

Generating More Interesting Responses in Neural Conversation Models with Distributional Constraints

Ashutosh Baheti¹, Alan Ritter¹, Jiwei Li^{2,3}, and Bill Dolan³

¹Ohio State University, OH, USA

²Shannon.AI and Renmin University of China, Beijing, China

³Microsoft Research, Redmond, WA, USA

{baheti.3, ritter.1492}@osu.edu

jiweil@stanford.edu

billdol@microsoft.com

Abstract

Neural conversation models tend to generate safe, generic responses for most inputs. This is due to the limitations of likelihood-based decoding objectives in generation tasks with diverse outputs, such as conversation. To address this challenge, we propose a simple yet effective approach for incorporating side information in the form of distributional constraints over the generated responses. We propose two constraints that help generate more content rich responses that are based on a model of syntax and topics (Griffiths et al., 2005) and semantic similarity (Arora et al., 2016). We evaluate our approach against a variety of competitive baselines, using both automatic metrics and human judgments, showing that our proposed approach generates responses that are much less generic without sacrificing plausibility. A working demo of our code can be found at <https://github.com/abaheti95/DC-NeuralConversation>.

1 Introduction

Recent years have seen growing interest in neural generation methods for data-driven conversation. This approach has the potential to leverage massive conversational datasets on the web to learn *open-domain* dialogue agents, without relying on hand-written rules or manual annotation. Such response generation models could be combined with traditional dialogue systems to enable more natural and adaptive conversation, in addition to new applications such as predictive response suggestion (Kannan et al., 2016), however many challenges remain.

A major drawback of neural conversation generation is that it tends to produce too many “safe” or generic responses, for example: “*I don’t know*” or “*What are you talking about?*”. This is a pervasive problem that has been independently reported

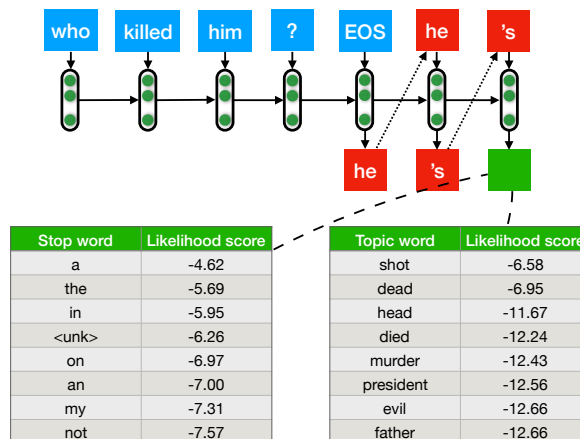


Figure 1: Illustration of the dull response problem in maximum likelihood neural conversation generation using an example from the OpenSubtitles corpus. Function (stop) words tend to receive higher log probabilities than content (topic) words. The highest likelihood stop words and topic words in this context are listed.

by multiple research groups (Li et al., 2016a; Serban et al., 2016; Li et al., 2016c).¹ The effect is due to the use of conditional likelihood as a decoding objective – maximizing conditional likelihood is a suitable choice for text-to-text generation tasks such as machine translation, where the source and target are semantically equivalent, however, in conversation there are many acceptable ways to respond. Simply choosing most predictable reply often leads to very dull conversation.

Figure 1 illustrates the problem with conditional likelihood using an example. After encoding the source message using a bidirectional LSTM with attention, and fixing the first two words of the response, we show the highest ranked words (according to log-likelihood scores) taken from a list

¹<https://research.googleblog.com/2015/11/computer-respond-to-this-email.html>

of stop words² in contrast to those selected from a list of topic words.³ As illustrated in the figure, response generation that is based on maximum likelihood is biased towards stop-words and therefore results in responses that are *safe* (likely to be plausible in the context of the input), but also *bland* (don't contribute any new information to the conversation). This motivates the need for augmenting the decoding objective to encourage the use of more content words.

To address the *dull-response* problem in neural conversation, in this paper, we propose a new decoding objective that flexibly incorporates side-information in the form of distributional constraints. We explore two constraints, one which encourages the distribution over topics and syntax in the response to match that found in the user's input. To estimate these distributions, we leverage the unsupervised model of topics and syntax proposed by Griffiths and Steyvers (2005). The second constraint encourages generated responses to be semantically similar to the user's input; semantic similarity is measured using fixed-dimensional sentence embeddings (Arora et al., 2016).

