Title Page

Title

Usefulness of pystan and numpyro in Bayesian item response theory

Authors and Affiliations

Mizuho Nishio, MD, PhD[1*], Eiji Ota, MS[2], Hidetoshi Matsuo, MD, PhD[1], Takaaki Matsunaga, MD[1], Aki Miyazaki, MD[1], Takamichi Murakami, MD, PhD[1]

1 Department of Radiology, Kobe University Graduate School of Medicine, 7-5-2 Kusunoki-cho, Chuo-ku, Kobe, Hyogo 650-0017 JAPAN

2 Futaba Numerical Technologies, 18-502 Iruma riverside apartment, 603-1 Bushi, Iruma, Saitama 358-0053 JAPAN

*Corresponding Author

Mizuho Nishio, MD, PhD

Department of Radiology, Kobe University Graduate School of Medicine, 7-5-2 Kusunoki-cho,

Chuo-ku, Kobe, Hyogo 650-0017 JAPAN

Tel.: +81-78-382-6104. Fax: +81-78-382-6129

e-mail: nishiomizuho@gmail.com

Manuscript Type

Original article

Funding information

This work was supported by JSPS KAKENHI (Grant Number: 22K07665 and 23K17229).

Conflict of Interest

There are no conflicts of interest to declare.

Abstract

Purpose: The purpose of this study is to compare two libraries dedicated to Markov chain Monte Carlo method: pystan and numpyro.

Materials and methods: Bayesian item response theory (IRT), 1PL-IRT and 2PL-IRT, were implemented with pystan and numpyro. Then, the Bayesian 1PL-IRT and 2PL-IRT were applied to two types of medical data obtained from a published paper. The same prior distributions of latent parameters were used in both pystan and numpyro. Estimation results of latent parameters of 1PL-IRT and 2PL-IRT were compared between pystan and numpyro. Additionally, the computational cost of Markov chain Monte Carlo method was compared between the two libraries. To evaluate the computational cost of IRT models, simulation data were generated from the medical data and numpyro.

Results: For all the combinations of IRT types (1PL-IRT or 2PL-IRT) and medical data types, the mean and standard deviation of the estimated latent parameters were in good agreement between pystan and numpyro. In most cases, the sampling time using Markov chain Monte Carlo method was shorter in numpyro than that in pystan. When the large-sized simulation data were used, numpyro with a graphics processing unit was useful for reducing the sampling time.

Conclusion: Numpyro and pystan were useful for applying the Bayesian 1PL-IRT and 2PL-IRT.

Keywords

item response theory; Markov chain Monte Carlo; graphics processing unit.

1 Introduction

Item response theory (IRT) is a statistical framework used for analyzing test results and

evaluating test items and test takers quantitatively. While IRT is commonly used in educational

and psychological research (1–3), there are several applications of IRT to medical research. For

example, Choi et al. used IRT for constructing the computer adaptive testing system of short-

form patient-reported outcome measures with the data from the Patient-Reported Outcomes

Measurement Information System project (4). Gershon et al. used IRT to build the quality of life

item banks for adults with neurological disorders (5).

Generally, IRT is applied to the results of binary responses to the test items (e.g., correct

and incorrect answers). In medical diagnosis, the results of various diagnostic procedures are

frequently defined as binary responses. Therefore, it is possible to apply IRT to the data of

medical diagnosis. To apply IRT to the data of medical diagnosis, the following correspondence

is assumed: (i) the patient as the test item, (ii) the doctor as the test taker, and (iii) the results of

the binary responses obtained through medical diagnosis as test results. For example, Nishio et

al. used IRT for analyzing the results of medical diagnoses by radiologists (6).

The Bayesian IRT can be implemented using probabilistic programming languages or

dedicated libraries (e.g., JAGS, Stan, pystan, and numpyro) (7–10). For example, previous

studies used Stan for the implementation of the Bayesian IRT, graded response model, and

nominal response model (6,11,12). The recent advances in hardware and software make it

possible to use the Bayesian IRT efficiently. However, there is no study comparing the

efficiency of the Bayesian IRT from the viewpoint of computational cost.

