

# Raven’s Progressive Matrices Completion with Latent Gaussian Process Priors

Fan Shi, Bin Li\*, Xiangyang Xue

Shanghai Key Laboratory of Intelligent Information Processing  
School of Computer Science, Fudan University  
{fshi19, libin, xyxue}@fudan.edu.cn

## Abstract

Abstract reasoning ability is fundamental to human intelligence. It enables humans to uncover relations among abstract concepts and further deduce implicit rules from the relations. As a well-known abstract visual reasoning task, Raven’s Progressive Matrices (RPM) are widely used in human IQ tests. Although extensive research has been conducted on RPM solvers with machine intelligence, few studies have considered further advancing the standard answer-selection (classification) problem to a more challenging answer-painting (generating) problem, which can verify whether the model has indeed understood the implicit rules. In this paper we aim to solve the latter one by proposing a deep latent variable model, in which multiple Gaussian processes are employed as priors of latent variables to separately learn underlying abstract concepts from RPMs; thus the proposed model is interpretable in terms of concept-specific latent variables. The latent Gaussian process also provides an effective way of extrapolation for answer painting based on the learned concept-changing rules. We evaluate the proposed model on RPM-like datasets with multiple continuously-changing visual concepts. Experimental results demonstrate that our model requires only few training samples to paint high-quality answers, generate novel RPM panels, and achieve interpretability through concept-specific latent variables.

## Introduction

A major performance of human abstract thinking is the ability to discover relations between abstract concepts and then summarize underlying rules. Abstract visual reasoning tasks reflect such abstract thinking capabilities. The Raven’s Progressive Matrices (RPM) problem, a kind of nonverbal abstract visual reasoning task proposed by John C. Raven (Raven and Court 1938), are usually used to test human fluid intelligence (Bilker et al. 2012). As an example shown in Figure 1, participants of the RPM problem should observe a  $3 \times 3$  context panel with the right-bottom cell invisible. They are expected to find out underlying concept-changing rules and choose an answer (Raven and Court 1938) from a selection panel. The RPM problem predominantly emphasizes two aspects of intelligence: visual representation and

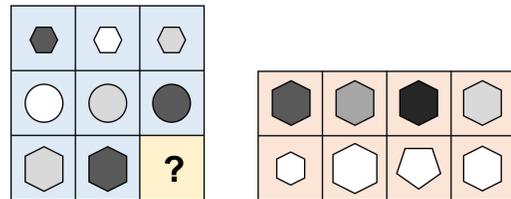


Figure 1: An example of the RPM panel from the RAVEN dataset (Zhang et al. 2019a). The context panel (left) consists of eight observed cells in blue background and one missing cell in yellow background; the selection panel (right) consists of eight cells in orange background.

abstract reasoning. From the perspective of visual representation, our vision and nervous system tend to analyze independent visual concepts from an observed object: such as *color*, *shape*, and *size*. From the perspective of abstract reasoning, when conducting RPM IQ tests, our strong abstract reasoning ability is based on the comprehension of independent visual concepts and also, concept-changing rules.

In the field of machine learning, enabling computational models with the abstract reasoning ability is challenging. Some similarity-based models (Barrett et al. 2018) measure similarity between every two cell content to carry out the reasoning process. Recently, as researchers show more interest in deep learning, a considerable amount of literature has been published on methods utilizing deep neural networks (Barrett et al. 2018; Zhang et al. 2019b,a). For a particular downstream task, the choice of feature representation has a great impact on performance. If feature representations are well adapted to the current task, the accuracy and robustness of models can be improved (Bengio, Courville, and Vincent 2013). The variational autoencoder (VAE) (Kingma and Welling 2013) has emerged as a functional tool in the unsupervised representation learning. To enhance the interpretability in representations, the notion of disentanglement is put forward. By decomposing and regularizing the evidence lower bound (ELBO) of VAEs, a number of disentangled latent variable models have been proposed (Higgins et al. 2017; Kumar, Sattigeri, and Balakrishnan 2017; Kim and Mnih 2018; Chen et al. 2018). These regularization-based models force each dimension (or

\*Corresponding author

multiple non-overlapping dimensions) of latent variables to code a single visual concept, which endows these dimensions clear meanings. Some recent studies (Steenbrugge et al. 2018; van Steenkiste et al. 2019) point out that replacing deep representations in the Wild Relational Network (WReN) (Barrett et al. 2018) with disentangled representations will improve the selection accuracy and data efficiency of the model when solving the RPM problem.

The related studies indicate substantial progress in the RPM problem. However, these methods still have downsides: (1) Similarity-based RPM solvers fail to capture concept-changing rules and lack the explicit generative process of entire panels; (2) Methods that introduce representations from the regularization-based disentangled VAE struggle to encode reliable concept-specific latent variables (Locatello et al. 2019). Both drawbacks spoil the interpretability of models. In fact, human cognition and abstract reasoning abilities are not limited in the original selective RPM problem. While tested with easily recognized rules, a human can paint the correct answer after inspecting context panels. Most RPM solvers nowadays focus on the regular cell selection problem. However, the entire panel completion is a more challenging task that further evaluates the abstract reasoning ability to find concept-changing rules.

In this paper, we propose an interpretable deep latent variable model that explicitly defines the generative and inference process of entire panels as well as regards the prior of concept-changing rules as 2D Gaussian processes (GPs) with the learnable kernel (Wilson et al. 2016) to solve panel completing tasks. Structure information within context panels is introduced to guide the learning of concept-specific representations. Due to the continuity of GPs and the concept-oriented characteristic of representations, it is possible to extrapolate answers from independent visual concepts of the observed cell contents. Deduced answer representations and context cell representations are finally decoded to construct complete panel images.

