

Empowering Language Understanding with Counterfactual Reasoning

Fuli Feng^{1,2}, Jizhi Zhang³, Xiangnan He^{3*}, Hanwang Zhang⁴, Tat-Seng Chua²

¹NExT-Sea Joint Lab, ²National University of Singapore

³University of Science and Technology of China, ⁴Nanyang Technological University

fulifeng93@gmail.com, cdzhangjizhi@mail.ustc.edu.cn
xiangnanhe@gmail.com, hanwangzhang@ntu.edu.sg, dcscts@nus.edu.sg

Abstract

Present language understanding methods have demonstrated extraordinary ability of recognizing patterns in texts via machine learning. However, existing methods indiscriminately use the recognized patterns in the testing phase that is inherently different from us humans who have *counterfactual thinking*, e.g., to scrutinize for the hard testing samples. Inspired by this, we propose a *Counterfactual Reasoning Model*, which mimics the counterfactual thinking by learning from few counterfactual samples. In particular, we devise a *generation module* to generate representative counterfactual samples for each factual sample, and a *retrospective module* to retrospect the model prediction by comparing the counterfactual and factual samples. Extensive experiments on sentiment analysis (SA) and natural language inference (NLI) validate the effectiveness of our method.

1 Introduction

Language understanding (Ke et al., 2020) is a central theme of artificial intelligence (Chomsky, 2002), which empowers a wide spectral of applications such as sentiment evaluation (Feldman, 2013), commonsense inference (Bowman et al., 2015). The models are trained on labeled data to recognize the textual patterns closely correlated to different labels. Owing to the extraordinary representational capacity of deep neural networks, the models can well recognize the pattern and make prediction accordingly (Devlin et al., 2019). However, the cognitive ability of these data-driven models is still far from human beings due to lacking *counterfactual thinking* (Pearl, 2019).

Counterfactual thinking is a high-level cognitive ability beyond pattern recognition (Pearl, 2019). In addition to observing the patterns within factual

samples, counterfactual thinking calls for comparing the fact with imaginations, so as to make better decision. For instance, given a factual sample “*What do lawyers do when they die? Lie still.*”, the intuitive evaluation of its sentiment based on the textual patterns will recognize “*Lie still*” as an objective description of body posture which is *neutral*. By scrutinizing that the “*still*” could be intentionally postponed, we can imagine a counterfactual sample “*What do lawyers do when they die? Still lie.*” and uncover the negative sarcastic pun, whose sentiment is more accurate.

Recent work (Kaushik et al., 2019; Zeng et al., 2020) shows that incorporating counterfactual samples into model training improves the generalization ability. However, these methods follow the standard machine learning paradigm that uses the same procedure (e.g., a forward propagation) to make prediction in the testing phase. That is, making decision for testing samples according to their relative positions to the model decision boundary. The indiscriminate procedure focuses on the textual patterns occurred in the testing sample and treats all testing samples equally, which easily fails on hard samples (cf. Figure 1). On the contrary, humans can discriminate hard samples and ponder the decision with a rational system (Daniel, 2017), which imagines counterfactual and adjusts the decision.

The key to bridge this gap lies in imitating the counterfactual thinking ability of humans, i.e., learning a decision making procedure to serve for the testing phase. That is a procedure of: 1) constructing counterfactual samples for a target factual sample; 2) calling the trained language understanding model to make prediction for the counterfactual samples; and 3) comparing the counterfactual and factual samples to retrospect the model prediction. However, the procedure is non-trivial to achieve for two reasons: 1) the space of counterfactual sample is huge since any variant from the target factual

*Corresponding author.

sample can be a counterfactual sample. It is thus challenging to search for suitable counterfactual samples that can facilitate the decision making. 2) The mechanism of how we retrospect the decision is still unclear, making it hard to be imitated.

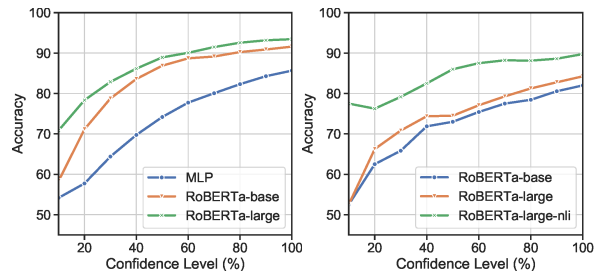
Towards the target, we propose a *Counterfactual Reasoning Model* (CRM), which is a two-phase procedure consisting a *generation module* and a *retrospection module*. In particular, given a factual sample in the testing phase, the generation module constructs representative counterfactual samples by imagining *what would the content be if the label of the sample is y* . To imitate the unknown retrospection mechanism of humans, we build the retrospection module as a carefully designed deep neural network that separately compares the latent representation and the prediction of the factual and counterfactual samples. The proposed CRM forms a general paradigm that can be applied to most existing language understanding models without constraint on the format of the language understanding task. We select two language understanding tasks: SA and NLI, and test CRM on three representative models for each task. Extensive experiments on benchmark datasets validate the effectiveness of CRM, which achieves performance gains ranging from 5.1% to 15.6%.

The main contributions are as follow:

- We propose the *Counterfactual Reasoning Model* to enlighten the language understanding model with counterfactual thinking.
- We devise a generation module and a retrospection module that are task and model agnostic.
- We conduct extensive experiments, which validate the rationality and effectiveness of the proposed method.

