

DVQA: Understanding Data Visualizations via Question Answering

Kushal Kafle^{1,*} Brian Price² Scott Cohen² Christopher Kanan¹
¹Rochester Institute of Technology ²Adobe Research
¹{kk6055, kanan}@rit.edu ²{bprice, scohen}@adobe.com

Abstract

Bar charts are an effective way to convey numeric information, but today’s algorithms cannot parse them. Existing methods fail when faced with even minor variations in appearance. Here, we present DVQA, a dataset that tests many aspects of bar chart understanding in a question answering framework. Unlike visual question answering (VQA), DVQA requires processing words and answers that are unique to a particular bar chart. State-of-the-art VQA algorithms perform poorly on DVQA, and we propose two strong baselines that perform considerably better. Our work will enable algorithms to automatically extract numeric and semantic information from vast quantities of bar charts found in scientific publications, Internet articles, business reports, and many other areas.

1. Introduction

Data visualizations, e.g., bar charts, pie charts, and plots, contain large amounts of information in a concise format. These visualizations are specifically designed to communicate data to people, and are not designed to be machine interpretable. Nevertheless, making algorithms capable to make inferences from data visualizations has enormous practical applications. Here, we study systems capable of answering open-ended questions about bar charts, which we refer to as data visualization question answering (DVQA). DVQA would enable vast repositories of charts within scientific documents, web-pages, and business reports to be queried automatically. Example DVQA images and questions grouped by the different tasks are shown in Fig. 1.

Besides practical benefits, DVQA can also serve as a challenging proxy task for generalized pattern matching, attention, and multi-step reasoning systems. Answering a question about a chart requires multi-step attention, memory, measurement, and reasoning that poses significant challenges to the existing systems. For example, to answer the question ‘What is the accuracy of algorithm vice on the

*A portion of this research was done while Kushal Kafle was an intern at Adobe Research.

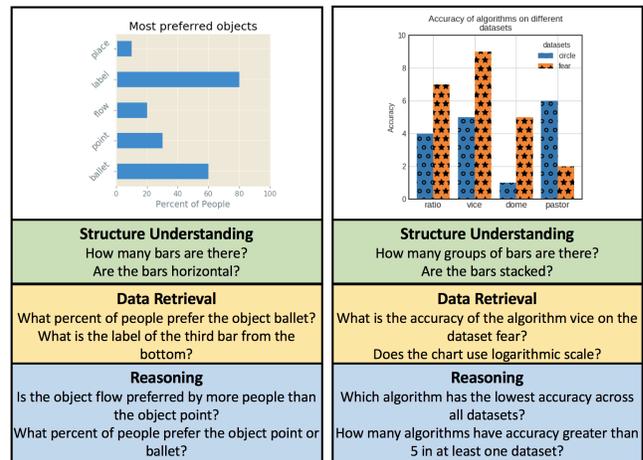


Figure 1: DVQA involves answering questions about diagrams. We present a dataset for DVQA with bar charts that exhibit enormous variety in appearance and style. We show that VQA systems cannot answer many DVQA questions and we describe more effective algorithms.

dataset fear?’ in Fig. 1 can require finding the appropriate color and hatching that represents the dataset *fear*, finding the group of bars that represent the algorithm *vice*, measuring the height of the bar based on the y-axis, and if necessary interpolating between two neighboring values.

DVQA is related to visual question answering (VQA) [27, 5], which deals with answering open-ended questions about images. VQA is usually treated as a classification problem, in which answers are categories that are inferred using features from image-question pairs. DVQA poses three major challenges that are overlooked by existing VQA datasets with natural images. First, VQA systems typically assume two fixed vocabulary dictionaries: one for encoding words in questions and one for producing answers. In DVQA, assuming a fixed vocabulary makes it impossible to properly process many questions or to generate answers unique to a bar chart, which are often labeled with proper nouns, abbreviations, or concatenations (e.g., ‘Jan-Apr’). Our models demonstrate two ways for

