

Keywords and Instances: A Hierarchical Contrastive Learning Framework Unifying Hybrid Granularities for Text Generation

Mingzhe Li^{1,2,*}, Xiexiong Lin^{3,*}, Xiuying Chen⁴, Jinxiong Chang³, Qishen Zhang³, Feng Wang³, Taifeng Wang³, Zhongyi Liu³, Wei Chu³, Dongyan Zhao^{1,2,†}, Rui Yan^{5,†}

¹Wangxuan Institute of Computer Technology, Peking University, Beijing, China

²Center for Data Science, AAIS, Peking University, Beijing, China

³Ant Group ⁴Computational Bioscience Reseach Center, KAUST

⁵Gaoling School of Artificial Intelligence, Renmin University of China

li_mingzhe@pku.edu.cn, xiexiong.lxx@antfin.com

Abstract

Contrastive learning has achieved impressive success in generation tasks to mitigate the “exposure bias” problem and discriminatively exploit the different quality of references. Existing works mostly focus on contrastive learning on the instance-level without discriminating the contribution of each word, while keywords are the gist of the text and dominant the constrained mapping relationships. Hence, in this work, we propose a hierarchical contrastive learning mechanism, which can unify hybrid granularities semantic meaning in the input text. Concretely, we first propose a keyword graph via contrastive correlations of positive-negative pairs to iteratively polish the keyword representations. Then, we construct intra-contrasts within instance-level and keyword-level, where we assume words are sampled nodes from a sentence distribution. Finally, to bridge the gap between independent contrast levels and tackle the common contrast vanishing problem, we propose an inter-contrast mechanism that measures the discrepancy between contrastive keyword nodes respectively to the instance distribution. Experiments demonstrate that our model outperforms competitive baselines on paraphrasing, dialogue generation, and storytelling tasks.

1 Introduction

Generation tasks such as storytelling, paraphrasing, and dialogue generation aim at learning a certain correlation between text pairs that maps an arbitrary-length input to another arbitrary-length output. Traditional methods are mostly trained with “teacher forcing” and lead to an “exposure bias” problem (Schmidt, 2019). Incorporating the generation method with contrastive learning achieved impressive performance on tackling such issues, which takes an extra consideration of synthetic negative samples contrastively (Lee et al., 2021).

* Equal Contribution

† Corresponding authors: Rui Yan and Dongyan Zhao

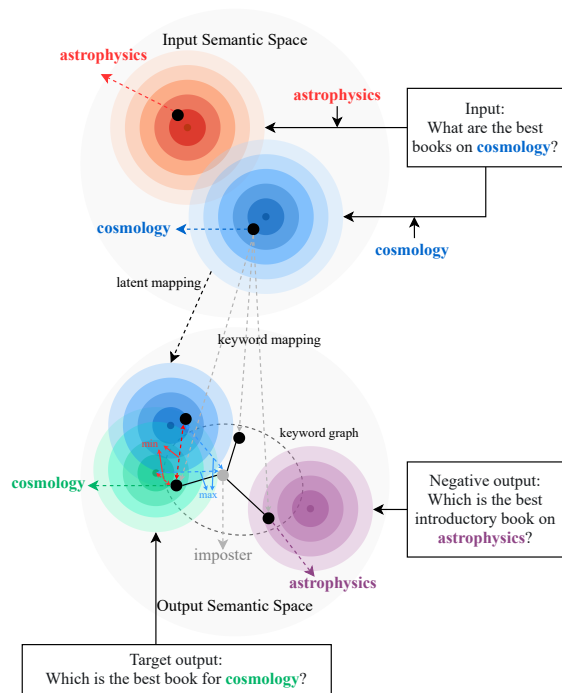


Figure 1: The semantic meaning of the sentence “what are the best books on cosmology?” would be greatly changed if the keyword “cosmology” is changed to “astrophysics”.

Existing contrastive mechanisms are mainly focused on the instance level (Lee et al., 2021; Cai et al., 2020). However, word-level information is also of great importance. Take the case shown in the upper part of Figure 1 for example, the keyword covers the gist of the input text and determines the embedding space of the text. The text representation will be significantly affected if adding a slight perturbation on the keyword, *i.e.*, changing “cosmology” to “astrophysics”. In addition, as shown on the bottom part, under some circumstances, it is too easy for the model to do the classification since the semantic gap between contrastive pairs is huge. Thus, the model fails to distinguish the actual discrepancy, which causes a “contrast vanishing”

problem at both instance-level and keyword-level.

Based on the above motivation, in this paper, we propose a hierarchical contrastive learning method built on top of the classic CVAE structure. We choose CVAE due to its ability in modeling global properties such as syntactic, semantic, and discourse coherence (Li et al., 2015; Yu et al., 2020). We first learn different granularity representations through two independent contrast, *i.e.*, instance-level and keyword-level. Specifically, we use the universal and classic TextRank (Mihalcea and Tarau, 2004) method to extract keywords from each text, which contain the most important information and need to be highlighted. On the instance-level, we treat the keyword in the input text as an additional condition for a better prior semantic distribution. Then, we utilize Kullback–Leibler divergence (Kullback and Leibler, 1951) to reduce the distance between prior distribution and positive posterior distribution, and increase the distance with the negative posterior distribution. While on the keyword-level, we propose a keyword graph via contrastive correlations of positive-negative pairs to learn informative and accurate keyword representations. By treating the keyword in the output text as an anchor, the imposter keyword is produced by neighboring nodes of the anchor keyword and forms the keyword-level contrast, where the similarity between the imposter keyword and the anchor keyword is poorer than the positive keyword.