2 Pilot Study

Decisions are usually accompanied by confidence, a feeling of being wrong or right (Boldt et al., 2019). From the perspective of model confidence, we investigate the performance of language understanding models across different testing samples. We estimate the model confidence on a sample as the widely used *Maximum Class Probability* (MCP) (Corbière et al., 2019), which is the probability over the predicted class. A lower value of MCP means less confidence and “hard” sample. According to the value of MCP, we rank the testing



(a) Sentiment analysis (b) Natural language inference

Figure 1: Prediction performance of the language understanding models over testing samples at different confidence levels.

samples in ascending order and split them into ten groups, *i.e.*, confidence level from 1 to 10.

Figure 1 shows the performance of representative models over samples at different model confidence levels on the SA and NLI tasks (see Section 4.1 for model and dataset descriptions). From the figures, we can observe a clear increasing trend of classification accuracy as the confidence level increases from 1 to 10 in all cases. In other words, these models fail to predict accurately for the hard samples. It is thus essential to enhance the standard inference with a more precise decision making procedure.

3 Methodology

In this section, we first formulate the task of learning a decision making procedure for the testing phase (Section 3.1), followed by introducing the proposed CRM (Section 3.2) and the paradigm of building language understanding solutions with CRM (Section 3.3).

3.1 Problem Formulation

As discussed in the previous work (Wu et al., 2020; Li et al., 2020, 2019), language understanding tasks can be abstracted as a classification problem where the input is a text and the target is to make decision across a set of candidates of interests. We follow the problem setting with consideration of counterfactual samples (Kaushik et al., 2019; Liang et al., 2020), where the training data are twofold: 1) *factual samples* $\mathcal{T} = \{(\mathbf{x}, y)\}$ where $y \in [1, C]$ denotes the class or the target decision of the text; $\mathbf{x} \in \mathbb{R}^D$ is the latent representation of the text, which encodes the textual contents¹. 2) *counterfactual samples*

¹The input is indeed the plain text which is projected to a latent representation by an encoder (*e.g.*, a Transformer (Devlin et al., 2019)) in the cutting edge solutions. We omit the

$\mathcal{T}^* = \{(\mathbf{x}_c^*, c) | (\mathbf{x}, y) \in \mathcal{T}, c \in [1, C] \& c \neq y\}$ where (\mathbf{x}_c^*, c) is a counterfactual sample in class c corresponds to the factual sample (\mathbf{x}, y) ². We assume that a classification model (e.g., BERT (Devlin et al., 2019)) has been trained over the labeled data. Formally,

$$\hat{\theta} = \min_{\theta} \sum_{(\mathbf{x}, y) \in \mathcal{T} / \mathcal{T}^*} l(y, f(\mathbf{x} | \theta)) + \alpha \|\theta\|, \quad (1)$$

where $\hat{\theta}$ is the learned parameters of the model $f(\cdot)$; $l(\cdot)$ is a classification loss such as cross-entropy (Kullback, 1997), and α is a hyper-parameter to adjust the regularization.

The target is to build a decision making procedure to perform counterfactual reasoning when serving for the testing phase. Given a testing sample \mathbf{x} , the core is a policy of generating counterfactual samples and retrospecting the decision, which is formulated as:

$$\mathbf{y} = h(\mathbf{x}, \{\mathbf{x}^*\} | \eta, \hat{\theta}), \quad \{\mathbf{x}^*\} = g(\mathbf{x} | \omega),$$

$\mathbf{y} \in \mathbb{R}^C$ denotes the final prediction for the testing sample \mathbf{x} , which is a distribution over the classes; \mathbf{x}^* is one of the generated counterfactual samples for \mathbf{x} . The generation module $g(\cdot)$ parameterized by ω is expected to construct a set of representative counterfactual samples for the target factual sample, which provide signals for the retrospection module $h(\cdot)$ parameterized by η to retrospect the prediction $f(\mathbf{x} | \hat{\theta})$ given by the trained classification model. In particular, $h(\cdot)$ and $g(\cdot)$ will be learned from the factual and counterfactual training samples, respectively.

3.2 Counterfactual Reasoning Model

Figure 2 illustrates the process of CRM where the arrows in grey color represent the standard inference of trained classification model, and arrows in red color represent the retrospection with consideration of counterfactual samples.

3.2.1 Retrospection Module

We devise the retrospection module with one key consideration—distilling signals for making final decision by comparing both the latent representation and the prediction of the counterfactual samples with the factual sample. To achieve the target,

encoder for brevity since focusing on the decision making.

²Given the labeled factual sample, counterfactual samples can be constructed either manually (Kaushik et al., 2019) or automatically (Chen et al., 2020) by conducting minimum changes on \mathbf{x} to swap its label from y to c

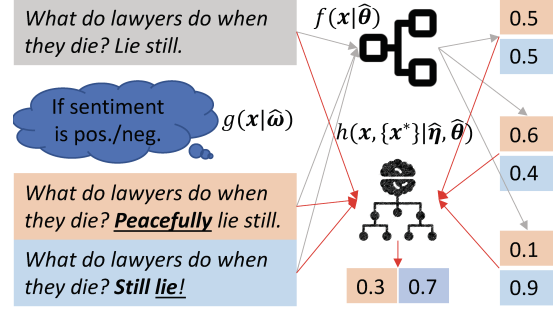


Figure 2: Illustration of the proposed CRM.

we devise three key building blocks for retrospection, which successively perform *representation comparison*, *prediction comparison*, and *fusion*. In particular, the module first compares the representation of each counterfactual sample with the factual sample; then compares their predictions accordingly; and fuses the comparison across the counterfactual samples.

