SRSM 2019 Title and Abstract

AUTHORS: Adam R. Hafdahl (presenter), Michael T. Brannick

TITLE: Less-parametric prediction intervals for meta-analysis via mixtures and resampling

ABSTRACT: An increasingly popular result from frequentist random-effects meta-analysis is a prediction interval (PrI), which may be 2-sided or 1-sided (i.e., prediction bound). Such a PrI is a set of plausible values for the effect-size (ES) parameter in another study from the same universe; for instance, 95% of random 95% PrIs contain a random new study's ES parameter. Frequentist PrIs are similar to but distinct from credibility intervals in validity generalization as well as Bayesian PrIs from a predictive distribution. In this work we propose a frequentist PrI meant to be less sensitive to violated assumptions about the ES parameter's distribution and the sampling distribution of estimators of that ES-parameter distribution's (hyper)parameters (e.g., mean, variance). To motivate this, we view a primary study's classical Student-*t* PrI for a new observation as two quantiles of the following prediction distribution: a normal whose mean and variance parameters follow a distribution that reflects sampling error in their estimates. This location-scale mixture distribution is just a normal averaged (i.e., integrated, marginalized) over its parameters' distribution, as if we sampled a mean and variance then sampled from their normal distribution.

Our meta-analytic PrI for a univariate ES is just the appropriate 1 or 2 quantiles of a prediction distribution constructed in 4 steps: 1. Estimate the between-studies mean and variance (BSMV) with a random-effects method. 2. Use those point estimates to approximate the ES parameter's distribution as a mixture (over studies) of normals—1 for each study's ES parameter, akin to empirical Bayes posteriors. 3. Simulate a distribution of BSMVs by resampling many simulated meta-analyses' ES estimates via ES parameters from Step 2's mixture, using each resample to estimate the BSMV, and obtaining random BSMVs from sampling-error estimates. 4. Use each resample's random BSMV to obtain a mixture of normals as in Step 2, then combine resamples' mixtures into a double mixture (over studies and resamples). We can obtain this prediction distribution's quantile(s), such as .025 and .975 for an equal-tail 95% PrI or .10 for a lower 90% prediction bound, by root finding. Unlike most meta-analytic PrIs, this method's double mixture incorporates between-studies and sampling variation with only mild distributional assumptions.

We evaluated our proposed PrI, a few simpler variants of it, and a few easy or popular methods using a Monte Carlo study. In conditions crossing number of studies, size of between-studies and within-study variances, and other factors, we simulated meta-analyses with generic ideal ESs (i.e., normal, unbiased ES estimator with known variance) and popular variance-stabilized ES metrics, drawing ES parameters from normal and non-normal distributions. On PrI coverage probability and other criteria, no method considered here outperformed all others consistently across all conditions, but unlike other methods our proposed PrI was usually among

the best and rarely among the worst. We mention extensions (e.g., study-level covariates, multivariate ESs, heteroscedastic ESs, functions of ESs) and other avenues for further development. This is joint work with Michael T. Brannick.