
Amortized Inference of Variational Bounds for Learning Noisy-OR

Yiming Yan

University of Southern California

Melissa Ailem

University of Southern California

Fei Sha

Google Research

Abstract

Classical approaches for approximate inference depend on cleverly designed variational distributions and bounds. Modern approaches employ amortized variational inference, which uses a neural network to approximate any posterior without leveraging the structures of the generative models. In this paper, we propose Amortized Conjugate Posterior (ACP), a hybrid approach taking advantages of both types of approaches. Specifically, we use the classical methods to derive specific forms of posterior distributions and then learn the variational parameters using amortized inference. We study the effectiveness of the proposed approach on the NOISY-OR model and compare to both the classical and the modern approaches for approximate inference and parameter learning. Our results show that the proposed method outperforms or are at par with other approaches.

1 INTRODUCTION

Classical techniques in probabilistic graphical models exploit (tractable) structures heavily for approximate inference (Koller et al., 2009). Well-studied examples are mean-field approaches by assuming factorized forms of posteriors (Jordan et al., 1999), (conjugate) variational bounds on likelihoods (Jaakkola and Jordan, 1999, 2000), and others (Saul et al., 1996; Saul and Jordan, 1996). In these methods, the forms of the approximate posteriors depend on the model structures, the definitions of the conditional probability tables or distributions, and the priors. For instance, deriving variational bounds often requires identifying special properties such as convexity and concavity of likelihood (or partition) functions. And in some cases, the

derivation also depends on whether an upper-bound or lower-bound is needed (Jebara and Pentland, 2001).

In contrast, recent approaches in amortized variational inference (AVI) use neural networks to represent posteriors and reparameterization tricks to compute the likelihoods with Monte Carlo samples (Kingma and Welling, 2013; Mnih and Gregor, 2014). Even earlier, neural networks were applied to approximate posterior distribution in a supervised manner (Morris, 2001). What is appealing in this type of methods is that selecting the inference neural network requires significantly reduced efforts, and the structure of the generative model (and the corresponding likelihood function) does not directly come into play in determining the inference network. In other words, the inference network (*i.e.* the encoder) and the generative model (*i.e.* the decoder) are parameterized independently, without explicitly sharing information. As such, with a large amount of training data, a high-capacity inference network is able to approximate the posterior well, and learning the generative model can be effective as the variance of the Monte Carlo sampling can be reduced. However, when the amount of the training data is small, the inference network can overfit and estimating the generative model has a high variance.

Is there a way to combine these two different types of approaches? In this paper, we take a step in this direction. Our main idea is to use the above-mentioned classical methods to derive approximate but tractable posteriors (and approximate likelihood functions) and then identify the optimal variational parameters by learning a neural network.

The key difference from the classical approaches is that the variational parameters are not optimized to give the tightest bounds on the likelihood functions but to maximize the evidence lower bound (ELBO). On the other hand, the key difference from the AVI is that the posterior and the generative model share information such that the posterior contributes explicitly to the gradient of the ELBO with respect to the generative model parameters.

We apply this new hybrid approach to the NOISY-OR

model, which has been studied for binary observations and latent variables. Earlier work such as (Jaakkola and Jordan, 1999; Singliar and Hauskrecht, 2006) proposed classical variational inference approaches. More recently, polynomial-time algorithms are designed to learn the structure and parameters of NOISY-OR model (Halpern and Sontag, 2013; Jernite et al., 2013). However they require strong assumptions on the structure of the graph. Halpern (Halpern, 2016) studied semi-supervised method with stochastic variational learning on NOISY-OR and achieved good performance in parameter recovering. Moreover, stochastic variational inference (SVI) has been designed for hierarchical NOISY-OR (Ji et al., 2019) for faster convergence and achieved on-par performance with its batch counterpart. Taking advantages of the conjugate bounds, the proposed approach generalizes better than AVI and SVI when there is limited data for fitting while remaining equally competitive when there is more data.

In the following, we will describe related work by reviewing different types of variational inference techniques in Section 2 for the NOISY-OR model. We describe our approach in Section 3. In Section 4, we show our experiment results, followed by discussion in section 5.

2 VARIATIONAL INFERENCE FOR NOISY-OR

2.1 Basic Ideas

We are interested in modeling a random variable \mathbf{x} with a generative latent variable model, where \mathbf{z} denotes the latent variables. θ denotes the model parameter and N denotes the total number of data points in the training set $\mathbf{X} = \{\mathbf{x}^{(n)}, n = 1, 2, \dots, N\}$. The log-likelihood is defined as

$$\begin{aligned} L &= \sum_n \log \int p(\mathbf{x}^{(n)}, \mathbf{z}; \theta) d\mathbf{z} \\ &= \sum_n \log \int p(\mathbf{x}^{(n)} | \mathbf{z}; \theta) p(\mathbf{z}; \theta) d\mathbf{z} \end{aligned} \quad (1)$$

For discrete latent variables, the integral is interpreted as summing over all possible configurations of \mathbf{z} .

