

A Neural Divide-and-Conquer Reasoning Framework for Image Retrieval from Linguistically Complex Text

Yunxin Li¹, Baotian Hu^{1*}, Yuxin Ding¹, Lin Ma², Min Zhang¹

¹Harbin Institute of Technology, Shenzhen, China, ²Meituan, Beijing

{hubaotian, yxding, zhangmin2021}@hit.edu.cn

liyunxin987@163.com, forest.linma@gmail.com

Abstract

Pretrained Vision-Language Models (VLMs) have achieved remarkable performance in image retrieval from text. However, their performance drops drastically when confronted with linguistically complex texts that they struggle to comprehend. Inspired by the Divide-and-Conquer (Smith, 1985) algorithm and dual-process theory (Groves and Thompson, 1970), in this paper, we regard linguistically complex texts as compound proposition texts composed of multiple simple proposition sentences and propose an end-to-end Neural Divide-and-Conquer Reasoning framework, dubbed NDCR. It contains three main components: 1) *Divide*: a proposition generator divides the compound proposition text into simple proposition sentences and produces their corresponding representations, 2) *Conquer*: a pretrained VLMs-based visual-linguistic interactor achieves the interaction between decomposed proposition sentences and images, 3) *Combine*: a neural-symbolic reasoner combines the above reasoning states to obtain the final solution via a neural logic reasoning approach. According to the dual-process theory, the visual-linguistic interactor and neural-symbolic reasoner could be regarded as analogical reasoning System 1 and logical reasoning System 2. We conduct extensive experiments on a challenging image retrieval from contextual descriptions data set. Experimental results and analyses indicate NDCR significantly improves performance in the complex image-text reasoning problem. Code link: <https://github.com/YunxinLi/NDCR>.

1 Introduction

Image-text retrieval tasks have made remarkable progress owing to pretrained Vision-Language Models (VLMs) such as LXMERT (Tan and Bansal, 2019), UNITER (Chen et al., 2020), OSCAR (Li et al., 2020b; Zhang et al., 2021), ViLBERT (Lu et al., 2019), CLIP (Radford et al., 2021),

* Corresponding author.

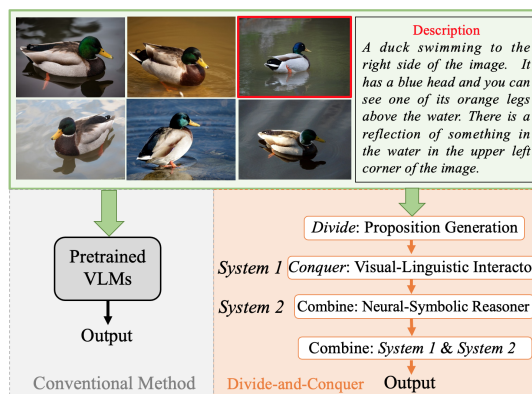


Figure 1: An example from the IMAGECODE (Krojer et al., 2022) data set, where the description is linguistically complex and images are minimally contrastive. The target image is in red and others are incorrect frames. The bottom part depicts the conventional method and the neural divide-and-conquer reasoning framework.

and many others. These VLMs are usually trained on the large-scale short text-image corpus by cross-modal semantic alignment methods. They are capable of essential perceptual computing capability and excel at retrieving images from sentences with few objects and simple linguistic, e.g., “*There is a duck swimming in the pond*”. However, when pretrained VLMs meet the case of retrieving the accurate image from similar candidates based on a linguistically complex text, as the example shown in Figure 1, previous works (Krojer et al., 2022; Talmor et al., 2021a; Thrush et al., 2022) show that they struggle to understand the elaborate description and perform complex cross-modal reasoning.

According to the dual-process theory for human thinking (Groves and Thompson, 1970; Evans, 2003; Pelaccia et al., 2011), human brains contain two thinking systems: System 1 performs analogical reasoning well, which is fast yet unconscious; System 2 is capable of abstract logical reasoning, which is slow yet conscious and well-suitable for complex reasoning problems. The theory could

also hold for the image-text retrieval tasks, and the widely adopted models (e.g., VLMs) focus on analogical reasoning as System 1 based on the analysis of deep learning networks (Bengio, 2017, 2019; Bengio et al., 2021). For the linguistically complex description that contains multiple conditions, they have inferior performance, and we need to introduce logical reasoning System 2 more to cover and logically incorporate the scattered information in the description based on System 1. Inspired by the above investigations and classical Divide-and-Conquer (Smith, 1985) algorithm, we design an end-to-end Neural Divide-and-Conquer Reasoning framework named NDCR. As shown in Figure 1, our key idea is to regard the complex description as compound proposition text and solve the challenging retrieval problem in three steps: divide, conquer, and combine.

Specifically, **Divide**: NDCR first utilizes a proposition generator to divide the complex compound text and produce the global representation of simple proposition sentences with visually printing them. **Conquer**: we devise a visual-linguistic interactor to achieve the interaction between decomposed proposition sentences and images, which resembles System 1. It uses the Transformer (Vaswani et al., 2017)-based contextual interactor to achieve the inter-learning of different proposition-image pairs. Considering the incorrectness or information loss of simple proposition representation, we also present a modifier to incorporate the context reasoning information to improve their cross-modal reasoning states. **Combine**: we design a learnable neural-symbolic reasoner to integrate reasoning information of simple propositions logically. It first employs a negation executor to obtain a simple proposition sentence’s negational reasoning hidden state and corresponding confidence score. Then, we use the global reasoning information of compound proposition text as the query signal to perform the conjunction operation across simple propositions and their negational information. Finally, as shown in Figure 1, we also combine the inferred results of the neural-symbolic reasoner (resembles System 2) and visual-linguistic interactor (resembles System 1) to obtain the final solution. In this way, the whole framework integrate the capabilities of Systems 1 and 2 to obtain better performance.