We evaluate the model in RPM-like datasets with continuously changing visual concepts. To evaluate both the reconstruction quality and comprehension of underlying concept-changing rules, we calculate the averaged MSE over panel cells by leaving out and predicting one cell in turn. For all datasets, the proposed model outperforms the VAE-GAN inpainting model (Yu et al. 2018). In addition to the missing cell prediction, our model can also generate novel panels according to deduced concept-changing rules, which has not been investigated in previous works. Further experiments about the multi-cell prediction and visual concept visualization indicate that our model realizes interpretability by means of the concept-specific latent variables and human-like rule reasoning process.

## Related Work

### RPM Problem Solver

To solve the RPM problem, some computational models are proposed (Lovett and Forbus 2017; Little, Lewandowsky, and Griffiths 2012; Lovett, Forbus, and Usher 2010) to extract symbolic representations from the context panel and

perform inference based on the particular rules. With the extensive application of deep neural networks, more and more models use convolutional neural networks as image feature extractors. A novel method called WReN (Barrett et al. 2018), based on the Relation Network (Santoro et al. 2017) and similarity scores, achieves competitive performance in the RPM problem. To evaluate the WReN, an RPM-like dataset named Procedurally Generated Matrices (PGM) is built. CoPINet (Zhang et al. 2019b) introduces the contrast effect from humans and animals to enhance the reasoning ability. In (Zhang et al. 2019a), authors release a more complex dataset RAVEN and use a tree-structured neural module called Dynamic Residual Tree (DRT) to represent the intrinsic structure of images. However, there exist obvious patterns in RAVEN’s selection panels, which leads to the high classification accuracy even without observing context panels (Hu et al. 2020). To overcome this defect, an unbiased Balanced-RAVEN dataset is proposed (Hu et al. 2020). Then the V-PROM dataset (Teney et al. 2020) extends the image style inside cells to more complex visual scenes. On the other hand, some studies on the visual representation (Steenbrugge et al. 2018; van Steenkiste et al. 2019) reveal that disentangled representations can significantly improve model performance in the RPM problem.

### Disentangled VAE

The disentangled VAE provides an idea for the concept-specific representation learning.

Regularization-based methods impose regularization constraints to the posterior distribution or aggregate posterior distribution in ELBO.  $\beta$ -VAE (Higgins et al. 2017) reformulates maximization of the objective function to an optimization problem with inequality constraints. While reconstructing the real data, these constraints force the posterior distribution closed to an isotropic standard Gaussian prior. DIP-VAE (Kumar, Sattigeri, and Balakrishnan 2017) adopts an extra regularization term to reduce the distance between the aggregate posterior and prior by moment matching. FactorVAE (Kim and Mnih 2018) and  $\beta$ -TCVAE (Chen et al. 2018) deal with the mutual information over-penalization problem in  $\beta$ -VAE. Both methods further decompose the ELBO and exactly exert a regularization penalty in the total correlation to ensure the independence between dimensions.  $\beta$ -TCVAE estimates the total correlation by minibatch-weighted sampling, while FactorVAE uses a density ratio trick that requires an additional discriminator.

In (Locatello et al. 2019), the authors point out the importance of explicit inductive bias and supervision information in the disentangled representation learning. By contrast, some methods introduce inductive bias, such as distinctive structures in data and network architectures, into models. VITAE (Skafte and Hauberg 2019) combines a symmetric transformation and an inverse transformation operations in the computation process to disentangle the appearance and perspective of objects. CascadeVAE (Jeong and Song 2019) and JointVAE (Dupont 2018) assume that the latent space consists of discrete and continuous parts and disentangle both discrete and continuous representations. Spatial Broadcast Decoder (Watters et al. 2019) broadcasts and concate-

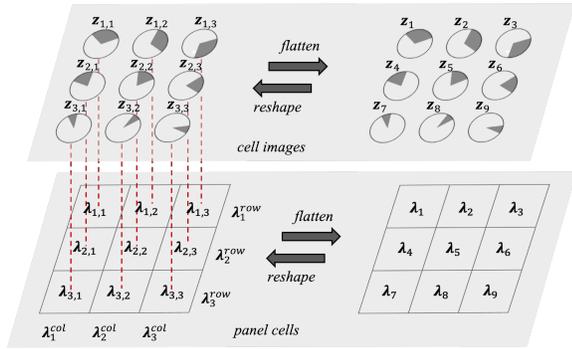


Figure 2: An illustration of the RPM panel and its notations.

nates latent variables to improve the disentanglement ability.

### GP with Deep Kernels

A GP is a set of random variables whose finite subsets satisfy the multivariate Gaussian distribution (Williams and Rasmussen 2006), which can be expressed as  $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), \mathbf{K}_{XX})$  where  $m(\mathbf{x})$  is the mean function usually set to zero function and  $\mathbf{K}_{XX}$  denotes the covariance matrix calculated from predefined kernel functions.

To improve the flexibility and avoid the empirical selection of kernels, some studies (Wilson et al. 2016; Sun et al. 2018) incorporate the GP kernel with deep neural networks. In the function-space view, a GP defines the particular distribution over functions whose basic characters (e.g., smoothness) are determined via kernels. Therefore, constructing or choosing kernels is an essential work in the GP prediction. In (Wilson et al. 2016), the model encodes original locations with a deep structure consisting of nonlinear mappings whose output is used in the spectral mixture base kernel. Neural Kernel Network (Sun et al. 2018) exploits composition rules to form more complicated kernels by considering each unit of a neural network as one basic kernel. Most works attempt to use the GP with deep kernels to create function spaces mapping data to predictive values. However, in this paper, we regard deep kernel GPs as the prior of concept-changing rules whose input and output reside in latent variable spaces.

### The Proposed Method

Inspired by the discussion in previous works about disentangled representations in the original RPM problem (van Steenkiste et al. 2019), we propose a deep latent variable model to solve the RPM painting problem. The model firstly encodes context cell images (objects inside the panel cells) to independent concept-specific latent variables. The latent variables of missing cell images are then derived via the GP rule reasoning where GPs are taken as the priors of underlying concept-changing rules. Finally, all latent variables are concatenated and sent to a decoder to reconstruct the original context panels as well as missing cell images. In this section, we will introduce the GP rule reasoning, generative model, inference model, and parameter learning method in detail.

