Sign-flipping score based confidence for testing in generalized linear models

Antonio Calcagnì*, Jesse Hemerik†, Jelle Goeman†, and Livio Finos*

* University of Padova
† Leiden University Medical Center

NetIBS-IR 2019
Napoli, July 7-10

Introduction

GLMs can suffer from misspecifications due to:

- overdispersion
- heteroscedasticity
- nuisance parameters

$\rightarrow$ hypothesis testing is often problematic in these conditions

Introduction

Permutation tests offer a way out.

With respect to the parametric counterpart, they:

- require less assumptions
- provide exact control of Type I error
- asymptotically converge to the parametric tests
- easily work in the multivariate case (e.g., multiplicity correction)

Some limitations:

- presence of continuous confounders as nuisance (observations are no longer exchangeable)

Introduction

A novel solution for GLMs with continuous nuisance confounders:

Robust testing in generalized linear models by sign-flipping score contributions

Jesse Hemerik*†
Jelle Goeman*and
Livio Finos‡

September 5, 2018

Abstract

Generalized linear models are often misspecified due to overdispersion, heteroscedasticity and ignored nuisance variables. Existing quasi-likelihood methods for testing in misspecified models often do not provide satisfactory type-I error rate control. We provide a novel semi-parametric test, based on sign-flipping individual score contributions. The tested parameter is allowed to be multi-dimensional and even high-dimensional. Our test is often robust against the mentioned forms of misspecification and provides better type-I error control than its competitors. When nuisance parameters are estimated, our basic test becomes conservative. We show how to take nuisance estimation into account to obtain an asymptotically exact test. Our proposed test is asymptotically equivalent to its parametric counterpart.

Sign-flipping score

Let $(y_1, \ldots, y_n)$ be iid random realizations from $(Y_1, \ldots, Y_n)$ with density (exp. family):

$$f(y_i|\theta_i) = h(y_i, \tau) \exp\left(\frac{y_i\theta_i - b(\theta_i)}{a_i(\tau)}\right)$$

Then:

$$g(\mathbb{E}[Y_i]) = b'(\theta_i) = g^{-1}(\mathbf{x}_i\boldsymbol{\beta} + \mathbf{z}_i\gamma) \qquad \mathbf{z}_i \text{ continuous confounder}$$
$$\mathbb{VAR}[Y_i] = b''(\theta_i)a_i(\tau)$$

Sign-flipping effective score test *in a nutshell*:

$$H_0 : \beta = \beta_0|\hat{\gamma}$$
$$s^*_{\hat{\gamma}} = \sum_{i=1}^{n} \pm \left(\frac{\partial}{\partial\beta} \log f(Y_i|\beta,\gamma)|_{\beta=\beta_0} - \hat{\mathcal{J}}_{xz}^T \hat{\mathcal{J}}_{zz}^{-1} \frac{\partial}{\partial\gamma} \log f(Y_i|\beta,\gamma)|_{\beta=\beta_0}\right)$$

$\mathcal{J}_{()}$ block of the observed Fisher information under H_0

Sign-flipping score

Let $(y_1, \ldots, y_n)$ be iid random realizations from $(Y_1, \ldots, Y_n)$ with density (exp. family):

$$f(y_i|\theta_i) = h(y_i, \tau) \exp\left(\frac{y_i\theta_i - b(\theta_i)}{a_i(\tau)}\right)$$

Then:

$$g(\mathbb{E}[Y_i]) = b'(\theta_i) = g^{-1}(\mathbf{x}_i\boldsymbol{\beta} + \mathbf{z}_i\gamma) \qquad \mathbf{z}_i \text{ continuous confounder}$$
$$\mathbb{VAR}[Y_i] = b''(\theta_i)a_i(\tau)$$

Sign-flipping effective score test *in a nutshell*:

- robust against misspecification of the model
- better control of type I error for small and large sample size

Permutation confidence intervals

Few studies have investigated the problem of how constructing **permutation confidence intervals** (pCIs).

Some examples:

- Pesarin & Salmaso (2010): grid-search algorithm
- Garthwaite & Buckland (1992): stochastic Robbins-Monro algorithm
- Pauly, Asendorf, & Konietschke (2014): permutation-based range preserving CIs for the Behrens-Fisher problem (inference for the AUC)

Permutation confidence intervals

Few studies have investigated the problem of how constructing **permutation confidence intervals** (pCIs).