We conduct extensive experiments on a large-scale image retrieval from contextual descriptions

data set, IMAGECODE (Krojer et al., 2022). The experimental results indicate that NDCR achieves the state-of-the-art performance and the ablation and case studies verify the effectiveness of different modules.

Our contributions are as follows:

- We propose a divide-and-conquer reasoning framework for image retrievals from linguistically complex text, where we first attempt to combine the perceptually analogical reasoning System 1 and neural-symbolic logic reasoning System 2 to solve the complex multimodal reasoning problem.
- We design a proposition generator capable of producing the global representation of decomposed simple proposition sentences for linguistically complex texts and visually printing them as text.
- Experimental results indicate our approach remarkably improves the performance, and we obtain the first place on the leaderboard¹. Ablation and case studies confirm the effectiveness of introducing and combining logical reasoning System 2 based on System 1.

2 Related Works

Pretrained Vision-Language Models for Cross Modal Matching. Owing to the success of Transformer (Vaswani et al., 2017) architecture equipped with pretrain-finetuning (Erhan et al., 2010) learning method, pretrained VLMs have made a remarkable performance in cross-modal matching or reasoning tasks (Talmor et al., 2021b), especially image-text retrieval. Early pretrained VLMs utilize BERT (Devlin et al., 2019)-like single encoder architecture to encode and fuse the image-text information, then perform image-text reasoning such as ViLBERT (Lu et al., 2019), VisualBERT (Li et al., 2019), and Oscar (Li et al., 2020b). In addition, dual-encoder architecture such as CLIP (Radford et al., 2021), and ALBERT (Li et al., 2021), performs better than single-encoder architecture on image-text matching tasks and is widely used in industry because of its efficiency.

Divide-and-Conquer for Question Answering. The divide-and-conquer algorithm (Smith, 1985)

¹<https://mccgill-nlp.github.io/imagecode>

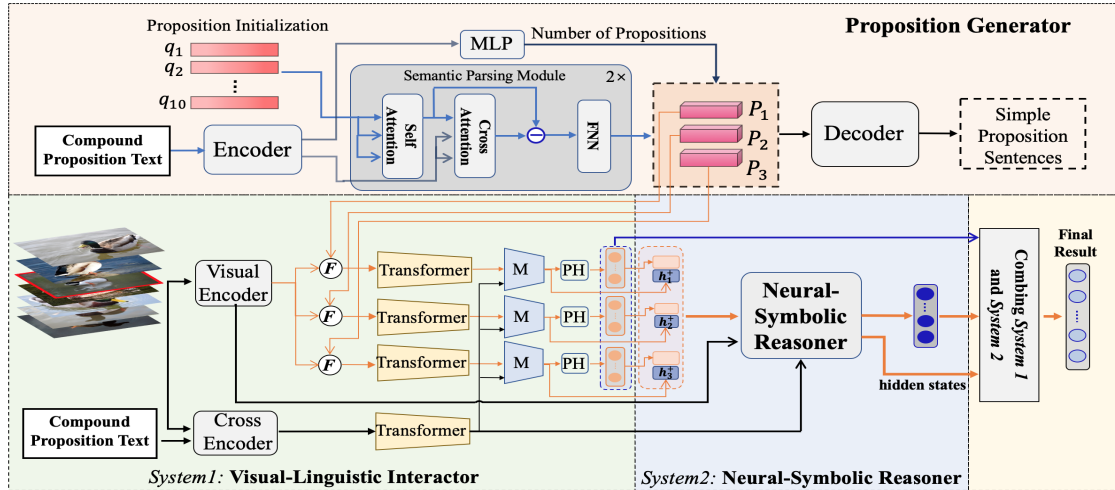


Figure 2: The overall architecture of neural divide-and-conquer reasoning framework.

aims to divide the complex problem into multiple simple problems and then combine the sub-problem results to achieve the final solution. This idea has been used in complex question-answering tasks in the natural language processing area. Zhang et al. (2019) proposed to utilize the decomposition of complex questions for semantic parsing. Min et al. (2019) adopt the question decomposition and rescoring method to perform multi-hop reading comprehension, which makes the reasoning path interpretable and robust. Wolfson et al. (2022) utilized the QDMR structures of complex questions to conduct the decompose-synthesize text-to-SQL transformation. Previous pipeline approaches may lead to error cascades in the upper inference process due to the incompleteness or error of decomposed text. The image-text retrieval task has strict requirements on the correctness of text semantic understanding, thus we propose an end-to-end divide-and-conquer method for alleviating the error cascade issue via the whole learning process.

Dual-Process Theory. The dual-process theory shows that human brains have two different thinking Systems. System 1 performs analogical reasoning, and System 2 performs conscious logical reasoning. Combining this theory with practical tasks, some researchers designed various approaches. Mittal et al. (2017) believed that combining vector space models with external knowledge graphs could be regarded as thinking ‘fast’ in vector space along with thinking ‘slow’ and ‘deeply’ by reasoning over the knowledge graph. Anthony et al. (2017) also proposed to use a deep learning network with a tree search engine as System 1 and System 2, respectively, for sequential decision-

making problems. Bengio (2017, 2019) advocated the design of a conscious network to achieve the leap from System 1 to System 2. Liu et al. (2022) designed a neural-symbolic system for natural language understanding tasks, which combines the explicit symbolic calculation-based System 2 and fast deep learning network-based System 1. For complex multi-modal reasoning problem, e.g., image retrieval from linguistically complex text, humans usually combine System 1 and System 2 to obtain the final solution. However, current methods relying mainly on deep learning networks resemble System 1 and lack the logical reasoning capability, thus suffering from image-text reasoning with the complex description. In this light, we make the first attempt to combine System 1 and System 2 to tackle this issue by designing a neural divide-and-conquer reasoning framework. We introduce a neural-symbolic reasoner in System 2 to conduct the logical operation. The overall framework contains analogical and logical reasoning as humans think, making appreciable gains.

