

A Brief Study on the Effects of Training Generative Dialogue Models with a Semantic loss

Prasanna Parthasarathi^{*,1,4} Mohamed Abdelsalam^{+,2,4} Sarath Chandar^{3,4,5} Joelle Pineau^{1,3,5}

¹ School of Computer Science, McGill University ² University of Montréal

³ École Polytechnique de Montréal,

⁴ Quebec Artificial Intelligence Institute (Mila), ⁵ Canada CIFAR AI Chair

Abstract

Neural models trained for next utterance generation in dialogue task learn to mimic the n-gram sequences in the training set with training objectives like negative log-likelihood (NLL) or cross-entropy. Such commonly used training objectives do not foster generating alternate responses to a context. But, the effects of minimizing an alternate training objective that fosters a model to generate alternate response and score it on semantic similarity has not been well studied. We hypothesize that a language generation model can improve on its diversity by learning to generate alternate text during training and minimizing a semantic loss as an auxiliary objective. We explore this idea on two different sized data sets on the task of next utterance generation in goal oriented dialogues. We make two observations (1) minimizing a semantic objective improved diversity in responses in the smaller data set (Frames) but only as-good-as minimizing the NLL in the larger data set (MultiWoZ) (2) large language model embeddings can be more useful as a semantic loss objective than as initialization for token embeddings.

1 Introduction

Data for language generation tasks in goal-oriented dialogue has semantically diverse samples, where the diversity can be observed from the dialogue topics to the utterances used for getting information on specific slot-values from the user. But, in many niche domains, collecting a large high-quality annotated data set is costly, and often a small data set focused on specific tasks (Wei et al., 2018; Asri et al., 2017) is used for training. This restricts the model to only learn task-specific frequent contexts and seldom learn semantically

similar context due to the lack of sufficient samples (Vinyals and Le, 2015; Serban et al., 2015; Li et al., 2017; Parthasarathi and Pineau, 2018).

Optimizing only on objectives like negative log-likelihood (NLL), and Cross-Entropy (CE) losses foster learning by making the models mimic targets at the token level (Dušek et al., 2020). The models, hence, mostly generate only the observable patterns in the targets in training set (Huang et al., 2017). This can be attributed to the training procedure being uninformative about the semantic similarity of responses. To better understand, consider Target: Would you like to travel to Paris ?, R1: How about Paris as your destination ?, R2: Would she like to read to me ? . R2 has 4 tokens in the same position as in the target but R1 is semantically *similar* to the target. However, the NLL/CE loss for predicting R2 will be lower than predicting R1. This is a common occurrence when training a language generation model, and training on a small data set can exacerbate this issue even further.

Word embeddings from large language models like GloVe (Pennington et al., 2014) , BERT (Devlin et al., 2018) or fastText (Bojanowski et al., 2017) have been shown to have nice properties that preserve some of the linguistic structures (Sinha et al., 2020) that help in understanding semantic and temporal structures in dialogue. We make use of the semantics in the large word embeddings by computing a distance heuristic between the sampled text from model distribution and the target during training. This auxiliary semantic loss ¹ encourages the model in generating sentences that are similar to the target and thereby potentially diversifying the model responses. Although the results are on dialogue generation tasks, the results

* Corresponding author (pparth2@cs.mcgill.ca)

+ Equal authorship

¹<https://github.com/ppartha03/Semantic-Loss-Dialogue-Generation>

are comparable to any broad conditional language generation tasks like caption generation (Vinyals et al., 2015), text summarization (Luhn, 1958) and others (Gatt and Krahrmer, 2018).

Our contributions in the paper are:

- Comprehensively evaluate the proposed semantic loss on two differently sized data sets.
- Show that minimizing a semantic loss on the sampled responses as a training objective improves text generation diversity in limited data setting.
- Show that language model embeddings are useful as semantic loss than word embedding initialization.

2 Conditional Language Generation

In an encoder-decoder architecture, the encoder neural network (Lang et al., 1990) encodes a textual summary of previous utterance exchanges between a user and an agent, H_{i-1} , and the current user utterance u_i . The encoded summary is used by a decoder network to generate the corresponding agent response ($a_i^* = (w_1^i, w_2^i, \dots, w_T^i)$).

