Skepthical review: *Cross-Spectral Wiener Filtering for Optimal Thermal Sunyaev-Zel'dovich Signal Extraction and Galaxy Cluster Detection*

Summary

This manuscript develops and tests a multi-frequency Wiener / LMMSE filtering framework to reconstruct the thermal Sunyaev–Zel'dovich (tSZ) Compton-y field from multi-band CMB maps in the presence of strong, spatially correlated foregrounds—especially the cosmic infrared background (CIB). Using FLAMINGO-based simulations with six channels spanning **90–857 GHz** (SO LAT–like + Planck HFI–like; Sec. 2.1), the authors estimate full auto- and cross-frequency power spectra and construct a harmonic-space linear estimator (Sec. 2.2.1) that treats the CIB as correlated noise in the data covariance rather than attempting deterministic nulling (contrasted to an ILC baseline; Sec. 2.2.2). Performance is evaluated both in harmonic space via transfer functions / correlations (Sec. 2.3.1, 3.2) and end-to-end via a matched-filter cluster detection + photometry pipeline (Sec. 2.3.2, 3.3). The main reported outcome is that the MWF yields more controlled small-scale behavior and substantially higher cluster-catalog purity (notably in CIB-bright environments), with a tighter mass–observable relation but with a predictable negative photometric bias attributed to Wiener suppression (Sec. 3.3.2–3.3.3). The contribution is timely and potentially impactful for forthcoming survey analyses; however, key elements of the observation model, covariance construction/validation, baseline ILC definition, and the cluster-selection/photometric calibration details are currently under-specified, limiting reproducibility and making it difficult to assess robustness and generalizability to real data (Sec. 2–4).

Strengths

- Addresses a central practical challenge for tSZ cluster science in current/upcoming surveys: CIB contamination that is both bright and spatially correlated with cluster environments (Sec. 1).
- Uses cross-frequency (auto + cross) covariance information in a principled linear MMSE/Wiener framework, which is a natural way to suppress correlated contaminants without hard spectral nulls (Sec. 2.2.1).
- Evaluation goes beyond map-level diagnostics to end-to-end cluster finding and photometry, reporting completeness/purity behavior and mass–observable trends that are directly relevant to cosmological use (Sec. 2.3, 3.3).
- The purity degradation of the ILC in CIB-bright regions (Fig. 5; Sec. 3.3.2) is a clear and practically useful diagnostic that motivates covariance-aware approaches.
- The manuscript acknowledges numerical conditioning challenges from extreme dynamic range and documents stabilization steps (float64, standardization, Tikhonov regularization), which is important for implementers (Sec. 2.1, 2.2.1, 3.1).

Major issues

1. **The core observation model and the definitions/units of the key Wiener-filter objects are not specified rigorously enough to verify correctness and enable reproduction (Sec. 2.2.1; also impacts Sec. 3.1–3.3).** In particular, Eq. (2) is presented as $\mathbf{W}\ell=(\mathbf{S}\ell+\mathbf{N}\ell)^{-1}\,\mathbf{s}}}, with\mathbf{s{\ell,{\rm tSZ}} describedasacross − spectrumvector, butthepaperdoesnotwriteanexplicitmulti − frequencymixingmodel(spectralresponses, beamtransferfunctions, mapunits)thatmapsthesimulatedcomponentsintothechannelma \mathbf{S}\ell+\mathbf{N}\ellequalsthetotaldatacovariance\mathbf{C}andthereconstructed\hat y.}(\ell) required by standard LMMSE, (ii) what exactly is included in "signal" vs "noise" bookkeeping (CMB/kSZ/tSZ/CIB/noise), (iii) whether spectra are computed on beam-convolved maps or after beam equalization/deconvolution, and (iv) whether the standardization of the y-map (Sec. 2.1) changes the meaning/units of $\mathbf{s}_{\ell,{\rm tSZ}

