

PMM Amplitude Estimation for SDR Baseband Under Non-Gaussian Radio-Frequency Noise

Zabolotnii S. V.^{1,2,3} , Koval V. V.² , Chepynoha A. V.⁴ 

¹Cherkasy State Business College, Cherkasy 18028, Ukraine

²State Scientific Research Institute of Armament and Military Equipment Testing and Certification, Cherkasy, Ukraine

³Uzhhorod National University, Uzhhorod, Ukraine

⁴Cherkasy State Technological University, Cherkasy, Ukraine

E-mail: zabolotnii.serhii@csbc.edu.ua

Amplitude estimation from Software-Defined Radio (SDR) baseband recordings becomes unreliable when the additive Radio-Frequency (RF) noise departs from the Gaussian model assumed by Ordinary Least Squares (OLS). Rather than comparing estimators one by one, this paper treats OLS, the Polynomial Maximization Method of orders two and three (PMM2, PMM3) and the fractional branch of the Parametrically Adaptive Transition Polynomial (PATP) as members of a single continuously parameterised family, and states the problem as the synthesis of a decision rule that selects a member from the cumulants of one finite block, minimising the regret with respect to an oracle choice. The protocol combines synthetic polygaussian noise, real RF residuals from RadioML2018 over-the-air recordings, controlled heavy-tailed and platykurtic regimes, and leptokurtic blocks from two drone recordings. Per-block PMM2 is unstable at $N = 1024$ because the skewness estimate has a low signal-to-noise ratio; pooling the cumulants over calibration blocks returns it to the OLS baseline. On real platykurtic RF noise PMM3 lowers the mean squared error from 3.28×10^{-3} to 2.76×10^{-3} , and Newey–West HAC standard errors raise the empirical 95% coverage of all three point estimators from about 0.72 to 0.920–0.924 — PMM3 included (0.922) — which is still about three percentage points below nominal. In the heavy-tailed regime relative efficiency is reported as a ratio of medians of squared errors with paired bootstrap intervals, because the mean-based figure does not converge under infinite variance: at $\alpha_s=1.2$ the oracle bound PATP* reaches 38.8 [33.8; 44.5] while the practically achievable adaptive rule AUTO_EXT reaches 6.41 [4.96; 8.26]. The family is two-sided — it gains on both non-Gaussian edges and returns exactly OLS on the Gaussian controls — unlike the one-sided Huber M-estimator, which loses 2–12% on platykurtic noise. Finally, we quantify what the dispatch rule costs: against a per-block oracle it selects a suboptimal member in 65.0% of the synthetic and 69.9% of the real-noise blocks, with regret 16.19 and 1.220 respectively, and on the synthetic set it is $10.7\times$ worse than plain OLS. The resulting operational recommendation is therefore conditional, not unconditional dispatch: PMM3 for nearly symmetric platykurtic RF residuals, pooled calibration for PMM2, the fractional branch for heavy tails, and OLS whenever the block cumulants do not clearly indicate otherwise.

Keywords: PMM; PATP; amplitude estimation; SDR; non-Gaussian noise; heavy-tailed noise; fractional moments; cumulants; HAC standard errors; estimator selection; regret; median-based relative efficiency; RadioML2018

DOI: [10.64915/RADAP.2026.105.39-51](https://doi.org/10.64915/RADAP.2026.105.39-51)

Introduction

Amplitude estimation from Software-Defined Radio (SDR) baseband recordings is one of the elementary operations of radio signal processing. It enters receiver-chain calibration, channel-coefficient estimation, feature normalization prior to detection or classification, and the broader pipelines of radio monitoring and spectrum analysis. In the simplest formulation the observed signal is written as the linear model

$$y[n] = As[n] + \xi[n], \quad n = 0, \dots, N-1, \quad (1)$$

where A is the unknown amplitude, $s[n]$ a known template, and $\xi[n]$ an additive noise term. For independent Gaussian noise the natural and computationally convenient solution is Ordinary Least Squares (OLS), which in this case coincides with the maximum-likelihood estimator and attains minimum variance within the classical model.

In real Radio-Frequency (RF) recordings, however, the Gaussian assumption is often only a technical idealization. RF residuals may be shaped by co-channel interference, impulsive disturbances (including the Middleton Class A/B models [1]) [2], hardware nonli-

nearities, residual synchronization errors, receiver-chain filtering, and the over-the-air channel. Their distribution may therefore be asymmetric, heavy-tailed (as in the α -stable or impulsive models that are also observed empirically in modern wireless systems [3, 4]) or, conversely, platykurtic; the residual samples are moreover frequently autocorrelated in time. For such data OLS usually remains unbiased under zero-mean noise but loses efficiency, and the standard uncertainty estimates may yield confidence intervals that are too narrow. It is therefore important, in the SDR setting, to assess not only the point error of the amplitude estimate but also the reliability of the statistical inference.

Existing approaches to non-Gaussian noise fall broadly into three groups. The first imposes an explicit probability model on the noise — α -stable, Middleton-type, exponential-power, or mixture distributions. This route supports parametric maximum-likelihood estimators but requires correct specification of the probability density and estimation of additional nuisance parameters. The second group relies on robust criteria such as LAD, L_p regression, M-estimators [5, 6], and sign- or quantile-based methods; a survey of these techniques in signal processing is given in [7]. They limit the influence of outliers and heavy tails but often do not explicitly exploit the distributional-shape information carried by higher-order cumulants. The third route is intermediate: rather than the full density, it describes non-Gaussianity by a finite set of moments or cumulants and uses them to build statistically more efficient estimators; L-moments [8] are a related example.

The Polynomial Maximization Method (PMM), introduced by Kunchenko [9], belongs to this third line. In spirit it is close to maximum likelihood, since it also maximizes a sample statistic in a neighbourhood of the true parameter. The essential difference is that PMM does not require the full noise density, but uses a finite moment–cumulant description instead. Earlier work showed that second-order PMM (PMM2) reduces the variance of linear-regression estimates under asymmetric non-Gaussian errors [10], with the gain governed by skewness γ_3 and excess kurtosis γ_4 ; the method has been extended to polygaussian mixtures and to autoregressive processes with non-Gaussian innovations. At the same time, PMM2 degenerates to OLS for symmetric distributions, so a third-order polynomial is needed for nearly symmetric non-Gaussian errors.

Third-order PMM (PMM3) was developed for symmetric non-Gaussian errors [11] and uses fourth- and sixth-order information. Its theoretical variance reduction is governed by the cumulants γ_4 and γ_6 . The platykurtic regime is especially relevant to the SDR context: it is precisely where a cubic score has room to act, because a flat-topped density carries fourth- and sixth-order structure that a quadratic score cannot exploit. Section 5.2 quantifies this directly on uniform,

Beta(2,2) and generalized-Gaussian noise. This makes PMM3 a natural candidate for RF residuals that are nearly symmetric but depart from the Gaussian in their kurtosis.