3 Method

3.1 Overview

Image retrieval from contextual descriptions (Krojer et al., 2022) aims to infer the correct image given a linguistically complex text $Y = (y_1, \dots, y_N)$ and similar images $I = (I_1, \dots, I_L)$, where y_i , N , I_i , and L represent the i th token, the total length of text, i th image, and the number of images, respectively. We propose a novel divide-and-conquer reasoning framework to tackle such a task. It consists of three components, namely,

Proposition Generator, Visual-Linguistic Interactor, and Neural-Symbolic Reasoner, which are coupled and trained in an end-to-end manner. Specifically, the proposition generator divides the complex description into multiple proposition sentences, allowing it to convert the complex matching problem to simple ones. Afterwards, the visual-linguistic interactor achieves the interaction between decomposed proposition sentences and images, resembling System 1, to perform the essential analogical reasoning. Subsequently, the neural-symbolic reasoner that relies on the reasoning state output by the visual-linguistic interactor resembles System 2 to perform logical reasoning. Finally, we also combine the output results of System 1 and System 2 to obtain the final solution.

3.2 Proposition Generator

The proposition generator is a sequence-to-sequence model based on the pretrained language model BART. As shown in Figure 2, it employs the encoder to obtain the text representation $\mathbf{H}_Y = (h_{cls}, h_{y_1}, \dots, h_{y_N})$ where h_{y_i} represents the i th token hidden state. Subsequently, we design a two-layer semantic parsing module to gain the global representation of simple proposition sentences. Concretely, we set the maximum number of simple propositions to 10 and randomly initialize them. The initial vectors are fed to the semantic parsing module to interact with the compound text representation. Take the first layer as an example; the calculation process is following,

$$\begin{aligned} \mathbf{h}_s^T &= \text{Self-Attention}(\mathbf{h}^I), \\ \mathbf{h}_c^T &= \text{Cross-Attention}(\mathbf{h}_s^T, \mathbf{H}_Y), \\ \mathbf{h}_f^T &= \text{FNN}(\mathbf{h}_c^T - \mathbf{h}_s^T), \end{aligned} \quad (1)$$

where \mathbf{h}^I is the randomly initial proposition representations. Attention and FNN calculation subnetworks are identical to the transformer (Vaswani et al., 2017) architecture. Different from the transformer, we let the output of Cross-Attention layer subtract the output of Self-Attention layer, aiming to achieve information differences across propositions.

By doing the same two-layer calculation, we obtain ten global hidden states of simple propositions. Due to context containing different numbers of simple proposition, we use a MLP to predict the target number of simple proposition sentences. It only attends to the global hidden state h_{cls} of compound proposition text. Suppose that the predicted

number M of simple propositions is 3 (same as Figure 2), we adopt the first-three hidden states of the semantic parsing module as the global representation of the targeted simple proposition. As shown in Figure 2, for explaining what simple propositions represent, we also use the decoder of BART to generate the simple proposition sentence with only attending to their global representations.

3.3 System 1: Visual-Linguistic Interactor

After obtaining the global representations of simple proposition sentences, we introduce the visual-linguistic interactor to mine the interaction of image-proposition pairs. Specifically, we use a pretrained visual encoder to obtain the image encoding representations $\mathbf{H}_I = (\mathbf{h}_{I_1}, \dots, \mathbf{h}_{I_L})$ and fuse them with the simple proposition representation via the dot-product way (as the “ F ” shown in Figure 2). The two-modal fusion process is $\mathbf{H}(p) = \lambda \cdot \text{Norm}(\mathbf{P}) \cdot \text{Norm}(\mathbf{H}_I)$, where λ is the hyperparameter set to enlarge the scale of fused vectors. We denote the fused sequence representation of proposition-image pairs to $\mathbf{H}(p) = (\mathbf{H}(p_1), \dots, \mathbf{H}(p_M))$ where $\mathbf{H}(p_1)$ indicates the sequential representation of first proposition combined with images.

Then, we employ a two-layer transformer to perform the contextual information interaction for fused sequential representations $\mathbf{H}(p)$ and obtain the initial reasoning states of simple proposition on images. Considering the incorrectness or information loss of simple proposition representation obtained by the proposition generator, we introduce a MLP-based modifier to incorporate the reasoning state of compound proposition text to enhance previous initial reasoning states of simple propositions. The whole process is performed as Eq. 2,

$$\begin{aligned} \mathbf{H}_P^{S_1} &= \text{Transformer}(\mathbf{H}(p) + PE), \\ \mathbf{H}_C^{sg} &= \text{Transformer}(\mathbf{H}_C + PE), \\ \mathbf{H}^{S_1} &= \mathbf{W}^{M_1} \text{ReLU}(\mathbf{W}^{M_2} [\mathbf{H}_P^{S_1}, \mathbf{H}_C^{sg}]), \end{aligned} \quad (2)$$

where \mathbf{H}_C indicates the fusion information of the compound proposition text and images, gained by the cross-modal encoder (arr. cross encoder as shown in Figure 2). $\mathbf{W}^{M_1} \in \mathbb{R}^{2d \times d}$ and $\mathbf{W}^{M_2} \in \mathbb{R}^{2d \times 2d}$ are learnable parameters. Before feeding $\mathbf{H}(p)$ into the transformer, we introduce the learnable position embeddings PE to facilitate it pay attention to the contextual information across images. After obtaining the final reasoning state $\mathbf{H}^{S_1} = (\mathbf{h}_1^+, \dots, \mathbf{h}_M^+)$ of simple propositions