Language generation models are mostly trained with NLL objective as defined in Equation 1,

$$\mathbb{L}_{MLE} = - \sum_{t=1}^T \log P(w_t^i | w_{<t}^i, H_{i-1}, u_i) \quad (1)$$

where T is the number of tokens generated in the response (a_i^*), w_t^i is the t -th word in the i -th utterance, and $w_{<t}^i$ denote tokens generated till step t .

3 Semantic Loss

We introduce training with a semantic loss computed with word embeddings from any trained language model. The semantic loss to be minimized is computed in three steps: (1) $a_i^{sampled} = (w_1^i, w_2^i, \dots, w_{T'}^i)$ is generated by sampling tokens from decoder’s distribution over the vocabulary at every step. (2) Average the word vectors of the sampled ($\hat{b}^{a_i^{sampled}}$) and ground truth responses (\hat{b}^{a_i}) with the embeddings from large language models like BERT, GloVe or fastText. Then, compute L2 distance between the two as shown in Equation 2.

$$d_{SEM}^i = || \hat{b}^{a_i^{sampled}} - \hat{b}^{a_i} ||_2 \quad (2)$$

(3) Minimize d_{SEM}^i calculated with the non-differentiable sampling operation, we use REINFORCE (Williams, 1992) to compute \mathbb{L}_{SEM}

(Equation 3).

$$\mathbb{L}_{SEM} = -(-d_{SEM}^i - r(b)) \sum_{t=1}^{T'} \log P(w_t^i) \quad (3)$$

where T' is the number of tokens in $a_i^{sampled}$ and $r(b)$ is the reward baseline computed with average over a moving window of previous rewards to reduce the variance. The model minimizes \mathbb{L}_{Train} as shown in Equation 4.

$$\mathbb{L}_{Train} = \mathbb{L}_{MLE} + \alpha * \mathbb{L}_{SEM} \quad (4)$$

where $\alpha \in \mathbb{R}^+$ is a hyperparameter to specify the strength of the regularization by \mathbb{L}_{SEM} , the optimal value for α depends on the data set. Note: \mathbb{L}_{Train} prefers $R1$ over $R2$ from the example in Section 1, unlike \mathbb{L}_{MLE} .

4 Experiments

We experiment on two differently sized data sets – Frames (Asri et al., 2017) and MultiWoZ 2.0 (Budzianowski et al., 2018) – which are relatively small and large. We compute \mathbb{L}_{SEM} using the commonly used language model embeddings BERT-Base (Devlin et al., 2018), GloVe (Pennington et al., 2014) and fastText (Bojanowski et al., 2017) to compare the benefit of using different embeddings.

Evaluation Metrics: We measure the performance on overlap based metric BLEU (Papineni et al., 2002); and diversity in the generated text by computing the fraction of distinct-1 and distinct-2 grams, similar to Welleck et al. (2019); Li et al. (2015), on validation set. Also, as a proxy to evaluate generalization to generating n-grams that the decoder was never trained to, we measure the fraction of bigrams generated by the model during validation that were not in the training targets, as % *Unseen*. Also, to measure the effects of minimizing the semantic loss on language quality, we perform human evaluation for comparing the different training techniques. Further we compare the improvements in diversity between using BERT for initialization of word embeddings and using it in a semantic loss objective.

4.1 Quantitative Evaluation

Experimental result in Figure 1(a) shows that performance of the model trained with \mathbb{L}_{Train} decreases on the overlap based metric, BLEU. This is explained by the \mathbb{L}_{Train} trained models, with

greedy decoding, generating a greater fraction of unique bigrams (Figure 1(b)) on the validation set than the \mathbb{L}_{MLE} trained model: measured with metrics distinct-1 and distinct-2 (Li et al., 2015). As the model learns to discover semantically similar bigrams, the performance on overlap based metric decreases. Further, % Unseen metric measured in Figure 1(c) shows that \mathbb{L}_{Train} fosters generation of new bigrams.

In the experiments, we observed a pattern of % Unseen spiking at regular intervals, indicating that the loss helped the model to periodically discover newer bigrams, which increased the NLL in training as the syntax around the bigram has to be re-learned by minimizing the now higher NLL objective. This is different from beam search as beam

USER : of those 3 options , i would prefer 11 days any other hotel options can you check if there are other hotel options for september 1 - 20 ? what are the departure and return dates for this.

TARGET : sept 13th through the 19th.

SEMANTIC BEAM1 : i ' m sorry i have nothing from santiago .