Due to the logarithm before the integral, estimating θ is intractable. We introduce a variational distribution $q(\mathbf{z} | \mathbf{x}; \phi)$ to approximate the posterior $p(\mathbf{z} | \mathbf{x})$. This gives rise to maximizing the ELBO,

$$\begin{aligned} \mathcal{L}(\mathbf{x}^{(n)}; \theta, \phi) &= \mathbb{E}_{q(\mathbf{z} | \mathbf{x}^{(n)}; \phi)} [\log p(\mathbf{x}^{(n)} | \mathbf{z}; \theta)] \\ &\quad - D_{KL}(q(\mathbf{z} | \mathbf{x}^{(n)}; \phi) \| p(\mathbf{z}; \theta)) \leq \log p(\mathbf{x}^{(n)}; \theta) \end{aligned} \quad (2)$$

where the $D_{KL}(q \| p)$ denotes the Kullback-Leibler (KL) divergence between the distributions q and p , and ϕ is




Figure 1: NOISY-OR model in plate notations.

the variational parameter. To avoid notation cluttering, we will drop the subscript $^{(n)}$ and omit θ and ϕ whenever the context is clear.

Classical approaches restrict the distribution family for $q(\mathbf{z} | \mathbf{x}; \phi)$ so that the expected conditional likelihood (the first term in the ELBO) can be computed. For example, mean-field approximation assumes a factorized form of the variational distribution. Unfortunately, even for some common likelihood functions $p(\mathbf{x} | \mathbf{z})$, the factorized form does not turn the expectation tractable. We describe one such model and then describe how the classical and recent approaches tackle such challenges.

2.2 NOISY-OR

NOISY-OR is a bipartite directed graph modeling the dependencies among binary observations. The structure is shown in Fig. 1, where $\mathbf{x} \in \{0, 1\}^D$ represents the observed D -dimensional data, $\mathbf{z} \in \{0, 1\}^{K+1}$ are the latent variables (with $z_0 = 1$). The model defines the distribution

$$p(\mathbf{x}, \mathbf{z}) = \prod_{i=1}^D p(x_i | \mathbf{z}) \prod_{k=0}^K p(z_k) \quad (3)$$

where $p(x_i | \mathbf{z})$ and $p(z_k)$ are Bernoulli distributions. In particular,

$$p(x_i = 0 | \mathbf{z}) = (1 - p_{i0}) \prod_{k=1}^K (1 - p_{ik})^{z_k} \quad (4)$$

where $p_{ik} = p(x_i = 1 | z_k = 1)$. Redefining $\theta_{ik} = -\log(1 - p_{ik})$, we have

$$p(x_i = 0 | \mathbf{z}) = e^{-\theta_{i0} - \sum_{k=1}^K \theta_{ik} z_k} = e^{-\theta_i^T \mathbf{z}} \quad (5)$$

where we slightly abuse the notation on $\theta_i = \{\theta_{i0}, \theta_{i1}, \dots, \theta_{iK}\}$ and $\mathbf{z} = \{1, z_1, z_2, \dots, z_K\}$. In more compact form, we have

$$p(x_i | \mathbf{z}) = p(x_i = 1 | \mathbf{z})^{x_i} p(x_i = 0 | \mathbf{z})^{1-x_i}. \quad (6)$$

where x_i takes value of either 0 or 1.

The form of the likelihood of positive observation $p(x_i = 1 | \mathbf{z})$ makes it intractable for computing its expectation of logarithm, even when the posterior is approximated in factorized form.

2.3 Conjugate Dual Inference

The conjugate dual for approximating $p(x_i = 1|\mathbf{z})$ of NOISY-OR model was first introduced in (Jaakkola and Jordan, 1999), where an approximate upper-bound of likelihood is:

$$p(x_i = 1|\mathbf{z}) = e^{\log(1-e^{-a_i})} = e^{f(a_i)} \leq e^{\psi_i a_i - g(\psi_i)} = \tilde{p}(x_i = 1|\mathbf{z}, \psi_i) \quad (7)$$

Here $f(a_i)$ is a concave function with $a_i = \boldsymbol{\theta}_i^T \mathbf{z}$, ψ_i is a variational parameter associated with a_i , and $g(\cdot)$ is the conjugate dual function of $f(\cdot)$. For $f(s) = \log(1-e^{-s})$, we have

$$g(t) = -t \log t + (t + 1) \log(t + 1) \quad (8)$$

The detailed derivation can be found in (Jaakkola and Jordan, 1999).

The resulting upper-bound $\log \tilde{p}(x_i = 1|\mathbf{z}, \psi_i)$ has an appealing property that it is linear in a_i (hence, \mathbf{z}). Thus after applying the Bayes rule, we achieve a factorized (variational) posterior distribution

$$q(\mathbf{z}|\mathbf{x}, \boldsymbol{\psi}) \propto \prod_{i:x_i=1}^D \tilde{p}(x_i = 1|\mathbf{z}, \psi_i) \cdot \prod_{i:x_i=0}^D p(x_i = 0|\mathbf{z}) \prod_{k=0}^K p(z_k) \quad (9)$$

Note that we only use the upper-bound for positive observations where $x_i = 1$. For negative observations, we use the true likelihood. After re-organizing the terms, we observe that

$$q(\mathbf{z}|\mathbf{x}, \boldsymbol{\psi}) = \prod_{k=1}^K q(z_k|\mathbf{x}, \boldsymbol{\psi}) \quad (10)$$

Namely, the variational posterior factorizes. And each factor is

$$q(z_k = 1|\mathbf{x}, \boldsymbol{\psi}) = \sigma \left(\sum_{i:x_i=1} \psi_i \theta_{ik} - \sum_{i:x_i=0} \theta_{ik} + \log \frac{\mu_k}{1 - \mu_k} \right) \quad (11)$$

where $\sigma(\cdot)$ is the sigmoid function and μ_k is the parameter of prior distribution $\mu_k = p(z_k = 1)$. The detailed derivation can be found in supplementary material.

