Evaluating for Diversity in Question Generation over Text

Michael Sejr Schlichtkrull[*]
University of Amsterdam
University of Edinburgh
`m.s.schlichtkrull@uva.nl`

Weiwei Cheng
Amazon
`weiweic@amazon.com`

Abstract

Generating diverse and relevant questions over text is a task with widespread applications. We argue that commonly-used evaluation metrics such as BLEU and METEOR are not suitable for this task due to the inherent diversity of reference questions, and propose a scheme for extending conventional metrics to reflect diversity. We furthermore propose a variational encoder-decoder model for this task. We show through automatic and human evaluation that our variational model improves diversity without loss of quality, and demonstrate how our evaluation scheme reflects this improvement.

1 Introduction

Question generation has widespread applications in online education (Lindberg et al., 2013), search (Rothe et al., 2014), and question answering (Yang et al., 2017; Lewis and Fan, 2019). Generating a single question per context can be inadequate as questions are naturally diverse in aspect, answer, and phrasing. Following Li et al. (2016), the first two correspond to *semantic diversity*, while the third one corresponds to *lexical diversity*.

Evaluation in question generation uses metrics including BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), and ROUGE (Lin, 2004). Common to these metrics is an implicit assumption that *references are paraphrases*. In question generation this assumption does not hold, as multiple *diverse* questions may be equally relevant for a given context. To measure the ability of a system to produce diverse sets of questions, a new evaluation metric is needed. We propose a function composition framework consisting of meta-metrics relying on existing evaluation metrics. As we will show in Section 3, our framework generalizes F_1, a widely-used measure in information retrieval and multivariate analysis.

---

[*] Work done as an intern at Amazon.

Context: Amazon.com, Inc. is an American electronic commerce and cloud computing company based in Seattle, Washington, founded by Jeff Bezos on July 5, 1994.

Question 1:   Which Seattle-based company was founded in 1994?
                                                            — Amazon.com

Question 2:   What is the name of the company that Jeff Bezos founded?
                                                            — Amazon.com

Question 3:   When was Amazon.com founded?
                                                            — July 5, 1994

Figure 1: Example context with questions naturally diverse in subject, structure, and answer.

Recent papers on text generation over images have sought to increase diversity through generative modeling such as variational autoencoders (Jain et al., 2017). To demonstrate the usefulness of our F-like metrics, we build a conditional variational autoencoder for text-to-text generation inspired by these ideas and based on Zhang et al. (2016). We achieve significant improvements in diversity without performance loss, in terms of both automatic and human evaluation. We show that our F-like approach better captures the improvement in diversity offered by this model, and that the breakdown into "precision" and "recall" enabled by the meta-metric helps illustrate the strengths and weaknesses of different systems.

2 Related Work

Early work on natural language question generation focuses primarily on rule-based models (Heilman and Smith, 2010; Agarwal and Mannem, 2011) and template-based slot filling (Lindberg et al., 2013). End-to-end neural models were introduced by Du et al. (2017), and extende in (Chan and Fan, 2019; Bao et al., 2020). Generating diverse questions is discussed in a series of recent papers (Zhou et al.,

2017; Harrison and Walker, 2018; Song et al., 2018; Yao et al., 2018; Shen et al., 2020) attempting to improve *semantic diversity* by conditioning on (potential) answer positions or question types. While showing promising results, such prior information may not always be practically available. Sultan et al. (2020) recently applied nucleus sampling (Holtzman et al., 2020) to diversify question generation models, demonstrating improved performance on a downstream question answering task.

Diverse text generation has been studied for other tasks, including image question generation (Jain et al., 2017), conversation modeling (Li et al., 2016), machine translation (Zhang et al., 2016; Schulz et al., 2018), and image captioning (Vijayakumar et al., 2016; Pu et al., 2016; Dai et al., 2017). Models typically rely either on conditional variational autoencoders (CVAE) (Kingma and Welling, 2013; Sohn et al., 2015) or conditional generative adversarial networks (CGAN) (Mirza and Osindero, 2014), and evaluation has relied on BLEU, ROUGE, and METEOR.

