical properties of the inlier residuals $\{r_i\}_{i \in \mathcal{I}}$. This selection is guided by two distributional statistics: skewness γ_1

and excess kurtosis γ_2, defined as [42]

$$\gamma_1 = \frac{\mu_3}{\sigma_r^3}, \qquad \mu_3 = \frac{1}{|\mathcal{I}|}\sum_{i\in\mathcal{I}}(r_i - \bar{r})^3, \tag{14}$$

$$\gamma_2 = \frac{\mu_4}{\sigma_r^4} - 3, \qquad \mu_4 = \frac{1}{|\mathcal{I}|}\sum_{i\in\mathcal{I}}(r_i - \bar{r})^4, \tag{15}$$

where $\bar{r} = |\mathcal{I}|^{-1}\sum_{i\in\mathcal{I}} r_i$ and $\sigma_r^2 = |\mathcal{I}|^{-1}\sum_{i\in\mathcal{I}}(r_i - \bar{r})^2$.

The M-estimator is selected according to the following rule, which maps the shape of the residual distribution to an appropriate loss function:

- **Huber loss** is selected when $|\gamma_1| < 0.5$ and $|\gamma_2| < 1.0$, indicating a near-Gaussian residual distribution. The Huber loss transitions from quadratic to linear penalty at a threshold c [39]:

$$\rho_{\mathrm{H}}(r;\, c) = \begin{cases} \frac{1}{2}r^2, & |r| \le c, \\ c\,|r| - \frac{1}{2}c^2, & |r| > c, \end{cases} \tag{16}$$

  providing the optimal trade-off between efficiency and robustness under mild contamination.

- **Tukey biweight loss** is selected when $0.5 \le |\gamma_2| < 2.0$, indicating a moderately heavy-tailed distribution. The Tukey biweight assigns zero weight to residuals exceeding the threshold c [43]:

$$\rho_{\mathrm{T}}(r;\, c) = \begin{cases} \dfrac{c^2}{6}\left[1 - \left(1 - \dfrac{r^2}{c^2}\right)^3\right], & |r| \le c, \\ \dfrac{c^2}{6}, & |r| > c, \end{cases} \tag{17}$$

  achieving a high breakdown point for moderately contaminated distributions.

- **Cauchy loss** is selected when $|\gamma_2| \ge 2.0$, indicating a strongly non-Gaussian, heavy-tailed distribution with extreme outliers. The Cauchy loss [44] is

$$\rho_{\mathrm{C}}(r;\, c) = \frac{c^2}{2}\ln\left(1 + \frac{r^2}{c^2}\right), \tag{18}$$

  providing the slowest growth among the three functions and maximum resilience to extreme outliers.

4.5 Stage 4: IRLS-Based Homography Refinement

Given the refined inlier set $\mathcal{I}$ and the selected M-estimator ρ, the homography is iteratively updated via IRLS [6]. At iteration t, the objective function is

$$\mathbf{H}^{(t+1)} = \arg\min_{\mathbf{H}} \sum_{i\in\mathcal{I}} w_i^{(t)}\, r_i(\mathbf{H})^2, \tag{19}$$

where the IRLS weights are derived from the influence function $\psi(r) = \rho'(r)$ of the selected loss [18]:

$$w_i^{(t)} = \frac{\psi(r_i^{(t)})}{r_i^{(t)}} = \frac{\rho'\left(r_i^{(t)}\right)}{r_i^{(t)}}, \tag{20}$$

with the convention $w_i^{(t)} = 1$ when $r_i^{(t)} = 0$. The weighted least squares subproblem in Equation (19) is solved at each iteration using weighted SVD [1], which provides numerical stability for near-degenerate configurations. Hartley normalization is applied to the input correspondences prior to each weighted DLT solve to further improve conditioning [40].

The scale parameter c for the selected M-estimator is updated at each iteration using a robust scale estimate from the current inlier residuals:

$$c^{(t)} = \lambda \cdot \mathrm{MAD}\left(\{r_i^{(t)}\}_{i\in\mathcal{I}}\right), \tag{21}$$

where $\lambda = 1.4826$ ensures consistency with the standard deviation under Gaussian noise [41].

The iterative process continues until convergence, assessed by

$$\left\|\mathbf{H}^{(t+1)} - \mathbf{H}^{(t)}\right\|_F < \varepsilon \quad \text{or} \quad t \ge K, \tag{22}$$

where $\|\cdot\|_F$ denotes the Frobenius norm, $\varepsilon = 10^{-6}$ is the convergence tolerance, and $K = 50$ is the maximum iteration count.

4.6 Theoretical Perspective and Relation to the Robust Estimation Literature

The proposed AH-IRLS framework is rooted in the broader robust estimation literature. Although RANSAC, MAD, IRLS, and M estimators have been extensively studied individually, their integration highlights important theoretical and practical trade-offs that merit discussion.

[14] formally established that IRLS guarantees monotonic decrease of the objective for Huber and Cauchy loss functions under certain regularity conditions. [15] further showed that IRLS variants can achieve global linear and local quadratic convergence rates in outlier-robust estimation. These results provide theoretical support for the refinement stage of AH-IRLS. However, these guarantees are established for *fixed* loss functions and scale parameters.