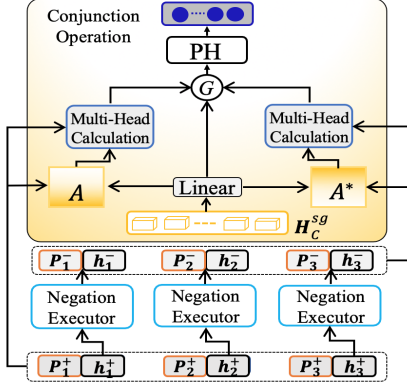


Figure 3: The detailed workflow of Neural-Symbolic Reasoner. It contains the underlying negation executor and upper conjunction operation.

in System 1, we adopt a linear prediction head to produce the confidence score of each proposition to images, which are defined as $P^{S_1} = (p_1^+, \dots, p_M^+)$ and $p_M^+ \in \mathbb{R}^{1 \times L}$.

3.4 System 2: Neural-Symbolic Reasoner

For complex reasoning problems, the logical reasoning process usually plays a more significant role for intelligent machines and human reasoning (Benigio, 2019), which the visual-linguistic interactor is not capable of. Instead of combining the inferring results in System 1 via rule-based methods such as mean pooling, inspired by Shi et al. (2020); Chen et al. (2021), we devise a learnable Neural-Symbolic Reasoner (NSR) to perform logical reasoning based on System 1 as shown in Figure 2. As depicted in Figure 3, it contains a negation executor to obtain the negational reasoning states and a conjunction operation to acquire the result of logical reasoning with attention to the positive and negational reasoning information.

Negation Executor. The negation executor is a module that takes the reasoning state of a simple proposition as input and produces the corresponding reasoning state of its negation as output. Its aim is to obtain useful cross-modal reasoning states for the negation of a proposition. We regard \mathbf{H}^{S_1} as the positive reasoning state and use a two-layer MLP with the ReLU activation function to obtain the negational reasoning state. The calculation process is given in Eq. 3,

$$\text{NEG}(\mathbf{H}^{S_1}) = W_2^n \text{ReLU}(W_1^n \mathbf{H}^{S_1} + b_1^n) + b_2^n, \quad (3)$$

where $W_2^n, W_1^n \in \mathbb{R}^{d \times d}$, $b_1^n, b_2^n \in \mathbb{R}^{1 \times d}$ are learnable parameters. We define the output of negation executor to $\mathbf{H}^N = (\mathbf{h}_1^-, \dots, \mathbf{h}_M^-)$, contrast

to \mathbf{H}^{S_1} . The negational proposition has a different cross-modal reasoning state \mathbf{H}^N than the corresponding positive proposition \mathbf{H}^{S_1} . We use the same linear prediction head as System 1 to produce the corresponding confidence score on images, which are presented to $P^N = (p_1^-, \dots, p_M^-)$. To make the negation executor effective, we will define a negational feedback loss to locally optimize it.

Conjunction Operation. Firstly, we define a new joint representation that incorporates reasoning hidden states and corresponding confidence scores as the initial state of conjunction operation. The process is presented in Eq. 4,

$$\begin{aligned} \mathbf{P}_i^+ &= \text{Softmax}(p_i^+) \cdot \mathbf{H}_I, i = 1, \dots, M, \\ \mathbf{H}_{p_i^+}^{ns} &= [\mathbf{P}_i^+, \mathbf{h}_i^+], i = 1, \dots, M, \end{aligned} \quad (4)$$

where $[\cdot, \cdot]$ indicates the concat calculation and \mathbf{H}_I is the representation of images. $\mathbf{H}_{p_i^+}^{ns}$ represents the positive joint representation of i th proposition. We use the same calculation method as Eq. 4 to obtain the initialized negational representation $\mathbf{H}_{p_i^-}^{ns}$. Then, we utilize the reasoning state of compound proposition text \mathbf{H}_C^{sg} (Eq. 2) as the signal to drive the conjunction calculation via the method of multi-head attention equipped with gate fusion, as shown in Figure 3. The whole calculation process is presented in Eq. 5,

$$\begin{aligned} \mathbf{H}^+ &= \text{MultiHead}(W^s \mathbf{H}_C^{sg}, \mathbf{H}_{p_i^+}^{ns}), \\ \mathbf{H}^- &= \text{MultiHead}(W^s \mathbf{H}_C^{sg}, \mathbf{H}_{p_i^-}^{ns}), \\ g^+ &= W^g [\mathbf{H}^+, W^s \mathbf{H}_C^{sg}] + b^g, \\ g^- &= W^g [\mathbf{H}^-, W^s \mathbf{H}_C^{sg}] + b^g, \\ \mathbf{H}^f &= W^{S_2} (g^+ \mathbf{H}^+ + g^- \mathbf{H}^-), \end{aligned} \quad (5)$$

where $W^s \in \mathbb{R}^{2d \times 2d}$, $W^g \in \mathbb{R}^{1 \times 4d}$, $W^{S_2} \in \mathbb{R}^{2d \times d}$ are the learnable parameters and $\mathbf{H}^f \in \mathbb{R}^{1 \times L \times d}$. We also utilize another linear prediction head to obtain the final confidence score of neural-symbolic reasoner, which is defined as $P^{S_2} \in \mathbb{R}^{1 \times L}$.