Several schemes for evaluating with diversity in image captioning were proposed in (Alihosseini et al., 2019). Their proposals rely either on statistical divergence between language modeling of the generated and reference sets, or on Jaccard index computed at the n-gram level. In Dhingra et al. (2019), another metric is proposed for scoring systems where generated sentences overlapping with items on the *source* side rather than the *target* side should also be scored highly.

3 Evaluating for Diversity

Conventional metrics evaluate the correctness of proposed questions in relation to a set of reference questions; they do not, however, measure the degree to which the proposed questions *cover* the set of reference questions. Consider the sentence *"Germany won the 2014 world cup"*, paired with two reference questions *"Who won the 2014 world cup?"* and *"Which event did Germany win in 2014?"*. Using traditional metrics, a system that always generates *"Who won the 2014 world cup?"* and a system that alternates between generating the two would be given equal, perfect scores, when in fact the first system has only learned half the task. Moreover, systems that generate a sentence usings parts of both questions may wrongly be scored highly (see Appendix A for an example).

In this work, we propose a framework to extend

commonly used scoring functions to account for coverage over reference questions. Given a set of reference questions $R \subset \Omega$, a set of predicted questions $P \subset \Omega$, and a scoring function $s : \Omega \times \Omega \to \mathbb{R}$, we propose the following two functions for comparing P and R:

$$u(P, R; s) = \frac{1}{|P|} \sum_{p \in P} \max_{r \in R} s(p, r) \quad (1)$$

$$v(P, R; s) = \frac{1}{|R|} \sum_{r \in R} \max_{p \in P} s(p, r) \quad (2)$$

To compute the function u, we identify for each predicted question the best match in the references w.r.t. the scoring function. We then compute how "close" the predictions are to the references by summing up scores between each predicted question and its respective best match. The function v is computed analogously in Eq. (2).[1] We combine u and v with their harmonic mean. This leads to an overall measure to compare P and R,

$$F(P, R) = \frac{2 \times u(P, R) \times v(P, R)}{u(P, R) + v(P, R)} . \quad (3)$$

In the special case of a binary scoring function $s : \Omega \times \Omega \to \{0, 1\}$, where 1 is given to exact matches, u and v are identical to conventional *precision* and *recall*, respectively, and Eq. (3) is equal to F_1. Therefore, we can interpret u and v as generalized precision and recall, respectively.

To the best of our knowledge, our construction of the F function is new. The "best match" idea used in comparing P and R, i.e., examining the optimal score an item in the set can attain w.r.t. another set, has been applied in local community detection in network analysis (Clauset, 2005; Lancichinetti et al., 2009) and behavior research in social networks (Adali et al., 2010). In machine learning, the work by Goldberg et al. (2010) on clustering analysis bears the closest resemblance to this idea. However, their assumption of a symmetric s-function does not hold for many NLP applications.

4 Variational Question Generation

Our approach extends the encoder-decoder model proposed by Du et al. (2017) with a simple latent variable following Zhang et al. (2016). Given a training corpus of context-question pairs $\mathcal{D} = \{(X_1, Y_1), (X_2, Y_2), \ldots, (X_D, Y_D)\}$ such

---

[1]Note that $v(P, R; s) \neq u(R, P; s)$ since s is not necessarily symmetric, i.e., $s(x, y) \neq s(y, x)$.

that $X_i = \{x_{i1}, \ldots, x_{ij}\}$ and $Y_i = \{y_{i1}, \ldots, y_{ik}\}$, the objective is to minimize the negative log-likelihood of the training data. We first encode each token in the context paragraph using either GloVe embeddings (Pennington et al., 2014) or ELMo embeddings (Peters et al., 2018). The entire sentence is then encoded through a 600-dimensional BiLSTM, giving forwards, backwards, and concatenated representations $\overrightarrow{b_t}, \overleftarrow{b_t}, b_t = \left[\overrightarrow{b_t}, \overleftarrow{b_t}\right]$ for each timestep t.

