Multivariate random-effects meta-analysis for sparse data using `smvmeta`

Christopher James Rose

Center for Epidemic Interventions Research
and
Cluster for Reviews and Health Technology Assessments
Division for Health Services
Norwegian Institute of Public Health
Oslo, Norway

Northern European Stata Conference, Oslo, Norway
10 September 2024

Conflicts of interest

I have no conflicts of interest to declare [1]

[1] I'm the author of the method and `smvmeta` command I'm talking about today

Multivariate meta-analysis, sparse data, and `smvmeta`

Meta-analysis

- Systematic review — identify and synthesize all relevant evidence
- Meta-analysis — "pool" estimates from multiple studies into a single estimate
- Univariate (pairwise) meta-analysis — synthesize multiple estimates of one variate

Meta-analysis

- Systematic review — identify and synthesize all relevant evidence
- Meta-analysis — "pool" estimates from multiple studies into a single estimate
- Univariate (pairwise) meta-analysis — synthesize multiple estimates of one variate
- Example using Stata's `bcgset` data on efficacy of BCG vaccine for tuberculosis

Study	BCG Vaccine Yes	BCG Vaccine No	Control Yes	Control No	Risk ratio with 95% CI	Weight (%)
Aronson, 1948	4	119	11	128	0.41 [0.13, 1.26]	5.06
Ferguson & Simes, 1949	6	300	29	274	0.20 [0.09, 0.49]	6.36
Rosenthal et al., 1960	3	228	11	209	0.26 [0.07, 0.92]	4.44
Hart & Sutherland, 1977	62	13,536	248	12,619	0.24 [0.18, 0.31]	9.70
Frimodt-Moller et al., 1973	33	5,036	47	5,761	0.80 [0.52, 1.25]	8.87
Stein & Aronson, 1953	180	1,361	372	1,079	0.46 [0.39, 0.54]	10.10
Vandiviere et al., 1973	8	2,537	10	619	0.20 [0.08, 0.50]	6.03
TPT Madras, 1980	505	87,886	499	87,892	1.01 [0.89, 1.14]	10.19
Coetzee & Berjak, 1968	29	7,470	45	7,232	0.63 [0.39, 1.00]	8.74
Rosenthal et al., 1961	17	1,699	65	1,600	0.25 [0.15, 0.43]	8.37
Comstock et al., 1974	186	50,448	141	27,197	0.71 [0.57, 0.89]	9.93
Comstock & Webster, 1969	5	2,493	3	2,338	1.56 [0.37, 6.53]	3.82
Comstock et al., 1976	27	16,886	29	17,825	0.98 [0.58, 1.66]	8.40
Overall					0.49 [0.34, 0.70]	

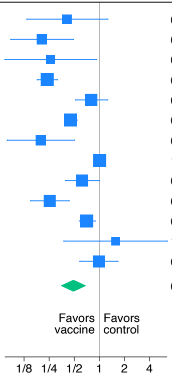

Heterogeneity: $\tau^2 = 0.31$, $I^2 = 92.22\%$, $H^2 = 12.86$

Test of $\theta_i = \theta_j$: $Q(12) = 152.23$, $p = 0.00$

Test of $\theta = 0$: $z = -3.97$, $p = 0.00$

Random-effects REML model

Note: "effect size" often used in place of variate, estimate, etc.

Heterogeneity and random-effects meta-analysis

- Simplest meta-analysis model assumes all studies share the same estimation target
- However, this is rarely the case in biomedicine and other fields
  (e.g., some differences between populations, treatments, outcomes, etc.)
- Differences in estimation targets is called *heterogeneity*

Heterogeneity and random-effects meta-analysis

- Simplest meta-analysis model assumes all studies share the same estimation target
- However, this is rarely the case in biomedicine and other fields
  (e.g., some differences between populations, treatments, outcomes, etc.)
- Differences in estimation targets is called *heterogeneity*
- Heterogeneity is often dealt with using *random-effects* (RE) meta-analysis
- Estimation target in RE meta-analysis is a *distribution* of study-level targets
- Usually modelled using a normal distribution, $N(\mu, \tau^2)$

Multivariate meta-analysis

- Multivariate case — each study reports a multivariate estimate
- The total number of variates studied by all included studies is p
- Each study can report estimates for between 1 and p variates

Multivariate meta-analysis

- Multivariate case — each study reports a multivariate estimate
- The total number of variates studied by all included studies is p
- Each study can report estimates for between 1 and p variates
- There may be heterogeneity and the variates may be *correlated*
- The classic example is diagnostic test accuracy:
  - the variates are sensitivity and specificity
  - sensitivity and specificity are correlated (increasing one tends to decrease the other)
- Failing to account for correlation $\rightarrow$ bias
  (Riley, 2009; Riley, Thompson, and Abrams, 2008)
- Higher-dimensional applications include the study of risk factors