The purpose of the current study was to compare the results of the Bayesian IRT

implemented with two dedicated libraries (pystan and numpyro). In the current study, the

Bayesian 1PL-IRT and 2PL-IRT implemented with pystan and numpyro were applied to the two

types of medical data obtained from the published paper (6). The estimation results of latent

parameters and the computational cost were compared between pystan and numpyro. For

reproducibility, our implementation of the Bayesian IRT in pystan and numpyro used in the

current study is disclosed as open source through GitHub

(https://github.com/jurader/irt_pystan_numpyro).

2 Materials and methods

Because this study used the medical data obtained from the published paper, institutional review

board approval or informed consent of patients was not necessary.

2.1 Medical Data

The two types of medical data (BONE and BRAIN data) were obtained from the published

paper (6). Table 1 shows the characteristics of the two types of medical data. The BONE data

include binary responses from 60 patients (test items) and 7 radiologists (test takers), and the

BRAIN data include those from 42 patients and 14 radiologists. The total numbers of the binary

responses were 420 and 588 in the BONE and BRAIN data, respectively.

2.2 1PL-IRT

IRT is a statistical model for analyzing the results of binary responses. While there are several

types of IRT models (13), 1PL-IRT and 2PL-IRT were used. In the current study, latent

parameters of IRT are estimated based on the results of medical diagnoses by test takers.

In 1PL-IRT, one latent parameter (β_i) is used to represent the difficulty of test item i, and

another latent parameter (θ_j) is used to represent the ability of test taker j. The following

equations represent 1PL-IRT.

$$\Pr\left(r_{ij} = 1\right) = \frac{1}{1 + \exp\left(-z_{ij}\right)}$$

$$z_{ij} = \theta_j - \beta_i$$

Here,

- $\Pr(r_{ij} = 1)$ represents the probability that the response of test taker j to test item i is correct,

- β_i is the difficulty parameter of test item i,

- θ_j is the ability parameter of test taker j.

2.3 2PL-IRT

In 2PL-IRT, two latent parameters (α_i and β_i) are used to represent test item i. The following equations represent 2PL-IRT.

$$\Pr(r_{ij} = 1) = \frac{1}{1 + \exp(-z_{ij})}$$

$$z_{ij} = \alpha_i(\theta_j - \beta_i)$$

Here,

- α_i and β_i are the discrimination and difficulty parameters of test item i.

2.4 Experiments

We used Google Colaboratory to run the experiments. The following software packages were used on Google Colaboratory: pystan, version 3.3.0; jax, version 0.4.4; jaxlib, version

0.4.4+cuda11.cudnn82; numpyro, version 0.10.1. Two cores of Intel(R) Xeon(R) (2.20 GHz)

and NVIDIA(R) Tesla T4(R) were used as CPU and a graphics processing unit (GPU),

respectively.

2.4.1 Experiments for agreement of latent parameters

The Bayesian 1PL-IRT and 2PL-IRT, implemented with pystan and numpyro, were applied to

the BONE and BRAIN data. For 1PL-IRT, the following prior distributions were used:

- $\beta_i \sim N(0, 2)$,

- $\theta_j \sim N(0, 2)$,

where N represents a normal distribution in which the first and second arguments are the

average and variance of normal distribution, respectively. For 2PL-IRT, the following prior

distribution was used in addition to those of 1PL-IRT:

- $log(\alpha_i) \sim N(0.5, 1)$.

The same prior distributions of the latent parameters were used in both pystan and numpyro.

The following parameters were used for sampling using Markov chain Monte Carlo method in both pytan and numpyro: num_chains=6, num_samples=8000, num_warmup=2000. For numpyro GPU version, chain_method='parallel' was used. After the sampling, the posterior distributions of the latent parameters were obtained for the Bayesian 1PL-IRT and 2PL-IRT. The posterior distributions of the latent parameters were then compared between pystan and numpyro.

