

Auto-Encoding Variational Neural Machine Translation

Bryan Eikema & Wilker Aziz

Institute for Logic, Language and Computation

University of Amsterdam

b.eikema@uva.nl, w.aziz@uva.nl

Abstract

We present a deep generative model of bilingual sentence pairs for machine translation. The model generates source and target sentences jointly from a shared latent representation and is parameterised by neural networks. We perform efficient training using amortised variational inference and reparameterised gradients. Additionally, we discuss the statistical implications of joint modelling and propose an efficient approximation to maximum a posteriori decoding for fast test-time predictions. We demonstrate the effectiveness of our model in three machine translation scenarios: in-domain training, mixed-domain training, and learning from a mix of gold-standard and synthetic data. Our experiments show consistently that our joint formulation outperforms conditional modelling (i.e. standard neural machine translation) in all such scenarios.

1 Introduction

Neural machine translation (NMT) systems (Kalchbrenner and Blunsom, 2013; Sutskever et al., 2014; Cho et al., 2014b) require vast amounts of labelled data, i.e. bilingual sentence pairs, to be trained effectively. Oftentimes, the data we use to train these systems are a byproduct of mixing different sources of data. For example, labelled data are sometimes obtained by putting together corpora from different domains (Sennrich et al., 2017). Even for a single domain, parallel data often result from the combination of documents independently translated from different languages by different people or agencies, possibly following different guidelines. When resources are scarce, it is not uncommon to mix in some synthetic data, e.g. bilingual data artificially obtained by having a model translate target monolingual data to the source language (Sennrich et al., 2016a). Translation direction,

original language, and quality of translation are some of the many factors that we typically choose not to control for (due to lack of information or simply for convenience).¹ All those arguably contribute to making our labelled data a mixture of samples from various data distributions.

Regular NMT systems do not explicitly account for latent factors of variation, instead, given a source sentence, NMT models a single conditional distribution over target sentences as a fully supervised problem. In this work, we introduce a deep generative model that generates source and target sentences jointly from a shared latent representation. The model has the potential to use the latent representation to capture global aspects of the observations, such as some of the latent factors of variation just discussed. The result is a model that accommodates members of a more complex class of marginal distributions. Due to the presence of latent variables, this model requires posterior inference, in particular, we employ the framework of amortised variational inference (Kingma and Welling, 2014). Additionally, we propose an efficient approximation to maximum a posteriori (MAP) decoding for fast test-time predictions.

Contributions We introduce a deep generative model for NMT (§3) and discuss theoretical advantages of joint modelling over conditional modelling (§3.1). We also derive an efficient approximation to MAP decoding that requires only a single forward pass through the network for prediction (§3.3). Finally, we show in §4 that our proposed model improves translation performance in at least three practical scenarios: i) in-domain

¹Also note that this list is by no means exhaustive. For example, Rabinovich et al. (2017) show influence of factors such as personal traits and demographics in translation. Another clear case is presented by Johnson et al. (2017), who combine parallel resources for multiple languages to train a single encoder-decoder architecture.

training on little data, where test data are expected to follow the training data distribution closely; ii) mixed-domain training, where we train a single model but test independently on each domain; and iii) learning from large noisy synthetic data.

2 Neural Machine Translation

In machine translation our observations are pairs of random sequences, a source sentence $x = \langle x_1, \dots, x_m \rangle$ and a target sentence $y = \langle y_1, \dots, y_n \rangle$, whose lengths m and n we denote by $|x|$ and $|y|$, respectively. In NMT, the likelihood of the target given the source

$$P(y|x, \theta) = \prod_{j=1}^{|y|} \text{Cat}(y_j | f_\theta(x, y_{<j})) \quad (1)$$

factorises without Markov assumptions (Sutskever et al., 2014; Bahdanau et al., 2015; Cho et al., 2014a). We have a fixed parameterised function f_θ , i.e. a neural network architecture, compute categorical parameters for varying inputs, namely, the source sentence and target prefix (denoted $y_{<j}$).

Given a dataset \mathcal{D} of i.i.d. observations, the parameters θ of the model are point-estimated to attain a local maximum of the log-likelihood function, $\mathcal{L}(\theta|\mathcal{D}) = \sum_{(x,y) \in \mathcal{D}} \log P(y|x, \theta)$, via stochastic gradient-based optimisation (Robbins and Monro, 1951; Bottou and Cun, 2004).

Predictions For a trained model, predictions are performed by searching for the target sentence y that maximises the conditional $P(y|x)$, or equivalently its logarithm, with a greedy algorithm

$$\arg \max_y P(y|x, \theta) \approx \text{greedy} \log P(y|x, \theta) \quad (2)$$

such as beam-search (Sutskever et al., 2014), possibly aided by a manually tuned length penalty. This *decision rule* is often referred to as MAP decoding (Smith, 2011).

3 Auto-Encoding Variational NMT

To account for a latent space where global features of observations can be captured, we introduce a random sentence embedding $z \in \mathbb{R}^d$ and model the joint distribution over observations as a marginal of $p(z, x, y|\theta)$.² That is, $(x, y) \in \mathcal{D}$ is assumed to be sampled from the distribution

$$P(x, y|\theta) = \int p(z)P(x, y|z, \theta)dz. \quad (3)$$

²We use uppercase $P(\cdot)$ for probability mass functions and lowercase $p(\cdot)$ for probability density functions.

where we impose a standard Gaussian prior on the latent variable, i.e. $Z \sim \mathcal{N}(0, I)$, and assume $X \perp Y|Z$. That is, given a sentence embedding z , we first generate the source conditioned on z ,

$$P(x|z, \theta) = \prod_{i=1}^{|x|} \text{Cat}(x_i | g_\theta(z, x_{<i})), \quad (4)$$

then generate the target conditioned on x and z ,

$$P(y|x, z, \theta) = \prod_{j=1}^{|y|} \text{Cat}(y_j | f_\theta(z, x, y_{<j})). \quad (5)$$

Note that the source sentence is generated without Markov assumptions by drawing one word at a time from a categorical distribution parameterised by a recurrent neural network g_θ . The target sentence is generated similarly by drawing target words in context from a categorical distribution parameterised by a sequence-to-sequence architecture f_θ . This essentially combines a neural language model (Mikolov et al., 2010) and a neural translation model (§2), each extended to condition on an additional stochastic input, namely, z .

3.1 Statistical considerations

Modelling the conditional directly, as in standard NMT, corresponds to the statistical assumption that the *distribution* over source sentences can provide no information about the distribution over target sentences given a source. That is, conditional NMT assumes independence of β determining $P(y|x, \beta)$ and α determining $P(x|\alpha)$. Scenarios where this assumption is unlikely to hold are common: where x is noisy (e.g. synthetic or crowdsourced), poor quality x should be assigned low probability $P(x|\alpha)$ which in turn should inform the conditional. Implications of this assumption extend to parameter estimation: updates to the conditional are not sensitive to how exotic x is.

Let us be more explicit about how we parameterise our model by identifying 3 sets of parameters $\theta = \{\theta_{\text{emb-x}}, \theta_{\text{LM}}, \theta_{\text{TM}}\}$, where $\theta_{\text{emb-x}}$ parameterises an embedding layer for the source language. The embedding layer is shared between the two model components

$$P(x, y|z, \theta) = P(x|z, \underbrace{\theta_{\text{emb-x}}, \theta_{\text{LM}}}_{\alpha}) P(y|x, z, \underbrace{\theta_{\text{emb-x}}, \theta_{\text{TM}}}_{\beta}) \quad (6)$$

and it is then clear by inspection that $\alpha \cap \beta = \{z, \theta_{\text{emb-x}}\}$. In words, we break the independence

assumption in two ways, namely, by having the two distributions share parameters and by having them depend on a shared latent sentence representation z . Note that while the embedding layer is deterministic and global to all sentence pairs in the training data, the latent representation is stochastic and local to each sentence pair.

Now let us turn to considerations about latent variable modelling. Consider a model $P(x|\theta_{\text{emb-x}}, \theta_{\text{LM}})P(y|x, \theta_{\text{emb-x}}, \theta_{\text{TM}})$ of the joint distribution over observations that does not employ latent variables. This alternative, which we discuss further in experiments, models each component directly, whereas our proposed model (3) requires marginalisation of latent embeddings z . Marginalisation turns our directed graphical model into an undirected one inducing further structure in the marginal. See Appendix B, and Figure 2 in particular, for an extended discussion.