To unify individual intra-contrasts and tackle the “contrast vanishing” problem in independent contrastive granularities, we leverage an inter-contrast, the Mahalanobis contrast, to investigate the contrastive enhancement based on the Mahalanobis distance (De Maesschalck et al., 2000), a measure of the distance between a point and a distribution, between the instance distribution and the keyword representation. Concretely, we ensure the distance from the anchor instance distribution to the ground-truth keyword vector is closer than to the imposter keyword vector. The Mahalanobis contrast plays an intermediate role that joins the different granularities contrast via incorporating the distribution of instance with the representation of its crucial part, and makes up a more comprehensive keyword-driven hierarchical contrastive mechanism, so as to ameliorate the generated results.

We empirically show that our model outperforms CVAE and other baselines significantly on three generation tasks: paraphrasing, dialogue genera-

tion, and storytelling.

Our contributions can be summarized as follows:

- To our best knowledge, we are the first to propose an inter-level contrastive learning method, which unifies instance-level and keyword-level contrasts in the CVAE framework.
- We propose three contrastive learning measurements: KL divergence for semantic distribution, cosine distance for points, and Mahalanobis distance for points with distribution.
- We introduce a global keyword graph to obtain polished keyword representations and construct imposter keywords for contrastive learning.

2 Related Work

2.1 Contrastive Learning

Contrastive learning is used to learn representations by teaching the model which data points are similar or not. Due to the excellent performance on self-supervised and semi-supervised learning, it has been widely used in natural language processing (NLP). Firstly, Mikolov et al. (2013) proposed to predict neighboring words from context with noise-contrastive estimation. Then, based on word representations, contrastive learning for sentence has been utilized to learn semantic representations. Lee et al. (2021) generated positive and negative examples by adding perturbations to the hidden states. Cai et al. (2020) augmented contrastive dialogue learning with group-wise dual sampling. Moreover, contrastive learning has also been utilized in caption generation (Mao et al., 2016), summarization (Liu and Liu, 2021) and machine translation (Yang et al., 2019). Our work differs from previous works in focusing on hierarchical contrastive learning on hybrid granularities.

2.2 Mahalanobis Distance

The Mahalanobis distance is a measure of the distance between a point and a distribution (De Maesschalck et al., 2000). The distance is zero if the point is on the distribution. Recently, Mahalanobis distance is popularly applied to the NLP tasks (Tran et al., 2019). Podolskiy et al. (2021) showed that while Transformer is capable of constructing homogeneous representations of in-domain utterances, the Mahalanobis distance captures geometrical disparity from out of domain utterances. Further, Ren et al. (2021) considered that the raw density from deep generative models may fail at out-of-domain detection and proposed to fix this using a likeli-

hood ratio between two generative models as a confidence score.

2.3 Conditional Variational Auto-Encoder

Variational autoencoder (VAE) was proposed by Kingma and Welling (2013), and has been widely used in various tasks such as headline generation (Li et al., 2021), dialogue generation (Serban et al., 2017) and story generation (Yu et al., 2020). Based on VAE, a more advanced model, Conditional VAE (CVAE), was proposed to generate diverse images conditioned on certain attributes, which was also applied to generate diverse outputs in NLP tasks (Zhao et al., 2017; Qiu et al., 2019). Existing works concentrate on generating diverse outputs, and we take one step further to utilize prior and posterior latent distribution to compare positive and negative samples, which helps to learn more accurate semantic information.

3 Method

3.1 Background

VAE: Variational auto-encoder (VAE) is a typical encoder-decoder structural model with certain types of latent variables. Given an input x , VAE models the latent variable z through the prior distribution $p_\theta(z)$, and the observed data x is reconstructed by the generative distribution $p_\theta(x|z)$ which is the likelihood function that generates x conditioned on z . Since z is unknown, it should be estimated according to the given data x as $p_\theta(z|x)$. While the posterior density $p_\theta(z|x) = p_\theta(x|z)p_\theta(z)/p_\theta(x)$ is intractable, VAE introduces a recognition posterior distribution $q_\phi(z|x)$ approximates to the true posterior $p_\theta(z|x)$. Thus, VAE is trained by optimizing the lower bound on the marginal likelihood of data x as:

$$\begin{aligned} \log p_\theta(x) &\geq \mathbb{E}_{z \sim q_\phi(z|x)} [\log p_\theta(x|z)] \\ &\quad - D_{KL}(q_\phi(z|x) || p_\theta(z)), \end{aligned} \quad (1)$$

where D_{KL} is the Kullback–Leibler divergence.

