

LRG at SemEval-2021 Task 4: Improving Reading Comprehension with Abstract Words using Augmentation, Linguistic Features and Voting

Abheesht Sharma*

Dept. of CS&IS
BITS Pilani, Goa Campus

f20171014@goa.bits-pilani.ac.in

Harshit Pandey*

Dept. of Computer Science
Pune University

hp2pandey1@gmail.com

Gunjan Chhablani*

Dept. of CS&IS
BITS Pilani, Goa Campus

chhablani.gunjan@gmail.com

Yash Bhartia

Dept. of CS&IS
BITS Pilani, Goa Campus

f20190151@goa.bits-pilani.ac.in

Tirtharaj Dash

APP Centre for AI Research
BITS Pilani, Goa Campus

tirtharaj@goa.bits-pilani.ac.in

Abstract

In this article, we present our methodologies for SemEval-2021 Task-4: Reading Comprehension of Abstract Meaning. Given a fill-in-the-blank-type question and a corresponding context, the task is to predict the most suitable word from a list of 5 options. There are three sub-tasks within this task: Imperceptibility (subtask-I), Non-Specificity (subtask-II), and Intersection (subtask-III). We use encoders of transformers-based models pre-trained on the masked language modelling (MLM) task to build our Fill-in-the-blank (FitB) models. Moreover, to model imperceptibility, we define certain linguistic features, and to model non-specificity, we leverage information from hypernyms and hyponyms provided by a lexical database. Specifically, for non-specificity, we try out augmentation techniques, and other statistical techniques. We also propose variants, namely *Chunk Voting* and *Max Context*, to take care of input length restrictions for BERT, etc. Additionally, we perform a thorough ablation study, and use Integrated Gradients to explain our predictions on a few samples. Our best submissions achieve accuracies of 75.31% and 77.84%, on the test sets for subtask-I and subtask-II, respectively. For subtask-III, we achieve accuracies of 65.64% and 62.27%. The code is available [here](#).

1 Introduction

A very common assessment in schools is question-answering based on a given “comprehension passage”. Students are given a comprehension passage, from which they are supposed to glean necessary information, and answer short questions (such as fill-in-the-blanks-type question) based on what they have garnered from the given passage. While trying to find the most appropriate word for the blank, the children look at the words surrounding the blank

(“context”). The word should be such that when the word fills the blank, the sentence makes sense and it is grammatically correct. Inspired by this, and perhaps, after the enormous success of Transformers (Vaswani et al., 2017), researchers at Google came up with a large number of “pretraining tasks” and built knowledge-heavy language models which could be fine-tuned on various natural language processing (NLP) downstream tasks. One of the earlier pretraining tasks was “Masked Language Modelling (MLM)”, one of the two pretraining tasks of the breakthrough model, BERT (Devlin et al., 2019). The approach here was similar to how kids are taught language at school: some tokens in the text were randomly “masked” and the model was trained to predict these masked tokens.

SemEval-2021 Task-4 (Zheng et al., 2021) focuses on a similar idea. Every sample has an article, and a corresponding question. The question has a blank which the model is supposed to predict from a set of 5 options. The novelty in the task lies in its 3 subtasks: Imperceptibility (subtask-I), Non-Specificity (subtask-II), and Intersection (subtask-III). A description of these subtasks is given in Section 3. In this work, we propose using BERT and its derivative models such as DistilBERT (Sanh et al., 2019), ALBERT (Lan et al., 2019) and RoBERTa (Liu et al., 2019). Further, we propose 2 BERT variants: (1) *BERT Voting*; (2) *BERT Max. Context*. Most importantly, we also model the concepts of imperceptibility and non-specificity. For imperceptibility, we create statistical embeddings using features that have a high correlation with concreteness. For non-specificity, we propose two approaches: (1) we augment the dataset by replacing some nouns in the article by their hypernyms; and (2) we use the options’ hyponyms to decide the most appropriate option. We also experiment with GA-Reader (Dhingra et al., 2017b) and GSAMN-based approaches (Lai et al., 2019) by trying out

* Equal contribution. Author ordering determined by coin flip.

their various combinations with BERT.

In Section 2, we perform a succinct literature survey. Section 3 elucidates our approach, including the modelling aspect, the various variants of the base model, and the different ways we model imperceptibility and non-specificity. In Section 4, we perform an extensive ablation and comparative study.