After introducing distributional constraints into the decoding objective, we empirically demonstrate, in an evaluation that is based on human judgments, that our approach generates more content-rich responses when compared with two competitive baselines: Maximum Mutual Information (MMI) (Li et al., 2016a), in addition to an approach that conditions on topic models as additional context in neural conversation (Xing et al., 2017). While encouraging the model to generate less bland responses can be risky, we find that our approach achieves comparable plausibility while introducing significantly more content.

2 Neural Conversation Generation

As a starting point for our approach we leverage the Seq2Seq model (Sutskever et al., 2014; Bahdanau et al., 2014) which has been used as a basis for a broad range of recent work on neural conversation (Kannan et al., 2016; Li et al., 2016a; Serban et al., 2016; Shao et al., 2017). This model consists of two parts, an *encoder* and a *decoder* both of which are typically stacked LSTM layers. The encoder reads the input sequence and creates

²<https://www.ranks.nl/stopwords>

³The top 10 topic words were taken from each of the 50 topics inferred by an HMM-LDA model (after removing stop words).

a hidden representation. The decoder conditions on this representation, using attention, and generates the response using a neural network language model (Bengio et al., 2003; Sutskever et al., 2011).

3 Distributional Topic and Semantic Similarity Constraints

Neural generation models select a response, \hat{Y} by maximizing over a decoding objective, typically using greedy beam search from left to right over partially completed responses, which are scored using the decoder RNN language model. A commonly used decoding objective is the conditional likelihood of the target given the source, $P(Y|X)$:

$$\begin{aligned}\hat{Y} &= \arg \max_Y \{\log P(Y|X)\} \\ &= \arg \max_{w_1, \dots, w_n} \left\{ \sum_{i=1}^n \log P(w_i | w_1, \dots, w_{i-1}, X) \right\}\end{aligned}\quad (1)$$

As discussed in Section 1, models trained to maximize conditional likelihood tend to assign low probability to content words as compared to (more frequent) function words, leading to bland, generic responses most of the time. To ameliorate this, we introduce distributional constraints in the form of additional terms in the decoding objective that favor hypotheses containing more content words that are similar to the source in the *Topical* and *Semantic* sense.

For the constraint in the topic domain, we are interested in the topic probability distributions of the source, X , and target Y , $P(T|X)$ and $P(T|Y)$, where T is a random variable defined over k topics. Then we can modify the decoding objective from Eq 1:

$$\begin{aligned}\hat{Y}^T &= \arg \max_Y \{ \log P(Y|X) + \\ &\quad \alpha \times \Delta(P(T|X), P(T|Y)) \}\end{aligned}\quad (2)$$

Here, Δ is a similarity function between the two probability distributions and α is a tunable hyperparameter to adjust impact of this constraint.

Much recent work has investigated how to encode the semantic meaning of a sentence into a fixed high dimensional embedding space (Kiros et al., 2015; Wieting and Gimpel, 2017). Given such an embedding representation of X and Y , one can find the semantic similarity between the two and similar to Eq 2 we can add a semantic similarity constraint to the likelihood objective as

follows:

$$\hat{Y}^{Emb} = \arg \max_Y \{ \log(P(Y|X)) + \beta \times \nabla(Emb(X), Emb(Y)) \} \quad (3)$$

where, $Emb()$ is a function that maps an utterance to a semantic vector representation, ∇ is a function that computes similarity of the two embeddings and β is a tunable parameter.

Both of the constraint terms from Eq 2 and Eq 3 are additive in nature and thus can be combined in a straightforward fashion. This formulation allows us to systematically combine information from three different models to produce better responses in terms of topic and semantic relevance. Conceptually, the likelihood term governs the grammatical structure of the response while the topic and semantic constraints drive content selection (Nenkova and Passonneau, 2004; Barzilay and Lapata, 2005).

4 Decoding with Distributional Constraints

In Section 3, we defined two constraints (one topic constraint and one semantic) for use in the decoding objective. Incorporating these constraints during decoding requires that they factorize in a way that is compatible with left-to-right beam search over words in the response. The standard approach to computing posterior distributions in topic models requires a probabilistic inference procedure over the entire source and target. Furthermore, computing semantic representations can involve the use of complex neural architectures. Both of these procedures are difficult to integrate into decoding, because they are computationally expensive and would need to be called repeatedly within the inner loop of the decoder. Furthermore, when performing left-to-right beam search, as is common practice in neural generation, the complete response is generally not available. To address these challenges, we propose using simple additive variants of these methods that factorize over words and which we found to enable efficient decoding without sacrificing performance.