CVAE: The conditional variational auto-encoder (CVAE) is the supervised version of VAE with an additional output variable. Giving a dataset $\{x_i, y_i\}_{i=1}^N$ consisting of N samples, CVAE is trained to maximize the conditional log-likelihood, and the variational lower bound of the model is written as follows:

$$\begin{aligned} \log p_\theta(y|x) &\geq \mathbb{E}_{z \sim q_\phi(z|x,y)} [\log p_\theta(y|x,z)] \\ &\quad - D_{KL}(q_\phi(z|x,y) || p_\theta(z|x)). \end{aligned} \quad (2)$$

Assuming the type of latent variable obeys Gaussian distribution, the first right-hand side term can be approximated by drawing samples $\{z_i\}_{i=1}^N$ from the recognition posterior distribution $q_\phi(z|x,y)$, where $z \sim N(\mu, \sigma^2 I)$, and then objective of the CVAE with Gaussian distribution can be written as:

$$\begin{aligned} \mathcal{L}_{cvae}(x, y; \theta, \phi) &= -\frac{1}{N} \sum_{i=1}^N \log p_\theta(y|x, z_i) \\ &\quad + D_{KL}(q_\phi(z|x, y) || p_\theta(z|x)), \end{aligned} \quad (3)$$

where $z_i = g_\phi(x, y, \epsilon_i)$, $\epsilon_i \sim \mathcal{N}(0, I)$. The distribution $q_\phi(z|x, y)$ is reparameterized with a differentiable function g_ϕ , which enables the model trainable via stochastic gradient descent.

Inspired by Wu et al. (2019), we add keyword u as an additional condition to the prior distribution to control the generation process, which turns the $p_\theta(z|x)$ in Equation 3 into $p_\theta(z|x, u)$.

3.2 Hierarchical Contrastive Learning

In this section, we introduce our hierarchical contrastive learning method, which is comprised of three parts: instance-level contrast based on KL divergence (sec.3.2.1), keyword-level contrast based on keyword graph (sec.3.2.2), and inter-contrast: Mahalanobis contrast (sec.3.2.3).

3.2.1 Instance-level Contrastive Learning

To tackle the ‘‘exposure bias’’ problem and discriminatively exploit the different quality of references, instance-level contrastive learning is introduced to learn discrepancies of targets. Specifically, in addition to the observed input data x and positive output y^+ , a negative output y^- is added to construct a contrastive pair $\{(x, y^+), (x, y^-)\}$. In this case, the prior distribution $p_\theta(z|x)$ is learned from a prior network, which is denoted as $f_\theta(x)$. The approximate posteriors $q_\phi(z|x, y^+)$ and $q_\phi(z|x, y^-)$ are learned from a posterior network and represented as $f_\phi(x, y^+)$ and $f_\phi(x, y^-)$, respectively. The objective here is to make the distance between a prior distribution and positive posterior distribution closer than with the negative posterior distribution. Thus, the instance-level contrastive loss function can be written as:

$$\mathcal{L}_{ins} = -\mathbb{E}_{f_\phi} \left[\log \left(1 - \frac{e^{h(f_\phi(x, y^+), f_\theta(x)) / \tau}}{\sum_{y^* \in Y} e^{h(f_\phi(x, y^*), f_\theta(x)) / \tau}} \right) \right],$$

where the $y^* \in Y$ can be positive sample y^+ or negative sample y^- , and the τ is a temperature

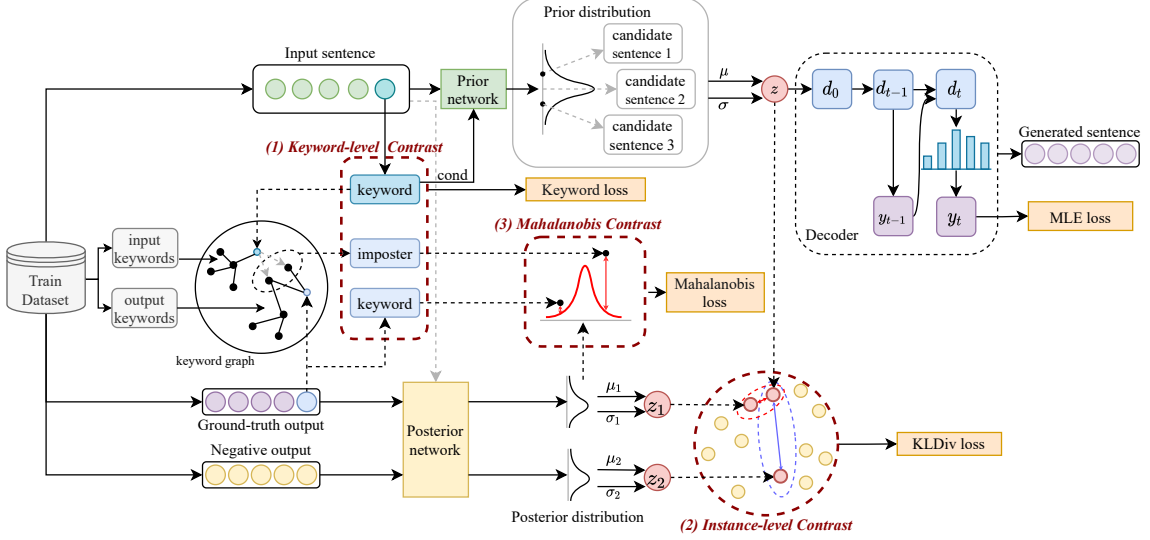


Figure 2: The architecture of hierarchical contrastive learning, which consists of three parts: (1) Keyword-level contrast from keyword graph; (2) Instance-level contrast based on KL divergence for semantic distribution; and (3) Mahalanobis contrast between instance-level and keyword-level.

parameter to control push and pull force. The function $h(\cdot)$ denotes the distance between elements, which is set as Kullback–Leibler divergence (Kullback and Leibler, 1951) in instance-level contrast, $D_{KL}(f_\phi(x, y^*) || f_\theta(x))$, to measure the difference between two distributions.