   Recommendation: In Sec. 2.2.1 (or an appendix), add an explicit harmonic-space data model, e.g. $d_{i,\ell m} = B_i(\ell)\,[g_i\,y_{\ell m} + a_i^{\rm CMB} c_{\ell m} + a_i^{\rm kSZ} k_{\ell m} + a_i^{\rm CIB} f_{\ell m}] + n_{i,\ell m}$, defining all coefficients (tSZ spectral factor g_i, assumed CMB/kSZ scaling, CIB convention), map units ($K_{\rm CMB}$ vs intensity), and how beams/bandpasses enter. Then define $\mathbf{C}{dd}(\ell)and\mathbf{C}}(\ell) andexplicitlymapthemontoyour\mathbf{S\ell,\mathbf{N}\ell,\mathbf{s}{\ell,{\rm tSZ}} notation.Clarifywhetherspectraarecomputedwithbeamsincluded, oraftersmoothingtoacommonbeam.Finally, stateexplicitlyhowy − standardizationisappliedinthepipeline : whether\mathbf{s}isbuiltfromstandardizedvsphysicaly, andhow\hat y$ is de-standardized before integrated-Y photometry (Sec. 2.3.2, 3.3.3).}

2. **Treatment of CIB as "noise" vs its physical correlation with tSZ/halos is not disentangled, which matters for bias interpretation and for the claim that the observed ∼20–30% negative photometric bias is "Wiener attenuation" (Sec. 1, 3.3.3, 4).** Because the CIB is correlated with large-scale structure and cluster environments, a correlated foreground can introduce bias terms beyond a simple (single-channel) $S/(S+N)$ intuition; additionally, if y–CIB correlations are present in the simulations, they may leak into $\mathbf{s}{\ell,{\rm tSZ}}\equiv\mathbf{C}}(\ell) dependingonhowitisestimated.WithoutclarifyingwhetherthesimulationincludesexplicittSZ–CIBcross − correlationandhowitpropagatesinto\mathbf{C{dy}and\mathbf{C}$, it is hard to interpret whether the method is suppressing purely contaminating CIB, suppressing cluster-associated correlated emission in a mass/redshift/environment-dependent way, or partially fitting correlated CIB as "signal."

   Recommendation: Clarify in Sec. 2.1 and Sec. 2.2.1 whether the simulated skies include non-zero y–CIB cross-spectra (one-halo/two-halo terms), and explicitly report/plot $C_\ell^{y,{\rm CIB}_\nu}$ for at least a few ν. Then, in Sec. 3.3.3, decompose the recovered-Y bias into (i) multiplicative attenuation from the filter response/transfer function and (ii) bias from correlated foreground structure (if present). If feasible, add a controlled test in Sec. 3: re-run with y–CIB correlation turned off (or scaled) to isolate how much of the gain (purity/scatter) and how much of the bias arises from covariance suppression vs true correlation with the target. Also revise the text in Sec. 3.3.3 to avoid the scalar $S/(S+N)$ statement for the multi-frequency case unless you explicitly state assumptions under which it reduces to that form.

3. **Covariance estimation and use ($\mathbf{S}\ell, \mathbf{N}\ell, and/or total \mathbf{C}_{dd}$) is under-specified and may be "simulation-trained" in a way that overstates robustness (Sec. 2.2.1, 3.1).** The manuscript states auto/cross spectra are estimated from 500 simulated patches, but does not provide patch geometry/area, masking/apodization, multipole binning, pseudo-C_ℓ vs direct Fourier estimation, whether the covariance-estimation simulations are disjoint from those used for evaluation (risk of circularity), nor how spectra are smoothed/regularized. Given the method's main selling point is exploiting detailed cross-frequency CIB covariance, sensitivity to finite-sample noise and to modest covariance mismatch is a central concern for real-data applicability.

   Recommendation: Expand Sec. 2.2.1 and Sec. 3.1 with a reproducible covariance pipeline description: patch size/shape, number of patches, sky fraction, mask/apodization, ℓ-range and binning, estimator used for spectra (and any debiasing), and any smoothing/fit applied to enforce positive-definiteness. State explicitly whether filter design uses simulations/patches disjoint from evaluation maps. Add at least one robustness study in Sec. 3.1–3.3: construct the MWF using one set of realizations/regions and apply it to another; and/or perturb the assumed covariances (e.g., scale CIB power by ±10–20%, alter cross-band correlation coefficients, misestimate noise) and quantify impact on transfer functions (Sec. 3.2) and catalog purity/completeness and Y-bias/scatter (Sec. 3.3).