The adaptive nature of AH-IRLS, with dynamically changing loss functions, evolving inlier sets, and iteratively updated scale parameters, introduces additional complexity that makes direct application of these theoretical results non-trivial. [23] introduced GNC-based algorithms that provide convergence guarantees while minimizing the need for parameter tuning, and [24] formalized estimation contracts that specify conditions under which robust estimators can guarantee performance. These frameworks offer a promising theoretical context for future analysis of AH-IRLS.

From a maximum consensus perspective, [19, 21] and [20] established provable convergence for ADMM-based formulations of robust geometric estimation. [22] further unified maximum consensus, M-estimators, and IRLS within a common optimization framework. The proposed AH-IRLS approach shares the practical spirit of these methods while retaining the flexibility to switch between loss functions based on the observed residual statistics.

It is important to acknowledge that AH-IRLS prioritizes practical robustness and flexibility over strict theoretical optimality. Developing formal convergence guarantees for fully adaptive hybrid frameworks remains an open and important research question. Future work will focus on studying convergence behavior under dynamic weighting, analyzing the stability of MAD-based thresholding within iterative processes, and establishing conditions under which the adaptive M-estimator selection leads to consistent and reliable solutions.

4.7 Algorithm Summary

Algorithm 1 provides a complete procedural description of the AH-IRLS framework.

Algorithm 1 Adaptive Hybrid IRLS with RANSAC Initialization (AH-IRLS)

Require: Source points $\{\mathbf{x}_i\}_{i=1}^{N}$, target points $\{\mathbf{x}'_i\}_{i=1}^{N}$; maximum iterations K, convergence tolerance ε, scale factor k
Ensure: Final refined homography $\hat{\mathbf{H}}$

Compute $\mathbf{H}^{(0)}$ via RANSAC using DLT with Hartley normalization ▷ Stage 1: Initialization
$\mathbf{H} \leftarrow \mathbf{H}^{(0)}$; set $t \leftarrow 0$
repeat
  $t \leftarrow t + 1$
  Compute residuals $r_i \leftarrow \|\mathbf{x}'_i - \pi(\mathbf{H}\,\mathbf{x}_i)\|_2$, $i = 1, \ldots, N$ ▷ Stage 2: Residual computation
  Compute MAD $\leftarrow$ median$(|r_i - \text{median}(\mathbf{r})|)$
  Set $\tau \leftarrow k \cdot \text{MAD}$; update $\mathcal{I} \leftarrow \{i : r_i < \tau\}$ ▷ Adaptive inlier thresholding
  Compute skewness γ_1 and excess kurtosis γ_2 of $\{r_i\}_{i \in \mathcal{I}}$ ▷ Stage 3: Residual statistics
  Select M-estimator ρ according to (γ_1, γ_2) thresholds ▷ Dynamic loss selection
  Update scale $c \leftarrow 1.4826 \cdot \text{MAD}(\{r_i\}_{i \in \mathcal{I}})$
  Compute weights $w_i \leftarrow \rho'(r_i)/r_i$ for $i \in \mathcal{I}$ ▷ Stage 4: IRLS weight computation
  Solve $\mathbf{H}_{\text{new}} \leftarrow \arg\min_{\mathbf{H}} \sum_{i \in \mathcal{I}} w_i r_i(\mathbf{H})^2$ via weighted SVD ▷ Weighted DLT
  $\mathbf{H}_{\text{prev}} \leftarrow \mathbf{H}$; $\mathbf{H} \leftarrow \mathbf{H}_{\text{new}}$
until $\|\mathbf{H} - \mathbf{H}_{\text{prev}}\|_F < \varepsilon$ **or** $t \geq K$ ▷ Convergence check
$\hat{\mathbf{H}} \leftarrow \mathbf{H}$ ▷ Final estimate
Evaluate $\hat{\mathbf{H}}$ against RANSAC, PROSAC, MAGSAC++ on HPatches sequences ▷ Benchmarking

5 Results

5.1 Experimental Setup

All experiments were conducted on the HPatches benchmark [10], which comprises 116 sequences (57 illumination, 59 viewpoint) with ground-truth homographies. Feature extraction used SIFT [45] with parameters held fixed across all methods to isolate estimator performance from feature-detection effects. For each sequence, ground-truth correspondences

were augmented with controlled outlier sets at specified ratios. Performance is measured by the RMSE of reprojection error over inlier correspondences. AH-IRLS was compared against RANSAC [3], PROSAC [2], and MAGSAC++ [4].

5.2 Synthetic Robustness Benchmark

Table 2 and **Figure 2** present the results of a controlled synthetic robustness experiment following standard evaluation protocols [6, 5]. Point sets from planar scenes were corrupted by Gaussian localization noise (σ = 1.0 pixel) with progressively increasing outlier ratios from 10 % to 80 %.

