

Did the Model Understand the Question?

Pramod K. Mudrakarta
University of Chicago
pramodkm@uchicago.edu

Ankur Taly
Google Brain
{ataly, mukunds, kedar}@google.com

Mukund Sundararajan
Google

Kedar Dhamdhere
Google

Abstract

We analyze state-of-the-art deep learning models for three tasks: question answering on (1) images, (2) tables, and (3) passages of text. Using the notion of *attribution* (word importance), we find that these deep networks often ignore important question terms. Leveraging such behavior, we perturb questions to craft a variety of adversarial examples. Our strongest attacks drop the accuracy of a visual question answering model from 61.1% to 19%, and that of a tabular question answering model from 33.5% to 3.3%. Additionally, we show how attributions can strengthen attacks proposed by [Jia and Liang \(2017\)](#) on paragraph comprehension models. Our results demonstrate that attributions can augment standard measures of accuracy and empower investigation of model performance. When a model is accurate but for the wrong reasons, attributions can surface erroneous logic in the model that indicates inadequacies in the test data.

1 Introduction

Recently, deep learning has been applied to a variety of question answering tasks. For instance, to answer questions about images (e.g. [Kazemi and Elqursh, 2017](#)), tabular data (e.g. [Neelakantan et al., 2017](#)), and passages of text (e.g. [Yu et al., 2018](#)). Developers, end-users, and reviewers (in academia) would all like to understand the capabilities of these models.

The standard way of measuring the goodness of a system is to evaluate its error on a test set. High accuracy is indicative of a good model only if the test set is representative of the underlying real-world task. Most tasks have large test and training sets, and it is hard to manually check that they are representative of the real world.

In this paper, we propose techniques to analyze the sensitivity of a deep learning model to question words. We do this by applying *attribution* (as discussed in section 3), and generating adversarial questions. Here is an illustrative example: recall Visual Question Answering ([Agrawal et al., 2015](#)) where the task is to answer questions about images. Consider the question “*how symmetrical are the white bricks on either side of the building?*” (corresponding image in Figure 1). The system that we study gets the answer right (“*very*”). But, we find (using an attribution approach) that the system relies on only a few of the words like “*how*” and “*bricks*”. Indeed, we can construct adversarial questions about the same image that the system gets wrong. For instance, “*how spherical are the white bricks on either side of the building?*” returns the same answer (“*very*”). A key premise of our work is that most humans have expertise in question answering. Even if they cannot manually check that a dataset is representative of the real world, they can identify important question words, and anticipate their function in question answering.

1.1 Our Contributions

We follow an analysis workflow to understand three question answering models. There are two steps. First, we apply Integrated Gradients (henceforth, IG) ([Sundararajan et al., 2017](#)) to attribute the systems’ predictions to words in the questions. We propose visualizations of attributions to make analysis easy. Second, we identify weaknesses (e.g., relying on unimportant words) in the networks’ logic as exposed by the attributions, and leverage them to craft adversarial questions.

A key contribution of this work is an *overstability* test for question answering networks. [Jia and Liang \(2017\)](#) showed that reading comprehension networks are overly stable to semantics-altering edits to the passage. In this work, we find that

such overstableness also applies to questions. Furthermore, this behavior can be seen in visual and tabular question answering networks as well. We use attributions to define a general-purpose test for measuring the extent of the overstableness (sections 4.3 and 5.3). It involves measuring how a network’s accuracy changes as words are systematically dropped from questions.

We emphasize that, in contrast to model-independent adversarial techniques such as that of Jia and Liang (2017), our method exploits the strengths and weaknesses of the model(s) at hand. This allows our attacks to have a high success rate. Additionally, using insights derived from attributions we were able to improve the attack success rate of Jia and Liang (2017) (section 6.2). Such extensive use of attributions in crafting adversarial examples is novel to the best of our knowledge.

Next, we provide an overview of our results. In each case, we evaluate a pre-trained model on new inputs. We keep the networks’ parameters intact.

Visual QA (section 4): The task is to answer questions about images. We analyze the deep network in Kazemi and Elqursh (2017). We find that the network ignores many question words, relying largely on the image to produce answers. For instance, we show that the model retains more than 50% of its original accuracy even when every word that is not “color” is deleted from all questions in the validation set. We also show that the model under-relies on important question words (e.g. nouns) and attaching content-free prefixes (e.g., “in not many words, . . .”) to questions drops the accuracy from 61.1% to 19%.