4. **The baseline ILC is not defined precisely enough to judge fairness/representativeness, and some reported ILC failure modes could reflect implementation choices rather than intrinsic limitations of modern ILC variants (Sec. 2.2.2, 3.2–3.3, 4).** It is unclear whether the ILC is global vs ℓ-dependent, whether it is localized (patch/needlet), how $\mathbf{C}_\ell$ is estimated, whether beams are equalized, whether CMB is nulled (often done in y-map construction), what regularization is used, and whether weights are derived from the same covariance-estimation setup as the MWF. Without these details, the comparison risks being interpreted as "MWF vs ILC in general," which would be too broad.

   Recommendation: In Sec. 2.2.2, give the explicit ILC weight formula and implementation domain (per-ℓ, binned-ℓ, patch-wise, needlet/NILC-like). Specify how $\mathbf{C}\ell is estimated (same patches/masks/bins as MWF or not), how beams/noise are handled (common-resolution smoothing vs including beam factors in \mathbf{C}\ell$), and whether additional constraints are applied (e.g., CMB null, dust/CIB constraints). In Sec. 4, narrow the claim scope: state clearly that conclusions are relative to this specific baseline. If feasible, add at least one more competitive baseline (e.g., constrained ILC with CMB nulling and/or a scale-localized/needlet ILC) to better contextualize the gains.

5. **The simulated-observation specification is incomplete, making it hard to map results onto actual SO/Planck performance and to reproduce the study (Sec. 2.1).** The manuscript does not provide a per-channel table of central frequencies, beam FWHM, pixelization/resolution, noise levels (e.g., μK-arcmin or noise N_ℓ), sky coverage/masking, or whether noise is white/anisotropic/1/f. It is also unclear which astrophysical components are included/omitted (radio point sources, Galactic dust, beam/calibration uncertainties, bandpass mismatch), and this matters because covariance-based methods can be sensitive to missing/non-stationary components (Sec. 4).

   Recommendation: Add a concise table (main text or appendix) in Sec. 2.1 listing for each channel: ν, beam FWHM, map unit, pixel size (HEALPix $N_{\rm side}$ or flat-sky pixel scale), and noise model (amplitude and ℓ-dependence). State sky geometry (full-sky vs tiles), masking/apodization, and how SO+Planck are combined spatially. Enumerate included components (CMB, tSZ, kSZ, CIB, instrumental noise) and explicitly list omitted ones (radio sources, Galactic dust, etc.). In Sec. 4, add a limitations paragraph assessing how these omissions could affect covariance estimation and cluster purity/bias in real data, and how the framework could incorporate additional components.

6. **Cluster detection/photometry methodology and selection-function characterization are not yet detailed/quantified enough for cosmology-facing interpretation (Sec. 2.3.2, 3.3.1–3.3.3).** Key missing items include the explicit matched-filter definition (and which noise power spectrum enters), template/profile construction and size grid, SNR normalization, peak-finding/deblending, halo matching criteria, halo mass definition (e.g., M_{500c}), and mass/redshift ranges used. The fixed "3-pixel aperture" integrated-Y needs an angular-scale mapping and justification across varying cluster sizes; additionally, the paper reports a sizeable negative Y bias but does not present a practical calibration workflow (e.g., transfer-function correction or simulation-derived mapping) or quantify residual scatter after calibration.

   Recommendation: In Sec. 2.3.2 and Sec. 3.3.1, write the matched-filter equations explicitly (Fourier/harmonic form), define the noise $C_\ell^{\rm noise}$ used in the denominator, and state whether the filter is recomputed separately for MWF vs ILC maps. Describe the cluster template (beam-convolved profile, parameter grid) and SNR normalization. Specify peak-finding/deblending and matching-to-halo rules (matching radius, handling of multiple matches). In Sec. 3.3.2–3.3.3, report completeness and purity as functions of true mass (and, if feasible, redshift) at the chosen SNR thresholds, not only vs recovered Y, and provide numerical scatter metrics (e.g., $\sigma_{\ln Y}$) with uncertainties. State the pixel size so the 3-pixel aperture corresponds to a well-defined angular scale; ideally test a size-adaptive aperture or profile-based flux estimator and verify conclusions are unchanged. Finally, outline and demonstrate a calibration scheme for the MWF attenuation/bias (e.g., fit $Y_{\rm true}$ vs $Y_{\rm rec}$ in simulations, or correct using the measured transfer function), and report the post-calibration scatter relevant for cosmological analyses.