Example multivariate meta-analysis

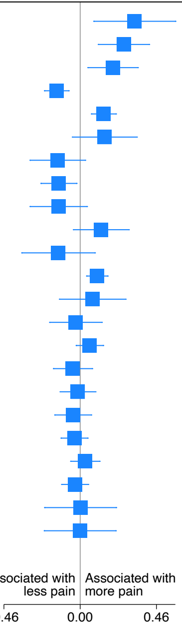

Risk Factor	Correlation	95% CI	*p*-value	*k*	*P*-score	I^2
ROCF Recall	0.34	0.09 to 0.55	0.005	1	92.4%	.
Catastroph.	0.28	0.12 to 0.42	0.000	2	89.1%	59.9%
Temp. Summ.	0.21	0.05 to 0.36	0.005	2	77.2%	0.0%
K-L Grade	-0.15	-0.23 to -0.08	0.000	3	65.3%	0.0%
S. Joints	0.15	0.07 to 0.23	0.000	2	65.1%	0.0%
Heat Thold	0.16	-0.05 to 0.36	0.173	1	61.9%	.
Urban/Semi	-0.15	-0.32 to 0.03	0.125	1	60.1%	.
C. Retain.	-0.14	-0.25 to -0.02	0.009	1	60.0%	.
Education	-0.14	-0.32 to 0.04	0.156	1	58.4%	.
Support	0.14	-0.04 to 0.31	0.172	1	57.1%	.
Expectation	-0.14	-0.37 to 0.10	0.341	1	56.9%	.
Pain	0.11	0.04 to 0.18	0.000	8	50.5%	97.3%
Warm Thold	0.08	-0.14 to 0.29	0.654	1	44.5%	.
Pain S.-E.	-0.03	-0.20 to 0.14	0.913	1	38.9%	.
BMI	0.06	-0.03 to 0.15	0.182	4	35.2%	70.9%
Men. Health	-0.05	-0.17 to 0.08	0.646	5	34.9%	79.6%
Function	-0.02	-0.13 to 0.10	0.942	2	34.8%	90.1%
Pat. Resur.	-0.05	-0.17 to 0.07	0.603	1	33.7%	.
Male	-0.04	-0.13 to 0.05	0.508	4	29.2%	0.0%
Comorbidity	0.03	-0.06 to 0.13	0.676	3	29.0%	0.0%
Older Age	-0.03	-0.12 to 0.05	0.577	4	28.5%	0.0%
Kinesophob.	0.00	-0.23 to 0.23	1.000	1	23.6%	.
Surg. Dur.	-0.00	-0.23 to 0.23	1.000	1	23.6%	.

Sparse multivariate random-effects meta-analysis model

Meta-analysis and meta-regression in Stata

- `smvmeta` is a new add-on command for random-effects multivariate meta-analysis: (Rose, C. J., The Stata Journal 24:2, 2024)
- Stata already has excellent built-in support for meta-analysis:
  - `meta esize` for computing effect sizes from summaries (e.g., big and little *N*s)
  - `meta summarize` and `meta forestplot` for meta-analysis
  - `meta regress` for meta-regression
  - `meta meregress` and `meta multilevel` for multilevel meta-regression
  - `meta mvregress` for multivariate meta-regression
- Where does `smvmeta` fit in?

Multivariate meta-analysis with sparse data

- In the most general case, multivariate meta-analysis models are parameterized by:
  - a p-vector of mean effect sizes β
  - a $p \times p$ correlation matrix $\mathbf{\Phi}$
  - a $p \times p$ variance-covariance matrix $\mathbf{\Psi}$ that models heterogeneity
- Such a model will have $p + p^2$ parameters (note that this is quadratic in p)

Multivariate meta-analysis with sparse data

- In the most general case, multivariate meta-analysis models are parameterized by:
  - a p-vector of mean effect sizes β
  - a $p \times p$ correlation matrix $\mathbf{\Phi}$
  - a $p \times p$ variance-covariance matrix $\mathbf{\Psi}$ that models heterogeneity
- Such a model will have $p + p^2$ parameters (note that this is quadratic in p)
- The numbers of papers reporting multivariate results tends to be small
- We are likely to have far fewer study estimates than parameters
- This is what I mean by sparsity — it makes estimation challenging