3.5 Combining System 1 and System 2

In addition, we combine inferring confidence scores in System 1 and System 2 to obtain the final solution, achieving the complementarity of System 1 and System 2. First, we need to acquire the whole

representation of \mathbf{H}^{S_1} and \mathbf{H}^f as follows:

$$\begin{aligned}\mathbf{H}_W^f &= (W^l \mathbf{H}^f + b^l)^T \mathbf{H}^f, \\ \mathbf{H}_W^{S_1} &= (W^l \mathbf{H}^{S_1} + b^l)^T \mathbf{H}^{S_1},\end{aligned}\quad (6)$$

where $W^l \in \mathbb{R}^{d \times 1}$, b^l are learnable parameters. $\mathbf{H}_W^{S_1} = (\mathbf{h}_{w_1}^+, \dots, \mathbf{h}_{w_M}^+) \in \mathbb{R}^{M \times d}$ and $\mathbf{H}_W^f \in \mathbb{R}^{1 \times d}$ are used to gain the final solution via Eq. 7,

$$\begin{aligned}\mathbf{h}_c &= \sum_{j=0}^M (W^a \mathbf{H}_W^f + W^b \mathbf{h}_{w_j}^+ + b^c), \\ \hat{S}_j &= V(W^a \mathbf{H}_W^f + W^b \mathbf{h}_{w_j}^+ + b^c) + b^v, \\ sig &= f(W^f [\mathbf{H}_W^f, \mathbf{h}_c] + b^f), \\ P^f &= sig \cdot \left(\sum_{j=0}^M \hat{S}_j p_j^+ \right) + (1 - sig) \cdot P^{S_2},\end{aligned}\quad (7)$$

where $W^a, W^b \in \mathbb{R}^{d \times d}$, $b^c \in \mathbb{R}^d$, $V \in \mathbb{R}^{d \times 1}$, $W^f \in \mathbb{R}^{2d \times d}$, $b^v, b^f \in \mathbb{R}^1$ are learnable parameters and $f(\cdot)$ indicates the sigmoid activation function. This way, we can obtain the final result via taking the maximum one of the confidence score $P^f \in \mathbb{R}^{1 \times L}$.

3.6 Training Strategies

To make the proposition generator perform proposition decomposition and generation effectively, we train it on a large-scale corpus solely and then train the whole NDCR framework on the specific training data. The two training phases are as follows:

Phase 1. We first pretrain the proposition generator on the released large-scale complex text simplification data set MinWikiSplit (Niklaus et al., 2019), which is composed of 203K pairs of aligned complex source and simplified target sentences. We adopt the cross entropy generation loss \mathcal{L}_g for the decoder output. Similar to SimCSE (Gao et al., 2021), we employ the contrastive learning loss \mathcal{L}_c to make the global representation of simple proposition sentence different. In addition, we use a cross-entropy multi-label classification loss \mathcal{L}_p to train the prediction head of numbers of propositions, where the label is the number of simple sentences in the pretraining corpus. The whole training loss:

$$\mathcal{L}_{phrase1} = \mathcal{L}_g + \mathcal{L}_c + \mathcal{L}_p. \quad (8)$$

Phase 2. While training NDCR, we employ the proposition sentence-image confidence score to calculate the classification loss. The loss will cover the output of System 1, System 2 and final solution,

which is defined as follows:

$$\mathcal{L}_{match} = \sum_{i=0}^{M+2} \text{cross-entropy}(p_i, q), \quad (9)$$

where $p_i \in \mathbb{R}^{1 \times L}$ and q is the golden label. To make the negation executor effective, we devise a negational feedback loss \mathcal{L}_{neg} to optimize it. We take the prediction result of modifier in System 1 as the positive distribution and make the belief distribution output by the negation executor on the image candidates be far away from positive distribution. The loss calculation method is shown in Eq. 10,

$$\mathcal{L}_{neg} = \sum_{z=0}^M \max(\theta - \text{KL}(p_z^-, p_z^+), 0.0), \quad (10)$$

where KL indicates the K-L Divergence (Kullback and Leibler, 1951). θ is a super-parameter used to expand the positive and negational interval, which is set to 0.2. Hence, the whole optimization target is $\mathcal{L}_{match} + \mathcal{L}_{neg}$.

4 Experiments

4.1 Dataset

We conduct extensive experiments on a challenging data set IMAGECODE (Krojer et al., 2022), which contains 94,020 images, and they are divided into 9,402 sets. The overall images are collected from four released data sets: MSR-VTT (Xu et al., 2016), Video-Storytelling (Li et al., 2020a), YouCook (Das et al., 2013), and Open Images V6 (Kuznetsova et al., 2020). It consists of 21,202 human-writing complex descriptions and manually labelling corresponding golden images, which are divided into 16,594, 2,302, and 2,306 for training, validating, and testing, respectively. The image sources in the overall data set include video frames and static images.