Moreover, computing the expectation of variational upper-bound $\tilde{p}(x_i = 1|\mathbf{z}, \psi_i)$ with respect to the factorized (variational) posterior distribution $q(\mathbf{z}|\mathbf{x}, \boldsymbol{\psi})$ is trivial. And the gradients can be computed directly without sampling. Hence we take $\tilde{p}(x_i = 1|\mathbf{z}, \psi_i)$ as the approximation to $p(x_i = 1|\mathbf{z})$. Similarly, the likelihood conjugate lower-bound can be applied for variational

inference. Interested readers could refer to (Jaakkola and Jordan, 1999; Šingliar and Hauskrecht, 2006).

The conjugate dual inference (CDI) choose the best ψ_i to achieve the tightest upper-bound

$$\begin{aligned} \psi_i^* &= \arg \min_{\psi_i} \{\tilde{p}(x_i = 1)\} \\ &= \arg \min_{\psi_i} \{\mathbb{E}_{p(\mathbf{z})} [\tilde{p}(x_i = 1|\mathbf{z}, \psi_i)]\} \end{aligned} \quad (12)$$

However, in the next section, we will show that we can choose ψ_i^* differently, leading to a different type of inference technique.

2.4 Stochastic Variational Inference

Recently, SVI has been proposed for (hierarchical) NOISY-OR model (Ji et al., 2019). Similar to CDI, the conjugate lower-bound in (Jaakkola and Jordan, 1999; Šingliar and Hauskrecht, 2006) is applied to approximate the positive likelihood for its tractability of taking expectation. Different from CDI, where the variational posterior is constructed as eq. (12), the parameters of variational posterior are treated as free-parameters and optimized directly. In other words, two sets of variational parameters are introduced:

- $\boldsymbol{\psi}$: variational parameters introduced by the conjugate bound (eq. (7)).
- $\boldsymbol{\phi}$: parameters of variational posterior distribution $q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi}) = \prod_{k=1}^K q(z_k|\mathbf{x}; \boldsymbol{\phi})$.

Comparing to CDI, the unconstrained posterior gives higher inference capacity. Meanwhile, SVI can handle large-scale datasets and converges faster by taking mini-batches while CDI considers the whole batch during learning.

2.5 Amortized Variational Inference

Under the auto-encoding variational Bayes framework (Kingma and Welling, 2013), AVI uses global parameters to predict the parameters of approximate posterior distribution directly. For instance, (Kingma and Welling, 2013) predicts the parameters $\boldsymbol{\phi}$ of $q(\mathbf{z}|\mathbf{x}^{(n)}; \boldsymbol{\phi})$ as $\boldsymbol{\phi} = f(\mathbf{x}^{(n)}; \boldsymbol{\lambda})$. Here, $\boldsymbol{\lambda}$ is the global trainable parameter, which is shared across all $\mathbf{x}^{(n)}$, $n = 1, \dots, N$. As a special case, $f(\cdot)$ can be a neural network and $\boldsymbol{\phi} = \text{MLP}(\mathbf{x}; \boldsymbol{\lambda})$.

The expectation of log-likelihood requires sampling the latent variable $\mathbf{z} \sim q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})$. However, the gradients cannot be back-propagated through stochastic random variable \mathbf{z} . Hence for certain types of distributions $q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})$, the reparameterization trick can be applied to reparameterize the random variable \mathbf{z} using a differentiable function (such as neural network)

$\mathbf{z} = g_{\lambda}(\mathbf{x}, \epsilon)$, where $\epsilon \sim p(\epsilon)$ (a known and easily sampled distribution). Then the expected log-likelihood can be rewritten as

$$\mathbb{E}_{q(\mathbf{z}|\mathbf{x};\phi)}[\log p(\mathbf{x}|\mathbf{z})] = \mathbb{E}_{p(\epsilon)}[\log p(\mathbf{x}|g_{\lambda}(\mathbf{x}, \epsilon))] \quad (13)$$

which can then be computed with Monte Carlo sampling from $p(\epsilon)$.

AVI utilizes the advantages of deep neural networks as inference model without explicitly deriving complex conjugate dual functions for likelihoods, and has the power of approximating posterior flexibly. However, when there is not enough training data, the inference model might overfit and leads to a large variance in estimating the generative model. We will introduce our method in the next section to tackle this problem.

3 AMORTIZED CONJUGATE POSTERIOR

In this section, we introduce our inference strategy, Amortized Conjugate Posterior (ACP), a hybrid approach combining the classical CDI and the recent AVI approaches.