QA on tables (section 5): We analyze a system called Neural Programmer (henceforth, NP) (Neelakantan et al., 2017) that answers questions on tabular data. NP determines the answer to a question by selecting a sequence of operations to apply on the accompanying table (akin to an SQL query; details in section 5). We find that these operation selections are more influenced by content-free words (e.g., “in”, “at”, “the”, etc.) in questions than important words such as nouns or adjectives. Dropping all content-free words reduces the validation accuracy of the network from 33.5%¹ to 28.5%. Similar to Visual QA, we

¹This is the single-model accuracy that we obtained on training the Neural Programmer network. The accuracy reported in the paper is 34.1%.

show that attaching content-free phrases (e.g., “in not a lot of words”) to the question drops the network’s accuracy from 33.5% to 3.3%. We also find that NP often gets the answer right for the wrong reasons. For instance, for the question “*which nation earned the most gold medals?*”, one of the operations selected by NP is “first” (pick the first row of the table). Its answer is right only because the table happens to be arranged in order of rank. We quantify this weakness by evaluating NP on the set of perturbed tables generated by Pasupat and Liang (2016) and find that its accuracy drops from 33.5% to 23%. Finally, we show an extreme form of overstableness where the table itself induces a large bias in the network regardless of the question. For instance, we found that in tables about Olympic medal counts, NP was predisposed to selecting the “prev” operator.

Reading comprehension (Section 6): The task is to answer questions about paragraphs of text. We analyze the network by Yu et al. (2018). Again, we find that the network often ignores words that should be important. Jia and Liang (2017) proposed attacks wherein sentences are added to paragraphs that ought not to change the network’s answers, but sometimes do. Our main finding is that these attacks are more likely to succeed when an added sentence includes all the question words that the model found important (for the original paragraph). For instance, we find that attacks are 50% more likely to be successful when the added sentence includes top-attributed nouns in the question. This insight should allow the construction of more successful attacks and better training data sets.

In summary, we find that all networks ignore important parts of questions. One can fix this by either improving training data, or introducing an inductive bias. Our analysis workflow is helpful in both cases. It would also make sense to expose end-users to attribution visualizations. Knowing which words were ignored, or which operations the words were mapped to, can help the user decide whether to trust a system’s response.

2 Related Work

We are motivated by Jia and Liang (2017). As they discuss, “*the extent to which [reading comprehension systems] truly understand language remains unclear*”. The contrast between Jia and Liang

(2017) and our work is instructive. Their main contribution is to fix the evaluation of reading comprehension systems by augmenting the test set with adversarially constructed examples. (As they point out in Section 4.6 of their paper, this does not necessarily fix the model; the model may simply learn to circumvent the specific attack underlying the adversarial examples.) Their method is independent of the specification of the model at hand. They use crowdsourcing to craft passage perturbations intended to fool the network, and then query the network to test their effectiveness.

In contrast, we propose improving the analysis of question answering systems. Our method peeks into the logic of a network to identify high-attribution question terms. Often there are several important question terms (e.g., nouns, adjectives) that receive tiny attribution. We leverage this weakness and perturb questions to craft targeted attacks. While Jia and Liang (2017) focus exclusively on systems for the reading comprehension task, we analyze one system each for three different tasks. Our method also helps improve the efficacy Jia and Liang (2017)’s attacks; see table 4 for examples. Our analysis technique is specific to deep-learning-based systems, whereas theirs is not.

We could use many other methods instead of Integrated Gradients (IG) to attribute a deep network’s prediction to its input features (Baehrens et al., 2010; Simonyan et al., 2013; Shrikumar et al., 2016; Binder et al., 2016; Springenberg et al., 2014). One could also use model agnostic techniques like Ribeiro et al. (2016b). We choose IG for its ease and efficiency of implementation (requires just a few gradient-calls) and its axiomatic justification (see Sundararajan et al. (2017) for a detailed comparison with other attribution methods).

Recently, there have been a number of techniques for crafting and defending against adversarial attacks on image-based deep learning models (cf. Goodfellow et al. (2015)). They are based on oversensitivity of models, i.e., tiny, imperceptible perturbations of the image to change a model’s response. In contrast, our attacks are based on models’ over-reliance on few question words even when other words should matter.

We discuss task-specific related work in corresponding sections (sections 4 to 6).

3 Integrated Gradients (IG)

We employ an attribution technique called Integrated Gradients (IG) (Sundararajan et al., 2017) to isolate question words that a deep learning system uses to produce an answer.