SEMANTIC BEAM2 : i ' m sorry i have nothing from santiago . is there another destination and would you be interested

SEMANTIC BEAM3 : i ' m sorry i have nothing from santiago . is there another destination ?

SEMANTIC BEAM4 : i ' m sorry i have nothing from santiago . is there another destination you would like to go

SEMANTIC BEAM5 : i ' m sorry i have nothing from santiago . is there another destination you would like to be

USER : of those 3 options , i would prefer 11 days any other hotel options can you check if there are other hotel options for september 1 - 20 ? what are the departure and return dates for this.

TARGET : sept 13th through the 19th.

BEAM 1 : i can i do not have to help , sorry , i sorry , sorry , i sorry ,

BEAM 2 : i can i do n't have to help , sorry , i sorry , i sorry , i sorry

BEAM 3 : i can i do not have to help , sorry , i sorry , i sorry , i sorry

BEAM 4 : i can i do not have for that sorry , i sorry , i sorry , i sorry ,

BEAM 5 : i can i do not have to help , sorry , i sorry , i sorry , sorry ,

Table 1: Comparing the diversity in beam search between the model trained with \mathbb{L}_{Train} (top) and with \mathbb{L}_{MLE} (bottom)

sampling conforms to the distribution learnt with

USER : i will also need a taxi to pick me up by 24:30 . i need the contact number and car type please.

BEST BLEU : i have booked you a yellow lexus . the contact number is 07346991147.

DIVERGED : okay pull d assisting joining botanic gardens , good and good bye.

Table 2: Aggressively exploring with dropping larger fraction of tokens in a sentence lead to divergence in language generation in MultiWoZ as shown.

\mathbb{L}_{MLE} , whereas \mathbb{L}_{Train} allows to learn a distribution that allows learning to use valid alternatives in the training. This allows a better beam search, as shown in the example Table 1.

4.2 BERT Initialization vs BERT Semantic loss

We construct 4 different models by combining the two different loss functions (Loss1: \mathbb{L}_{MLE} , Loss2: \mathbb{L}_{Train}) with two different initializations (Init1: random, and Init2: BERT) for the word embeddings. Diversity measured with *distinct-2* (Figure 1(d)) showed that *Init1;Loss2* model showed greater improvements than *Init2;Loss1* or *Init2;Loss1*. The result suggests that BERT can be more useful in \mathbb{L}_{Train} than embedding initialization. This could be reasoned by the strong regularization enforced by the word embedding that is unyielding to exploration in generating sequences in addition to the \mathbb{L}_{MLE} objective.

4.3 Negative Result in MultiWoZ

We observed that the model trained with \mathbb{L}_{Train} performed only as good as training with \mathbb{L}_{MLE} on our defined evaluation metrics (Figure 1(e),1(f)) in MultiWoZ. The overlap based metric and unique bigrams generated did not have as much improvement as it had in Frames data set (Figures 1(b), 1(f)).

To overcome this issue, during training, we increased the model's exploration to newer tokens by masking tokens in the decoder output at random before sampling a response. This helped the model in discovering newer bigrams eventually. This technique generated larger fraction of unseen bigrams but the randomness in dropping tokens generated more noise in the text generated (Table 2). Making the random exploration useful with additional constraints to keep the syntax from diverging is a potential future work.

