Bootstrap Inference for Multiple Imputation Under Uncongeniality

Jonathan Bartlett
www.thestatsgeek.com
www.missingdata.org.uk

Department of Mathematical Sciences
University of Bath, UK

Joint Statistical Meetings, 1st August 2018

Acknowledgement

This research made use of the Balena High Performance Computing (HPC) Service at the University of Bath.

Outline

Motivation

- MI is very popular, for many reasons, part of which are the simplicity of Rubin's rules.
- If imputations are 'proper' and imputation and analysis models are 'congenial':
  - Rubin's variance estimator is asymptotically unbiased
  - Confidence intervals attain nominal coverage
- Under uncongeniality, Rubin's variance estimator can be biased upwards or downwards, depending on setting - Meng 1994 [2], Wang and Robins 1998 [6].

Motivation

- When the imputer and analyst are the same, but we do not have congeniality, in some settings we may want to obtain the sharpest (valid) inference possible.
- e.g. using control based MI for missing data in confirmatory phase 3 randomised clinical trials.
- Here Rubin's rule variance estimator is biased upwards.
- For particular settings, we may be able to derive valid analytical variance estimators.
- For continuous endpoints analysed using mixed models, Tang 2017 [4] derived the following delta method variance estimator...

Tang 2017 [4]

4.2. *Alternative Variance Estimator via the Delta Method*

The asymptotic sampling variance of the MI estimator can be obtained via the delta method

LEMMA 7. $\mathrm{var}(\bar{\alpha}_p^\infty) = \sum_{t=1}^{p} \frac{\partial \bar{\alpha}_p^\infty}{\partial \underline{\theta}_t} \mathrm{var}(\hat{\underline{\theta}}_t)(\frac{\partial \bar{\alpha}_p^\infty}{\partial \underline{\theta}_t})' + V_\delta^e (X_f' X_f)^{-1}[X_a' X_a](X_f' X_f)^{-1}$, *where* $\mathrm{var}(\hat{\underline{\theta}}_t) = \sigma_t^2 (Z_{o_t}' Z_{o_t})^{-1}$, *and* $\frac{\partial \bar{\alpha}_p^\infty}{\partial \underline{\theta}_t} = l_{pt}[\Upsilon_t - n_{1_+}(X_f' X_f)^{-1}(\sum_{s=0}^{p-1} \pi_{1s} \bar{x}_{1s} J_{t_s}')]$.

Let $\frac{\partial \bar{\delta}_p^\infty}{\partial \underline{\theta}_t} = l_{pt}(J_{\delta_t}' - \sum_{s=0}^{p-1} \pi_{1s} v_s J_{t_s}')$. The variance of the treatment effect at visit p is

$$\mathrm{var}(\bar{\delta}_p^\infty) = \sum_{t=1}^{p} \left(\frac{\partial \bar{\delta}_p^\infty}{\partial \underline{\theta}_t}\right) \mathrm{var}(\hat{\underline{\theta}}_t) \left(\frac{\partial \bar{\delta}_p^\infty}{\partial \underline{\theta}_t}\right)' + V_\delta^e \left[\frac{1}{n_{1+}} + h_x\right]. \tag{16}$$

Bootstrap alternatives

- Deriving and implementing such variance estimators is hard, and model specific.
- What other options do we have?
- Recently Schomaker and Heumann 2018 [3] investigated four combinations of bootstrap with MI.
- von Hippel 2018 [5] has also proposed a bootstrap MI combination approach.
- We investigate which are valid under uncongeniality, and of these, which are computationally efficient.
- We will assume sample size is sufficiently large such that the MI estimator is normally distributed.

Outline

Rubin's rules MI

- Parameter of interest θ.
- Impute M times, and estimate θ, yielding $\hat{\theta}_m$, $m = 1, .., M$.
- $\hat{\theta}_M = M^{-1} \sum_{m=1}^{M} \hat{\theta}_m$.
- Imputation specific estimates follow

$$\hat{\theta}_m = \hat{\theta}_\infty + a_m$$

where $\hat{\theta}_\infty = \lim_{M \to \infty} \hat{\theta}_M$, $\text{Var}(\hat{\theta}_\infty) = \sigma^2_\infty$, $E(a_m) = 0$, $\text{Var}(a_m) = \sigma^2_{\text{btw}}$