handling out-of-vocabulary (OOV) words. Second, the language utilized in VQA systems represent fixed semantic concepts that are immutable over images, *e.g.*, phrases such as ‘A large shiny red cube’ used in CLEVR [13] represent a fixed concept; once the word ‘red’ is associated with the underlying semantic concept, it is immutable. By contrast, the words utilized in labels and legends in DVQA can often be arbitrary and could refer to bars of different position, size, texture, and color. Third, VQA’s natural images exhibit regularities that are not present in DVQA, *e.g.* to infer the answer to ‘What is the weather like?’ for the image in Fig. 2, an agent could use color tones and overall brightness to infer ‘sunny.’ Changing the color of the fire hydrant only changes the local information that impacts questions about the fire hydrant’s properties. However, in bar charts, a small change, *e.g.*, shuffling the colors of the legend in Fig. 2, completely alters the chart’s information. This makes DVQA an especially challenging problem.

This paper makes three major contributions:

1. We describe the DVQA dataset, which contains over 3 million image-question pairs about bar charts. It tests three forms of diagram understanding: a) structure understanding; b) data retrieval; and c) reasoning. The DVQA dataset will be publicly released.
2. We show that both baseline and state-of-the-art VQA algorithms are incapable of answering many of the questions in DVQA. Moreover, existing classification based systems based on a static and predefined vocabulary are incapable of answering questions with unique answers that are not encountered during training.
3. We describe two DVQA systems capable of handling words that are unique to a particular image. One is an end-to-end neural network that can read answers from the bar chart. The second is a model that encodes a bar chart’s text using a dynamic local dictionary.

2. Related Work

2.1. Automatically Parsing Bar Charts

Extracting data from bar charts using computer vision has been extensively studied [2, 6, 19, 29, 31]. Some focus on extracting the visual elements from the bar charts [29], while others focus on extracting the data from each bar directly [31, 19]. Most of these approaches use fixed heuristics and make strong simplifying assumptions, *e.g.*, [31] made several simplifying assumptions about bar chart appearance (bars are solidly shaded without textures or gradients, no stacked bars, etc.). Moreover, they only tested their data extraction procedure on a total of 41 bar charts.

Our DVQA dataset has variations in bar chart appearance that go far beyond the capabilities of any of the aforementioned works. Moreover, DVQA requires more than just data extraction. Correctly answering DVQA questions re-

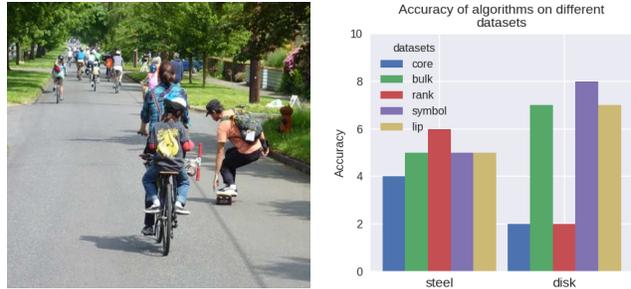


Figure 2: Natural images vs. bar charts. **Left:** Small changes in an image typically have little impact on a question in VQA. **Right:** Bar charts convey information using a sparse, but precise, set of visual elements. Even small changes can completely alter the information in the chart.

quires basic language understanding, attention, concept of working short-term memory and reasoning.

2.2. VQA with Natural Images

Over the past three years, multiple VQA datasets containing natural images have been publicly released [27, 5, 30, 25, 15]. The most popular dataset is The VQA Dataset [8, 5]. It is much larger and more varied than earlier VQA datasets, such as COCO-QA [30] and DAQUAR [27]. However, the first version of the dataset, VQA 1.0, suffered from extreme language bias, resulting in many questions not requiring the image to correctly answer them [8]. In the second version, VQA 2.0, this bias was greatly reduced; however, VQA 2.0 still suffers from heavily skewed distribution in the *kinds* of questions present in the dataset [15].

Numerous VQA algorithms have been proposed, ranging from Bayesian approaches [14, 27], methods using spatial attention [34, 33, 26, 28], compositional approaches [3, 4], and bilinear pooling schemes [23, 7]. Almost all VQA algorithms pose it as a classification problem in which each class is synonymous with a particular answer. For more extensive reviews see [16] and [32].