The integer-order moment apparatus of PMM2/PMM3 is a special case of a wider, continuously parameterized family of stochastic polynomials — the Parametrically Adaptive Transition Polynomial (PATP) [12]. In this family the basis $\varphi_i(\xi; \alpha) = \text{sign}(\xi) |\xi|^{p_i(\alpha)}$ continuously interpolates between a fractional-moment regime ($\alpha \rightarrow 0$), suited to heavy-tailed impulsive noise where higher-order integer moments diverge, and a classical power regime ($\alpha = 1$) whose integer representatives are PMM2 and PMM3. In particular, the second-order theoretical variance-reduction coefficient $g_2 = 1 - \gamma_3^2/(2 + \gamma_4)$ corresponds to $g_2(\alpha)$ at $\alpha = 1$. In this paper we treat both wings of the family: the finite-moment platykurtic regime ($\alpha = 1$, the cubic PMM3 track) and the fractional branch, which extends the method to the heavy-tailed impulsive RF noise where the integer tracks no longer apply.

Despite the existing PMM theory for regression models, its transfer to SDR amplitude estimation is not automatic. First, SDR windows are of finite length, and at $N = 1024$ the cumulant estimates may have substantial variance. This is especially critical for PMM2: if the true skewness is small relative to the standard error of $\hat{\gamma}_3$, the per-block PMM weights become noise-driven and may increase rather than reduce MSE. Second, real RF residuals are not necessarily independent; their autocorrelation affects the standard errors and may lead to under-coverage of confidence intervals. For practical SDR use one must therefore verify not only the theoretical variance-reduction coefficient, but also the stability of cumulant estimation and the quality of the interval inference.

We study PMM in exactly this practical setting, and we state the problem as one of estimator selection inside a single family rather than as a pairwise comparison of methods. The contribution is threefold. First, a formal problem statement (Section 1) and an explicit selection rule with N -scaled, cumulant-standard-error-based thresholds (Algorithm 1). Second, a quantitative characterisation of that rule — misclassification rate and regret against a per-block oracle, a sweep of the threshold scale, and the two regimes in which the rule fails — which turns “adaptive dispatch” from a design claim into a measured quantity. Third, a well-posed measure of efficiency in the infinite-variance regime, together with the interval inference (Newey–West HAC) for the estimator we actually recommend. The paper does not claim that PMM is universally better than OLS; it states the conditions under which each family member is safe to use, and the price of choosing between them automatically.

1 Problem Statement

Model. A baseband SDR block of length N is observed according to (1), where $\xi[n]$ is additive noise with zero mean and unknown distribution F_ξ , in general non-Gaussian and possibly autocorrelated in time.

Class of estimators. We consider the class of M-estimators generated by the PATP stochastic polynomials: the estimate $\hat{A}$ solves

$$\sum_n s[n] \psi(y[n] - \hat{A} s[n]) = 0, \quad (2)$$

where the score function is built on the basis $\varphi_i(\varepsilon; \alpha) = \text{sign}(\varepsilon) |\varepsilon|^{p_i(\alpha)}$ with the exponent map $p_i(\alpha)$ of (11). The parameter $\alpha \in [0, 1]$ continuously connects three regimes: fractional ($\alpha = 0$), degenerate linear ($\alpha = 1/2$, identically OLS) and power ($\alpha = 1$), whose integer representatives are PMM2 and PMM3. Hence OLS, PMM2, PMM3 and PMM_FRAC are not competing methods but *members of one parameterised family* $\mathcal{M} = \{\text{OLS}, \text{PMM2}, \text{PMM3}, \text{PMM_FRAC}(\alpha)\}$.

Criterion. For a member $m \in \mathcal{M}$ define the relative efficiency $RE(m; F_\xi) = \text{MSE}_{\text{OLS}}/\text{MSE}_m$. The oracle choice $m^*(F_\xi) = \arg \max_m RE(m; F_\xi)$ requires knowledge of F_ξ — unavailable in practice — and, in a Monte Carlo study, knowledge of the true amplitude A .

The problem addressed here. The oracle is unattainable because the noise distribution is unknown and must be estimated from a finite block: at $N = 1024$ the standard errors of the sample cumulants are $\text{SE}(\hat{\gamma}_3) = \sqrt{6/N} \approx 0.077$ and $\text{SE}(\hat{\gamma}_4) = \sqrt{24/N} \approx 0.153$, so the estimates $\hat{\gamma}_3, \hat{\gamma}_4$ that any practical rule must use are themselves noisy. The problem is therefore stated as follows: construct a decision rule $\delta : (\hat{\gamma}_3, \hat{\gamma}_4) \mapsto \mathcal{M}$, acting only on the cumulant estimates of the current block, that minimises the expected regret

$$R(\delta) = \frac{\text{MSE}(\delta)}{\text{MSE}(m^*)} \quad (3)$$

subject to a *no-harm target* $RE(\delta) \geq 1$, i.e. the rule should not lose to OLS in any noise regime. Whether a given rule meets this target is an empirical question, and Section 5.3 shows that the rule studied here meets it on real-noise data and violates it on the synthetic polygaussian set.

This formulation makes the subject of the paper the synthesis and quantitative characterisation of the rule δ rather than a comparison of known methods, together with two accompanying finite-block questions: (a) whether pooled cumulant calibration stabilises the PMM2 weights at $N = 1024$; and (b) whether the Newey–West HAC correction is required for reliable interval inference under autocorrelated RF residuals.

2 PMM Estimators

The OLS estimate of the amplitude in (1) is

$$\hat{A}_{\text{OLS}} = \frac{\mathbf{y}^\top \mathbf{s}}{\mathbf{s}^\top \mathbf{s}}. \quad (4)$$

OLS remains unbiased under zero-mean non-Gaussian noise, but it is not generally efficient.

For the asymmetric track, PMM2 uses the quadratic score

$$\psi(\varepsilon) = c_1 \varepsilon + c_2 (\varepsilon^2 - m_2), \quad (5)$$

where m_2 is the noise variance. The optimal weights are determined by skewness γ_3 and excess kurtosis γ_4 :

$$c_1 = \frac{2 + \gamma_4}{2 + \gamma_4 - \gamma_3^2}, \quad c_2 = \frac{-\gamma_3}{2 + \gamma_4 - \gamma_3^2}. \quad (6)$$

The corresponding theoretical variance-reduction coefficient is

$$g_2 = 1 - \frac{\gamma_3^2}{2 + \gamma_4}. \quad (7)$$

Thus PMM2 can help only when the asymmetry is strong enough to be estimated reliably.

For symmetric platykurtic noise, PMM3 uses the cubic score

$$\psi(\varepsilon) = \varepsilon(\kappa - \varepsilon^2), \quad (8)$$

with

$$\kappa = \frac{m_6 - 3m_4m_2}{m_4 - 3m_2^2}. \quad (9)$$

The theoretical coefficient for this track is

$$g_3 = 1 - \frac{\gamma_4^2}{6 + 9\gamma_4 + \gamma_6}. \quad (10)$$

PMM3 is activated for nearly symmetric platykurtic residuals. Which family member is actually applied to a given block is decided by the rule stated in Algorithm 1 (Section 3), whose thresholds are multiples of the standard errors of the cumulant estimates.

