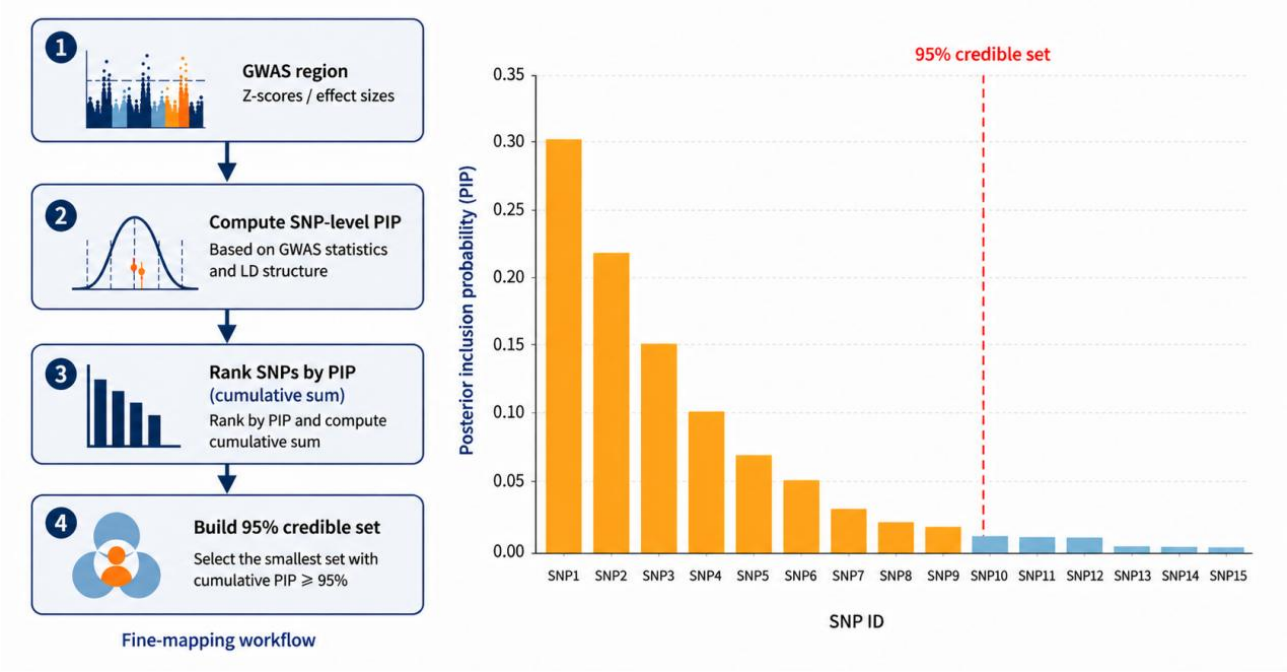

Figure 1 Workflow and illustrative example of credible set construction based on posterior inclusion probability (PIP)
Note: The left panel illustrates the core statistical workflow of fine-mapping: starting from GWAS summary statistics (e.g., Z-scores or effect sizes) within a genomic region, SNP-level posterior inclusion probabilities (PIPs) are computed, ranked, and cumulatively summed to construct a credible set under a predefined coverage threshold (e.g., 95%). The right panel presents an example PIP distribution, where each bar represents the PIP of a candidate SNP, and the dashed line indicates the cutoff at which the cumulative probability reaches 95%

In practical applications, the size of a credible set is not fixed, but is jointly determined by the complexity of local LD structure and the number of underlying causal variants. When LD relationships are relatively simple and only a single causal variant exists within a region, the credible set is usually compact. By contrast, under high-LD backgrounds or in the presence of multiple causal variants, the candidate set often expands substantially, which essentially reflects increased statistical non-identifiability (Hutchinson et al., 2019). In recent years, calibrated credible sets have further improved coverage estimation, bringing it closer to the underlying causal architecture (Shrestha et al., 2024). Therefore, this probability-set-based analytical framework not only enhances the interpretability of fine-mapping results, but also provides a relatively unified input form for functional annotation integration and cross-omics analysis.

2.2 Paradigm shift from single-variant testing to probabilistic inference

From the perspective of statistical inference, the difference between fine-mapping and traditional single-variant GWAS testing is not merely a refinement of analytical steps, but a fundamental shift in both the inferential target and the object of inference. Traditional GWAS mainly relies on single-SNP tests and determines whether a locus is associated with a trait through a significance threshold (Schaid et al., 2018). This approach is efficient for genome-wide scanning, but it rests on a strong simplifying assumption, namely, that the most significant SNP can serve as a proxy for the causal variant. This assumption often fails in the presence of complex LD structure or multiple causal signals, thereby leading to incorrect attribution or bias in the prioritization of candidate loci.

In contrast, Bayesian fine-mapping does not attempt to select a single “best” locus directly from significance results. Rather, it reformulates the problem as the estimation of the posterior distribution over all possible causal configurations, conditional on the observed data and the LD structure. Within this framework, PIP provides a continuous probabilistic characterization of the causal plausibility of an individual locus, whereas the credible set defines the candidate space under a prespecified probability constraint. For example, if three SNPs within a region have PIPs of 0.60, 0.25, and 0.10, respectively, they jointly constitute a 95% credible set, rather than only the