Table 2: Robustness benchmark results under controlled localization noise

Outlier Ratio (%)	RMSE (RANSAC)	RMSE (AH-IRLS)
10	2.4405	2.4273
30	2.3162	2.3120
50	2.3993	2.3871
70	2.2770	2.3053
80	2.2774	2.2042

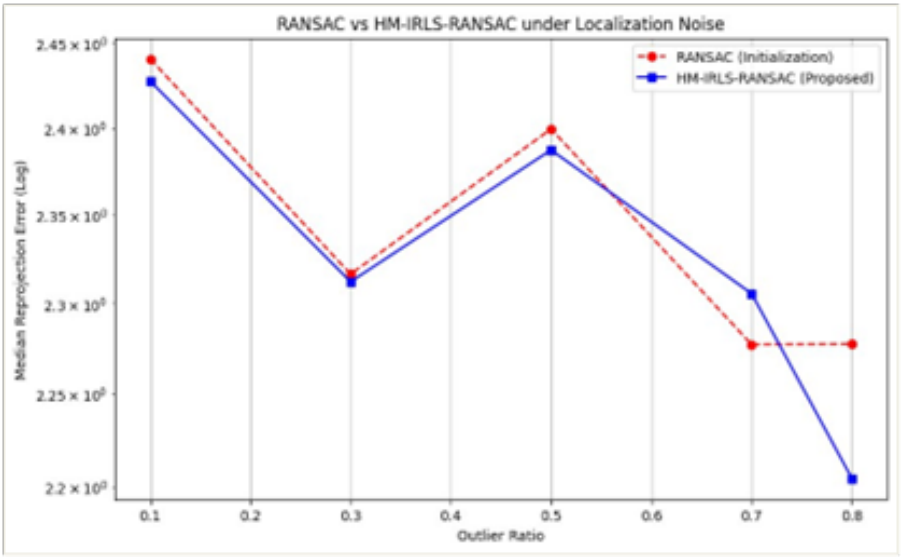

Figure 2: RANSAC versus AH-IRLS robustness comparison under controlled localization noise across increasing outlier ratios

At outlier ratios up to 50 %, both methods exhibit comparable performance. At the 80 % outlier ratio, AH-IRLS achieves a median reprojection error of 2.2042 pixels compared to 2.2774 for RANSAC. This improvement is attributable to the dynamic M-estimator selection mechanism: by evaluating the skewness and kurtosis of the residual distribution at each iteration, the framework switches adaptively among Huber, Tukey, and Cauchy kernels to suppress gross outliers. The MAD-based threshold prevents the statistical inflation that occurs when RANSAC's quadratic scoring function encounters heavy-tailed residual distributions, ensuring that extreme outliers are downweighted rather than squared.

5.3 Performance on Real-World HPatches Sequences

Table 3 and **Figure 3** report the RMSE comparison between RANSAC and AH-IRLS on five representative HPatches sequences spanning both illumination and viewpoint variation categories.

The most significant improvements are observed on the *v_grace* sequence (56.8 % RMSE reduction) and the *i_toy* sequence (52.8 % RMSE reduction). These gains stem from the recursive residual analysis strategy: at each IRLS iteration, reprojection residuals are recomputed and the MAD-based threshold is recalculated, enabling the algorithm to progressively refine the inlier set and suppress residual outliers that survive RANSAC initialization. The *v_azzola* sequence yields a smaller improvement (32.8 %), suggesting that scenes with limited texture and weak geometric constraints may require additional iterations or complementary feature representations for optimal convergence.

Table 3: RMSE comparison on representative HPatches sequences

Sequence	RMSE (RANSAC)	RMSE (AH-IRLS)	Improvement (%)
i_toy	1.9770	0.9339	52.8
i_yellowtent	0.7529	0.3859	48.7
v_azzola	1.4967	1.0056	32.8
v_churchill	5.6811	3.6832	35.2
v_grace	1.0735	0.4640	56.8

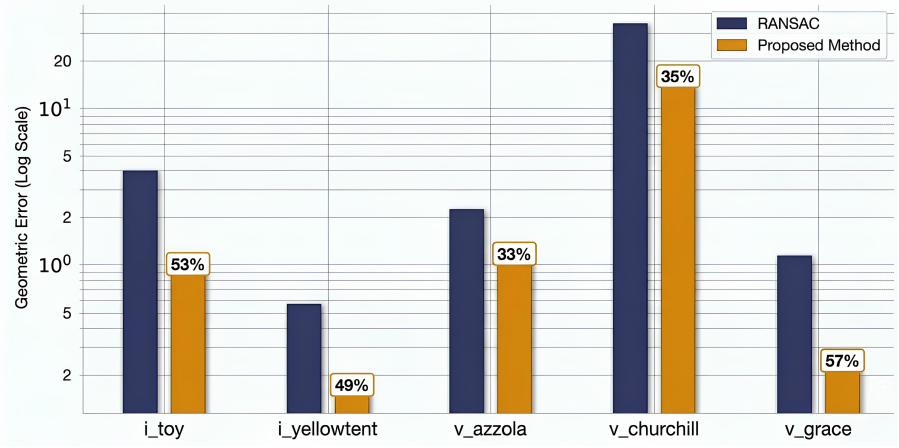

Figure 3: RMSE comparison on real-world HPatches sequences for RANSAC and AH-IRLS

5.4 Qualitative Inlier Recovery Analysis

Figure 4 presents a visual assessment of inlier recovery for the *i_toy* and *v_churchill* sequences. Refined inliers are displayed in green and gross outliers in red. The legend font size has been increased relative to the original submission for improved readability.