We construct context representations for each decoding step using bilinear attention as presented by Luong et al. (2015). That is, given a query vector k_t corresponding to decoding step t:

$$a_{i,t} = \frac{\exp(k_t^T W_a b_i)}{\sum_j \exp(k_t^T W_a b_j)} \quad (4)$$

$$r_t = \sum_i^{|X|} a_{i,t} b_i \quad (5)$$

$$c_t = \tanh(W_c [r_t, k_t]) \quad (6)$$

We compute query vectors using another BiLSTM. The input to each step t consists of the concatenation of the GloVe embedding e_t of the target-side token generated at the previous timestep $t-1$, and the previous context vector c_{t-1}. That is

$$k_t = \text{LSTM}([e_t, c_{t-1}], k_{t-1}). \quad (7)$$

We compute c_0 as $\left[\overrightarrow{b_{|X|}}, \overleftarrow{b_1}\right]$. Finally, the distribution over target-side tokens at decoder step t is computed as

$$p(\hat{y}_t \mid X, y_1, \ldots, y_{t-1}) = \text{softmax}(W_o c_t + b_o). \quad (8)$$

At training time, we rely on teacher forcing to provide the token to be embedded in e_t. That is, we use the gold token y_{t-1}. At inference time, we use the generated token $\hat{y}_{t-1}$ at the previous timestep. We decode using beam search with a beam size of three, following Du et al. (2017).

Following (Zhang et al., 2016; Jain et al., 2017), we introduce a latent variable $z \in \mathbb{R}^d$ conditional on x in the decoder to model the underlying semantic space. We redefine the query vectors used for attention, Eq. (7), as

$$k_t' = \text{LSTM}([e_t, c_{t-1}, z], k_{t-1}'). \quad (9)$$

We introduce an approximation $q_\phi(z \mid x, y)$ for the intractable true posterior $p(z \mid x, y)$. Instead of the

true log-likelihood, we optimize the *evidence lower bound* (ELBO), a key idea underpinning variational autoencoders (Kingma and Welling, 2013; Sohn et al., 2015):

$$\mathcal{L}(x, y; \theta, \phi) = -\text{KL}\left(q_\phi(z \mid x, y) \| p_\theta(z \mid x)\right) + \mathbb{E}_{q_\phi(z \mid x, y)} \left[\log p_\theta(y \mid z, x)\right] \quad (10)$$

Following Zhang et al. (2016), we define $q_\phi(z \mid x, y)$ through a neural network with parameters ϕ as a Gaussian of the form

$$q_\phi(z \mid x, y) = \mathcal{N}\left(z; \mu(x, y), \sigma(x, y)^2 \mathbf{I}\right). \quad (11)$$

We obtain a representation of the context paragraph used in the posterior by mean-pooling over the context encoder states. Similarly, we represent the target question using mean-pooled ELMo vectors (Peters et al., 2018). That is,

$$h_c = \frac{1}{|x|} \sum_{i=1}^{|x|} b_i, \quad h_q = \frac{1}{|y|} \sum_{i=1}^{|y|} \text{ELMo}(y_i). \quad (12)$$

We then obtain the Gaussian parameters

$$h_z = \text{ReLU}\left(W_z [h_c, h_q] + b_z\right) \quad (13)$$

$$\mu = W_\mu h_z + b_\mu, \quad \log \sigma = W_\sigma h_z + b_\sigma. \quad (14)$$

Similarly, we model the conditional prior $p(z \mid x)$ using a neural network

$$p_\theta(z \mid x) = \mathcal{N}(z; \mu'(x), \sigma'(x)^2 \mathbf{I}) \quad (15)$$

$$h_z' = \text{ReLU}(W_z' h_c + b_z') \quad (16)$$

$$\mu' = W_\mu' h_z' + b_\mu', \quad \log \sigma' = W_\sigma' h_z' + b_\sigma'. \quad (17)$$

5 Experiments and Evaluations

We compare our model to the deterministic baseline from Du et al. (2017)[2] on the SQuAD dataset (Rajpurkar et al., 2016). We demonstrate comparable performance using conventional metrics, and better performance on our proposed metrics. We further corroborate this finding through human evaluation. Inspired by Vijayakumar et al. (2016), we also evaluate a random beam selection (RBS) heuristic, where we induce diversity by sampling from the top b beams. Details of hyperparameters are given in Appendix B.