Representation comparison. Given a pair of counterfactual sample \mathbf{x}^* and factual sample \mathbf{x} , we believe the signals meaningful for making final decision lie in the difference of the samples and how the difference affects the classification. To distill such signals, we devise the representation comparison block as $\mathbf{y}_\Delta = f(\mathbf{x} - \mathbf{x}^* | \hat{\theta})$, where $\mathbf{y}_\Delta \in \mathbb{R}^C$ denotes the prediction of the representation difference $\mathbf{x} - \mathbf{x}^*$ given by the trained classification model. Note that we leverage the trained model to enlighten how the content difference affects the classification since the model is trained to capture the connection between the textual patterns and the classes. It should be noted that we use a duplicate of the trained classification model for the representation comparison. That is to say, the training of the retrospection module will not affect the classification model.

Prediction comparison. To retrospect the prediction $f(\mathbf{x} | \hat{\theta})$, we devise a prediction comparison block to compare the predictions of each counterfactual and factual sample pair and distill patterns from $f(\mathbf{x} | \hat{\theta})$, $f(\mathbf{x}^* | \hat{\theta})$, and \mathbf{y}_Δ . Inspired by the success of convolutional neural network (CNN) in capture local-region patterns, the block is devised as a CNN, which is formulated as:

$$\mathbf{y}^* = \text{CNN}(f(\mathbf{x} | \hat{\theta}), f(\mathbf{x}^* | \hat{\theta}), \mathbf{y}_\Delta), \quad (2)$$

where \mathbf{y}^* denotes the retrospected prediction when comparing to \mathbf{x}^* . In particular, a *stack layer* first stacks the three predictions as a matrix, which serves as an “image” to facilitate “observing” pat-

terns. Formally, $\mathbf{Y} = \left[f(\mathbf{x}|\hat{\theta}), f(\mathbf{x}^*|\hat{\theta}), \mathbf{y}_\Delta \right]$ where $\mathbf{Y} \in \mathbb{R}^{C \times 3}$. \mathbf{Y} is then fed into an 1D convolution layer to capture the intra-class patterns across the predictions, which is formulated as:

$$\mathbf{H} = \sigma(\mathbf{Y} * \mathbf{F}), \mathbf{H}_{ij} = \sigma(\mathbf{Y}_{:i} \mathbf{F}_j), \quad (3)$$

where $\mathbf{F} \in \mathbb{R}^{3 \times K}$ denotes the filters in the convolution layer, and $\sigma(\cdot)$ is an activation function such as *GELU* (Hendrycks and Gimpel, 2016). $\mathbf{Y}_{:i}$ and \mathbf{F}_j represent the i -th row of \mathbf{Y} and the j -th column of \mathbf{F} , respectively. The filter \mathbf{F}_j can learn rules for conducting retrospection. For instance, a filter $[1, -1, 0]$ means deducting the prediction of the counterfactual sample from that of the factual sample. The output $\mathbf{H} \in \mathbb{R}^{C \times K}$ is then flattened as a vector and fed into a fully-connected (FC) layer to capture the inter-class patterns. Formally,

$$\mathbf{y}^* = \mathbf{W} \text{flatten}(\mathbf{H}) + \mathbf{b}, \quad (4)$$

where \mathbf{W} and \mathbf{b} are model parameters.

Fusion. The target is to fuse the retrospected predictions $\{\mathbf{y}^*\}$ into a final decision \mathbf{y} . Inspired by the success of pooling function in reading out patterns, we devise the block as $\mathbf{y} = \text{pooling}(\{\mathbf{y}^*\})$. As the fusion is performed after the pairwise comparison, we term it as *late fusion*.

Training. We update the parameters of the retrospection module by minimizing the classification loss over the factual training samples, which is:

$$\hat{\eta} = \min_{\eta} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{T}} l(\mathbf{y}, \mathbf{y}) + \lambda \|\eta\|. \quad (5)$$

where λ denotes the hyper-parameter to adjust the weight of the regularization term.

It should be noted that no existing research has uncovered the specific mechanism of retrospection in our brain, *i.e.*, the order of comparison and fusion is unclear. As such, we further devise two fusion strategies: *middle fusion* and *early fusion*, which performs fusion within the CNN, *i.e.*, during comparison, and before the CNN, respectively.

- *Middle fusion* performs aggregation between the convolution layer and the FC layer. This fusion first calculates the latent comparison signals \mathbf{H} for each pair of counterfactual and factual samples according to Equation 3. The aggregated signals $\text{pooling}(\{\mathbf{H}\})$ are then fed into the FC layer (Equation 4) to obtain the final decision \mathbf{y} .

- *Early fusion* aggregates the counterfactual samples before performing comparison, which is

formulated as $\tilde{\mathbf{x}}^* = \text{pooling}(\{\mathbf{x}^*\})$. In this way, the retrospection module is formulated as: $\mathbf{y} = \text{CNN}(f(\mathbf{x}|\hat{\theta}), f(\tilde{\mathbf{x}}^*|\hat{\theta}), f(\tilde{\mathbf{x}}^* - \mathbf{x}|\hat{\theta}))$. For all the three fusion methods, we can use either regular pooling function without parameter or parameterized pooling function (Ying et al., 2018) to enhance the expressiveness of the retrospection module. In our experiments, using a simple mean pooling achieves a performance that is comparable to the parameterized one in most cases (*cf.* Table 3).

3.2.2 Generation Module

The target is to construct counterfactual samples that are informative for retrospecting the decision on the target factual sample \mathbf{x} . As the task involves making decision among C candidate classes, we believe that the key to generate representative counterfactual samples lies in imagining “*what would the content be if the sample belongs to class c* ”, *i.e.*, generating C counterfactual samples $\{\mathbf{x}_c^*\}$. With the C classes as the targets, the searching space of samples can also be largely narrowed down. Toward this end, we devise the generation module with two main considerations: 1) decomposing the factual sample \mathbf{x} to distill contents irrelevant to the label of the sample $\mathbf{u} = d(\mathbf{x}|\omega)$; 2) injecting class c into \mathbf{u} to form the counterfactual sample \mathbf{x}_c^* .