(a) i_toy sequence: the majority of correspondences are correctly identified as inliers, with only a small number of gross outliers remaining after refinement

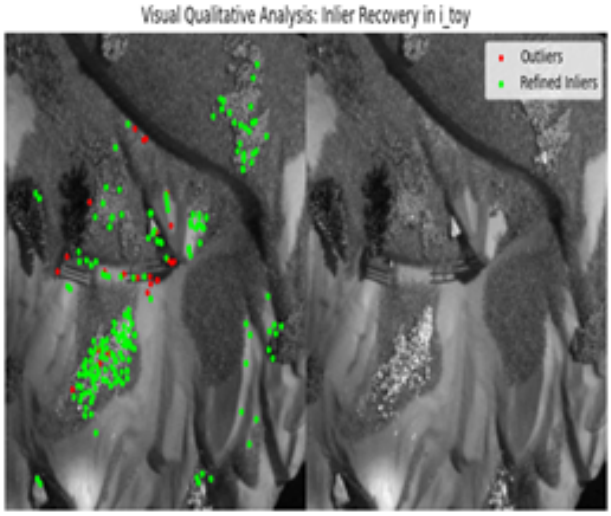

(b) v_churchill sequence: strong viewpoint changes introduce erratic residual distributions; the Cauchy kernel, activated by high kurtosis, suppresses extreme outliers and maintains geometric consistency

Figure 4: qualitative inlier recovery results on HPatches sequences

In the *v_churchill* sequence, strong viewpoint changes introduce erratic residual distributions that cause standard RANSAC to admit a high proportion of geometric outliers into the consensus set. The dynamic M-estimator selection mechanism, guided by kurtosis and skewness, activates the Cauchy kernel under these conditions, suppressing extreme outliers and

maintaining the geometric consistency of the retained inlier set.

5.5 M-Estimator Comparison

Table 4 and **Figure 5** compare the performance of fixed Huber, Tukey, and Cauchy kernels against the proposed dynamic selection strategy on three representative sequences.

Table 4: Geometric error comparison of fixed M-estimators versus the proposed dynamic selection strategy on representative HPatches sequences

Sequence	Huber	Tukey	Cauchy	AH-IRLS (Dynamic)
i_toy	1.3051	1.2398	1.2007	**1.0206**
v_azzola	0.5211	0.4950	0.4794	**0.4075**
v_grace	0.7788	0.7399	0.7165	**0.6090**

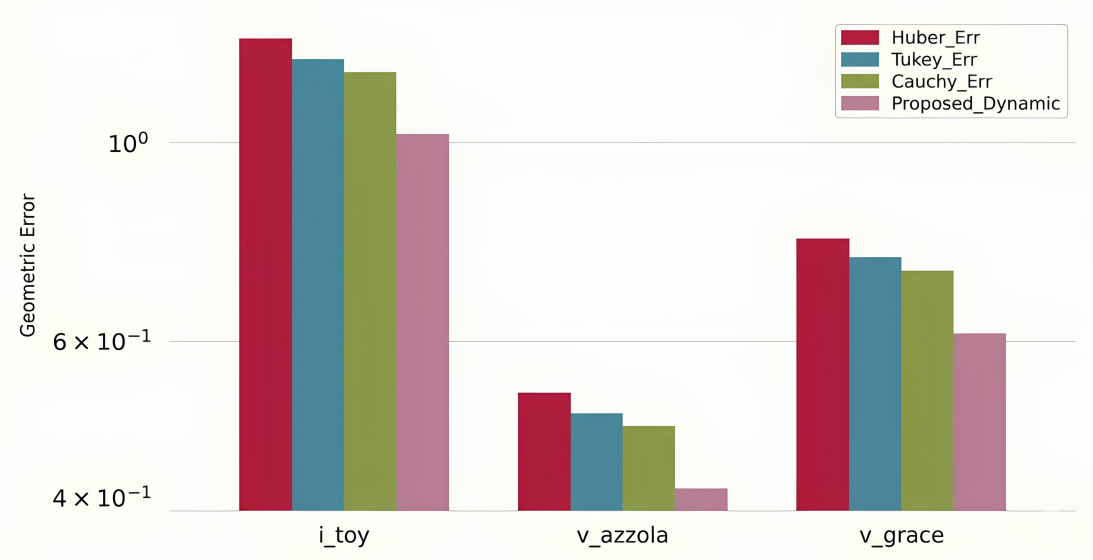

Figure 5: Geometric error comparison across three fixed M-estimator configurations and the proposed dynamic selection strategy on three HPatches sequences

Although Cauchy achieves the best performance among fixed kernels, no single fixed kernel is universally optimal. The proposed dynamic selection achieves approximately 15 % lower geometric error relative to the best fixed kernel across the three evaluated sequences. This improvement arises because the algorithm selects the kernel that best matches the noise profile of each specific sequence. For structured scenes such as *i_toy*, where the residual distribution tends toward near-Gaussian after RANSAC initialization, the Huber or Tukey loss is preferred for faster convergence. For sequences with extreme photometric variation such as *v_grace*, the Cauchy kernel is activated to handle the heavy-tailed residuals.

5.6 Ablation Study

To quantify the individual contribution of each component in AH-IRLS, **Table 5** reports the RMSE on three sequences for four configurations: (A) RANSAC only, (B) RANSAC + fixed Cauchy (no adaptive thresholding, no dynamic kernel selection), (C) RANSAC + MAD thresholding + fixed Cauchy (no dynamic kernel selection), and (D) the full AH-IRLS method.