2 Background

The advent of large-scale question answering systems began with straightforward tasks, like the one introduced by the SimpleQuestions Dataset (Bordes et al., 2015), which consisted of knowledge-base fact triples which were later used to answer questions. However, this dataset would only judge a model based on the ability to relate the facts to the question at hand. The purpose of NLP research is to be able to create a generalised model that may answer questions based on any context, thus datasets like the CNN Daily Mail (Hermann et al., 2015) and SQuAD (Rajpurkar et al., 2016) were created. In a typical question-answering dataset, an original and anonymised context is provided before each question. Before transformers, methods consisting of LSTM/GRUs were used to achieve good results on the aforementioned tasks. These datasets however, always had answers in the passage.

The CLOTH (Xie et al., 2018) dataset focuses on passages from middle-school and high-school text, with multiple fill-in-the-blanks in the passage. The ReCAM (Zheng et al., 2021) dataset puts a twist to archetypal fill-in-the-blank datasets by providing answer choices that are abstract in some form and which are not available in the passage itself. The models created for the QA task have to take into account semantic relations between the options and the context. GA-Reader (Dhingra et al., 2017b), is one such model, which utilises a multi-hop architecture with a novel attention mechanism, that serves as a baseline to this task.

3 Methodology

3.1 MLM-Based Transformers for Cloze-Style QA

The first model we employ follows a cloze-style question answering approach, in which we use various pretrained transformer models as encoders, followed by a decoder layer, which helps us to select the correct answer.

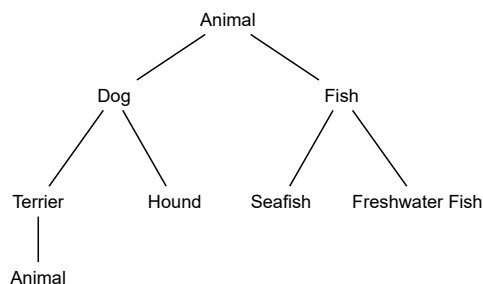


Figure 1: An example of a Hypernymy Tree

Specifically, we leverage BERT along with some of its popular and successful variants such as: DistilBERT, ALBERT, and RoBERTa. In the MLM task, tokens in the text are randomly masked, and the model is trained in a self-supervised way to predict these masked tokens. Conceptually, these transformers-based models are expected to take care of bidirectional context while predicting the masked token.

In our method, firstly, the transformer model learn the contextual embeddings of the article and the question. For the next block, the embedding of the masked token (i.e., the blank) is passed through a fully-connected layer, of which, the number of outputs corresponds to the size of the vocabulary space for the pretrained model. Each candidate option is first tokenised using WordPiece tokeniser (Wu et al., 2016), and mapped to the vector in the output vocabulary space. If the candidate option generates multiple tokens, we average the mapped scores. The model chooses the option with the highest logit value. An overview of the model is given in Figure 2.

3.2 Improvement Approaches

3.2.1 Imperceptibility:

Nouns can be clearly demarcated into two broad categories: Concrete Nouns, and Abstract Nouns. Concrete Nouns are words that represent tangible concepts, i.e., any noun referring to a name, place, object, material, etc. is considered a concrete word. Concrete words refer to concepts that can be felt by 5 human senses: Sight, Sound, Smell, Taste, and Touch. In contrast, any noun alluding to an abstract concept that cannot be experienced by our senses is an abstract word (Spreeen and Schulz, 1966). In subtask-I, the model has to predict the most accurate and the most imperceptible word from the given options. To model the imperceptibility of




Figure 2: Architecture of Transformer-based FitB Model

every word, we incorporate certain linguistic features which are highly correlated with the notion of “imperceptibility”. These linguistic features are listed below:

Length and Frequency of the Word In existing literature, authors have claimed that there exists strong evidence that concrete words are, in general, shorter than abstract words (Tanaka et al., 2013). A reasonable justification provided is that more frequently used words tend to be short (Feng et al., 2011) and since humans have a penchant for describing objects, places, or things near them, these frequently used words are generally concrete nouns. It is rather intuitive that humans would prefer ease in the pronunciation of oft-used words. Moreover, many abstract words in the English language are formed by adding suffixes to the root word, such as “coarse” becomes “coarseness”, “forget” becomes “forgetfulness” and so on (Tanaka et al., 2013).

Number of Senses of the Word In Linguistics, polysemy refers to the capacity for a word to have multiple meanings or senses. Abstract nouns are observed to be more “polysemous” than concrete nouns (Tanaka et al., 2013). For example, in WordNet (Fellbaum, 1998), the word “dog” has 8 senses, while the word “love” has 10 senses.