3.2.2 Keyword-level Contrastive Learning

Since the instance-level contrast focuses on learning high-level information and fails to discriminate the contribution of each word, we incorporate it with a keyword-level contrast to pay more attention to the specific keyword.

Keyword Graph: Given an input-output text pair (x, y) , keywords k_x, k_y can be extracted from x and y , respectively. For an input text x_i with keyword $k_{x,i}$, input texts that contain the same keyword are gathered into a cluster $C_i = \{x_j\}_{j=1}^n, k_{x,j} \in x_j$, where n is the number of texts in C_i . Each text $x_j \in C_i$ has a positive-negative output text pair $\{(y_j^+, y_j^-)\}$ containing a positive output keyword $k_{y,j}^+$ and a negative one $k_{y,j}^-$, respectively. Thus, spreading to the entire cluster C_i , for the output text y_i , there exists positive relations $r_{i,j}^+$ between its keyword $k_{y,i}$ and each of the surrounded positive keywords $\{k_{y,j}^+\}_{j=1}^n$. Likewise, negative relations $r_{i,j}^-$ correlates the output keyword $k_{y,i}$ and the surrounded negative ones $\{k_{y,j}^-\}_{j=1}^n$.

Based on these keywords as nodes and their relations as edges (Chen et al., 2021), the key-

word graph \mathcal{G}_k is constructed. Each node representation h_i^0 is initialized as the average BERT embedding (Devlin et al., 2018) of texts in the cluster C_i with the same corresponding keyword $k_{x,i}$. Then, the relation edge r_{ij}^0 that connects node i and node j is learned via a feedforward layer $r_{ij}^0 = \text{FFN}([h_i^0; h_j^0])$.

Then, the representations of nodes and relation edges are iteratively updated with their connected nodes via the graph attention (GAT) layer and the feed-forward (FFN) layer. In the t -th iteration, we first update each edge representation by paying attention to the connected nodes, denoted as:

$$\beta_{r_*}^t = \text{softmax}\left(\frac{(r_{ij}^t W_p)(h_*^t W_h)^T}{\sqrt{d}}\right), \quad (4)$$

$$p_{ij}^t = \beta_{r_i}^t h_i^t + \beta_{r_j}^t h_j^t, \quad (5)$$

$$r_{ij}^{t+1} = \text{FFN}(r_{ij}^t + p_{ij}^t), \quad (6)$$

where h_*^t can be h_i^t or h_j^t .

Then, based on the obtained edge representation r_{ij}^{t+1} , we update the node representations considering both the related nodes and relation edges by the graph attention layer, $\text{GAT}(h_i^t, h_j^t, r_{ij}^{t+1})$, which is designed as:

$$e_{ij}^t = \frac{(h_i^t W_q)(h_j^t W_k + r_{ij}^{t+1} W_r)^T}{\sqrt{d}}, \quad (7)$$

$$\alpha_{ij}^t = \frac{\exp(e_{ij}^t)}{\sum_{l \in N_i} \exp(e_{il}^t)}, \quad (8)$$

$$u_i^t = \sum_{j \in N_i} \alpha_{ij}^t (h_j^t W_v + r_{ij}^{t+1}), \quad (9)$$

where W_q, W_k, W_r and W_v are all learnable parameters, and the α_{ij}^t is the attention weight between h_i^t and h_j^t . Besides, to avoid gradient vanishing after several iterations, a residual connection is added to the output u_i^t and the updated node representations h_i^{t+1} is obtained. In this way, the new representation of each keyword node consists of the relation dependency information from neighbor nodes N_i . We take the node representations from the last iteration as the final keyword representations, denoted as u for brevity.

Keyword-level Contrast: The keyword-level contrastive learning arises from input keywords against positive output keywords and negative impostor keywords. The input keyword u_{in} is extracted from the input text as an anchor, and the output keyword u_{out} is extracted from ground-truth output text. While the impostor keyword is calculated from the negative neighbours of the output keyword u_{out} , written as $u_{imp} = \sum_i W_i u_i$, where u_i is the representation of keyword node which is obtained by the keyword graph learning procedure described above. In this way, with the help of neighbour nodes in the graph, we can obtain a more indistinguishable and difficult negative sample. The loss of keyword level contrastive learning thus can be written as:

$$\mathcal{L}_{keyword} = -\mathbb{E}[\log \frac{e^{h(u_{in}, u_{out})/\tau}}{\sum_{u_* \in U} e^{h(u_{in}, u_*)/\tau}}], \quad (10)$$

where $u_* \in U$ denotes the positive output keyword u_{out} or impostor keyword u_{imp} . In keyword-level contrast, $h(\cdot)$ utilizes cosine similarity to calculate the distance between points.