Minor issues

1. Harmonic-space evaluation metrics need clearer definitions, uncertainty estimates, and interpretation (Sec. 2.3.1, 3.2). The transfer function $T(\ell) = C_\ell^{\hat{y}y}/C_\ell^{yy}$ is an amplitude response (not a "power fraction"), and r_ℓ is referenced but not explicitly defined. Differences between methods are discussed qualitatively without error bands or concise quantitative summaries (e.g., characteristic ℓ where T drops).

   Recommendation: In Sec. 2.3.1, define all metrics explicitly, including $r_\ell = C_\ell^{\hat{y}y}/\sqrt{C_\ell^{\hat{y}\hat{y}}C_\ell^{yy}}$, and clarify that Eq. (3) measures linear response. In Sec. 3.2, estimate uncertainties from patch-to-patch variance or multiple realizations and show error bands (or report typical errors in text). Consider also plotting $C_\ell^{\hat{y}\hat{y}}$ and an effective residual power (e.g., $C_\ell^{\hat{y}\hat{y}} - T(\ell)^2 C_\ell^{yy}$) to separate multiplicative attenuation from additive contamination.

2. Regularization and numerical conditioning choices are not fully dimensionally/quantitatively documented (Sec. 2.2.1, 3.1). The Tikhonov term $\epsilon\mathbf{I}$ with $\epsilon = 10^{-5}$ is not defined relative to a covariance scale (units), and the effect of ϵ (and any ℓ-binning) on bias/variance is not quantified.

   Recommendation: Define ϵ in a dimensionally consistent way (e.g., $\epsilon\,\mathrm{Tr}(\mathbf{C})/N$ or after an explicit normalization that makes $\mathbf{C}$ dimensionless). Add a short stability check (Sec. 3.1): vary ϵ over a plausible range (e.g., 10^{-6}–10^{-4}) and show that key results (transfer function, purity, Y-bias) are stable. If $\mathbf{W}_\ell$ is computed in ℓ-bins, state bin widths and verify insensitivity.

3. The "local CIB intensity" variable used to demonstrate environment-dependent purity (Fig. 5; Sec. 3.3.2) is not defined operationally, limiting interpretability and real-data portability.

   Recommendation: In Sec. 3.3.2, specify exactly how local CIB intensity is computed: which frequency channel(s), whether it uses the pure simulated CIB component or a multi-component map, what beam/smoothing is applied, what aperture/scale is used, and whether it is evaluated at the detection position or averaged in a neighborhood. If it relies on simulation-only truth, briefly discuss an observational proxy (e.g., high-frequency intensity map after masking bright sources).

4. Discussion of related component-separation context could be broadened/clarified (Sec. 1, 2.2.2, 4). The text contrasts MWF with "standard ILC" but does not clearly position against constrained ILC, NILC/needlet approaches, matched multi-filters, or parametric maximum-likelihood methods.

   Recommendation: Add a short literature/context paragraph (Sec. 1 or Sec. 4) citing constrained ILC / NILC and related multi-frequency estimators, and clarify which subclass your baseline ILC represents. Emphasize where covariance-driven MWF is expected to excel (when covariances are reliable) and where alternatives may be competitive (strong non-Gaussianity, spatially varying statistics, poorly modeled components).

5. Figure 1 (and related text) lacks sufficient unit/definition clarity and could better support the numerical-conditioning narrative (Sec. 2.1; Fig. 1). The variance axis lacks units, the meaning of "per-channel variance" is not specified (mask, beam, inclusion of noise), and the linear scaling obscures large dynamic range. The claim of improved conditioning is not supported by a quantitative diagnostic (e.g., eigenvalue spectrum/condition number).