4.2 Baselines

We compare NDCR with various types of pre-trained VLMs and other designed models based on the specific condition of this task. Specifically, ViLBERT (Lu et al., 2019) is a cross encoder where language and vision interact in the transitional layer via cross attention calculation. CLIP (Radford et al., 2021) is a two-stream vision-language encoder with two independent visual and textual encoders. UNITER (Chen et al., 2020) is a single-stream encoder where visual representations and

text tokens are concatenated and interact via the same transformer. OFA (Wang et al., 2022) is a unified cross-modal and unimodal encoder and has achieved impressive performance on multiple cross modal reasoning tasks. Krojer et al. (2022) also designed a contextual module to improve the interaction across different text-image fusion representations, achieving state-of-the-art performance.

4.3 Implementation Details

The L , λ , and d equal 10, 1000, and 512, respectively. For the proposition generator, we adopt a two-layer semantic parsing module and the pre-trained parameters of BART-base version. We set the maximum number of propositions to 10 and trained the proposition generator for 15 epochs on the MinWikiSplit data set. In addition, we set the depth of transformer block to 2 in the visual-linguistic interactor and utilized the finetuned visual encoder of CLIP (ViT-B/16) to encode images. For the cross encoder, we adopt the OFA-large architecture and first finetune it for two epochs before training the overall structure of NDCR. We froze the cross encoder, proposition generator, and visual encoder to prevent overfitting while training NDCR. While training all models, we set the batch size, initial learning rate, and dropout rate to 36, $6 \times 1e^{-5}$, and 0.1, respectively. The maximum training epoch is set to 30, and we employ the Adam Optimizer (Kingma and Ba, 2014) with the initial learning rate declining linearly to train all models. We use the validation set to select the best-performing model.

4.4 Main Results

Overall Performance. We present the performance of NDCR and comparative models on the test set in Table 1. '†' indicates that the pretrained VLMs are equipped with the contextual module and temporal embedding to enhance the contextual semantic interaction across similar images. This variant shows its effectiveness on the case of video frame according to the comparative performances such as CLIP vs. CLIP†. Table 1 reports that the proposed method achieves new state-of-the-art performance on the whole test set and significantly outperforms previous strong baseline (34.1 vs. 29.9, \uparrow 4.2). NDCR improves performances both on video frames and static images, especially static images (\uparrow 4.3), which shows its generalization on different cases. We observe that all models perform poorly on the testing samples whose images are from the

Method ↓ Type →	All	Video	Static
CLIP (Radford et al., 2021)	28.4	20.0	<u>60.0</u>
CLIP† (Krojer et al., 2022)	<u>29.9</u>	<u>22.0</u>	59.8
UNITER (Chen et al., 2020)	24.8	17.4	52.8
UNITER† (Krojer et al., 2022)	25.7	19.1	50.5
ViLBERT (Lu et al., 2019)	20.9	15.0	42.7
ViLBERT† (Krojer et al., 2022)	24.5	18.0	49.3
NDCR (ours)	34.1	26.1	64.3

Table 1: Model performance (accuracy) on **original testing set**. The results of CLIP, UNITER, ViLBERT, and their variants(†) are reported by Krojer et al. (2022). The underscore and bold indicate the second highest value and best performance (same as following tables). We report results for all examples and two disjoint subsets: video frames and static images.

Method ↓ Type →	All	Video	Static
OFA (Wang et al., 2022)	29.0	22.1	54.8
OFA†	<u>30.0</u>	<u>23.6</u>	54.6
CLIP (Radford et al., 2021)	27.4	19.7	<u>56.5</u>
CLIP† (Krojer et al., 2022)	27.6	20.8	53.2
NDCR (ours)	32.8	25.7	59.2
System 2	32.4	25.3	59.3
System 2 w/o Negation	32.0	25.3	57.3
System 1	31.6	24.5	58.3
System 1 w/o Modifier	19.3	16.4	30.3

Table 2: Ablation experiments on the **testing*** set, where we manually label the testing set to conduct ablation studies. 'Negation' and 'Modifier' indicate the negation executor and modifier. We adopt the mean pooling method to aggregate the predicted results of simple proposition in System 1 and w/o Modifier.

video clips, which may be attributed to the high similarity across video frames. Hence, there is a big room to improve the whole performance on the challenging multi-modal reasoning task.

4.5 Ablation Study

Effectiveness of Modules. To study the effectiveness of different modules, we re-annotate the test sample with the help of eight related workers (original test labels are not released). The experimental results are presented in Table 2. The performances of reproduced baselines and NDCR have a slight decline, which is because the labelling process for most examples is difficult. There are specific quality differences across human-labelling results, yet it does not affect testing and comparing model perfor-

Method ↓ Nums_of_props →	1	2	3	4	5
Total Number	61	863	1239	126	16
CLIP [†]	27	264	302	41	3
OFA [†]	28	299	340	34	5
System 1	30	297	364	36	4
System 2	31	304	373	36	3
NDCR	31	304	380	37	3
Δ	3	5	40	2	-2

Table 3: The number of samples accurately inferred on different numbers of simple proposition sentences. 'Nums_of_props' indicates the number of simple propositions. Δ represents the difference in the number of samples that NDCR and OFA[†] accurately predict.

mances. For the fairness of model comparison, the random seeds of all ablation experiments are set to the same value 10. Firstly, NDCR achieves the best performance and significantly surpasses other models on two test sets. When we add System 2 based on System 1, the overall performance improves by about 1.0, suggesting the neural-symbolic reasoner's effectiveness. Comparing System 2 and System 2 w/o negation, we observe that the negation executor improves the performance of the neural-symbolic reasoner, mainly in the case of static images. In addition, comparing System 1 and System 1 w/o modifier, we observe that introducing the context reasoning information is a very useful way to enhance the reasoning state representation of decomposed simple proposition sentences. Compared to the best baseline OFA-large (470M), the total parameter size of NDCR is about 440M. NDCR has fewer parameters yet significantly outperforms it (as shown in Table 2). This suggests that the overall performance improvement of NDCR is not due to having larger parameters.