3.2.3 Mahalanobis Contrastive Learning

Note that there exists a space gap between the instance-level contrast and the keyword-level contrast, which disturbs the completeness of this hierarchical contrastive architecture. Besides, the contrastive values vanish when the distance metric is hard to measure the actual discrepancy between positive and negative merely in instance distributions or in keyword representations. To mitigate such problems, we design a Mahalanobis contrastive mechanism to correlate the instance distribution and keyword representation, where the objective is to minimize the margin between the output keyword u_{out} and the posterior semantic distribution $q_\phi(z|x, y) \triangleq f_\phi(x, y)$ and maximize the margin between the impostor keyword u_{imp} and

the posterior distribution $f_\phi(x, y)$:

$$\mathcal{L}_{ma} = -\mathbb{E}_{f_\phi}[\log(1 - \frac{e^{h(f_\phi(x, y), u_{out})/\tau}}{\sum_{u_* \in U} e^{h(f_\phi(x, y), u_*)/\tau}})], \quad (11)$$

where $u_* \in U$ can be the positive output keyword u_{out} or negative impostor keyword u_{imp} . In Mahalanobis contrast, $h(\cdot)$ utilizes Mahalanobis distance (De Maesschalck et al., 2000) to measure the similarity from keyword point to the instance distribution. In the univariate Gaussian case, $z \sim p(z|x, y) = N(\mu, \sigma^2)$, then the $h(f_\phi(x, y), u_*) \triangleq D_{MA}(p_\theta(z|x, y)||u_*) = (u_* - \mu)\sigma^2 I(u_* - \mu)$.

Finally, we equip the CVAE model with the proposed hierarchical contrastive learning framework to unify hybrid granularities by adding \mathcal{L}_{ins} , $\mathcal{L}_{keyword}$ and \mathcal{L}_{ma} to the reconstructed loss of Equation 3.

4 Experiment

4.1 Tasks and Datasets

We conduct experiments on three public datasets QQP, Douban, RocStories for paraphrasing, dialogue generation, and storytelling task, respectively. The details of the datasets are as follows:

Dialogue (Douban) Douban (Cai et al., 2020) consists of Chinese daily conversations between pairs of speakers, collected from a popular social network website, Douban group¹. The dataset contains 218,039/10,000/10,000 context-response pairs for training/validation/test, with an average of 3.94 turns per context and 38.32 characters per utterance. We concatenate historical dialogues and turn it into a single-turn dialogue training corpus.

Paraphrasing (QQP) QQP (Iyer et al., 2017; Wang et al., 2019) is a dataset published by the community question-answering website Quora on whether a pair of questions is semantically consistent. To adapt it to the contrastive learning task, we only keep question pairs that have positive and negative rewriting for the same input. Thus, there remain 44,949 samples in the dataset, which are split into training/validation/test sets of 40,441/2,254/2,254 samples.

Storytelling (RocStories) RocStories consists of 98,163 high-quality hand-crafted stories, which capture causal and temporal commonsense relations of daily events (Mostafazadeh et al., 2016).

¹<https://www.douban.com/>

	Models	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Extrema	Average	Greedy
(a).QQP	CVAE	0.4562	0.2150	0.0962	0.0496	0.6606	0.8371	0.8406
	Seq2Seq	0.4510	0.2117	0.0950	0.0497	0.6543	0.8243	0.8533
	Transformer	0.4832	0.2339	0.1086	0.0590	0.6523	0.8274	0.8531
	Seq2Seq-DU	0.5613	0.2781	0.1334	0.0763	0.6679	0.8302	0.8590
	DialoGPT	0.5749	0.2845	0.1337	0.0749	0.6658	0.8393	0.8597
	BERT-GEN	0.5452	0.2781	0.1343	0.0764	0.6673	0.8299	0.8586
	T5	0.6172	0.3301	0.1730	0.1019	0.6679	0.8408	0.8601
	Group-wise	0.5201	0.2472	0.1112	0.0582	0.6576	0.8337	0.8569
	T5-CLAPS	0.6264	0.3394	0.1787	0.1058	0.6683	0.8430	0.8612
	Ours	0.6430	0.3517	0.1845	0.1153	0.6701	0.8495	0.8661
(b).Douban	CVAE	0.0640	0.0259	0.0102	0.0047	0.4473	0.4814	0.6006
	Seq2Seq	0.0542	0.0218	0.0086	0.0039	0.4388	0.4802	0.5960
	Transformer	0.0531	0.0210	0.0081	0.0036	0.4401	0.4807	0.5989
	Seq2Seq-DU	0.0887	0.0333	0.0123	0.0050	0.4591	0.4972	0.6083
	Dialo-gpt	0.0953	0.0363	0.0136	0.0057	0.4633	0.4997	0.6108
	BERT-GEN	0.0823	0.0314	0.0119	0.0050	0.4630	0.5001	0.6095
	T5	0.1007	0.0379	0.0140	0.0057	0.4659	0.5056	0.6147
	Group-wise	0.0581	0.0239	0.0101	0.0053	0.4489	0.4903	0.6002
	T5-CLAPS	0.1162	0.0413	0.0165	0.0062	0.4667	0.5071	0.6159
	Ours	0.1398	0.0516	0.0188	0.0074	0.4691	0.5091	0.6179
(c).RocStories	CVAE	0.2581	0.0969	0.0360	0.0148	0.5135	0.5632	0.6011
	Seq2Seq	0.2324	0.0896	0.0340	0.0150	0.5133	0.5649	0.6007
	Transformer	0.2552	0.0596	0.0354	0.0145	0.5129	0.5637	0.6006
	Seq2Seq-DU	0.3089	0.1131	0.0384	0.0156	0.5166	0.5773	0.6116
	DialoGPT	0.3126	0.1105	0.0398	0.0158	0.5178	0.5860	6103
	BERT-GEN	0.3040	0.1071	0.0385	0.0153	0.5173	0.5866	0.6098
	T5	0.3274	0.1177	0.0416	0.0164	0.5194	0.5864	0.6118
	Group-wise	0.2748	0.1017	0.0366	0.0152	0.5139	0.5651	0.6009
	T5-CLAPS	0.3420	0.1261	0.0452	0.0176	0.5232	0.5880	0.6134
	Ours	0.3552	0.1341	0.0485	0.0184	0.5218	0.5884	0.6131

Table 1: Automatic evaluation results on (a) QQP for paraphrasing, (b) Douban for dialogue generation, and (c) RocStories for storytelling. The best results in each group are highlighted with **bold**.