While there are significant similarities between VQA and DVQA, one critical difference is that many DVQA questions require directly reading text from a chart to correctly answer them. This demands being able to handle words that are unique to a particular chart, which is a capability that is not needed by algorithms operating on existing VQA datasets with natural images.

2.3. Reasoning, Synthetic Scenes, and Diagrams

While VQA is primarily studied using natural images, several datasets have been proposed that use synthetic scenes or diagrams to test reasoning and understanding [13, 21, 22]. The CLEVR [13] dataset has complex reasoning questions about synthetically created scenes, and systems that perform well on popular VQA datasets perform poorly

on CLEVR. The TQA [22] and AI2D [21] datasets both involve answering science questions about text and images. Both datasets are relatively small, *e.g.*, AI2D only contains 15,000 questions. These datasets require more than simple pattern matching and memorization. Similar to our work, their creators showed that state-of-the-art VQA systems for natural image datasets performed poorly on their datasets. However, there are key differences between these datasets and DVQA. First, none of these datasets contain questions specific to bar charts. Second, their datasets use multiple-choice schemes that reduce the problem to a ranking problem, rather than the challenges posed by having to generate open-ended answers. Studies have shown that multiple-choice schemes have biases that models will learn to exploit [12]. In contrast, we treat DVQA as an open-ended question answering task.

Concurrent to our work, FigureQA [18] also explores question answering for charts, however, with following major limitations compared to our DVQA dataset: 1) it contains only yes/no type questions; 2) it does not contain questions that require numeric values as answers; 3) it has fixed labels for bars across different figures (*e.g.*, a red bar is always labeled 'red'); and 4) it avoids the OOV problem.

3. DVQA: The Dataset

DVQA is a challenging synthetic dataset that tests multiple aspects of bar chart understanding that cause state-of-the-art methods for VQA to fail, which we demonstrate in experiments. Synthetically generating DVQA gave us precise control over the positions and appearances of the visual elements. It also gave us access to meta-data about these components, which would not be available with real data. This meta-data contains all information within the chart, including the precise position of each drawing element, the underlying data used to create the chart and location of all text-elements. This data can be used as an additional source of supervision or to ensure that an algorithm is “attending” to relevant regions. As shown in Fig. 3, the DVQA dataset contains a large variety of typically available styles of bar chart. The questions in the dataset require the ability to reason about the information within a bar chart (see Fig. 1). DVQA contains 3,487,194 total question answer pairs for 300,000 images divided into three major question types. Tables 1 and 2 show statistics about the DVQA dataset. Additional statistics are given in the supplemental materials.

3.1. Appearance, Data, and Question Types

DVQA consists of bar charts with question-answer pairs that are generated by selecting a visual style for a chart, choosing data for a chart, and then generating questions for that chart. Here, we briefly explain how this was done. Additional details are provided in the supplemental materials.

Visual Styles: We use python’s popular drawing tool, Matplotlib to generate our charts since it offers unparalleled programmatic control over each of the element drawn. As shown in Fig. 3, DVQA’s bar charts contain a wide variability in both appearance and style that can capture the common styles found in scientific documents and the Internet. Some of these variations include the difference in the number of bars and groups; presence or absence of grid lines; difference in color, width, spacing, orientation, and texture of the bars; and difference in the orientation and the location of labels and legends.

To label individual bars and legend entries, we select the 1000 most frequent nouns in the Brown Corpus using NLTK’s part-of-speech tagging for our training set and our ‘easy’ test set *Test-Familiar*. To measure a system’s ability to scale to unknown answers, we also created a more difficult test set *Test-Novel*, in which we use 500 new words that are not seen during training.

Underlying Data: DVQA has three bar chart data types: linear, percentage, and exponential. For each of these data value types, the bars can take any of the 10 randomly chosen values in the range 1 – 10 for linear data, 10 – 100 for percentage, and 1 - 10^{10} for exponential data type. A small percentage of bars are allowed to have a value of zero which appears as a missing bar in the chart.