3.2 Parameter estimation

The marginal in Equation (3) is clearly intractable, thus precluding maximum likelihood estimation. Instead, we resort to variational inference (Jordan et al., 1999; Blei et al., 2017) and introduce a variational approximation $q(z|x, y, \lambda)$ to the intractable posterior $p(z|x, y, \theta)$. We let the approximate posterior be a diagonal Gaussian

$$\begin{aligned} Z|\lambda, x, y &\sim \mathcal{N}(\mathbf{u}, \text{diag}(\mathbf{s} \odot \mathbf{s})) \\ \mathbf{u} &= \mu_\lambda(x, y) \\ \mathbf{s} &= \sigma_\lambda(x, y) \end{aligned} \quad (7)$$

and predict its parameters (i.e. $\mathbf{u} \in \mathbb{R}^d, \mathbf{s} \in \mathbb{R}_{>0}^d$) with neural networks whose parameters we denote by λ . This makes the model an instance of a variational auto-encoder (Kingma and Welling, 2014). See Figure 1 in Appendix B for a graphical depiction of the generative and inference models.

We can then jointly estimate the parameters of both models (generative θ and inference λ) by maximising the ELBO (Jordan et al., 1999), a lowerbound on the marginal log-likelihood,

$$\begin{aligned} \log P(x, y|\theta) &\geq \mathcal{E}(\theta, \lambda|x, y) = \\ &\mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} [\log P(x, y|z = \mathbf{u} + \epsilon \odot \mathbf{s}, \theta)] \\ &- \text{KL}(\mathcal{N}(z|\mathbf{u}, \text{diag}(\mathbf{s} \odot \mathbf{s}))||\mathcal{N}(z|0, I)), \end{aligned} \quad (8)$$

where we have expressed the expectation with respect to a fixed distribution—a reparameterisation available to location-scale families such as the Gaussian (Kingma and Welling, 2014; Rezende

et al., 2014). Due to this reparameterisation, we can compute a Monte Carlo estimate of the gradient of the first term via back-propagation (Rumelhart et al., 1986; Schulman et al., 2015). The KL term, on the other hand, is available in closed form (Kingma and Welling, 2014, Appendix B).

3.3 Predictions

In a latent variable model, MAP decoding (9a) requires searching for y that maximises the marginal $P(y|x, \theta) \propto P(x, y|\theta)$, or equivalently its logarithm. In addition to approximating exact search with a greedy algorithm, other approximations are necessary in order to achieve fast prediction. First, rather than searching through the true marginal, we search through the evidence lowerbound. Second, we replace the approximate posterior $q(z|x, y)$ by an auxiliary distribution $r(z|x)$. As we are searching through the space of target sentences, not conditioning on y circumvents combinatorial explosion and allows us to drop terms that depend on x alone (9b). Finally, instead of approximating the expectation via MC sampling, we condition on the expected latent representation and search greedily (9c).

$$\arg \max_y \log P(y|x) \quad (9a)$$

$$\approx \arg \max_y \mathbb{E}_{r(z|x)} [\log P(y|z, x)] \quad (9b)$$

$$\approx \text{greedy} \log P(y|\mathbb{E}_{r(z|x)}[z], x) \quad (9c)$$

Together, these approximations enable prediction with a single call to an arg max solver, in our case a standard greedy search algorithm, which leads to prediction times that are very close to that of the conditional model. This strategy, and (9b) in particular, suggests that a good auxiliary distribution $r(z|x)$ should approximate $q(z|x, y)$ closely.

We parameterise this *prediction model* using a neural network and investigate different options to estimate its parameters. As a first option, we restrict the approximate posterior to conditioning on x alone, i.e. we approach posterior inference with $q_\lambda(z|x)$ rather than $q_\lambda(z|x, y)$, and thus, we can use $r(z|x) = q_\lambda(z|x)$ for prediction.³ As a second option, we make $r_\phi(z|x)$ a diagonal Gaussian and estimate parameters ϕ to make $r_\phi(z|x)$ close to the approximate posterior $q_\lambda(z|x, y)$ as measured by

³Note that this does not stand in contrast to our motivation for joint modelling, as we still tie source and target through z in the generative model, but it does limit the context available for posterior inference.

$D(r_\phi, q_\lambda)$. For as long as $D(r_\phi, q_\lambda) \in \mathbb{R}_{\geq 0}$ for every choice of ϕ and λ , we can estimate ϕ jointly with θ and λ by maximising a modified ELBO

$$\log P(x, y | \theta) \geq \mathcal{E}(\theta, \lambda | x, y) - D(r_\phi, q_\lambda) \quad (10)$$

which is loosened by the gap between r_ϕ and q_λ . In experiments we investigate a few options for $D(r_\phi, q_\lambda)$, all available in closed form for Gaussians, such as $\text{KL}(r_\phi || q_\lambda)$, $\text{KL}(q_\lambda || r_\phi)$, as well as the Jensen-Shannon (JS) divergence.

Note that r_ϕ is used only for prediction as a decoding *heuristic* and as such need not be stochastic. We can, for example, design $r_\phi(x)$ to be a point estimate of the posterior mean and optimise

$$\mathcal{E}(\theta, \lambda | x, y) - \|r_\phi(x) - \mathbb{E}_{q_\lambda(z|x,y)}[z]\|_2^2 \quad (11)$$

which remains a lowerbound on log-likelihood.

4 Experiments

We investigate two translation tasks, namely, WMT’s translation of news (Bojar et al., 2016) and IWSLT’s translation of transcripts of TED talks (Cettolo et al., 2014), and concentrate on translations for German (DE) and English (EN) in either direction. In this section we aim to investigate scenarios where we expect observations to be representative of various data distributions. As a sanity check, we start where training conditions can be considered in-domain with respect to test conditions. Though note that this does not preclude the potential for appreciable variability in observations as various other latent factors still likely play a role (see §1). We then mix datasets from these two remarkably different translation tasks and investigate whether performance can be improved across tasks with a single model. Finally, we investigate the case where we learn from synthetic data in addition to gold-standard data. For this investigation we derive synthetic data from observations that are close to the domain of the test set in an attempt to avoid further confounders.

Data For bilingual data we use News Commentary (NC) v12 (Bojar et al., 2017) and IWSLT 2014 (Cettolo et al., 2014), where we assume NC to be representative of the test domain of the WMT News task. The datasets consist of 255,591 training sentences and 153,326 training sentences respectively. In experiments with synthetic data, we subsample 10^6 sentences from the News Crawl 2016 articles (Bojar et al., 2017) for either German

or English depending on the target language. For the WMT task, we concatenate `newstest2014` and `newstest2015` for validation/development (5,172 sentence pairs) and report test results on `newstest2016` (2,999 sentence pairs). For IWSLT, we use the split proposed by Ranzato et al. (2016) who separated 6,969 training instances for validation/development and reported test results on a concatenation of `dev2010`, `dev2012` and `tst2010-2012` (6,750 sentence pairs).

Pre-processing We tokenized and truecased all data using standard scripts from the Moses toolkit (Koehn et al., 2007), and removed sentences longer than 50 tokens. For computational efficiency and to avoid problems with closed vocabularies, we segment the data using BPE (Sennrich et al., 2016b) with 32,000 merge operations independently for each language. For training the truecaser and the BPEs we used a concatenation of all the available bilingual and monolingual data for German and all bilingual data for English.

Systems We develop all of our models on top of Tensorflow NMT (Luong et al., 2017). Our baseline system is a standard implementation of conditional NMT (COND) (Bahdanau et al., 2015). To illustrate the importance of latent variable modelling, we also include in the comparison a simpler attempt at JOINT modelling where we do not induce a shared latent space. Instead, the model is trained in a fully-supervised manner to maximise what is essentially a combination of two nearly independent objectives,

$$\begin{aligned} \mathcal{L}(\theta | \mathcal{D}) = & \sum_{(x,y) \in \mathcal{D}} \sum_{i=1}^{|x|} \log P(x_i | x_{<i}, \theta_{\text{emb-x}}, \theta_{\text{LM}}) \\ & + \sum_{j=1}^{|y|} \log P(y_j | x, y_{<j}, \theta_{\text{emb-x}}, \theta_{\text{TM}}), \quad (12) \end{aligned}$$

namely, a language model and a conditional translation model. Note that the two components of the model share very little, i.e. an embedding layer for the source language. Finally, we aim at investigating the effectiveness of our auto-encoding variational NMT (AEVNMT).⁴ Appendix A contains a detailed description of the architectures that parameterise our systems.⁵

⁴Code available from github.com/Roxot/AEVNMT.

