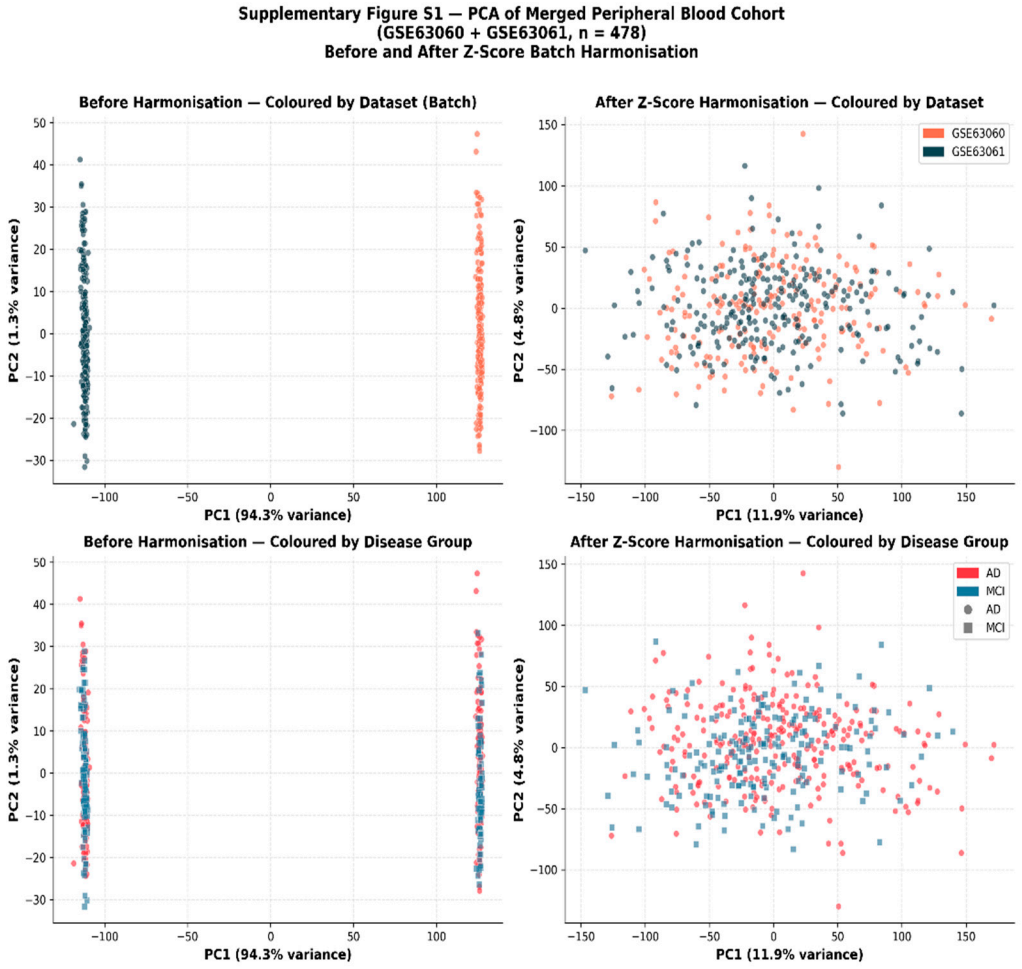



**Figure S1.** Principal Component Analysis (PCA) of the Merged Peripheral Blood Cohort (GSE63060 + GSE63061, n = 478) Before and After Z-Score Batch Harmonisation.

| Condition | Factor | F-statistic | $R^2$ (variance explained) | p-value | Permutations | Significance |
|---|---|---|---|---|---|---|
| Before z-score harmonisation | Dataset (batch) | 7903.324 | 100.0% | 0.0010 | 999 | ** |
| After z-score harmonisation | Dataset (batch) | 0.000 | 0.0% | 1.0000 | 999 | n.s. |

**Table S1**. PERMANOVA results quantifying batch effect before and after z-score harmonisation in the merged peripheral blood cohort (GSE63060 + GSE63061, n = 478; 999 permutations).

| k (filter size) | Cumulative ω-mass (%) | BalAcc Mean (%) | BalAcc Std (%) | Note |
|---|---|---|---|---|
| 10 | 0.3% | 61.69 | ±3.77 | |
| 15 | 0.4% | 64.12 | ±3.58 | proposed |
| 20 | 0.5% | 65.30 | ±4.75 | |
| 30 | 0.7% | 68.41 | ±4.88 | |
| 50 | 1.1% | 69.50 | ±4.49 | |