Question Types: DVQA contains three types of questions: 1) structure understanding, 2) data retrieval, and 3) reasoning. To generate these questions, we use fixed templates. Based on the context of the chart reflected through its title and labels, the questions will vary along the template. Below, we will show a random assortment of these questions with further details presented in the supplementary materials.

Structure Understanding. Structure understanding questions test a system’s ability to understand the overall structure of a bar chart. These questions include:

1. How many bars are there?
2. How many groups/stacks of bars are there?
3. How many bars are there per group?
4. Does the chart contain any negative values?
5. Are the bars horizontal?
6. Does the chart contain stacked bars?
7. Is each bar a single solid color without patterns?

Data Retrieval. Data retrieval questions test a system’s ability to retrieve information from a bar chart by parsing the chart into its individual components. These questions often require paying attention to specific region of the chart. These questions include:

1. Are the values in the chart presented in a logarithmic scale?
2. Are the values in the chart presented in a percentage scale?
3. What percentage of people prefer the object **O**?
4. What is the label of the third bar from the left?
5. What is the label of the first group of bars from the left?

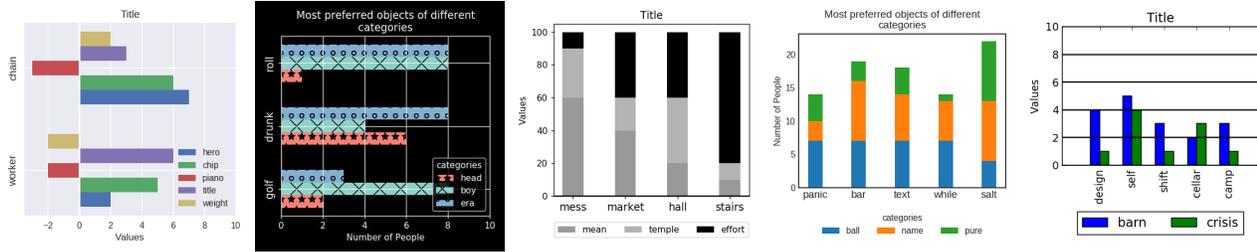


Figure 3: Example bar chart images from DVQA. DVQA contains significant variation in appearance and style.

6. What is the label of the second bar from the left in each group?
7. What element does the C color represent?
8. How many units of the item I were sold in the store S?

Reasoning. Reasoning questions test a model’s ability to collect information from multiple components of a bar chart and perform operations on them. These include:

1. Which algorithm has the highest accuracy?
2. How many items sold more than N units?
3. What is the difference between the largest and the smallest value in the chart?
4. How many algorithms have accuracies higher than N ?
5. What is the sum of the values of $L1$ and $L2$?
6. Did the item $I1$ sold less units than $I2$?
7. How many groups of bars contain at least one bar with value greater than N ?
8. Which item sold the most units in any store?
9. Which item sold the least number of units summed across all the stores?
10. Is the accuracy of the algorithm $A1$ in the dataset $D1$ larger than the accuracy of the algorithm $A2$ in the dataset $D2$?

3.2. Post-processing to Minimize Bias

Several studies in VQA have shown that bias in datasets can impair performance evaluation and give inflated scores to systems that simply exploit statistical patterns [12, 15, 1]. In DVQA, we have taken several measures to combat such biases. To ensure that there is no correlation between styles, colors, and labels, we randomize the generation of charts. Some questions can have strong priors, *e.g.*, the question ‘Does the chart contain stacked bar?’ has a high probability of the correct answer being ‘no’ because these stacked charts are uncommon. To compensate for this, we randomly remove these questions until yes/no answers are balanced for each question type where yes/no is an answer. A similar scheme was used to balance other structure understanding questions as well as the first two data retrieval questions.

4. DVQA Algorithms & Models

In this section, we describe two novel deep neural network algorithms along with five baselines. Our proposed algorithms are able to read text from bar charts, giving them

Table 1: Dataset statistics for different DVQA splits for different question types.