⁵In comparison to COND, AEVNMT requires additional components: a source language model, an inference network,

	NC	IWSLT
Dropout	30%	30%
Word dropout rate	10%	20%
KL annealing steps	80,000	80,000
KL($q(z) p(z)$) on EN-DE	5.94	8.01

Table 1: Strategies to promote use of latent representation along with the validation KL achieved.

Hyperparameters Our recurrent cells are 256-dimensional GRU units (Cho et al., 2014b). We train on batches of 64 sentence pairs with Adam (Kingma and Ba, 2015), learning rate 3×10^{-4} , for at least T updates. We then perform convergence checks every 500 batches and stop after 20 checks without any improvement measured by BLEU (Papineni et al., 2002). For in-domain training we set $T = 140,000$, and for mixed-domain training, as well as training with synthetic data, we set $T = 280,000$. For decoding we use a beam width of 10 and a length penalty of 1.0. We investigate the use of dropout (Srivastava et al., 2014) for the conditional baseline with rates from 10% to 60% in increments of 10%. Best validation performance on WMT required a rate of 40% for EN-DE and 50% for DE-EN, while on IWSLT it required 50% for either translation direction. To spare resources, we also use these rates for training the simple JOINT model.

Avoiding collapsing to prior Many have noticed that VAEs whose observation models are parameterised by *strong generators*, such as recurrent neural networks, learn to ignore the latent representation (Bowman et al., 2016; Higgins et al., 2017; Sønderby et al., 2016; Alemi et al., 2018). In such cases, the approximate posterior “collapses” to the prior, and where one has a fixed prior, such as our standard Gaussian, this means that the posterior becomes independent of the data, which is obviously not desirable. Bowman et al. (2016) proposed two techniques to counter this effect, namely, “KL annealing”, and target word dropout. KL annealing consists in incorporating the KL term of Equation (8) into the objective gradually, thus allowing the posterior to move away from the prior more freely at early stages of training. After

and possibly a prediction network. However, this does not add much sequential computation: the inference network can run in parallel with the source encoder, and the source language model runs in parallel with the target decoder.

Objective	BLEU \uparrow
ELBO _{x,y} - KL($r_\phi(z x) q_\lambda(z x,y)$)	14.7
ELBO _{x,y} - KL($q_\lambda(z x,y) r_\phi(z x)$)	14.8
ELBO _{x,y} - JS($r_\phi(z x) q_\lambda(z x,y)$)	14.9
ELBO _{x,y} - $\ r_\phi(x) - \mathbb{E}_{q_\lambda(z x,y)}[Z]\ _2^2$	14.8
ELBO _x	14.9

Table 2: EN-DE validation results for NC training. ELBO_x means we condition on the source alone for posterior inference, i.e. the variational approximation $q_\lambda(z|x)$ is used for training and for predictions. In all other cases, we condition on both observations for training, i.e. $q_\lambda(z|x,y)$, and train either a distribution $r_\phi(z|x)$ or a point estimate $r_\phi(x)$ for predictions.

a number of annealing steps, the KL term is incorporated in full and training continues with the actual ELBO. In our search we considered annealing for 20,000 to 80,000 training steps. Word dropout consists in randomly masking words in observed target prefixes at a given rate. The idea is to harm the potential of the decoder to capitalise on correlations internal to the structure of the observation in the hope that it will rely more on the latent representation instead. We considered rates from 20% to 40% in increments of 10%. Table 1 shows the configurations that achieve best validation results on EN-DE. To spare resources, we reuse these hyperparameters for DE-EN experiments. With these settings, we attain a non-negligible validation KL (see, last row of Table 1), which indicates that the approximate posterior is different from the prior at the end of training.

ELBO variants We investigate the effect of conditioning on target observations for posterior inference during training against a simpler variant that conditions on the source alone. Table 2 suggests that conditioning on x is sufficient and thus we opt to continue with this simpler version. Do note that when we use both observations for posterior inference, i.e. $q_\lambda(z|x,y)$, and thus train an approximation r_ϕ for prediction, we have additional parameters to estimate (e.g. due to the need to encode y for q_λ and x for r_ϕ), thus it may be the case that for these variants to show their potential we need larger data and/or prolonged training.

4.1 Results

In this section we report test results in terms of BLEU (Papineni et al., 2002) and BEER (Stanojević and Sima’an, 2014), but in Appendix E

Task	Model	EN-DE		DE-EN	
		BLEU \uparrow	BEER \uparrow	BLEU \uparrow	BEER \uparrow
IWSLT14	COND	23.0 (0.1)	58.6 (0.1)	27.3 (0.2)	59.8 (0.1)
	JOINT	23.2	58.7	27.5	59.8
	AEVNMT	23.4 (0.1)	58.8 (0.1)	28.0 (0.1)	60.1 (0.1)
WMT16	COND	17.8 (0.2)	53.1 (0.1)	20.1 (0.1)	53.7 (0.1)
	JOINT	17.9	53.4	20.1	53.7
	AEVNMT	18.4 (0.2)	53.5 (0.1)	20.6 (0.2)	53.6 (0.1)

Table 3: Test results for in-domain training on IWSLT (top) and NC (bottom): we report average (1std) across 5 independent runs for COND and AEVNMT, but a single run of JOINT.

we additionally report METEOR (Denkowski and Lavie, 2011) and TER (Snover et al., 2006). We de-truecase and de-tokenize our system’s predictions and compute BLEU scores using SacreBLEU (Post, 2018).⁶ For BEER, METEOR and TER, we tokenize the results and test sets using the same tokenizer as used by SacreBLEU. We make use of BEER 2.0, and for METEOR and TER use MULTEVAL (Clark et al., 2011). In Appendix D we report validation results, in this case in terms of BLEU alone as that is what we used for model selection. Finally, to give an indication of the degree to which results are sensitive to initial conditions (e.g. random initialisation of parameters), and to avoid possibly misleading significance testing, we report the average and standard deviation of 5 independently trained models. To spare resources we do not report multiple runs for JOINT, but our experience is that its performance varies similarly to that of the conditional baseline.

We start with the case where we can reasonably assume training data to be in-domain with respect to test data. Table 3 shows in-domain training performance. First, we remark that our conditional baseline for the IWSLT14 task (IWSLT training) is very close to an external baseline trained on the same data (Bahdanau et al., 2017).⁷ The results on IWSLT show benefits from joint modelling and in particular from learning a shared latent space. For the WMT16 task (NC training), BLEU shows a similar trend, namely, joint modelling with a shared latent space (AEVNMT) outperforms both conditional modelling and the simple joint model.

⁶Version string: BLEU+case.mixed+numrefs.1+smooth.exp+tok.13a+version.1.2.12

⁷Bahdanau et al. (2017) report 27.56 on the same test set for DE-EN, though note that they train on words rather than BPEs and use a different implementation of BLEU.

We now consider the scenario where we know for a fact that observations come from two different data distributions, which we realise by training our models on a concatenation of IWSLT and NC. In this case, we perform model selection once on the concatenation of both development sets and evaluate the same model on each domain separately. We can see in Table 4 that conditional modelling is never preferred, JOINT performs reasonably well, especially for DE-EN, and that in every comparison our AEVNMT outperforms the conditional baseline both in terms of BLEU and BEER.

Another common scenario where two very distinct data distributions are mixed is when we capitalise on the abundance of monolingual data and train on a concatenation of gold-standard bilingual data (we use NC) and synthetic bilingual data derived from target monolingual corpora via back-translation (Sennrich et al., 2016a) (we use News Crawl). In such a scenario the latent variable might be able to inform the translation model of the amount of noise present in the source sentence. Table 5 shows results for both baselines and AEVNMT. First, note that synthetic data greatly improves the conditional baseline, in particular translating into English. Once again AEVNMT consistently outperforms conditional modelling and joint modelling without latent variables.