Decomposition. To distill \mathbf{u} , we need to recognize the connection between the content of the factual sample and each class. We thus account for class representations in the decomposition function. To align the sample space of the generation module with the retrospection module $h(\cdot)$ and the classification model $f(\cdot)$, we extract the parameters from the prediction layer of the trained classification model as the class representations. In particular, we extract the mapping matrix $\mathbf{W} \in \mathbb{R}^{C \times D}$ where the c -th row corresponds to class c . Note that we assume that the prediction layer has the same dimensionality as the latent representation, which is a common setting in most cutting edge language understanding models. The decomposition function is devised as a CNN to capture both the intra-dimension and inter-dimension connections between the factual sample and the classes.

- *Stack layer.* The stack layer stacks the factual sample, class representations, and the element-wise product between sample and each class, which is formulated as: $\mathbf{X} = [\mathbf{x}, \mathbf{W}^T, \mathbf{x} \odot \mathbf{W}^T]$. $\mathbf{x} \odot \mathbf{W}^T \in \mathbb{R}^{D \times C}$ shed lights on how closely each dimension of \mathbf{x} connect to each class, where large

absolute value indicates closer connections.

- *Convolution layer.* This layer uses 1D horizontal filters to learn patterns of deducting class relevant contents from the factual sample, which is formulated as $\mathbf{h} = \text{pooling}(\sigma(\mathbf{X} * \mathbf{F}^g))$. $\mathbf{F}^g \in \mathbb{R}^{(2C+1) \times L}$ denotes the filters where L is the total number of filters. The output $\mathbf{h} \in \mathbb{R}^D$ is a hidden representation.

- *FC layers.* We use two FC layers to capture the inter-dimension connections. Formally, $\mathbf{u} = \mathbf{W}^2 \sigma(\mathbf{W}^1 \mathbf{h} + \mathbf{b}^1) + \mathbf{b}^2$, where $\mathbf{W}^2 \in \mathbb{R}^{D \times M}$, $\mathbf{W}^1 \in \mathbb{R}^{M \times D}$, $\mathbf{b}^2 \in \mathbb{R}^D$, and $\mathbf{b}^1 \in \mathbb{R}^M$ are learnable parameters. M is a hyper-parameter to adjust the complexity of the decomposition function. Note that we can stack more layers to enhance the expressiveness of the function, whereas using two layers according to the universal approximation theorem (Hornik, 1991).

We learn the parameters of the decomposition function from the counterfactual training samples by optimizing the following objective:

$$\min_{\omega} \sum_{(\mathbf{x}_c^*, c) \in \mathcal{T}^*} r(\mathbf{u}_c^*, \tilde{\mathbf{u}}_c) + \gamma l(c, f(\mathbf{x}_c^* - \mathbf{u}_c^* | \hat{\theta})) + r(\mathbf{u}, \tilde{\mathbf{u}}_c) + \gamma l(y, f(\mathbf{x} - \mathbf{u} | \hat{\theta})), \quad (6)$$

where $\mathbf{u}_c^* = d(\mathbf{x}_c^* | \omega)$ and $\mathbf{u} = d(\mathbf{x} | \omega)$ are the decomposition results of the counterfactual sample \mathbf{x}_c^* and the corresponding factual sample \mathbf{x} ; $\tilde{\mathbf{u}}_c = \frac{1}{2}(\mathbf{x} + \mathbf{x}_c^*)$ denotes the target value of the decomposition. The two terms $r(\cdot)$ and $l(\cdot)$ are Euclidean distance (Dattorro, 2010) and classification loss. By minimizing the two terms, we encourage the decomposition result: 1) to be close to the target value $\tilde{\mathbf{u}}_c$; and 2) if being deducted from the original sample (*e.g.*, $\mathbf{x} - \mathbf{u}$), the classification cannot be influenced. γ is a hyper-parameter to balance the two terms.

The rationality of setting $\tilde{\mathbf{u}}_c = \frac{1}{2}(\mathbf{x} + \mathbf{x}_c^*)$ as the target class irrelevant content of \mathbf{x} and \mathbf{x}_c^* comes from the parallelogram law (Nash, 2003). Note that this pair of samples belong to two different classes where a decision boundary (a hyperplane) lies between the two classes y and c . Considering that the sample \mathbf{x} corresponds to a vector in the hidden space, we can decompose the vector into two components that are orthogonal and parallel to the decision boundary, *i.e.*, $\mathbf{x}_c^* = \mathbf{o}_c^* + \mathbf{p}_c^*$ and $\mathbf{x} = \mathbf{o} + \mathbf{p}$. Since the two samples belong to different classes, their orthogonal components are in opposite directions and their addition will only retain the parallel components, which are irrelevant

to judging the class between y and c ³.

Injection. Accordingly, given a testing sample \mathbf{x} , we can inject the orthogonal components towards class c via $\mathbf{x}_c^* = 2 * d(\mathbf{x} | \hat{\omega}_c) - \mathbf{x}$, which is the imagined content of the sample if it belongs to class c . In this way, for each testing sample, we conduct the injection over all the classes and construct C counterfactual samples $\{\mathbf{x}_c^*\}$, which are then used in the retrospection module⁴.