Formally, suppose a function $F : \mathbb{R}^n \rightarrow [0, 1]$ represents a deep network, and an input $x = (x_1, \dots, x_n) \in \mathbb{R}^n$. An attribution of the prediction at input x relative to a baseline input x' is a vector $A_F(x, x') = (a_1, \dots, a_n) \in \mathbb{R}^n$ where a_i is the *contribution* of x_i to the prediction $F(x)$. One can think of F as the probability of a specific response. x_1, \dots, x_n are the question words; to be precise, they are going to be vector representations of these terms. The attributions a_1, \dots, a_n are the influences/blame-assignments to the variables x_1, \dots, x_n on the probability F .

Notice that attributions are defined relative to a special, uninformative input called the *baseline*. In this paper, we use an empty question as the baseline, that is, a sequence of word embeddings corresponding to padding value. Note that the context (image, table, or passage) of the baseline x' is set to be that of x ; only the question is set to empty. We now describe how IG produces attributions.

Intuitively, as we interpolate between the baseline and the input, the prediction moves along a trajectory, from uncertainty to certainty (the final probability). At each point on this trajectory, one can use the gradient of the function F with respect to the input to attribute the change in probability back to the input variables. IG simply aggregates the gradients of the probability with respect to the input along this trajectory using a path integral.

Definition 1 (Integrated Gradients) *Given an input x and baseline x' , the integrated gradient along the i^{th} dimension is defined as follows.*

$$\text{IG}_i(x, x') ::= (x_i - x'_i) \times \int_{\alpha=0}^1 \frac{\partial F(x' + \alpha \times (x - x'))}{\partial x_i} d\alpha$$

(here $\frac{\partial F(x)}{\partial x_i}$ is the gradient of F along the i^{th} dimension at x).

Sundararajan et al. (2017) discuss several properties of IG. Here, we informally mention a few desirable ones, deferring the reader to Sundararajan et al. (2017) for formal definitions.

IG satisfies the condition that the attributions sum to the difference between the probabilities at

the input and the baseline. We call a variable un-influent if all else fixed, varying it does not change the output probability. IG satisfies the property that un-influent variables do not get any attribution. Conversely, influential variables always get some attribution. Attributions for a linear combination of two functions F_1 and F_2 are a linear combination of the attributions for F_1 and F_2 . Finally, IG satisfies the condition that symmetric variables get equal attributions.

In this work, we validate the use of IG empirically via question perturbations. We observe that perturbing high-attribution terms changes the networks’ response (sections 4.4 and 5.5). Conversely, perturbing terms that receive a low attribution does not change the network’s response (sections 4.3 and 5.3). We use these observations to craft attacks against the network by perturbing instances where generic words (e.g., “a”, “the”) receive high attribution or contentful words receive low attribution.

4 Visual Question Answering

4.1 Task, model, and data


The Visual Question Answering Task (Agrawal et al., 2015; Teney et al., 2017; Kazemi and Elqursh, 2017; Ben-younes et al., 2017; Zhu et al., 2016) requires a system to answer questions about images (fig. 1). We analyze the deep network from Kazemi and Elqursh (2017). It achieves 61.1% accuracy on the validation set (the state of the art (Fukui et al., 2016) achieves 66.7%). We chose this model for its easy reproducibility.

The VQA 1.0 dataset (Agrawal et al., 2015) consists of 614,163 questions posed over 204,721 images (3 questions per image). The images were taken from COCO (Lin et al., 2014), and the questions and answers were crowdsourced.

The network in Kazemi and Elqursh (2017) treats question answering as a classification task wherein the classes are 3000 most frequent answers in the training data. The input question is tokenized, embedded and fed to a multi-layer LSTM. The states of the LSTM attend to a featurized version of the image, and ultimately produce a probability distribution over the answer classes.

4.2 Observations

We applied IG and attributed the top selected answer class to input question words. The baseline for a given input instance is the image and an



Question: **how** symmetrical **are** the **white** **bricks** on **either** side of the building

Prediction: very

Ground truth: very

Figure 1: Visual QA (Kazemi and Elqursh, 2017): Visualization of attributions (word importances) for a question that the network gets right. Red indicates high attribution, blue negative attribution, and gray near-zero attribution. The colors are determined by attributions normalized w.r.t the maximum magnitude of attributions among the question’s words.

empty question². We omit instances where the top answer class predicted by the network remains the same even when the question is emptied (i.e., the baseline input). This is because IG attributions are not informative when the input and the baseline have the same prediction.

A visualization of the attributions is shown in fig. 1. Notice that very few words have high attribution. We verified that altering the low attribution words in the question does not change the network’s answer. For instance, the following questions still return “very” as the answer: “*how spherical are the white bricks on either side of the building*”, “*how soon are the bricks fading on either side of the building*”, “*how fast are the bricks speaking on either side of the building*”.