The PMM2/PMM3 score functions are the integer (power) representatives of the wider, continuously parameterized PATP family [12], in which the basis $\varphi_i(\varepsilon; \alpha) = \text{sign}(\varepsilon) |\varepsilon|^{p_i(\alpha)}$ is governed by an exponent $\alpha \in [0, 1]$ through the map

$$p_i(\alpha) = \frac{1}{i} + \left(4 - i - \frac{3}{i}\right)\alpha + \left(2i - 4 + \frac{2}{i}\right)\alpha^2, \quad (11)$$

which interpolates three regimes: fractional ($\alpha = 0 \Rightarrow p_i = 1/i$, low-order fractional moments for heavy tails in the spirit of FLOM [13]), degenerate linear ($\alpha = 1/2 \Rightarrow p_i = 1$, coinciding with OLS), and power ($\alpha = 1 \Rightarrow p_i = i$). For $S = 2$ the leading fractional term is realized as the score $\text{sign}(\varepsilon) |\varepsilon|^{p_2(\alpha)}$, and the classical coefficient g_2 corresponds to the value at $\alpha = 1$. This yields a single estimator that continuously spans the range from heavy-tailed impulsive noise (where integer moments diverge) to the platykurtic regime.

3 Methodology and the Selection Rule

3.1 Two dispatch rules

Two dispatch rules appear in this paper and must not be confused. The *base rule*, used for the AUTO column of Table 2 and for the audit of Section 5.3, is the dispatcher shipped with the PMM R package (see Code availability): it routes to PMM2 when $|\hat{\gamma}_3| > 1$, or when $|\hat{\gamma}_3| > 0.3$ and the theoretical factor $g_2 = 1 - \hat{\gamma}_3^2/(2 + \hat{\gamma}_4)$ is below 0.95, to PMM3 when $|\hat{\gamma}_3| \leq 0.3$ and $\hat{\gamma}_4 < -0.7$, and to OLS otherwise. The *extended rule* AUTO_EXT, used for the heavy-tailed and real-RF experiments of Sections 5.2 and 5.4, adds a guard against divergent moments and a fractional branch; it is stated in full as Algorithm 1. The two rules coincide on Gaussian and platykurtic blocks and differ where the sample kurtosis is extreme.

3.2 Algorithm 1: AUTO_EXT

Algorithm 1. AUTO_EXT: selection of a family member from the block cumulants.

Input: noise block $x[1..N]$; pooled exponent α_{pool} (or NONE); threshold scales k_3, k_4 ; divergence bound D ; heavy-tail threshold F ; marginal-gain threshold G (default 0.95).

Output: member $m \in \{\text{OLS}, \text{PMM2}, \text{PMM3}, \text{PMM_FRAC}(\alpha)\}$

```

1  if  $N < 2$ : return OLS                                (cumulants undefined)
2   $\hat{\gamma}_3 \leftarrow \text{skew}(x)$ ;  $\hat{\gamma}_4 \leftarrow \text{kurt}(x)$           (Fisher, bias-corrected)
3   $\tau_3 \leftarrow k_3\sqrt{6/N}$ ;  $\tau_4 \leftarrow k_4\sqrt{24/N}$       (N-scaled gates)
4  if  $\hat{\gamma}_4$  not finite or  $\hat{\gamma}_4 > D$ : (Level 0: divergent moments)
5    return PMM_FRAC( $\alpha_{\text{pool}}$ ) if  $\alpha_{\text{pool}} \neq \text{NONE}$ , else OLS
6  if  $|\hat{\gamma}_3| > \tau_3$ : (Level 1: asymmetric)
7     $g_2 \leftarrow 1 - \hat{\gamma}_3^2/(2 + \hat{\gamma}_4)$ 
8    return OLS if  $g_2 \geq G$  (predicted gain  $< 1 - G$ )
9    return PMM2
10 if  $\hat{\gamma}_4 < -\tau_4$ : return PMM3 (Level 2: platykurtic)
11 if  $\hat{\gamma}_4 > F$ : (Level 3: heavy-tailed)
12   return PMM_FRAC( $\alpha_{\text{pool}}$ ) if  $\alpha_{\text{pool}} \neq \text{NONE}$ , else OLS
13 return OLS (Level 4: intermediate / Gaussian)

```

Threshold rationale. The thresholds are not free constants but multiples of the standard error of the corresponding cumulant, so that they scale automatically with the block length: $\tau_3 = k_3\sqrt{6/N}$ and $\tau_4 = k_4\sqrt{24/N}$. The values 0.3 and -0.7 used here correspond to $k_3 = 3.92$ and $k_4 = 4.57$ at $N = 1024$, i.e. roughly four standard errors: a branch is activated only when the departure from Gaussianity is reliably distinguishable from the estimation error. The optimality of k_3, k_4 is verified empirically by a sweep (Section 5.3) rather than postulated.

The marginal-gain guard. Step 8 routes a block back to OLS when the predicted variance reduction $1 - g_2$ is below $1 - G = 5\%$. Measured offline over the 37500 blocks of the fifteen Monte Carlo sets used in this paper, the guard reroutes 5.46% of blocks and changes the estimate by less than 10^{-15} in relative terms, because at $g_2 \geq 0.95$ the PMM2 and OLS estimates are numerically

indistinguishable. Its role is to remove the risk of applying noisy cumulant weights where there is no gain to be had, not to improve accuracy.

Safety by construction is not a guarantee.

Every branch that cannot compute reliable weights — too short a block, divergent moments with no pooled exponent available, or an intermediate regime — returns OLS. This makes OLS the default of the rule by construction, but it does not by itself guarantee $RE \geq 1$: a branch that fires may still select a member that is worse than OLS on that particular block. Section 5.3 measures how often this happens.

Idealised input. In all experiments the rule is evaluated on the noise block itself rather than on OLS residuals. For the one-regressor model (1) the difference is a projection on a single direction, i.e. a relative shift of order $O(1/N)$, but the reported figures are therefore those of a dispatcher with ideal access to the noise; a deployed pipeline would compute the cumulants from $\hat{\varepsilon} = y - \hat{A}s$.

3.3 Monte Carlo protocol and uncertainty quantification

All PMM computations were performed through the R package EstemPMM v0.4.0 (see Code availability); the fractional PATP branch is implemented by a separate IRLS solver. Each Monte Carlo trial used one signal block of length $N = 1024$. The true amplitude was sampled from Uniform(0.5, 2.0), and the same realization was passed to all estimators, producing paired comparisons.

Six estimators were evaluated: OLS, per-block PMM2, per-block PMM3, automatic dispatch, PMM2 with pooled cumulants, and PMM3 with pooled cumulants. The pooled variants estimate the distributional parameters from 200 calibration blocks of the same SNR bucket and then keep the PMM weights fixed for the evaluated block. This reduces the standard error of the cumulant estimates by approximately a factor of $\sqrt{200}$.

Definition of relative efficiency. RE is reported in two forms. Where all moments are finite (Sections 5.1 and 5.2, platykurtic regime) $RE = \text{MSE}_{\text{OLS}}/\text{MSE}$ is used. Where the noise has infinite variance the sample MSE does not converge, and RE is defined as the ratio of medians of squared errors, $\text{med}(e_{\text{OLS}}^2)/\text{med}(e^2)$, with a paired percentile bootstrap over noise blocks (2000 resamples, the same block indices applied to numerator and denominator, so that the common-random-number pairing is preserved). Replacing a mean-based dispersion measure by a median-based one is the standard remedy when the underlying moment need not exist [14]. A 1% upper-trimmed-mean variant is reported in Table S5 of the supplementary material and preserves the ordering of the methods.