Instead of seeking the tightest upper-bounds for the likelihood function (for positive observations), we propose to optimize ψ to maximize the ELBO in eq. (2). Moreover, to amortize the inference, we parameterize ψ as in AVI,

$$\psi = \text{MLP}(\mathbf{x}; \phi) \quad (14)$$

where the parameters ϕ of the neural network are shared by all data points. Namely, for each $\mathbf{x}^{(n)}$, its variational distribution is

$$\begin{aligned} q(\mathbf{z}|\mathbf{x}^{(n)}; \psi^{(n)}) &= q(\mathbf{z}|\mathbf{x}^{(n)}; \text{MLP}(\mathbf{x}^{(n)}; \phi)) \\ &\equiv q(\mathbf{z}|\mathbf{x}^{(n)}; \phi) \end{aligned} \quad (15)$$

As in AVI, to optimize both θ and ϕ , we use Monte Carlo sampling to compute the ELBO (and its gradients with respect to the parameters). Note that in NOISY-OR, the ELBO can be written as

$$\begin{aligned} \mathcal{L}(\mathbf{x}; \theta, \phi) &= \sum_{i:x_i=1}^D \mathbb{E}_{q(\mathbf{z}|\mathbf{x};\phi)}[\log p(x_i = 1|\mathbf{z}; \theta)] \\ &+ \sum_{i:x_i=0}^D \mathbb{E}_{q(\mathbf{z}|\mathbf{x};\phi)}[\log p(x_i = 0|\mathbf{z}; \theta)] \\ &- D_{KL}(q(\mathbf{z}|\mathbf{x}; \phi) || p(\mathbf{z}; \theta)) \leq \log p(\mathbf{x}; \theta) \end{aligned} \quad (16)$$

where the expectation of the positive log-likelihood (the first term in eq. (16)) is intractable while the negative log-likelihood (the second term) and the KL divergence can be computed analytically. Hence, we

Table 1: Contrast of variational inference approaches

	objective	constrained posterior	amortized
AVI	ELBO	no	yes
LB-CDI	LB	yes (LB)	no
UB-CDI	UB	yes (UB)	no
SVI	LB	no	no
ACP (ours)	ELBO	yes (UB)	yes

need to estimate the first term using Monte Carlo sampling.

Specifically, for each training data point \mathbf{x} , we use eq. (14) to compute the variational parameters, and use the form of posterior as eq. (11). We then sample from the posterior — since $q(\mathbf{z}|\mathbf{x}; \phi)$ is Bernoulli, we use the Gumbel-Softmax reparameterization trick (Jang et al., 2016). The samples are then used to compute the expectation of *true* positive likelihoods $p(x_i = 1|\mathbf{z})$ (and their gradients). The pseudo-code for parameter estimation of ACP is presented in the supplementary material.

The key difference from AVI is that the (variational) posterior has also dependency on the generative model parameters θ . While they contribute indirectly to the gradients of the expected conditional likelihoods for positive observations in the ELBO, the expected likelihoods for negative ones and the KL divergence between the posterior and the prior are analytically tractable and the gradients with respect to θ are directly used to optimize the parameters. In AVI, the expected conditional likelihood is a constant with respect to the generative model parameters and it does not contribute to the update.

In Table 1, we summarize various approaches. LB and UB in the parenthesis stand for whether the constrained posterior is derive from the lower-bound or upper-bound of the ELBO. Comparing to AVI, the specific form (eq. (11)) of the variational posterior – how evidence \mathbf{x} is incorporated and the architecture is formed – contains model-specific knowledge, which could prevent overfitting and improve generalization. Comparing to CDI and SVI, ACP optimizes ELBO directly instead of optimizing the lower-bound or upper-bound of ELBO, which gives better performance. Our empirical results support this claim.

4 EXPERIMENTS

We compare the performance of different inference methods on both synthetic and real-world datasets. The synthetic datasets are described in section 4.1. We use them as the ground-truth is known. We investigate how the characteristics of different approaches vary with respect to the amount of training data in section 4.2-4.4. The real-world dataset is described in

Table 2: Parameters for synthetic datasets of *patterned* weights

Dataset	D	K	μ_k	N_{test}	Sparsity
SYN-PATTERN	64	8	0.125	1000	89.0%
MULTI-MNIST	784	10	0.2	5000	80.3%

 Table 3: Parameters for synthetic datasets of *random* weights (SYN-RANDOM)

size	α_θ	β_θ	α_μ	β_μ	s	N_{test}	Sparsity
$D = 50$	1	5	1	10	0.95	1000	94.2%
$K = 100$	2	5	2	5	0.95	1000	71.8%
$D = 500$	1	5	1	20	0.995	2000	98.4%
$K = 500$	1	20	1	20	0.95	2000	95.3%
$D = 500$	1	10	1	10	0.95	2000	73.6%
$K = 100$	1	5	1	5	0.95	5000	89.2%

section 4.5 and is used to illustrate one of the practical uses of the inference methods in topic modeling. Our experiments show that the proposed ACP method outperforms other variational approaches in most cases.

4.1 Synthetic Datasets

For reproducibility, we detail the learning and optimization hyperparameters in the Supplementary Materials.