2.4.2 Experiments for computational cost

To evaluate the computation cost of 1PL-IRT and 2PL-IRT implemented with pystan and numpyro, simulation data were generated from the medical data (BONE and BRAIN data) and numpyro. The computer simulation was performed in the following steps: (i) estimating the posterior distributions of the latent parameters of 1PL-IRT and 2PL-IRT for the two types of medical data, and (ii) generating binary responses from the IRT equations and the estimated posterior distributions for the two types of medical data. The total number of binary responses in the simulation data were as follows: 420, 840, 2100, 4200, 8400, 21000, 42000, 84000, 210000, and 420000 for the BONE data; 588, 1176, 2940, 5880, 11760, 29400, 58800, 117600, 294000, and 588000 for the BRAIN data. To evaluate the computational cost in pystan and numpyro, the sampling time using Markov chain Monte Carlo method was measured. In numpyro, both CPU

version and GPU version were used for the sampling. The following parameters were used for the sampling in both pytan and numpyro: num_chains=2, num_samples=3000, num_warmup=500. For numpyro GPU version, chain_method='parallel' was used. Due to the limitation of Google Colaboratory, it was not possible to evaluate the sampling time in several large-sized simulation data.

3 Results

In the current study, we focused on the ability parameters of test takers, and the estimation results of test items were omitted. Tables 2–5 present the estimation results of the ability parameters of test takers. In addition, Figures 1 and 2 show representative scatter plots of the estimation results between pytan and numpyro, which are obtained from values of Tables 2 and 5, respectively. Tables 2–5 show the mean, standard deviation, and credible interval (94% highest density interval) as the estimation results of the ability parameters of test takers. Based on Tables 2–5 and Figures 1 and 2, we found that there was good agreement between pystan and numpyro for 1PL-IRT and 2PL-IRT of the BONE and BRAIN data.

From Tables 2–5, Lin's concordance correlation coefficients (CCC) of the estimated mean of the ability parameters were calculated between (a) pystan v.s. numpyro CPU version,

(b) pystan v.s. numpyro GPU version, and (c) numpyro CPU version v.s. numpyro GPU version

(14). The following criteria were used to evaluate CCC (15,16); low CCC values (< 0.900) were

considered to represent poor agreement, whereas higher CCC values represented moderate

(0.900–0.950), substantial (0.951–0.990), and almost perfect agreement (> 0.990). In the current

study, the CCC values of the estimated mean of the ability parameters were as follows: (a)

1.000, (b) 1.000, and (c) 1.000 for 1PL-IRT and 2PL-IRT of the BONE and BRAIN data,

indicating almost perfect agreement.

Figures 3–6 show the sampling time for the simulation data of the BONE and BRAIN

data. When original-size simulation data were used, the sampling time was shorter in pystan

than numpyro CPU version. However, in the simulation data of the BONE and BRAIN data

except for the original size, the sampling time was shorter in numpyro CPU version than pystan.

Moreover, when the large-sized simulation data (total number of binary responses >30000–

50000) were used, the sampling time was shorter in numpyro GPU version than numpyro CPU

version.

4 Discussion

The current study aimed to compare the estimation results of the ability parameter of test takers

using two different libraries (pystan and numpyro) for two different types of IRT models and medical data (BONE and BRAIN data). The study found that there was good agreement between pystan and numpyro for all the combinations of the IRT models and medical data; there was almost perfect agreement between pystan and numpyro in the CCC values of the estimated mean of ability parameters. The current study also compared the sampling time for the simulation data of the BONE and BRAIN data. We found that while the sampling time was shorter in pystan than numpyro CPU version for the original-size data, it was shorter in numpyro CPU version than pystan for the simulation data except for the original size. For the large-sized simulation data, the sampling time was shorter in numpyro GPU version than numpyro CPU version.

Our results show that there was almost perfect agreement in the ability parameters of the Bayesian IRT between pystan and numpyro. This suggests that researchers can choose either library for implementing the Bayesian IRT. While numpyro requires only Python, both Python and Stan (two different programming languages) are necessary for pystan. Many practitioners and researchers may find numpyro to be simple and straightforward.

Although we used the simulation data, our results of sampling time show that the fastest libraries differed based on the total number of binary responses. Specifically, pystan was the fastest for the original-size simulation data, while numpyro CPU version was the fastest for the

small-sized and medium-sized data. For the large-sized simulation data, the sampling time was shorter in numpyro GPU version than numpyro CPU version. This implies that practitioners and researchers should select either pystan or numpyro based on the data size.