Silent fallback to OLS. When the iterative solver fails to converge, the implementation falls back to the

OLS estimate for that block. This is reported explicitly wherever it matters: on α -stable noise with $\alpha_s=1.2$ the fractional solver converges on 38.3% of blocks, so on the remaining 61.7% AUTO_EXT is literally OLS.

Standard errors. For uncertainty quantification, the sandwich variance of an M-estimator is

$$\sigma^2(\hat{A}) = \frac{\sum_i \psi^2(\hat{\varepsilon}_i) s_i^2}{(\sum_i \psi'(\hat{\varepsilon}_i) s_i^2)^2}. \quad (12)$$

When the residuals are autocorrelated, the numerator is replaced by a Newey–West HAC estimate with the Bartlett kernel:

$$\sigma_{\text{HAC}}^2(\hat{A}) = \frac{\sum_{j=-L}^L \left(1 - \frac{|j|}{L+1}\right) \sum_t z_t z_{t-j}}{(\sum_t s_t^2)^2}, \quad (13)$$

where $z_t = s_t \psi(\hat{\varepsilon}_t)$ for the PMM tracks and $z_t = s_t \hat{\varepsilon}_t$ for OLS [15, 16]. For the PMM3 score (8) the derivative entering (12)–(13) is $\psi'(\varepsilon) = \kappa - 3\varepsilon^2$. The lag parameter was set to $L = \lfloor 4(N/100)^{2/9} \rfloor$, which gives $L = 6$ for $N = 1024$; data-driven alternatives to this fixed rule are available [17], and the finite-sample behaviour of kernel HAC estimators, including their downward bias and the resulting under-coverage at moderate N , is analysed in [18, 19].

4 Datasets

The datasets used cover qualitatively different noise regimes. The first two are the basis for PMM2/PMM3: a synthetic polygaussian set with leptokurtic mixture noise (which probes the PMM2 branch) and a realnoise set built from RadioML2018 over-the-air recordings [20] by subtracting the clean reference signal from the noisy recording (which preserves the empirical platykurtic RF residuals without a parametric model). To test both wings of the PATP family we add controlled synthetic regimes: heavy-tailed α -stable ($\alpha_s \in \{1.2, 1.5, 1.8, 2.0\}$) and Student-t ($\nu \in \{3, 5, 10, \infty\}$), and platykurtic uniform / Beta(2,2) / generalized-Gaussian ($\beta=4$). The real impulsive RF consists of leptokurtic blocks ($\hat{\gamma}_4 > 3$, $|\hat{\gamma}_3| < 0.5$) extracted from over-the-air I/Q recordings

of two independent drones (DJI Mavic and Parrot Disco, OcuSync, Tampere).

Table 1 Dataset comparison

Dataset	Blocks	$\hat{\gamma}_3$	
Synthetic polygauss	100k	$+0.08 \pm 0.66$	
RadioML2018 realnoise	100k	$+0.00 \pm 0.09$	
α -stable $\alpha_s=1.2-1.8$	8k	≈ 0	
Student-t $\nu=3, 5, 10$	8k	≈ 0	
unif. / Beta(2,2) / GG($\beta=4$)	8k	≈ 0	$-1.2,$
$\alpha_s=2.0$ / $\nu=\infty$	8k	≈ 0	
DJI Mavic / Parrot Disco (real)	661 / 693	$+0.0 \pm 0.2$	

[†]The two drone sets come from one over-the-air recording per drone ($n_{\text{rec}} = 1$). Their blocks are therefore not independent — they share a channel, a receiver and a capture session — so the drone results are reported descriptively, with no p-values and no claim of generalisation across sources or receivers (Table S1 of the supplementary material).

The synthetic polygaussian residuals are leptokurtic; the RadioML residuals are nearly symmetric and platykurtic (the primary test of PMM3). The α -stable and Student-t sets define a controlled heavy-tailed regime (the fractional branch), the uniform/Beta/GG sets a platykurtic one (the power branch of PMM3), and the real drone blocks an empirical impulsive RF case; $\alpha_s=2.0$ and $\nu=\infty$ are Gaussian controls used to verify the absence of degradation.

5 Experimental Results

5.1 Finite-block behaviour of PMM2 and PMM3

Table 2 summarizes the main MSE and coverage results. On synthetic polygaussian noise, OLS remains the best or tied-best estimator. Per-block PMM2 degrades strongly because $\hat{\gamma}_3$ is too noisy at $N = 1024$. This is not a contradiction of PMM theory; it is a finite-block cumulant-estimation failure. Once the cumulants are pooled over calibration blocks, PMM2 recovers to OLS-level MSE.

Table 2 MSE and empirical 95% coverage. HAC coverage uses Newey–West/Bartlett standard errors with $L = \lfloor 4(N/100)^{2/9} \rfloor = 6$ at $N = 1024$; “—” means no HAC variant is computed for that estimator. AUTO is the base dispatch rule of Section 3 (not AUTO_EXT)

Method	Synthetic			Realnoise		
	MSE	Cov.	Cov. (HAC)	MSE	Cov.	Cov. (HAC)
OLS	6.77×10^{-4}	0.942	0.984	3.28×10^{-3}	0.724	0.924
PMM2	1.88×10^{-2}	0.914	0.924	3.51×10^{-3}	0.721	0.920
PMM3	7.54×10^{-4}	0.970	0.981	2.76×10^{-3}	0.725	0.922
AUTO (base rule)	7.27×10^{-3}	0.928	—	2.96×10^{-3}	0.724	—
PMM2_POOL	6.77×10^{-4}	0.972	—	3.28×10^{-3}	0.724	—
PMM3_POOL	7.20×10^{-4}	0.974	—	2.92×10^{-3}	0.729	—

Monte-Carlo standard error of a coverage estimate at $n = 2500$ is ≈ 0.005 ; the 95% interval for the PMM3 HAC coverage on realnoise is $[0.912, 0.933]$. The 2500 blocks come from a single real-noise dataset, so these are within-dataset Monte-Carlo statements about coverage, not statements about generalisation across recordings. On the synthetic set the residuals are essentially uncorrelated, so the HAC correction is unnecessary and over-covers (0.981–0.984 against a nominal 0.95): the correction is needed where autocorrelation is present, and is merely conservative where it is not.




Fig. 1. MSE versus SNR for amplitude estimators on the synthetic polygaussian dataset. Per-block PMM2 is unstable, whereas pooled PMM2 recovers to the OLS level

The AUTO column of Table 2 is the base dispatcher of Section 3; its 7.27×10^{-3} on the synthetic set is an order of magnitude above OLS, which Section 5.3 traces to the same per-block PMM2 instability.

On real RF residuals, PMM3 gives the lowest MSE: 2.76×10^{-3} versus 3.28×10^{-3} for OLS. On the platykurtic subset ($\hat{\gamma}_4 < -0.7$, $|\hat{\gamma}_3| < 0.3$, $n = 551$),

the relative efficiency reaches 1.74. This confirms that PMM3 is useful when the residual distribution matches its symmetric platykurtic operating region. The same threshold, however, is shown in Section 5.3 to be too strict for these data: the regret-minimising gate is wider, $\hat{\gamma}_4 < -0.50$ instead of -0.70 .




Fig. 2. Relative efficiency of PMM3 versus estimated excess kurtosis on the realnoise dataset. PMM3 is beneficial in the platykurtic region and should not be used as a universal replacement for OLS

5.2 The family is two-sided

We now treat PMM/PATP not as separate tracks but as a single family spanning the full range of non-Gaussian noise, and we compare it with a standard robust baseline — the Huber M-estimator [5] ($c = 1.345$). All runs are Monte Carlo with 2500 realizations per dataset.