4.1 Topic Similarity

Estimating the topic distribution of the source, $P(T|X)$, and response, $P(T|Y)$, is a key step in implementing the topic-similarity constraint. HMM-LDA is a generative model that is able to separate topic and syntax words, by inferring topic

distributions in a corpus while flexibly modeling function words. We briefly summarize this model before describing our implementation.

4.1.1 Syntax-Topics model

Griffiths et. al. (2005) suggested an unsupervised generative model that simultaneously labels each word in a document with a syntax (c) and topic (z) state. They modify the Latent Dirichlet Allocation model to include a syntactic component akin to a Hidden Markov Model (HMM). In LDA, each topic (z) is associated with a probability distribution over the vocabulary $\phi^{(z)}$. HMM-LDA adds additional distributions over words for each syntactic class (c) as $\phi^{(c)}$. A special class, $c = 0$, is reserved for topics. The transition model between classes c_{i-1} to c_i follows a multinomial distribution $\pi^{(c_{i-1})}$. Each document has an associated distribution over topics $\theta^{(d)}$; each word, w_j , in the document has an associated latent topic variable, z_j , that is drawn from $\theta^{(d)}$ and c_j is drawn from $\pi^{(c_{j-1})}$. If $c_j = 0$, then w_j is drawn from $\phi^{(z_j)}$, otherwise it is drawn from $\phi^{(c_j)}$. Markov Chain Monte Carlo inference (MCMC) is used to infer values for the hidden topic and syntax variables associated with a given document collection. To estimate topic and syntax distributions, we performed collapsed Gibbs sampling over our training corpus of conversations, where each conversation is treated as a document. One sample of the hidden variables was used to estimate model parameters after 2,500 iterations of burn in. Our code for training the HMM-LDA model is [available online](#)⁴.

4.1.2 Estimating Topic Distributions with HMM-LDA

To compute distributional topic constraints in neural response generation, we first need an efficient method for estimating topic distributions that factorizes over words, given a point estimate of an HMM-LDA model's parameters. We would like to estimate topic distributions based on content words contained in a sentence and ignore function words. HMM-LDA provides us with topic, $\phi^{(z)}$, and syntax, $\phi^{(c)}$, distributions over the vocabulary of words, $w \in V$. Treating a sentence as a bag-of-words we can estimate its distribution over topics as a sum of topic distributions over all words normalized by sentence length. However, we found this approach does not to work well in

⁴<https://github.com/abaheti95/HMM-LDA>

practice because it gives equal weight to topic and syntax words. To address this issue, we weighted each word’s topic distribution $P(T|w)$ by its probability of being generated by the topic component of the HMM-LDA model (i.e. $P(C = 0|w)$). The topic distribution of a sentence, S , is estimated as:

$$P(T|S) = \frac{1}{Z} \sum_{w \in S} P(T|w) P(C = 0|w) \quad (4)$$

where $Z = \sum_{w \in S} P(C = 0|w)$ is a normalizing constant that corresponds to the expected number of content words in the sentence. As mentioned earlier, a more accurate estimate of the topic distribution could be obtained using MCMC inference or by applying the forward-backward algorithm. However, these methods are computationally expensive and not well-suited to the decoding framework used in neural generation.

The method described above allows us to efficiently compute the topic distribution of a sentence for use in the topic constraint in Eq 2. For a similarity function, Δ , we simply use the vector dot product, which is closely related to cosine similarity. This formulation has the advantage that it enables memoization during decoding. Another advantage is that it captures the ratio of topic to syntax words due to the weights $P(C = 0|w)$.⁵ Therefore, the overall constraint has the effect of keeping the syntax-topics ratio in generated hypothesis similar to the source.

4.2 Semantic Similarity

To define the semantic similarity constraint we first encode a semantic representation of the source and target into a fixed dimensional embedding space. There are many sentence embedding methods that could be used, however we want this encoding to be relatively efficient as it will be used many times during beam search.

Arora et. al. (2016) recently proposed a simple sentence embedding method, which was shown to have competitive performance across a variety of tasks. Their approach uses a weighted average of word embeddings where each word is weighted by $\frac{a}{a+P(w)}$; here, $P(w)$ is the unigram probability and a is a hyperparameter. Such a weighting scheme reduces the impact of frequent words (typ-

⁵Assuming topic distribution of syntax words to be uniform, a sentence with more syntax words will dampen modes in the distribution. Alternately, with less syntax words the overall distribution will be more peaked.

ically function words) in the overall sentence embedding. Next the first principal component of all the sentence embeddings in the corpus is removed. (Arora et al., 2016) points that the first principal component has high cosine similarity with common function words. Removing this component gives sentence embeddings that encapsulate the semantic meaning of the sentence. We use this technique in our implementation of $Emb()$ in Eq 3. For the similarity function, Δ , we use the dot product. Analogous to the topic constraint described above, this approach to measuring semantic similarity also decomposes over words and works well in the decoding framework.