Tables 3 and 4 show that the total number of latent parameters may affect the usefulness of GPU for reducing the sampling time. In complex models, such as 2PL-IRT used in this study, numpyro GPU version may tend to be faster than numpyro CPU version. The effect of GPU on the sampling time should be evaluated in future studies.

This study had several limitations. First, we evaluated only Bayesian 1PL-IRT and 2PL-IRT. Further studies are needed to investigate the effectiveness of pystan and numpyro in other types of Bayesian models. Second, although we evaluated the sampling time, we used the simulation data instead of real-world data. Future studies should use real-world data to evaluate the sampling time. Third, we used Google Colaboratory. Although Google Colaboratory has several merits (e.g., ease of use and availability), our experiments were performed using limited types of hardware provided by Google Colaboratory.

In conclusion, the current study demonstrated that both pystan and numpyro were effective in the estimation for 1PL-IRT and 2PL-IRT of the BONE and BRAIN data. Moreover, the study provides useful information about the sampling time for the different data and the libraries using the computer simulations. The results of this study may be helpful for researchers and

practitioners who use these models for large-scale data. Overall, the findings of this study

contribute to the growing body of research on the application of the Bayesian methods in

medical data.

References

1. Cappelleri JC, Jason Lundy J, Hays RD. Overview of classical test theory and item response theory for the quantitative assessment of items in developing patient-reported outcomes measures. Clin Ther. Clin Ther; 2014;36(5):648–662. doi: 10.1016/J.CLINTHERA.2014.04.006.

2. Embretson SE, Reise SP. Item response theory for psychologists. Item Response Theory for Psychologists. Taylor and Francis; 2000. doi: 10.4324/9781410605269/ITEM-RESPONSE-THEORY-PSYCHOLOGISTS-SUSAN-EMBRETSON-STEVEN-REISE.

3. Hays RD, Morales LS, Reise SP. Item response theory and health outcomes measurement in the 21st century. Med Care. 2000;38(9 Suppl):II28-42.

4. Choi SW, Reise SP, Pilkonis PA, Hays RD, Cella D. Efficiency of static and computer adaptive short forms compared to full-length measures of depressive symptoms. Qual Life Res. NIH Public Access; 2010;19(1):136. doi: 10.1007/S11136-009-9560-5.

5. Gershon RC, Lai JS, Bode R, et al. Neuro-QOL: quality of life item banks for adults with neurological disorders: item development and calibrations based upon clinical and general population testing. Qual Life Res. Qual Life Res; 2012;21(3):475–486. doi: 10.1007/S11136-011-9958-8.

6. Nishio M, Akasaka T, Sakamoto R, Togashi K. Bayesian Statistical Model of Item Response Theory in Observer Studies of Radiologists. Acad Radiol. Elsevier; 2020;27(3):e45–e54. doi: 10.1016/J.ACRA.2019.04.014.

7. pystan · PyPI. . https://pypi.org/project/pystan/. Accessed March 21, 2023.

8. Carpenter B, Gelman A, Hoffman MD, et al. Stan: A Probabilistic Programming Language. J Stat Softw. 2017;76(1):1–32. doi: 10.18637/jss.v076.i01.

9. Depaoli S, Clifton JP, Cobb PR. Just Another Gibbs Sampler (JAGS). Journal of Educational and Behavioral Statistics. 2016;41(6):628–649. doi: 10.3102/1076998616664876.

10. Phan D, Pradhan N, Jankowiak M. Composable Effects for Flexible and Accelerated Probabilistic Programming in NumPyro. 2019;

11. Luo Y, Jiao H. Using the Stan Program for Bayesian Item Response Theory: Educational and Psychological Measurement. SAGE PublicationsSage CA: Los Angeles, CA; 2017;78(3):384–408. doi: 10.1177/0013164417693666.

12. Nishio M, Kobayashi D, Matsuo H, Urase Y, Nishioka E, Murakami T. Bayesian multidimensional nominal response model for observer study of radiologists. Jpn J Radiol. Jpn J Radiol; 2022; doi: 10.1007/S11604-022-01366-Y.

13. Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB. Bayesian Data Analysis. Bayesian Data Analysis. Chapman and Hall/CRC; 2013. doi: 10.1201/B16018.

