

---

# Probabilistically-autoencoded horseshoe-disentangled multidomain item-response theory models

---

**Joshua C. Chang**  
Epidemiology and Biostatistics  
Rehabilitation Medicine, CC, NIH  
& team med $\varepsilon_{rrata}$   
josh@mederrata.com

Shashaank Vattikuti & Carson C. Chow  
Laboratory of Biological Modeling, NIDDK, NIH  
& team med $\varepsilon_{rrata}$   
{shashaank.vattikuti, carson.chow}@nih.gov

## Abstract

Item response theory (IRT) is a non-linear generative probabilistic paradigm for using exams to identify, quantify, and compare latent traits of individuals, relative to their peers, within a population of interest. In pre-existing multidimensional IRT methods, one requires a factorization of the test items. For this task, linear exploratory factor analysis is used, making IRT a posthoc model. We propose skipping the initial factor analysis by using a sparsity-promoting horseshoe prior to perform factorization directly within the IRT model so that all training occurs in a single self-consistent step. Being a hierarchical Bayesian model, we adapt the WAIC to the problem of dimensionality selection. IRT models are analogous to probabilistic autoencoders. By binding the generative IRT model to a Bayesian neural network (forming a probabilistic autoencoder), one obtains a scoring algorithm consistent with the interpretable Bayesian model. In some IRT applications the black-box nature of a neural network scoring machine is desirable. In this manuscript, we demonstrate within-IRT factorization and comment on scoring approaches.

## 1 Introduction

### 1.1 Item response theory (IRT)

IRT [7, 10] is the dominant statistical paradigm for using tests in order to quantify latent traits for individuals in the context of their populations. In both form and function, IRT models resemble autoencoders. Like in an autoencoder, the traits are a nonlinear factor representation of input data, and the latent representation is a major analytical product of interest. While additional latent variables exist both in IRT and in neural network-based autoencoders, of these, only in an IRT model are the latent variables interpretable. These parameters constitute item calibrations – difficulty and discrimination parameters – that correspond to the information content of each test question. In effect, IRT models are able to separate individual effects from test-specific effects in understanding responses. In contrast, the weights of a neural network lack mechanistic interpretability.

Model interpretability is a weakness for deep learning because artificial neural networks are intrinsically black boxes. It is difficult to understand how neural networks encode relationships between variables in their multidimensional mappings. In practical terms, this opaqueness has made it difficult to determine whether neural networks are making good objective inferences, or are simply codifying biases inherent in their training data.

Unfortunately, these issues are an impediment to using autoencoders in testing, where testing outcomes have real-life consequences. Disentangled autoencoders like the  $\beta$ -variational autoencoder [16] ( $\beta$ -VAE), and related methods [17, 3, 20, 6] show promise – they provide improved interpretability

by producing disentangled representations. What is meant by disentanglement, however, is specific to the method. Disentanglement typically refers to statistical independence of these latent factors, and not on dependencies to the input data itself.

Typically, for IRT, one desires disentanglement of original test items through partitioning, so that each partition informs a different dimension of the latent trait. The partitioning is performed prior to training of the IRT model, using exploratory factor analysis. Subdomains of the IRT model are then fit independently, conditional on the partition, using various empirical Bayesian [5] or Bayesian methods [31, 28].

Hence, IRT models are dependent on the initial partitioning of items into domains. The domain partitioning is a linear covariance analysis, where various criteria are commonly-used for accepting and rejecting items within each factor domain, without attention to consistency with the nonlinear IRT model. Additionally, partitional uncertainty is not propagated into the IRT model.

We propose factorization directly within IRT models using sparse coding. Additionally, we interpret IRT as a probabilistic decoder, and couple it to a neural network encoder. In training the encoder, the neural network learns the inversion of the generative model. Hence, one obtains a scoring machine that is consistent with the learned posterior of the IRT model. Testing is a rare application in which black box scoring machines may have value. The uninterpretable nature of the neural network can help impede gaming of exams if parameters from the scoring algorithm become partially compromised.

## 2 Modeling

In the parlance of neural networks, we present the overall method under the framework of a probabilistic autoencoder, depicted in Fig. 1. Item response theory (IRT) provides a substitute for the less-structured decoder neural network present in vanilla autoencoders.

### 2.1 Hierarchical Bayesian item response theory (IRT)




Figure 1: **Architecture:** The encoder (left) computes via an uninterpretable neural network an interpretable representation of the input data consistent with the mechanistic Bayesian hierarchical model (right).

In item response theory (IRT), a person’s test responses are modeled using both personal traits and item-specific parameters. The item parameters relate to the difficulty of the item and the discrimination of the item, or the degree to which the question’s responses are determined by personal traits. The two types of attributes are linked together via item response functions.