   Recommendation: Add explicit axis units (and confirm Compton-y is dimensionless) and specify how variance is computed (mask/region, beam/pixel resolution, components included). Use a log y-axis (or numeric annotations) for the variance panel. Add a small inset/table reporting the covariance condition number or eigenvalue spread before/after standardization to substantiate the conditioning claim.

Very minor issues

1. Notation/formatting consistency issues (e.g., Compton-y vs Compton-y, Compton-Y capitalization, unit formatting like 857 GHz) and occasional long sentences reduce readability (Sec. 1, 3.3.1, 4; figures/captions).

   Recommendation: Proofread for consistent notation and unit formatting throughout. Consider splitting a few long multi-clause sentences in Sec. 1, Sec. 3.3.1, and Sec. 4 for clarity. Add in-figure panel labels (a/b/c) where relevant so captions remain unambiguous when figures are reused.

2. Some harmonic- vs map-space shorthand could confuse readers (Sec. 2.2.2). The text alternates between $\hat y_{\ell m}=\mathbf{w}_\ell^\dagger\mathbf{d}$ and $\hat y=\sum_i w_i d_i$ without explicitly stating the latter is per-(ℓ,m) (or per-ℓ).

   Recommendation: Add a brief sentence in Sec. 2.2.2 clarifying the domain: e.g., that ILC weights are applied in harmonic space per ℓ (and thus per ℓ, m), consistent with Eq. (1). Also state whether † is used purely for complex harmonic notation and whether weights are real.

Key statements and references

- • **The simulations reveal that in the six-channel SO+Planck setup, the CIB-dominated 857 GHz map has a variance approximately 10^{15} times larger than the tSZ Compton-y signal, and the high-frequency (545 and 857 GHz) channels have variances about five orders of magnitude larger than the lower-frequency CMB-dominated channels, necessitating float64 precision and target-map standardization to maintain numerical stability in the Wiener filter covariance inversion.**
- *Reference(s):* (none)

- • **In the harmonic domain, the Internal Linear Combination reconstruction exhibits a transfer function $T(\ell) > 1$ at intermediate multipoles ($1000 < \ell < 3000$) and strong noise amplification at $\ell > 3000$ due to enforcing a foreground null in noisy channels, whereas the Multi-Frequency Wiener Filter maintains $T(\ell) \approx 1$ on large scales and smoothly suppresses power on small, noise-dominated scales, consistent with the behavior of an optimal linear estimator that minimizes mean-squared error.**
- *Reference(s):* (none)
- • **In the matched-filter cluster detection analysis, the catalog derived from the Multi-Frequency Wiener Filter Compton-y map maintains purity $> 80\%$ across a wide range of recovered integrated-Y values and shows purity largely independent of local CIB intensity, while the Internal Linear Combination catalog's purity degrades sharply in regions of bright CIB emission due to misidentifying CIB fluctuations as clusters; correspondingly, ILC photometry exhibits a large positive bias in recovered integrated Y, whereas the MWF shows a predictable negative bias of roughly 20–30% arising from the Wiener suppression factor $S/(S+N)$.**
- *Reference(s):* (none)

Mathematical consistency audit

This section audits **symbolic/analytic** mathematical consistency (algebra, derivations, dimensional/unit checks, definition consistency).

Maths relevance: light

The paper's analytic content centers on a harmonic-space multi-frequency Wiener (LMMSE) linear estimator for reconstructing a tSZ Compton-y field from multi-channel observations, plus a brief ILC baseline and diagnostic definitions (transfer function). Only three numbered equations are provided; key derivations and the explicit observation/mixing model are omitted, which limits verifiability of the central Wiener-weight expression and its interpretation under standardization and correlated foregrounds.