Number of Hyponyms Tanaka et al. 2013 find a direct correlation between the abstractness of a noun and the number of hyponyms the word has. We consider the number of hyponyms of the most commonly occurring sense of the word, and the average number of hyponyms of all the senses of the word.

Score-based Features Abstract nouns evoke emotions in humans. SentiWordNet (Baccianella et al., 2010), another lexical database like WordNet, gives scores based on the how positive, negative or objective they are. Abstract words have a higher positive/negative score, while concrete words have a higher objective score. Again, here, we consider these scores for the most commonly occurring sense, and the average scores of all the senses of the word.

Depth in Hypernymy Tree This feature is more suited for non-specificity. However, we include this as a feature of imperceptibility since the concepts of imperceptibility and non-specificity are related. For example, consider the words “money” and “property”. The latter is more imperceptible and non-specific than the former. Moreover, this is particularly useful for Subtask-III. Therefore, the depth of a word in the hypernymy tree is directly proportional to the concreteness of the word.

From the features above, we have a 13-dimensional vector for every word in the lexicon. The embedding is created so that every dimension is directly proportional to the concreteness of the word. For example, the length of a word is in general, indirectly proportional to the concreteness of the word, so we take the length dimension of the vector as *large value – length of word*, where we take 10,000 as the *large value*. The large value chosen was the same for all features which are indirectly proportional to concreteness.

Towards improving the trained model, we use a method which we term as the *Difference Method*. If the difference of the top-2 probabilities predicted by the model is greater than a certain threshold, this implies that the model is sure of the prediction it has made. However, if the difference is less than the tunable threshold, the model is ambivalent about whether the option with the highest probability or the option with the second highest probability is correct. In this case, we compute for how many dimensions the value of the linguistic embedding of the second word is less than the value of the linguistic embedding of the first word. If the majority of the values (i.e., 7) are less, we change the prediction of the model to the second-most probable option. The threshold is tuned on the dev set. Furthermore, we use a *Threshold Method* towards improving the model performance. If the highest probability is less than a tunable threshold, the model is unsure of its predictions and we consider the improvement approaches on the option with the second-highest probability.

3.2.2 Non-Specificity

According to [Spreen and Schulz, 1966](#), a highly specific word refers to a very particular instance, while a non-specific word refers to a generic concept, i.e., it encompasses many classes/instances. For example, consider the words “animal”, “bird” and “eagle”. The words are listed in increasing order of specificity.

We find parallels between the definition of specificity/non-specificity and the linguistic phenomenon of hypernymy. [Schreuder and Baayen, 1995](#) define a hypernym as “a word with a general meaning that has basically the same meaning of a more specific word”. The more specific word is the corresponding hyponym. In simpler terms, each word is related to some super-types and sub-types, called as hypernyms and hyponyms, respectively. In linguistics, hyponymy is a semantic relation be-

tween a hyponym denoting a subtype and a hypernym denoting a supertype.

For example, in figure 1, as we traverse up the hypernymy tree, assuming we consider the word “dog”, we find that its hypernym is “animal”, which is much broader than “dog”. On the other hand, as we go down the hypernymy tree, we find more specific terms for the word “dog” such as “terrier”. Essentially, hyponyms represent “IS-A” relationships. For example, “terrier” is a “dog”. We leverage the hypernymy property of words to help the model in deciding the most non-specific option. The two methods which we implement are:

Hypernym Augmentation Method In order to infuse a sense of non-specificity (other than training on the given dataset for non-specificity), we augment the dataset for subtask-I. We randomly select n nouns from the article by using a basic POS Tagging pipeline. For each noun, we use the Lesk algorithm ([Lesk, 1986](#)) to find the most appropriate sense of the word based on its context. For this sense of the word, we find its hypernyms, pick a hypernym uniformly at random from this list of hypernyms and replace the noun in the article with the hypernym. We do this for all 2^n combinations, i.e., corresponding to every sample, we have 2^n augmented samples. Furthermore, we randomly mask tokens in this dataset and train BERT on the MLM task, on this dataset. This serves a dual purpose. Firstly, it serves as a sort of domain adaptation, and secondly, it infuses a sense of non-specificity in the model.