Parameter	Value
RNN Type	Bi-LSTM
Layers	4
Hidden layer dim.	1000
Learning rate	0.1
max. grad. norm.	1
Optimization	Adadelta
Parameter Init	(-0.08, 0.08) (uniform)

Table 1: Hyperparameter setting for training

Bucket	#dialogues	#test
b1 (3-6 words)	10994	334
b2 (7-15 words)	15794	333
b3 (16-25 words)	5167	333
total	31955	1000

Table 2: Test set from Cornell Movie Dialogue Corpus. Column 2 shows the total number of dialogues that we got after all pre-processing and Column 3 shows the number of sampled dialogues in the test set.

5 Datasets

For training purposes we use OpenSubtitles (Tiedemann, 2009), a large corpus of movie subtitles (roughly 60M-70M lines) that is freely available and has been used in a broad range of recent work on data-driven conversation. OpenSubtitles does not contain speaker annotations on the dialogue turns, so as previously noted when used for learning data-driven conversation models the data is somewhat noisy. Nonetheless, it is possible to create a useful corpus of conversations from this data by assuming each line corresponds to a full speaker turn. Although this assumption is often violated, prior work has successfully trained

and evaluated neural conversation models using this corpus. In our experiments we used a preprocessed version of this dataset distributed by Li et. al. (2016a).⁶ The dataset contains large number of two turn dialogues out of which we sampled 23M to use as our training set and 10k as a validation set.

Due to the noisy nature of the OpenSubtitles conversations we do not use them for evaluation. Instead, we leverage the Cornell Movie Dialogue Corpus (Danescu-Niculescu-Mizil and Lee, 2011) which is much smaller but contains accurate speaker annotations. We extracted all two turn conversations (source target pair) from this corpus and removed those with less than three and more than 25 words. After this, we divided the remaining conversations into three buckets based on source length. The numbers can be found in Table 2. From each bucket we randomly sampled ≈ 333 dialogues for a total of 1000 dialogues in our test set. We evaluate all models on this test set. Since automatic metrics do not correlate with human judgment, we manually tuned the hyperparameters (α and β) on a small development set (4 dialogues from each bucket to create a small 12 sentence development set; disjoint from test set). We manually inspected the responses generated by the model on the development set for different values of α and β and choose those that performed best.

6 Experimental Conditions and Baselines

During learning we use the same hyperparameters for all models; these are displayed in Table 1, and are based on those reported by Li et. al. (2016a).⁷ We compare our approach with the following baselines:

MMI: We re-implemented the MMI-bidi method proposed by Li et. al. (2016a). MMI is a particularly appropriate baseline for comparison, as it encourages responses that have higher relevance to the input in contrast to conditional likelihood, which tends to favor responses with higher unconditional probability. MMI-bidi generates B candidates using Beam search on a Seq2Seq model trained to maximize conditional likelihood of the target given the source, $P(Y|X)$, then re-ranks them using a separately trained source given tar-

get model, $P(X|Y)$. Combining both directions in this way has the effect of maximizing mutual information (Li et al., 2016a).

TA-Seq2Seq: Another relevant baseline is the TA-Seq2Seq model of Xing et. al. (2017) that integrates information from a pre-trained topic model into neural response generation using an attention mechanism to condition on relevant topic words. They evaluate their model on a dataset of Chinese forum posts. Unfortunately we could not use the code provided by the authors due to data-mismatch (their model makes use of user identities which are not available in the OpenSubtitles corpus). We therefore compare with a re-implementation of their approach in which we modify each source sentence to include a list of the 20 most relevant topic words from HMM-LDA and then train using the same Seq2Seq framework with attention. This enables the model to condition on the relevant topic words. In addition to incorporating attention over topics, Xing et. al. also introduced an approach to *biased generation* - to replicate this we add a constant factor to all topic words during the prediction.