Rubin's rules MI

- The variance of $\hat{\theta}_M$ is thus

$$\mathrm{Var}(\hat{\theta}_M) = \sigma^2_\infty + \frac{\sigma^2_{\mathrm{btw}}}{M}$$

- Under congeniality $\sigma^2_\infty = \sigma^2_{\mathrm{btw}} + \sigma^2_{\mathrm{wtn}}$, which leads to Rubin's variance estimator:

$$(1 + M^{-1})\frac{1}{M-1}\sum_{m=1}^{M}(\hat{\theta}_m - \hat{\theta}_M)^2 + M^{-1}\sum_{m=1}^{M}\widehat{\mathrm{Var}}(\hat{\theta}_m)$$

Outline

MI boot Rubin

1. Impute M times
2. For $m = 1, .., M$, generate B nonparametric bootstraps
3. $\hat{\theta}_{m,b}$ estimate from imputation m, bootstrap b
4. For imputation m, then estimate σ^2_{wtn} by

$$\widehat{\text{Var}}_{bs}(\hat{\theta}_m) = (B-1)^{-1} \sum_{b=1}^{B} (\hat{\theta}_{m,b} - \tilde{\theta}_m)^2$$

   where $\tilde{\theta}_m = B^{-1} \sum_{b=1}^{B} \hat{\theta}_{m,b}$
5. Rubin's rules applied to $\hat{\theta}_m$ and $\widehat{\text{Var}}_{bs}(\hat{\theta}_m)$, $m = 1, .., M$

Inference is based on Rubin's rules, so we don't expect unbiased variance estimates under uncongeniality

MI boot pooled

As per MI boot Rubin, except at the final stage, a $(1-2\alpha)\%$ percentile confidence interval for θ is formed by taking the α and $1-\alpha$ empirical percentiles of the pooled MB sample of $\hat{\theta}_{m,b}$ values.

Assuming the estimator is unbiased, point estimates follow

$$\hat{\theta}_{m,b} = \hat{\theta}_{\infty} + a_m + b_b$$

where $\text{Var}(a_m) = \sigma^2_{\text{btw}}$ and $\text{Var}(b_b) = \sigma^2_{\text{wtn}}$.

MI boot pooled

For large B the corresponding MI boot pooled variance estimator is approximately unbiased for

$$(1 - M^{-1})\sigma^2_{\text{btw}} + \sigma^2_{\text{wtn}}$$

Thus for large M and B this will be close to Rubin's variance estimator, and hence be unbiased under congeniality.

However, for small M, it is biased downwards and intervals expected to undercover (under congeniality), as Schomaker and Heumann found.

Inference is again based (essentially) on Rubin's rules, so we don't expect unbiased variance estimates under uncongeniality

Outline

Boot MI

1. Bootstrap B times
2. For $b = 1, .., B$, impute M times
3. Let $\hat{\theta}_b = M^{-1} \sum_m \hat{\theta}_{b,m}$
4. Form percentile intervals based on $\hat{\theta}_b$, or alternatively a Wald interval based on

$$\text{Var}_{BootMI} = (B-1)^{-1} \sum_{b=1}^{B} (\hat{\theta}_b - \hat{\theta}_{BM})^2 \tag{1}$$

where $\hat{\theta}_{BM} = B^{-1} \sum_{b=1}^{B} \hat{\theta}_b$

Boot MI

The point estimates $\hat{\theta}_{bm}$ now follow

$$\hat{\theta}_{bm} = \hat{\theta}_{\infty} + c_b + a_m$$

with $\mathrm{Var}(c_b) = \sigma^2_{\infty}$ and $\mathrm{Var}(a_m) = \sigma^2_{\text{btw}}$

It follows that Var_{BootMI} is unbiased for $\sigma^2_{\infty} + \frac{\sigma^2_{\text{btw}}}{M}$.

We expect unbiased variance estimation under congeniality or uncongeniality

Boot MI pooled

The same as Boot MI, but form percentile intervals based on pooled sample of $\hat{\theta}_{b,m}$.

Schomaker and Heumann found this overcovered in simulations (under congeniality).

For large B and M, the variance of the pooled sample estimates $\sigma^2_\infty + \sigma^2_{\text{btw}}$, and hence is biased upwards, explaining the overcoverage.

We would not expect nominal coverage, under congeniality or uncongeniality

Boot MI for inference under uncongeniality

Boot MI is the only approach we expect to give unbiased variance estimates under uncongeniality.