We focus on the parametric graded response model (GRM), which is used to model tests with ordinal responses. According to the GRM, the probability of a response of  $j \in \{1, 2, \dots, J\}$  to item  $i$  for person  $p$  is the likelihood function

$$\begin{aligned} \Pr(X_{pi} = j | \theta_p, \tau_{i,j}, \tau_{i,j+1}, \lambda_i) &= \Pr(X_{pi} \geq j | \theta_p, \tau_{i,j}, \lambda_i) - \Pr(X_{pi} \geq j + 1 | \theta_p, \tau_{i,j+1}, \lambda_i) \\ &= S(\lambda_i(\theta_p - \tau_{i,j})) - S(\lambda_i(\theta_p - \tau_{i,j+1})), \end{aligned} \quad (1)$$

where  $S$  is the sigmoid function

$$S(x) = \frac{1}{1 + e^{-x}}.$$

The item discrimination parameters  $\lambda_i$  control the scale of the response functions, or the strengths of the relationships between trait and response:  $\lambda_i = 0$  implies that a item is non-informative. The item-specific parameters  $\tau_{ij} < \tau_{i,j+1}$  and the personal traits  $\theta_p$  set the locations of the response distribution. For notational simplicity, we set  $\tau_{i,0} = -\infty$  and  $\tau_{i,J} = \infty$ .

Extending the GRM to multidimensional traits, we consider a mixture model where each component of the trait is informed by a weighted combination of the original items. Our objective in setting priors for the training problem is to simultaneously regularize an IRT model and factorize the original test items. To this end, we encourage sparsity in the mappings between the latent representation and the original items, by using hierarchical shrinkage on the item-domain discrimination parameters.

To encourage item-to-domain sparsity, we set horseshoe priors [24] on the discrimination parameters, yielding the hierarchical model equations

<b>Likelihood</b>	<b>Prior</b>
$\Pr(X_{pi} = j   \{\theta_p^{(d)}\}_d, \{\{\tau_{ij}^{(d)}\}_j\}_d, \{\lambda_i^{(d)}\}_d)$ $= \sum_{d=1}^D w_{id} \Pr(X_{pi} = j   \theta_p^{(d)}, \tau_{i,j}^{(d)}, \tau_{i,j+1}^{(d)}, \lambda_i^{(d)})$ $w_{id} = \left(\lambda_i^{(d)}\right)^\nu / \sum_{d=0}^{D-1} \left(\lambda_i^{(d)}\right)^\nu$ $\vec{\theta}_p   \mathcal{W} = \text{encoder}(\vec{X}_p; \mathcal{W})$	$\lambda_i^{(d)}   \eta_i, \xi_i^{(d)}, \kappa^{(d)} \sim \mathcal{N}^+(0, \eta_i \xi_i^{(d)} \kappa^{(d)})$ $\xi_i^{(d)} \sim C^+(0, \xi_0)$ $\eta_i \sim C^+(0, \eta_0)$ $\kappa^{(d)} \sim C^+(0, \kappa_0^{(d)})$ $\tau_{i,0}^{(d)} \sim \mathcal{N}(\mu_i^{(d)}, 1)$ $\mu_i^{(d)} \sim \mathcal{N}(0, 1)$ $\tau_{i,j}^{(d)}   \tau_{i,j-1}^{(d)} \sim \mathcal{N}^+(\tau_{i,j-1}^{(d)}, 1)$ $\theta_p, \mathcal{W} \sim \mathcal{N}(\theta_p; 0, 1) p(\mathcal{W}, \dots),$ <div style="text-align: right;">(2)</div>

defined for  $i \in 1, \dots, I$   $d \in 1, \dots, D$ ,  $p \in 1, \dots, P$ , and  $\mathcal{W}$  represents neural-network parameters.

The half-Cauchy hyperparameters determining the variance of the discrimination parameters constitute a variant of the horseshoe prior. In this context, the horseshoe prior factors items into domains by encouraging sparsity in the discrimination parameters. We also use horseshoe priors in regularizing the probabilistic encoder. When applied to weights of neural networks, the horseshoe prior has been shown to improve out of sample performance [14, 13], as it is also known to do in the general setting of Bayesian hierarchical models [25, 26].