While finetuning BERT MLM on the augmented dataset, we freeze two layers, due to time and computational constraints. We replace the normal BERT Encoder in our BERT FitB model with the BERT Encoder fine-tuned on the augmented dataset.

Hyponyms Options Method Here, we use the *Difference Method/Threshold Method*. If the model is sure of its prediction, we keep the prediction of the model. Otherwise, we generate hyponyms for each option using WordNet. After the hyponyms are tokenised, we use the trained model’s output and map each hyponym token to the output vocabulary space and get the corresponding scores. We then take the maximum score amongst all of the hyponyms as the predicted probability for that option. The reason for incorporating this approach pertains to how the transformer models were pre-

trained. Consider the following sentence: “He had a [MASK] and it was bitter”. Now, suppose that we have two options: “beer” and “drink”. Generally, our transformer-based model would look at the word “bitter” and predict “beer”. However, “drink” is more non-specific than “beer”.

3.2.3 BERT Fill-in-the-blanks Variants

To address the limitations of the vanilla transformer-based models, we attempt multiple modifications to the proposed baseline transformer models, specifically for BERT. The major limitation of the pretrained BERT model that we’ve used, is the restriction on the length of the tokenised inputs. Only 512 tokens from a sample can be processed by BERT in one parse and hence, some articles end up getting truncated and context is lost. The following are some of the modifications we’ve made to improve the performance of our models:

Voting We tokenise the question and the article. We split the article into chunks and pair each chunk with the question such that the length of the tokenised $(chunk, question)$ pair is 512. While splitting the article into chunks, we keep a max-overlap stride of 128 so that the context of the previous chunk is not lost. We train the model on these newly formed $(chunk, question)$ pairs. During inference, we take the weighted sum of the logits. For *BERT FitB Voting (Similarity)*, the weights are calculated as:

$$weight_{ij} = \frac{u_i \cdot v_j}{\|u_i\| \|v_j\|} \quad (1)$$

where u_i is the embedding of the question in the i^{th} sample, and v_j is the embedding of the j^{th} chunk of the sample’s article. To find the embeddings, we extract the *[CLS]* embedding from a pretrained BERT encoder.

We also try out an alternate way of defining the weights:

$$weight_{ij} = \frac{|\{q_i \text{ toks.}\} \cap \{chunk_j \text{ toks.}\}|}{|\{chunk_j \text{ toks.}\}|} \quad (2)$$

where $\{q_i \text{ toks.}\}$ is the set of tokens in the i^{th} sample’s question, and $\{chunk_j \text{ toks.}\}$ is the set of tokens in the j^{th} chunk of the sample. $|\cdot|$ represents the cardinality of a set. We call the method *BERT FitB Voting (Exact Matching)*.

We normalise the computed weights:

$$norm_weight_{ij} = \frac{weight_{ij}}{\sum_{j=1}^{n_i} weight_{ij}} \quad (3)$$

where n_i is the number of chunks in the i^{th} sample.

The idea behind this is that higher the similarity between the question and the article’s chunk, higher is the weight assigned to the logits returned by the trained model with the question-chunk pair as input. In Equation 2, we find the fraction of tokens common between the question and chunk.

Max Context This method is a slight modification of the Voting Method. Instead of training the model on all $(chunk, question)$ pairs for a particular sample, we train the model on the pair with the highest weight. The weights are calculated as described in Equation 2.

3.2.4 GA-Reader-based Approaches

We propose a few modifications to the baseline, namely GA-Reader (Dhingra et al., 2017a) provided by the organisers.

GA-Reader BERT We use GA-Reader on top of BERT embeddings. This could lead to potential improvement in performance for subtask-I as BERT embeddings are more feature-rich than GloVe embeddings.

GA-BERT Based on the Gated-Attention Reader, we came up with an approach that uses Gated-Attention across two-BERT streams. The first stream takes in the question input, and works like the regular BERT model. The second stream takes the article input. Assume the layer outputs for layer L are Q_L and A_L , respectively, for question and article streams. Then, to the layer $L + 1$ for question stream, Q_L is passed as input, while to layer $L + 1$ for article stream, $GA(Q_L, A_L)$ is passed, where GA is the Gated-Attention function. This is done for all 12 layers of BERT-BASE. Finally, on this model, two types of heads are attached - Selection and Pooling (similar to BERT FitB), and Attention Classification (similar to GA-Reader). The logits for each head are concatenated and a fully-connected layer is added on top. Since this is a major change in the architecture of BERT, this model needs a significant amount of pretraining.