3.3 Learning Paradigm with CRM

The existing work (Kaushik et al., 2019; Zeng et al., 2020) for language understanding typically follows the standard learning paradigm, *i.e.*, training a classification model over labeled data. Applying the proposed CRM indeed forms a new learning paradigm for constructing language understanding solutions. Algorithm 1 illustrates the procedure of the new paradigm.

Algorithm 1 Learning paradigm with CRM

Input: Training data $\mathcal{T}, \mathcal{T}^*$.
/ Training */*
1: Optimize Equation 1; ▷ Classification model training
2: Optimize Equation 6; ▷ Generation module training
3: Optimize Equation 5; ▷ Retrospection module training
4: Return $\hat{\theta}$, $\hat{\omega}_c$, and $\hat{\eta}$.
/ Testing */*
5: Calculate $f(\mathbf{x} | \hat{\theta})$; ▷ Classification model inference
6: **for** $c = 1 \rightarrow C$ **do**
7: $\mathbf{x}_c^* = 2 * g(\mathbf{x} | \hat{\omega}_c) - \mathbf{x}$; ▷ Generation
8: **end for**
9: Calculate $h(\mathbf{x}, \{\mathbf{x}_c^*\} | \hat{\eta}, \hat{\theta})$; ▷ Retrospection

4 Experiments

We conduct experiments on two representative language understanding tasks, SA and NLI, to answer the following research questions:

- **RQ1:** To what extent counterfactual reasoning improves language understanding?
- **RQ2:** How does the design of the retrospection module affect the proposed CRM?
- **RQ3:** How effective are the counterfactual samples generated by the proposed generation module?

4.1 Experiment Settings

Datasets. We adopt the same datasets in (Kaushik et al., 2019) for both tasks. The SA data are reviews

³Note that we normalize all samples to be unit vectors in the decomposition function. Moreover, inspired by (Parascandolo et al., 2018), we train a decomposition function for each class, *i.e.*, class-specific parameters $\hat{\omega}_c$

⁴The generation module consists of C decomposition functions $d(\mathbf{x} | \hat{\omega}_c)$ and the non-parametric injection function.

from IMDb, which are labeled as either positive or negative. For each factual review, the dataset contains a manually constructed counterfactual sample where the crowd workers are asked to manipulate the text to reverse the label with the constraint of no gratuitous change. NLI is a three-way classification task with two sentences as inputs and the target of detecting their relation within entailment, contradiction, and neutral. For each factual sample, four counterfactual samples are given, which are constructed by editing either the first or the second sentence with target relations different to the label of the factual sample.

Classification models. Owing to the extraordinary representational capacity of language model, fine-tuning pre-trained language model has become the emergent technique for solving language understanding tasks (Devlin et al., 2019). We select the widely used RoBERTa-base⁵ and RoBERTa-large⁶ for the consideration of the robustness of the RoBERTa (Liu et al., 2019) and our limited computation resources. For SA, we also test the classical Multi-Layer Perceptron (MLP) (Teney et al., 2020) with *tf-idf* text features (Schütze et al., 2008) as inputs. For NLI, we further test RoBERTa-large-nli⁷, which has been fine-tuned on the large-scale MultiNLI dataset (Williams et al., 2018).

Baselines. As the proposed CRM leverages counterfactual samples, we compare CRM with three representative methods using counterfactual samples in language understanding tasks: 1) **+CF** (Kaushik et al., 2019), which uses counterfactual samples as data augmentation for model training; 2) **+GS** (Teney et al., 2020), which compares the factual and counterfactual samples in model training through regularizing their gradients; and 3) **+CL** (Liang et al., 2020), which compares the factual and counterfactual samples through a contrastive loss. Moreover, we report the performance of the testing model under *Normal Training*, i.e., training over factual samples only.

Implementation. We implement the proposed CRM with PyTorch 1.7.0 based on Hugging Face Transformer⁸, which is released at: https://github.com/fulifeng/Counterfactual_Reasoning_Model. In all cases, we follow the setting of +CF for training the classification model, which is a standard fine-tuning in (Liu et al., 2019). We then use

⁵<https://huggingface.co/roberta-base>.

⁶<https://huggingface.co/roberta-large>.

⁷<https://huggingface.co/roberta-large-mnli>.

⁸<https://github.com/huggingface/transformers>.

adam (Kingma and Ba, 2014) with learning rate of 0.001 to optimize the retrospection module and the generation module. For the retrospection module, we set the number of filters in the convolution layer K as 10, the weight for regularization λ as 0. As to the generation module, we set the number of convolution filters as 10, the size of the hidden layer M as 256, and the weight for balancing Euclidean distance and classification loss γ as 15. We report the average classification accuracy over 5 different runs. For each repeat, we train the model with 20 epochs and select the model with the best performance on the validation set.

4.2 Performance Comparison (RQ1)

We first use the handcrafted counterfactual samples to demonstrate the effectiveness of counterfactual reasoning in the inference stage of language understanding model, which can be seen as using a golden standard generation module to provide counterfactual samples for the retrospection module. Note that we do not use the label of counterfactual samples in the testing set. Table 1 shows the performance of the compared methods on the two tasks. From the table, we observe that:

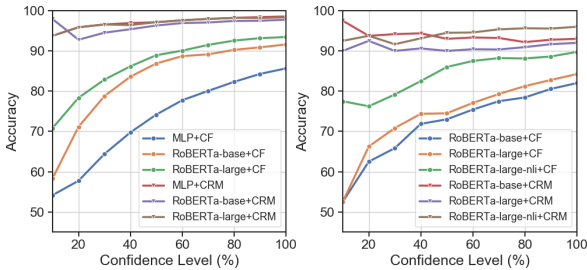
- **+CRM** largely outperforms all the baseline methods in all cases. As compared to +CF, the same classification model without CRM in the testing phase, +CRM achieves relative performance improvement up to 15.6%. The performance gain is attributed to the retrospection module, which justifies the rationality and effectiveness of incorporating counterfactual thinking into the inference stage of language understanding model. In other words, by comparing the factual sample with its counterfactual samples, the retrospection module indeed makes more accurate decisions.
- On the SA task, a huge gap (85.3 ↔ 93.4) lies in the performance of the shallow model MLP and the deep RoBERTa-base/RoBERTa-large. When applying +CRM, MLP achieves a performance that is comparable to the deep models. The result indicates that counterfactual reasoning can compensate for the disadvantages caused by the insufficient model representational capacity. In addition, the result reflects that CRM brings cognitive ability beyond recognizing textual patterns. If the retrospection module only facilitates capturing the correlation between textual patterns and classes, such simple model cannot bridge the

Sentiment Classification						
Backbone	Normal Training	+CF	+GS	+CL	+CRM	RI
MLP	86.9±0.5	85.3±0.3	84.6±0.4	-	98.6±0.2	15.6%
RoBERTa-base	93.2±0.6	92.3±0.7	92.2±0.9	91.8±1.1	97.5±0.3	5.7%
RoBERTa-large	93.6±0.6	93.4±0.4	93.1±0.5	94.1±0.4	98.2±0.3	5.1%
Natural Language Inference						
Backbone	Normal Training	+CF	+GS	+CL	+CRM	RI
RoBERTa-base	83.5±0.8	83.4±0.9	83.8±1.7	84.1±1.1	91.5±1.6	9.7%
RoBERTa-large	87.9±1.7	85.8±1.2	86.2±1.2	86.5±1.6	93.8±1.9	9.3%
RoBERTa-large-nli	89.4±0.7	88.2±1.0	87.2±1.4	88.2±1.0	94.4±1.2	7.1%

Table 1: Performance of the proposed CRM (Early Fusion) and baselines on the SA and NLI tasks. RI means the relative performance improvement achieved by +CRM over the classification model without CRM, *i.e.*, +CF.

huge gap of representational capacity between MLP and RoBERTa-large.

- The performance of baseline methods are comparable to each other in most cases, *i.e.*, incorporating counterfactual samples into model training does not necessarily improve the testing performance on factual samples. This result is consistent with (Kaushik et al., 2019), which is reasonable since these methods are devised for enhancing the generalization ability, especially for the out-of-distribution testing samples, which can sacrifice the performance on normal testing samples. Besides, the result indicates that training with counterfactual samples is insufficient for achieving counterfactual thinking, which reflects the rationality of enhancing the inference paradigm with a decision making procedure.



(a) Sentiment analysis (b) Natural language inference

Figure 3: Prediction performance of +CF and +CRM over testing samples at different confidence levels.

Performance on hard samples. Furthermore, we investigate whether the proposed CRM facilitate dealing with hard samples. Recall that we split the testing samples into 10 groups according to the confidence of the classification model, *i.e.*, +CF (*cf.* Section 2). We perform group-wise comparison between +CF and +CRM. Figure 3 shows the performance of all the classification models with +CF and +CRM. From the figures, 1) we observe that

the performance of +CRM is stable across different confidence levels, whereas the performance of the classification model shows a clear decreasing trend as the confidence level decreases from 10 to 1. The result indicates that the retrospection module is insensitive to the confidence of the classification model. 2) In all cases, +CRM achieves the largest performance gain at the first group with confidence level of 1, *i.e.*, the hardest group to the classification model. For instance, the improvement reaches 85.7% on the RoBERTa-base model for the NLI task. The large improvements further justifies the effectiveness of the retrospection module, *i.e.*, comparing the prediction of factual samples to counterfactual samples indeed facilitates dealing with hard samples.

Sentiment Classification		
Backbone	Implicit	+CRM
MLP	79.3±0.2	98.6±0.2
RoBERTa-base	94.7±0.6	97.5±0.3
RoBERTa-large	98.0±0.4	98.2±0.3
Natural Language Inference		
Backbone	Implicit	+CRM
RoBERTa-base	81.9±3.5	91.5±1.6
RoBERTa-large	87.4±2.2	93.8±1.9
RoBERTa-large-nli	88.8±1.6	94.4±1.2

Table 2: Performance comparison of implicit modeling (end-to-end model) and explicit modeling (CRM) of counterfactual thinking.

CRM V.S. implicit modeling. According to the uniform approximation theorem (Hornik, 1991), the CRM can also be approximated by a deep neural network. We thus investigate whether counterfactual thinking can be learned in an implicit manner. In particular, we evaluate a model that takes both the factual sample and counterfactual samples as inputs to make prediction for the fac-

tual one. Table 2 shows the performance, where we have the following observations: 1) The implicit modeling performs much worse than the proposed CRM in most cases, which justifies the effectiveness of the retrospection module and the rationality of modeling comparison explicitly. 2) On the NLI task, RoBERTa-base+CRM outperforms RoBERTa-large (implicit), which means that the superior performance of CRM is not because of the additional model parameters introduced by the retrospection module, but the explicit comparison between factual and counterfactual samples.

4.3 In-depth Analysis

Effects of retrospection module design (RQ2).