---

[2]The results for our implementation differ slightly from the ones in their paper. We suspect this is due to differences between our from-scratch TensorFlow implementation and their Torch-based extension of the OpenNMT framework.

system	METEOR	P-METEOR	R-METEOR	F-METEOR
baseline	15.98	15.97	14.02	14.90
CVAE (25d)	16.32	16.29	14.13	15.02
CVAE (50d)	16.31	16.28	14.42	15.11
CVAE (100d)	16.34	16.32	14.76	15.31
baseline + RBS	16.19	16.19	14.55	15.19
CVAE + RBS	**16.71**	**16.71**	**15.43**	**15.93**

Table 1: METEOR and our extensions P-, R-, and F-METEOR. METEOR scores are first averaged per context across questions and then in total across all contexts.

Table 1 reports the results where METEOR is chosen as the s-function. METEOR has been shown to correlate well with human judgments (Banerjee and Lavie, 2005). In the interest of space, results of BLEU and ROUGE are given in Appendix C. For METEOR, we average first per context c and subsequently over the dataset $\mathcal{D}$ to prevent overweighting contexts with more reference questions:

$$s(\mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{(c,Y) \in \mathcal{D}} \frac{1}{|Y|} \sum_{y \in Y} s(c,y). \qquad (18)$$

We report results of the variational model using 25, 50, and 100 dimensional latent variables, along with the random beam selection extension of both systems the two-layer baseline and the 100-dimensional CVAE. The 100-dimensional version of our variational model performs favorably, especially in terms of F-metrics. Through our recall score, we can identify exactly where the improvement occurs –our models show larger improvement on recall, e.g. they match more references.

Increasing the dimensionality of the latent variable z strictly improves the recall of our models, and consequently the F-metrics. This supports our intuition that the degree of diversity the model can express is controlled by the amount of information that is encoded in the latent variable. The RBS heuristic adds a small but rather consistent gain to both the baseline *and* the CVAE model. Example outputs of our system can be found in Appendix D.

To confirm our findings, we use Amazon Mechanical Turk to conduct human evaluation. We presented the annotators a context paragraph with questions generated by the baseline and our CVAE model. We tasked them to rate the example with three criteria, *fluency*, *relevancy*, and *diversity*.

Table 2 shows the two systems, with their RBS extension, performed comparably in terms of flu-

system	fluency	relevancy	diversity
baseline	2.49	2.62	1.11
baseline + RBS	2.56	2.64	1.57
CVAE	**2.57**	2.64	1.78
CVAE + RBS	2.47	**2.65**	**1.86**

Table 2: Human evaluation on the test set. For each criterion, the score ranges from 1 to 3, indicating, e.g., *not fluent*, *almost fluent*, *fluent*. Each example is rated by three annotators.

ency and relevancy, whereas the CVAE demonstrated a significant improvement in diversity. (With two-sample t-tests, $p < 0.05$ for both CVAE vs. baseline and CVAE+RBS vs. baseline+RBS. More details are in the annotation file.) The finding is consistent with the automatic evaluation.

6 Conclusion

We have introduced a framework to extend existing evaluation metrics into F_1-like scoring functions explicitly rewarding diversity and enabling detailed comparison in terms of precision and recall. Furthermore, we have presented the first variational autoencoder for question generation over general text. Our model shows comparable results to the baseline in terms of conventional evaluation metrics, while producing significantly more diverse questions according to human and automatic evaluation. Our experiments suggest that the modeling of diversity is an important aspect of question generation systems, both to generate engaging questions and to better model the inherently diverse training inputs, and our development of a family of metrics suitable for evaluating models in this setting represents a step in that direction.