Data with Patterned Weights We created this type of synthetic datasets with the following recipe:

- Select the parameters μ_k of the prior distributions $p(z_k = 1) = \mu_k$, $k = 1 \cdots K$.
- Select the generative model parameter $\theta \in \mathbb{R}^{D \times K}$ and the “leak” probability $\theta_0 \in \mathbb{R}^D$. θ and θ_0 needs to be non-negative.
- Sample N latent variables $\mathbf{z}^{(n)}$, $n = 1, \dots, N$ from the prior distribution $p(\mathbf{z})$.
- Sample N observed data points $\mathbf{x}^{(n)}$ from the conditional probability $p(\mathbf{x}^{(n)} | \mathbf{z}^{(n)})$.

where K is the number of latent variables and D is the number of observed variables. The two types of datasets are described below.

We generated two datasets: SYN-PATTERN and MULTI-MNIST. The configurations of D , K and μ_k are specified in Table 2. The model parameters θ and θ_0 are reshaped from the patterns depicted in Fig. 2a (SYN-PATTERN) and 2b (MULTI-MNIST). Each pattern is reshaped to a D -dimensional vector representing $\theta_k \in \mathcal{R}^D$, where the white pixels in the k th pattern indicates the corresponding parameters θ_{ik} to be 0. And

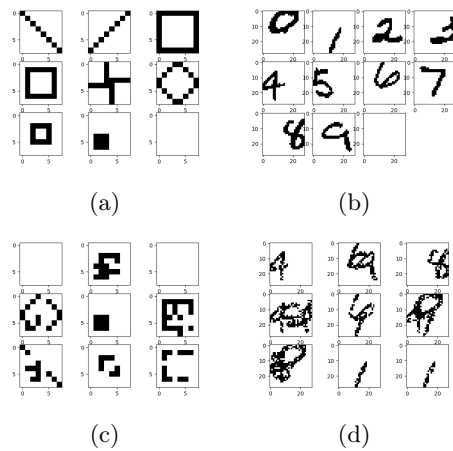


Figure 2: (a) and (b). The patterns of the parameters θ for the SYN-PATTERN and MULTI-MNIST datasets, respectively. (c) and (d). Several observed \mathbf{x} , sampled from the SYN-PATTERN and MULTI-MNIST datasets.

the last pattern refers to the pattern of “leak” probability. The values of non-zero θ_{ik} (*i.e.*, black pixels) are $-\log(1 - 0.8)$ which means $p(x_i = 1 | z_k = 1) = 0.8$. Fig. 2c and 2d show some data points sampled from the two datasets. Each data point is a combination of their parameter patterns with missing parts. We use N_{test} data points for validation and another N_{test} for testing. The values of N_{test} are reported in Table 2.

Data with Random Weights We also built a synthetic datasets with random sampled weights, namely SYN-RANDOM. Concretely, θ , θ_0 and μ are sampled from distributions $\text{BETA}(\alpha_\theta, \beta_\theta)$ and $\text{BETA}(\alpha_\mu, \beta_\mu)$. Moreover, the sparsity of the dataset is controlled by constraining the sparsity degree s to be $0 < s < 1$, which can be viewed as the probability of removing the connection between z_k and x_i by setting $\theta_{ik} = 0$, to enforce the sparse connections between latent and observed variables. The random removal of those connections can create orphaned variables without any connections. Thus, we randomly add connections to the latent and observed variables which do not have any connection. We describe the configurations of SYN-RANDOM in Table 3.


The SPARSITY is defined as the percentage of negative observations in the dataset $\text{SPARSITY} = \frac{1}{ND} \sum_{n=1}^N \sum_{i=1}^D \mathbb{I}(x_i^{(n)} = 0)$ where lower SPARSITY value indicates denser dataset.

4.2 Inference

Herein, we compare ACP with baselines on their abilities of accurately approximating the posterior distribution. We fix the generative model with its ground truth parameters, and evaluate the accuracy of infer-

N_{train}	Method	NELBO	F1	EM
1000	AVI	14.0	94.5	91.0
	ACP	14.4	94.3	90.4
100	AVI	18.4	86.6	76.8
	ACP	17.4	87.1	81.1
20	AVI	37.2	49.2	47.6
	ACP	22.2	76.1	64.0
-	SVI	22.6	62.5	43.0
	LB-CDI	22.4	63.6	45.4
	UB-CDI	96.6	24.6	39.0

(a)



(b)

Figure 3: Inference results on SYN-PATTERN. (a). NELBO: negative ELBO (lower is better). Higher F1 and EM are better. (b). The learning curve of UB-CDI. It shows that tighter variational bound does not imply better approximate posterior.

ence by computing the ELBO with respect to different approximate posterior distributions on the *held-out* data. Since the generative model is fixed, the ELBO is maximized when $q(\mathbf{z}|\mathbf{x}; \phi) = p(\mathbf{z}|\mathbf{x}; \theta)$. Hence higher ELBO indicates better inference performance. Moreover, we compare the ground truth $\mathbf{z}^{(n)}$ and $\hat{\mathbf{z}}^{(n)} \sim q(\mathbf{z}|\mathbf{x}^{(n)}; \phi)$, $n = 1, \dots, N$ using macro F1 and Exact Match (EM) scores as inference accuracy.

We analyzed two types of CDI, namely UB-CDI and LB-CDI as CDI with variational Upper-Bound and Lower-Bound respectively. The parameters in CDI are optimized following the optimization strategy in (Jaakkola and Jordan, 1999; Singliar and Hauskrecht, 2006), where we find the tightest likelihood upper-bound (or lower-bound) using fix-point optimization and use it to compute posterior (eq. (12)). Similarly, SVI is trained to maximize the lower-bound of the ELBO.