### GP Rule Reasoning

In Figure 2, a complete RPM panel contains 3 rows and columns. We use  $\lambda^{row} = \{\lambda_1^{row}, \lambda_2^{row}, \lambda_3^{row}\}$  and  $\lambda^{col} = \{\lambda_1^{col}, \lambda_2^{col}, \lambda_3^{col}\}$  to denote row and column latent variables. By combining  $\lambda^{row}$  and  $\lambda^{col}$ , we can represent the panel cell in the  $i$ th row and  $j$ th column by a tuple  $(\lambda_i^{row}, \lambda_j^{col})$ . For convenience of notations, we define that  $\lambda_{i,j} = (\lambda_i^{row}, \lambda_j^{col})$ . Then we define panel's underlying concept-changing rules as 2D functions  $f^d: R \times R \rightarrow R$  that map  $\lambda_{i,j}$  to cell image latent variables  $z_{i,j}$  and we have  $f^d(\lambda_{i,j}) = z_{i,j}^d$ , where  $i, j \in \{1, 2, 3\}$  and  $d$  is the index of  $z_{i,j}$ 's dimension varying from 1 to  $D$ . Considering that each dimension corresponds to a single visual concept, the proposed method tends to learn different concept-changing rules for every dimension. To introduce rule functions into the Bayesian learning framework, we intuitively choose GPs as the priors. Another index  $k = 1, \dots, 9$  is used to denote the flattened 2D latent variables  $\lambda$  and  $z$ . We have  $k = 3(i-1) + j$  and  $i, j \in \{1, 2, 3\}$  (e.g.,  $z_6 = z_{2,3}$  in Figure 2). The  $d$ th dimension of  $z$  satisfies a multivariate Gaussian distribution.

$$\begin{aligned} \lambda_{1:9} &= \text{flatten}(\lambda) \\ \mathbf{K}_\theta^d &= \begin{bmatrix} \kappa_\theta^d(\lambda_1, \lambda_1) & \cdots & \kappa_\theta^d(\lambda_1, \lambda_9) \\ \vdots & \ddots & \vdots \\ \kappa_\theta^d(\lambda_9, \lambda_1) & \cdots & \kappa_\theta^d(\lambda_9, \lambda_9) \end{bmatrix} \\ z_{1:9}^d &\sim \mathcal{N}(\mathbf{0}, \mathbf{K}_\theta^d) \\ z^d &= \text{reshape}(z_{1:9}^d) \end{aligned} \quad (1)$$

where  $\kappa_\theta^d(\cdot, \cdot)$  is the kernel function utilizing a neural network  $g_\theta^d$  and RBF kernel parameterized with  $l$  and  $\sigma$  as in (Wilson et al. 2016).

$$\kappa_\theta^d(\lambda_m, \lambda_n) = l^2 \exp\left(-\frac{\|g_\theta^d(\lambda_m) - g_\theta^d(\lambda_n)\|_2^2}{2\sigma^2}\right) \quad (2)$$

It is the learnable parameters in  $\kappa_\theta^d(\cdot, \cdot)$  that realize the automatic function space adaption for real data. *flatten* denotes a function that flattens a matrix to a vector and *reshape* denotes an inverse procedure that deforms a vector into a matrix. The complete process described in Eq.(1) is referred to as  $z^d = \text{panelSampling}(\lambda, d)$ . By means of the GP priors, we can perform the following operations readily.

**Complete Panels Sampling.** As described above, we can calculate kernels  $\mathbf{K}_\theta^d$  via  $\lambda$  and then separately sample  $z^d$  from dimension-wise multivariate Gaussian distributions defined in Eq.(1). In the remaining generative process,  $z^d$  are concatenated through dimensions and then decoded to form a complete panel.

**Conditional Cells Prediction.** One significant facility of GPs is the ability to predict function values at novel locations. In the same way, we can infer the dimension-wise answer latent variable  $z_{3,3}^d$  based on the latent variables of the observed cell images  $z_{-3,3}^d$ , where the notation  $-3, 3$  means leaving the element indexed 3, 3 out. This process gives the proposed method ability of abstract reasoning, which is utilized in the inference phase.

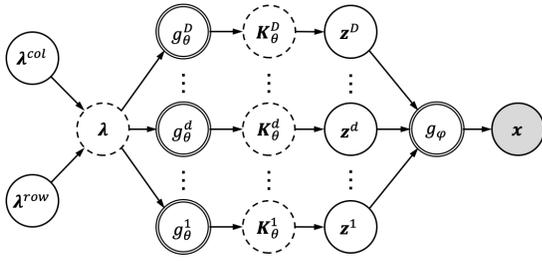


Figure 3: The diagram of the generative model. Solid circles are random variables including both latent and observed variables. Double solid circles mean neural networks and dashed circles denote determinate variables.