Note that the order of comparison and fusion in the retrospection mechanism of us humans is still unclear. We investigate how the fusion strategies influence the effectiveness of the proposed CRM. Table 3 shows the performance of CRM based on the early fusion (EF), late fusion (LF), and middle fusion (MF) on the NLI task. We omit the comparison on the SA task since the dataset only has one counterfactual sample for the target factual sample. For both EF and LF, we use the mean pooling as the pooling function. As to MF, we use a pooling function that is equipped with self-attention (Vaswani et al., 2017). The reasons of this setting are twofold: 1) using mean pooling will make LF and MF equivalent since the FC layer in the retrospection module is a linear mapping. Note that LF performs pooling after the FC layer, while the pooling function of MF is just before the FC layer. 2) The comparison between the LF and MF can thus shed light on whether parameterized pooling function can benefit the retrospection.

From the table, we can observe that, in most cases, CRM based on different fusion strategies achieve performance comparable to each other. It indicates that the retrospection is insensitive to the order of fusion and the comparison between counterfactual and factual samples. Considering that MF with mean pooling is equivalent to LF, we can see that the benefit of parameterized pooling function is limited. In particular, MF only performs better than LF on one of the three testing models.

Effects of generation module (RQ3). We then investigate whether the proposed generation module constructs useful counterfactual samples for retrospection. We train and test the retrospection module (using EF) with the generated samples on

RoBERTa-large on the SA task. We omit the experiments of other settings for saving computation resources. In this way, the model achieves an accuracy of 94.5 which is better than +CF (93.4) but worse than +CRM with manually constructed counterfactual samples (98.2) (*cf.* Table 1). The result indicates that the generated samples indeed facilitate the retrospection while the generation quality can be further improved. Moreover, on the testing samples at confidence level of 1, using the generated samples achieves an accuracy of 81.3 which is much better than +CF (70.8) (*cf.* Figure 3). The generated samples indeed benefit the decision making over hard testing samples.

5 Related Work

Counterfactual sample. Constructing counterfactual samples has become an emergent data augmentation technique in natural language processing, which has been used in a wide spectral of language understanding tasks, including SA (Kaushik et al., 2019; Yang et al., 2020), NLI (Kaushik et al., 2019), named entity recognition (Zeng et al., 2020) question answering (Chen et al., 2020), dialogue system (Zhu et al., 2020), vision-language navigation (Fu et al., 2020). Beyond data augmentation under the standard supervised learning paradigm, a line of research explores to incorporate counterfactual samples into other learning paradigms such as adversarial training (Zhu et al., 2020; Fu et al., 2020; Teney et al., 2020) and contrastive learning (Liang et al., 2020). This work lies in an orthogonal direction that incorporates counterfactual samples into the decision making procedure of model inference.

Counterfactual inference. A line of research attempts to enable deep neural networks with counterfactual thinking by incorporating counterfactual inference (Yue et al., 2021; Wang et al., 2021; Niu et al., 2021; Tang et al., 2020; Feng et al., 2021). These methods perform counterfactual inference over the model predictions according to a pre-defined causal graph. Due to the requirement of causal graph, such methods are hard to be generalized to different tasks. Our method does not suffer from such limitation since working on the counterfactual samples which can be generated without a comprehensive causal graph.

Hard sample. A wide spectral of machine learning techniques are related to dealing with the hard samples in language understanding. For instance,

Backbone	+CF	EF	<i>RI</i>	LF	<i>RI</i>	MF	<i>RI</i>
RoBERTa-base	83.4±0.9	91.5±1.6	9.7%	92.8±1.8	11.3%	89.6±2.0	7.4%
RoBERTa-large	85.8±1.2	93.8±1.9	9.3%	95.3±0.7	11.1%	93.4±1.7	8.9%
RoBERTa-large-nli	88.2±1.0	94.4±1.2	7.1%	93.8±0.4	6.4%	94.7±1.3	7.4%

Table 3: Performance of the proposed CRM based on early fusion (EF), late fusion (LF), or middle fusion (MF) on the NLI task. *RI* represents the relative performance improvement over the +CF method.

adversarial training (Khashabi et al., 2020) enhances the model robustness against perturbations and attacks, which are hard samples for normally trained models. Debiased training (Tu et al., 2020; Utama et al., 2020) eliminates the spurious correlation or bias in training data to enhance the generalization ability and deal with out-of-distribution samples. In addition to the training phase, a few inference techniques might improve the model performance on hard samples, including posterior regularization (Srivastava et al., 2018) and causal inference (Yu et al., 2020; Niu et al., 2021). However, both techniques require domain knowledge such as prior or causal graph tailored for specific applications. On the contrary, this work provides a general paradigm that can be used for most language understanding tasks.

6 Conclusion

In this work, we pointed out the issue of standard inference of existing language understanding models. We proposed a Counterfactual Reasoning Model which empowers the trained model with a high-level cognitive ability, counterfactual thinking. By applying the proposed CRM, we formed a new paradigm for building language understanding solutions. We conducted extensive experiments, which validate the effectiveness of our proposal, especially in dealing with hard samples.

This work opens up a new research direction about the decision making procedure in testing phase. In the future, we will explore sequential decision procedure to resolve the constraint on the number of constructed counterfactual samples. In addition, we will investigate generation module for language understanding with unsupervised generative techniques (Sauer and Geiger, 2021).

Acknowledgments

This research is supported by the Sea-NExT Joint Lab, Singapore MOE AcRF T2, National Natural Science Foundation of China (U19A2079) and National Key Research and Development Program of China (2020AAA0106000). Our thanks also go

to all the anonymous reviewers for their valuable suggestions.

