

Bayesian sample size determination for Type A uncertainty evaluation

Jörg Martin
joint work with Clemens Elster

Physikalisch-Technische Bundesanstalt
EMUE Workshop - Paris

January 22nd 2020



What is Sample size determination (SSD)?

Sample size determination (SSD)

Find required number of datapoints n

What is Sample size determination (SSD)?

Sample size determination (SSD)

Find required number of datapoints n to meet a certain *criterion*

What is Sample size determination (SSD)?

Experimental design

Sample size determination

What is Sample size determination (SSD)?

Bayesian experimental design

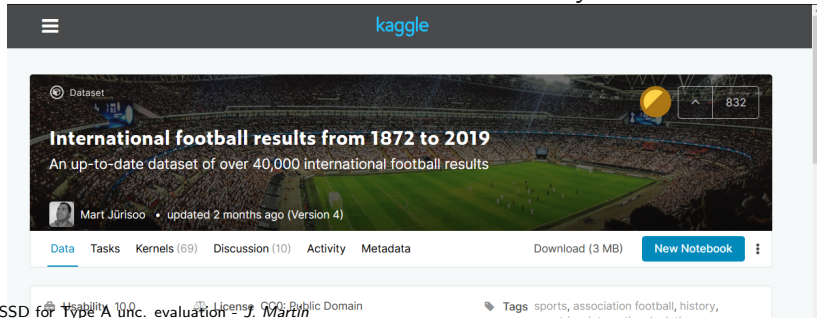
Bayesian sample size determination

Toy example: Football

Q: *Average number of goals in a football match?* $\Rightarrow \theta$

SSD: Minimal number of games to answer "Q" $\Rightarrow n$

Dataset: International football matches in the years 2014-2019



The screenshot shows the Kaggle dataset page for "International football results from 1872 to 2019". The page features a header with the Kaggle logo and a hamburger menu. Below the header is a large banner image of a football stadium at night. The banner contains the text "International football results from 1872 to 2019" and "An up-to-date dataset of over 40,000 international football results". Below the banner is a section for the dataset creator, Mart Jüriso, with a profile picture and the text "updated 2 months ago (Version 4)". Below this section is a row of tabs: "Data", "Tasks", "Kernels (69)", "Discussion (10)", "Activity", and "Metadata". To the right of these tabs are the options "Download (3 MB)" and "New Notebook". At the bottom of the page, there is a section for "Tags" which includes "sports, association football, history, countries, international relations".

Toy example: Football

Q: *Average number of goals in a football match?* $\Rightarrow \theta$

SSD: **Minimal number of games to answer "Q"** $\Rightarrow n$


Dataset: International football matches in the years 2014-2019



The screenshot shows the Kaggle dataset page for "International football results from 1872 to 2019". The page features a header with the Kaggle logo and a hamburger menu. Below the header, there is a large banner image of a football stadium at night. The banner contains the text "International football results from 1872 to 2019" and "An up-to-date dataset of over 40,000 international football results". Below the banner, there is a section for the dataset creator, Mart Jüriso, with a profile picture and the text "updated 2 months ago (Version 4)". At the bottom of the page, there is a navigation bar with links for "Data", "Tasks", "Kernels (69)", "Discussion (10)", "Activity", and "Metadata". On the right side of the navigation bar, there is a "Download (3 MB)" button and a "New Notebook" button. Below the navigation bar, there is a section for "Tags" with the text "sports, association football, history, countries, international relations".

Distribution of goals between 2014-2019


Q: Average number of goals in a football match?



Full dataset: $n = 5784$ games

Subsample of 7 games

Q: Average number of goals in a football match?



$n = 7$ games - **Is this enough to answer Q?**

What is Sample size determination (SSD)?

Sample size determination (SSD)

Find required number of datapoints n to meet a certain *criterion*

What is Sample size determination (SSD)?