For AVI and ACP, the variational parameters ϕ are optimized to maximize the ELBO. We report the results of learned ϕ under different amount of training data. For non-amortized methods SVI and CDI, we optimize the variational parameter on samples from held-out set directly.

Fig. 3a shows the experiment results, where N_{train} indicates the number of training data. We observe that with a sufficient amount of training data ($N_{train} = 1000$), AVI achieves slightly better performance due to its high flexibility to approximate the posterior. Yet when we have only limited amount of training data, ACP gains huge advantages over AVI.

We observed that UB-CDI achieves very poor performance. The learning curve of UB-CDI is depicted in Fig. 3b, where LL-UB indicates the Log-Likelihood Upper-Bound:

$$\log \tilde{p}(\mathbf{x}|\phi) = \log \sum_{\mathbf{z}} \prod_{i:x_i=1} \tilde{p}(x_i = 1|\mathbf{z}, \phi) \cdot \prod_{j:x_j=0} p(x_j = 0|\mathbf{z})p(\mathbf{z}).$$

In the first 10 rounds of fix-point optimization, LL-UB becomes tighter as we optimize more iterations and then converges. However, with a tighter upper-bound,

the ELBO first drops quickly, and then improves *only slightly* during optimization. The final ELBO after convergence is much worse, even comparing to the initial point. This observation indicates that a tighter likelihood upper-bound is not necessarily equivalent to a better approximate posterior. Thus we will not compare to UB-CDI in the rest of our experiments.

Although LB-CDI is also optimized to obtain the tightest bound, its performance is much better than UB-CDI. The reason is that LB-CDI optimizes the lower-bound of ELBO. As it optimizes the lower-bound, the ELBO is pushed up. However the performance of LB-CDI is still much worse than amortized methods when training data is sufficient. Hence, optimizing the ELBO directly is very helpful comparing to optimizing its approximation. The low performance of SVI confirms this observation.

Note that while SVI performance does not rely on the number of training samples, ACP (and AVI) can improve rapidly with the increasing number of training data. In particular, since ground-truth parameters are generally not known and need to be learned from data, we anticipate ACP and AVI have more advantages in the learning setting, which we describe below.

4.3 Parameter Estimation

To further analyze the properties of ACP, we jointly train the generative and inference models on SYN-PATTERN and MULTI-MNIST dataset. The goal is to recover the patterns of θ in a fully unsupervised way, and compare the performance of ACP over the baselines with different amount of training data.

Fig. 4 shows the results for AVI and ACP on SYN-PATTERN¹. From Fig. 4a and 4b, we observe that both AVI and ACP recover all the patterns and achieve similar performance. When reducing the amount of training data to 200 (Fig. 4c and 4d), we observe that ACP still reconstructs all the patterns with slightly worse performance. In contrast, the performance of

¹Due to the low performance of SVI and LB-CDI, their results are moved to the Supplementary Materials.


(a) AVI, $N_{train} = 1000$ (b) ACP, $N_{train} = 1000$ (c) AVI, $N_{train} = 200$ (d) ACP, $N_{train} = 200$

Figure 4: The recovered parameters of SYN-PATTERN after training with 1000 and 200 data points using AVI and ACP.

AVI degrades more severely. Specifically, two patterns out of the eight are not recovered (*i.e.* the middle left and middle right patterns in Fig. 4c). Additionally, some patterns are merged (*i.e.* the upper right and middle patterns in Fig. 4c). This result indicates that model dependent posterior form is helpful in structured inference for learning useful latent representations, especially when we have small amount of training data. Similarly, we perform experiments on MULTI-MNIST dataset. The results are reported in the Supplementary Materials.

4.4 Generative Modeling with Synthetic Datasets

While the previous sections focus on comparing methods on inference accuracy and parameter recovery, this section focuses on generative modeling and compares methods in the metric of negative ELBO on held-out set. According to eq. (11), sparse data requires less parameter ψ to be approximated. Thus, in addition to varying the amount of training data, we also control the sparsity of the dataset and evaluate how it would affect the generative modeling. We use SYN-RANDOM datasets described in Table 2 for these studies.

Fig. 5 shows the experiment results with different amount of training data in various degrees of sparsity. ACP consistently outperforms competing methods. It performs on-par or slightly better than SVI, AVI and LB-CDI when using a large amount of data but performs noticeably better when using a small amount of data. Intuitively, denser data will result in worse approximation as the variational bound is applied to an increased number of positive observations. However, in Fig. 5, we observe that ACP is not affected much by this phenomenon. An explanation is that ACP is not optimized to achieve the tightest variational bound, yet to improve the overall performance with respect to the ELBO.

Table 4 shows the modeling results on SYN-RANDOM with 500 observed and 100 latent variables. Here we did not experiment LB-CDI with large training size (greater than 5k) due to the computational complexity. From Table 4, we can observe that our ACP achieves

Table 4: Generative modeling on SYN-RANDOM, where $D = 500$, $K = 100$, *sparsity* = 89.2%

N_{train}	AVI	SVI	LB-CDI	ACP
40k	150.7	156.3	-	149.5
20k	151.0	156.5	-	149.8
10k	156.0	156.9	-	150.3
5k	157.6	157.4	-	151.2
1k	159.2	158.0	158.8	157.2

on-par or slightly better performance comparing with AVI when we have sufficient training data, while SVI is significantly worse than both amortized approaches. When we have limited training data, the performance of AVI degrades faster than ACP. The large variance of AVI leads to the worst performance when the training set is limited. Our ACP method outperforms AVI and SVI with both sufficient and limited training data.