References

- Annika Boldt, Anne-Marika Schiffer, Florian Waszak, and Nick Yeung. 2019. Confidence predictions affect performance confidence and neural preparation in perceptual decision making. *Scientific reports*, 9(1):1–17.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. 2020. Counterfactual samples synthesizing for robust visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10800–10809.
- Noam Chomsky. 2002. *Syntactic structures*. Walter de Gruyter.
- Charles Corbière, Nicolas Thome, Avner Bar-Hen, Matthieu Cord, and Patrick Pérez. 2019. Addressing failure prediction by learning model confidence. In *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*, pages 2898–2909. Curran Associates, Inc.
- Kahneman Daniel. 2017. Thinking, fast and slow.
- Jon Dattorro. 2010. *Convex optimization & Euclidean distance geometry*. Lulu. com.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186. ACL.
- Ronen Feldman. 2013. Techniques and applications for sentiment analysis. *Communications of the ACM*, 56(4):82–89.
- Fuli Feng, Weiran Huang, Xin Xin, Xiangnan He, and Tat-Seng Chua. 2021. Should graph convolution trust neighbors? a simple causal inference method. In *SIGIR*.

- Tsu-Jui Fu, Xin Eric Wang, Matthew F Peterson, Scott T Grafton, Miguel P Eckstein, and William Yang Wang. 2020. Counterfactual vision-and-language navigation via adversarial path sampler. In *European Conference on Computer Vision*, pages 71–86. Springer.
- Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Kurt Hornik. 1991. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257.
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2019. Learning the difference that makes a difference with counterfactually-augmented data. In *International Conference on Learning Representations*.
- Pei Ke, Haozhe Ji, Siyang Liu, Xiaoyan Zhu, and Minlie Huang. 2020. Sentilare: Linguistic knowledge enhanced language representation for sentiment analysis. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6975–6988.
- Daniel Khashabi, Tushar Khot, and Ashish Sabharwal. 2020. More bang for your buck: Natural perturbation for robust question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 163–170.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Solomon Kullback. 1997. *Information theory and statistics*. Courier Corporation.
- Xiaoya Li, Jingrong Feng, Yuxian Meng, Qinghong Han, Fei Wu, and Jiwei Li. 2020. A unified mrc framework for named entity recognition. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5849–5859.
- Xiaoya Li, Fan Yin, Zijun Sun, Xiayu Li, Arianna Yuan, Duo Chai, Mingxin Zhou, and Jiwei Li. 2019. Entity-relation extraction as multi-turn question answering. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1340–1350.
- Zujie Liang, Weitao Jiang, Haifeng Hu, and Jiaying Zhu. 2020. Learning to contrast the counterfactual samples for robust visual question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3285–3292.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. *Roberta: A robustly optimized bert pretraining approach*. *arXiv e-prints*.
- Alan Nash. 2003. A generalized parallelogram law. *The American mathematical monthly*, 110(1):52–57.
- Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. 2021. Counterfactual vqa: A cause-effect look at language bias. In *CVPR*.
- Giambattista Parascandolo, Niki Kilbertus, Mateo Rojas-Carulla, and Bernhard Schölkopf. 2018. Learning independent causal mechanisms. In *International Conference on Machine Learning*, pages 4036–4044. PMLR.
- Judea Pearl. 2019. The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3):54–60.
- Axel Sauer and Andreas Geiger. 2021. Counterfactual generative networks. *The International Conference on Learning Representations*.
- Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. 2008. *Introduction to information retrieval*, volume 39. Cambridge University Press Cambridge.
- Shashank Srivastava, Igor Labutov, and Tom Mitchell. 2018. Zero-shot learning of classifiers from natural language quantification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 306–316.
- Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. 2020. Long-tailed classification by keeping the good and removing the bad momentum causal effect. *NeurIPS*, 33.
- Damien Teney, Ehsan Abbasnejad, and Anton van den Hengel. 2020. Learning what makes a difference from counterfactual examples and gradient supervision. *arXiv preprint arXiv:2004.09034*.
- Lifu Tu, Garima Lalwani, Spandana Gella, and He He. 2020. An empirical study on robustness to spurious correlations using pre-trained language models. *Transactions of the Association for Computational Linguistics*, 8:621–633.
- Prasetya Ajie Utama, Nafise Sadat Moosavi, and Iryna Gurevych. 2020. Mind the trade-off: Debiasing nlu models without degrading the in-distribution performance. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8717–8729.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30:5998–6008.
- Wenjie Wang, Fuli Feng, Xiangnan He, Hanwang Zhang, and Tat-Seng Chua. 2021. ”click” is not equal to” like”: Counterfactual recommendation for mitigating clickbait issue. In *SIGIR*.

- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122.
- Wei Wu, Fei Wang, Arianna Yuan, Fei Wu, and Jiwei Li. 2020. Corefqa: Coreference resolution as query-based span prediction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6953–6963.
- Linyi Yang, Eoin Kenny, Tin Lok James Ng, Yi Yang, Barry Smyth, and Ruihai Dong. 2020. Generating plausible counterfactual explanations for deep transformers in financial text classification. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6150–6160.
- Rex Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L Hamilton, and Jure Leskovec. 2018. Hierarchical graph representation learning with differentiable pooling. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 4805–4815.
- Sicheng Yu, Yulei Niu, Shuohang Wang, Jing Jiang, and Qianru Sun. 2020. Counterfactual variable control for robust and interpretable question answering. *arXiv preprint arXiv:2010.05581*.
- Zhongqi Yue, Tan Wang, Hanwang Zhang, Qianru Sun, and Xian-Sheng Hua. 2021. Counterfactual zero-shot and open-set visual recognition. In *CVPR*.
- Xiangji Zeng, Yunliang Li, Yuchen Zhai, and Yin Zhang. 2020. Counterfactual generator: A weakly-supervised method for named entity recognition. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7270–7280.
- Qingfu Zhu, Weinan Zhang, Ting Liu, and William Yang Wang. 2020. Counterfactual off-policy training for neural dialogue generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3438–3448.