References

Sibel Adali, Robert Escriva, Mark Goldberg, Mykola Hayvanovych, Malik Magdon-Ismail, Boleslaw Szymanski, William Wallace, and Gregory Williams. 2010. Measuring behavioral trust in social networks. In *IEEE International Conference on Intelligence and Security Informatics*, pages 150–152. IEEE.

Manish Agarwal and Prashanth Mannem. 2011. Automatic gap-fill question generation from text books. In *Proceedings of the Sixth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 56–64, Portland, Oregon. Association for Computational Linguistics.

Danial Alihosseini, Ehsan Montahaei, and Mahdieh Soleymani Baghshah. 2019. Jointly measuring diversity and quality in text generation models. In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 90–98, Minneapolis, Minnesota. Association for Computational Linguistics.

Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.

Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Songhao Piao, Jianfeng Gao, Ming Zhou, et al. 2020. Unilmv2: Pseudo-masked language models for unified language model pre-training. *arXiv preprint arXiv:2002.12804*.

Ying-Hong Chan and Yao-Chung Fan. 2019. BERT for question generation. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 173–177, Tokyo, Japan. Association for Computational Linguistics.

Aaron Clauset. 2005. Finding local community structure in networks. *Physical Review E*, 72:026132.

Bo Dai, Sanja Fidler, Raquel Urtasun, and Dahua Lin. 2017. Towards diverse and natural image descriptions via a conditional GAN. In *Proceedings of the 2017 IEEE International Conference on Computer Vision*, pages 2989–2998. IEEE.

Bhuwan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William Cohen. 2019. Handling divergent reference texts when evaluating table-to-text generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4884–4895, Florence, Italy. Association for Computational Linguistics.

Xinya Du, Junru Shao, and Claire Cardie. 2017. Learning to ask: Neural question generation for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1342–1352, Vancouver, Canada. Association for Computational Linguistics.

Mark Goldberg, Mykola Hayvanovych, and Malik Magdon-Ismail. 2010. Measuring similarity between sets of overlapping clusters. In *Proceedings of the 2010 IEEE Second International Conference on Social Computing*, pages 303–308. IEEE Computer Society.

Vrindavan Harrison and Marilyn Walker. 2018. Neural generation of diverse questions using answer focus, contextual and linguistic features. *arXiv prepring arXiv:1809.02637*.

Michael Heilman and Noah A. Smith. 2010. Good question! statistical ranking for question generation. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 609–617, Los Angeles, California. Association for Computational Linguistics.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. The curious case of neural text degeneration. In *International Conference on Learning Representations*.

Unnat Jain, Ziyu Zhang, and Alexander Schwing. 2017. Creativity: Generating diverse questions using variational autoencoders. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 5415–5424. IEEE.

Diederik P Kingma and Max Welling. 2013. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.

Andrea Lancichinetti, Santo Fortunato, and János Kertész. 2009. Detecting the overlapping and hierarchical community structure in complex networks. *New Journal of Physics*, 11(3):033015.

Mike Lewis and Angela Fan. 2019. Generative question answering: Learning to answer the whole question. In *The 7th International Conference on Learning Representations*.

Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119. Association for Computational Linguistics.

Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

David Lindberg, Fred Popowich, John Nesbit, and Phil Winne. 2013. Generating natural language questions to support learning on-line. In *Proceedings of the 14th European Workshop on Natural Language Generation*, pages 105–114, Sofia, Bulgaria. Association for Computational Linguistics.

Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.

Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.

Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.

Yunchen Pu, Zhe Gan, Ricardo Henao, Xin Yuan, Chunyuan Li, Andrew Stevens, and Lawrence Carin. 2016. Variational autoencoder for deep learning of images, labels and captions. In *Advances in Neural Information Processing Systems*, pages 2352–2360. Curran Associates, Inc.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.

Anselm Rothe, Brenden M Lake, and Todd M Gureckis. 2014. Asking and evaluating natural language questions. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, pages 2051–2056. Cognitive Science Society, Curran Associates, Inc.