Answer-Attention Since GA-Reader also attends to the candidate answer embeddings, we also attempt an approach where we pass the options to the BERT model. On the option embeddings and the *[MASK]* token embeddings, we apply multiplicative attention (dot product) to get attention

Model Information		Imperceptibility		Non-Specificity	
Model	Variant	Val Acc.	Test Acc.	Val Acc.	Test Acc.
BERT Fill-in-the-Blank	base	67.03%	66.77%	64.39%	65.74%
BERT Fill-in-the-Blank	large	74.79%	75.30%	72.73%	75.16%
DistilBERT Fill-in-the-Blank	base	67.03%	66.02%	63.69%	62.67%
RoBERTa Fill-in-the-Blank	base	52.45%	51.11%	33.73%	35.99%
RoBERTa Fill-in-the-Blank	large	51.02%	52.44%	33.14%	34.95%
ALBERT Fill-in-the-Blank	base-v2	31.42%	30.46%	31.84%	31.14%
ALBERT Fill-in-the-Blank	large-v2	31.06%	30.76%	30.08%	33.27%
GA-Reader (baseline)	-	21.23%	21.51%	21.50%	21.86%

Table 1: Results of the Vanilla Fill-in-the-Blank(FitB) Models and GA-Reader

scores. These scores are directly used as logits for the prediction.

3.2.5 GSAMN-based Approaches

BERT-GSAMN-Cloze Lai et al. (2019) propose a combination of Gated-Attention and Self-Attention - Gated Self-Attention (GSA). They show improvements on smaller datasets compared to Compare-Aggregate Approaches. We use two GSA layers on top of BERT Embeddings, and use the same decoder and selection method as BERT FitB.

4 Experimental Setup

In all our experiments, we use the PyTorch implementations of the transformers-based models provided by the HuggingFace (Wolf et al., 2019). The metric for all the 3 subtasks is accuracy. For subtask-I, to obtain the linguistic features mentioned in 3.2, and to obtain the hypernyms and hyponyms for subtask-II, we use the lexical database, WordNet provided by NLTK (Bird and Loper, 2004), a library in Python. For both subtasks, we train our models on train + trial dataset, and evaluate them on the dev set.

The training and the evaluation of systems was on Google Colaboratory’s free GPU (NVIDIA K80/P100). The training time varies with the models. It is around 1-2 hours for the base variants and 2-4 hours for the large models, which is well within the 12 hour limit of Colab. DistilBERT took about half an hour for training.

For finetuning the BERT FitB Hypr Aug Model on the augmented dataset on the MLM task, we use Nvidia-DGX Station with the following specifications: four 32 GB Tesla V100 GPUs, 256 GB RAM and forty Intel Xeon 2.20GHz processors since it is a computationally intensive task.

4.1 Hyperparameters

For all our experiments, we use Adam Optimiser (Kingma and Ba, 2017) and Cross Entropy Loss. For choosing the optimal set of hyperparameters, we run a Grid Search on our models. We zero in on a learning rate of 1e-5. Schedulers such as Linear Scheduler, Cosine Annealing Scheduler, etc. seem to have a negative impact on the results. For the FitB models, we keep all the layers unfrozen. Additionally, the maximum input length is kept as 512. We train our models for 4 epochs, keeping a batch size of 2.

4.2 Ablation Study/Results

Among the vanilla models, BERT FitB Large performs the best. This is understandable when it comes to DistilBERT and ALBERT, since these models are pruned and distilled for faster computation. Notably, DistilBERT gives comparable performance to BERT FitB Base. A slightly surprising observation was that there is a degradation in accuracy on using RoBERTa. This could be because even though it was pretrained more robustly than BERT on the MLM task, it was not pretrained on the Next Sentence Prediction Task, and hence, might perform worse on Textual Entailment tasks. A peculiar observation is that the large variants of ALBERT FitB and RoBERTa FitB models perform worse than their base variants. This may imply that more training data is needed to train the large variants. For subtask-I, in table 2, we also demonstrate the results of BERT Ensemble, in which we ensemble (i.e., averaging over the predictions) two checkpoints saved during the training process. When it comes to the *Difference Method* using Linguistic Features for imperceptibility, we observe an improvement on the dev set, but a slight fall is observed while evaluating it on the test set. This might be solved by careful tuning of the threshold.