## Generative Model

According to Figure 3, the process to generate the cell images  $\mathbf{x}$  of the entire panel from  $\mathbf{z}$ ,  $\lambda^{row}$ , and  $\lambda^{col}$  is

$$\begin{aligned}
 \lambda^{row} &\sim \mathcal{N}(\mathbf{0}, \sigma_\lambda^2 \mathbf{I}) \\
 \lambda^{col} &\sim \mathcal{N}(\mathbf{0}, \sigma_\lambda^2 \mathbf{I}) \\
 \mathbf{z}^d &= \text{panelSampling}(\lambda, d), \quad d = 1, 2, \dots, D \quad (3) \\
 \mathbf{z} &= \text{concat}(\mathbf{z}^1, \dots, \mathbf{z}^D) \\
 \mathbf{x}_{i,j} &\sim \mathcal{N}(g_\varphi(\mathbf{z}_{i,j}), \sigma_x^2 \mathbf{I}), \quad i, j \in \{1, 2, 3\}
 \end{aligned}$$

$\lambda^{row}$  and  $\lambda^{col}$  are respectively panel's rows and columns that form the 2D locations  $\lambda$  of rule functions. Their priors are isotropic multivariate Gaussian distributions parameterized with  $\sigma_\lambda$ . In function *panelSampling* (defined in the previous section GP Rule Reasoning), dimension-wise covariance matrices  $\mathbf{K}_\theta^d$  are calculated based on the learnable kernel functions. Then rule functions can be sampled from the GPs. Since no Gaussian noise is introduced to observations,  $\mathbf{z}^d$  is equal to the function value sampled from the multivariate Gaussian distribution. Once  $\mathbf{z}^d$  is stacked dimension by dimension, a neural network  $g_\varphi$  will decode  $\mathbf{z}$  to means of the Gaussian distribution over  $\mathbf{x}$ . Variances of the Gaussian distribution over  $\mathbf{x}$  are artificially determined by a constant  $\sigma_x$ . The joint probability of latent variables  $\eta = \{\mathbf{z}, \lambda^{row}, \lambda^{col}\}$ , and  $\mathbf{x}$  is

$$\begin{aligned}
 p_{\varphi, \theta}(\mathbf{x}, \eta) &= \prod_{i=1}^3 \prod_{j=1}^3 p_\varphi(\mathbf{x}_{i,j} | \mathbf{z}_{i,j}) p(\lambda^{row}) \\
 & p(\lambda^{col}) \prod_{d=1}^D p_\theta(\mathbf{z}^d | \lambda^{row}, \lambda^{col}) \quad (4)
 \end{aligned}$$

## Variational Inference

In the framework of VAE, since the exact posterior distribution is hard to derive, a variational distribution is used to approximate the posterior distribution. Following the diagram in Figure 4, the variational posterior distribution can be ex-

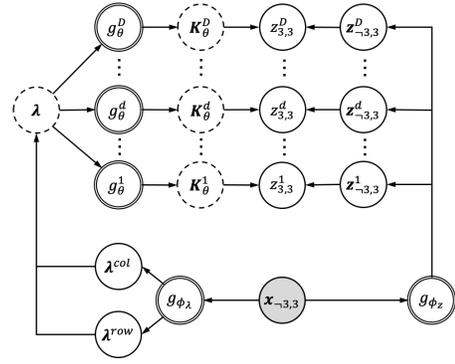


Figure 4: The diagram of the inference model. The meaning of distinct circles are the same as those in Figure 3.

pressed as

$$\begin{aligned}
 q_{\phi, \theta}(\eta | \mathbf{x}_{-3,3}) &= \prod_{d=1}^D q_\theta(\mathbf{z}_{3,3}^d | \mathbf{z}_{-3,3}^d, \lambda^{row}, \lambda^{col}) \\
 q_{\phi_\lambda}(\lambda^{row} | \mathbf{x}_{-3,3}) q_{\phi_\lambda}(\lambda^{col} | \mathbf{x}_{-3,3}) & \quad (5) \\
 \prod_{(i,j) \in \Omega} q_{\phi_z}(\mathbf{z}_{i,j} | \mathbf{x}_{i,j})
 \end{aligned}$$

where  $\phi = \{\phi_\lambda, \phi_z\}$  denote the neural network  $g_{\phi_\lambda}$  and  $g_{\phi_z}$ , respectively;  $\Omega = \{1, 2, 3\}^2 \setminus (3, 3)$  is the set for the index  $(i, j)$ . Variational distributions  $q_{\phi_\lambda}(\lambda^{row} | \mathbf{x}_{-3,3})$  and  $q_{\phi_\lambda}(\lambda^{col} | \mathbf{x}_{-3,3})$  are both Gaussian distributions denoting rows and columns of the context panel. Means and standard deviations of the distribution are outputs of the neural network  $g_{\phi_\lambda}$ . The variational distribution of  $\mathbf{z}_{i,j}$  is also the Gaussian distribution and parameters of  $q_{\phi_z}(\mathbf{z}_{i,j} | \mathbf{x}_{i,j})$  are derived from the neural network  $g_{\phi_z}$ . For inference, only cell images  $\mathbf{x}_{-3,3}$  in the context panel are available. The most essential part of the variational distribution is  $q_\theta(\mathbf{z}_{3,3}^d | \mathbf{z}_{-3,3}^d, \lambda^{row}, \lambda^{col})$  which follows the same calculation manner as the noise-free prediction with GPs. Denoting the Gaussian distribution of  $\mathbf{z}_{3,3}^d$  as  $\mathcal{N}(\mu_\theta, \sigma_\theta^2)$ , where

$$\begin{aligned}
 \mu_\theta &= \mathbf{R}_\theta^\top \mathbf{C}_\theta^{-1} \mathbf{t}, \quad \sigma_\theta^2 = d_\theta - \mathbf{R}_\theta^\top \mathbf{C}_\theta^{-1} \mathbf{R}_\theta, \\
 \lambda_{1:9} &= \text{flatten}(\lambda), \quad \mathbf{z}_{1:9}^d = \text{flatten}(\mathbf{z}^d) \\
 \mathbf{C}_\theta &= \begin{bmatrix} \kappa_\theta^d(\lambda_1, \lambda_1) & \cdots & \kappa_\theta^d(\lambda_1, \lambda_8) \\ \vdots & \ddots & \vdots \\ \kappa_\theta^d(\lambda_8, \lambda_1) & \cdots & \kappa_\theta^d(\lambda_8, \lambda_8) \end{bmatrix} \quad (6) \\
 \mathbf{R}_\theta &= [\kappa_\theta^d(\lambda_1, \lambda_9), \dots, \kappa_\theta^d(\lambda_8, \lambda_9)]^\top \\
 \mathbf{t} &= [z_1^d, \dots, z_8^d]^\top, \quad d_\theta = \kappa_\theta^d(\lambda_9, \lambda_9)
 \end{aligned}$$

To simplify the above equation, we omit the dimension symbol  $d$  in matrices. Eq.(6) indicates that when inferring  $\mathbf{z}_{3,3}^d$ , we need first to obtain  $\mathbf{z}_{-3,3}^d$ ,  $\lambda^{row}$  and  $\lambda^{col}$  from cell images in the context panel. Considering the  $\lambda^{row}$  and  $\lambda^{col}$  as locations and  $\mathbf{z}_{-3,3}^d$  as function values in the rule function, the distribution over  $\mathbf{z}_{3,3}^d$  can be eventually acquired.