**Table S2**: NPO optimisation; increasing k in the full pipeline is constrained by NPO search-space capacity.

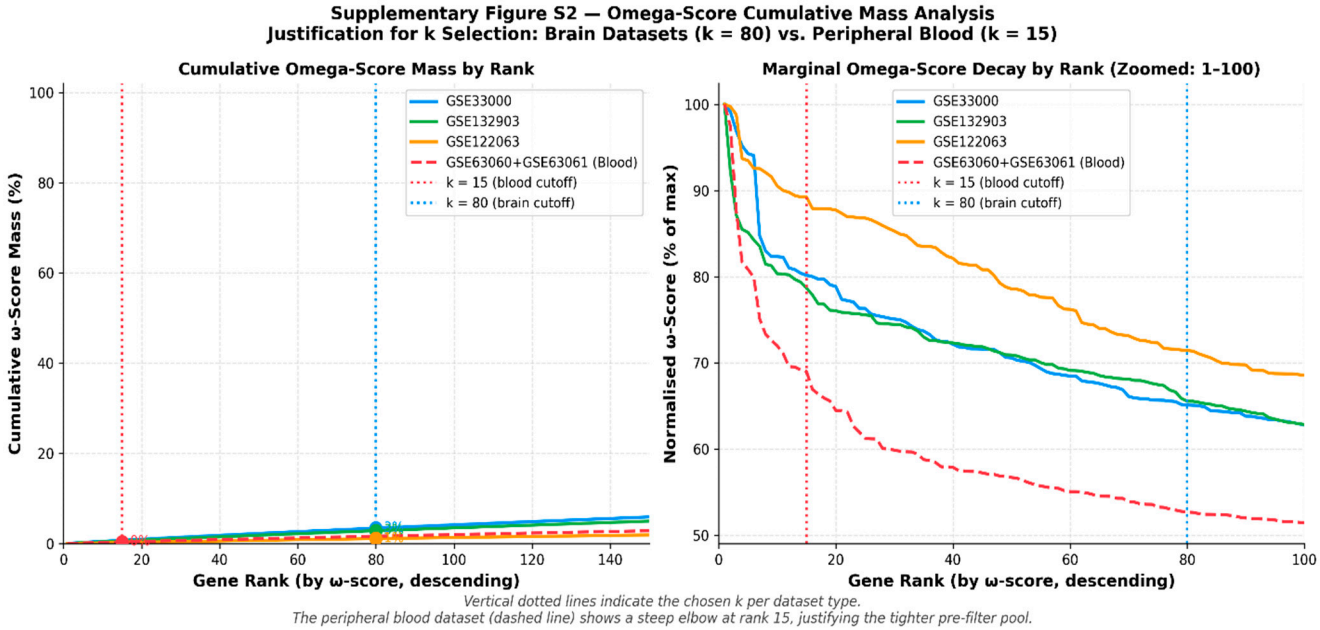


**Figure S2**. Omega-Score Cumulative Mass Analysis: Empirical Justification for Filter Pool Size Selection (k = 80 for Brain Datasets; k = 15 for Peripheral Blood).

| Weight Configuration [w_F, w_MI, w_W] | GSE33000 BalAcc (%) | GSE132903 BalAcc (%) | GSE122063 BalAcc (%) | Grand Mean (%) | Δ vs. Proposed |
|---|---|---|---|---|---|
| [0.40, 0.30, 0.30] Proposed | 85.13 | 84.61 | 93.03 | 87.59 | — |
| [0.33, 0.33, 0.33] Equal | 85.28 | 84.63 | 93.94 | 87.95 | +0.36 |
| [0.50, 0.25, 0.25] | 84.64 | 83.58 | 93.03 | 87.08 | −0.50 |
| [0.25, 0.50, 0.25] | 85.28 | 84.61 | 94.55 | 88.14 | +0.56 |
| [0.25, 0.25, 0.50] | 84.64 | 84.63 | 95.05 | 88.11 | +0.52 |
| [0.60, 0.20, 0.20] | 84.97 | 84.61 | 91.92 | 87.16 | −0.42 |
| [0.20, 0.60, 0.20] | 85.13 | 84.63 | 94.34 | 88.03 | +0.45 |

Table S3. Omega-score weight perturbation analysis across seven configurations on three brain datasets.

All configurations evaluated using the 5-member soft-voting ensemble (ExtraTrees + SVM-RBF + LR + RF + GradientBoosting) with stratified 5-fold cross-validation. The omega filter was fitted inside each training fold (leakage-free). The NPO wrapper was not applied so as to isolate the effect of filter-weight variation. Wilcoxon signed-rank test (proposed vs. equal weights): W = 21.5, p = 0.54 (not significant); n = 15 fold evaluations (5 folds × 3 datasets). w_F = Fisher Discriminant Ratio weight; w_MI = Mutual Information weight; w_W = Welch's t-test weight.

| Dataset | True-label stable probes (≥60%) | True-label max frequency | Null mean freq. (100 permutations) | Null SD | Z-score |
|---|---|---|---|---|---|
| GSE33000 | 23 | 80% | 20.8% | 2.9% | >10 |
| GSE132903 | 15 | 100% | 20.8% | 2.9% | >10 |
| GSE122063 | 14 | 100% | 20.8% | 2.9% | >10 |

*Table S4. Permutation control for biomarker selection frequency.*