Data 102 Lecture 12:

Uncertainty quantification for GLMs

Why uncertainty quantification?

Lecture overview

- Definitions
  - Types of uncertainty quantification
    - Confidence intervals vs credible intervals
    - Prediction intervals
  - Desirable properties
- Algorithms
  - Credible intervals for Bayesian GLM
  - Model-based confidence intervals for frequentist GLM
  - Bootstrap
    - Bootstrap in general
    - Bootstrap for GLM
  - (Optional) Prediction intervals for frequentist GLM

Definitions

Confidence intervals vs credible intervals

See whiteboard notes

Prediction intervals

See whiteboard notes

Coverage, width, validity

See whiteboard notes

Algorithms

Credible intervals for Bayesian GLM

```
[13]: arviz.plot_trace(poisson_trace)
```

```
[13]: array([[<AxesSubplot:title={'center':'Intercept'}>,
              <AxesSubplot:title={'center':'Intercept'}>],
             [<AxesSubplot:title={'center':'x'}>,
              <AxesSubplot:title={'center':'x'}>],
             [<AxesSubplot:title={'center':'mu'}>,
              <AxesSubplot:title={'center':'mu'}>]], dtype=object)
```

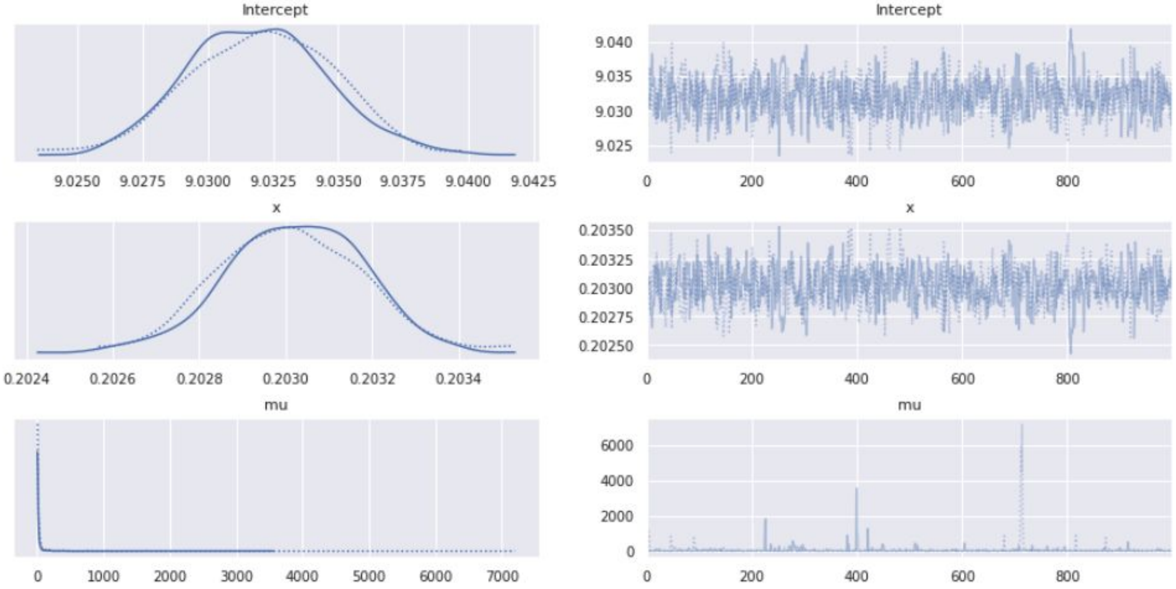

Credible intervals for Bayesian GLM

Model-based confidence intervals for frequentist GLM

Assume that the model specification is **correct**, i.e. there is a true β_o such that the data is generated from

$$y_i = g^{-1}(X_i^T \beta_0) + \epsilon_i, \qquad \mathbb{E}[\epsilon_i | X_i] = 0$$

Then we have statistical theory that controls the **estimation error** $\hat{\beta} - \beta_0$

Asymptotically,

$$\sqrt{n}(\hat{\beta}_n - \beta_0) \Rightarrow \mathcal{N}(0, I_n(\beta_0)^{-1})$$

Model-based confidence intervals for frequentist GLM

```python
negbin_model = sm.GLM(
    ok_turbines.totals, sm.add_constant(ok_turbines.year),
    family=sm.families.NegativeBinomial()
)
negbin_results = negbin_model.fit()
print(negbin_results.summary())
```

```
                 Generalized Linear Model Regression Results
==============================================================================
Dep. Variable:                 totals   No. Observations:                   17
Model:                            GLM   Df Residuals:                       15
Model Family:        NegativeBinomial   Df Model:                            1
Link Function:                    log   Scale:                          1.0000
Method:                          IRLS   Log-Likelihood:                -134.14
Date:                Wed, 17 Feb 2021   Deviance:                       7.1483
Time:                        12:51:51   Pearson chi2:                     1.90
No. Iterations:                    11
Covariance Type:            nonrobust
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
const          4.2059      0.544      7.725      0.000       3.139       5.273
year           0.2389      0.043      5.514      0.000       0.154       0.324
==============================================================================
```

$$\beta_{0,year} = 0.24 \pm 0.08$$

Bootstrap

Bootstrap demo: Distribution of sample mean

See notebook demo

Bootstrap demo takeaways

As N increases, bootstrap does better and better

How well it does depends on

- Smoothness of functional
- How well-behaved the distribution is

Limitations

- Confidence intervals only asymptotically valid
- Require smooth functional (bad e.g. maximum of a dist)

Bootstrap demo: GLMs

See notebook demo

Bootstrap demo takeaways

Advantages

- Can give valid CIs even when certain assumptions are dropped
- Bootstrap CIs are less overconfident than model-based CIs in general

Limitations

- Doesn’t work at all when the independence noise assumption fails
- Does not account for model bias, i.e. if regression function part of model is wrong, then bootstrap CI gives uncertainty for the projected model
- Cannot compute prediction intervals