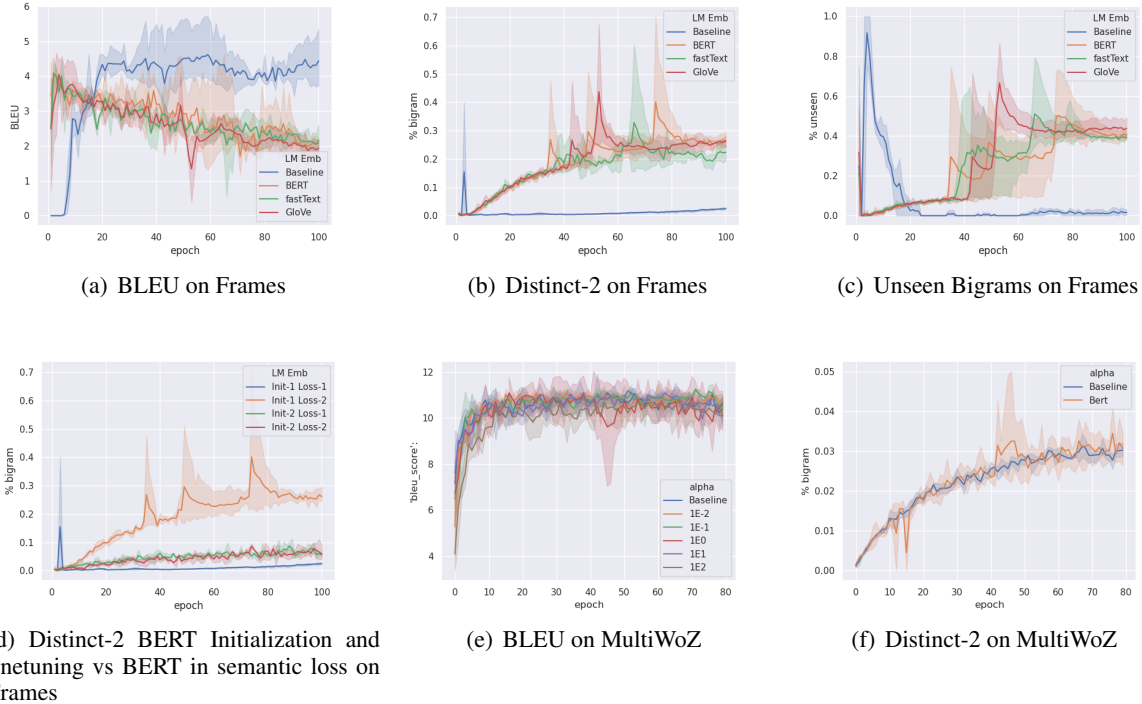


Figure 1: Quantitative comparison of different loss functions, and initialization on Frames and MultiWoZ data sets.

4.4 Human Evaluation

We perform two human studies (Appendix B.2) with two sets of 100 randomly sampled contexts from test set of Frames data set with 3 scorers per pair. In Study 1, the volunteers were

Metric	% Wins	% Losses	% Ties
<i>Diversity</i>	65	16	19
<i>Relevance</i>	45	38	17

Table 3: **Study 1:** %Wins denote the #times the scorers picked *Init1;Loss2*'s response and %Loss is when it was the *Init1;Loss1*'s response.

shown the responses generated with *Init1;Loss1* and *Init1;Loss2*. Like in (Li et al., 2015), we ask the volunteers to select the one that is *relevant* to the context, and the one that is *interesting/diverse* in two separate questions. We allow ties in both

Metric	% Wins	% Losses	% Ties
<i>Diversity</i>	63	24	13
<i>Relevance</i>	41	31	28

Table 4: **Study 2:** %Wins denote the #times the scorers picked *Init1;Loss2*'s response and %Loss is when scorers picked the *Init2;Loss1*.

the questions. In Study 2, we compare *Init2;Loss1* and *Init1;Loss2* with questions as in Study 1.

The results of Study 1 and Study 2 shown in Table 3 and 4 show that, despite the lower BLEU scores, minimizing \mathbb{L}_{Train} indirectly fosters diversity in responses; human scorers found the model trained with the proposed semantic loss objective to be diverse/interesting on an average of 65% and 63% in studies 1 and 2 respectively. This verifies again in a different experiment that BLEU scores do not correlate well with human scores (Liu et al., 2016). The regularization from the BERT initialization is not promoting diversity which, from the experiments, depends on minimizing the semantic objective. The relevance of the response is not significantly higher than the baseline, which was expected as the semantic loss was expected only to improve the diversity.

5 Conclusion

Training with a semantic loss has a positive effect in a smaller data set and that reflects on the model's improvement in diversity measuring metrics. But, the semantic loss was not very effective in a large data set due to the lack of diversity within and a hard bias dictated by the samples in the data set. The results obtained in the paper

shows that training with semantic loss can be effective in low data setting.

Acknowledgements

We would like to acknowledge Compute Canada and Calcul Quebec for providing computing resources used in this work. The authors would also like to thank members of Chandar Research Lab, Mila for helping with the code reviews and reviewing the manuscripts. Sarath Chandar and Joelle Pineau are supported by Canada CIFAR AI Chair, and Sarath Chandar is also supported by an NSERC Discovery Grant.