## 2.2 Probabilistic encoder

Referring to Fig. 1, we use the graded response model as a probabilistic decoder. Note that the decoder model is itself sufficient for calibrating and implementing item response theory. However, by coupling a neural network encoder to the decoder, we obtain a test scoring algorithm for a new subject's responses

$$\text{encoder} : \text{new responses} \rightarrow \mathbb{E}_{\text{posterior parameters}} \Pr \left( \text{traits} \left| \begin{array}{l} \text{training} \\ \text{responses,} \end{array} \right. \begin{array}{l} \text{calibrated} \\ \text{model} \end{array} \right), \quad (3)$$

consistent with the Bayesian inference performed in training the IRT model. Note that Eq. 3 is from the viewpoint that the encoder represents probability densities. There are many ways to accomplish this representation. One way of doing so is to use weight-uncertain Bayesian neural networks [4], where the encoder is an ensemble of neural networks, in a manner true to Fig. 1. Other computationally-efficient approximate methods also exist such as dropout [30, 11], or simply using arbitrary neural network architectures to express either parametric [18], semiparametric, or non-parametric probability distributions [19, 21].

Looking at implementation details, there are some further similarities between the GRM and neural networks. The GRM (Eq. 1) has scaling degeneracies that are absolved by our hierarchical priors.

Similarly, neural networks are also scale-degenerate, in the absence of scale-enforcing priors or regularization. Additionally, there is a permutation degeneracy in the GRM, that compromises identifiability, that we ameliorate by setting  $\kappa_0^{(d+1)} < \kappa_0^{(d)}$ . There also exists the same degeneracy in DNNs: each dense layer of  $n$  nodes has  $n!$  computationally-identical configurations, suggesting a similar remedy.

### 2.3 Training

This model makes use of the horseshoe prior, for which the Cauchy distribution plays a starring role. The heavy peak and fat tails of Cauchy priors give it excellent shrinkage properties, however, they make it difficult to sample from directly. In lieu of direct sampling, we use the following auxiliary parameterization [22],

$$x \sim \mathcal{C}^+(0, \sigma) \iff \begin{aligned} x^2 &\sim \text{Inverse-Gamma}\left(\frac{1}{2}, \frac{1}{\lambda}\right) \\ \lambda &\sim \text{Inverse-Gamma}\left(\frac{1}{2}, \frac{1}{\sigma^2}\right). \end{aligned} \quad (4)$$

Directly training the conjoined model of Fig. 1 is challenging. Our strategy has been to proceed in blocks, by initially training the IRT model on its own, and then training the neural network separately, conditional on the IRT model. For both sets of training, we used mean-field ADVI with inverse-Gamma distributions parameterizing the scale parameters.

### 2.4 Model evaluation

To decide on a metric for model comparison, we consider the real-world use-case of IRT models. In practice, IRT models are first calibrated to a training set of responses from which both person- and item-specific model parameters are learned. Then, conditional on these pre-trained item-specific parameters, new respondents are scored by learning their person-specific parameters (traits). Hence, the scoring procedure finds the traits of new respondents relative to those of the calibration sample.

Implicitly, the assumption is that the calibrated model is predictive of how new respondents interact with the test. Hence, the out-of-sample predictive power of any calibrated IRT model is important. The general way to assess this power is to use held-out data, through k-fold cross validation, which requires multiple model fits. To save from having to pay the added computational cost of refitting the model, we adapt an approximation of Bayesian leave-one-out cross validation (LOO) known as the Widely Applicable Information Criterion (WAIC) [29, 12, 23, 27].

The WAIC is a simple-to-compute asymptotic approximation of LOO under minimal regularity conditions. In particular, the WAIC does not assume that the mapping between data and model parameters is bijective. The WAIC consists of the balance between two simple terms which often have exact algebraic expressions [8]. The first term is the log pointwise predictive density (lppd), which is the ‘‘pointwise’’ sum of the logarithm of the marginal likelihoods for the training data, or

$$\text{lppd} = \sum_p \log \left\{ \mathbb{E}_{\text{posterior}}^{\text{param}} \left[ \Pr(\mathbf{X}_p | \text{params}) \right] \right\} \approx \sum_p \log \left\{ \frac{1}{S} \sum_{s=0}^{S-1} \left[ \Pr(\mathbf{X}_p | \text{param}_{\text{sample } s}) \right] \right\}, \quad (5)$$

where we treat each person  $p$ ’s collection of responses  $\mathbf{X}_p$  as a statistically interchangeable point and assume that the joint posterior distribution of the calibration parameters is represented by a sample of size  $S$ . The second term in the WAIC is a penalty term of which there are two commonly used variants. The second variant has lower variance than the first [12], so we adapt it to our purposes in computing

$$\text{pwaic} = \sum_p \text{Var}_{\text{posterior}}^{\text{param}} \left[ \log \Pr(\mathbf{X}_p | \text{param}_{\text{sample } s}) \right]. \quad (6)$$

Eqs 5 and 6 combine together to constitute the WAIC,

$$\text{WAIC} = -2 \times (\text{lppd} - \text{pwaic}), \quad (7)$$

which is on the same deviance scale as used in the AIC.