Each story paragraph contains 5 sentences with an average of 43 words. Following the previous work Yu et al. (2021), we split the dataset into 8:1:1 for training, validation, and test.

For the above three datasets, in order to construct different levels of contrastive learning, we performed the same preprocessing of extracting keywords. We utilize the TextRank model (Mihalcea and Tarau, 2004) to extract keywords from each input and output sample, respectively. Besides, the vocabulary size of both datasets is the same as BERT (Devlin et al., 2018) setting.

4.2 Implementation Details

Our experiments are implemented in TensorFlow (Abadi et al., 2016) on an NVIDIA Tesla P100 GPU. For our model and all baselines, we follow the same setting as described below. We pad or cut the input to 100, 20, 100 words for dialogue generation, paraphrasing, and storytelling, respectively. The truncation length is decided based on the observation that there is no significant improvement when increasing input length. The minimum decoding step is 5, and the maximum step is 20 for

all tasks. Experiments were performed with a batch size of 256, and we use Adam optimizer (Kingma and Ba, 2015) as our optimizing algorithm. During the test stage, the beam-search size is set to 4 for all methods and the checkpoint with the smallest validation loss is chosen. Note that for better performance, our model is built based on BERT, and the decoding process is the same as Transformer (Vaswani et al., 2017). Finally, due to the limitation of time and memory, small settings are used in the pre-training baselines.

4.3 Compared Baselines

We compare our method against several traditional generation models, pretrained-based generation models, and contrastive learning models.

Traditional generation models: (1) **CVAE** (Zhao et al., 2017) generates sentences based on latent variables, sampling from potential semantic distribution. (2) **Seq2Seq** (Sutskever et al., 2014) is a sequence-to-sequence framework combined with attention mechanism and pointer network. (3) **Transformer** (Vaswani et al., 2017) is an abstract method based solely on attention mechanisms.

Models	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Extrema	Average	Greedy
ours	0.6430	0.3517	0.1845	0.1153	0.6701	0.8495	0.8661
w/o graph	0.6295	0.3333	0.1675	0.1001	0.6685	0.8455	0.8647
w/o keyword	0.5764	0.2993	0.1499	0.0892	0.6673	0.8450	0.8539
w/o MA	0.6013	0.3208	0.1628	0.0981	0.6605	0.8436	0.8524

Table 2: Ablation results on dataset QQP.

Pretrained-based generation models: (4) **Seq2Seq-DU** (Feng et al., 2021) is concerned with dialogue state tracking in a task-oriented dialogue system. (5) **DialogPT** (Zhang et al., 2020) proposes a large, tunable neural conversational response generation model trained on more conversation-like exchanges. (6) **BERT-GEN** (Devlin et al., 2018) augments Seq2Seq with BERT as the encoder. (7) **T5** (Raffel et al., 2020) introduces a unified framework that converts all text-based language problems into a text-to-text format.

Contrastive learning methods: (8) **Group-wise** (Cai et al., 2020) augments contrastive dialogue learning with group-wise dual sampling. (9) **T5-CLAPS** (Lee et al., 2021) generates negative and positive samples for contrastive learning by adding small and large perturbations, respectively.

4.4 Evaluation Metrics

To evaluate the performance of our model against baselines, we adopt the following metrics widely used in existing studies.

BLEU We utilize BLEU score (Papineni et al., 2002) to measure word overlap between the generated text and the ground-truth. Specifically, following the conventional setting of (Gu et al., 2019), we adopt BLEU-1~4 scores under the smoothing techniques (smoothing 7).

Embedding To evaluate our model more comprehensively, we also capture the semantic matching degrees between the bag-of-words (BOW) embeddings of generated text and reference (Gu et al., 2019). Particularly we adopt three metrics: 1) *Extrema*, cosine similarity between the largest extreme values among the word embeddings in the two texts; 2) *Average*, cosine similarity between the averaged word embeddings of generated text and reference; 3) *Greedy*, greedily matching words in the two texts based on cosine similarities.

4.5 Experimental Results

4.5.1 Overall Performance

Automatic Evaluation The experimental results are summarized in Table 1. The upper part lists the

	Flu	Mean	Diff
Seq2Seq-DU	2.03	2.12	1.76
DialogPT	2.18	2.04	1.97
T5-CLAPS	2.24	2.16	2.19
Ours	2.51[▲]	2.45[▲]	2.43[▲]

Table 3: Fluency (Flu), Meaningfulness (Mean), and Differential (Diff) comparison by human evaluation. [▲] shows the statistical significance tested by a two tailed paired t-test.

effects of traditional generation methods such as Seq2Seq and Transformer, and the lower part shows the latest pretrained-based methods including DialogPT and T5. Overall, pretrained-based methods generally outperform traditional methods, and this also proves the effectiveness of the pretrained language model on the generation tasks. Secondly, we can find that the performance is significantly improved after adding contrast learning. Finally, our method outperforms T5-CLAPS by 2.7%, 3.6% on QQP, by 20.3%, 24.9% on Douban, and by 3.9%, 6.3% on RocStories in terms of BLEU-1, BLEU-2, respectively, which proves the superiority of our model.