Checked items

1. ✔ **Harmonic-space linear estimator form** (Eq. (1), Sec. 2.2.1, p.3)
   - **Claim:** The reconstructed tSZ harmonic coefficient is a weighted linear combination of the multi-frequency data vector: $y\text{-hat}*{\ell m}* = *W*\ell^\dagger d_{\ell m}$.
   - **Checks:** shape/linear-algebra consistency, notation consistency
   - **Verdict:** PASS; confidence: high; impact: moderate
   - **Assumptions/inputs:** $d_{\ell m}$ is a 6-vector of frequency-channel harmonic coefficients at fixed (ℓ, m). W_ℓ is a 6-vector of weights depending only on ℓ (statistical isotropy).
   - **Notes:** If W_ℓ is a column vector, $W_\ell^\dagger$ is a row vector and the product yields a scalar y-**hat**$_{\ell m}$, consistent with the stated goal.
2. ⚠ **Multi-frequency Wiener weight expression** (Eq. (2), Sec. 2.2.1, p.4)
   - **Claim:** The Wiener weights are $W_\ell = (S_\ell + N_\ell)^{-1} s_{\ell,\mathrm{tSZ}}$, with S_ℓ the 6×6 covariance of (CMB,tSZ,kSZ) and N_ℓ the covariance of instrumental noise plus CIB auto/cross spectra; $s_{\ell,\mathrm{tSZ}}$ is the 6-vector cross-power between true tSZ and each observed channel.
   - **Checks:** definition consistency, dimensional/units consistency, missing-derivation check
   - **Verdict:** UNCERTAIN; confidence: medium; impact: critical
   - **Assumptions/inputs:** $S_\ell + N_\ell$ equals the total data covariance $C_{dd}(\ell)$. $s_{\ell,\mathrm{tSZ}}$ corresponds to the cross-covariance vector $C_{dy}(\ell)$ between data and the target y field in the same units/conventions as $C_{dd}(\ell)$. All spectra are defined consistently with any beam smoothing and map units.
   - **Notes:** Eq. (2) matches the standard LMMSE/Wiener structure $W = C_{dd}^{-1} C_{dy}$, but the paper does not provide the explicit mixing model (spectral responses g_i, beam factors $B_i(\ell)$, map units) needed to verify that S_ℓ, N_ℓ, and $s_{\ell,\mathrm{tSZ}}$ are defined in the correct space and combine to $C_{dd}(\ell)$. The label "desired astrophysical signals (CMB,tSZ,kSZ)" is also conceptually confusing because only tSZ is the reconstruction target, though this does not by itself break the algebra.
3. ✔ **Treatment of correlated CIB in the Wiener framework** (Sec. 2.2.1, p.4 (paragraph defining N_ℓ))
   - **Claim:** Including the full CIB auto- and cross-frequency spectra in N_ℓ treats CIB as correlated noise to be statistically suppressed.
   - **Checks:** conceptual/definition consistency with Eq. (2), sanity check on covariance usage
   - **Verdict:** PASS; confidence: medium; impact: moderate
   - **Assumptions/inputs:** CIB contributions are included in the total data covariance matrix used for inversion. Any correlation between CIB and tSZ is either negligible or correctly captured in $s_{\ell,\mathrm{tSZ}}$ if estimated empirically as $\langle dy \rangle$.
   - **Notes:** As stated, adding CIB covariance to N_ℓ is consistent with forming a total $C_{dd}(\ell)$. However, without an explicit statement of whether $\langle \mathrm{CIB} \cdot y \rangle$ is included/excluded in $s_{\ell,\mathrm{tSZ}}$, it is unclear how the method handles tSZ–CIB correlation; the core algebra still permits it.
4. ⚠ **Regularization term added before inversion** (Sec. 2.2.1, p.4 (Tikhonov regularization description))
   - **Claim:** Matrix invertibility is ensured by adding ϵI with $\epsilon = 10^{-5}$ to the diagonal of $(S_\ell + N_\ell)$ before inversion.

- **Checks:** dimensional/units consistency, definition clarity
- **Verdict:** UNCERTAIN; confidence: medium; impact: minor
- **Assumptions/inputs:** ϵ has compatible units with the covariance matrix entries or the covariance has been normalized to be dimensionless.
- **Notes:** A raw additive constant ϵ is only dimensionally consistent if the matrix has been normalized or if ϵ is defined relative to a covariance scale. The paper does not specify such scaling/normalization.