References

- Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. Frames: A corpus for adding memory to goal-oriented dialogue systems. *arXiv preprint*.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *TACL*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Ultes Stefan, Ramadan Osman, and Milica Gašić. 2018. Multiwoz - a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of EMNLP*.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2020. Evaluating the state-of-the-art of end-to-end natural language generation: The e2e nlg challenge. *Computer Speech & Language*.
- Albert Gatt and Emiel Kraemer. 2018. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*.
- Gabriel Huang, Hugo Berard, Ahmed Touati, Gauthier Gidel, Pascal Vincent, and Simon Lacoste-Julien. 2017. Parametric adversarial divergences are good task losses for generative modeling. *arXiv preprint*.
- Kevin J Lang, Alex H Waibel, and Geoffrey E Hinton. 1990. A time-delay neural network architecture for isolated word recognition. In *Neural networks*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Jiwei Li, Will Monroe, Tianlin Shi, Alan Ritter, and Dan Jurafsky. 2017. Adversarial learning for neural dialogue generation. In *arXiv preprint*.
- Chia-Wei Liu, Ryan Lowe, Iulian Vlad Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of EMNLP*.
- Hans Peter Luhn. 1958. The automatic creation of literature abstracts. *IBM Journal of research and development*.
- 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 on association for computational linguistics*. Association for Computational Linguistics.
- Prasanna Parthasarathi and Joelle Pineau. 2018. Extending neural generative conversational model using external knowledge sources. In *Proceedings of EMNLP*.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP*.
- Iulian Vlad Serban, Ryan Lowe, Laurent Charlin, and Joelle Pineau. 2015. A survey of available corpora for building data-driven dialogue systems. In *arXiv preprint*.
- Koustuv Sinha, Prasanna Parthasarathi, Jasmine Wang, Ryan Lowe, William L Hamilton, and Joelle Pineau. 2020. Learning an unreferenced metric for online dialogue evaluation. In *Proceedings of ACL*.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint*.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Proceedings of the CVPR*.
- Zhongyu Wei, Qianlong Liu, Baolin Peng, Huaixiao Tou, Ting Chen, Xuan-Jing Huang, Kam-Fai Wong, and Xiang Dai. 2018. Task-oriented dialogue system for automatic diagnosis. In *Proceedings of ACL*.

Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019. Neural text generation with unlikelihood training. *arXiv*.

Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*.

A Training and hyperparameters

- We used a 128 unit hidden size LSTM with a 128 unit input embedding dimension.
- The range of the α we tested in log-scale is $[-2, 2]$. And, the best alpha selected based on the early saturation of distinct-2 was $1E-1$ and used this for experiments in different language model embeddings used for computing \mathbb{L}_{SEM} .
- We use Adam optimizer with $4E-3$ as learning rate and other parameters as default.
- For the choice of word embeddings, we used 300 dimensional GloVe and fastText, and 768 dimensional BERT-Base.
- For REINFORCE with baseline, we computed the average for the last 20 samples as the baseline.
- We averaged the results over 5 different seeds. For the baseline model, we chose the best performing seed with respect to BLEU score and for the model trained with \mathbb{L}_{Train} based on early saturation on distinct-2 on the validation set for human evaluation.

B Frames Experiments

B.1 Word repeats

Evaluating generalization to unseen bigrams is tricky as there can be potentially many word repeats. To not count that, we looked at the fraction of bigrams that were word repeats, one of the most common errors by language generation models (Figure 2).

The result showed two interesting things: First, the word repeats are minimal but does happen when training with semantic loss, though the gain of discovering unseen bigrams is more useful. Second, the NLL trained model initially generates many word repeats along with a few unseen tokens and they both die down due to the strong MLE objective that overfits to the targets in the training.




Figure 2: Comparing d_{SEM} with different word embeddings on fraction of bigrams generated on the validation set that are word repeats on the Frames data set.

B.2 Human Evaluation

For human evaluation, we asked for English speaking graduate students as volunteers to take part in the two studies. To reduce the cognitive load on individual participants, we split the 100 samples in 4 sets of 25 samples. We computed the inter-annotators agreement with cohen-kappa coefficient (Cohen, 1960) in the sklearn package (Pedregosa et al., 2011).

	Q1:Relevance	Q2:Diversity
Study 1	0.28	0.22
Study 2	0.33	0.23

Table 5: Average of cohen kappa score averaged over the evaluation of annotators on the different sets of samples in the two studies.

The results shown in Table 5 that the annotators had a fair agreement in the two studies. The range of the scores is between -1 and 1, and a score above 0 indicates agreement amongst the annotators. The slightly lower agreement on Q2 is because of the ambiguity in the perception of "what is interesting".