14. Lin LI-K. A Concordance Correlation Coefficient to Evaluate Reproducibility. Biometrics. JSTOR; 1989;45(1):268. doi: 10.2307/2532051.

15. Nishio M, Koyama H, Ohno Y, et al. Emphysema Quantification Using Ultralow-Dose CT With Iterative Reconstruction and Filtered Back Projection. AJR Am J Roentgenol. AJR Am J Roentgenol; 2016;206(6):1184–1192. doi: 10.2214/AJR.15.15684.

16. Kojita Y, Matsuo H, Kanda T, et al. Deep learning model for predicting gestational age after the first trimester using fetal MRI. Eur Radiol. Eur Radiol; 2021;31(6):3775–3782. doi: 10.1007/S00330-021-07915-9.

Figure and figure legends

Note: (A) and (B) are obtained from values of Table 2.

(A)

(B)

Figure 2

Representative scatter plots of estimated ability parameters of 2PL-IRT between numpyro and

pystan for BONE data. (A) plot for mean of estimated ability parameters, (B) plot for SD of

estimated ability parameters.

Note: (A) and (B) are obtained from values of Table 5.

(A)

(B)

Figure 3

Sampling time of 1PL-IRT for simulation BONE data.

Figure 4

Sampling time of 2PL-IRT for simulation BONE data.

Sampling time of 1PL-IRT for simulation BRAIN data.

Figure 6

Sampling time of 2PL-IRT for simulation BRAIN data.

Tables

Table 1

Characteristics of two types of medical data.

Name of data	Number of test takers	Number of test items	Total number of binary responses
BONE data	7	60	420
BRAIN data	14	42	588

Table 2

Estimation results of ability parameters of test takers in 1PL-IRT for BONE data.

Pystan		estimation				diagnostics		
		mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
	theta[0]	2.201	0.501	1.257	3.133	40373	36621	1
	theta[1]	3.387	0.61	2.25	4.531	47056	35326	1
	theta[2]	3.105	0.583	2.055	4.233	44783	36038	1
	theta[3]	2.845	0.556	1.8	3.882	44444	36888	1
	theta[4]	1.383	0.452	0.511	2.205	38698	37129	1
	theta[5]	2.401	0.517	1.443	3.378	40061	36388	1
	theta[6]	2.2	0.5	1.273	3.153	40777	38464	1
numpyro CPU		estimation				diagnostics		
		mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
	theta[0]	2.2	0.508	1.277	3.185	32921	34003	1
	theta[1]	3.391	0.615	2.26	4.567	34162	35373	1
	theta[2]	3.099	0.586	2	4.201	35639	34188	1
	theta[3]	2.845	0.553	1.818	3.888	34989	35124	1
	theta[4]	1.382	0.454	0.517	2.225	29675	34343	1
	theta[5]	2.4	0.519	1.449	3.392	33463	34351	1
	theta[6]	2.198	0.506	1.267	3.174	33284	35760	1

numpyro GPU	estimation				diagnostics		
	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
theta[0]	2.199	0.506	1.242	3.134	32251	34882	1
theta[1]	3.384	0.615	2.252	4.556	34198	36149	1
theta[2]	3.1	0.585	2.015	4.212	36366	35577	1
theta[3]	2.843	0.556	1.786	3.875	34938	33051	1
theta[4]	1.379	0.451	0.538	2.236	30942	34819	1
theta[5]	2.4	0.521	1.466	3.419	32772	34828	1
theta[6]	2.199	0.506	1.262	3.163	32300	36210	1

Table 3

Estimation results of ability parameters of test takers in 1PL-IRT for BRAIN data.