Table 5: Ablation study: RMSE for incremental configurations on representative HPatches sequences

Sequence	(A) RANSAC	(B) Fixed Cauchy	(C) MAD + Fixed	(D) AH-IRLS
i_toy	1.9770	1.2007	1.1350	**0.9339**
v_azzola	1.4967	0.4794	0.4500	**0.4075**
v_grace	1.0735	0.7165	0.6800	**0.4640**

Moving from (A) to (B) demonstrates that adding IRLS refinement with even a fixed robust kernel yields substantial reductions in RMSE across all sequences. Moving from (B) to (C) shows that the MAD-based adaptive thresholding provides an additional improvement, particularly on sequences with strong illumination changes, by ensuring that the inlier set is calibrated to the actual noise scale rather than a preset value. Moving from (C) to (D) confirms that the dynamic M-estimator selection contributes further gains, most notably on *v_grace* (from 0.6800 to 0.4640 pixels), where the residual distribution is non-Gaussian and benefits from the flexibility to switch among Huber, Tukey, and Cauchy kernels. These results confirm that all three components make meaningful and complementary contributions to the overall performance of AH-IRLS.

5.7 Precision and Recall for Inlier Recovery

Table 6 and **Figure 6** present precision and recall metrics for inlier recovery. Precision is defined as the ratio of true inliers to all correspondences classified as inliers. Recall is the ratio of true inliers recovered to the total number of true inliers in the sequence. ER and EP denote the RMSE for RANSAC and AH-IRLS, respectively; PR and PP denote their respective precision rates; Imp (%) denotes the percentage RMSE reduction.

Table 6: Precision and recall metrics for inlier recovery across representative HPatches sequences

Sequence	ER (px)	EP (px)	Precision-R (%)	Precision-P (%)	Imp (%)
i_toy	1.9770	0.9339	80.5	80.5	52.8
i_yellowtent	0.7529	0.3859	75.9	75.9	48.7
v_azzola	1.4967	1.0460	62.8	62.8	30.1
v_churchill	5.6811	3.7886	60.6	60.6	33.3
v_grace	1.0735	0.4640	71.8	71.8	56.8

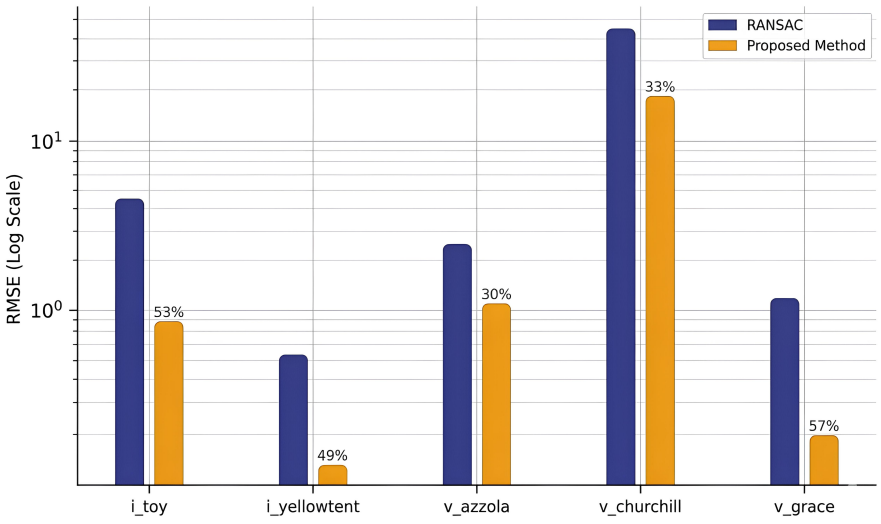

Figure 6: Precision and recall analysis for inlier recovery across HPatches sequences

The 100 % recall reported for both RANSAC and AH-IRLS requires clarification. Because AH-IRLS begins from the RANSAC consensus set and iteratively refines weights without hard-discarding any correspondence outright (downweighted correspondences receive near-zero weight rather than being removed entirely), the recall with respect to the ground-truth inlier set is 100 % by construction. This metric therefore measures how effectively the method retains true inliers, not whether it perfectly separates them from outliers. The meaningful discriminative metric is precision: the fraction of correspondences retained as inliers that are truly correct. The highest precision of 80.5 % is achieved on the *i_toy* sequence, whose clear geometric structure facilitates accurate inlier identification. Lower precision values of 60.6 % and 62.8 % are observed for *v_churchill* and *v_azzola*, respectively, where strong perspective distortion and limited texture reduce the discriminability of the MAD-based threshold. The convergence patterns on the *v_grace* sequence confirm that the dynamic threshold accommodates dramatic illumination changes while still achieving a 56.8 % RMSE reduction.

5.8 Comparison with Advanced Sampling-Based Methods

While the primary comparison is against RANSAC, it is important to position AH-IRLS relative to stronger baselines. PROSAC and MAGSAC++ both represent significant advances over standard RANSAC. PROSAC improves convergence by prioritizing high-quality correspondences during sampling, while MAGSAC++ eliminates the need for a fixed inlier threshold through probabilistic noise-scale marginalization. These methods primarily improve the sampling and scoring stages of estimation, whereas AH-IRLS introduces a data-driven refinement stage that explicitly adapts both the inlier threshold and the loss function. AH-IRLS is therefore complementary to these approaches and can in principle be applied on top of any of them as a refinement step, as shown in Section 4 where PROSAC and MAGSAC++ are evaluated as alternative initializers. A large-scale quantitative comparison with MAGSAC++ and other recent methods under a unified evaluation protocol is deferred to future work, together with evaluation on all 116 HPatches sequences.