	Total Questions	Unique Answers
Structure	471,108	10
Data	1,113,704	1,538
Reasoning	1,613,974	1,576
Grand Total	3,487,194	1,576

Table 2: DVQA dataset statistics for different splits.

	Images	Questions	Unique Answers
Train	200,000	2,325,316	1,076
Test-Familiar	50,000	580,557	1,075
Test-Novel	50,000	581,321	577
Grand Total	300,000	3,487,194	1,576

the ability to answer questions with chart-specific answers or requiring chart-specific information.

All of the models that process images use the ImageNet pre-trained ResNet-152 [10] CNN with 448×448 images resulting in a $14 \times 14 \times 2048$ feature tensor, unless otherwise noted. All models that process questions use a 1024 unit single layer LSTM to encode questions, where each word in the question is embedded into a dense 300 dimensional representation. Training details are given in Sec. 4.4.

4.1. Baseline Models

We evaluate five baseline models for DVQA:

1. **YES:** This model answers ‘YES’ for all questions, which is the most common answer in DVQA by a small margin over ‘NO’.
2. **IMG:** A question-blind model. Images are encoded using Resnet using the output of its final convolutional layer after pooling, and then the answer is predicted from them by an MLP with one hidden-layer that has 1,024 units and a softmax output layer.
3. **QUES:** An image-blind model. It uses the LSTM encoder to embed the question, and then the answer is predicted by an MLP with one hidden-layer that has 1,024 units and a softmax output layer.

4. **IMG+QUES**: This is a combination of the QUES and IMG models. It concatenates the LSTM and CNN embeddings, and then feeds them to an MLP with one 1024-unit hidden layer and a softmax output layer.
5. **SAN-VQA**: The Stacked Attention Network (SAN) [34] for VQA. We use our own implementation of SAN as described by [20], where it was shown that upgrading the original SAN’s image features and a couple small changes produces state-of-the-art results on VQA 1.0 and 2.0. SAN operates on the last CNN convolutional feature maps, where it processes this map attentively using the question embedding from our LSTM-based scheme.

4.2. Multi-Output Model (MOM)

Our Multi-Output Model (MOM) for DVQA uses a dual-network architecture, where one of its sub-networks is able to generate chart-specific answers. MOM’s classification sub-network is responsible for *generic* answers. MOM’s optical character recognition (OCR) sub-network is responsible for *chart-specific* answers that must be read from the bar chart. The classification sub-network is identical to the SAN-VQA algorithm described earlier in Sec. 4.1. An overview is given in Fig. 4.

MOM’s OCR sub-network tries to predict the bounding box containing the correct label and then applies a character-level decoder to that region. The bounding box predictor is trained as a regression task using a mean-squared-error (MSE) loss. An image patch is extracted from this region, which is resized to 128×128 , and then a small 3-layer CNN is applied to it. Since the orientation of the text in the box will vary, we employ an N -step spatial attention mechanism to encode the relevant features for each of the N possible characters in the image patch, where N is the largest possible character-sequence ($N = 8$ in our experiments). These N features are encoded using a bi-directional gated recurrent unit (GRU) to capture the character level correlations found in naturally occurring words. The GRU encoding is followed by a classification layer that predicts the character sequence, which is trained using connectionist temporal classification (CTC) loss [9].

MOM must determine whether to use the classification sub-network (*i.e.* SAN-VQA) or the OCR sub-network to answer a question. To determine this, we train a separate binary classifier that determines which of the outputs to trust. This classifier takes the LSTM question features as input to predict whether the answer is generic or chart-specific. For our DVQA dataset this classifier is able to predict the correct branch with perfect accuracy on the test data.

4.3. SANDY: SAN with DYnamic Encoding Model

MOM handles chart-specific answers by having a sub-network capable of generating unique strings; however, it

has no explicit ability to visually read bar chart text and its LSTM question encoding cannot handle chart-specific words. To explore overcoming these limitations, we modified SAN to create SANDY, SAN with DYnamic encoding model. SANDY uses a dynamic encoding model (DEM) that explicitly encodes chart-specific words in the question, and can directly generate chart-specific answers. The DEM is a dynamic local dictionary for chart-specific words. This dictionary is used for encoding words as well as answers.