C MultiWoZ Experiments

C.1 Negative Result

We observed that the semantic loss was not as useful as it was in the smaller data set. The bigram distribution of the two data sets (Table 6 and 7) showed that the bigrams in the context on an average occurs 92 times in MultiWoZ as compared to only 17 times in Frames. Similarly, a bigram in the target occurs 13 times in MultiWoZ compared to only 5.4 times in Frames.

From the analysis on the distribution of bigrams in the two data sets, we arrived at the following conjecture: With a simplistic assumption, consider the following sentences: I want to leave from London, I want to leave on Tuesday, I want to leave from Florida occur 3, 2, and 5 times respectively in a small data set and 30, 20, and 50 times in a relatively larger data set. The language model of the decoder, after generating I want to leave, will sample one of the three bigrams, on Tuesday, to London, from Florida.

data set	Unique Bigrams	Total Bigrams
<i>Frames</i>	30K	0.5M
<i>MultiWoZ</i>	40K	3.6M

Table 6: Count of Bigrams from only the contexts of the two data sets

data set	Unique Bigrams	Total Bigrams
<i>Frames</i>	22K	127k
<i>MultiWoZ</i>	71K	900k

Table 7: Count of Bigrams from only the targets of the two data sets

The output of the encoder-decoder at every step being a multinomial distribution over the vocabulary, the architecture can be abstracted for our understanding to maintain a Dirichlet distribution that is generalizable.

The bias of sampling from Florida is much higher in a large data set and relatively much lower in a smaller data set, which can even generate I want to leave from Florida to London on Tuesday with a relatively higher probability. As sampling from the decoder is still dependent on \mathbb{L}_{MLE} , the diversity in sampling is decreased when training with NLL on a large data set.

But then, as the larger data set has 7 times more support for a bigram than in the smaller data set, out of distribution sampling is difficult.

C.2 Out-of-NLL Sampling

To break the rigid sampling distribution, with a non-zero probability we dropped words from the vocabulary before sampling the tokens in $a_i^{sampled}$.

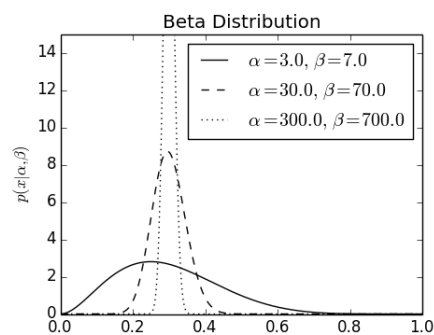


Figure 3: Beta distribution with differently scaled α and β values. The lower values correspond to the smaller data sets and the higher values correspond to the larger data sets.

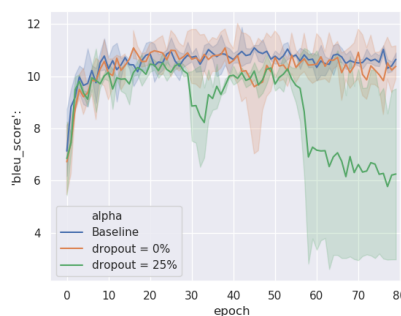


Figure 4: Effect of dropout on automatic evaluation metric. The drop in BLEU is due to the model generating newer bigrams.

With the semantic loss providing non-binary scores, the model gets feedback for all sampled responses, even those unlikely to be sampled but are sampled due to the masking of the vocabulary. That lead to a sharp divergence of training (Table 2) even before the model learnt to appropriately diversify its responses (Figure 5).

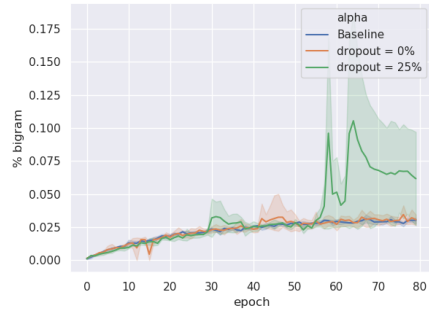


Figure 5: The distinct-2 grams when trained with dropout and random substitution indeed helped the model to sample out-of-NLL distribution. But, the overwhelming noise diverged the training and the model responses degenerated.

The % *unseen* and distinct-1 and 2 scores keep increasing (Figures 5) but due to the high amount of diversity in the tokens generated, many of the responses were not legible as seen in Table 2.