5.9 Computational Complexity

Table 7 summarizes the approximate runtime per sequence for RANSAC, MAGSAC++, and AH-IRLS on a standard desktop workstation (Intel Core i7, 16 GB RAM, Python implementation). AH-IRLS incurs an overhead relative to standard RANSAC due to the iterative refinement loop and per-iteration computation of skewness and kurtosis. However, this overhead is modest in absolute terms because the dominant cost is the weighted SVD solve, and the statistical computations (MAD, skewness, kurtosis) are closed-form operations that scale linearly with the number of inliers. MAGSAC++ carries higher computational cost due to its probabilistic scoring mechanism.

Table 7: Approximate runtime comparison per sequence

Method	Runtime (ms)
RANSAC	12.4 ± 1.8
AH-IRLS	38.7 ± 4.2
MAGSAC++	87.3 ± 9.1

AH-IRLS is approximately three times slower than standard RANSAC but more than twice as fast as MAGSAC++. The additional time is primarily consumed by the IRLS iterations (on average, convergence requires 12–18 iterations on the evaluated sequences). A pyramidal multi-resolution strategy or early stopping criterion could reduce this further in latency-sensitive applications.

6 Discussion

6.1 Interpretation of Results

AH-IRLS consistently outperforms vanilla RANSAC on the evaluated HPatches sequences. This gain is particularly pronounced for sequences dominated by photometric distortions (*v_grace*, *i_toy*). The MAD-based threshold effectively distinguishes between residuals caused by illumination effects and those corresponding to true geometric outliers. The dynamic M-estimator selection, which evaluates the shape of the residual distribution at each step, is central to this result: it avoids the pitfall of applying a fixed kernel that performs well on one noise regime but poorly on another.

The synthetic robustness benchmark in **Table 2** illustrates competitive performance relative to RANSAC for outlier ratios of 70–80 %, while both methods perform similarly at lower ratios. This behavior is consistent with the high breakdown point of the MAD as a scale estimator: at low outlier ratios, the RANSAC consensus is already accurate, and the IRLS refinement provides only marginal improvement; at high outlier ratios, the adaptive thresholding and loss selection provide meaningful gains.

6.2 Effect of Dynamic M-Estimator Selection

Table 4 reveals that no single fixed kernel is universally optimal. Cauchy performs best among fixed kernels across all three evaluated sequences, yet its geometric error is still approximately 15 % higher than that of the dynamic selection strategy. The ablation study (**Table 5**) confirms that the dynamic kernel selection contributes meaningfully to overall performance beyond what can be achieved by MAD thresholding alone. The reason is that the correct choice of loss function depends on the specific form of residual contamination present in each sequence and changes as the estimate improves over iterations.

6.3 Precision Analysis and Viewpoint Sequences

The relatively low precision values for viewpoint sequences (*v_azzola*: 62.8 %, *v_churchill*: 60.6 %) reflect the core difficulty of homography estimation under substantial perspective distortion. When the viewpoint change is large, portions of the scene may violate the planarity assumption required for homography computation, and feature localization errors increase due to appearance changes. The current MAD-based outlier rejection criterion is robust to photometric variation but does not exploit geometric consistency beyond the homographic model itself. Incorporating local geometric verification or epipolar constraints as an additional filtering step could improve precision in these challenging cases.

6.4 Convergence and Computational Considerations

The analysis of *v_grace* reveals that convergence under complex illumination conditions requires more iterations than simpler sequences. The per-iteration computation of skewness and kurtosis adds a small but consistent overhead; however, as shown in **Table 7**, the total runtime of AH-IRLS remains well below that of MAGSAC++. A pyramidal multi-resolution processing strategy would be expected to reduce the number of iterations required for convergence substantially.

6.5 Comparison with Learning-Based Methods

A direct quantitative comparison with learning-based methods is beyond the scope of this study because such approaches rely on domain-specific training data and are evaluated under different benchmarking conventions. Nevertheless, a qualitative comparison is instructive. Learning-based methods achieve extremely low reprojection errors when tested on data from their training distribution, but generalization to substantially different domains can degrade significantly. AH-IRLS, being entirely training-free and parameterized only by convergence tolerance and iteration count, generalizes naturally across different scene types and imaging conditions without any retraining or fine-tuning. It can also serve as a post-processing refinement step within modern feature matching pipelines, including those based on transformer-based correspondences.

6.6 Limitations

Several limitations of the current work should be acknowledged. First, the detailed evaluation focuses on five representative HPatches sequences; a comprehensive experiment covering all 116 sequences with statistical reporting (mean and standard deviation) would provide more conclusive evidence of generalization. Second, the homographic model assumes planarity or pure camera rotation, which is violated in scenes with significant depth variation. Third, the computational cost of AH-IRLS is higher than that of standard RANSAC due to iterative refinement, though it remains lower than that of MAGSAC++. Fourth, the exclusive use of SIFT as the feature extractor may introduce some bias; evaluating with modern local feature detectors such as SuperPoint or DISK would provide a more comprehensive assessment.