5. ✔ **ILC constraint statement** (Sec. 2.2.2, p.4)
   - **Claim:** ILC minimizes variance of $y\text{-hat} = \Sigma_i w_i d_i$ subject to $\Sigma_i w_i g_i = 1$ (unit response to tSZ spectral shape).
   - **Checks:** constraint/normalization consistency
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** g_i is the tSZ spectral response in the same units as d_i.
   - **Notes:** The stated constraint is the standard unit-gain condition for ILC. The closed-form solution is not provided but omission does not create an internal contradiction.
6. ✔ **Transfer function definition** (Eq. (3), Sec. 2.3.1, p.4)
   - **Claim:** Transfer function is $T(\ell) = P_{\text{cross}}(y\text{-hat}, y_{\text{true}})/P_{\text{auto}}(y_{\text{true}})$.
   - **Checks:** algebraic sanity check, interpretation consistency
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** Power spectra P_{cross} and P_{auto} are defined with consistent conventions (same ℓ-binning, mask corrections if any, etc.). Interpretation assumes residuals are approximately uncorrelated with y_{true} for $T(\ell)$ to estimate linear response.
   - **Notes:** If $y\text{-hat} = \alpha(\ell) y_{\text{true}} + n_{\text{res}}$ with $\langle n_{\text{res}} y_{\text{true}} \rangle \approx 0$, then $P_{\text{cross}} = \alpha P_{\text{auto}}$ and $T(\ell) = \alpha$, so the formula is internally coherent.
7. ✖ **Transfer function described as 'fractional power'** (Sec. 2.3.1 around Eq. (3), p.4)
   - **Claim:** Text describes $T(\ell)$ as fractional response to true signal power.
   - **Checks:** definition/interpretation consistency
   - **Verdict:** FAIL; confidence: high; impact: minor
   - **Assumptions/inputs:** No additional definition of 'power response' is given beyond Eq. (3).
   - **Notes:** Eq. (3) yields an amplitude (linear) response under standard assumptions; a 'power fraction' would typically scale like α^2. The equation itself is fine; the wording is inconsistent with the mathematical meaning.
8. ✖ **Claimed Wiener attenuation factor** $S/(S+N)$ (Sec. 3.3.3, p.9 (discussion of negative bias))
   - **Claim:** MWF suppression is described as Wiener attenuation by a factor $S/(S+N)$.
   - **Checks:** consistency with earlier definitions (vector vs scalar), limiting-case reasoning
   - **Verdict:** FAIL; confidence: medium; impact: moderate
   - **Assumptions/inputs:** Estimator is the multi-frequency vector Wiener filter of Eq. (2), not a single-channel scalar Wiener filter.
   - **Notes:** $S/(S+N)$ is a scalar single-channel form; with 6 channels the response depends on the full matrix $(S_\ell + N_\ell)^{-1}$ and the cross-covariance vector. The paper already computes an empirical $T(\ell)$; that is the appropriate general measure of suppression.
9. ⚠ **Standardization of y map vs filter outputs/photometry** (Sec. 2.1 (standardization), p.3 and Sec. 2.3.2 / 3.3.3 (integrated-Y & bias), pp.5–9)
   - **Claim:** Standardizing y improves conditioning; later integrated Compton-Y photometry and biases are discussed as if in physical units.
   - **Checks:** units/scale consistency across sections, missing-steps check
   - **Verdict:** UNCERTAIN; confidence: medium; impact: critical
   - **Assumptions/inputs:** If y is standardized, either all relevant spectra and outputs are in standardized units or a de-standardization step exists.
   - **Notes:** No explicit statement connects standardized y used for filtering to the $y\text{-hat}$ used for integrated-Y aperture sums and bias plots. Without a documented rescaling, the interpretation of photometric bias and the comparison between methods is analytically ambiguous.

Limitations

- The paper provides no explicit linear mixing/observation equation defining how each component enters each frequency channel (spectral responses, beam factors, units), which prevents a complete symbolic verification of Eq. (2).
- No derivation is shown for the Wiener solution (normal equations), so verification is limited to checking plausibility and internal definition consistency rather than step-by-step algebra.
- Power-spectrum conventions (normalization, masking, beam deconvolution) are not defined; these affect dimensional consistency checks for S_ℓ, N_ℓ, and $s_{\ell,\text{tSZ}}$.