By mixing different sources of data we are trying to diagnose whether the generative model we propose is robust to unknown and diverse sources of variation mixed together in one training set (e.g. NC + IWSLT or gold-standard + synthetic data). However, note that a point we are certainly not trying to make is that the model has been designed to perform *domain adaptation*. Nonetheless, in Appendix C we try to shed light on what happens when we use the model to translate genres

Training	Model	WMT16		IWSLT14	
		BLEU \uparrow	BEER \uparrow	BLEU \uparrow	BEER \uparrow
EN-DE	COND	17.6 (0.4)	53.9 (0.2)	23.9 (0.3)	59.3 (0.1)
	JOINT	18.1	54.3	24.2	59.5
	AEVNMT	18.4 (0.2)	54.5 (0.2)	24.1 (0.3)	59.5 (0.2)
DE-EN	COND	21.6 (0.2)	55.5 (0.2)	29.1 (0.2)	60.9 (0.1)
	JOINT	22.3	55.6	29.2	61.2
	AEVNMT	22.3 (0.1)	55.6 (0.1)	29.2 (0.1)	61.1 (0.1)

Table 4: Test results for mixed-domain training: we report average (1std) across 5 independent runs for COND and AEVNMT, but a single run of JOINT.

it has never seen. On a dataset covering various unseen genres, we observe that both COND and AEVNMT perform considerably worse showing that without taking domain adaptation seriously both models are inadequate. In terms of BLEU, differences range from -0.3 to 0.8 (EN-DE) and 0.3 to 0.7 (DE-EN) and are mostly in favour of AEVNMT (17/20 comparisons).

Remarks It is intuitive to expect latent variable modelling to be most useful in settings containing high variability in the data, i.e. mixed-domain and synthetic data settings, though in our experiments AEVNMT shows larger improvements in the in-domain setting. We speculate two reasons for this: i) it is conceivable that variation in the mixed-domain and synthetic data settings are too large to be well accounted by a diagonal Gaussian; and ii) the benefits of latent variable modelling may diminish as the amount of available data grows.

4.2 Probing latent space

To investigate what information the latent space encodes we explore the idea of training simple *linear probes* or *diagnostic classifiers* (Alain and Bengio, 2017; Hupkes et al., 2018). With simple Bayesian logistic regression we have managed to predict from $Z \sim q(z|x)$ domain indicators (i.e. newswire vs transcripts) and gold-standard vs synthetic data at performance above 90% accuracy on development set. However, a similar performance is achieved from the deterministic average state of the bidirectional encoder of the conditional baseline. We have also been able to predict from $Z \sim q(z|x)$ the level of noise in back-translated data measured on the development set at the sentence level by an automatic metric, i.e. METEOR, with performance above what can be done with

random features. Though again, the performance is not much better than what can be done with a conditional baseline. Still, it is worth highlighting that these aspects are rather coarse, and it is possible that the performance gains we report in §4.1 are due to far more nuanced variations in the data. At this point, however, we do not have a good qualitative assessment of this conjecture.

5 Related Work

Joint modelling In similar work, Shah and Barber (2018) propose a joint generative model whose probabilistic formulation is essentially identical to ours. Besides some small differences in architecture, our work differs in two regards: motivation and strategy for predictions. Their goal is to jointly learn from multiple language pairs by sharing a single polyglot architecture (Johnson et al., 2017). Their strategy for prediction is based on a form of stochastic hill-climbing, where they sample an initial z from the standard Gaussian prior and decode via beam search in order to obtain a draft translation $\tilde{y} = \text{greedy}_y P(y|z, x)$. This translation is then iteratively refined by encoding the pair $\langle x, \tilde{y} \rangle$, re-sampling z , though this time from $q(z|x, \tilde{y})$, and re-decoding with beam search. Unlike our approach, this requires multiple calls to the inference network and to beam search. Moreover, the inference model, which is trained on gold-standard observations, is used on noisy target sentences.

Cotterell and Kreutzer (2018) interpret back-translation as a single iteration of a wake-sleep algorithm (Hinton et al., 1995) for a joint model of bitext $P(x, y|\theta) = P(y|x, \theta)P_*(x)$. They sample directly from the data distribution $P_*(x)$ and learn two NMT models, a generative $P(y|x, \theta)$ and an auxiliary model $Q(x|y, \phi)$, each trained

WMT16	EN-DE		DE-EN	
	BLEU \uparrow	BEER \uparrow	BLEU \uparrow	BEER \uparrow
COND	17.8 (0.2)	53.1 (0.1)	20.1 (0.1)	53.7 (0.1)
+ synthetic data	22.3 (0.3)	57.0 (0.2)	26.9 (0.2)	58.5 (0.1)
JOINT + synthetic data	22.2	57.0	26.7	58.6
AEVNMT + synthetic data	22.5 (0.2)	57.0 (0.1)	27.4 (0.2)	58.8 (0.1)

Table 5: Test results for training on NC plus synthetic data (back-translated News Crawl): we report average (1std) across 5 independent runs for COND and AEVNMT, but a single run of JOINT.

on a separate objective. Zhang et al. (2018) propose a joint model of bitext trained to incorporate the back-translation heuristic as a trainable component in a formulation similar to that of Cotterell and Kreutzer (2018). In both cases, joint modelling is done without a shared latent space and without a source language model.

Multi-task learning An alternative to joint learning is to turn to multi-task learning and explore parameter sharing across models trained on different, though related, data with different objectives. For example, Cheng et al. (2016) incorporate both source and target monolingual data by multi-tasking with a non-differentiable auto-encoding objective. They jointly train a source-to-target and target-to-source system that act as encoder and decoder respectively. Zhang and Zong (2016) combine a source language model objective with a source-to-target conditional NMT objective and shared the source encoder in a multi-task learning fashion.

Variational LMs and NMT Bowman et al. (2016) first proposed to augment a neural language model with a prior over latent space. Our source component is an instance of their model. More recently, Xu and Durrett (2018) proposed to use a hyperspherical uniform prior rather than a Gaussian and showed the former leads to better representations. Zhang et al. (2016) proposed the first VAE for NMT. They augment the conditional with a Gaussian sentence embedding and model observations as draws from the marginal $P(y|x, \theta) = \int p(z|x, \theta)P(y|x, z, \theta)dz$. Their formulation is a conditional deep generative model (Sohn et al., 2015) that does not model the source side of the data, where, rather than a fixed standard Gaussian, the latent model is itself parameterised and depends on the data. Schulz et al. (2018) extend the model of Zhang et al. (2016) with a

Markov chain of latent variables, one per timestep, allowing the model to capture greater variability.

Latent domains In the context of statistical MT, Cuong and Sima’an (2015) estimate a joint distribution over sentence pairs while marginalising discrete latent domain indicators. Their model factorises over word alignments and is not used directly for translation, but rather to improve word and phrase alignments, or to perform data selection (Hoang and Sima’an, 2014), prior to training. There is a vast literature on domain adaptation for statistical machine translation (Cuong and Sima’an, 2017), as well as for NMT (Chu and Wang, 2018), but a full characterisation of this exciting field is beyond the scope of this paper.

6 Discussion and Future Work

We have presented a joint generative model of translation data that generates both observations conditioned on a shared latent representation. Our formulation leads to questions such as *why joint learning?* and *why latent variable modelling?* to which we give an answer based on statistical facts about conditional modelling and marginalisation as well as empirical evidence of improved performance. Our model shows moderate but consistent improvements across various settings and over multiple independent runs.

In future work, we shall investigate datasets annotated with demographics and personal traits in an attempt to assess how far we can go in capturing fine grained variation. Though note that if such factors of variation vary widely in distribution, it may be naïve to expect we can model them well with a simple Gaussian prior. If that turns out to be the case, we will investigate mixing Gaussian components (Miao et al., 2016; Srivastava and Sutton, 2017) and/or employing a hierarchical prior (Goyal et al., 2017).

Acknowledgements



This project has received funding from the Dutch Organization for Scientific Research VICI Grant No 277-89-002 and from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825299 (GoURMET). We also thank Philip Schulz, Khalil Sima'an, and Joost Bastings for comments and helpful discussions. A Titan Xp card used for this research was donated by the NVIDIA Corporation.