Philip Schulz, Wilker Aziz, and Trevor Cohn. 2018. A stochastic decoder for neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1243–1252, Melbourne, Australia. Association for Computational Linguistics.

Sheng Shen, Yaliang Li, Nan Du, Xian Wu, Yusheng Xie, Shen Ge, Tao Yang, Kai Wang, Xingzheng Liang, and Wei Fan. 2020. On the generation of medical question-answer pairs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:8822–8829.

Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, pages 3483–3491. Curran Associates, Inc.

Linfeng Song, Zhiguo Wang, Wael Hamza, Yue Zhang, and Daniel Gildea. 2018. Leveraging context information for natural question generation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 2, pages 569–574. Association for Computational Linguistics.

Md Arafat Sultan, Shubham Chandel, Ramón Fernandez Astudillo, and Vittorio Castelli. 2020. On the importance of diversity in question generation for QA. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5651–5656, Online. Association for Computational Linguistics.

Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. 2016. How to train deep variational autoencoders and probabilistic ladder networks. *arXiv preprint arXiv:1602.02282*.

Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*.

Zhilin Yang, Junjie Hu, Ruslan Salakhutdinov, and William W Cohen. 2017. Semi-supervised QA with generative domain-adaptive nets. *arXiv preprint arXiv:1702.02206*.

Kaichun Yao, Libo Zhang, Tiejian Luo, Lili Tao, and Yanjun Wu. 2018. Teaching machines to ask questions. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4546–4552. International Joint Conferences on Artificial Intelligence.

Biao Zhang, Deyi Xiong, Jinsong Su, Hong Duan, and Min Zhang. 2016. Variational neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*,

pages 521–530, Austin, Texas. Association for Computational Linguistics.

Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2017. Neural question generation from text: A preliminary study. In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 662–671. Springer.

A   In-Between Response

Our evaluation framework gives lower scores to in-between questions. We illustrate this now with an example in Section 5.

Suppose we have a context with the reference set $R = \{r_1, r_2\}$, where

- r_1: *who won the 2014 world cup*

- r_2: *which event did Germany win in 2014*

And the system outcome is a set P containing a single element

- p_1: *which event did the 2014 world cup*

The values of METEOR, ROUGE, BLEU (sentence-level) and the corresponding F-metrics are

- BLEU: 0.5946   F-BLEU: 0.2867

- ROUGE: 0.6240   F-ROUGE: 0.5987

- METEOR: 0.3773   F-METEOR: 0.3516

B   Hyperparameters

We implemented our models in TensorFlow. Wherever possible, we kept the hyperparameters identical to the ones used by Du et al. (2017). For the variational model, we experiment with different dimensionalities for the latent variable z, choosing from $\{25, 50, 100\}$.

Following Sønderby et al. (2016), we anneal the KL term, with a scaling factor starting at 0 and increasing at a rate of 0.03 per iteration. This rate is determined through greedy search from the set $\{0.01, 0.03, 0.05, 0.1\}$. We apply dropouts to the latent variable z on the example level, dropping out the variable entirely with a probability 0.2, selected from the set $\{0.1, 0.2, 0.3, 0.4, 0.5\}$. For the baseline, we saw slight improvements from stacking two LSTM layers in the encoder and the decoder (with dropouts of probability 0.3 applied in-between), while for the variational model multiple layers had no significant impact. For the RBS heuristic, we selected the number of beams $b = 2$ on the development set as well.

C   Results of Additional Metrics

Results with ROUGE and BLEU are summarized in Table 3 and 4.

D   System Outputs

Some example outputs from the baseline system and our best-performing CVAE model are given below.