Because the α -stable law with $\alpha_s < 2$ has infinite variance, the sample MSE of OLS is not a convergent quantity: at $\alpha_s=1.2$ a single replicate out of 2500 carries 68% of the OLS sum of squared errors, and the mean-based RE ranges from 2.2×10^3 to 4.3×10^4 across five nominally identical Monte Carlo cells. We therefore use the median-based RE of Section 3. Under this well-defined functional the heavy-tail gains are large but finite: at $\alpha_s=1.2$ the oracle bound PATP* reaches $RE = 38.8$ [33.8; 44.5] and the Huber baseline 74.1 [64.0; 87.4], whereas the practically achievable adaptive rule AUTO_EXT — which selects α from pooled cumulants rather than from the true amplitude — reaches 6.41 [4.96; 8.26]. PATP* must be read as an oracle upper bound: α^* is chosen post hoc as the argmax of RE over the sweep grid, so it bounds what the fractional branch can deliver, not what it delivers in deployment, and the remaining gap to AUTO_EXT — a factor of about six at $\alpha_s=1.2$ and below two for $\alpha_s \geq 1.8$ and for Student-t — is the price of estimating α from data. Under the same functional the Huber M-estimator remains the strongest single method in the heavy tails, which is expected:

any sublinear score beats OLS almost without bound on infinite-variance noise.

The distinction is on the other side. On the *platykurtic* side Huber is significantly below unity on all three sets (0.88–0.98, upper confidence limits below 1), because its robust down-weighting attenuates large residuals exactly where short tails call for the opposite, whereas the PMM/PATP family gains through the power term of PMM3 (up to 2.96 [2.62; 3.41] on the uniform set) and returns exactly OLS on both Gaussian controls.

Within the fractional branch the gain rises monotonically with tail weight and α^* shifts from the linear edge ($\alpha=1/2$, Gaussian) towards the fractional edge ($\alpha \rightarrow 0$); optimising α outperforms the fixed cumulant-based PMM_FRAC (6.26 [5.41; 7.05] versus 2.29 [2.02; 2.53] at $\alpha_s=1.5$), at the cost of being an oracle. On the platykurtic side, where all moments are finite, the median- and mean-based numbers agree to within about 6% (uniform: 3.15 versus 2.96), which confirms that the instability is confined to the infinite-variance regime and leaves the two-sidedness conclusion unchanged.

The same recomputation also corrects the integer branch: the mean-based $RE = 1.11$ of PMM3 at $\alpha_s=1.2$ was an artifact of the few replicates on which the solver falls back to OLS, and its median-based efficiency is 5.08 [4.45; 5.87]; conversely, the fractional solver converges on only 38.3% of the $\alpha_s=1.2$ blocks, so AUTO_EXT coincides with OLS on the remaining 61.7% — a limitation of the practical branch that we state explicitly.

Table 3 Relative efficiency vs. OLS, heavy-tailed regime (fractional branch). RE is the *ratio of medians of squared errors*, $\text{med}(e_{\text{OLS}}^2)/\text{med}(e^2)$; brackets give the 95% paired bootstrap interval (2000 resamples over noise blocks). PATP* is an *oracle upper bound* (α^* selected post hoc as the argmax of RE over the sweep grid, knowing the true amplitude); AUTO_EXT is the *practically achievable* adaptive rule (α from pooled cumulants)

Distribution	med. $\hat{\gamma}_4$	Huber	PMM3	PMM_FRAC	PATP* (oracle)	α^*	AUTO_EXT (practical)
α -stable $\alpha_s=1.2$	326	74.1 [64.0; 87.4]	5.08 [4.45; 5.87]	6.41 [4.93; 8.35]	38.8 [33.8; 44.5]	0.0	6.41 [4.96; 8.26]
α -stable $\alpha_s=1.5$	156	8.21 [6.97; 9.29]	2.74 [2.32; 3.10]	2.29 [2.02; 2.53]	6.26 [5.41; 7.05]	0.0	2.29 [2.03; 2.54]
α -stable $\alpha_s=1.8$	32	1.94 [1.74; 2.19]	1.57 [1.40; 1.74]	1.22 [1.10; 1.34]	1.77 [1.59; 2.00]	0.0	1.20 [1.10; 1.32]
Student-t $\nu=3$	12	1.84 [1.64; 2.12]	1.33 [1.17; 1.51]	1.33 [1.21; 1.48]	1.79 [1.59; 2.02]	0.0	1.28 [1.16; 1.40]
Student-t $\nu=5$	2.9	1.27 [1.16; 1.39]	1.18 [1.09; 1.30]	1.25 [1.13; 1.37]	1.23 [1.14; 1.34]	0.0	1.08 [1.02; 1.20]
Student-t $\nu=10$	0.9	1.04 [0.96; 1.13]	1.06 [0.98; 1.14]	1.00 [0.95; 1.07]	1.00 [0.95; 1.07]	0.3	1.01 [0.98; 1.05]
$\alpha_s=2.0$ (control)	≈ 0	0.98 [0.90; 1.04]	0.99 [0.94; 1.02]	1.00	1.00	0.5	1.00
$\nu=\infty$ (control)	≈ 0	0.92 [0.86; 0.98]	1.01 [0.97; 1.05]	1.00 [0.99; 1.01]	1.00	0.5	1.00

α^* uses the same grid and the same selection rule as in the mean-based version; re-selecting α^* under the median criterion changes none of the heavy-tailed rows ($\alpha^* = 0$ either way) and differs only on the near-Gaussian controls, where $RE \approx 1$ and the argmax is noise-driven. At $\alpha_s=1.2$ and 1.5 the PMM_FRAC and AUTO_EXT columns are the same estimator — the rule routes every block to the fractional branch — so their point estimates coincide and the two intervals differ only by bootstrap resampling noise; at $\alpha_s=1.8$ the rule routes part of the blocks elsewhere and the columns genuinely differ. The Gaussian control was a single merged row in the previous version; it is split here into the two actual sets, whose median-based numbers differ (Huber 0.98 and 0.92). A 1% upper-trimmed variant preserves the ordering of the methods but is still unstable at $\alpha_s=1.2$ for PMM3 (21.8 [3.5; 25.5]); the median-based figure is therefore the headline and the trimmed one is reported in Table S5 of the supplementary material.

Table 4 Relative efficiency vs. OLS, platykurtic regime (power branch); same median-based definition and bootstrap as Table 3. All moments are finite here, so mean- and median-based numbers agree closely

Distribution	$\hat{\gamma}_4$	Huber	PMM3	PATP-frac* (oracle)	AUTO_EXT (practical)
Uniform	-1.20	0.983 [0.945; 0.995]	2.96 [2.62; 3.41]	1.00	2.96 [2.62; 3.43]
Beta(2,2)	-0.85	0.88 [0.82; 0.91]	1.39 [1.22; 1.54]	1.00	1.38 [1.21; 1.53]
Gen. Gaussian $\beta=4$	-0.81	0.91 [0.86; 0.93]	1.35 [1.22; 1.50]	1.00	1.33 [1.21; 1.47]
Gaussian	≈ 0	0.92 [0.86; 0.99]	1.01 [0.97; 1.05]	1.00	1.00

Huber is significantly below unity on all three platykurtic sets (upper confidence limits below 1), while PMM3/AUTO_EXT deliver 1.33–2.96 and exactly 1.00 on the Gaussian control. The two-sidedness conclusion is unchanged by the switch to the robust functional.