The WAIC, being a model statistic, is a stochastic quantity that one estimates conditional on the on-hand data. Hence, it should not be treated as certain and one must also take its estimation error into account. To this end, one may approximate the standard error of the WAIC using the expression

$$\widehat{\text{se}}(\text{WAIC}) \approx \sqrt{P \times \text{Var}_{\text{people}} \left[ \log \left\{ \mathbb{E}_{\text{posterior}}^{\text{param}} \left[ \Pr(\mathbf{X}_p | \text{params}) \right] \right\} - \text{Var}_{\text{posterior}}^{\text{param}} \left[ \log \Pr(\mathbf{X}_p | \text{sample } s) \right] \right]}. \quad (8)$$

### 3 Experiments

As a testing ground for the overall method, we used the Right Wing Authoritarianism (RWA) scale [2], for which there exists an openly available dataset [1] representing 9881 college-aged respondents. The scale consists of twenty two questions on cultural and societal beliefs, with responses reported on a 9-level Likert scale.

Fig. 2 demonstrates factorization of the scale into two dimensions, learned using mean field ADVI, with inverse-gamma distributions to approximate all scale parameters. In these examples, we set the following scale parameters:  $\eta_0 = 0.01$ ,  $\xi_0 = 0.01$ ,  $\kappa_0^{(d)} = 0.01 \times (0.1)^d$ . For comparison, the linear factor loadings are also given.

Fig. 3 demonstrates factorization of the scale into three dimensions, also learned using mean field ADVI, with inverse-gamma distributions to approximate all scale parameters. Also in these examples, we set the following scale parameters:  $\eta_0 = 0.01$ ,  $\xi_0 = 0.01$ ,  $\kappa_0^{(d)} = 0.01 \times (0.1)^d$ .

In both Figs. 2 and 3, we initialized the ADVI algorithm so that the mean of the discrimination parameters is one plus the corresponding factor loading. We found this choice to lead to stable training using the Adam optimizer with an initial step size ranging between  $10^{-3}$  and  $10^{-4}$ , and with an exponentially decaying learning rate.

The WAIC statistics for the two and three-dimensional models were  $\text{WAIC}_{D=2} \approx 5.67 \pm 0.025 \times 10^5$  and  $\text{WAIC}_{D=3} \approx 5.66 \pm 0.025 \times 10^5$  respectively, mildly supporting the adaptation of a three-dimensional factorization on purely predictive grounds.

We performed all experiments using TensorFlow probability [9]. The source code is available upon request. The python package `factor-analyzer` provided the algorithm for our exploratory factor analysis results that we contrasted against the within-IRT factorizations.

## 4 Discussion

In this manuscript we have defined a Bayesian hierarchical model for performing factorization of test items in item response theory, making an analogy to probabilistic autoencoders. We demonstrated the method on the Right Wing Authoritarianism scale and used Bayesian predictive model comparison to contrast two and three-dimensional factorizations on the basis of prediction.

### 4.1 Within-IRT versus linear factorization

In practice, people use factor analysis as a first step in item response theory, in order to segregate items into separate domains for fitting IRT. Factor analysis results in a non-sparse matrix of loadings. To use factor analysis for this purpose, one must adapt an ultimately arbitrary cut-off criteria. A common criterion used in the literature is to set the cutoff at an absolute value of 0.4, though other advised cutoffs dependent on sample size exist [15].

For the RWA factorization, an absolute cut-off of  $|F| < 0.4$  in the factor analysis leads to a different factorization from that obtained within-IRT, though with some similarities. Looking at the two dimensional factorization of Fig. 2, the within-IRT factorization assigns fewer overall items into the first dimension relative the linear factor analysis. Looking at the item separation within-IRT, there are clear differences between the two overall factor dimensions. The first dimension tends to pick up on items that describe attitudes towards openness to liberal cultural attitudes. The second dimension tends to pick up on items describing adherence to traditional values and deference to authorities.