Multivariate meta-analysis with sparse data

- In the most general case, multivariate meta-analysis models are parameterized by:
  - a p-vector of mean effect sizes β
  - a $p \times p$ correlation matrix $\mathbf{\Phi}$
  - a $p \times p$ variance-covariance matrix $\mathbf{\Psi}$ that models heterogeneity
- Such a model will have $p + p^2$ parameters (note that this is quadratic in p)
- The numbers of papers reporting multivariate results tends to be small
- We are likely to have far fewer study estimates than parameters
- This is what I mean by sparsity — it makes estimation challenging
- Previous models address this by requiring us to specify or assume how variates are correlated (e.g., using published values or assumed matrix structures)
- Studies typically do not publish correlation estimates; we may be unable or unwilling to make these assumptions — this is where `smvmeta` fits in

Multivariate meta-analysis with sparse data

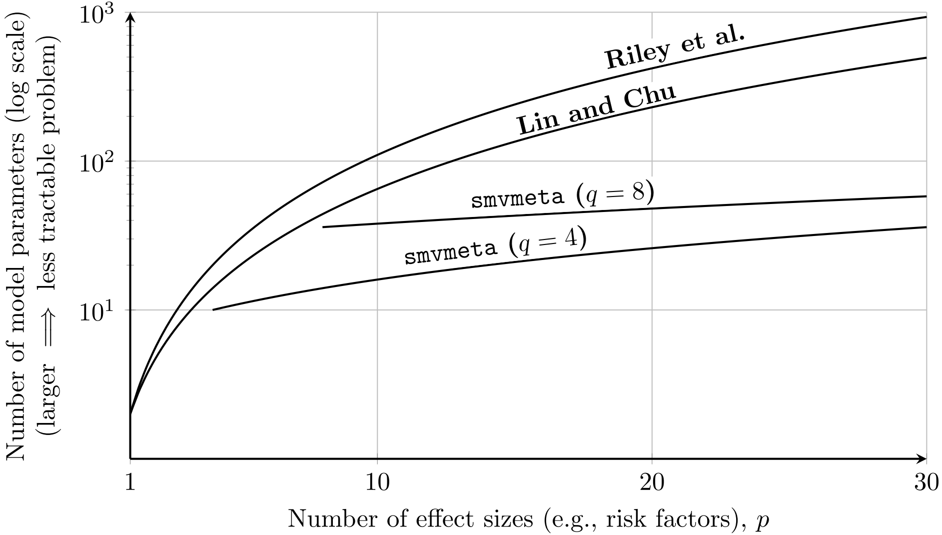

In previous models (e.g., Riley, Lin and Chu), the number of model parameters scales quadratically with p. The model used by `smvmeta` scales linearly with p.

How to use `smvmeta`

The `pain12` dataset

- Systematic review on risk factors for pain and function after total knee arthroplasty (TKA) (Olsen et al., 2020, 2022, 2023)
- $\sim$ 20% of patients experience pain and poor function after surgery
- Identifying risk factors could lead to better outcomes

The `pain12` dataset

- Systematic review on risk factors for pain and function after total knee arthroplasty (TKA) (Olsen et al., 2020, 2022, 2023)
- $\sim 20\%$ of patients experience pain and poor function after surgery
- Identifying risk factors could lead to better outcomes
- Studies typically estimate associations on a range of metrics (e.g., RRs, ORs, MDs, correlations)
- Meta-analysis must be performed on a common metric:
  1. Extracted the reported estimates
  2. Imputed the corresponding tetrachoric correlations
  3. Fisher *z*-transformed [2] the correlations

[2] Inverse hyperbolic tangent function

Install `smvmeta`, ancillary files (incl. the `pain12` dataset), and list:

```
. net install st0749.pkg , all

. use pain12, clear

. list in 1/10
```

	study	factor	z	z_se
1.	Lingard 2007	Men. Health	.0242904	.038518
2.	Papakostidou 2012	Pain	.1760553	.0713318
3.	Papakostidou 2012	Older Age	-.0094856	.0745588
4.	Papakostidou 2012	Male	-.0968132	.0735721
5.	Papakostidou 2012	BMI	-.0395895	.0744007
6.	Papakostidou 2012	Support	.136985	.0725892
7.	Papakostidou 2012	Education	-.1422099	.0724378
8.	Papakostidou 2012	Urban/Semi	-.1482175	.0722573
9.	Sullivan 2011	Male	-.0012088	.0945578
10.	Sullivan 2011	Older Age	-.0012088	.0945578