References

- G. Alain and Y. Bengio. 2017. Understanding intermediate layers using linear classifier probes. In *ICLR, 2017*, Toulon, France.
- A. Alemi, B. Poole, I. Fischer, J. Dillon, R. A. Saurous, and K. Murphy. 2018. [Fixing a broken ELBO](#). In *Proceedings of ICML, 2018*, pages 159–168, Stockholm, Sweden.
- D. Bahdanau, P. Brakel, K. Xu, A. Goyal, R. Lowe, J. Pineau, A. Courville, and Y. Bengio. 2017. An actor-critic algorithm for sequence prediction. In *ICLR, 2017*, Toulon, France.
- D. Bahdanau, K. Cho, and Y. Bengio. 2015. [Neural Machine Translation by Jointly Learning to Align and Translate](#). In *ICLR, 2015*, San Diego, USA.
- D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. 2017. Variational inference: A review for statisticians. *JASA*, 112(518):859–877.
- O. Bojar, R. Chatterjee, C. Federmann, Y. Graham, B. Haddow, S. Huang, M. Huck, P. Koehn, Q. Liu, V. Logacheva, C. Monz, M. Negri, M. Post, R. Rubino, L. Specia, and M. Turchi. 2017. [Findings of the 2017 conference on machine translation \(wmt17\)](#). In *Proceedings of WMT, 2017*, pages 169–214, Copenhagen, Denmark.
- O. Bojar, R. Chatterjee, C. Federmann, Y. Graham, B. Haddow, M. Huck, A. Jimeno Yepes, P. Koehn, V. Logacheva, C. Monz, M. Negri, A. Neveol, M. Neves, M. Popel, M. Post, R. Rubino, C. Scarton, L. Specia, M. Turchi, K. Verspoor, and M. Zampieri. 2016. [Findings of the 2016 conference on machine translation](#). In *Proceedings of WMT, 2016*, pages 131–198, Berlin, Germany.
- L. Bottou and Y. L. Cun. 2004. Large scale online learning. In S. Thrun, L. K. Saul, and B. Schölkopf, editors, *NIPS, 2004*, pages 217–224. Vancouver, Canada.
- S. R. Bowman, L. Vilnis, O. Vinyals, A. Dai, R. Jozefowicz, and S. Bengio. 2016. Generating sentences from a continuous space. In *Proceedings of CoNLL, 2016*, pages 10–21, Berlin, Germany.
- M. Cettolo, J. Niehues, S. Stüker, L. Bentivogli, and M. Federico. 2014. Report on the 11th iwslt evaluation campaign, iwslt 2014. In *Proceedings of IWSLT, 2014*, Lake Tahoe, USA.
- Y. Cheng, W. Xu, Z. He, W. He, H. Wu, M. Sun, and Y. Liu. 2016. [Semi-supervised learning for neural machine translation](#). In *Proceedings of ACL, 2016*, pages 1965–1974, Berlin, Germany.
- K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio. 2014a. [On the Properties of Neural Machine Translation: Encoder-Decoder Approaches](#). In *Proceedings of SSST, 2014*, pages 103–111, Doha, Qatar.
- K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. 2014b. [Learning phrase representations using rnn encoder–decoder for statistical machine translation](#). In *Proceedings of EMNLP, 2014*, pages 1724–1734, Doha, Qatar.
- C. Chu and R. Wang. 2018. [A survey of domain adaptation for neural machine translation](#). In *Proceedings of COLING, 2018*, pages 1304–1319, Santa Fe, USA.
- J. H. Clark, C. Dyer, A. Lavie, and N. A. Smith. 2011. [Better hypothesis testing for statistical machine translation: Controlling for optimizer instability](#). In *Proceedings of ACL, 2011*, pages 176–181, Portland, USA.
- R. Cotterell and J. Kreutzer. 2018. Explaining and generalizing back-translation through wake-sleep. *arXiv preprint arXiv:1806.04402*.
- H. Cuong and K. Sima'an. 2015. [Latent domain word alignment for heterogeneous corpora](#). In *Proceedings of NAACL-HLT, 2015*, pages 398–408, Denver, Colorado.
- H. Cuong and K. Sima'an. 2017. A survey of domain adaptation for statistical machine translation. *Machine Translation*, 31(4):187–224.
- H. Cuong, K. Sima'an, and I. Titov. 2016. [Adapting to all domains at once: Rewarding domain invariance in smt](#). *TACL*, 4:99–112.
- M. Denkowski and A. Lavie. 2011. [Meteor 1.3: Automatic metric for reliable optimization and evaluation of machine translation systems](#). In *Proceedings of WMT, 2011*, pages 85–91, Edinburgh, Scotland.
- P. Goyal, Z. Hu, X. Liang, C. Wang, E. P. Xing, and C. Mellon. 2017. Nonparametric variational auto-encoders for hierarchical representation learning. In *ICCV, 2017*, pages 5104–5112, Venice, Italy.

- I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. 2017. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *ICLR, 2017*, Toulon, France.
- G. E. Hinton, P. Dayan, B. J. Frey, and R. M. Neal. 1995. The "wake-sleep" algorithm for unsupervised neural networks. *Science*, 268(5214):1158–1161.
- C. Hoang and K. Sima'an. 2014. Latent domain translation models in mix-of-domains haystack. In *Proceedings of COLING, 2014*, pages 1928–1939, Dublin, Ireland.
- D. Hupkes, S. Veldhoen, and W. Zuidema. 2018. Visualisation and 'diagnostic classifiers' reveal how recurrent and recursive neural networks process hierarchical structure. *JAIR*, 61:907–926.
- M. Johnson, M. Schuster, Q. Le, M. Krikun, Y. Wu, Z. Chen, N. Thorat, F. a. Viégas, M. Wattenberg, G. Corrado, M. Hughes, and J. Dean. 2017. Google's multilingual neural machine translation system: Enabling zero-shot translation. *TACL*, 5:339–351.
- M. Jordan, Z. Ghahramani, T. Jaakkola, and L. Saul. 1999. An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233.
- N. Kalchbrenner and P. Blunsom. 2013. Recurrent continuous translation models. In *Proceedings of EMNLP, 2013*, pages 1700–1709, Seattle, USA.
- D. P. Kingma and J. Ba. 2015. Adam: A method for stochastic optimization. In *ICLR, 2015*, San Diego, USA.
- D. P. Kingma and M. Welling. 2014. Auto-encoding variational bayes. In *ICLR, 2014*, Banff, Canada.
- P. Koehn, H. Hoang, A. Birch, C. Callison-Burch, M. Federico, N. Bertoldi, B. Cowan, W. Shen, C. Moran, R. Zens, C. Dyer, O. Bojar, A. Constantin, and E. Herbst. 2007. Moses: open source toolkit for statistical machine translation. In *Proceedings of ACL, 2007*, pages 177–180, Prague, Czech Republic.
- D. Koller and N. Friedman. 2009. *Probabilistic Graphical Models*. MIT Press.
- A. Kucukelbir, D. Tran, R. Ranganath, A. Gelman, and D. M. Blei. 2017. Automatic differentiation variational inference. *JMLR*, 18(1):430–474.
- M. Luong, E. Brevdo, and R. Zhao. 2017. Neural machine translation (seq2seq) tutorial. <https://github.com/tensorflow/nmt>.
- Y. Miao, L. Yu, and P. Blunsom. 2016. Neural variational inference for text processing. In *ICML, 2016*, pages 1727–1736, New York, USA.
- T. Mikolov, M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur. 2010. Recurrent neural network based language model. In *ISCA, 2010*, Kyoto, Japan.
- V. Nair and G. E. Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of ICML, 2010*, Haifa, Israel.
- K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of ACL, 2002*, pages 311–318, Philadelphia, USA.
- M. Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of WMT, 2018*, pages 186–191, Brussels, Belgium.
- E. Rabinovich, R. N. Patel, S. Mirkin, L. Specia, and S. Wintner. 2017. Personalized machine translation: Preserving original author traits. In *Proceedings of EACL, 2017*, pages 1074–1084, Valencia, Spain.
- M. Ranzato, S. Chopra, M. Auli, and W. Zaremba. 2016. Sequence level training with recurrent neural networks. In *ICLR, 2016*, San Juan, Puerto Rico.
- D. J. Rezende, S. Mohamed, and D. Wierstra. 2014. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of ICML, 2014*, 2, pages 1278–1286, Beijing, China.
- H. Robbins and S. Monro. 1951. A stochastic approximation method. *Ann. Math. Statist.*, 22(3):400–407.
- D. E. Rumelhart, G. E. Hinton, and R. J. Williams. 1986. Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1. *Nature*, 323.
- J. Schulman, N. Heess, T. Weber, and P. Abbeel. 2015. Gradient estimation using stochastic computation graphs. In *NIPS, 2015*, pages 3528–3536, Montreal, Canada.
- P. Schulz, W. Aziz, and T. Cohn. 2018. A stochastic decoder for neural machine translation. In *Proceedings of ACL, 2018*, Melbourne, Australia.
- R. Sennrich, A. Birch, A. Currey, U. Germann, B. Haddow, K. Heafield, A. V. Miceli Barone, and P. Williams. 2017. The university of edinburgh's neural mt systems for wmt17. In *Proceedings of WMT, 2017*, pages 389–399, Copenhagen, Denmark.
- R. Sennrich, B. Haddow, and A. Birch. 2016a. Improving neural machine translation models with monolingual data. In *Proceedings of ACL, 2016*, pages 86–96, Berlin, Germany.
- R. Sennrich, B. Haddow, and A. Birch. 2016b. Neural machine translation of rare words with subword units. In *Proceedings of ACL, 2016*, pages 1715–1725, Berlin, Germany.