The polls are already years overdue and were scheduled for Sunday . They were postponed because of an ongoing stalemate between the government and a group of opposition senators over an electoral law . Haiti is the poorest country in the region and is still struggling to recover from a 2010 earthquake . Protesters lit piles of wood in the central neighbourhood of Bel Aire before marching to a wealthy hillside neighbourhood , where riot police guarded hotels , shops and Haiti 's elections office . Some demanded President Michel Martelly 's resignation for his " inability to organise elections in the country " . Two opposition activists who had organised the protest were arrested by police for " public unrest and inciting violence " . Mid - term senate elections in Haiti had been due in May 2012 , while the municipal poll is three years behind schedule as Haiti slowly emerges from the earthquake which left much of the country devastated in 2010 . In June , President Michel Martelly decreed that the elections be held on 26 October . The date was set after lengthy talks mediated by the president of Haiti 's Bishops ' Conference , Cardinal Chibly Langlois , intended to overcome the political deadlock between the opposition and the government . But after the National Assembly failed to pass an electoral law in time , the office of Mr Martelly announced another postponement on Sunday . No new date has been set , but the statement said that " President Michel Martelly , in his constant concern to guarantee political stability , promises to pursue consultations with the different sectors of national life in order to hold the elections as soon as possible " . Opposition politicians accuse President Martelly of wanting to rule by decree - a likely scenario if no elections are held before the lower chamber 's term runs out in January . The government argues that opposition politicians are also dragging their feet in the hope of extending their time in office without elections . Thousands of Haitians marched in the capital Port - au - Prince on Sunday in protest at a delay in the country 's [MASK] and municipal elections .

Options: Local, Annual, **Legislative**, Municipal, Devastating

Figure 3: Explanation of a Correctly Classified Sample from Subtask-I (Imperceptibility). The correct option is highlighted in green.

Model	Variant	Val Acc.	Test Acc.
BERT FitB LF	large	75.75%	75.06%
DistilBERT FitB LF	base	68.10%	65.73%
BERT FitB ENS	large	75.15%	77.28%
BERT FitB ENS LF	large	75.87%	75.26%
BERT FitB EM	large	76.58%	76.35%
BERT FitB EM LF	large	76.82%	76.10%
BERT FitB VS	large	76.58%	76.54%
BERT FitB VS LF	large	76.82%	76.20%
BERT MC	large	74.07%	73.76%

Table 2: Results and Ablation Study of the Improvement Methods on Subtask-I⁰

In the future, we aspire to learn embeddings using these Linguistic Features as input to common models such as Word2Vec (Mikolov et al., 2013).

For non-specificity, with the hypernym augmentation method, BERT FitB achieves lower accuracy. A possible reason for this could be that replacing the nouns with their hypernyms in some contexts changes the meaning of the sentence (even though we use Lesk Algorithm for WSD, not all hypernyms make sense). For example, the word "drink" is replaced with "food". For the hyponyms method, we can improve our results by recursively generating hyponyms for a particular option, instead of taking the immediate hyponyms. Again, threshold tuning may help.

In Table 3, a positive sign for the *Difference Method* or the *Threshold Method* is the improvement in the results of BERT FitB Voting (Exact Matching) when we consider the hyponyms. The accuracy jumps from 72.86% to 75.79% on the dev set and from 77.83% to 78.98% on the test set. This reinforces our claim that with more careful tuning of the threshold, we might get improvements on the test set in other methods.

Model	Variant	Val Acc.	Test Acc.
BERT FitB Hypo	large	75.09%	72.83%
BERT FitB Hypr Aug	large	62.26%	60.78%
BERT FitB Hypr Aug Hypo	large	64.51%	55.52%
BERT FitB EM	large	72.86%	77.83%
BERT FitB EM Hypo	large	75.79%	78.98%
BERT FitB VS	large	73.09%	77.59%
BERT FitB VS Hypo	large	75.56%	78.63%
BERT MC	large	71.33%	71.21%

Table 3: Results and Ablation Study of the Improvement Methods on Subtask-II⁰

BERT FitB Voting performs better than vanilla BERT FitB on both subtasks. This is intuitive since in the latter, we truncate the article to 512 tokens without any consideration of how much context is lost. Voting, on the other hand, considers all contexts and hence, gives a superior performance.

For GA-Reader-BERT, when compared with the GA-Reader baseline, the accuracy improves from 21% to 39% on subtask-I dev set. Due to computational restrictions, we couldn't pretrain GA-BERT, and only fine-tuned it for subtask-I to get an idea about its performance, which was sub-optimal (19%). The Answer-Attention system gave us a dev score of $\approx 61\%$ on subtask-I, which is much higher than the baseline.