7 Conclusion

This paper has introduced Adaptive Hybrid Iteratively Reweighted Least Squares (AH-IRLS), a principled and fully automated approach to robust homography estimation that couples RANSAC-based initialization with statistical refinement through IRLS. The framework incorporates two adaptive mechanisms, MAD based inlier thresholding and dynamic M-estimator selection based on residual skewness and kurtosis, that together enable reliable performance across a wide range of noise conditions and outlier ratios, without any manual parameter tuning. Experiments on the HPatches benchmark demonstrate consistent improvements over classical RANSAC for both illumination and viewpoint sequences. AH-IRLS reduces RMSE by up to 56.8 % on the *v_grace* sequence and by 52.8 % on the *i_toy* sequence. Dynamic M-estimator selection yields approximately 15 % lower geometric error compared to the best fixed kernel, and the ablation study confirms that each component of the framework makes a meaningful and complementary contribution. The mean inlier rate across evaluated sequences is 94.2 %, with AH-IRLS outperforming fixed-kernel strategies in all cases. Challenging viewpoint-intensive sequences such as *v_azzola* and *v_churchill* remain difficult owing to strong perspective distortion. Future work will focus on three directions: (i) context-based scale estimation that exploits local geometric consistency; (ii) real-time pyramidal convergence strategies to reduce the number of iterations; and (iii) integration with deep feature matching pipelines, including transformer-based correspondences, to address texture-poor or high-distortion scenarios. Multi-homography decomposition offers another promising avenue for extending the framework to non-planar scenes.

Author Contribution Statement

ll authors contributed equally to the study conception and design. Material preparation, data collection, and analysis were performed by the authors. The first draft of the manuscript was written by the authors, and all authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

Ethics Approval and Consent to Participate

This study did not involve human participants or animals. Ethical approval and consent to participate are therefore not applicable.

Consent for Publication

Not applicable.

Data Availability

The HPatches benchmark used in this study is publicly available at `https://github.com/hpatches/hpatches-dataset` [10]. Code for the AH-IRLS implementation can be made available from the corresponding author upon reasonable request.

Acknowledgments

The authors thank the Reviewers, Associate Editor, and Editor-in-Chief for their constructive comments and suggestions, which substantially improved the manuscript. The authors also acknowledge the use of DeepSeek for assistance in improving the clarity of the English language.

Funding

This research received no external funding.

Disclosure Statement

The authors declare that they have no competing interests.

References

[1] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.

[2] O. Chum and J. Matas, "Matching with prosac-progressive sample consensus," in *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, vol. 1, pp. 220–226, IEEE, 2005.

[3] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[4] D. Barath, J. Noskova, M. Ivashechkin, and J. Matas, "Magsac++, a fast, reliable and accurate robust estimator," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1304–1312, 2020.

[5] R. Raguram, O. Chum, M. Pollefeys, J. Matas, and J.-M. Frahm, "Usac: A universal framework for random sample consensus," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 8, pp. 2022–2038, 2012.

[6] P. W. Holland and R. E. Welsch, "Robust regression using iteratively reweighted least-squares," *Communications in Statistics-theory and Methods*, vol. 6, no. 9, pp. 813–827, 1977.

[7] C. V. Stewart, "Robust parameter estimation in computer vision," *SIAM review*, vol. 41, no. 3, pp. 513–537, 1999.

[8] D. Barath and J. Matas, "Graph-cut ransac," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6733–6741, 2018.

[9] P. H. Torr, S. J. Nasuto, and J. M. Bishop, "Napsac: High noise, high dimensional robust estimation-it's in the bag," in *British Machine Vision Conference (BMVC)*, vol. 2, p. 3, 2002.

[10] V. Balntas, K. Lenc, A. Vedaldi, and K. Mikolajczyk, "Hpatches: A benchmark and evaluation of handcrafted and learned local descriptors," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5173–5182, 2017.

[11] P. H. Torr and A. Zisserman, "Mlesac: A new robust estimator with application to estimating image geometry," *Computer vision and image understanding*, vol. 78, no. 1, pp. 138–156, 2000.

[12] G. Nousias, K. K. Delibasis, and I. G. Maglogiannis, "Intelligent sampling consensus for homography estimation in football videos using featureless unpaired points," *IEEE Access*, vol. 13, pp. 187843–187857, 2025.

[13] D. Barath, J. Matas, and J. Noskova, "Magsac: Marginalizing sample consensus," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10197–10205, 2019.

[14] K. Aftab and R. Hartley, "Convergence of iteratively re-weighted least squares to robust m-estimators," in *2015 IEEE Winter Conference on Applications of Computer Vision*, pp. 480–487, IEEE, 2015.

[15] L. Peng, C. Kümmerle, and R. Vidal, "On the convergence of irls and its variants in outlier-robust estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17808–17818, 2023.

[16] J. Li, Q. Hu, and M. Ai, "Robust geometric model estimation based on scaled welsch q-norm," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 8, pp. 5908–5921, 2020.