Human Evaluation We also assessed system performance by eliciting human judgments on 100 randomly selected test instances on QQP dataset. Three annotators are asked to rate paraphrasing questions generated by T5-CLAPS, DialogPT, Seq2Seq-DU, and our model according to Fluency (Flu), Meaningfulness (Mean), and Differential (Diff). The rating score ranges from 1 to 3, with 3 being the best. Table 3 lists the average scores of each model, showing that our model outperforms other baselines among all metrics, which indicates that our model generates paraphrasing sentences more readable successfully. The kappa statistics are 0.53, 0.61, and 0.56 for fluency, meaningfulness, and differential, respectively, which indicates the moderate agreement between annotators.

4.5.2 Ablation Study

We conduct ablation tests to assess the importance of the keyword graph architecture (w/o graph), keyword (w/o keyword), as well as the Maha-

Input text: What one exercise will help me lose belly fat?
Reference paraphrasing text: How do i remove belly fat?
Keyword extracted by TextRank: belly fat
Generated text 1: What is the best exercise way to lose belly fat?
Keyword from random-selected: disposable
Generated text 2: Can one-off exercise lose belly fat?
Seq2Seq-DU: What are the best ways to lose weight?
DialoGPT: Which exercise helps to lose weight?
T5-CLAPS: How can I lose weight?

Table 4: Case study to verify the influence of sampling different keywords. The texts in red and blue indicate the parts corresponding to the extracted keyword and random-selected keyword, respectively.

lanobis contrast (w/o MA contrast), and the results are shown in Table 2. Concretely, after removing the keywords (w/o keyword), using only instance-level contrastive, the effect of our model is greatly reduced by about 10.4%, which illustrates the desirability of considering the contributions of words in a sentence. On this basis, adding keyword contrastive learning with removing the keyword graph, the effect of the model has been improved but is still lower than our model by 2.1%. This shows that keywords are indeed conducive to capturing important information, and it also illustrates the significance of a keyword graph. Finally, the experiment of removing the Mahalanobis contrastive loss indicates that only with granularity independent contrast is not sufficient, and the Mahalanobis contrast plays a critical intermediate role.

4.5.3 Visualization of Different Levels of Contrastive Learning

To study the hierarchical contrastive learning, we visualize the vectors of keyword, input text, positive and negative output text on randomly sampled cases from QQP dataset, as shown in Figure 3. For visualization purposes, we reduce the dimension of the latent vector with t-SNE (Maaten and Hinton, 2008). It can be observed that the input sentence representation is located close to the keyword, which shows that the keyword, as the most important information in the sentence, determines the semantic distribution. Moreover, in contrastive learning, it can be seen that after training, the position of the input sentence is close to the positive samples and far away from the negative samples. This suggests that contrastive learning can correct the semantic distribution.

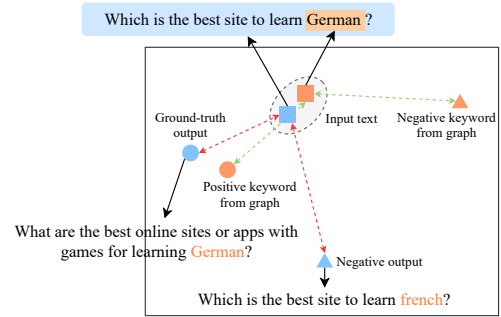


Figure 3: Visualization of contrastive learning. The square, circle and triangle represents the input text, positive output sample, and negative output sample, respectively. Blue represents the sentence, and yellow represents the keyword.

4.5.4 Analysis of Different Keywords

We finally investigate the influence of sampling different keywords. As shown in Table 4, for an input question, we provide keywords extracted by TextRank and randomly-selected keywords as the condition to control the semantic distribution and examine the quality of the generated text. As the most important information unit, different keywords lead to different semantic distributions and will result in different generated texts. The more properly the keywords are selected, the more accurately the sentences will be generated. When utilizing the keywords extracted by TextRank as a condition, the information “belly fat” is focused during the generation of paraphrasing questions, and the generated sentences are more accurate. On the contrary, after adding the random-selected keyword “disposable”, the generated question emphasizes “one-off exercise”, which brings incorrect information.

We also compare our model with several baselines in Table 4. Most baselines can generate fluent questions in this case. However, they focus on “lose weight”, and miss the significant information “belly fat”. Based on the above analysis, we can observe that keywords can emphasize and protect the highlight information in sentences, and affect the semantic distribution of as a condition.

5 Conclusion

In this paper, we propose a hierarchical contrastive learning mechanism, which consists of intra-contrasts within instance-level and keyword-level and inter-contrast with Mahalanobis contrast. The experimental results yield significant out-performance over baselines when applied in

the CVAE framework. In the future, we aim to extend the contrastive learning mechanism to different basic models, and will explore contrastive learning methods based on external knowledge.