pystan		estimation				diagnostics		
		mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
	theta[0]	2.062	0.49	1.128	2.976	14702	24025	1
	theta[1]	3.086	0.605	1.966	4.231	17011	25730	1
	theta[2]	2.677	0.539	1.664	3.687	16745	26368	1
	theta[3]	2.585	0.571	1.534	3.658	16199	25432	1
	theta[4]	1.252	0.412	0.482	2.029	14322	22764	1
	theta[5]	2.096	0.536	1.102	3.111	13843	23047	1
	theta[6]	2.169	0.511	1.206	3.125	15188	26449	1
numpyro CPU		estimation				diagnostics		
		mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
	theta[0]	2.057	0.486	1.159	2.991	15715	25398	1
	theta[1]	3.086	0.607	1.943	4.196	17702	27769	1
	theta[2]	2.676	0.546	1.679	3.723	16704	26124	1
	theta[3]	2.581	0.576	1.515	3.674	15881	25012	1
	theta[4]	1.248	0.409	0.494	2.037	15187	26085	1
	theta[5]	2.095	0.533	1.106	3.094	15093	24705	1
	theta[6]	2.167	0.511	1.194	3.104	15545	27676	1

numpyro GPU	estimation				diagnostics		
	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
theta[0]	2.063	0.487	1.159	2.988	14397	25263	1
theta[1]	3.09	0.601	1.975	4.226	16891	27636	1
theta[2]	2.68	0.546	1.658	3.701	15578	27217	1
theta[3]	2.585	0.576	1.525	3.682	14965	24243	1
theta[4]	1.251	0.411	0.503	2.046	14136	23883	1
theta[5]	2.096	0.538	1.105	3.112	13277	23510	1
theta[6]	2.171	0.513	1.237	3.162	14117	24946	1

Table 4

Estimation results of ability parameters of test takers in 2PL-IRT for BONE data.

pystan	estimation				diagnostics		
	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
theta[0]	1.615	0.567	0.569	2.695	19440	27845	1
theta[1]	0.929	0.53	-0.044	1.946	17105	28114	1
theta[2]	1.617	0.57	0.561	2.698	19394	29146	1
theta[3]	0.926	0.527	-0.075	1.902	17568	27120	1
theta[4]	1.139	0.537	0.145	2.163	18614	28549	1
theta[5]	1.144	0.543	0.117	2.147	16985	28225	1
theta[6]	-0.641	0.476	-1.557	0.23	14932	25140	1
theta[7]	-1.251	0.474	-2.144	-0.362	15336	26550	1
theta[8]	1.615	0.566	0.548	2.677	19382	29906	1
theta[9]	0.721	0.515	-0.234	1.702	17127	28476	1
theta[10]	1.615	0.566	0.568	2.693	19343	29784	1
theta[11]	1.615	0.567	0.568	2.695	19363	27743	1
theta[12]	0.166	0.492	-0.766	1.078	15641	26534	1
theta[13]	0.346	0.499	-0.607	1.264	16040	27328	1

numpyro CPU	estimation				diagnostics		
	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat

	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
theta[0]	1.615	0.562	0.613	2.722	18474	28342	1
theta[1]	0.921	0.527	-0.056	1.924	17348	27983	1
theta[2]	1.612	0.565	0.545	2.667	19804	27694	1
theta[3]	0.921	0.527	-0.062	1.922	17238	27787	1
theta[4]	1.141	0.539	0.122	2.139	17777	28747	1
theta[5]	1.136	0.538	0.129	2.153	17642	27878	1
theta[6]	-0.643	0.473	-1.51	0.256	15441	24593	1
theta[7]	-1.253	0.474	-2.141	-0.36	15186	24014	1
theta[8]	1.614	0.567	0.555	2.678	19272	27497	1
theta[9]	0.718	0.515	-0.253	1.69	17100	26777	1
theta[10]	1.611	0.566	0.519	2.648	19627	29707	1
theta[11]	1.612	0.568	0.556	2.701	19678	28018	1
theta[12]	0.163	0.491	-0.752	1.093	15522	25928	1
theta[13]	0.339	0.499	-0.584	1.294	15960	25830	1

numpyro GPU			estimation			diagnostics	
	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
theta[0]	1.618	0.561	0.581	2.689	18841	29675	1
theta[1]	0.925	0.527	-0.06	1.914	17399	26758	1
theta[2]	1.615	0.568	0.564	2.697	19508	27450	1
theta[3]	0.923	0.527	-0.057	1.929	17223	27357	1
theta[4]	1.143	0.537	0.151	2.166	17784	29717	1
theta[5]	1.139	0.541	0.157	2.196	17700	28834	1
theta[6]	-0.641	0.474	-1.524	0.247	15313	25672	1
theta[7]	-1.252	0.476	-2.139	-0.354	15033	26341	1

theta[8]	1.617	0.567	0.551	2.672	19673	26825	1
theta[9]	0.722	0.515	-0.275	1.67	17220	26307	1
theta[10]	1.613	0.566	0.559	2.689	19741	28897	1
theta[11]	1.614	0.566	0.546	2.681	19503	27841	1
theta[12]	0.166	0.492	-0.754	1.097	15657	26378	1
theta[13]	0.342	0.498	-0.576	1.307	15644	25302	1