Use smvmeta set to specify generic effect sizes, SEs, and risk factors:

```
. smvmeta set z z_se factor
Meta-analysis setting information
  Data
    No. observations: 51
              Sparse: Yes
  Effect size
                Type: Generic
               Label: Correlation with pain (Fisher z-transformed)
            Variable: z
             Missing: 0
  Precision
                Type: Standard error
               Label: SE on correlation
            Variable: z_se
             Missing: 0
  Factor variable
               Label: Risk Factor
            Variable: factor
              Levels: 23
             Missing: 0
  Model and method
               Model: Sparse multivariate random-effects meta-analysis (smvmeta)
              Method: Penalized maximum likelihood
```

Perform the meta-analysis using smvmeta estimate:

```
. smvmeta estimate, dim(3) nolog transform(corr) superior(big) sort(_Pscore, descending)

Sparse multivariate meta-analysis (smvmeta)      Number of obs    =         51
Factor variable label : Risk Factor              Num. factors     =         23
Optimization          : Penalized ML             Dimensions (q)   =          3
------------------------------------------------------------------------------
            |  Coef.   Eff. SE     P>u  [95% Conf. Int.]    k   P-score    I2
------------+-----------------------------------------------------------------
ROCF Recall |  0.342     0.126   0.005   0.089     0.554    1      92.4     .
Catastroph. |  0.278     0.074   0.000   0.118     0.424    2      89.1  59.9
Temp. Summ. |  0.211     0.076   0.005   0.050     0.361    2      77.2   0.0
K-L Grade   | -0.154     0.036   0.000  -0.230    -0.075    3      65.3   0.0
S. Joints   |  0.153     0.032   0.000   0.073     0.230    2      65.1   0.0
Heat Thold  |  0.159     0.117   0.173  -0.052     0.355    1      61.9     .
Urban/Semi  | -0.147     0.097   0.125  -0.318     0.033    1      60.1     .
...
Comorbidity |  0.032     0.077   0.676  -0.063     0.126    3      29.0   0.0
Older Age   | -0.035     0.062   0.577  -0.121     0.053    4      28.5   0.0
Kinesophob. |  0.001     0.095   1.000  -0.229     0.232    1      23.6     .
Surg. Dur.  | -0.001     0.095   1.000  -0.232     0.229    1      23.6     .
------------------------------------------------------------------------------
Estimates have been tanh-transformed.
Untransformed effect sizes with larger magnitudes are superior.
```

Use smvmeta forestplot to make the forest plot I showed earlier

How `smvmeta` works

The `smvmeta` model

- Recall that in previous models, the number of model parameters is quadratic in p
- The number of parameters in the `smvmeta` model is *linear* in p, not quadratic:
  - `smvmeta` does not attempt to decompose correlation and heterogeneity
    $\rightarrow$ only need to estimate a single covariance matrix $\boldsymbol{\Psi}$
  - `smvmeta` reduces the dimensionality of the problem by approximating $\boldsymbol{\Psi}$ in a low-dimensional space using *random projection*

The `smvmeta` model

- Random projection is similar to principal components analysis (PCA)
  - establish a random (i.e., arbitrary) low-dimensional orthonormal basis in $\mathbb{R}^{q<p}$ (i.e., a random matrix $\boldsymbol{R}$ of orthogonal unit vectors)
  - let $\boldsymbol{\Sigma}$ be a $q \times q$ covariance matrix (cf. $p \times p$)
  - approximate $\boldsymbol{\Psi}$ as $\boldsymbol{R\Sigma R}^\top$ (i.e., project $\boldsymbol{\Sigma}$ from $\mathbb{R}^q$ up into $\mathbb{R}^p$)
- Estimate an approximation to the distribution of study-level targets $N\left(\boldsymbol{\beta}, \boldsymbol{R\Sigma R}^\top\right)$
- Estimation performed using penalized maximum likelihood (see paper)
- Like PCA, good approximation with small q (e.g., $q \approx 6$ rather than $p \approx 25$)
- Use domain knowledge to choose q, or let `smvmeta` choose it automatically

Why and when you should trust `smvmeta`

Simulation-based validation

- Some studies may not report estimates for all variates
- An estimate for a particular variate may be missing from a study:
  - completely at random (MCAR)
  - because of its value or “significance” (MNAR)

Simulation-based validation

- Some studies may not report estimates for all variates
- An estimate for a particular variate may be missing from a study:
  - completely at random (MCAR)
  - because of its value or "significance" (MNAR)
- We may have — or lack — domain knowledge for choosing q

Simulation-based validation

- Some studies may not report estimates for all variates
- An estimate for a particular variate may be missing from a study:
  - completely at random (MCAR)
  - because of its value or "significance" (MNAR)
- We may have — or lack — domain knowledge for choosing q
- Simulated a large number of systematic reviews with known groundtruth to
  - estimate empirical coverage of `smvmeta`'s 95% CIs
  - compare `smvmeta`'s bias and precision to random-effects meta-regression

  under 3 scenarios:

Scenario	**Missingness mechanism**	**Is p known?**
1	MCAR	No
2	MCAR	Yes
3	MNAR	No

Summary of validation experiment results

- Results for the 3 scenarios are very similar
- Empirical coverage of 95% CIs is 94.4% (95% CI 94.1% to 94.7%)
- Bias is 1.03 (95% CI 1.02 to 1.04) times larger for `smvmeta` compared to meta-regression (but negligible in absolute terms)
- `smvmeta`'s estimates are more precise than for meta-regression (e.g., 95% CIs are about 90% as wide)

Summary of validation experiment results

- Results for the 3 scenarios are very similar
- Empirical coverage of 95% CIs is 94.4% (95% CI 94.1% to 94.7%)
- Bias is 1.03 (95% CI 1.02 to 1.04) times larger for `smvmeta` compared to meta-regression (but negligible in absolute terms)
- `smvmeta`'s estimates are more precise than for meta-regression (e.g., 95% CIs are about 90% as wide)
- `smvmeta` tends to perform better as p increases:
  - confidence intervals tend to get shorter compared to meta-regression
  - bias can be very high if p is small ($p \approx 10$)
  - not surprising, given `smvmeta` is based on random projection
- The method appears to be robust to the MNAR scenario
- Recommend `smvmeta` for sparse problems with $p \geq 15$ variates

Summary

Summary

- Introduced multivariate meta-analysis and the problem of sparsity
- Explained how to use `smvmeta` for multivariate random-effects meta-analysis
- Outlined how the method works (see the paper for more details)
- Summarized the results of simulation-based validation experiments

Thanks for listening

References

- Olsen, Unni et al. (2020). “Predictors of chronic pain and level of physical function in total knee arthroplasty: a protocol for a systematic review and meta-analysis”. In: *BMJ Open* 10.9, e037674.
- Olsen, Unni et al. (2022). “Factors Correlated With Physical Function 1 Year After Total Knee Arthroplasty in Patients With Knee Osteoarthritis: A Systematic Review and Meta-analysis”. In: *JAMA Network Open* 5.7, e2219636.
- — (2023). “Factors correlated with pain after total knee arthroplasty: A systematic review and meta-analysis”. In: *PLOS ONE* 18.3, e0283446.
- Riley, Richard D (2009). “Multivariate meta-analysis: the effect of ignoring within-study correlation”. In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 172.4, pp. 789–811.
- Riley, Richard D, John R Thompson, and Keith R Abrams (2008). “An alternative model for bivariate random-effects meta-analysis when the within-study correlations are unknown”. In: *Biostatistics* 9.1, pp. 172–186.
- Rücker, Gerta and Guido Schwarzer (2015). “Ranking treatments in frequentist network meta-analysis works without resampling methods”. In: *BMC Medical Research Methodology* 15.1, pp. 1–9.

Questions?

The `smvmeta` model

n-vector of effect size estimates (known)

$n \times p$ design matrix indicating the variates (known)

$\mathbb{R}^p \leftrightarrow \mathbb{R}^q$ (constructed)

$$\boldsymbol{y} \sim N\left(\boldsymbol{X}\boldsymbol{\beta}, \boldsymbol{\Lambda} + \boldsymbol{X}\boldsymbol{R}\boldsymbol{\Sigma}\boldsymbol{R}^\top\boldsymbol{X}^\top\right)$$

$q \times q$ covariance matrix (unknown)

p-vector of mean effect sizes (unknown)

$n \times n$ diagonal matrix of sampling variances for the effect sizes (known)

- $\boldsymbol{X}$ could include covariates for meta-regression (not yet implemented)
- $\boldsymbol{\beta}$ and $\boldsymbol{\Sigma}$ are estimated using penalized maximum likelihood (inverse-Wishart penalty applied to $\boldsymbol{\Sigma}$ to prevent trivial solutions; see paper)

Assessing superiority

- Natural research questions include "Which risk factor is most important?" and "Which treatment is best?" (superiority)
- `smvmeta` assesses superiority using *P*-scores (Rücker and Schwarzer, 2015)
  - measures the mean extent of certainty that a given effect size is superior to all others
  - *P*-scores are distinct from *p*-values
- Some preferable properties over posterior probabilities of superiority (e.g., `pbest( )` option of `mvmeta`; see paper)
  - particularly useful for meta-analyses of correlations
  - can be computed almost instantaneously (no MCMC needed)
- `smvmeta` provides an option to specify what superiority means in your meta-analysis