D.1

Context: according to the doctrine of impermanence , life embodies this flux in the aging process , the cycle of rebirth (samsara), and in any experience of loss .
Du et al.: what is the cycle of rebirth ?
CVAE 1: what is the process of impermanence ?
CVAE 2: according to the doctrine of impermanence , what does the cycle of rebirth mean ?
Du et al. + RBS 1: how does the cycle of rebirth ?
Du et al. + RBS 2: how does the cycle of rebirth ?
CVAE + RBS 1: what does the term of impermanence refer to ?
CVAE + RBS 2: according to the doctrine of cycle , what does the cycle of impermanence refer to ?

D.2

Context: the world trade center path station , which opened on july 19 , 1909 as the hudson terminal , was also destroyed in the attack .
Du et al.: on what date was the world trade center ?
CVAE 1: what was destroyed in the attack ?
CVAE 2: when was the world trade center station opened ?
Du et al. + RBS 1: when did the world trade center rail open ?
Du et al. + RBS 2: when was the world trade center station opened ?
CVAE + RBS 1: what was the hudson terminal ?
CVAE + RBS 2: when did the hudson terminal open ?

D.3

Context: michiru ōshima created orchestral arrangements for the three compositions , later to be performed by an ensemble conducted by yasuzo takemoto .
Du et al.: who created musical arrangement for the three compositions ?
CVAE 1: what was the name of the orchestral ōshima created ?

system	ROUGE	P-ROUGE	R-ROUGE	F-ROUGE
baseline	37.77	37.61	33.02	34.44
CVAE (25d)	37.54	37.19	33.77	34.91
CVAE (50d)	37.42	37.24	34.01	35.32
CVAE (100d)	38.30	38.10	35.07	36.32
baseline + RBS	37.84	37.72	34.07	35.44
CVAE + RBS	**38.32**	**38.11**	**35.24**	**36.59**

Table 3: ROUGE-L and our extensions P-, R-, and F-ROUGE.

system	cBLEU	BLEU	P-BLEU	R-BLEU	F-BLEU
baseline	11.02	5.17	4.57	3.10	3.55
CVAE (25d)	9.67	4.92	4.45	3.22	3.62
CVAE (50d)	9.78	5.00	4.50	3.34	3.71
CVAE (100d)	10.23	5.15	4.67	3.78	3.98
baseline + RBS	**11.04**	**5.18**	4.56	3.12	3.59
CVAE + RBS	10.39	5.15	**4.69**	**3.82**	**4.01**

Table 4: Corpus-level BLEU-4, sentence-level BLEU-4, and our extensions P-, R-, and F-BLEU computed using the sentence-level metric. For the corpus-level metric, we sample from the CVAE once per context rather than once per reference question.

CVAE 2: who created compositions arrangements ?
CVAE 3: who created created arrangements ?
Du et al. + RBS 1: who performed the ensemble ōshima ?
Du et al. + RBS 2: who performed a ensemble ōshima ?
Du et al. + RBS 3: who performed the ensemble ōshima ?
CVAE + RBS 1: who performed the musical arrangements for the three compositions ?
CVAE + RBS 2: who created the compositions to be performed by michiru ensemble ?
CVAE + RBS 3: who performed musical musical in the three compositions ?

D.4

Context: examples include a concert on 23 march 1833 , in which chopin , liszt and hiller performed -lrb- on pianos -rrb- a concerto by j.s. bach for three keyboards ; and , on 3 march 1838 , a concert in which chopin , his pupil adolphe gutmann , charles-valentin alkan , and alkan 's teacher joseph zimmermann performed alkan 's arrangement , for eight hands , of two movements from beethoven 's 7th symphony .
Du et al.: who performed alkan 's ?

CVAE 1: who performed alkan 's 8th march ?
CVAE 2: who performed alkan 's 8th march ?
CVAE 3: how many movements did joseph play in the concert ?
Du et al. + RBS 1: who performed alkan 's arrangement ?
Du et al. + RBS 2: who performed alkan 's first solo ?
Du et al. + RBS 3: who performed alkan 's arrangement ?
CVAE + RBS 1: who performed a solo concert ?
CVAE + RBS 2: who performed the concert at the march of 1876 ?
CVAE + RBS 3: who performed liszt 's arrangement 23 ?