To create a local word dictionary, DEM assumes it has access to an OCR system that gives it the positions and strings for all text-areas in a bar chart. Given this collection of boxes, DEM assigns each box a unique numeric index. It assigns an index of 0 to the box in the lower-left corner of the image. Then, it assigns the box with the position closest to the first box with an index of 1. The box closest to 1 that is not yet assigned an index is then assigned the index of 2, and so on until all boxes in the image are assigned an index. In our implementation, we assume that we have a perfect (oracle) OCR system for input, and we use the dataset’s annotations for this purpose. No chart in the training data had more than 30 text labels, so we set the local dictionary to have at most $M = 30$ elements.

The local dictionary augments the N element global dictionary. This enables DEM to create $(M + N)$ -word dictionary that are used to encode each word in a question. The local dictionary is also used to augment the L element global answer dictionary. This is done by adding M extra classes to the classifier representing the dynamic words. If these classes are predicted, then the output string is assigned using the local dictionary’s appropriate index.

We test two versions of SANDY. The oracle version directly uses annotations from the DVQA dataset to build a DEM. The OCR version uses the output of the open-source Tesseract OCR. Tesseract’s output is pre-processed in three ways: 1) we only use words with alphabetical characters in them, 2) we filter word detections with confidence less than 50%, and 3) we filter single-character word detections.

4.4. Training the Models

All of the classification based systems, except SANDY and the OCR branch of MOM, use a global answer dictionary from training set containing 1076 words, so they each have 1076 output units. MOM’s OCR branch contains 27 output units; 1 for each alphabet and 1 reserved for blank character. Similarly, SANDY’s output layer contains 107 units, with the indices 31 through 107 are reserved for common answers and indices 0 through 30 are reserved for the local dictionary.

For a fair comparison, we use the same training hyper-parameters for all the models and closely follow the architecture for SAN models from [20] wherever possible. SAN portion for all the models are trained using early stopping

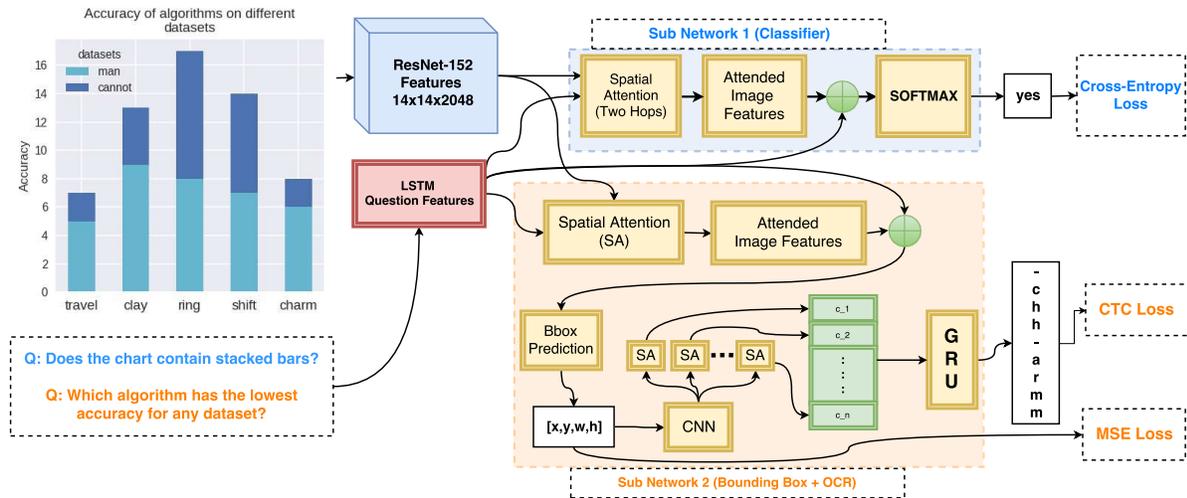


Figure 4: Overview of our Multi-Output Model (MOM) for DVQA. MOM uses two sub-networks: 1) classification sub-network that is responsible for *generic answers*, and 2) OCR sub-network that is responsible for *chart-specific answers*.

and regularized using dropout of 0.5 on inputs to all convolutional, fully-connected and LSTM units. All models use Adam [24] optimizer with an initial learning rate of 0.001.