On analyzing attributions across examples, we find that most of the highly attributed words are words such as “*there*”, “*what*”, “*how*”, “*doing*”—they are usually the less important words in questions. In section 4.3 we describe a test to measure the extent to which the network depends on such words. We also find that informative words in the question (e.g., nouns) often receive very low attribution, indicating a weakness on part of the network. In Section 4.4, we describe various attacks that exploit this weakness.

4.3 Overstability test

To determine the set of question words that the network finds most important, we isolate words that most frequently occur as top attributed words in questions. We then drop all words except these and compute the accuracy.

Figure 2 shows how the accuracy changes as the size of this isolated set is varied from 0 to 5305.

²We do not black out the image in our baseline as our objective is to study the influence of just the question words for a given image

We find that just one word is enough for the model to achieve more than 50% of its final accuracy. That word is “color”.




Figure 2: VQA network (Kazemi and Elqursh, 2017): Accuracy as a function of vocabulary size, relative to its original accuracy. Words are chosen in the descending order of how frequently they appear as top attributions. The X-axis is on logscale, except near zero where it is linear.

Note that even when empty questions are passed as input to the network, its accuracy remains at about 44.3% of its original accuracy. This shows that the model is largely reliant on the image for producing the answer.

The accuracy increases (almost) monotonically with the size of the isolated set. The top 6 words in the isolated set are “color”, “many”, “what”, “is”, “there”, and “how”. We suspect that generic words like these are used to determine the type of the answer. The network then uses the type to choose between a few answers it can give for the image.

4.4 Attacks

Attributions reveal that the network relies largely on generic words in answering questions (section 4.3). This is a weakness in the network’s logic. We now describe a few attacks against the network that exploit this weakness.

Subject ablation attack

In this attack, we replace the subject of a question with a specific noun that consistently receives low attribution across questions. We then determine, among the questions that the network originally answered correctly, what percentage result in the same answer after the ablation. We repeat this process for different nouns; specifically, “fits”, “childhood”, “copyrights”, “mornings”, “disorder”, “importance”, “topless”, “critter”, “jumper”, “tweet”, and average the result.

Prefix	Accuracy
in not a lot of words	35.5%
in not many words	32.5%
what is the answer to	31.7%
<i>Union of all three</i>	19%
Baseline prefix	
tell me	51.3%
answer this	55.7%
answer this for me	49.8%
<i>Union of baseline prefixes</i>	46.9%

Table 1: VQA network (Kazemi and Elqursh, 2017): Accuracy for prefix attacks; original accuracy is 61.1%.

We find that, among the set of questions that the network originally answered correctly, 75.6% of the questions return the same answer despite the subject replacement.

Prefix attack

In this attack, we attach content-free phrases to questions. The phrases are manually crafted using generic words that the network finds important (section 4.3). Table 1 (top half) shows the resulting accuracy for three prefixes —“*in not a lot of words*”, “*what is the answer to*”, and “*in not many words*”. All of these phrases nearly halve the model’s accuracy. The union of the three attacks drops the model’s accuracy from 61.1% to 19%.

We note that the attributions computed for the network were crucial in crafting the prefixes. For instance, we find that other prefixes like “*tell me*”, “*answer this*” and “*answer this for me*” do not drop the accuracy by much; see table 1 (bottom half). The union of these ineffective prefixes drops the accuracy from 61.1% to only 46.9%. Per attributions, words present in these prefixes are not deemed important by the network.

4.5 Related work

Agrawal et al. (2016) analyze several VQA models. Among other attacks, they test the models on question fragments of telescopically increasing length. They observe that VQA models often arrive at the same answer by looking at a small fragment of the question. Our stability analysis in section 4.3 explains, and intuitively subsumes this; indeed, several of the top attributed words appear in the prefix, while important words like “color” often occur in the middle of the question. Our analysis enables additional attacks, for instance, replacing question subject with low attri-

bution nouns. Ribeiro et al. (2016a) use a model explanation technique to illustrate overstabity for two examples. They do not quantify their analysis at scale. Kafle and Kanan (2017); Zhang et al. (2016) examine the VQA data, identify deficiencies, and propose data augmentation to reduce over-representation of certain question/answer types. Goyal et al. (2016) propose the VQA 2.0 dataset, which has pairs of similar images that have different answers on the same question. We note that our method can be used to improve these datasets by identifying inputs where models ignore several words. Huang et al. (2017) evaluate robustness of VQA models by appending *questions* with semantically similar questions. Our prefix attacks in section 4.4 are in a similar vein and perhaps a more natural and targeted approach. Finally, Fong and Vedaldi (2017) use saliency methods to produce image perturbations as adversarial examples; our attacks are on the question.