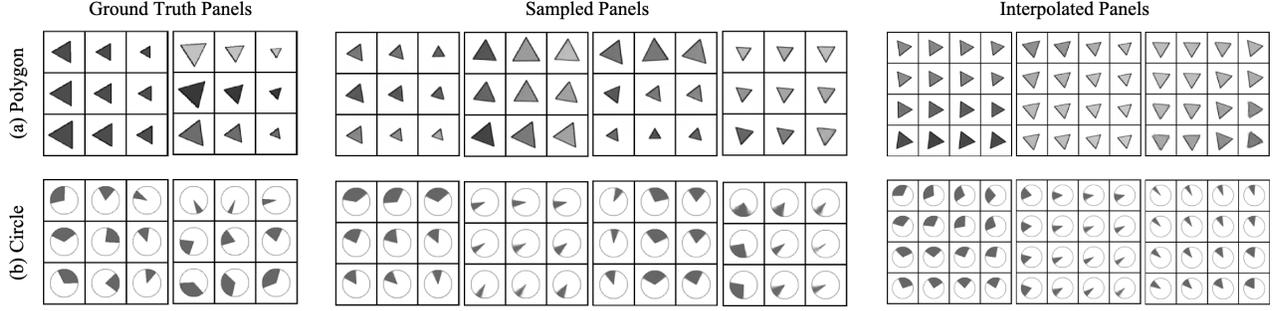


Figure 5: Panels sampled from datasets and the generative model. The left side are ground truth panels from datasets. Panels sampled from the generative model is displayed in the middle. The right side contains panels sampled via interpolation.

## Parameter Learning

In this paper, we optimize the ELBO  $\mathcal{L}$  instead of the log-likelihood  $\log p_{\varphi, \theta}(\mathbf{x})$  by amortized variational inference.

$$\begin{aligned}
\mathcal{L} &= \mathbb{E}_{q_{\varphi, \theta}(\boldsymbol{\eta} | \mathbf{x}_{-3,3})} \left[ \log \frac{p_{\varphi, \theta}(\mathbf{x}, \boldsymbol{\eta})}{q_{\varphi, \theta}(\boldsymbol{\eta} | \mathbf{x}_{-3,3})} \right] \\
&= \mathbb{E}_{q_{\varphi, \theta}(\mathbf{z}, \boldsymbol{\lambda} | \mathbf{x}_{-3,3})} \left[ \log p_{\varphi}(\mathbf{x} | \mathbf{z}) \right] \\
&\quad - \beta D_{KL}(q_{\phi}(\boldsymbol{\lambda} | \mathbf{x}_{-3,3}) \| p(\boldsymbol{\lambda})) \\
&\quad - \gamma \mathbb{E}_{q_{\phi}(\boldsymbol{\lambda} | \mathbf{x}_{-3,3})} \left[ \sum_{d=1}^D D_{KL}(q_{\theta}(z_{3,3}^d | z_{-3,3}^d, \boldsymbol{\lambda}) \right. \\
&\quad \left. q_{\phi}(z_{-3,3}^d | \mathbf{x}_{-3,3}) \| p_{\theta}(z^d | \boldsymbol{\lambda})) \right]
\end{aligned} \tag{7}$$

where  $\beta$  and  $\gamma$  are hyperparameters employed to control the penalty in KL divergence to impose the dimensional independence on latent variables (Higgins et al. 2017). To estimate the expectation of the first and third terms in Eq.(7), we sample  $\boldsymbol{\lambda}$  and  $\mathbf{z}$  from the variational posterior with reparameterization trick (Kingma and Welling 2013). Then the gradient-based backpropagation algorithm is executed to adjust the parameters of involved neural networks. However, when calculating KL divergence in the third term, the explicit form of distribution  $q_{\theta}(z_{3,3}^d | z_{-3,3}^d, \boldsymbol{\lambda}) q_{\phi}(z_{-3,3}^d | \mathbf{x}_{-3,3})$  is hard to derive. Alternatively, as the marginal distributions of  $p_{\theta}(z^d | \boldsymbol{\lambda})$  are computationally tractable, we propose a marginal matching method to approximate the non-computable KL divergence. To match marginal distributions of the posterior and prior, we substitute the prior with the product of its marginal distributions  $p_{\theta}(z_{3,3}^d | z_{-3,3}^d, \boldsymbol{\lambda}) p_{\theta}(z_{-3,3}^d | \boldsymbol{\lambda})$ . By means of this approximation method, the third term can be transformed into a summary over the KL divergence between marginal distributions

$$\sum_{(i,j) \in \Omega} D_{KL}(q_{\phi}(z_{i,j}^d | \mathbf{x}_{i,j}) \| p_{\theta}(z_{i,j}^d | \boldsymbol{\lambda})) \tag{8}$$

By maximizing the approximated ELBO, our model can be trained end to end. The detailed derivation process is described in the supplementary material.