PATP α -transition: RE peaks at the fractal edge as tails get heavier ($RE > 1 \Rightarrow$ beats OLS)




Fig. 3. $RE(\alpha)$ curves for the α -stable and Student-t families: the RE peak at the fractional edge ($\alpha \rightarrow 0$) is higher the heavier the tails; all curves cross at $\alpha=1/2$ (PATP=OLS)





Fig. 4. α^* versus measured $\hat{\gamma}_4$: a data-driven shift of the optimum towards the fractional edge as the tails become heavier



Median-based RE with 95% paired block-bootstrap intervals; controlled Monte-Carlo regimes only (real RF is descriptive, see Table S1).

Fig. 5. Regime map under the median-based RE with 95% paired block-bootstrap intervals: Huber is one-sided — it is significantly below OLS on the platykurtic side and on the Gaussian control — whereas the PMM/PATP family gains on both non-Gaussian edges and returns exactly OLS at the centre. Controlled Monte-Carlo regimes only; real RF is descriptive and appears in Table S1 of the supplementary material

5.3 What the selection rule costs

Section 3 justified the thresholds of Algorithm 1 as roughly four standard errors of the corresponding cumulant. This section measures what that choice costs.

Against a per-block oracle that always picks the estimator with the smallest squared error, the base dispatcher chooses a suboptimal member in 65.0% of the synthetic blocks and 69.9% of the real-noise blocks. The resulting regret, $\text{MSE}(\text{AUTO})/\text{MSE}(\text{oracle})$, is 1.220 [1.188; 1.261] on real noise and 16.19 [9.79; 24.24] on

the synthetic set. The per-block oracle is optimistically biased, since it partly selects on estimation noise; an oracle restricted to one decision per $\hat{\gamma}_4$ decile — the same information the rule itself has — gives 1.075 and 11.28 respectively. Sweeping k over [0.5, 8] shows that on real noise the rule is close to its own optimum: the minimum regret, 1.166, is attained at $k^* = 3.25$, only 4.6% below the value at the thresholds used here, so those thresholds are not tuned to the data — but they are slightly too conservative, and the optimum moves towards a *wider* PMM3 region ($\hat{\gamma}_4 < -0.50$ instead of -0.70).

Two negative findings must be stated explicitly. First, on real noise the rule is beaten by the trivial constant policy “always PMM3” (MSE ratio to OLS 0.902 versus 0.841), because the $\hat{\gamma}_4 < -0.7$ gate rejects 922 blocks on which PMM3 is in fact the best estimator. Second, on the synthetic polygaussian set no value of k makes the rule competitive: regret decreases monotonically to the edge of the grid, and the dispatcher remains $10.7\times$ worse than plain OLS, because it inherits the per-block PMM2 instability at $N = 1024$ documented in Section 5.1; this ratio is tail-driven — ten of 2500 blocks carry 55.9% of the total squared error, and the median-based ratio is 1.37.

Misclassification is concentrated where the theory predicts: on real noise it is 35.6% in the deeply platykurtic decile ($\hat{\gamma}_4 < -0.93$) and rises to 77–83% in the deciles adjacent to the threshold, where the distance to the threshold is comparable to $\text{SE}(\hat{\gamma}_4) = \sqrt{24/N} = 0.153$. A hard threshold on a noisy cumulant is thus the weak point of the current rule. The safe operating recommendation that follows is conditional rather than automatic: per-block PMM3 on nearly symmetric platykurtic RF residuals, pooled cumulants for PMM2 whenever PMM2 is used at all, and OLS elsewhere. The full audit — the confusion matrices against the oracle, the misclassification rate broken down by $\hat{\gamma}_4$ decile, and the constant-policy comparison — is given in Tables S2–S4 and Fig. S1 of the supplementary material.

5.4 Interval inference and real RF

On the real-noise dataset, naive sandwich standard errors give only about 0.72 empirical coverage for every point estimator. With the Newey–West HAC correction, coverage rises to 0.924 for OLS, 0.920 for PMM2 and 0.922 for PMM3 ($n = 2500$; Monte-Carlo standard error of a coverage estimate ≈ 0.005). The three HAC coverages are statistically indistinguishable (exact McNemar on paired blocks: $p = 0.60$ for PMM3 versus OLS, $p = 0.66$ for PMM3 versus PMM2), so the recommended estimator, PMM3, is *not* paid for by a loss of interval validity: it attains the same coverage as OLS while reducing the MSE from 3.28×10^{-3} to 2.76×10^{-3} . The correction itself is large — the median HAC-to-sandwich standard-error ratio is 1.75 for PMM3 and about 1.81 for OLS and PMM2 — which confirms that the RF residuals are strongly autocorrelated.

A residual gap of roughly three percentage points below the nominal 0.95 remains for all three estimators, and at the lowest SNR (−5 dB) coverage falls further, to 0.908–0.916. This gap is a property of the finite-sample Newey–West estimator at $N = 1024$ (finite truncation lag $L = \lfloor 4(N/100)^{2/9} \rfloor$, downward bias of the long-run variance [18, 19]), not of the PMM weighting; intervals computed on single $N = 1024$ blocks should therefore be read as slightly optimistic. No tuning was performed: L follows the standard rule and κ is taken from the PMM fit

as is. Closing the remaining three points would require a different long-run variance estimator (prewhitened HAC or a block bootstrap) and is left for future work.

Real impulsive RF. Table S1 of the supplementary material reports the leptokurtic OcuSync blocks of the two drones. These numbers are descriptive: all blocks of a given drone come from a single recording ($n_{\text{rec}} = 1$), so the intervals describe block-to-block dispersion within one recording and support no p-value and no claim of generalisation to other units or receivers. Read that way, three points stand out. First, the spread *between* functionals exceeds any confidence interval — Huber on the Parrot recording is 2.86 under the mean-based definition and 302 under the median-based one — because the OLS error on a real block is dominated by a deterministic per-block offset that any robust score removes on most blocks and none removes on a minority of them. We therefore report all three functionals and give no single headline number for real RF. Second, the gain is source-dependent: on the Parrot recording the fractional branch gains under every functional, whereas on the Mavic recording the gains are marginal on the mean-based and trimmed functionals (1.02–1.03). Third, the integer PMM3 is in the wrong regime on the Mavic blocks and loses (median-based $RE = 0.55$ [0.44; 0.70]), which is consistent with the regime map: leptokurtic blocks are not the operating region of the cubic score. The boundary is therefore narrower than a single number suggests: on real RF the fractional branch is useful, its magnitude depends on the structure of the source, and one recording per drone is not enough to say more.

6 Discussion

The results suggest four conservative rules for SDR practice.

First, PMM2 should not be applied with per-block cumulants when the skewness is weak relative to the sampling error of $\hat{\gamma}_3$. At $N = 1024$, the standard error of the skewness is approximately $\sqrt{6/N} = 0.077$, which makes small skewness estimates unreliable. Pooled calibration is therefore not a cosmetic refinement but a requirement for the stability of PMM2.