5 Question Answering over Tables

5.1 Task, model, and data

We now analyze question answering over tables based on the WikiTableQuestions benchmark dataset (Pasupat and Liang, 2015). The dataset has 22033 questions posed over 2108 tables scraped from Wikipedia. Answers are either contents of table cells or some table aggregations. Models performing QA on tables translate the question into a structured program (akin to an SQL query) which is then executed on the table to produce the answer. We analyze a model called Neural Programmer (NP) (Neelakantan et al., 2017). NP is the state of the art among models that are weakly supervised, i.e., supervised using the final answer instead of the correct structured program. It achieves 33.5% accuracy on the validation set.

NP translates the input into a structured program consisting of four operator and table column selections. An example of such a program is “reset (score), reset (score), min (score), print (name)”, where the output is the name of the person who has the lowest score.

5.2 Observations

We applied IG to attribute operator and column selection to question words. NP preprocesses inputs and whenever applicable, appends symbols *tm_token*, *cm_token* to questions that signify matches between a question and the accom-

panying table. These symbols are treated the same as question words. NP also computes priors for column selection using question-table matches. These vectors, *tm* and *cm*, are passed as additional inputs to the neural network. In the baseline for IG, we use an empty question, and zero vectors for column selection priors³.

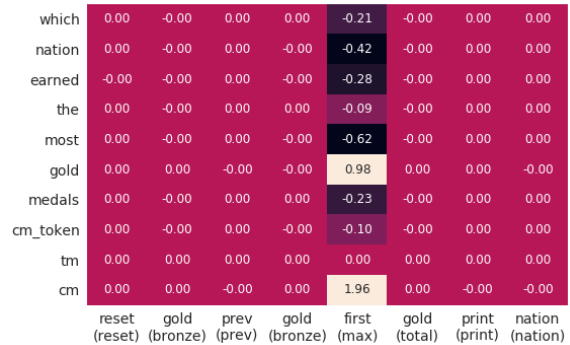


Figure 3: Visualization of attributions. Question words, preprocessing tokens and column selection priors on the Y-axis. Along the X-axis are operator and column selections with their baseline counterparts in parentheses. Operators and columns not affecting the final answer, and those which are same as their baseline counterparts, are given zero attribution.

We visualize the attributions using an alignment matrix; they are commonly used in the analysis of translation models (fig. 3). Observe that the operator “first” is used when the question is asking for a superlative. Further, we see that the word “gold” is a trigger for this operator. We investigate implications of this behavior in the following sections.

5.3 Overstability test

Similar to the test we did for Visual QA (section 4.3), we check for overstability in NP by looking at accuracy as a function of the vocabulary size. We treat table match annotations *tm_token*, *cm_token* and the out-of-vocab token (*unk*) as part of the vocabulary. The results are in fig. 4. We see that the curve is similar to that of Visual QA (fig. 2). Just 5 words (along with the column selection priors) are sufficient for the model to reach more than 50% of its final accuracy on the validation set. These five words are: “many”, “number”, “*tm_token*”, “after”, and “total”.

5.4 Table-specific default programs

We saw in the previous section that the model relies on only a few words in producing correct answers. An extreme case of overstability is when

³Note that the table is left intact in the baseline

Operator sequence	#	Triggers	Insights
reset, reset, max, print	109	[<i>unk</i> , date, position, points, name, competition, notes, no, year, venue]	sports
reset, prev, max, print	68	[<i>unk</i> , rank, total, bronze, gold, silver, nation, name, date, no]	medal tallies
reset, reset, first, print	29	[name, <i>unk</i> , notes, year, nationality, rank, location, date, comments, hometown]	player rankings
reset, mfe, first, print	25	[notes, date, title, <i>unk</i> , role, genre, year, score, opponent, event]	awards
reset, reset, min, print	17	[year, height, <i>unk</i> , name, position, floors, notes, jan, jun, may]	building info.
reset, mfe, max, print	14	[opponent, date, result, location, rank, site, attendance, notes, city, listing]	politics
reset, next, first, print	10	[<i>unk</i> , name, year, edition, birth, death, men, time, women, type]	census

Table 2: Attributions to column names for table-specific default programs (programs returned by NP on empty input questions). See supplementary material, table 6 for the full list. These results are indication that the network is predisposed towards picking certain operators solely based on the table.