4.5 Topic Models

We compare AVI and ACP on topic modeling. We use the titles of all the Neural Information Processing Systems (NeurIPS) papers from 1987 to 2016 (nip, 2019). Each observed data point $\mathbf{x}^{(n)}$ is a D -dimensional binary vector representing a paper’s title, where D is the size of the vocabulary. The value of $x_i^{(n)}$, $i = 1, \dots, D$ indicates the presence/absence of word w_i in the n -th title. After word lemmatization, removing stop words, the 5 most common words and the words with less than 5 occurrences in the whole corpus, we obtain a dataset with 7241 data points and 1216 unique words. The average length of the paper title is 4.49 after pre-processing. We use 1000 data points for validation, 1000 for testing, and the remaining ones for training. We model the data with $K = 20$ latent variables.

Each latent variable z_k is interpreted as a topic capturing a distribution of words. To further show the semantic coherence of words in each topic, we report the point-wise mutual information (PMI), which has been shown to be highly correlated with human judgment in assessing word relatedness (Newman et al., 2009), between word pairs of each topic. To do so we use the whole English WIKIPEDIA corpus, that consists of approximately 4 millions of documents and 2 billions

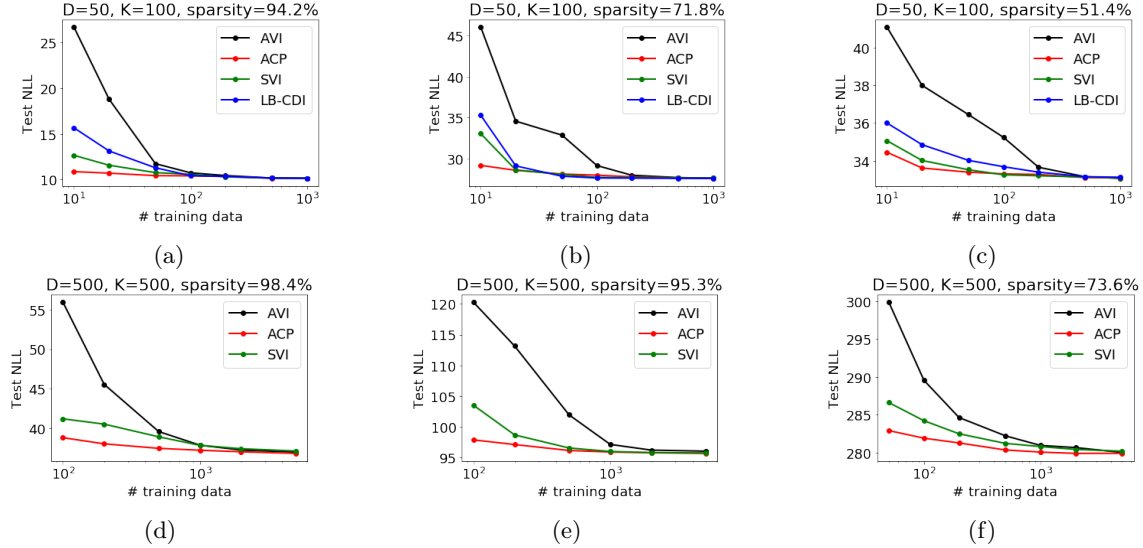


Figure 5: Comparison of AVI and ACP in NOISY-OR under different model structure and different data sparsity.

Table 5: Top 10 words inferred on NeurIPS Titles dataset for top 4 topics with 5241 and 2000 training data.

$N_{train} = 5241$, ACP (PMI = 2.78)				$N_{train} = 5241$, AVI (PMI = 2.49)			
Topic 1	Topic 2	Topic 3	Topic 4	Topic 1	Topic 2	Topic 3	Topic 4
classification	bayesian	sparse	application	process	inference	multi	regression
application	base	analysis	classification	inference	inference	bayesian	gaussian
via	optimization	deep	process	gaussian	bayesian	probabilistic	via
method	adaptive	datum	linear	markov	analysis	information	estimation
process	method	method	stochastic	datum	gaussian	inference	inference
multi	function	estimation	analysis	mixture	datum	dynamic	analysis
bayesian	object	multi	multi	analysis	mixture	approach	linear
kernel	estimation	feature	datum	bayesian	variational	function	process
feature	information	convex	time	variational	approach	application	optimization
image	datum	probabilistic	via	dynamic	probabilistic	process	
PMI : 3.82	PMI : 3.22	PMI : 3.18	PMI : 3.16	PMI : 3.74	PMI : 3.73	PMI : 3.69	PMI : 2.98
$N_{train} = 2000$, ACP (PMI = 2.75)				$N_{train} = 2000$, AVI (PMI = 2.55)			
Topic 1	Topic 2	Topic 3	Topic 4	Topic 1	Topic 2	Topic 3	Topic 4
image	kernel	method	inference	base	optimization	feature	process
feature	base	estimation	bayesian	adaptive	sparse	function	gaussian
optimization	image	sparse	process	process	recognition	base	via
application	process	non	classification	kernel	classification	datum	datum
recognition	dynamic	stochastic	kernel	method	method	information	inference
fast	estimation	optimization	multi	linear	base	adaptive	optimization
clustering	classification	application	analysis	datum	adaptive	reinforcement	time
via	deep	gradient	linear	system	linear	search	recognition
representation	optimal	multi	fast	sample	function	gaussian	base
reinforcement	sample	fast	application	dynamic	regression	bound	latent
PMI : 3.47	PMI : 3.40	PMI : 3.20	PMI : 3.19	PMI : 3.09	PMI : 3.00	PMI : 2.95	PMI : 2.91