- H. Shah and D. Barber. 2018. [Generative neural machine translation](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *NIPS, 2018*, pages 1352–1361. Montreal, Canada.
- N. A. Smith. 2011. *Linguistic Structure Prediction*. Morgan and Claypool.
- M. Snover, B. J. Dorr, R. Schwartz, L. Micciulla, and J. Makhoul. 2006. A study of translation edit rate with targeted human annotation. *Proceedings of AMTA, 2006*, pages 223 – 231.
- K. Sohn, H. Lee, and X. Yan. 2015. Learning structured output representation using deep conditional generative models. In *NIPS, 2015*, pages 3483–3491, Montreal, Canada.
- C. K. Sønderby, T. Raiko, L. Maaløe, S. K. Sønderby, and O. Winther. 2016. Ladder variational autoencoders. In *NIPS, 2016*, pages 3738–3746, Barcelona, Spain.
- A. Srivastava and C. Sutton. 2017. Autoencoding variational inference for topic models. In *ICLR, 2017*, Toulon, France.
- N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *JMLR*, 15:1929–1958.
- M. Stanojević and K. Sima'an. 2014. [Fitting sentence level translation evaluation with many dense features](#). In *Proceedings of EMNLP, 2014*, pages 202–206, Doha, Qatar.
- I. Sutskever, O. Vinyals, and Q. V. V. Le. 2014. Sequence to sequence learning with neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, editors, *NIPS, 2014*, pages 3104–3112. Montreal, Canada.
- J. Xu and G. Durrett. 2018. [Spherical latent spaces for stable variational autoencoders](#). In *Proceedings of EMNLP, 2018*, pages 4503–4513, Brussels, Belgium.
- B. Zhang, D. Xiong, j. su, H. Duan, and M. Zhang. 2016. [Variational neural machine translation](#). In *Proceedings of EMNLP, 2016*, pages 521–530, Austin, USA.
- J. Zhang and C. Zong. 2016. [Exploiting source-side monolingual data in neural machine translation](#). In *Proceedings of EMNLP, 2016*, pages 1535–1545, Austin, Texas.
- Z. Zhang, S. Liu, M. Li, M. Zhou, and E. Chen. 2018. [Joint training for neural machine translation models with monolingual data](#). In *Proceedings of AAAI, 2018*, pages 555–562, New Orleans, USA.

A Architectures

Here we describe parameterisation of the different models presented in §3. Rather than completely specifying standard blocks, we use the notation $\text{block}(\text{inputs}; \text{parameters})$, where we give an indication of the relevant parameter set. This makes it easier to visually track which model a component belongs to.

A.1 Source Language Model

The source language model consists of a sequence of categorical draws for $i = 1, \dots, |x|$

$$X_i | z, x_{<i} \sim \text{Cat}(g_\theta(z, x_{<i})) \quad (13)$$

parameterised by a single-layer recurrent neural network using GRU units:

$$\mathbf{f}_i = \text{emb}(x_i; \theta_{\text{emb-x}}) \quad (14a)$$

$$\mathbf{h}_0 = \tanh(\text{affine}(z; \theta_{\text{init-lm}})) \quad (14b)$$

$$\mathbf{h}_i = \text{GRU}(\mathbf{h}_{i-1}, \mathbf{f}_{i-1}; \theta_{\text{gru-lm}}) \quad (14c)$$

$$g_\theta(z, x_{<i}) = \text{softmax}(\text{affine}(\mathbf{h}_i; \theta_{\text{out-x}})). \quad (14d)$$

We initialise the GRU cell with a transformation (14b) of the stochastic encoding z . For the simple joint model baseline we initialise the GRU with a vector of zeros as there is no stochastic encoding we can condition on in that case.

A.2 Translation Model

The translation model consists of a sequence of categorical draws for $j = 1, \dots, |y|$

$$Y_j | z, x, y_{<j} \sim \text{Cat}(f_\theta(z, x, y_{<j})) \quad (15)$$

parameterised by an architecture that roughly follows Bahdanau et al. (2015). The encoder is a bidirectional GRU encoder (16b) that shares source embeddings with the language model (14a) and is initialised with its own projection of the latent representation put through a tanh activation. The decoder, also initialised with its own projection of the latent representation (16d), is a single-layer recurrent neural network with GRU units (16f). At any timestep the decoder is a function of the previous state, previous output word embedding, and a context vector. This context vector (16e) is a weighted average of the bidirectional source encodings, of which the weights are computed by a Bahdanau-style attention mechanism. The output of the GRU decoder is projected to the target vocabulary size and mapped to the simplex using a

softmax activation (17) to obtain the categorical parameters:

$$\mathbf{s}_0 = \tanh(\text{affine}(z; \theta_{\text{init-enc}})) \quad (16a)$$

$$\mathbf{s}_1^m = \text{BiGRU}(\mathbf{f}_1^m, \mathbf{s}_0; \theta_{\text{bigru-x}}) \quad (16b)$$

$$\mathbf{e}_j = \text{emb}(y_j; \theta_{\text{emb-y}}) \quad (16c)$$

$$\mathbf{t}_0 = \tanh(\text{affine}(z; \theta_{\text{init-dec}})) \quad (16d)$$

$$\mathbf{c}_j = \text{attention}(\mathbf{s}_1^m, \mathbf{t}_{j-1}; \theta_{\text{bahd}}) \quad (16e)$$

$$\mathbf{t}_j = \text{GRU}(\mathbf{t}_{j-1}, [\mathbf{c}_j, \mathbf{e}_{j-1}]; \theta_{\text{gru-dec}}), \quad (16f)$$

and

$$f_\theta(z, x, y_{<j}) = \text{softmax}(\text{affine}([\mathbf{t}_j, \mathbf{e}_{j-1}, \mathbf{c}_j]; \theta_{\text{out-y}})). \quad (17)$$

In baseline models, recurrent cells are initialised with a vector of zeros as there is no stochastic encoding we can condition on.