IRT weights		Question	Linear factors	
$w^{(0)}$	$w^{(1)}$		$F^{(0)}$	$F^{(1)}$
0.01	0.99	This country would work a lot better if certain groups of troublemakers would just shut up and accept their group's traditional place in society.	-0.14	0.72
1.0	0.0	Homosexuals and feminists should be praised for being brave enough to defy "traditional family values."	0.71	-0.1
1.0	0.0	There is no "one right way" to live life; everybody has to create their own way.	0.75	0.03
0.0	1.0	Our country will be great if we honor the ways of our forefathers, do what the authorities tell us to do, and get rid of the "rotten apples" who are ruining everything.	-0.04	0.83
0.9	0.1	A "woman's place" should be wherever she wants to be. The days when women are submissive to their husbands and social conventions belong strictly in the past.	0.75	0.01
0.01	0.99	There are many radical, immoral people in our country today, who are trying to ruin it for their own godless purposes, whom the authorities should put out of action.	-0.15	0.66
0.17	0.83	God's laws about abortion, pornography and marriage must be strictly followed before it is too late, and those who break them must be strongly punished.	-0.61	0.19
1.0	0.0	Some of the best people in our country are those who are challenging our government, criticizing religion, and ignoring the "normal way things are supposed to be done."	0.48	-0.3
0.01	0.99	What our country really needs is a strong, determined leader who will crush evil, and take us back to our true path.	0.06	0.84
0.94	0.06	You have to admire those who challenged the law and the majority's view by protesting for women's abortion rights, for animal rights, or to abolish school prayer.	0.75	-0.08
0.1	0.9	The "old-fashioned ways" and the "old-fashioned values" still show the best way to live.	-0.39	0.44
0.99	0.01	Everyone should have their own lifestyle, religious beliefs, and sexual preferences, even if it makes them different from everyone else.	0.81	0.06
0.07	0.93	Our country will be destroyed someday if we do not smash the perversions eating away at our moral fiber and traditional beliefs.	-0.3	0.57
0.87	0.13	Our country needs free thinkers who have the courage to defy traditional ways, even if this upsets many people.	0.48	-0.21
1.0	0.0	There is absolutely nothing wrong with nudist camps.	0.53	-0.15
0.01	0.99	The only way our country can get through the crisis ahead is to get back to our traditional values, put some tough leaders in power, and silence the troublemakers spreading bad ideas.	-0.06	0.82
0.88	0.12	Atheists and others who have rebelled against the established religions are no doubt every bit as good and virtuous as those who attend church regularly.	0.71	0.0
0.14	0.86	It is always better to trust the judgement of the proper authorities in government and religion than to listen to the noisy rabble-rousers in our society who are trying to create doubt in people's minds.	0.02	0.73
0.95	0.05	Gays and lesbians are just as healthy and moral as anybody else.	0.88	0.06
0.01	0.99	Our country desperately needs a mighty leader who will do what has to be done to destroy the radical new ways and sinfulness that are ruining us.	-0.1	0.77
0.15	0.85	Women should have to promise to obey their husbands when they get married.	-0.53	0.25
0.18	0.82	The established authorities generally turn out to be right about things, while the radicals and protestors are usually just "loud mouths" showing off their ignorance.	0.07	0.69

Figure 2: **Two-dimensional within-IRT factorization vs linear factorization of Right Wing Authoritarianism (RWA) scale test items**, colored by discrimination weights  $w_i^{(d)}$ . Shown are posterior expectations for the weights as estimated by mean-field ADVI. For comparison, linear factor loadings, colored by absolute value, are contrasted on the right hand side.

IRT weights			Question	Linear factors		
$w^{(0)}$	$w^{(1)}$	$w^{(2)}$		$F^{(0)}$	$F^{(1)}$	$F^{(2)}$
0.0	0.87	0.13	This country would work a lot better if certain groups of troublemakers would just shut up and accept their group's traditional place in society.	-0.12	0.59	0.18
1.0	0.0	0.0	Homosexuals and feminists should be praised for being brave enough to defy "traditional family values."	0.7	-0.13	0.01
1.0	0.0	0.0	There is no "one right way" to live life; everybody has to create their own way.	0.74	0.06	-0.04
0.0	0.88	0.12	Our country will be great if we honor the ways of our forefathers, do what the authorities tell us to do, and get rid of the "rotten apples" who are ruining everything.	-0.04	0.64	0.25
0.91	0.09	0.0	A "woman's place" should be wherever she wants to be. The days when women are submissive to their husbands and social conventions belong strictly in the past.	0.73	-0.14	0.15
0.0	0.86	0.14	There are many radical, immoral people in our country today, who are trying to ruin it for their own godless purposes, whom the authorities should put out of action.	-0.13	0.63	0.07
0.0	0.78	0.22	God's laws about abortion, pornography and marriage must be strictly followed before it is too late, and those who break them must be strongly punished.	-0.59	0.22	-0.02
1.0	0.0	0.0	Some of the best people in our country are those who are challenging our government, criticizing religion, and ignoring the "normal way things are supposed to be done."	0.5	0.09	-0.44
0.0	0.87	0.13	What our country really needs is a strong, determined leader who will crush evil, and take us back to our true path.	0.09	0.86	0.04
0.94	0.06	0.0	You have to admire those who challenged the law and the majority's view by protesting for women's abortion rights, for animal rights, or to abolish school prayer.	0.75	-0.01	-0.09
0.0	0.84	0.16	The "old-fashioned ways" and the "old-fashioned values" still show the best way to live.	-0.38	0.38	0.09
1.0	0.0	0.0	Everyone should have their own lifestyle, religious beliefs, and sexual preferences, even if it makes them different from everyone else.	0.8	0.04	-0.0
0.07	0.93	0.0	Our country will be destroyed someday if we do not smash the perversions eating away at our moral fiber and traditional beliefs.	-0.28	0.64	-0.04
0.99	0.0	0.0	Our country needs free thinkers who have the courage to defy traditional ways, even if this upsets many people.	0.49	0.07	-0.31
1.0	0.0	0.0	There is absolutely nothing wrong with nudist camps.	0.52	-0.04	-0.14
0.08	0.92	0.0	The only way our country can get through the crisis ahead is to get back to our traditional values, put some tough leaders in power, and silence the troublemakers spreading bad ideas.	-0.04	0.76	0.11
0.86	0.0	0.14	Atheists and others who have rebelled against the established religions are no doubt every bit as good and virtuous as those who attend church regularly.	0.71	0.03	-0.05
0.0	0.7	0.3	It is always better to trust the judgement of the proper authorities in government and religion than to listen to the noisy rabble-rousers in our society who are trying to create doubt in people's minds.	0.07	0.13	0.78
0.93	0.07	0.0	Gays and lesbians are just as healthy and moral as anybody else.	0.87	-0.12	0.19
0.09	0.91	0.0	Our country desperately needs a mighty leader who will do what has to be done to destroy the radical new ways and sinfulness that are ruining us.	-0.06	0.82	0.0
0.17	0.7	0.13	Women should have to promise to obey their husbands when they get married.	-0.51	0.37	-0.1
0.08	0.61	0.31	The established authorities generally turn out to be right about things, while the radicals and protestors are usually just "loud mouths" showing off their ignorance.	0.11	0.16	0.67