Table 5

Estimation results of ability parameters of test takers in 2PL-IRT for BRAIN data.

pystan	estimation				diagnostics		
	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
theta[0]	1.18	0.531	0.203	2.181	17825	28333	1
theta[1]	0.902	0.496	-0.008	1.843	15625	27710	1
theta[2]	1.254	0.518	0.296	2.235	18845	29397	1
theta[3]	0.704	0.538	-0.3	1.707	14418	26881	1
theta[4]	1.136	0.503	0.226	2.111	17740	28371	1
theta[5]	1.04	0.52	0.096	2.035	17265	28069	1
theta[6]	-0.509	0.394	-1.235	0.242	12631	24324	1
theta[7]	-1.27	0.433	-2.075	-0.446	15054	26335	1
theta[8]	1.458	0.572	0.369	2.512	17362	28080	1
theta[9]	0.94	0.531	-0.024	1.965	15271	26223	1
theta[10]	1.767	0.575	0.696	2.852	18638	29705	1
theta[11]	1.365	0.509	0.411	2.317	18758	29696	1
theta[12]	0.332	0.453	-0.51	1.193	13643	24442	1
theta[13]	0.324	0.465	-0.542	1.209	14035	24973	1
numpyro CPU	estimation				diagnostics		
	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat

	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
theta[0]	1.183	0.528	0.182	2.157	18389	29763	1
theta[1]	0.898	0.488	0.009	1.837	16627	28545	1
theta[2]	1.256	0.515	0.288	2.225	18436	28429	1
theta[3]	0.704	0.535	-0.291	1.711	14966	25748	1
theta[4]	1.134	0.501	0.19	2.071	17305	28056	1
theta[5]	1.044	0.518	0.089	2.031	16558	26045	1
theta[6]	-0.505	0.394	-1.225	0.262	12487	23078	1
theta[7]	-1.265	0.432	-2.078	-0.465	15003	25697	1
theta[8]	1.455	0.567	0.419	2.539	18819	28511	1
theta[9]	0.943	0.53	-0.035	1.942	16732	27505	1
theta[10]	1.766	0.573	0.644	2.801	17756	27724	1
theta[11]	1.365	0.506	0.426	2.326	19432	29606	1
theta[12]	0.334	0.453	-0.51	1.197	12867	24284	1
theta[13]	0.327	0.462	-0.523	1.204	13323	26250	1

numpyro GPU			estimation			diagnostics	
	mean	sd	hdi_3%	hdi_97%	ess_bulk	ess_tail	r_hat
theta[0]	1.186	0.527	0.226	2.197	17302	20916	1
theta[1]	0.897	0.488	0.001	1.827	16923	26620	1
theta[2]	1.251	0.517	0.292	2.22	15879	16665	1
theta[3]	0.702	0.532	-0.303	1.695	16956	26275	1
theta[4]	1.133	0.5	0.204	2.08	18464	28238	1
theta[5]	1.04	0.515	0.081	2.007	16628	28830	1
theta[6]	-0.507	0.392	-1.248	0.228	13277	24086	1
theta[7]	-1.264	0.433	-2.076	-0.458	14157	19916	1

theta[8]	1.455	0.565	0.395	2.504	17398	23571	1
theta[9]	0.942	0.528	-0.051	1.925	17510	22661	1
theta[10]	1.762	0.573	0.644	2.807	16888	24176	1
theta[11]	1.363	0.504	0.414	2.302	17175	12944	1
theta[12]	0.33	0.449	-0.503	1.18	13347	25251	1
theta[13]	0.33	0.463	-0.529	1.203	13968	25446	1