## Experiments

### RPM-like Datasets

We create new RPM-like datasets that remove selection panels and retain complete 9-cell panels. Visual concepts of the panel object in a dataset are divided into a changeable and unchangeable concept set. Concepts in the changeable concept set could change with rules but those in the unchangeable concept cannot. All concepts change in continuous spaces, and rules can be exerted in both the horizontal and vertical directions. Datasets are generated according to the following rules:

- **progress**: one concept changes with isometric steps horizontally or vertically (2nd ground truth in Figure 5(a));
- **biprogess**: one concept changes with isometric steps in *both* directions (1st ground truth in Figure 5(a));
- **multiprogress**: *multiple concepts* change with isometric steps horizontally or vertically (1st and 2nd ground truths in Figure 5(b)).

To create a panel, we first randomly choose a rule and a subset of the changeable concept set. As Figure 5 shows, the values of the chosen concept across rows are stochastically determined under constraints from the selected rule. The remaining concepts are kept fixed in a row (we consider only horizontal rules now). Finally, the panel will be transposed with the 50% probability to produce vertical rules.

In the following experiments, our model will be evaluated in two instances of the RPM-like datasets:

- **Polygon**: each cell has a centered polygon (one dataset contains only one type of polygon, and we choose the triangle here) with the changeable concepts *size* and *grayscale* as well as the unchangeable *rotation*. Panels can follow the rules *progress* and *biprogess*.
- **Circle**: there is a centered hollow circle in each panel cell with a black sector varied in the start position and size. We call the start position and size as the changeable concepts *position* and *size*, respectively, and there is no unchangeable concept. In the Circle dataset, we introduce the rules *progress* and *multiprogress*.

Results about more instances of the RPM-like dataset and the selective task accuracy compared with RPM solvers are exhibited in the supplementary material.

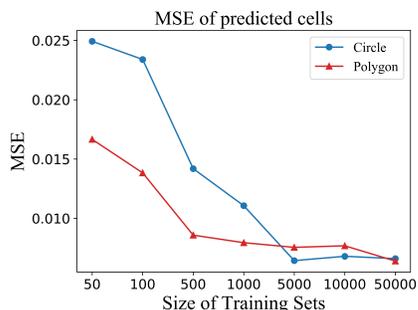


Figure 6: MSEs of our model on different-sized training sets.

	Polygon	Circle
VAE-GAN	5.16e-2	1.08e-1
Proposed	<b>7.55e-3 ± 3.9e-5</b>	<b>6.43e-3 ± 1.8e-5</b>

Table 1: MSEs of two models. Predictions of VAE-GAN are determinate, thus the repeated test is not performed.

Each instance of the dataset contains a disjoint training set, validation set, and test set. For further performance evaluation, we provide various-sized (from 50 to 50000) training sets and one fixed 10000-sample test set.

## Quantitative Results

**Evaluation Metrics.** To assess the painting performance, we use the *Mean Squared Error* (MSE) of predicted cell images. As only predicting right-bottom cells may neglect the relations between panel cells, we thus consider the varying-position MSE score. In a single prediction procedure, the tested model should predict a target cell given the rest 8 panel cells. The model will predict all 9 cells in turn, and calculate the average MSE between predictions and ground truths. This evaluation method estimates both the reconstruction error of predictions and the generalization ability to predict cells in novel positions.

**Data Efficiency.** By increasing the training set size from 50 to 50000, results in Figure 6 display satisfying results of our model on the Polygon training set with only 500 samples, and Circle training set with only 5000 samples. This phenomenon suggests that our model is data-efficient, which is beneficial from the interpretable GP rule reasoning. Therefore, in the following experiment, the training set size is set to 5000 except for additional statements.

**Model Comparison.** Since there exist very few previous works in the RPM painting problem, we choose to compare with a VAE-GAN inpainting model (referred to as VAE-GAN) (Hua and Kunda 2019), which is applied in the missing region inpainting of images (Yu et al. 2018). MSE scores on both two datasets are available in Table 1 where our model outperforms VAE-GAN trained with 50000 samples, indicating that our model’s painting ability can be automatically generalized from the right-bottom to other positions. In contrast, relatively low MSE scores of VAE-GAN implies its limited generalization ability for cell prediction. We pro-

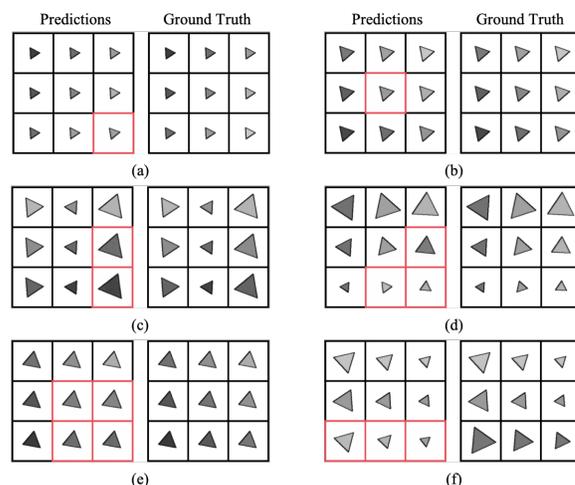


Figure 7: Prediction results on the Polygon dataset. Cell images bounded by red boxes are simultaneously drawn (predicted) by our model while other cells contain ground truth images.

vide the visualized prediction results of VAE-GAN in the supplementary material.

## Cell Prediction

For cell prediction, the varying-position and multi-cell prediction tasks emphasize the understanding of concept-changing rules. Figure 7 visualizes predictions of our model trained on the Polygon training set, which validates that our model can draw accurate predictions, and more importantly, understand concept-changing rules. Figure 7(a) and 7(b) substantially evidence the varying-position prediction ability of our model, that is, it is able to predict all cells in panels by training to predict only right-bottom cells. Furthermore, experiments about the multi-cell prediction in Figure 7(c) to 7(f) show that our model successfully captures the underlying concept-changing rules via GP priors. Predictions in Figure 7(c) are closed to ground truth images. However, in Figure 7(e), predictions deviate obviously in the concept *grayscale*. By further concealing a whole row from a panel as in Figure 7(f), our model derives the horizontal rule *progress* from remaining rows but the generated cell images have distinct sizes and colors compared to the ground truths. We deduce that there is enough context information for the concept-changing rule inference but inadequate information for the missing row reconstruction (because the whole row is removed). These results suggest that our model can execute robust and interpretable cell prediction.