Figure 3: **Three-dimensional within-IRT factorization vs linear factorization of Right Wing Authoritarianism (RWA) scale test items**, colored by discrimination weights  $w_i^{(d)}$ . Shown are posterior expectations for the weights as estimated by mean-field ADVI. For comparison, linear factor loadings, colored by absolute value, are contrasted on the right hand side.

The two methods clearly disagreed in their categorization of two items: *Women should have to promise to obey their husbands when they get married*, and *God’s laws about abortion, pornography and marriage must be strictly followed before it is too late, and those who break them must be strongly punished*. Factor analysis placed both of these items into the first factor, though with negative loading. Within-IRT factorization placed these two items into the second factor instead.

Extending the factorizations to three dimensions (Fig. 3), one sees that for the within-IRT factorization, some of the items that weighted nearly completely into the second dimension leak partially into the third dimension, while retaining the bulk of their weight in the second dimension. In contrast, for linear factorization, two of the items switch completely from the second factor to the third factor. This behavior suggests that the IRT factorization is inherently more-stable than linear factor analysis, demonstrating the partial pooling feature of Bayesian hierarchical modeling.

## 4.2 Selecting between two and three dimensions

Based purely on predictive accuracy, the three-dimensional model achieves a slight edge over the two-dimensional one, though the difference between the WAIC statistics for these models was within one standard error.

For practical purposes, the two models have very similar predictive capabilities, as estimated using the WAIC. Hence, other considerations must be used for deciding between the models. These considerations are highly context-dependent, and should be influenced by how one wishes to interpret or use the testing instrument.

The interpretability of the two-dimensional model is fairly clear, given the content of the items within the resulting factors. The two- and three-dimensional models have similar factorizations, except for some cases where the second dimension leaks some weight into the third dimension. The meaningfulness of the third dimension requires subject-matter knowledge to adjudicate.

## 4.3 Scoring

When training a Bayesian model using mean-field variational methods, one is breaking dependencies between the variables of interest. For this reason, the sampling distribution of the surrogate posterior distribution does not necessarily yield the correct prediction distribution, unless either the posterior parameter densities are heavily concentrated around an unimodal peak, or are factorizable.

The entire purpose of IRT is to obtain statistics of the prediction distribution for latent traits conditional on input responses. Mindful of this purpose, it is improper to directly use the surrogate posterior parameter distributions to perform scoring – this holds true regardless of whether one wishes to score using the neural network, or by marginalizing over the statistics of the IRT model. In either case, one must use either more-structured variational approximations, or direct Monte-Carlo samplings from the joint posterior density, to preserve at least some of the inter-variable associations.

Using direct Monte-Carlo simulations of the parameters is conceptually straightforward, yet computationally intensive, both at the time of training and at the time of prediction. One can reduce the computation needed at training by performing MCMC only after ADVI, but then still be left with the computation of propagating the parameter samples at the time of prediction.