5. Experiments

In this section, we describe the experimental results for models trained and tested on the DVQA dataset. DVQA’s extensive annotations are used to analyze the performance of each model on different question- and answer-types to reveal their respective strengths and weaknesses. In our experiments, we study the performance of algorithms on both familiar and novel chart labels, which are contained in two distinct test splits, Test-Familiar and Test-Novel. Every bar chart in Test-Familiar contains only labels seen during training. All of the models using the LSTM-encoder have entries in their word dictionaries for these familiar words, and all answers have been seen in the training set. The labels for the charts in Test-Novel are only seen in the test set, and no system has them in the dictionaries they use to encode words or to generate answers.

To measure performance, an algorithm gets a question correct only if it generates a string that is identical to the ground truth. To better assess MOM, we also measure its performance using edit distance, which is denoted MOM (± 1). This model is allowed to get a question correct as long as the answer it generates is within one edit distance or less compared to the correct answer.

5.1. General Observations

Overall performance of each method broken down based on question-type are shown in Table 3 and some qualitative examples are shown in Fig 5. Across all question-types,

NO, IMG, and QUES are the first, second, and third worst performing, respectively. Overall, SANDY performs best on both Test-Familiar and Test-Novel with SANDY-real following closely behind.

For structure questions, there is little difference across models for Test-Familiar and Test-Novel, which is expected because these questions ask about the general visual organization of a chart and do not require label reading. Performance increases greatly for IMG+QUES compared to either IMG or QUES, indicating structure questions demand combining image and question features.

For data retrieval and reasoning questions, SANDY and MOM both outperformed all baseline models. Both SANDY and SANDY-real outperformed MOM, and this gap was greater for Test-Novel.

5.2. Chart-specific Words in Questions and Answers

Many DVQA questions have chart-specific answers, *e.g.*, ‘Which algorithm has the highest accuracy?’ needs to be answered with the label of the bar with the highest value. These chart-specific answers are different than the generic answers that are shared across many bar charts, *e.g.*, ‘Does the chart contain stacked bars?’. Similarly, some DVQA questions refer to elements that are specific to a given chart, *e.g.*, ‘What is the accuracy of the algorithm A?’. To accurately answer these questions, an algorithm must be able to interpret the text-label **A** in the context of the given bar chart. Table 4 shows the accuracy of the algorithms for questions that have chart-specific labels in them (chart-specific questions) and questions whose answers contain chart-specific labels (chart-specific answers). As shown, whenever chart-specific labels appear in the answer, both IMG+QUES and SAN-VQA fail abysmally. While this is

Table 3: Overall results for models trained and tested on the DVQA dataset. Values are % of questions answered correctly.

	Test-Familiar				Test-Novel			
	Structure	Data	Reasoning	Overall	Structure	Data	Reasoning	Overall
YES	41.14	7.45	8.31	11.78	41.01	7.52	8.23	11.77
IMG	60.09	9.07	8.27	14.83	59.83	9.11	8.37	14.90
QUES	44.03	9.82	25.87	21.06	43.90	9.80	25.76	21.00
IMG+QUES	90.38	15.74	31.95	32.01	90.06	15.85	31.84	32.01
SAN-VQA	94.71	18.78	37.29	36.04	94.82	18.92	37.25	36.14
MOM	94.71	29.52	39.21	40.89	94.82	21.40	37.68	37.26
MOM (± 1)	94.71	38.20	40.99	45.03	94.82	29.14	39.26	40.90
SANDY (Oracle)	96.47	65.40	44.03	56.48	96.42	65.55	44.09	56.62
SANDY (OCR)	96.47	37.82	41.50	45.77	96.42	37.78	41.49	45.81

Table 4: Results for chart-specific questions and answers. State-of-the-art VQA algorithms struggle with both.