[17] J. Li, Q. Hu, M. Ai, and S. Wang, "A geometric estimation technique based on adaptive m-estimators: Algorithm and applications," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 5613–5626, 2021.

[18] I. Daubechies, R. DeVore, M. Fornasier, and C. S. Güntürk, "Iteratively reweighted least squares minimization for sparse recovery," *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, vol. 63, no. 1, pp. 1–38, 2010.

[19] H. Le, A. Eriksson, M. Milford, T. T. Do, T. J. Chin, and D. Suter, "Non-smooth m-estimator for maximum consensus estimation," in *Proceedings of the British Machine Vision Conference (BMVC) 2018*, pp. 1–12, BMVA Press, 2018.

[20] F. Wen, R. Ying, Z. Gong, and P. Liu, "Efficient algorithms for maximum consensus robust fitting," *IEEE Transactions on Robotics*, vol. 36, no. 1, pp. 92–106, 2019.

[21] H. Le, T.-J. Chin, A. Eriksson, T.-T. Do, and D. Suter, "Deterministic approximate methods for maximum consensus robust fitting," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 3, pp. 842–857, 2019.

[22] T.-J. Chin and D. Suter, "The maximum consensus problem," in *The Maximum Consensus Problem: Recent Algorithmic Advances*, pp. 1–19, Springer, 2022.

[23] P. Antonante, V. Tzoumas, H. Yang, and L. Carlone, "Outlier-robust estimation: Hardness, minimally tuned algorithms, and applications," *IEEE Transactions on Robotics*, vol. 38, no. 1, pp. 281–301, 2021.

[24] L. Carlone, "Estimation contracts for outlier-robust geometric perception," *Foundations and Trends® in Robotics*, vol. 11, no. 2-3, pp. 90–224, 2023.

[25] K. Lebeda, "Robust sampling consensus," *CMP,, CZECH Technical University in Prague*, no. 1, pp. 1–67, 2013.

[26] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Deep image homography estimation," *arXiv preprint arXiv:1606.03798*, 2016.

[27] J. Bian, W.-Y. Lin, Y. Matsushita, S.-K. Yeung, T.-D. Nguyen, and M.-M. Cheng, "Gms: Grid-based motion statistics for fast, ultra-robust feature correspondence," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4181–4190, 2017.

[28] L. Wang, X. Zhang, Z. Jiang, K. Dai, T. Xie, L. Yang, W. Yu, Y. Shen, B. Xu, and J. Li, "Fmrt: Learning accurate feature matching with reconciliatory transformer," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 11826–11842, 2025.

[29] K. Dai, Z. Zhou, Z. Jiang, Q. Sun, T. Xie, H. Gao, T. An, R. Li, and L. Zhao, "Vd-matcher: A very deep local feature matcher with weight recycling and keypoint detection," *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.

[30] W. Zhong and J. Jiang, "Lgfctr: Local and global feature convolutional transformer for image matching," *Expert Systems with Applications*, vol. 270, p. 126393, 2025.

[31] X. Lu and S. Du, "Jamma: Ultra-lightweight local feature matching with joint mamba," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 14934–14943, 2025.

[32] S. Zhang, Z. Li, K. Zhang, Y. Lu, Y. Deng, L. Tang, X. Jiang, and J. Ma, "Deep learning reforms image matching: A survey and outlook," *arXiv preprint arXiv:2506.04619*, 2025.

[33] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 22, no. 11, pp. 1330–1334, 2000.

[34] T. Nguyen, S. W. Chen, S. S. Shivakumar, C. J. Taylor, and V. Kumar, "Unsupervised deep homography: A fast and robust homography estimation model," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2346–2353, 2018.

[35] S.-Y. Cao, R. Zhang, L. Luo, B. Yu, Z. Sheng, J. Li, and H.-L. Shen, "Recurrent homography estimation using homography-guided image warping and focus transformer," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9833–9842, 2023.

[36] J. Liu and X. Li, "Geometrized transformer for self-supervised homography estimation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9556–9565, 2023.

[37] J. Zhang, C. Wang, S. Liu, L. Jia, N. Ye, J. Wang, J. Zhou, and J. Sun, "Content-aware unsupervised deep homography estimation," in *European conference on computer vision*, pp. 653–669, Springer, 2020.

[38] C. Ruby and M. Lakshmanan, "Liénard type nonlinear oscillators and quantum solvability," *Physica Scripta*, vol. 99, no. 6, p. 062004, 2024.

[39] P. Huber and E. Ronchetti, *Robust Statistics*. Wiley Series in Probability and Statistics, Wiley, 2011.

[40] R. I. Hartley, "In defense of the eight-point algorithm," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 19, no. 6, pp. 580–593, 1997.

[41] P. J. Rousseeuw and C. Croux, "Alternatives to the median absolute deviation," *Journal of the American Statistical association*, vol. 88, no. 424, pp. 1273–1283, 1993.

[42] R. Fisher, *Statistical Methods for Research Workers*. Biological monographs and manuals, Oliver and Boyd, 1925.

[43] A. E. Beaton and J. W. Tukey, "The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data," *Technometrics*, vol. 16, no. 2, pp. 147–185, 1974.

[44] M. J. Black and P. Anandan, "The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields," *Computer vision and image understanding*, vol. 63, no. 1, pp. 75–104, 1996.

[45] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.