Since Bayesian computation is approximate by nature, it is perhaps fruitful to forgo Bayesian neural networks and seek deterministic neural networks that directly parameterize Gaussian or other low-order approximations of the prediction distributions. Doing so would also come with the added benefit of making scoring incredibly fast. Direct Bayesian marginalization of the IRT model requires computation of a multidimensional integral either directly or through Monte Carlo methods. For this reason, the neural network architecture would be advantageous. Future work will focus on finding efficient architectures for this task.

## 5 Conclusion

We have demonstrated that one may perform multidimensional within-IRT factorization and model calibration in a single step, using sparse coding. By acknowledging the analogy between IRT and probabilistic autoencoders, one also obtains a versatile tool for inverting IRT model products to



score new test respondents. On the flip side, regularization used for Bayesian multidimensional IRT shows promise for taming Bayesian neural networks. Future work will concentrate on translating the pooling properties of Bayesian hierarchical models to neural networks.

### **Acknowledgements**

This work is supported by the Intramural Research Programs of the National Institutes of Health Clinical Center (CC) and the National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK), and the US Social Security Administration.

## References

- [1] Take a personality test - Open Source Psychometrics Project. URL <https://openpsychometrics.org/>.
- [2] Bob Altemeyer. *Enemies of freedom: Understanding right-wing authoritarianism*. Enemies of freedom: Understanding right-wing authoritarianism. Jossey-Bass, San Francisco, CA, US, 1988. ISBN 978-1-55542-097-0.
- [3] Javier Antoran and Antonio Miguel. Disentangling in Variational Autoencoders with Natural Clustering. *arXiv:1901.09415 [cs, stat]*, January 2019. URL <http://arxiv.org/abs/1901.09415>. arXiv: 1901.09415.
- [4] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight Uncertainty in Neural Networks. May 2015. URL <https://arxiv.org/abs/1505.05424v2>.
- [5] R. Darrell Bock, David Thissen, and Michele F. Zimowski. IRT Estimation of Domain Scores. *Journal of Educational Measurement*, 34(3):197–211, 1997. ISSN 1745-3984. doi: 10.1111/j.1745-3984.1997.tb00515.x. URL <https://www.onlinelibrary.wiley.com/doi/abs/10.1111/j.1745-3984.1997.tb00515.x>.
- [6] Christopher P. Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in  $\beta$ -VAE. *arXiv:1804.03599 [cs, stat]*, April 2018. URL <http://arxiv.org/abs/1804.03599>. arXiv: 1804.03599.
- [7] Chih-Hung Chang and Bryce B. Reeve. Item Response Theory and its Applications to Patient-Reported Outcomes Measurement. *Evaluation & the Health Professions*, 28(3):264–282, September 2005. ISSN 0163-2787. doi: 10.1177/0163278705278275. URL <https://doi.org/10.1177/0163278705278275>.
- [8] Joshua C. Chang. Predictive Bayesian selection of multistep Markov chains, applied to the detection of the hot hand and other statistical dependencies in free throws. *Royal Society Open Science*, 6(3):182174. doi: 10.1098/rsos.182174. URL <https://royalsocietypublishing.org/doi/full/10.1098/rsos.182174>.
- [9] Joshua V. Dillon, Ian Langmore, Dustin Tran, Eugene Brevdo, Srinivas Vasudevan, Dave Moore, Brian Patton, Alex Alemi, Matt Hoffman, and Rif A. Saurous. TensorFlow Distributions. *arXiv:1711.10604 [cs, stat]*, November 2017. URL <http://arxiv.org/abs/1711.10604>. arXiv: 1711.10604.
- [10] James F. Fries, James Witter, Matthias Rose, David Cella, Dinesh Khanna, and Esi Morgan-DeWitt. Item Response Theory, Computerized Adaptive Testing, and PROMIS: Assessment of Physical Function. *The Journal of Rheumatology*, 41(1):153–158, January 2014. ISSN 0315-162X, 1499-2752. doi: 10.3899/jrheum.130813. URL <http://www.jrheum.org/content/41/1/153>.
- [11] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. June 2015. URL <https://arxiv.org/abs/1506.02142v6>.
- [12] Andrew Gelman, Jessica Hwang, and Aki Vehtari. Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 24(6):997–1016, November 2014. ISSN 1573-1375. doi: 10.1007/s11222-013-9416-2. URL <https://doi.org/10.1007/s11222-013-9416-2>.
- [13] Soumya Ghosh and Finale Doshi-Velez. Model Selection in Bayesian Neural Networks via Horseshoe Priors. May 2017. URL <https://arxiv.org/abs/1705.10388v1>.
- [14] Soumya Ghosh, Jiayu Yao, and Finale Doshi-Velez. Structured Variational Learning of Bayesian Neural Networks with Horseshoe Priors. *arXiv:1806.05975 [cs, stat]*, June 2018. URL <http://arxiv.org/abs/1806.05975>. arXiv: 1806.05975.
- [15] J. F. Hair, R. E. Anderson, R. L. Tatham, and W. C. Black. *Multivariate Data Analysis, 5th international ed.* Prentice Hall, Upper Saddle River, 1998.
- [16] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner.  $\beta$ -VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. November 2016. URL <https://openreview.net/forum?id=Sy2fzU9gl>.