Sample size determination (SSD)

Find required number of datapoints n to meet a certain *criterion*

Criterion **here:** Uncertainty below threshold

- θ : Average number of goals in a match
- x : Number of goals in a game, e.g.

$$\text{Final score } 3 : 1 \quad \Rightarrow \quad x = 4$$

We assume $x \sim \text{Poi}(\theta)$

- For n games:

$$\mathbf{x} = (x_1, \dots, x_n),$$

where x_i number of goals in game i

- θ : Average number of goals in a match
- x : Number of goals in a game, e.g.

$$\text{Final score } 3 : 1 \quad \Rightarrow \quad x = 4$$

We assume $x \sim \text{Poi}(\theta)$

- For n games:

$$\mathbf{x} = (x_1, \dots, x_n),$$

where x_i number of goals in game i

We want to find an optimal n **before** we know \mathbf{x}

Results of n games



$$\mathbf{x} = (x_1, \dots, x_n)$$



Estimate for θ + Uncertainty

Example

Results of $n = 3$ games
(2:1, 0:0, 0:3)



$$\mathbf{x} = (3, 0, 3)$$



Estimate for θ + Uncertainty
 $\theta = 2.3, u(\theta|\mathbf{x}) = 0.9$

Example

Results of $n = 3$ games
(2:1, 0:0, 0:3)



$$\mathbf{x} = (3, 0, 3)$$



Estimate for θ + Uncertainty
 $\theta = 2.3$, $u(\theta|\mathbf{x}) = 0.9$

Example

Results of $n = 3$ games
(2:1, 0:0, 0:3)



$$\mathbf{x} = (3, 0, 3)$$



Estimate for θ + Uncertainty
 $\theta = 2.3$, $u(\theta|\mathbf{x}) = 0.9$

SSD: Choose n s.t. $u(\theta|\mathbf{x})$ is small enough (before knowing \mathbf{x})

- *Posterior distribution*

$$\begin{array}{ccccc} \pi(\theta|\mathbf{x}) & \propto & \pi(\theta) & \cdot & p(\mathbf{x}|\theta) \\ \text{posterior} & & \text{prior} & & \text{likelihood} \end{array}$$

(recall: $\mathbf{x} = (x_1, \dots, x_n)$)

- **Uncertainty:** *Posterior variance*

$$u^2(\theta|\mathbf{x}) := \text{Var}(\theta|\mathbf{x}) = \int (\theta - \mathbb{E}[\theta|\mathbf{x}])^2 \cdot \pi(\theta|\mathbf{x}) \, d\theta$$

What is Sample size determination (SSD)?

Sample size determination (SSD)

Find required number of datapoints n to meet a certain *criterion*


Criterion **here:** Uncertainty below threshold

Recipe for SSD

- Describe prior knowledge via *prior* $\pi(\theta)$
- Compute uncertainty $u^2(\theta|\mathbf{x})$ for **generic** \mathbf{x} (*we don't know \mathbf{x}*)

Prior knowledge

Prior knowledge about θ : “**Around 2.50 with unc. 1.0**”



$$\pi(\theta) = \text{Gamma}(\theta | \alpha = 6.25, \beta = 2.50)$$

$$u^2(\theta | \mathbf{x}) = \frac{\alpha + \sum_{i=1}^n x_i}{(n + \beta)^2}$$

Recipe for SSD

- Describe prior knowledge via *prior* $\pi(\theta)$
- Compute uncertainty $u^2(\theta|\mathbf{x})$ for **generic** \mathbf{x} (*we don't know \mathbf{x}*)

Recipe for SSD

- Describe prior knowledge via *prior* $\pi(\theta)$
- Compute uncertainty $u^2(\theta|\mathbf{x})$ for **generic** \mathbf{x} (*we don't know \mathbf{x}*)

Next task:

Choose sample size n so that uncertainty will be small enough

Want something like

$$\text{Var}(\theta|\mathbf{x}) = u^2(\theta|\mathbf{x}) \stackrel{!}{<} \varepsilon^2$$

but *before* measuring \mathbf{x} .

Want something like

$$\text{Var}(\theta|\mathbf{x}) = u^2(\theta|\mathbf{x}) \stackrel{!}{<} \varepsilon^2$$

but *before* measuring \mathbf{x} .

Idea: Average w.r.t. *prior predictive*

$$\pi_{\text{pred.}}(\mathbf{x}) = \int p(\mathbf{x}|\theta)\pi(\theta)d\theta$$

Average posterior variance criterion (**APVC**, e.g. [Wang2002])

$$\overline{u^2} = \mathbb{E}_{\mathbf{x} \sim \pi_{\text{pred.}}(\mathbf{x})} [u^2(\theta|\mathbf{x})] \stackrel{!}{<} \varepsilon^2$$

- $\overline{u^2}$: Expected (squared) uncertainty
- ε^2 : (subjective) threshold

Recipe for SSD (APVC)

- Describe prior knowledge via *prior* $\pi(\theta)$
- Compute uncertainty $u^2(\theta|\mathbf{x})$ for **generic** \mathbf{x} (*we don't know \mathbf{x}*)