7 Results and Analysis

Our proposed decoding objective constraints (topic and semantic) are complementary to the MMI objective, which encourages diversity and relevance to the source input. Therefore, in addition to comparing against the baselines described above, we evaluated three variants of our model: (1) maximum conditional likelihood combined with semantic and topic distributional constraints with a beam size of 10 (DC-10) (2) The same configuration with MMI-bidi re-ranking using a beam size of 10 DC-MMI10 and (3) MMI-bidi re-ranking with a beam size of 200 (DC-MMI200). We test all configurations on the 1000 conversations test set described in Section 5 and compare them on automatic metrics and also in a crowd-sourced human evaluation. We do not consider TA-200 (TA-Seq2Seq, Beam=200), DC-200 and MMI-10 for human evaluation as they appear to perform worse than other model variants in automatic metrics and also on our set of development sentences. Sample responses for all the remaining models are presented in Table 3.

7.1 Automatic Metrics

Following Li et. al. (2016a), we report distinct-1 and distinct-2, which measure the diversity of re-

⁶<http://nlp.stanford.edu/data/OpenSubData.tar>

⁷OpenNMT is used for training our models (Klein et al., 2017).

Source	Target (ground truth)	MMI	Approximate TA-Seq2Seq	Distributional Constraints
where does he work ?	in there , sir .	does he know where he is ?	what kind of place is this ?	he has work to do at the university
here 's your jacket !	uh , thanks ...	i don 't want it !	what kind of suit is this ?	you can stick your shirt in it !
what 's so damn funny ?	been to any good <unk>lately ?	what are you laughing at ?	what 's the matter with you ?	it 's sort of a funny thing you know ?
well , what exactly does our platoon do ? serve <unk>? process paperwork ?	our platoon has the best assignment of all . we 're the first into battle !	well as a matter of fact	the best thing to do ?	well the war is over but we don 't have time for this shit
left us here to rot . that 's what they 've done . heroes of the newspapers !	all right , all right . keep quiet .	they 'r e all dead !	it 's time to fight !	we are brothers and kings of war !
that 's it . you 're not setting foot off this ship until we return home .	but i haven 't even shown you my biology project ...	we 'r e not going anywhere	i hope you don 't mind	we have orders that the ship will return to earth immediately
i 'm han solo , captain of this vessel . who 's in charge then ?	i 'm ben kenobi. luke starkiller here is leading our expedition .	i don 't know sir sir	you 'r e a coward !	i am captain kirk commander of the ship
her grandmother said she 'd been threatening to run away . and i found the car at the miami bus terminal .	you don 't think sam could 've put it there ?	but she didn 't tell me	that 's one hell of a job	it was parked in a car crash near the road

Table 3: Sample responses of all the models on the dev set

sponses. These are the ratios of types to tokens for unigrams and bigrams, respectively. We also report BLEU-1 scores following previous work, however it should be noted that BLEU-1 is not generally accepted to correlate with human judgments in conversation generation tasks (Liu et al., 2016) as there are many acceptable ways to reply to an input which may not match a reference response. Lastly, we compare the percentage of **stop-words**⁸ of the responses generated by each model (smaller values, that are closer to the distribution of human conversations are preferred). The automatic evaluation is presented in Table 4.

⁸Long Stopword List from <https://www.ranks.nl/stopwords>. We appended punctuations to this list.

For brevity we define aliases for each system in the 2nd column of Table 4 which are used in subsequent discussion. The human responses are diverse and also generally longer than automatically generated responses. MMI200 has higher diversity than TA-Seq2Seq in terms of distinct-1 and distinct-2. This illustrates the importance of re-ranking using MMI. Our approach produces almost twice as many distinct unigrams and bigrams. We also observe MMI200 and TA-Seq2Seq achieve higher BLEU scores than our models, however this is not surprising since our models are designed to generate more interesting responses containing rarer content words that are less likely to appear in reference responses.

Model	Alias	distinct-1	distinct-2	BLEU -1	Avg. length	Stop- word%
Human responses	human	2381/0.176	7532/0.602	-	13.5	70.66
MMI (Beam=200)	MMI200	351/0.058	990/0.197	12.8	6.0	84.91
TA-Seq2Seq(Beam=10)	TA-10	237/0.036	524/0.095	12.9	6.5	79.40
Dist. Const. (Beam=10)	DC-10	710/0.097	2014/0.320	11.0	7.3	72.04
Dist. Const. + MMI (Beam=10)	DC-MMI10	732/0.099	2098/0.327	11.4	7.4	73.87
Dist. Const. + MMI (Beam=200)	DC-MMI200	850/ 0.116	2946/ 0.465	11.6	7.3	72.25

Table 4: Automatic metrics evaluation. The 3rd and 4th columns show the ratio of types to tokens for unigrams and bigrams respectively. 7th Column shows the % of stop-words generated by the models in their responses.