of words. The PMI between two words w_i and w_j is given by $\text{PMI}(w_i, w_j) = \log \frac{p(w_i, w_j)}{p(w_i)p(w_j)}$, where $p(w_i)$ is the probability that word w_i occurs in WIKIPEDIA, and $p(w_i, w_j)$ is the probability that words w_i and w_j co-occur in a 5-word window in any WIKIPEDIA document. Higher PMIs indicate higher semantic coherence.

Table 5 reports the results by each model. We report the best 4 topics in terms of PMI and visualize their top-10 words. The words w_i for topic z_k are selected with the highest parameters θ_{ik} , which corresponds to the highest value $p(x_i = 1 | z_k)$. We also report the average pairwise PMI between the top words within each topic, and the mean of the average pairwise PMI across all topics. With more training data ($N_{train} = 5241$), ACP and AVI achieve similar average PMI. When N_{train} is reduced to 2000, all four selected topics in ACP have better average pairwise PMI scores than AVI. In Supplementary Materials, we report additional

experimental results of document classification and latent representations of the learnt topics.

5 CONCLUSION

We proposed ACP for variational inference, which combines the classical techniques of deriving variational bounds over likelihoods and recent approaches using neural networks for amortized variational inference. We showed that by constraining the form of approximate posterior using classical methods and learning the variational parameters to maximize ELBO, ACP can generalize well even with a small amount of training data. While with large amount of training data, our approach retains the high capacity for better generative modeling. Empirical studies have shown the advantages of ACP. Our future direction is to extend this approach to hierarchical models.

Acknowledgements

Fei Sha is on leave from University of Southern California. We appreciate the feedback from the reviewers. This work is partially supported by NSF Awards IIS-1513966/ 1632803/1833137, CCF-1139148, DARPA Award#: FA8750-18-2-0117, DARPA-D3M - Award UCB-00009528, Google Research Awards, gifts from Facebook and Netflix, and ARO# W911NF-12-1-0241 and W911NF-15-1-0484.

References

- Neuralips titles 1987–2016. <https://www.kaggle.com/benhammer/nips-papers>, 2019.
- Yonatan Halpern. *Semi-Supervised Learning for Electronic Phenotyping in Support of Precision Medicine*. PhD thesis, New York University, 2016.
- Yonatan Halpern and David Sontag. Unsupervised learning of noisy-or bayesian networks. *arXiv preprint arXiv:1309.6834*, 2013.
- Tommi S Jaakkola and Michael I Jordan. Variational probabilistic inference and the qmr-dt network. *Journal of artificial intelligence research*, 10:291–322, 1999.
- Tommi S Jaakkola and Michael I Jordan. Bayesian parameter estimation via variational methods. *Statistics and Computing*, 10(1):25–37, 2000.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Tony Jebara and Alex Pentland. On reversing jensen’s inequality. In *Advances in Neural Information Processing Systems*, pages 231–237, 2001.
- Yacine Jernite, Yonatan Halpern, and David Sontag. Discovering hidden variables in noisy-or networks using quartet tests. In *Advances in Neural Information Processing Systems*, pages 2355–2363, 2013.
- Geng Ji, Dehua Cheng, Huazhong Ning, Changhe Yuan, Hanning Zhou, Liang Xiong, and Erik B Sudderth. Variational training for large-scale noisy-or bayesian networks. 2019.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Daphne Koller, Nir Friedman, and Francis Bach. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- Andriy Mnih and Karol Gregor. Neural variational inference and learning in belief networks. *arXiv preprint arXiv:1402.0030*, 2014.
- Quaid Morris. Recognition networks for approximate inference in bn20 networks. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pages 370–377. Morgan Kaufmann Publishers Inc., 2001.
- David Newman, Sarvnaz Karimi, and Lawrence Cavdon. External evaluation of topic models. In *in Australasian Doc. Comp. Symp.*, 2009. Citeseer, 2009.
- Lawrence K Saul and Michael I Jordan. Exploiting tractable substructures in intractable networks. In *Advances in neural information processing systems*, pages 486–492, 1996.
- Lawrence K Saul, Tommi Jaakkola, and Michael I Jordan. Mean field theory for sigmoid belief networks. *Journal of artificial intelligence research*, 4:61–76, 1996.
- Tomáš Šingliar and Miloš Hauskrecht. Noisy-or component analysis and its application to link analysis. *Journal of Machine Learning Research*, 7(Oct):2189–2213, 2006.