## Panel Sampling and Interpolation

Due to the explicit generative process that transforms cell latent variables  $\lambda$  to images  $x$ , it is possible to construct novel panels. The novel panels will be sampled via the learned latent GP priors. Due to learnable parameters in the kernels, rule function spaces are well-adapted to the dataset. In Figure 5(a), generated panels are diverse: the vertical

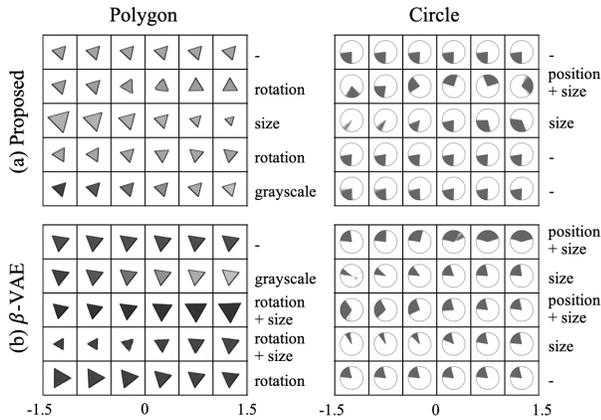


Figure 8: Traverse the dimensions of a cell image latent variable ranged from -1.5 to 1.5 on two datasets. On the right side of images, related visual concepts of each dimension are listed.

Factor VAE Score		
	Polygon	Circle
$\beta$ -VAE	$1.00 \pm 0.0$	$7.59e-1 \pm 1.8e-2$
Proposed	$1.00 \pm 0.0$	$1.00 \pm 0.0$
SAP Score		
	Polygon	Circle
$\beta$ -VAE	$4.40e-1 \pm 6.5e-3$	$9.41e-2 \pm 3.0e-3$
Proposed	$7.80e-1 \pm 3.8e-3$	$7.93e-1 \pm 4.0e-3$

Table 2: Factor VAE Scores and SAP Scores of the  $\beta$ -VAE and our model.

rule *progress* (3rd sample), horizontal rule *progress* (1st sample), and the rule *biprogess* (4th sample); the changing concept *size* (3rd sample) as well as *grayscale* (4th sample). The concept *rotation* of generated panels is not changed progressively in a row or column, which agrees to our definition about the unchangeable set. The generated panels in Figure 5(b) also satisfy the predefined creation principles about the Circle dataset.

For visual concepts, the latent GP priors offer a methodology for function value interpolation. Figure 5 suggest that a larger panel can be generated by interpolation between cell latent variables. Owing to the smoothness of continuous functions, interpolated cell latent variables will be decoded to continuous-varying images. Three subfigures in the right of Figure 5(a) display interpolation results on the Polygon dataset. By interpolating between corner cells, progressive changes in the concepts *grayscale* (all three samples) and *rotation* (1st and 3rd samples) are achieved. Values of the concepts *position* (1st and 2nd samples) and *size* (2nd and 3rd samples) change smoothly as well when interpolating on the Circle dataset.

### Interpretable Visual Concepts

The independence and interpretability of cell image latent variables are essential since the GP rule reasoning is respec-

tively exerted in each visual concept. Disentangled models are considered as a way to impose independence and meanings to dimensions. Because hyperparameters  $\beta$  and  $\gamma$  within Eq.(7) are similar to those in  $\beta$ -VAE loss, we compare our model with  $\beta$ -VAE (Higgins et al. 2017) to investigate the effect of the latent GP priors.  $\beta$ -VAE is trained with 50000 samples to avoid insufficient training.

**Evaluation Metrics.** *Factor VAE Score* (Kim and Mnih 2018) and *SAP Score* (Kumar, Sattigeri, and Balakrishnan 2017) are exploited as evaluation metrics. Factor VAE Score focuses on the modularity among dimensions by using a majority-vote classifier to map dimensions to visual concepts. When computing SAP Score, a correlation matrix of all dimensions and concepts is computed to acquire the difference between top-2 scores, which stresses the code compactness of representations. In short, to get a high Factor VAE Score, we should ensure one dimension to encode at most one concept. If a superior SAP Score is expected, the model has better code a concept with only one dimension. Since two scores inspect distinct aspects of the representation independence and interpretability, we integrate them to comprehensively assess the performance of our model.

Table 2 indicates that our model achieves higher scores than  $\beta$ -VAE on the Polygon and Circle datasets. Figure 8 gives more details through the dimension-wise latent space traverse. When evaluating our model on the Polygon dataset, nearly all dimensions encode one or less concept. Though Factor VAE Score of  $\beta$ -VAE on the Polygon dataset is competitive, visualized results suggest that the 3rd and 4th dimensions encode two concepts together. We consider that perturbations over the 5th dimension impose more deviation on the concept *rotation* compared with the 3rd and 4th dimensions, which supplies clues for the classifier to find the correct mapping. SAP Scores of two models are lower than Factor VAE Scores because both models fail to code one concept with merely one dimension. On the Circle dataset, our model encodes the concept *size* into two dimensions (2nd and 3rd). However,  $\beta$ -VAE applies four dimensions (1st to 4th) for the concept *size* and two dimensions (1st and 3rd) for *position*. It is evidenced by the above experiments that the GP rule reasoning allows our model to learn more reliable representations whose dimensions are independent and correspond to visual concepts (or have no meaning). The latent GP priors reveal the concept-changing rules in panels and bring more inductive biases for better disentanglement.