Recipe for SSD (APVC)


- Describe prior knowledge via *prior* $\pi(\theta)$
- Compute uncertainty $u^2(\theta|\mathbf{x})$ for **generic** \mathbf{x} (*we don't know \mathbf{x}*)
- Choose n such that $\overline{u^2} < \varepsilon^2$

recall: $\overline{u^2} = \mathbb{E}_{\mathbf{x} \sim \pi_{\text{pred.}}(\mathbf{x})} [u^2(\theta|\mathbf{x})]$

Say we want “ $u^2(\theta|\mathbf{x}) < 0.3^2$ ”:

$$\overline{u^2} < 0.30^2$$

Toy example



Toy example

Say we want “ $u^2(\theta|\mathbf{x}) < 0.3^2$ ”:

$$\overline{u^2} < 0.30^2 \quad \rightarrow \quad n = 26$$


Say we want " $u^2(\theta|\mathbf{x}) < 0.3^2$ ":

$$\overline{u^2} < 0.30^2 \quad \rightarrow \quad n = 26$$

Result for 4 random choices of 26 games:

- ① $\hat{\theta} = 2.36,$ $u(\theta|\mathbf{x}) = \mathbf{0.29}$
- ② $\hat{\theta} = 2.50,$ $u(\theta|\mathbf{x}) = \mathbf{0.30}$
- ③ $\hat{\theta} = 2.78,$ $u(\theta|\mathbf{x}) = \mathbf{0.31}$
- ④ $\hat{\theta} = 3.13,$ $u(\theta|\mathbf{x}) = \mathbf{0.33}$

Fluctuations of u^2



Average posterior variance criterion (APVC):

- puts $u^2(\theta|\mathbf{x})$ in right order of magnitude
- $u^2(\theta|\mathbf{x}) < \varepsilon^2$ only holds on average
- $u^2(\theta|\mathbf{x}) < \varepsilon^2$ is often violated

We propose: *Variation* of posterior variance criterion (VPVC)

Average of posterior variance criterion (APVC)

$$\overline{u^2} < \varepsilon^2$$

Variation of posterior variance criterion (VPVC)

$$\overline{u^2} + 2 \cdot \Delta u^2 < \varepsilon^2$$

VPVC:


$$\overline{u^2} + 2 \cdot \Delta u^2 < \varepsilon^2$$


where

$$\Delta u^2 = \mathbb{E}_{\theta \sim \pi(\theta)} \left[\text{Var}_{\mathbf{x} \sim p(\mathbf{x}|\theta)} (u^2(\theta|\mathbf{x})) \right]$$

Recipe for SSD (VPVC)

- Describe prior knowledge via *prior* $\pi(\theta)$
- Compute uncertainty $u^2(\theta|\mathbf{x})$ for **generic** \mathbf{x} (*we don't know \mathbf{x}*)
- Choose n such that $\overline{u^2} + 2 \cdot \Delta u^2 < \varepsilon^2$





Sample sizes

ε	APVC	VPVC	freq.
0.80	2	3	6
0.70	3	5	8
0.60	5	7	10
0.50	8	11	14
0.40	14	18	21
0.30	26	32	34
0.20	60	70	72
0.10	248	267	270

Sample sizes

ε	APVC	VPVC	freq.
0.80	2	3	6
0.70	3	5	8
0.60	5	7	10
0.50	8	11	14
0.40	14	18	21
0.30	26	32	34
0.20	60	70	72
0.10	248	267	270

APVC

<

VPVC

<

freq.



violates ε



no prior knowledge

- SSD = planning of sample size such that a certain condition is satisfied.
condition here: Restriction on the magnitude of uncertainty.
- Using the prior predictive $\pi_{\text{pred.}}$ this can be guaranteed “on average” (APVC)
- We propose a new method that limits the fluctuations (VPVC)
- Application of the VPVC yields substantially larger sample sizes than the APVC

- **This work:** J. Martin and C. Elster. "The variation of the posterior variance and Bayesian sample size determination." arXiv preprint arXiv:1907.12795 (2019).
- F. Wang and A. E. Gelfand. "A simulation-based approach to Bayesian sample size determination for performance under a given model and for separating models." *Statistical Science* 17.2 (2002): 193-208.
- T. Pham-Gia and N. Turkkan. "Sample size determination in Bayesian analysis." *Journal of the Royal Statistical Society: Series D (The Statistician)* 41.4 (1992): 389-397.