Figure 4: Accuracy as a function of vocabulary size. The words are chosen in the descending order of their frequency appearance as top attributions to question terms. The X-axis is on logscale, except near zero where it is linear. Note that just 5 words are necessary for the network to reach more than 50% of its final accuracy.

the operator sequences produced by the model are independent of the question. We find that if we supply an empty question as an input, i.e., the output is a function only of the table, then the distribution over programs is quite skewed. We call these programs table-specific default programs. On average, about 36.9% of the selected operators match their table-default counterparts, indicating that the model relies significantly on the table for producing an answer.

For each default program, we used IG to attribute operator and column selections to column names and show ten most frequently occurring ones across tables in the validation set (table 2).

Here is an insight from this analysis: NP uses the combination “reset, prev” to exclude the last row of the table from answer computation. The default program corresponding to “reset, prev, max, print” has attributions to column names such as “rank”, “gold”, “silver”, “bronze”, “nation”, “year”. These column names indicate medal tallies and usually have a “total” row. If the table happens not to have a “total” row, the model may

Attack phrase	Prefix	Suffix
in not a lot of words	20.6%	10.0%
if its all the same	21.8%	18.7%
in not many words	15.6%	11.2%
one way or another	23.5%	20.0%
<i>Union of above attacks</i>	3.3%	
Baseline		
please answer	32.3%	30.7%
do you know	31.2%	29.5%
<i>Union of baseline prefixes</i>	27.1%	

Table 3: Neural Programmer (Neelakantan et al., 2017): Left: Validation accuracy when attack phrases are concatenated to the question. (Original: 33.5%)

produce an incorrect answer.

We now describe attacks that add or drop content-free words from the question, and cause NP to produce the wrong answer. These attacks leverage the attribution analysis.

5.5 Attacks

Question concatenation attacks

In these attacks, we either suffix or prefix content-free phrases to questions. The phrases are crafted using irrelevant trigger words for operator selections (supplementary material, table 5). We manually ensure that the phrases are content-free.

Table 3 describes our results. The first 4 phrases use irrelevant trigger words and result in a large drop in accuracy. For instance, the first phrase uses “not” which is a trigger for “next”, “last”, and “min”, and the second uses “same” which is a trigger for “next” and “mfe”. The four phrases combined results in the model’s accuracy going down from 33.5% to 3.3%. The first two phrases alone drop the accuracy to 5.6%.

The next set of phrases use words that receive low attribution across questions, and are hence non-triggers for any operator. The resulting drop in accuracy on using these phrases is relatively

low. Combined, they result in the model’s accuracy dropping from 33.5% to 27.1%.

Stop word deletion attacks

We find that sometimes an operator is selected based on *stop words* like: “a”, “at”, “the”, etc. For instance, in the question, “what ethnicity is at the top?”, the operator “next” is triggered on the word “at”. Dropping the word “at” from the question changes the operator selection and causes NP to return the wrong answer.

We drop stop words from questions in the validation dataset that were originally answered correctly and test NP on them. The stop words to be dropped were manually selected⁴ and are shown in Figure 5 in the supplementary material.

By dropping stop words, the accuracy drops from 33.5% to 28.5%. Selecting operators based on stop words is not robust. In real world search queries, users often phrase questions without stop words, trading grammatical correctness for conciseness. For instance, the user may simply say “top ethnicity”. It may be possible to defend against such examples by generating synthetic training data, and re-training the network on it.

Row reordering attacks

We found that NP often got the question right by leveraging artifacts of the table. For instance, the operators for the question “which nation earned the most gold medals” are “reset”, “prev”, “first” and “print”. The “prev” operator essentially excludes the last row from the answer computation. It gets the answer right for two reasons: (1) the answer is not in the last row, and (2) rows are sorted by the values in the column “gold”.

In general, a question answering system should not rely on row ordering in tables. To quantify the extent of such biases, we used a perturbed version of WikiTableQuestions validation dataset as described in Pasupat and Liang (2016)⁵ and evaluated the existing NP model on it (there was no re-training involved here). We found that NP has only 23% accuracy on it, in contrast to an accuracy of 33.5% on the original validation dataset.

One approach to making the network robust to row-reordering attacks is to train against perturbed tables. This may also help the model generalize

⁴We avoided standard stop word lists (e.g. NLTK) as they contain contentful words (e.g. “after”) which may be important in some questions (e.g. “who ranked right after turkey?”)

⁵based on data at <https://nlp.stanford.edu/software/sempr/wikitables/dpd/>

better. Indeed, Mudrakarta et al. (2018) note that the state-of-the-art strongly supervised⁶ model on WikiTableQuestions (Krishnamurthy et al., 2017) enjoys a 7% gain in its final accuracy by leveraging perturbed tables during training.