We need relatively large B for reliable estimates of variance.

If we choose M small, point estimator is inefficient, and Monte-Carlo error may be larger than desired.

If we choose M large, BM is large, and computationally costly!

von Hippel's boot MI proposal

von Hippel [5] proposed using boot MI, with $\hat{\theta}_{BM}$ as the point estimator

Its variance is

$$\text{Var}(\hat{\theta}_{BM}) = (1 + B^{-1})\sigma^2_{\infty} + (BM)^{-1}\sigma^2_{\text{btw}}$$

We can fit a one way random intercepts model to the estimates $\hat{\theta}_{b,m}$ to estimate σ^2_{∞} and σ^2_{btw}, and insert into the preceding expression.

Since large B is required for reliable variance estimates, von Hippel suggested using $M = 2$.

With $M = 2$, the approach becomes computationally much less costly.

Outline

Simulation setup

Sample size $n = 500$.

Binary 'treatment' randomly assigned.

Y_1, Y_2 (baseline,follow-up) generated from correlated bivariate normal, with mean of Y_2 dependent on 'treatment'.

50% of Y_2 values made missing completely at random.

Analysis model is linear regression of Y_2 on treatment and Y_1, and interest focuses on the treatment coefficient.

10,000 simulations

Imputation methods

Each of the previously described combinations was used with $M = 10$ and $B = 200$

Except, Boot MI von Hippel, which used $B = 200$ and $M = 2$

First we imputed Y_2 using normal linear regression under MAR.

Next we impute Y_2 using the jump to reference MNAR approach, proposed by Carpenter *et al* [1]. This imputation model is uncongenial with the analysis model.

Results under congeniality (MAR imputation)

	Emp. SD	Est. SD	Med. CI width	CI coverage
MI Rubin	0.082	0.082	0.327	95.2
MI boot Rubin	0.082	0.082	0.327	95.1
MI boot pooled	0.082	0.078	0.301	93.4
Boot MI	0.082	0.082	0.321	95.1
Boot MI pooled	0.082	0.098	0.383	98.0
Boot MI von Hippel	0.080	0.080	0.315	95.1

MI boot pooled downward biased slightly, as expected.

Boot MI pooled biased upwards, as expected.

Results under uncongeniality (J2R imputation)

	Emp. SD	Est. SD	Med. CI width	CI coverage
MI Rubin	0.045	0.051	0.200	97.5
MI boot Rubin	0.045	0.051	0.200	97.5
MI boot pooled	0.045	0.050	0.197	97.3
Boot MI	0.045	0.044	0.175	94.8
Boot MI pooled	0.045	0.047	0.185	96.1
Boot MI von Hippel	0.044	0.044	0.174	94.9

Only Boot MI and Boot MI von Hippel are unbiased for the true repeated sampling variance.

All the others overestimate the variance, and hence CIs overcover.

Outline

Conclusions

- Under uncongeniality, bootstrap followed by MI can provide unbiased variance estimation and intervals which attain nominal coverage.
- von Hippel's version of this is attractive on computational efficiency grounds.
- Importantly, its application requires no customisation to the particular imputation/analysis model, unlike analytic alternatives.
- We have assumed:
  - the estimator is normally distributed
  - data are i.i.d. (c.f. stratified randomization)
- These slides at `www.thestatsgeek.com`

References

[1] J R Carpenter, J H Roger, and M G Kenward.
Analysis of longitudinal trials with protocol deviations: a framework for relevant, accessible assumptions and inference via multiple imputation.
Journal of Biopharmaceutical Statistics, 23:1352–1371, 2013.

[2] X L Meng.
Multiple-imputation inferences with uncongenial sources of input (with discussion).
Statistical Science, 10:538–573, 1994.

[3] M Schomaker and C Heumann.
Bootstrap inference when using multiple imputation.
Statistics in Medicine, 37(14):2252–2266, 2018.

[4] Y Tang.
On the multiple imputation variance estimator for control-based and delta-adjusted pattern mixture models.
Biometrics, 73(4):1379–1387, 2017.

[5] P. T. von Hippel.
Maximum likelihood multiple imputation: Faster, more efficient imputation without posterior draws.
ArXiv e-prints, 2018.
1210.0870v9.

[6] N Wang and J M Robins.
Large-sample theory for parametric multiple imputation procedures.
Biometrika, 85:935–948, 1998.