System 1 vs. System 2. We count the experimental results on the test set according to the number of simple proposition sentences into which compound proposition texts are divided. The results are shown in Table 3. The statistical results show that NDCR excels at image retrieval from complex text with medium length, especially for those containing three simple proposition sentences. It verifies the proposed method's effectiveness in handling the complex image-text reasoning problem. Compared to System 1, System 2 performs better on test samples containing 2 or 3 simple proposition sentences, which suggests that the neural-symbolic reasoner can improve the conjunction operation of predic-

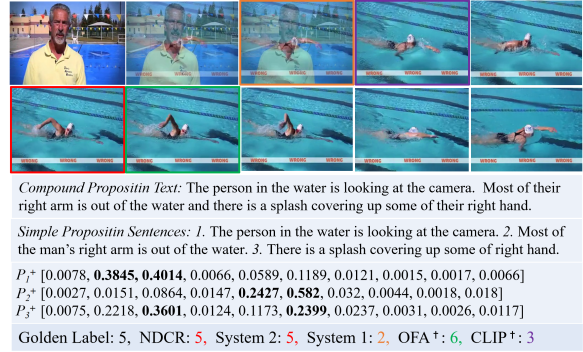


Figure 4: A case from the test set, where different colors correspond to the predicted result of models. $P_{1,2,3}^+$ represent the inferred confidence scores of simple proposition sentences in System 1 and are used to obtain the results in System 2 and final combination process.

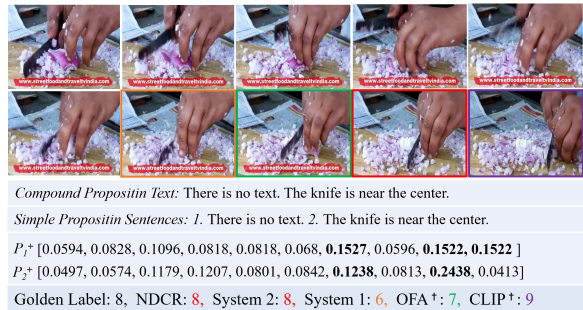


Figure 5: Another case from the test set, where it contains two simple proposition sentences.

tion results of decomposed propositions compared to rule-based methods such as mean pooling.

4.6 Case Study

We present two cases in Figure 4 and 5. For the first case (Figure 4), the proposition generator divides the complex text into three proposition sentences, and System 1 inferred the confidence scores ($P_{1,2,3}^+$) of them to ten images. Although these results of simple proposition sentences contain some errors due to having no explicit supervision signal to train, System 2 (neural-symbolic reasoner) could obtain the correct result with logical reasoning operation compared to the rule-based aggregation method in System 1. It indicates the robustness of System 2. In addition, we observe that the pretrained VLMs and System 1, which are capable of perceptual computing, often fail to cover all text semantics. It is easy for them to ignore pivotal text information (such as "there is no text" shown in Figure 5), which leads to inference errors. In conclusion, combining logical reasoning System 2 and powerful analogical reasoning System 1

(e.g., pretrained VLMs) has significant potential to take their advantages to address complex reasoning problems.

5 Conclusion

In this paper, inspired by the divide-and-conquer algorithm and dual-process theory, we introduced an end-to-end neural divide-and-conquer reasoning framework named NDCR to handle the challenging case of image retrievals from linguistically complex text. NDCR contains a proposition generator to divide the compound proposition text into multiple simple proposition sentences, then uses a visual-linguistic interactor to achieve the interaction of simple propositions and images. To improve the logical reasoning capability, we devise a neural-symbolic reasoner to gain the logical inferring result based on the output of the visual-linguistic interactor. This way, NDCR performs the low-level analogically perceptual computing in System 1 (visual-linguistic interactor) and high-level logical reasoning in System 2 (neural-symbolic reasoner). Finally, we combine the output result in Systems 1 and 2 to obtain the final solution.

Limitations

The proposed method NDCR has some limitations as follows: 1) The produced representation of simple proposition sentences in the proposition generator lies in a different space distribution with the image encoding, which affects the performance of their fused representation. Although we introduce the reasoning information of compound proposition text to alleviate this issue, we hope to solve it by improving the text understanding capability of pretrained VLMs. In addition, adopting the pretrained textual encoder of VLMs to perform proposition decomposition is inadequate due to that they present an inferior understanding for the discourse structure of long texts. 2) The performance of samples with highly similar images from video frames is quite different from that of humans. We may improve it from the perspective of image difference modelling. 3) The experimental results indicate that our method is effective at logical inference on examples with medium-length descriptions, but there is still room for improvement for longer descriptions.

Ethics Statement

IMAGECODE (Krojer et al., 2022) is an open data set used for scientific research. For ablation studies in the test set, we hired masters and undergraduate students from the research group to re-annotate the label of the test set. We have informed the creators of the data set and only conducted scientific research.

Acknowledge

We thank the anonymous reviewers for their constructive comments, and gratefully acknowledge the support of Natural Science Foundation of China (No.62006061, 61872107) and the Stable Support Program for Higher Education Institutions of Shenzhen (No.GXWD20201230155427003-20200824155011001).