A.3 Inference Network

The inference model $q(z|x, y, \lambda)$ is a diagonal Gaussian

$$Z|x, y \sim \mathcal{N}(\mathbf{u}, \text{diag}(\mathbf{s} \odot \mathbf{s})) \quad (18)$$

whose parameters are computed by an *inference network*. We use two bidirectional GRU encoders to encode the source and target sentences separately. To spare memory, we reuse embeddings from the generative model (19a-19b), but we prevent updates to those parameters based on gradients of the inference network, which we indicate with the function `detach`. To obtain fixed-size representations for the sentences, GRU encodings are averaged (19c-19d).

$$\mathbf{f}_1^m = \text{detach}(\text{emb}(x_1^m; \theta_{\text{emb-x}})) \quad (19a)$$

$$\mathbf{e}_1^n = \text{detach}(\text{emb}(y_1^n; \theta_{\text{emb-y}})) \quad (19b)$$

$$\mathbf{h}_x = \text{avg}(\text{BiGRU}(\mathbf{f}_1^m; \lambda_{\text{gru-x}})) \quad (19c)$$

$$\mathbf{h}_y = \text{avg}(\text{BiGRU}(\mathbf{e}_1^n; \lambda_{\text{gru-y}})) \quad (19d)$$

$$\mathbf{h}_{xy} = \text{concat}(\mathbf{h}_x, \mathbf{h}_y) \quad (19e)$$

$$\mathbf{h}_u = \text{ReLU}(\text{affine}(\mathbf{h}_{xy}; \lambda_{\text{u-hid}})) \quad (19f)$$

$$\mathbf{h}_s = \text{ReLU}(\text{affine}(\mathbf{h}_{xy}; \lambda_{\text{s-hid}})) \quad (19g)$$

$$\mathbf{u} = \text{affine}(\mathbf{h}_u; \lambda_{\text{u-out}}) \quad (19h)$$

$$\mathbf{s} = \text{softplus}(\text{affine}(\mathbf{h}_s; \lambda_{\text{s-out}})) \quad (19i)$$

We use a concatenation \mathbf{h}_{xy} of the average source and target encodings (19e) as inputs to compute the parameters of the Gaussian approximate posterior, namely, d -dimensional location and scale vectors. Both transformations use ReLU hidden activations (Nair and Hinton, 2010), but locations

live in \mathbb{R}^d and therefore call for linear output activations (19h), whereas scales live in $\mathbb{R}_{>0}^d$ and call for strictly positive outputs (19i), we follow [Kucukelbir et al. \(2017\)](#) and use softplus. The complete set of parameters used for inference is thus $\lambda = \{\lambda_{\text{gru-x}}, \lambda_{\text{gru-y}}, \lambda_{\text{u-hid}}, \lambda_{\text{u-out}}, \lambda_{\text{s-hid}}, \lambda_{\text{s-out}}\}$.

A.4 Prediction Network

The *prediction network* parameterises our prediction model $r(z|x, \phi)$, a variant of the inference model that conditions on the source sentence alone. In §4 we explore several variants of the ELBO using different parameterisations of r_ϕ . In the simplest case we do not condition on the target sentence during training, thus we can use the same network both for training and prediction. The network is similar to the one described in A.3, except that there is a single bidirectional GRU and we use the average source encoding (19c) as input to the predictors for \mathbf{u} and \mathbf{s} (20c-20d).

$$\mathbf{h}_u = \text{ReLU}(\text{affine}(\mathbf{h}_x; \lambda_{\text{u-hid}})) \quad (20a)$$

$$\mathbf{h}_s = \text{ReLU}(\text{affine}(\mathbf{h}_x; \lambda_{\text{s-hid}})) \quad (20b)$$

$$\mathbf{u} = \text{affine}(\mathbf{h}_u; \lambda_{\text{u-out}}) \quad (20c)$$

$$\mathbf{s} = \text{softplus}(\text{affine}(\mathbf{h}_s; \lambda_{\text{s-out}})) \quad (20d)$$

In all other cases we use $q(z|x, y, \lambda)$ parameterised as discussed in A.3 for training, and design a separate network to parameterise r_ϕ for prediction. Much like the inference model, the prediction model is a diagonal Gaussian

$$Z|x \sim \mathcal{N}(\hat{\mathbf{u}}, \text{diag}(\hat{\mathbf{s}} \odot \hat{\mathbf{s}})) \quad (21)$$

also parameterised by d -dimensional location and scale vectors, however in predicting $\hat{\mathbf{u}}$ and $\hat{\mathbf{s}}$ (22d-22e) it can only access an encoding of the source (22a).

$$\mathbf{h}_x = \text{avg}(\text{BiGRU}(\mathbf{f}_1^m; \phi_{\text{gru-x}})) \quad (22a)$$

$$\mathbf{h}_u = \text{ReLU}(\text{affine}(\mathbf{h}_x; \phi_{\text{u-hid}})) \quad (22b)$$

$$\mathbf{h}_s = \text{ReLU}(\text{affine}(\mathbf{h}_x; \phi_{\text{s-hid}})) \quad (22c)$$

$$\hat{\mathbf{u}} = \text{affine}(\mathbf{h}_u; \phi_{\text{u-out}}) \quad (22d)$$

$$\hat{\mathbf{s}} = \text{softplus}(\text{affine}(\mathbf{h}_s; \phi_{\text{s-out}})) \quad (22e)$$

The complete set of parameters is then $\phi = \{\phi_{\text{gru-x}}, \phi_{\text{u-hid}}, \phi_{\text{u-out}}, \phi_{\text{s-hid}}, \phi_{\text{s-out}}\}$. For the deterministic variant, we use $\hat{\mathbf{u}}$ (22d) alone to approximate \mathbf{u} (19h), i.e. the posterior mean of Z .

B Graphical models

Figure 1 is a graphical depiction of our AEVNMT model. Circled nodes denote random variables while uncircled nodes denote deterministic quantities. Shaded random variables correspond to observations and unshaded random variables are latent. The plate denotes a dataset of $|\mathcal{D}|$ observations.

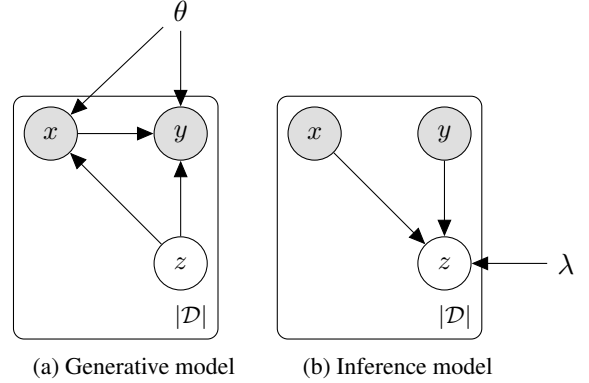




Figure 1: On the left we have AEVNMT, a generative model parameterised by neural networks. On the right we show an independently parameterised model used for approximate posterior inference.


In Figure 2a, we illustrate the precise statistical assumptions of AEVNMT. Here plates iterate over words in either the source or the target sentence. Note that the arrow from x_i to y_j states that the j th target word depends on all of the source sentence, not on the i th source word alone, and that is the case because x_i is within the source plate. In Figure 2b, we illustrate the statistical dependencies induced in the marginal distribution upon marginalisation of latent variables. Recall that the marginal is the distribution which by assumption produced the observed data. Now compare that to the distribution modelled by the simple JOINT model (Figure 2c). Marginalisation induces undirected dependencies amongst random variables creating more structure in the marginal distribution. In graphical models literature this is known as *moralisation* ([Koller and Friedman, 2009](#)).



(a) Joint distribution of AEVNMT



(b) Marginal distribution of AEVNMT



(c) Joint distribution modelled without latent variables

Figure 2: Here we zoom in into the model of Figure 1a to show the statistical dependencies between observed variables. In the joint distribution (top), we have the directed dependency of a source word on all of the previous source words, and similarly, of a target word on all of the previous target words in addition to the complete source sentence. Besides, all observations depend directly on the latent variable Z . Marginalisation of Z (middle) ties all variables together through undirected connections. At the bottom we show the distribution we get if we model the data distribution directly without latent variables.

C Robustness to out-of-domain data

We use our stronger models, those trained on gold-standard NC bilingual data and synthetic News data, to translate test sets in various unseen genres. These data sets are collected and distributed by TAUS,⁸ and have been used in scenarios of adaptation to all domains at once (Cuong et al., 2016). Table 6 shows the performance of AEVNMT and the conditional baseline. The first thing to note is the remarkable drop in performance showing that without taking domain adaptation seriously both models are inadequate. In terms of BLEU, differences range from -0.3 to 0.8 (EN-DE) and 0.3 to 0.7 (DE-EN) and are mostly in favour of AEVNMT, though note the increased standard deviations.

⁸TAUS Hardware, TAUS Software, TAUS Industrial Electronics, TAUS Professional & Business Services, and TAUS Legal available from TAUS data cloud <http://tausdata.org/>.