Model Alias	No(%)	Unsure(%)	Yes(%)
Plausible?			
human	19.807	23.448	56.745
MMI200	27.623	26.445	45.931
TA-10	26.981	26.874	46.146
DC-MMI200	30.835	24.41	44.754
Content Richness?			
human	16.488	19.914	63.597
MMI200	23.662	32.976	43.362
TA-10	31.799	30.086	38.116
DC-MMI200	20.021	26.660	53.319

Table 5: Human judgments for Plausibility of the different models. Each numerical cell contains a percentage value corresponding to its row truncated to 2 decimal precision.

As expected we observe that MMI200 and TA-10 have a higher percentage of stop-words than human responses. According to the human evaluation discussed in Section 7.2, these models were also found to have lower content richness.

7.2 Human Evaluation

We conducted a survey on the crowd-sourcing platform, Amazon Mechanical Turk. Every model response is scored on 2 categories: 1) **Plausibility** - *is the response plausible for the given source?* and 2) **Content Richness** - *does the response add new information to the conversation?* We asked the evaluators to respond on a 5-point scale to the questions above (Strongly Agree, Agree, Unsure, Disagree, Strongly Disagree). These were later collapsed to 3 categories (Agree, Unsure, Disagree). The results for plausibility and content richness of our model in addition to the MMI and TA-Seq2Seq baselines and human responses are presented in Table 5.

We observe that MMI200 and TA-10 models

Model Alias	No (%)	Unsure (%)	Yes (%)
Plausible?			
DC-10	36.617	27.944	35.439
DC-MMI10	33.619	28.694	37.687
DC-MMI200	30.835	24.41	44.754
Content Richness?			
DC-10	19.272	26.017	54.711
DC-MMI10	18.844	26.231	54.925
DC-MMI200	20.021	26.660	53.319

Table 6: Comparing the model variation by reducing beam size to 10 and also comparing decoder constraints without MMI reranking

achieve slightly better plausibility scores since they tend to generate safe, dull responses. However, we find that when using a beam size of 200 and MMI re-ranking, our approach which incorporates distributional constraints, DC-MMI200, achieves competitive plausibility, while achieving significantly higher content richness.

7.2.1 Statistical Significance of Results

To verify the statistical significance of our findings, we conducted a pairwise bootstrap test (Efron and Tibshirani, 1994; Berg-Kirkpatrick et al., 2012) comparing the difference between percentage of *Agree* annotations (Yes column in the Table 5). We computed p-values for each pair of models: MMI200 vs DC-MMI200 and TA vs DC-MMI200. For plausibility, we did not find a significant difference in either comparison (p-value ≈ 0.25) while for content richness, both differences were found to be significant (p-value $< 10^{-4}$). To summarize: our model significantly beats both baselines in terms of content richness while the difference in plausibility was not found to be statistically significant.

7.2.2 Pairwise Evaluation of Interestingness

To further validate our claims we also did a side by side comparison study between MMI200 and DC-MMI200. For every test case, we showed Mechanical Turk workers the source sentence along with responses generated by both systems and asked them select which is more interesting. We observe that in 56% out of 1000 cases, DC-MMI200 was rated as the more interesting response. The result is statistically significant with p-value $< 4 \times 10^{-4}$ (using an exact binomial test).

7.3 Model Variations

To see the effectiveness of our decoding constraints separately, we compare the best performing DC-MMI200 model with DC-10 and DC-MMI10, both of which use a beam size of 10 – DC-10 does not include MMI reranking. The results of Mechanical Turk evaluation, following the approach described in Section 7.2, are presented in Table 6. We observe that with a beam size of 10 our model is able to generate content rich responses, but suffers in terms of plausibility. The values in the table suggests the decoding constraints defined in this work successfully inject content words into candidate hypotheses and that MMI is able to effectively choose plausible candidates. In the case of DC-10 and DC-MMI10, both models generate the same candidates, but MMI is able to re-rank the results and thus improves plausibility.

8 Related Work

Conversational agents primarily fall into two categories: task oriented dialogue systems (Williams et al., 2013; Wen et al., 2015) and chatbots (Weizenbaum, 1966), although there have been some efforts to integrate the two (Dodge et al., 2015; Yu et al., 2017). Some of the earliest work on data-driven chatbots (Ritter et al., 2011) explored the use of phrase-based Statistical Machine Translation (SMT) on large numbers of conversations gathered from Twitter (Ritter et al., 2010). Subsequent progress on the use of neural networks in machine translation inspired the use of Sequence-to-Sequence (Seq2Seq) models for data-driven response generation (Shang et al., 2015; Sordoni et al., 2015; Li et al., 2016a).