Numerical results audit

This section audits **numerical/empirical** consistency: reported metrics, experimental design, baseline comparisons, statistical evidence, leakage risks, and reproducibility.

Six candidate numerical/internal-consistency checks were executed and all passed. Verified items include the enumeration and stated range of six frequency channels (90–857 GHz), dimensional consistency of $6 \times 6/6$-element objects in the MWF description, correct parsing/positivity of the regularization parameter $\epsilon = 10^{-5}$, ordering of SNR thresholds ($4 < 5$), and a sanity check relating a 3-pixel-radius aperture to its expected area/discrete pixel count.

Checked items

1. ✔ **CAND-001** (Page 3, Section 2.1 (Simulated sky model and observations))
   - **Claim:** The dataset comprises six frequency channels: three from SO LAT at 90, 150, and 217 GHz, and three from Planck HFI at 353, 545, and 857 GHz.
   - **Checks:** count_and_range_check
   - **Verdict:** PASS
   - **Notes:** Parsed listed frequencies (90, 150, 217, 353, 545, 857 GHz); verified count=6, min=90, max=857.
2. ✔ **CAND-002** (Page 1 (Abstract) and Page 10 (Conclusions))
   - **Claim:** Observations are across six frequency channels from 90 to 857 GHz.
   - **Checks:** range_consistency_with_enumeration
   - **Verdict:** PASS
   - **Notes:** Verified enumerated list has 6 elements, all within $[90, 857]$ GHz, and includes endpoints 90 and 857.
3. ✔ **CAND-003** (Page 4, Section 2.2.1 (Multi-frequency wiener filter))
   - **Claim:** S_ℓ is the 6×6 covariance matrix ... $s_{\ell,\mathrm{tSZ}}$ is a 6-element vector ... $(S_\ell + N_\ell)$ is regularized before inversion.
   - **Checks:** dimensionality_internal_consistency
   - **Verdict:** PASS
   - **Notes:** Checked $n_{\mathrm{matrix}} = 6$ and $n_{\mathrm{vector}} = 6$ are consistent for $(6 \times 6) \cdot (6)$ operations.
4. ✔ **CAND-004** (Page 4, Section 2.2.1 (Multi-frequency wiener filter))
   - **Claim:** Tikhonov regularization ... adding a small term, ϵI with $\epsilon = 10^{-5}$, to its diagonal before inversion.
   - **Checks:** numeric_parsing_and_magnitude_sanity
   - **Verdict:** PASS
   - **Notes:** Parsed unicode-minus expression as 1×10^{-5} and verified $\epsilon > 0$.
5. ✔ **CAND-005** (Page 5, Section 2.3.2 (Cluster detection and catalog analysis))
   - **Claim:** Cluster candidates were identified ... above thresholds of SNR > 4 and SNR > 5.
   - **Checks:** threshold_ordering_consistency
   - **Verdict:** PASS
   - **Notes:** Verified strict ordering of thresholds: $4 < 5$.
6. ✔ **CAND-006** (Page 5, Section 2.3.2 (Cluster detection and catalog analysis))
   - **Claim:** Integrated Compton-Y ... calculated by summing the pixel values within a 3-pixel radius aperture.
   - **Checks:** derived_quantity_area_check
   - **Verdict:** PASS
   - **Notes:** For $r = 3$ px: $\pi r^2 = 28.2743$; integer lattice count for $dx^2 + dy^2 \leq 9$ is 29; absolute difference 0.7257 within $abs_tol = 5$ (discrete inclusion convention dependent).

Limitations

- Only parsed text was available; numeric values embedded in figures are not explicitly retrievable without pixel/value extraction, which is out of scope.
- The paper provides very few explicit numeric totals/percentages; many quantitative claims are qualitative (e.g., 'orders of magnitude', '≈ 1') and require underlying simulation outputs to verify.
- No tables of results (purity/completeness counts, bias/scatter statistics) are present in the text; thus most performance claims cannot be fast-checked from explicit numbers.