6 Reading Comprehension

6.1 Task, model, and data

The reading comprehension task involves identifying a span from a context paragraph as an answer to a question. The SQuAD dataset (Rajpurkar et al., 2016) for machine reading comprehension contains 107.7K query-answer pairs, with 87.5K for training, 10.1K for validation, and another 10.1K for testing. Deep learning methods are quite successful on this problem, with the state-of-the-art F1 score at 84.6 achieved by Yu et al. (2018); we analyze their model.

6.2 Analyzing adversarial examples

Recall the adversarial attacks proposed by Jia and Liang (2017) for reading comprehension systems. Their attack ADDSENT appends sentences to the paragraph that resemble an answer to the question without changing the ground truth. See the second column of table 4 for a few examples.

We investigate the effectiveness of their attacks using attributions. We analyze 100 examples generated by the ADDSENT method in Jia and Liang (2017), and find that an adversarial sentence is successful in fooling the model in two cases:

First, a contentful word in the question gets low/zero attribution and the adversarially added sentence modifies that word. E.g. in the question, “Who did Kubiak take the place of after Super Bowl XXIV?”, the word “Super” gets low attribution. Adding “After *Champ* Bowl XXV, Crowton took the place of Jeff Dean” changes the prediction for the model. Second, a contentful word in the question that is not present in the context. For e.g. in the question “Where hotel did the Panthers stay at?”, “hotel”, is not present in the context. Adding “The Vikings stayed at Chicago hotel.” changes the prediction for the model.

On the flip side, an adversarial sentence is unsuccessful when a contentful word in the question having high attribution is not present in the added sentence. E.g. for “Where according to gross state product does Victoria rank in Australia?”, “Australia” receives high attribution. Adding “Accord-

⁶supervised on the structured program

Question	ADDSSENT attack that does not work	Attack that works
Who was Count of Melfi	Jeff Dean was the mayor of Bracco.	Jeff Dean was the mayor of Melfi.
What country was Abhisit Vejjajiva prime minister of , despite having been born in Newcastle ?	Samak Samak was prime minister of the country of Chicago, despite having been born in Leeds.	Abhisit Vejjajiva was chief minister of the country of Chicago, despite having been born in Leeds.
Where according to gross state product does Victoria rank in Australia ?	According to net state product, Adelaide ranks 7 in New Zealand	According to net state product, Adelaide ranked 7 in Australia. (as a prefix)
When did the Methodist Protestant Church split from the Methodist Episcopal Church ?	The Presbyterian Catholics split from the Presbyterian Anglican in 1805.	The Methodist Protestant Church split from the Presbyterian Anglican in 1805. (as a prefix)
What period was 2.5 million years ago ?	The period of Plasticean era was 2.5 billion years ago.	The period of Plasticean era was 1.5 billion years ago. (as a prefix)

Table 4: ADDSENT attacks that failed to fool the model. With modifications to preserve nouns with high attributions, these are successful in fooling the model. Question words that receive high attribution are colored red (intensity indicates magnitude).

ing to net state product, Adelaide ranks 7 in New Zealand.” does not fool the model. However, retaining “Australia” in the adversarial sentence does change the model’s prediction.

6.3 Predicting the effectiveness of attacks

Next we correlate attributions with efficacy of the ADDSENT attacks. We analyzed 1000 (question, attack phrase) instances⁷ where Yu et al. (2018) model has the correct baseline prediction. Of the 1000 cases, 508 are able to fool the model, while 492 are not. We split the examples into two groups. The first group has examples where a noun or adjective in the question has high attribution, but is missing from the adversarial sentence and the rest are in the second group. Our attribution analysis suggests that we should find more failed examples in the first group. That is indeed the case. The first group has 63% failed examples, while the second has only 40%.

Recall that the attack sentences were constructed by (a) generating a sentence that answers the question, (b) replacing all the adjectives and nouns with antonyms, and named entities by the nearest word in GloVe word vector space (Pennington et al., 2014) and (c) crowdsourcing to check that the new sentence is grammatically correct. This suggests a use of attributions to improve the effectiveness of the attacks, namely ensuring that question words that the model thinks are important are left untouched in step (b) (we note that other changes in should be carried out). In table 4,

⁷data sourced from <https://worksheets.codalab.org/worksheets/0xc86d3ebe69a3427d91f9aaa63f7d1e7d/>

we show a few examples where an original attack did not fool the model, but preserving a noun with high attribution did.

7 Conclusion

We analyzed three question answering models using an attribution technique. Attributions helped us identify weaknesses of these models more effectively than conventional methods (based on validation sets). We believe that a workflow that uses attributions can aid the developer in iterating on model quality more effectively.