Model	Domain	EN-DE					DE-EN				
		BLEU \uparrow	METEOR \uparrow	TER \downarrow	BEER \uparrow	BLEU \uparrow	METEOR \uparrow	TER \downarrow	BEER \uparrow		
COND	Computer Hardware	8.7 (0.3)	23.6 (0.2)	82.3 (4.0)	44.2 (0.3)	13.0 (0.4)	20.8 (0.1)	73.5 (3.6)	49.2 (0.1)		
	Computer Software	7.9 (0.2)	22.9 (0.1)	83.6 (3.2)	44.1 (0.3)	11.8 (0.8)	20.4 (0.3)	80.8 (7.0)	48.6 (0.3)		
	Industrial Electronics	7.6 (0.3)	21.8 (0.2)	89.6 (5.5)	42.7 (0.3)	10.1 (1.0)	17.6 (0.4)	85.7 (9.7)	46.0 (0.4)		
	Professional & Business Services	8.1 (0.4)	23.3 (0.4)	75.6 (0.4)	43.1 (0.4)	11.1 (0.2)	20.5 (0.1)	73.8 (1.4)	44.9 (0.1)		
	Legal	12.3 (0.1)	29.7 (0.3)	75.3 (0.5)	49.6 (0.1)	14.2 (0.3)	23.0 (0.3)	67.3 (2.1)	50.8 (0.3)		
AEVNM	Computer Hardware	9.1 (0.3)	23.9 (0.4)	79.4 (1.3)	44.7 (0.2)	13.6 (0.3)	21.0 (0.2)	70.6 (1.4)	49.3 (0.2)		
	Computer Software	8.3 (0.1)	23.2 (0.2)	81.4 (0.7)	44.5 (0.2)	12.2 (0.5)	20.7 (0.3)	77.0 (3.7)	48.6 (0.3)		
	Industrial Electronics	8.1 (0.2)	22.2 (0.3)	85.2 (1.0)	43.1 (0.2)	10.8 (0.6)	17.9 (0.3)	79.9 (4.9)	46.3 (0.2)		
	Professional & Business Services	7.8 (0.7)	23.0 (0.7)	75.7 (1.0)	43.0 (0.5)	11.5 (0.2)	20.6 (0.2)	73.1 (0.7)	45.1 (0.1)		
	Legal	13.1 (0.2)	30.5 (0.3)	74.4 (0.6)	49.9 (0.2)	14.5 (0.2)	23.2 (0.2)	66.1 (1.3)	50.9 (0.3)		

Table 6: Performance of models trained with NC and back-translated News on various TAUS test sets: we report average (1std) across 5 independent runs.

D Validation results

	WMT16		IWSLT14	
	EN-DE	DE-EN	EN-DE	DE-EN
COND	14.5 (0.2)	16.9 (0.2)	25.1 (0.1)	30.8 (0.1)
JOINT	14.8	17.1	25.2	31.0
AEVNMT	14.8 (0.2)	17.4 (0.2)	25.7 (0.0)	31.4 (0.0)

Table 7: Validation results reported in BLEU for in-domain training on NC and IWSLT: we report average (1std) across 5 independent runs for COND and AEVNMT, but a single run of JOINT.

WMT & IWSLT	EN-DE	DE-EN
COND	20.5 (0.1)	25.9 (0.1)
JOINT	20.7	26.1
AEVNMT	20.8 (0.1)	26.1 (0.1)

Table 8: Validation results reported in BLEU for mixed-domain training: we report average (1std) across 5 independent runs for COND and AEVNMT, but a single run of JOINT. The validation set used is a concatenation of the development sets from WMT and IWSLT.

WMT16	EN-DE	DE-EN
COND	14.5 (0.2)	16.9 (0.2)
+ synthetic data	17.4 (0.1)	21.8 (0.1)
JOINT + synthetic data	17.3	21.8
AEVNMT + synthetic data	17.6 (0.1)	22.1 (0.1)

Table 9: Validation results reported in BLEU for training on NC plus synthetic data: we report average (1std) across 5 independent runs for COND and AEVNMT, but a single run of JOINT.

E Additional Metrics

IWSLT14	EN-DE				DE-EN			
	BLEU ↑	METEOR ↑	TER ↓	BEER ↑	BLEU ↑	METEOR ↑	TER ↓	BEER ↑
COND	23.0 (0.1)	42.4 (0.1)	56.0 (0.1)	58.6 (0.1)	27.3 (0.2)	30.3 (0.1)	52.4 (0.5)	59.8 (0.1)
JOINT	23.2	42.8	56.1	58.7	27.5	30.3	52.7	59.8
AEVNMT	23.4 (0.1)	42.8 (0.2)	55.5 (0.3)	58.8 (0.1)	28.0 (0.1)	30.6 (0.1)	51.2 (0.6)	60.1 (0.1)

Table 10: Test results for in-domain training on IWSLT: we report average (1std) across 5 independent runs for COND and AEVNMT, but a single run of JOINT.

WMT16	EN-DE				DE-EN			
	BLEU ↑	METEOR ↑	TER ↓	BEER ↑	BLEU ↑	METEOR ↑	TER ↓	BEER ↑
COND	17.8 (0.2)	35.9 (0.2)	65.2 (0.4)	53.1 (0.1)	20.1 (0.1)	26.0 (0.1)	62.0 (0.3)	53.7 (0.1)
JOINT	17.9	36.2	64.1	53.4	20.1	26.0	62.7	53.7
AEVNMT	18.4 (0.2)	36.6 (0.2)	64.0 (0.4)	53.5 (0.1)	20.6 (0.2)	26.0 (0.1)	60.7 (0.8)	53.6 (0.1)

Table 11: Test results for in-domain training on NC: we report average (1std) across 5 independent runs for COND and AEVNMT, but a single run of JOINT.

WMT16	EN-DE				DE-EN			
	BLEU ↑	METEOR ↑	TER ↓	BEER ↑	BLEU ↑	METEOR ↑	TER ↓	BEER ↑
COND	17.8 (0.2)	35.9 (0.2)	65.2 (0.4)	53.1 (0.1)	20.1 (0.1)	26.0 (0.1)	62.0 (0.3)	53.7 (0.1)
+ synthetic data	22.3 (0.3)	40.9 (0.2)	58.5 (0.5)	57.0 (0.2)	26.9 (0.2)	30.4 (0.1)	53.0 (0.5)	58.5 (0.1)
JOINT + synthetic data	22.2	40.8	58.1	57.0	26.7	30.2	52.1	58.6
AEVNMT + synthetic data	22.5 (0.2)	41.0 (0.1)	58.1 (0.2)	57.0 (0.1)	27.4 (0.2)	30.6 (0.1)	52.0 (0.1)	58.8 (0.1)

Table 12: Test results for training on NC plus synthetic data (back-translated News Crawl): we report average (1std) across 5 independent runs for COND and AEVNMT, but a single run of JOINT.

EN-DE		WMT16					IWSLT14						
Training	Model	BLEU ↑	METEOR ↑	TER ↓	BEER ↑	BLEU ↑	METEOR ↑	TER ↓	BEER ↑	BLEU ↑	METEOR ↑	TER ↓	BEER ↑
EN-DE	COND	17.6 (0.4)	35.7 (0.3)	61.9 (0.7)	53.9 (0.2)	23.9 (0.3)	43.1 (0.3)	54.7 (0.2)	59.3 (0.1)				
	JOINT	18.1	36.3	61.4	54.3	24.2	43.4	54.5	59.5				
	AEVNMNT	18.4 (0.2)	36.6 (0.3)	60.9 (0.4)	54.5 (0.2)	24.1 (0.3)	43.5 (0.3)	54.0 (0.3)	59.5 (0.2)				
DE-EN	COND	21.6 (0.2)	27.4 (0.1)	59.9 (0.7)	55.5 (0.2)	29.1 (0.2)	31.5 (0.1)	50.9 (0.5)	60.9 (0.1)				
	JOINT	22.3	27.4	58.8	55.6	29.2	31.5	49.2	61.2				
	AEVNMNT	22.3 (0.1)	27.5 (0.1)	57.8 (0.6)	55.6 (0.1)	29.2 (0.1)	31.5 (0.0)	49.2 (0.4)	61.1 (0.1)				

Table 13: Test results for mixed-domain training: we report average (1std) across 5 independent runs for COND and AEVNMNT, but a single run of JOINT.