BERT-GSAMN-Cloze achieves $\approx 31\%$ accuracy on subtask-I dev set. The reasons for this could be lack of pretraining, unlike the original paper, or different way to getting the output logits. We see improvement as we reduced number of layers to 1 ($\approx 38\%$) and to 0 ($\approx 73\%$). Hence, we discarded this approach.

⁰LF=Linguistic Features, ENS=Ensemble, FitB=Fill-in-the-Blank, EM=Exact Matching, VS=Voting (similarity), MC=Max Context, Aug=Augmentation, Hypr=Hypernym, Hypo=Hyponym

The Royal College of Physicians of Edinburgh warned that being overweight may now be considered "the norm". It claimed a tax would help fund the "spiralling" healthcare costs associated with the problem. The British Soft Drinks Association (BSDA) insisted that the case is "not compelling". It cited research which suggested a 20% tax would save just four calories per day. Liverpool University chair of clinical epidemiology, Simon Capewell, is due to speak at a conference on the issue in Edinburgh later, entitled: "Obesity: A 21st Century Epidemic". Professor Capewell will cite Mexico as one example where a 10% sugary drinks tax is believed to have contributed to a 10% reduction in the consumption of such beverages while Finland, France, Hungary, Latvia and the USA have also introduced sugar taxes. He said: "The revenues raised can then be invested back into initiatives to increase children's health in these countries, as is happening in Mexico." Scotland has an excellent track record in addressing public health issues. Notable achievements include smoke-free public places and proposals for minimum unit pricing for alcohol. We need to explore how these developments could be repeated with sugary drinks. Gavin Partington, BSDA director general, said: "The efforts by soft drinks companies including product reformulation, smaller pack sizes and increased promotion of low and no-calorie drinks have led to a 7% reduction in calories from soft drinks in the last three years." It's also worth noting that politicians in Belgium and Denmark rejected the notion of a tax in 2013 and the experience in France shows that while sales of soft drinks initially fell after a tax was introduced in 2012, they have increased since. Doctors have called for the introduction of a tax on sugary drinks and drinks to tackle what they describe as an "obesity epidemic".

Options: Food, Terms, Head, Unit, Snacks

Figure 4: Explanation of a Correctly Classified Sample from Subtask-II (Non-Specificity). The correct option is highlighted in green.

Imperceptibility		Non-Specificity	
Model	Test Acc.	Model	Test Acc.
BERT FitB	65.64%	BERT FitB	61.83%
DistilBERT FitB	52.16%	BERT FitB with Hyponyms	59.95%
DistilBERT FitB + Linguistic Features	51.61%	BERT FitB with Hypernym Augmentation	45.98%
BERT FitB + Linguistic Features	65.54%	BERT FitB Voting (Exact Matching)	62.27%
BERT FitB Ensemble + Linguistic Features	64.95%	-	-

Table 4: Submitted Results of Subtask-III: Testing the performance of a system that is trained on one subtask and evaluated on the other.

4.3 Analysis of BERT FitB using Integrated Gradients

We use the method of Integrated Gradients (Sundararajan et al., 2017). We follow Ramnath et al. (2020) to compute the word-wise attribution scores for BERT FitB for both subtasks. We compute the Integrated Gradients of the target with respect to the embedding outputs. The Riemann Right Approximation Method with $n_{steps} = 25$ is used. After obtaining the token-wise attribution scores, we obtain the word-wise attribution scores by using token-to-word offset mapping. We pick the top-10 word-wise attribution scores and normalise them. To implement IG, we use the Captum (Kokhlikyan et al., 2020) library. In favour of brevity, we present one example for each subtask.

In Fig. 3, the correct answer is "legislative". The attribution scores of words like *senate*, *senators*, *municipal* and *President* are high, as is demonstrated by the intensity of the colour. The word "legislative" is, in a sense, more imperceptible than any of the words mentioned above. The *senate* is the legislative branch of the government, and *senators* are its members; *municipal* refers to municipal corporations which are the grassroots governing bodies, etc. Moreover, other words such as *elections*, *political*, *country* also have high attribution scores. These words are related to "legislative"

which exhibits the fact that BERT FitB is not only able to learn the concept of imperceptibility, but is also able to predict a suitable word.