While the attacks in this paper may seem unrealistic, they do expose real weaknesses that affect the usage of a QA product. Under-reliance on important question terms is not safe. We also believe that other QA models may share these weaknesses. Our attribution-based methods can be directly used to gauge the extent of such problems. Additionally, our perturbation attacks (sections 4.4 and 5.5) serve as empirical validation of attributions.

Reproducibility

Code to generate attributions and reproduce our results is freely available at https://github.com/pramodkaushik/acl18_results.

Acknowledgments

We thank the anonymous reviewers and Kevin Gimpel for feedback on our work, and David Dohan for helping with the reading comprehension network. We are grateful to Jiří Šimša for helpful comments on drafts of this paper.

References

- Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. 2016. Analyzing the behavior of visual question answering models. *arXiv preprint arXiv:1606.07356*.
- Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C Lawrence Zitnick, Dhruv Batra, and Devi Parikh. 2015. Vqa: Visual question answering. *arXiv preprint arXiv:1505.00468*.
- David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus-Robert Müller. 2010. How to explain individual classification decisions. *Journal of Machine Learning Research*, pages 1803–1831.
- Hedi Ben-younes, Rémi Cadene, Matthieu Cord, and Nicolas Thome. 2017. Mutan: Multimodal tucker fusion for visual question answering. *arXiv preprint arXiv:1705.06676*.
- Alexander Binder, Grégoire Montavon, Sebastian Bach, Klaus-Robert Müller, and Wojciech Samek. 2016. Layer-wise relevance propagation for neural networks with local renormalization layers. *CoRR*.
- Ruth C Fong and Andrea Vedaldi. 2017. Interpretable explanations of black boxes by meaningful perturbation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3429–3437.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847*.
- Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *International Conference on Learning Representations*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2016. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. *arXiv preprint arXiv:1612.00837*.
- Jia-Hong Huang, Cuong Duc Dao, Modar Alfadly, and Bernard Ghanem. 2017. A novel framework for robustness analysis of visual qa models. *arXiv preprint arXiv:1711.06232*.
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*.
- Kushal Kafle and Christopher Kanan. 2017. An analysis of visual question answering algorithms. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1983–1991. IEEE.
- Vahid Kazemi and Ali Elqursh. 2017. Show, ask, attend, and answer: A strong baseline for visual question answering. *arXiv preprint arXiv:1704.03162*.
- Jayant Krishnamurthy, Pradeep Dasigi, and Matt Gardner. 2017. Neural semantic parsing with type constraints for semi-structured tables. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1516–1526.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Pramod Kaushik Mudrakarta, Ankur Taly, Mukund Sundararajan, and Kedar Dhamdhere. 2018. It was the training data pruning too! *arXiv preprint arXiv:1803.04579*.
- Arvind Neelakantan, Quoc V. Le, Martín Abadi, Andrew McCallum, and Dario Amodei. 2017. Learning a natural language interface with neural programmer.
- Arvind Neelakantan, Quoc V Le, and Ilya Sutskever. 2016. Neural programmer: Inducing latent programs with gradient descent. In *International Conference on Learning Representations ICLR*.
- Panupong Pasupat and Percy Liang. 2015. Compositional semantic parsing on semi-structured tables. In *In Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Panupong Pasupat and Percy Liang. 2016. Inferring logical forms from denotations. *arXiv preprint arXiv:1606.06900*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100, 000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016a. Nothing else matters: model-agnostic explanations by identifying prediction invariance. *arXiv preprint arXiv:1611.05817*.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016b. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144. ACM.

- Avanti Shrikumar, Peyton Greenside, Anna Shcherbina, and Anshul Kundaje. 2016. Not just a black box: Learning important features through propagating activation differences. *CoRR*.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *CoRR*.
- Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin A. Riedmiller. 2014. Striving for simplicity: The all convolutional net. *CoRR*.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. [Axiomatic attribution for deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 3319–3328.
- Damien Teney, Peter Anderson, Xiaodong He, and Anton van den Hengel. 2017. Tips and tricks for visual question answering: Learnings from the 2017 challenge. *arXiv preprint arXiv:1708.02711*.
- Adams Wei Yu, David Dohan, Quoc Le, Thang Luong, Rui Zhao, and Kai Chen. 2018. [Fast and accurate reading comprehension by combining self-attention and convolution](#). In *International Conference on Learning Representations*.
- Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2016. Yin and yang: Balancing and answering binary visual questions. In *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*, pages 5014–5022. IEEE.
- Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. 2016. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4995–5004.