Our approach, which incorporates distributional constraints into the decoding objective, is related to prior work on posterior regularization (Mann

and McCallum, 2008; Ganchev et al., 2010; Zhu et al., 2014). Posterior regularization introduces similar distributional constraints on expectations computed over unlabeled data using a model’s parameters. These are typically added to the learning objective for semi-supervised scenarios where available labeled data is limited. In contrast, our approach introduces distributional constraints into the decoding objective as a way to combine neural conversation models trained on large quantities of conversational data with separately trained models of topics and semantic similarity that can drive content selection.

There are numerous examples of related work on improving neural conversation models. Shao et al. (2017) introduced a stochastic approach to beam search that does segment-by-segment reranking to promote diversity. Zhang et al. (2018) develop models which converse while assuming a persona defined by a short description of attributes. Wang et al. (2017) suggested decoding methods that influence the style and topic of the generated response. Bosselut et al. (2018) develop discourse-aware rewards with reinforcement learning (RL) to generate long and coherent texts. Li et al. (2016c) applied deep reinforcement learning to dialogue generation to maximize long-term reward of the conversation, as opposed to directly maximizing likelihood of the response. This line of work was further extended with adversarial learning (Li et al., 2017) that rewards generated conversations that are indistinguishable from real conversations in the data. Lewis et al. (2017) applied reinforcement learning with dialogue rollouts to generate replies that maximize expected reward, while learning to generate responses from a crowdsourced dataset of negotiation dialogues. Choi et al. (2018) used crowd-workers to gather a corpus of 100K information-seeking QA dialogues that are answerable using text spans from Wikipedia. Niu and Bansal (2018) designed a number of weakly-supervised models that generate polite, neutral or rude responses. Their fusion model combines a language model trained on polite utterances with the decoder. In the second method they prepend the utterance with a politeness label and scale its embedding to vary politeness. The third model is Polite-RL which assigns a reward based on a politeness classifier. Gimpel et al. (2013) explored methods for increasing the diversity of N-best lists in machine translation by in-

roducing a pairwise dissimilarity function. Similar ideas have been explored in the context of neural generation models. (Vijayakumar et al., 2016; Li and Jurafsky, 2016; Li et al., 2016b)

Following previous work we evaluated our approach using a combination of automatic metrics and human judgments. Some recent work has explored the possibility of adversarial evaluation of neural conversation models (Lowe et al., 2017; Li et al., 2017).

9 Conclusions

We presented an approach to generate more interesting responses in neural conversation models by incorporating side information in the form of distributional constraints. When using maximum likelihood decoding objectives, neural conversation models tend to generate safe responses, such as “I don’t know” for most inputs. Our proposed approach provides a flexible method of incorporating a broad range of distributional constraints into the decoding objective. We proposed and empirically evaluated two constraints that factorize over words, and therefore naturally fit into the commonly used left-to-right beam search decoding framework. The first encourages the use of more relevant topic words in the response the second encourages semantic similarity between the source and target. We empirically demonstrated, through human evaluation, that when taken together these constraints lead to responses that contribute significantly more information to the conversation, while maintaining plausibility in the context of the input.

Acknowledgments

We thank the anonymous reviewers for their valuable feedback. This material is based upon work supported by the National Science Foundation under Grant No. IIS-1464128.

References

Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2016. A simple but tough-to-beat baseline for sentence embeddings.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Regina Barzilay and Mirella Lapata. 2005. Collective content selection for concept-to-text generation. In

Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing.

- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155.
- Taylor Berg-Kirkpatrick, David Burkett, and Dan Klein. 2012. An empirical investigation of statistical significance in nlp. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 995–1005. Association for Computational Linguistics.
- Antoine Bosselut, Asli Celikyilmaz, Xiaodong He, Jianfeng Gao, Po-Sen Huang, and Yejin Choi. 2018. Discourse-aware neural rewards for coherent text generation. *arXiv preprint arXiv:1805.03766*.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentaoh Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. Quac : Question answering in context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics, ACL 2011*.
- Jesse Dodge, Andreea Gane, Xiang Zhang, Antoine Bordes, Sumit Chopra, Alexander Miller, Arthur Szlam, and Jason Weston. 2015. Evaluating prerequisite qualities for learning end-to-end dialog systems. *arXiv preprint arXiv:1511.06931*.
- Bradley Efron and Robert J Tibshirani. 1994. *An introduction to the bootstrap*. CRC press.
- Kuzman Ganchev, Jennifer Gillenwater, Ben Taskar, et al. 2010. Posterior regularization for structured latent variable models. *Journal of Machine Learning Research*, 11(Jul):2001–2049.
- Kevin Gimpel, Dhruv Batra, Chris Dyer, and Gregory Shakhnarovich. 2013. A systematic exploration of diversity in machine translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1100–1111.
- Thomas L Griffiths, Mark Steyvers, David M Blei, and Joshua B Tenenbaum. 2005. Integrating topics and syntax. In *Advances in neural information processing systems*, pages 537–544.
- Anjali Kannan, Karol Kurach, Sujith Ravi, Tobias Kaufmann, Andrew Tomkins, Balint Miklos, Greg Corrado, László Lukács, Marina Ganea, Peter Young, et al. 2016. Smart reply: Automated response suggestion for email. In *Proceedings of the 22nd ACM SIGKDD International Conference on*

- Knowledge Discovery and Data Mining*, pages 955–964. ACM.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. Opennmt: Open-source toolkit for neural machine translation. In *Proc. ACL*.
- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning of negotiation dialogues. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. 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.
- Jiwei Li and Dan Jurafsky. 2016. Mutual information and diverse decoding improve neural machine translation. *arXiv preprint arXiv:1601.00372*.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016b. A simple, fast diverse decoding algorithm for neural generation. *arXiv preprint arXiv:1611.08562*.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016c. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017. Adversarial learning for neural dialogue generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2157–2169.
- Chia-Wei Liu, Ryan Lowe, Iulian 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 the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. Towards an automatic turing test: Learning to evaluate dialogue responses. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1116–1126.
- Gideon S Mann and Andrew McCallum. 2008. Generalized expectation criteria for semi-supervised learning of conditional random fields. *Proceedings of ACL-08: HLT*, pages 870–878.
- Ani Nenkova and Rebecca Passonneau. 2004. Evaluating content selection in summarization: The pyramid method. In *Proceedings of the human language technology conference of the north american chapter of the association for computational linguistics: Hlt-naacl 2004*.
- Tong Niu and Mohit Bansal. 2018. Polite dialogue generation without parallel data. *Transactions of the Association for Computational Linguistics (TACL)*.
- Alan Ritter, Colin Cherry, and Bill Dolan. 2010. Unsupervised modeling of twitter conversations. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Alan Ritter, Colin Cherry, and William B Dolan. 2011. Data-driven response generation in social media. In *Proceedings of the conference on empirical methods in natural language processing*, pages 583–593. Association for Computational Linguistics.
- Iulian Vlad Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, volume 16, pages 3776–3784.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1577–1586.
- Yuanlong Shao, Stephan Gouws, Denny Britz, Anna Goldie, Brian Strope, and Ray Kurzweil. 2017. Generating high-quality and informative conversation responses with sequence-to-sequence models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2210–2219.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 196–205.
- Ilya Sutskever, James Martens, and Geoffrey Hinton. 2011. Generating text with recurrent neural networks. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pages 1017–1024.

- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Jörg Tiedemann. 2009. News from opus—a collection of multilingual parallel corpora with tools and interfaces. In *Recent advances in natural language processing*, volume 5, pages 237–248.
- 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*.
- Di Wang, Nebojsa Jojic, Chris Brockett, and Eric Nyberg. 2017. Steering output style and topic in neural response generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2140–2150.
- Joseph Weizenbaum. 1966. Eliza—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- John Wieting and Kevin Gimpel. 2017. Revisiting recurrent networks for paraphrastic sentence embeddings. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 2078–2088.
- Jason Williams, Antoine Raux, Deepak Ramachandran, and Alan Black. 2013. The dialog state tracking challenge. In *Proceedings of the SIGDIAL 2013 Conference*, pages 404–413.
- Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *AAAI*, volume 17, pages 3351–3357.
- Zhou Yu, Alexander Rudnicky, and Alan Black. 2017. Learning conversational systems that interleave task and non-task content. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? *arXiv preprint arXiv:1801.07243*.
- Jun Zhu, Ning Chen, and Eric P Xing. 2014. Bayesian inference with posterior regularization and applications to infinite latent svms. *The Journal of Machine Learning Research*, 15(1):1799–1847.