- [17] Hyunjik Kim and Andriy Mnih. Disentangling by Factorising. In *International Conference on Machine Learning*, pages 2649–2658, July 2018. URL <http://proceedings.mlr.press/v80/kim18b.html>.
- [18] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv:1312.6114 [cs, stat]*, December 2013. URL <http://arxiv.org/abs/1312.6114>. arXiv: 1312.6114.
- [19] G.V. Kobayashi and A.V. Zamyatin. Conditional probability density estimation using artificial neural network. In *2015 9th International Conference on Application of Information and Communication Technologies (AICT)*, pages 441–445, October 2015. doi: 10.1109/ICAICT.2015.7338597. ISSN: null.
- [20] Yingzhen Li and Stephan Mandt. Disentangled Sequential Autoencoder. *arXiv:1803.02991 [cs]*, March 2018. URL <http://arxiv.org/abs/1803.02991>. arXiv: 1803.02991.
- [21] Aristidis Likas. Probability density estimation using artificial neural networks. *Computer Physics Communications*, 135(2):167–175, April 2001. ISSN 0010-4655. doi: 10.1016/S0010-4655(00)00235-6. URL <http://www.sciencedirect.com/science/article/pii/S0010465500002356>.
- [22] Enes Makalic and Daniel F. Schmidt. A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters*, 23(1):179–182, January 2016. ISSN 1070-9908, 1558-2361. doi: 10.1109/LSP.2015.2503725. URL <http://arxiv.org/abs/1508.03884>. arXiv: 1508.03884.
- [23] Juho Piironen and Aki Vehtari. Comparison of Bayesian predictive methods for model selection. *Statistics and Computing*, 27(3):711–735, May 2017. ISSN 1573-1375. doi: 10.1007/s11222-016-9649-y. URL <https://doi.org/10.1007/s11222-016-9649-y>.
- [24] Juho Piironen and Aki Vehtari. Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics*, 11(2):5018–5051, 2017. ISSN 1935-7524. doi: 10.1214/17-EJS1337SI. URL <http://arxiv.org/abs/1707.01694>. arXiv: 1707.01694.
- [25] Nicholas G. Polson and James G. Scott. On the half-Cauchy prior for a global scale parameter. *arXiv:1104.4937 [stat]*, April 2011. URL <http://arxiv.org/abs/1104.4937>. arXiv: 1104.4937.
- [26] Nicholas G. Polson and James G. Scott. *Shrink Globally, Act Locally: Sparse Bayesian Regularization and Prediction \**. Oxford University Press, October 2011. ISBN 978-0-19-173192-1. URL <https://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199694587.001.0001/acprof-9780199694587-chapter-17>.
- [27] Aki Vehtari, Andrew Gelman, and Jonah Gabry. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5):1413–1432, September 2017. ISSN 1573-1375. doi: 10.1007/s11222-016-9696-4. URL <https://doi.org/10.1007/s11222-016-9696-4>.
- [28] Walter P. Vispoel, Carrie A. Morris, and Murat Kilinc. Applications of generalizability theory and their relations to classical test theory and structural equation modeling. *Psychological Methods*, 23(1):1–26, 2018. ISSN 1939-1463(Electronic),1082-989X(Print). doi: 10.1037/met0000107.
- [29] Sumio Watanabe. Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory. *Journal of Machine Learning Research*, 11(Dec):3571–3594, 2010. ISSN 1533-7928. URL <http://www.jmlr.org/papers/v11/watanabe10a.html>.
- [30] Yeming Wen, Paul Vicol, Jimmy Ba, Dustin Tran, and Roger Grosse. Flipout: Efficient Pseudo-Independent Weight Perturbations on Mini-Batches. *arXiv:1803.04386 [cs, stat]*, March 2018. URL <http://arxiv.org/abs/1803.04386>. arXiv: 1803.04386.
- [31] Ruby Chiu-Hsing Weng and D. Stephen Coad. Real-Time Bayesian Parameter Estimation for Item Response Models. *Bayesian Analysis*, 13(1):115–137, March 2018. ISSN 1936-0975, 1931-6690. doi: 10.1214/16-BA1043. URL <https://projecteuclid.org/euclid.ba/1482138050>.