Some examples:

- Pesarin & Salmaso (2010): grid-search algorithm
- Garthwaite & Buckland (1992): stochastic Robbins-Monro algorithm
- Pauly, Asendorf, & Konietschke (2014): permutation-based range preserving CIs for the Behrens-Fisher problem (inference for the AUC)

Related problems in computing pCIs:

- constructing the search space for $\beta_{|H_0}$
- multidimensional $H_0^{(1:p)}$ treated as set of univariate $H_0^{(1)}, \ldots, H_0^{(p)}$
- multidimensional case $H_0^{(1:p)}$ can be computationally prohibitive

Random-search grid algorithm

In general:

- β_l and β_u computed separately
- $\{(\beta_\text{l}, \beta_\text{u}) \in \Theta : \text{p-value} = \alpha\}$

Our idea: algorithm for pCIs based on sign-flipping score test:

- given $\beta_0^{(1)}, \ldots, \beta_0^{(M)}$, a set of separated hypotheses are tested:
  $H_0 : \beta_0 = \beta_0^{(m)} \quad (m = 1, \ldots, M)$
- $\beta_0^{(1)}, \ldots, \beta_0^{(M)}$ is defined via random grid search (inspired by Bhat et al., 2018)
- early-stopping rule as soon $|T^*_{S(m)}| \geq T_\alpha$ (cost or stress function)
  $T^*_{S(m)}$ permuted sign-flipping score statistics, T_α observed alpha level
- adaptive restarting rule

Random-search grid algorithm

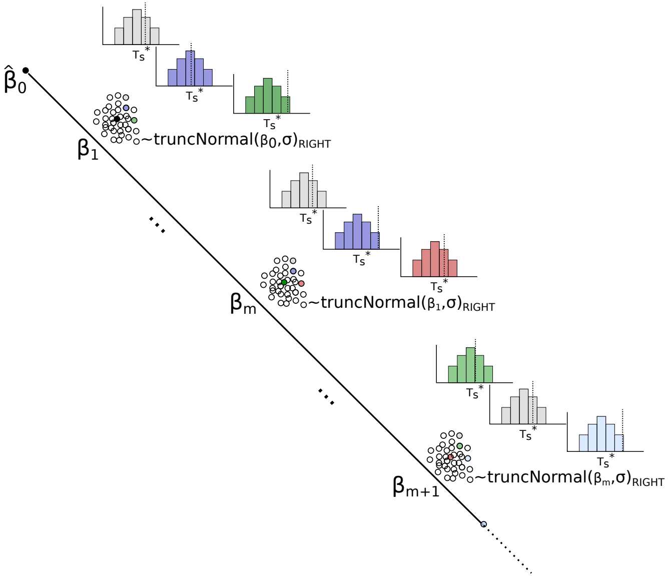

Random-search grid algorithm

Basic instructions for $\mathrm{pCI}_{\mathrm{lb}}$

```
double β = β̂                                   //initial estimate
int K                                          //number of candidates
double[K] x
int i = 0
double α = 0.05
do{
 i = i + 1
 x = truncnormal(K,β,σ,lb=-Inf,ub=β)

 for(k=1:K){
  pv[k] = effective_score(...,x[k])}

 [β,ρβ] = choose_best_beta(pv,x)              // x with min p-value

} while(ρβ < α)                                //stopping rule
```

Simulation

The algorithm is still under testing and preliminary results are available.

Short simulation study:

- $(y_1, \ldots, y_n) \sim \text{NegBinomial}\,(\mu_i, \phi)$, $n = 100$, $\phi = [1.2, 2.4]$
- $\mu_i = \exp\,([1\ x_i\ z_i]\boldsymbol{\beta})$, $\text{cor}(\mathbf{x}, \mathbf{z}) = 0.8$
- Inference and testing are conducted under the *wrong* Poisson model
- Outcome: sign-flipping pCIs, asymptotic Poisson and NegBinom CIs measured on *probability coverage* and *interval length* (the shortest, the better)
- Replicates: 500, Permutations: 1000

Simulation

Results:

			Poisson		*NegBinom*		*Flip*	
ϕ	β_0	$\hat{\beta}_0$	cover	length	cover	length	cover	length
1.20	0.195	0.198	0.564	0.315	0.932	0.472	0.956	0.497
2.40	1.59	1.57	0.688	0.590	0.936	0.943	0.928	1.071

$(1-\alpha) = 0.95$

Conclusions

- the method is robust to misspecification and overdispersion
- it works well with both LMs and GLMs
- it can be trivially extended to handle with multivariate hypotheses
- it requires parallel computing to speed up the search
- the algorithm is still under developing and further studies are needed

antonio.calcagni@unipd.it