Second, PMM3 is the practical winner for the real RF residuals considered here. The reason is structural: the residuals are nearly symmetric and platykurtic, so the cubic score matches the observed distributional regime. The gain is not uniform over all blocks; it is concentrated where the excess kurtosis is sufficiently negative. On these data the constant policy “always PMM3” is in fact better than the adaptive rule (MSE ratio to OLS 0.841 versus 0.902), which is the practical form of the same finding.

Third, confidence intervals in RF amplitude estimation cannot be interpreted without checking for autocorrelation. The under-coverage of the naive sandwich

intervals is large enough to change the conclusions about estimator reliability. HAC correction is a simple and large improvement — it recovers about 20 percentage points of coverage — but it does not fully restore the nominal level at $N = 1024$.

Fourth, PMM/PATP should be regarded as a *single two-sided family* that covers both non-Gaussian edges and degenerates at the centre. The power term (PMM3) gains on the platykurtic side, the fractional branch with α -optimisation gains in the heavy tails ($\alpha^* \rightarrow 0$), and on the Gaussian controls the family returns exactly OLS ($RE = 1.00$). This is a structural advantage over a one-sided robust baseline: the Huber M-estimator gains substantially in the heavy tails, but is significantly below OLS on the platykurtic side (median-based $RE = 0.88$ – 0.98) and loses 2–8% on the Gaussian controls, because it always down-weights large residuals.

The adaptive rule AUTO_EXT implements this two-sidedness — the divergent-moment guard prevents the inflated sample skewness of a heavy-tailed block from falsely triggering PMM2, and the marginal-gain guard keeps a negligible theoretical advantage from being paid for with noisy weights — but “adaptive” is not the same as “free”: measured against a per-block oracle the rule is suboptimal on about two thirds of the blocks, its regret is 1.22 on real RF noise, it is beaten there by the constant policy “always PMM3”, and on the synthetic polygaussian set it inherits the per-block PMM2 instability and ends up an order of magnitude worse than OLS. The weak point is identified: a hard threshold on a cumulant whose standard error is comparable to the distance to that threshold. The natural extension follows directly — a soft weighting between neighbouring members, or a threshold scale calibrated on $SE(\hat{\gamma}_4)$, would remove the near-threshold randomness. Two caveats remain: for $S = 2$ the fractional branch covers only moderately heavy tails ($\alpha_s \gtrsim 1.5$, where the $\nu_{3/2}$ moment exists; the Cauchy is excluded), and a formal central limit theorem for the PATP estimator remains open, so uncertainty there is assessed empirically (sandwich/HAC).

Limitations. Five limitations bound these conclusions. (a) The dispatch rule is evaluated on the noise block itself rather than on OLS residuals; for the one-regressor model the difference is $O(1/N)$, but the reported figures are those of an idealised input. (b) The audit of Section 5.3 characterises the base rule; AUTO_EXT is characterised only through its efficiency in Tables 3–4 and Table S1, and a per-block oracle audit of AUTO_EXT on heavy-tailed data remains to be done. (c) At $\alpha_s=1.2$ the fractional solver converges on 38.3% of blocks, so the reported AUTO_EXT efficiency is that of a rule which is OLS on the majority of blocks; improving the solver would raise it, and the number should be read as a lower bound for the branch and an upper bound for the current implementation. (d) The real-RF results rest on one recording per drone; they are descriptive

and carry no inference across sources or receivers, and the same caveat applies to the real-noise dataset, whose 2500 blocks come from a single RadioML2018 noise bank. (e) HAC coverage remains about three percentage points below nominal at $N = 1024$ for every estimator; a formal central limit theorem for the PATP estimator is still open, so its uncertainty is assessed empirically only.

These conclusions should not be over-generalized. The RadioML2018 residuals are a useful open RF source, but they do not cover all hardware chains, modulation settings, interference types, or field deployments. The next step is to repeat the same protocol on wider SDR recordings with receiver metadata and controlled channel conditions.

Conclusions

PMM can improve SDR amplitude estimation under non-Gaussian RF noise, but only when its distributional assumptions are respected, and the choice among family members is itself a measurable decision problem rather than a free lunch. PMM3 reduces MSE on real platykurtic RF residuals (2.76×10^{-3} versus 3.28×10^{-3}) and, with Newey–West HAC standard errors, attains the same empirical coverage as OLS (0.922 versus 0.924), so the point-estimate gain is not paid for with interval validity; the residual three-point gap below the nominal 0.95 applies to all estimators alike and is a finite-sample property of the HAC estimator. PMM2 is unsafe with per-block cumulants at $N = 1024$ unless the skewness is strong; pooled calibration restores it to the OLS baseline, not above it. The fractional PATP branch covers the heavy-tailed impulsive regime where the integer tracks do not apply, and under a well-defined median-based efficiency its practically achievable gain at $\alpha_s=1.2$ is 6.41 [4.96; 8.26] against an oracle upper bound of 38.8 [33.8; 44.5]. Together the two branches form a two-sided family that gains on both non-Gaussian edges and returns exactly OLS on the Gaussian controls, unlike a one-sided robust M-estimator.

The adaptive AUTO_EXT rule automates this selection, and we report the cost of the automation: a suboptimal member on 65–70% of blocks, regret 1.22 on real RF noise, and a failure mode on synthetic polygaussian noise where the rule is an order of magnitude worse than OLS. The safest operational strategy is therefore conditional rather than unconditional dispatch — per-block PMM3 on nearly symmetric platykurtic residuals, pooled cumulants whenever PMM2 is used, the fractional branch for heavy tails, and OLS by default — with a soft, SE-calibrated selection rule as the natural next step.

Code availability

All PMM computations were performed with the R package EstemPMM: Zabolotnii, S. (2026) *EstemPMM: Polynomial Maximization Method for Non-Gaussian Regression*. R package version 0.4.0. Available at: <https://CRAN.R-project.org/package=EstemPMM>. The results were verified to be numerically identical under package versions 0.3.2 and 0.4.0: the oracle reference fixture regenerated under 0.4.0 reproduces every value of the 0.3.2 fixture exactly, and the Monte Carlo results of Table 2 are bit-identical across the two versions. The experiment scripts and Monte Carlo configurations that reproduce every table and figure in this article are available at <https://github.com/SZabolotnii/sdr-kunchenko>.