References

- Thomas Anthony, Zheng Tian, and David Barber. 2017. Thinking fast and slow with deep learning and tree search. *Advances in Neural Information Processing Systems*, 30.
- Yoshua Bengio. 2017. The consciousness prior. *arXiv preprint arXiv:1709.08568*.
- Yoshua Bengio. 2019. From system 1 deep learning to system 2 deep learning. In *NeurIPS 2019*.
- Yoshua Bengio, Yann Lecun, and Geoffrey Hinton. 2021. Deep learning for ai. *Communications of the ACM*, 64(7):58–65.
- Hanxiong Chen, Shaoyun Shi, Yunqi Li, and Yongfeng Zhang. 2021. Neural collaborative reasoning. In *Proceedings of the Web Conference 2021*, pages 1516–1527.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *ECCV*.
- Pradipto Das, Chenliang Xu, Richard F. Doell, and Jason J. Corso. 2013. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2634–2641.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages

- 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. 2010. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*, 11(19):625–660.
- Jonathan St BT Evans. 2003. In two minds: dual-process accounts of reasoning. *Trends in cognitive sciences*, 7(10):454–459.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Philip M Groves and Richard F Thompson. 1970. Habituation: a dual-process theory. *Psychological review*, 77(5):419.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Benno Krojer, Vaibhav Adlakha, Vibhav Vineet, Yash Goyal, Edoardo Ponti, and Siva Reddy. 2022. Image retrieval from contextual descriptions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3426–3440, Dublin, Ireland. Association for Computational Linguistics.
- Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86.
- Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. 2020. The open images dataset v4. *International Journal of Computer Vision*, 128(7):1956–1981.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Deepak Gotmare, Shafiq Joty, Caiming Xiong, and Steven Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. In *NeurIPS*.
- Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S. Kankanhalli. 2020a. Video storytelling: Textual summaries for events. *IEEE Transactions on Multimedia*, 22(2):554–565.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Xiujun Li, Xi Yin, Chunyuan Li, Xiaowei Hu, Pengchuan Zhang, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. 2020b. Oscar: Object-semantics aligned pre-training for vision-language tasks.
- Zhixuan Liu, Zihao Wang, Yuan Lin, and Hang Li. 2022. A neural-symbolic approach to natural language understanding. *arXiv preprint arXiv:2203.10557*.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Sewon Min, Victor Zhong, Luke Zettlemoyer, and Han-naneh Hajishirzi. 2019. Multi-hop reading comprehension through question decomposition and rescoring. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6097–6109, Florence, Italy. Association for Computational Linguistics.
- Sudip Mittal, Anupam Joshi, and Tim Finin. 2017. Thinking, fast and slow: Combining vector spaces and knowledge graphs. *arXiv preprint arXiv:1708.03310*.
- Christina Niklaus, André Freitas, and Siegfried Handschuh. 2019. MinWikiSplit: A sentence splitting corpus with minimal propositions. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 118–123, Tokyo, Japan. Association for Computational Linguistics.
- Thierry Pelaccia, Jacques Tardif, Emmanuel Tribby, and Bernard Charlin. 2011. An analysis of clinical reasoning through a recent and comprehensive approach: the dual-process theory. *Medical education online*, 16(1):5890.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Shaoyun Shi, Hanxiong Chen, Weizhi Ma, Jiabin Mao, Min Zhang, and Yongfeng Zhang. 2020. Neural logic reasoning. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 1365–1374.
- Douglas R Smith. 1985. The design of divide and conquer algorithms. *Science of Computer Programming*, 5:37–58.
- Alon Talmor, Ori Yoran, Ronan Le Bras, Chandrasekhar Bhagavatula, Yoav Goldberg, Yejin Choi, and Jonathan Berant. 2021a. Commonsenseqa 2.0: Exposing the limits of ai through gamification. In *NeurIPS Datasets and Benchmarks*.

- Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang, Akari Asai, Gabriel Ilharco, Hananeh Hajishirzi, and Jonathan Berant. 2021b. [Multimodal{qa}: complex question answering over text, tables and images](#). In *International Conference on Learning Representations*.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, Hong Kong, China. Association for Computational Linguistics.
- Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. 2022. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5238–5248.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. [OFA: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 23318–23340. PMLR.
- Tomer Wolfson, Daniel Deutch, and Jonathan Berant. 2022. Weakly supervised text-to-SQL parsing through question decomposition. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2528–2542, Seattle, United States. Association for Computational Linguistics.
- Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5288–5296.
- Haoyu Zhang, Jingjing Cai, Jianjun Xu, and Ji Wang. 2019. Complex question decomposition for semantic parsing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4477–4486, Florence, Italy. Association for Computational Linguistics.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. Vinvl: Making visual representations matter in vision-language models.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
limitations
- A2. Did you discuss any potential risks of your work?
section 4.5 and limitations
- A3. Do the abstract and introduction summarize the paper's main claims?
Abstract and 1 introduction
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

comparative models in the experiment section.

- B1. Did you cite the creators of artifacts you used?
experiment
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
All dataset and comparative models are open and can be used in research purpose.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
experiments
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
we do not create any data set and conduct experiments on released public data set.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
we present comparative models and datasets in the experiments section.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
section 4.1 data sets

C Did you run computational experiments?

experiment

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
all models are run on the same device. we report the used version of pretrained models in the paper.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
implementation details
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
result analysis: table 1, 2, and 3.
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
our experiments do not use preprocessing tools and run all models on the python environment.
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
ethics statement or ablation study
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Not applicable. Left blank.
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
ethics statement
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
ethics statement
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Not applicable. we do not collect the new data and only re-annotate the label of test to conduct the ablation studies.
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
masters from the research group; in ethics statement