Similarly, in Fig. 4, the correct answer is "food". Note that "snacks" is also an option; however, food is more non-specific than "snacks" and hence, food is the correct option. Another interesting thing to note is the high attribution scores for words/phrases like *calories*, *beverages*, *sugar* and *sugary drinks*. This backs the fact that the model is able to learn the concept of non-specificity, i.e., the above mentioned words are essentially hyponyms of "food".

5 Conclusion

We tried out myriad approaches, taking care to not only focus on the architecture aspect, but also how we can quantify imperceptibility and non-specificity. Although we did not achieve favourable improvements in all approaches, we did observe gains in accuracy on the dev set. We reckon that with more careful tuning of parameters such as the threshold in the *Difference Method*, we will be able to achieve these gains on the test set.

We further interpreted the outputs of transformers-based models using Integrated Gradients, and demonstrated that transformer models are able to learn the concepts of imperceptibility and non-specificity. In the future, we intend

to solidify our proposed approaches and carry out further research in this interesting field.

Acknowledgments

We thank Rajaswa Patil¹ and Somesh Singh² for their support. We would also like to express our gratitude to our colleagues at the Language Research Group (LRG)³, who have been with us at every stepping stone.

References

- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. [SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Antoine Bordes, Nicolas Usunier, Sumit Chopra, and Jason Weston. 2015. [Large-scale Simple Question Answering with Memory Networks](#).
- 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.
- Bhuvan Dhingra, Hanxiao Liu, Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. 2017a. [Gated-attention readers for text comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1832–1846, Vancouver, Canada. Association for Computational Linguistics.
- Bhuvan Dhingra, Hanxiao Liu, Zhilin Yang, William W. Cohen, and Ruslan Salakhutdinov. 2017b. [Gated-Attention Readers for Text Comprehension](#).
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- Shi Feng, Zhiqiang Cai, and Danielle McNamara. 2011. Simulating human ratings on word concreteness.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching Machines to Read and Comprehend](#).
- Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A Method for Stochastic Optimization](#).
- Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, and Orion Reblitz-Richardson. 2020. [Captum: A unified and generic model interpretability library for pytorch](#).
- Tuan Lai, Quan Hung Tran, Trung Bui, and Daisuke Kihara. 2019. [A gated self-attention memory network for answer selection](#). 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 5953–5959, Hong Kong, China. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. [ALBERT: A lite BERT for self-supervised learning of language representations](#). *CoRR*, abs/1909.11942.
- Michael Lesk. 1986. Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone. In *Proceedings of the 5th annual international conference on Systems documentation*, pages 24–26.
- 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](#). *CoRR*, abs/1907.11692.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient Estimation of Word Representations in Vector Space](#).
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ Questions for Machine Comprehension of Text](#).
- Sahana Ramnath, Preksha Nema, Deep Sahni, and Mitesh M. Khapra. 2020. [Towards interpreting BERT for reading comprehension based QA](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3236–3242, Online. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *CoRR*, abs/1910.01108.
- Robert Schreuder and R. Harald Baayen. 1995. Modelling morphological processing. In Laurie B. Feldman, editor, *Morphological aspects of language processing*, pages 131–154. Erlbaum.

¹<https://rajaswa.github.io/>

²<https://someshsingh22.github.io/>

³<https://lrg.saidl.in/>

- Otfried Spreen and Rudolph W. Schulz. 1966. [Parameters of abstraction, meaningfulness, and pronounciability for 329 nouns](#). *Journal of Verbal Learning and Verbal Behavior*, 5(5):459 – 468.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. [Axiomatic attribution for deep networks](#). *CoRR*, abs/1703.01365.
- Shinya Tanaka, Adam Jatowt, Makoto Kato, and Katsumi Tanaka. 2013. [Estimating content concreteness for finding comprehensible documents](#). pages 475–484.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *CoRR*, abs/1706.03762.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *CoRR*, abs/1910.03771.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *CoRR*, abs/1609.08144.
- Qizhe Xie, Guokun Lai, Zihang Dai, and Eduard Hovy. 2018. [Large-scale cloze test dataset created by teachers](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2344–2356, Brussels, Belgium. Association for Computational Linguistics.
- Boyuan Zheng, Xiaoyu Yang, Yuping Ruan, Quan Liu, Zhen-Hua Ling, Si Wei, and Xiaodan Zhu. 2021. SemEval-2021 task 4: Reading Comprehension of Abstract Meaning.