References

- [1] Middleton D. (1999). Non-Gaussian noise models in signal processing for telecommunications: new methods and results for Class A and Class B noise models. *IEEE Transactions on Information Theory*, Vol. 45, No. 4, pp. 1129–1149. DOI: 10.1109/18.761256.
- [2] Clavier, L., Peters, G.W., Septier, F. *et al.* (2021). Impulsive noise modeling and robust receiver design. *J Wireless Com Network*, Vol. 2021, Article 13. DOI: 10.1186/s13638-020-01868-1.
- [3] Nikias C. L., Shao M. (1995). *Signal Processing with Alpha-Stable Distributions and Applications*. New York: Wiley-Interscience, 168 p. ISBN 0-471-10647-X.
- [4] Clavier L., Pedersen T., Rodriguez I., Lauridsen M., Egan M. (2021). Experimental evidence for heavy tailed interference in the IoT. *IEEE Communications Letters*, Vol. 25, No. 3, pp. 692–695. DOI: 10.1109/LCOMM.2020.3034430.
- [5] Huber P. J. (1964). Robust estimation of a location parameter. *Annals of Mathematical Statistics*, Vol. 35, No. 1, pp. 73–101. DOI: 10.1214/aoms/1177703732.
- [6] Huber P. J., Ronchetti E. M. (2009). *Robust Statistics*. 2nd ed. Hoboken, NJ: Wiley, 380 p. ISBN 978-0-470-12990-6.
- [7] Zoubir A. M., Koivunen V., Chakhchoukh Y., Muma M. (2012). Robust estimation in signal processing: a tutorial-style treatment of fundamental concepts. *IEEE Signal Processing Magazine*, Vol. 29, No. 4, pp. 61–80. DOI: 10.1109/MSP.2012.2183773.
- [8] Hosking J. R. M. (1990). L-moments: analysis and estimation of distributions using linear combinations of order statistics. *Journal of the Royal Statistical Society: Series B*, Vol. 52, No. 1, pp. 105–124. DOI: 10.1111/j.2517-6161.1990.tb01775.x.
- [9] Kunchenko Yu. P. (2002). *Polynomial Parameter Estimations of Close to Gaussian Random Variables*. Aachen: Shaker Verlag, 396 p. ISBN 3-8322-0032-0.
- [10] Zabolotnii, S., Warsza, Z. L., Tkachenko, O. (2018). Polynomial Estimation of Linear Regression Parameters for the Asymmetric PDF of Errors. In: Szewczyk, R., Zieliński, C., Kaliczyńska, M. (eds) Automation 2018. *AUTOMATION 2018. Advances in Intelligent Systems and Computing*, Vol 743, pp. 101–110. Springer, Cham. doi:10.1007/978-3-319-77179-3_75.
- [11] Zabolotnii, S. W., Warsza, Z. L., Tkachenko, O. (2020). Estimation of Linear Regression Parameters of Symmetric Non-Gaussian Errors by Polynomial Maximization Method. In: Szewczyk, R., Zieliński, C., Kaliczyńska, M. (eds) Automation 2019. *AUTOMATION 2019. Advances in Intelligent Systems and Computing*, Vol 920. Springer, Cham. doi: 10.1007/978-3-030-13273-6_59.
- [12] Zabolotnii S. (2026). Parametrically Adaptive Transition Polynomial: a Signed-Parity Continuous- α Extension of Kunchenko Stochastic Polynomials. *arXiv*:2605.14610. doi: 10.48550/arXiv.2605.14610.
- [13] Shao M., Nikias C. L. (1993). Signal processing with fractional lower order moments: stable processes and their applications. *Proceedings of the IEEE*, Vol. 81, No. 7, pp. 986–1010. DOI: 10.1109/5.231338.
- [14] Rousseeuw P. J., Croux C. (1993). Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, Vol. 88, No. 424, pp. 1273–1283. DOI: 10.1080/01621459.1993.10476408.
- [15] Newey W. K., West K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, Vol. 55, No. 3, pp. 703–708. DOI: 10.2307/1913610.
- [16] Zeileis A. (2004). Econometric computing with HC and HAC covariance matrix estimators. *Journal of Statistical Software*, Vol. 11, No. 10, pp. 1–17. DOI: 10.18637/jss.v011.i10.
- [17] Newey W. K., West K. D. (1994). Automatic lag selection in covariance matrix estimation. *The Review of Economic Studies*, Vol. 61, No. 4, pp. 631–653. DOI: 10.2307/2297912.
- [18] Andrews D. W. K. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, Vol. 59, No. 3, pp. 817–858. DOI: 10.2307/2938229.
- [19] Kiefer N. M., Vogelsang T. J. (2005). A new asymptotic theory for heteroskedasticity-autocorrelation robust tests. *Econometric Theory*, Vol. 21, No. 6, pp. 1130–1164. DOI: 10.1017/S0266466605050565.
- [20] O'Shea T. J., Roy T., Clancy T. C. (2018). Over-the-air deep learning based radio signal classification. *IEEE Journal of Selected Topics in Signal Processing*, Vol. 12, No. 1, pp. 168–179. DOI: 10.1109/JSTSP.2018.2797022.

Оцінювання амплітуди SDR-сигналу методом РММ за негаусівського радіочастотного шуму

Заболотний С. В., Коваль В. В., Чепинога А. В.

Оцінювання амплітуди за записами baseband-сигналів software-defined radio (SDR) стає ненадійним, коли адитивний радіочастотний (RF) шум відхиляється від гаусівської моделі, на якій базується ordinary least squares (OLS). Замість того щоб порівнювати оцінювачі попарно, у цій статті OLS, метод поліноміальної максимізації другого й третього порядків (PMM2, PMM3) та дробову гілку параметрично-адаптивного перехідного полінома (PATP) розглянуто як членів єдиного неперервно параметризованого сімейства, а задачу сформульовано як синтез вирішального правила, що обирає члена сімейства за кумулянтами одного скінченного блоку, мінімізуючи regret відносно оракульного вибору. Протокол поєднує

синтетичний полігаусівський шум, реальні RF-залишки з over-the-air записів RadioML2018, контрольовані важкохвостові та платикуртичні режими, а також лептокуртичні блоки з двох записів дронів. Поблоковий PMM2 нестійкий при $N = 1024$, бо оцінка асиметрії має низьке відношення сигнал/шум; pooled-усереднення кумулянтів за калібрувальними блоками повертає його до базового рівня OLS. На реальному платикуртичному RF-шумі PMM3 зменшує середньоквадратичну похибку з $3,28 \times 10^{-3}$ до $2,76 \times 10^{-3}$, а стандартні похибки Newey–West HAC піднімають емпіричне 95% покриття всіх трьох точкових оцінювачів приблизно з 0,72 до 0,920–0,924 – включно з PMM3 (0,922), – що все ще приблизно на три відсоткові пункти нижче номінального рівня. У важкохвостовому режимі відносну ефективність подано як відношення медіан квадратів похибок із парними bootstrap-інтервалами, бо середнє-базована величина не збігається за нескінченної дисперсії: при $\alpha_s=1,2$ оракульна межа PATP* сягає 38,8 [33,8; 44,5], тоді як практично досягне адаптивне

правило AUTO_EXT – 6,41 [4,96; 8,26]. Сімейство є двобічним: воно виграє на обох негаусівських краях і повертає точно OLS на гаусівських контролях – на відміну від однібічного M-оцінювача Хубера, який втрачає 2–12% на платикуртичному шумі. Нарешті, кількісно оцінено, чого коштує dispatch-правило: проти оракула на блок воно обирає неоптимального члена сімейства у 65,0% синтетичних і 69,9% реальношумових блоків, із regret 16,19 та 1,220 відповідно, а на синтетичному наборі воно у $10,7\times$ гірше за звичайний OLS. Тому операційна рекомендація є умовною, а не безумовною диспетчеризацією: PMM3 – для майже симетричних платикуртичних RF-залишків, pooled-калібрація – для PMM2, дробова гілка – для важких хвостів, а OLS – завжди, коли кумулянти блоку не вказують явно на інше.

Ключові слова: PMM; PATP; оцінювання амплітуди; SDR; негаусівський шум; важкохвостий шум; дробові моменти; кумулянти; HAC; вибір оцінювача; regret; медіанна відносна ефективність; RadioML2018