6 Acknowledgments

We would like to thank the anonymous reviewers for their constructive comments. This work was supported by National Key Research and Development Program of China (No. 2020AAA0106600), National Natural Science Foundation of China (NSFC Grant No. 62122089, No. 61832017 & No. 61876196), and Beijing Outstanding Young Scientist Program No. BJJWZYJH012019100020098. This work was also supported by Alibaba Group through Alibaba Research Intern Program.

7 Ethics Impact

In this paper, we propose an inter-level contrastive learning method, which unifies instance-level and keyword-level contrasts in the CVAE framework. The positive impact lies in that it can help improve the capability of generation models on paraphrasing, dialogue generation, and storytelling tasks. The negative impact may be that the generation process of the system is not fully controllable, so it is possible to generate inaccurate or unreasonable content in some extreme cases. Hence, extra processing steps might be needed if this method were to be used in scenarios where high accuracy is required.

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek Gordon Murray, Benoit Steiner, Paul A. Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zhang. 2016. Tensorflow: A system for large-scale machine learning. In *OSDI*.
- Hengyi Cai, Hongshen Chen, Yonghao Song, Zhuoye Ding, Yongjun Bao, Weipeng Yan, and Xiaofang Zhao. 2020. [Group-wise contrastive learning for neural dialogue generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 793–802. Association for Computational Linguistics.
- Xiuying Chen, Hind Alamro, Mingzhe Li, Shen Gao, Xi-angliang Zhang, Dongyan Zhao, and Rui Yan. 2021. Capturing relations between scientific papers: An abstractive model for related work section generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6068–6077.
- Roy De Maesschalck, Delphine Jouan-Rimbaud, and Désiré L Massart. 2000. The mahalanobis distance. *Chemometrics and intelligent laboratory systems*, 50(1):1–18.
- 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*.
- Yue Feng, Yang Wang, and Hang Li. 2021. A sequence-to-sequence approach to dialogue state tracking. *ACL 2021*.
- Xiaodong Gu, Kyunghyun Cho, Jung-Woo Ha, and Sunghun Kim. 2019. [DialogWAE: Multimodal response generation with conditional wasserstein auto-encoder](#). In *International Conference on Learning Representations*.
- Shankar Iyer, Nikhil Dandekar, and Kornel Csernai. 2017. [First quora dataset release: Question pairs](#).
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *ICLR*.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86.
- Seanie Lee, Dong Bok Lee, and Sung Ju Hwang. 2021. [Contrastive learning with adversarial perturbations for conditional text generation](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Jiwei Li, Minh-Thang Luong, and Dan Jurafsky. 2015. A hierarchical neural autoencoder for paragraphs and documents. *arXiv preprint arXiv:1506.01057*.
- Mingzhe Li, Xiuying Chen, Min Yang, Shen Gao, Dongyan Zhao, and Rui Yan. 2021. The style-content duality of attractiveness: Learning to write eye-catching headlines via disentanglement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13252–13260.
- Yixin Liu and Pengfei Liu. 2021. Simcls: A simple framework for contrastive learning of abstractive summarization. *arXiv preprint arXiv:2106.01890*.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.

- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pages 404–411.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318. ACL.
- Alexander Podolskiy, Dmitry Lipin, Andrey Bout, Ekaterina Artemova, and Irina Piontkovskaya. 2021. Revisiting mahalanobis distance for transformer-based out-of-domain detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13675–13682.
- Lisong Qiu, Juntao Li, Wei Bi, Dongyan Zhao, and Rui Yan. 2019. Are training samples correlated? learning to generate dialogue responses with multiple references. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3826–3835.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Jie Ren, Stanislav Fort, Jeremiah Liu, Abhijit Guha Roy, Shreyas Padhy, and Balaji Lakshminarayanan. 2021. A simple fix to mahalanobis distance for improving near-ood detection. *arXiv preprint arXiv:2106.09022*.
- Florian Schmidt. 2019. Generalization in generation: A closer look at exposure bias. *EMNLP-IJCNLP 2019*, page 157.
- Iulian Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112.
- Thanh Tran, Renee Sweeney, and Kyumin Lee. 2019. Adversarial mahalanobis distance-based attentive song recommender for automatic playlist continuation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 245–254.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In the Proceedings of ICLR.
- Bowen Wu, Mengyuan Li, Zongsheng Wang, Yifu Chen, Derek Wong, Qihang Feng, Junhong Huang, and Baoxun Wang. 2019. Guiding variational response generator to exploit persona. *arXiv preprint arXiv:1911.02390*.
- Zonghan Yang, Yong Cheng, Yang Liu, and Maosong Sun. 2019. Reducing word omission errors in neural machine translation: A contrastive learning approach.
- Meng-Hsuan Yu, Juntao Li, Zhangming Chan, Dongyan Zhao, and Rui Yan. 2021. Content learning with structure-aware writing: A graph-infused dual conditional variational autoencoder for automatic storytelling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6021–6029.
- Meng-Hsuan Yu, Juntao Li, Danyang Liu, Dongyan Zhao, Rui Yan, Bo Tang, and Haisong Zhang. 2020. Draft and edit: Automatic storytelling through multi-pass hierarchical conditional variational autoencoder. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1741–1748.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. Dialogpt: Large-scale generative pre-training for conversational response generation. In *ACL, system demonstration*.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. *arXiv preprint arXiv:1703.10960*.