## Conclusion

In this paper, we propose an interpretable deep latent variable model to solve the RPM painting problem. Our model is capable of learning visual concepts from datasets, GP rule reasoning in latent variable spaces, and painting missing cell images in RPM panels. For the GP rule reasoning, we define underlying concept-changing rules as 2D functions whose priors are the latent GPs. Independent and interpretable latent variables facilitate the dimension-wise abstract reasoning, while inference with entire panels in turn helps the learning of visual concepts. Due to the latent GP priors, the proposed model is not exclusively limited in predicting cell contents, it is able to naturally generate RPM panels.

## Acknowledgments

This research was supported in part by STCSM Projects (20511100400, 18511103104), Shanghai Municipal Science and Technology Major Projects (2017SHZDZX01, 2018SHZDZX01), Shanghai Research and Innovation Functional Program (17DZ2260900), and the Program for Professor of Special Appointment (Eastern Scholar) at Shanghai Institutions of Higher Learning.

## References

- Barrett, D. G.; Hill, F.; Santoro, A.; Morcos, A. S.; and Lillicrap, T. 2018. Measuring abstract reasoning in neural networks. *arXiv preprint arXiv:1807.04225* .
- Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* 35(8): 1798–1828.
- Bilker, W. B.; Hansen, J. A.; Brensinger, C. M.; Richard, J.; Gur, R. E.; and Gur, R. C. 2012. Development of abbreviated nine-item forms of the Raven’s standard progressive matrices test. *Assessment* 19(3): 354–369.
- Chen, R. T.; Li, X.; Grosse, R. B.; and Duvenaud, D. K. 2018. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, 2610–2620.
- Dupont, E. 2018. Learning disentangled joint continuous and discrete representations. In *Advances in Neural Information Processing Systems*, 710–720.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2017. beta-vaes: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*.
- Hu, S.; Ma, Y.; Liu, X.; Wei, Y.; and Bai, S. 2020. Hierarchical Rule Induction Network for Abstract Visual Reasoning. *arXiv preprint arXiv:2002.06838* .
- Hua, T.; and Kunda, M. 2019. Modeling Gestalt Visual Reasoning on the Raven’s Progressive Matrices Intelligence Test Using Generative Image Inpainting Techniques. *arXiv preprint arXiv:1911.07736* .
- Jeong, Y.; and Song, H. O. 2019. Learning discrete and continuous factors of data via alternating disentanglement. *arXiv preprint arXiv:1905.09432* .
- Kim, H.; and Mnih, A. 2018. Disentangling by factorising. *arXiv preprint arXiv:1802.05983* .
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* .
- Kumar, A.; Sattigeri, P.; and Balakrishnan, A. 2017. Variational inference of disentangled latent concepts from unlabeled observations. *arXiv preprint arXiv:1711.00848* .
- Little, D. R.; Lewandowsky, S.; and Griffiths, T. L. 2012. A Bayesian model of rule induction in Raven’s Progressive Matrices. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 34.
- Locatello, F.; Bauer, S.; Lucic, M.; Raetsch, G.; Gelly, S.; Schölkopf, B.; and Bachem, O. 2019. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, 4114–4124.
- Lovett, A.; and Forbus, K. 2017. Modeling visual problem solving as analogical reasoning. *Psychological review* 124(1): 60.
- Lovett, A.; Forbus, K.; and Usher, J. 2010. A structure-mapping model of Raven’s Progressive Matrices. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 32.
- Raven, J. C.; and Court, J. 1938. *Raven’s progressive matrices*. Western Psychological Services Los Angeles, CA.
- Santoro, A.; Raposo, D.; Barrett, D. G.; Malinowski, M.; Pascanu, R.; Battaglia, P.; and Lillicrap, T. 2017. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, 4967–4976.
- Skafta, N.; and Hauberg, S. 2019. Explicit Disentanglement of Appearance and Perspective in Generative Models. In *Advances in Neural Information Processing Systems*, 1018–1028.
- Steenbrugge, X.; Leroux, S.; Verbelen, T.; and Dhoedt, B. 2018. Improving generalization for abstract reasoning tasks using disentangled feature representations. *arXiv preprint arXiv:1811.04784* .
- Sun, S.; Zhang, G.; Wang, C.; Zeng, W.; Li, J.; and Grosse, R. 2018. Differentiable compositional kernel learning for Gaussian processes. *arXiv preprint arXiv:1806.04326* .
- Teney, D.; Wang, P.; Cao, J.; Liu, L.; Shen, C.; and van den Hengel, A. 2020. V-PROM: A Benchmark for Visual Reasoning Using Visual Progressive Matrices. In *AAAI*, 12071–12078.
- van Steenkiste, S.; Locatello, F.; Schmidhuber, J.; and Bachem, O. 2019. Are Disentangled Representations Helpful for Abstract Visual Reasoning? In *Advances in Neural Information Processing Systems*, 14245–14258.
- Watters, N.; Matthey, L.; Burgess, C. P.; and Lerchner, A. 2019. Spatial broadcast decoder: A simple architecture for learning disentangled representations in vaes. *arXiv preprint arXiv:1901.07017* .
- Williams, C. K.; and Rasmussen, C. E. 2006. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.
- Wilson, A. G.; Hu, Z.; Salakhutdinov, R.; and Xing, E. P. 2016. Deep kernel learning. In *Artificial intelligence and statistics*, 370–378.
- Yu, J.; Lin, Z.; Yang, J.; Shen, X.; Lu, X.; and Huang, T. S. 2018. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5505–5514.
- Zhang, C.; Gao, F.; Jia, B.; Zhu, Y.; and Zhu, S.-C. 2019a. Raven: A dataset for relational and analogical visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5317–5327.

Zhang, C.; Jia, B.; Gao, F.; Zhu, Y.; Lu, H.; and Zhu, S.-C. 2019b. Learning perceptual inference by contrasting. In *Advances in Neural Information Processing Systems*, 1075–1087.