	Test-Familiar			Test-Novel		
	Chart-specific Questions	Chart-specific Answers	Overall	Chart-specific Questions	Chart-specific Answers	Overall
YES	17.71	0.00	11.78	17.58	0.00	11.77
IMG	17.61	0.00	14.83	17.88	0.00	14.90
QUES	23.17	0.10	21.06	22.97	0.00	21.00
IMG+QUES	25.52	0.09	32.01	25.49	0.00	32.01
SAN-VQA	26.54	0.10	36.04	26.32	0.00	36.14
MOM	26.54	12.78	40.89	26.32	2.93	37.26
MOM (± 1)	26.54	23.62	45.03	26.32	12.47	40.90
SANDY (Oracle)	27.80	52.55	56.48	27.77	52.70	56.62
SANDY (OCR)	26.60	25.19	45.77	26.43	25.12	45.81

expected for Test-Novel, they perform no better on Test-Familiar. Likewise, all of the models except SANDY also face difficulty for questions with chart-specific labels. Overall, they fail to meaningfully outperform the QUES baseline. We believe that the small gain in accuracy by IMG+QUES and SAN-VQA over QUES is only because the image information, such as the type of scale used (linear, percentage, or logarithmic), enables these methods to guess answers with higher precision.

In chart-specific answers, SANDY showed highest accuracy. Moreover, its performance for Test-Novel is similar to that for Test-Familiar. In comparison, while MOM outperforms the baselines, its accuracy on Test-Novel is much lower than its accuracy on Test-Familiar. This could be because MOM’s string generation system is unable to produce accurate results with novel words. Supporting this, MOM often makes small string generation errors, as shown by the improved performance of MOM (± 1), which is evaluated using edit distance. MOM’s output is also dependent on the precise prediction of the bounding box containing the answer which could further affect the final accuracy. MOM’s localization performance is explored in more detail in the supplemental materials.

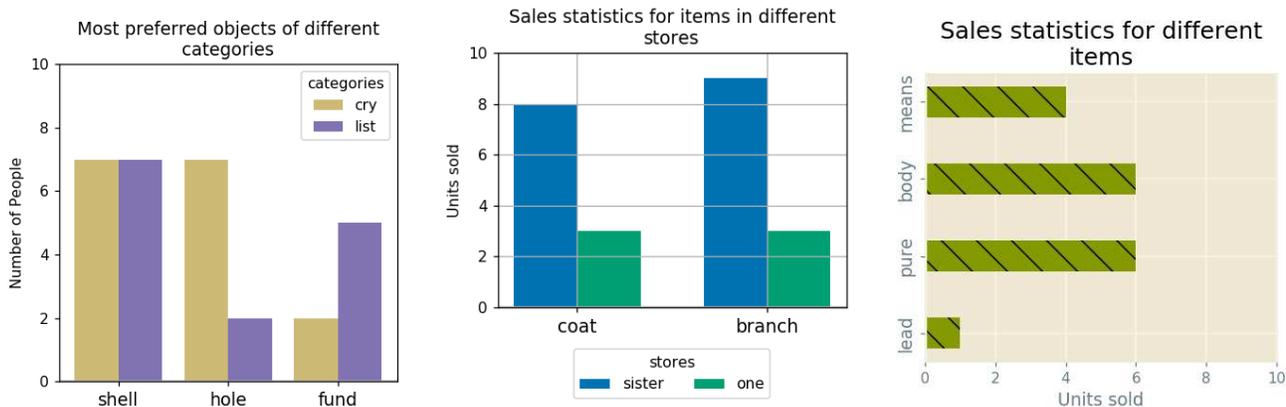
In addition to SANDY’s ability to predict chart-specific answer tokens, it can also be used to properly tokenize the chart-specific words in questions. An LSTM based question

encoder using a fixed vocabulary will not be able to encode the questions properly, especially when encoding questions with unknown words in Test-Novel. For questions with chart-specific labels on them, SANDY shows improvement in properly encoding the questions with the chart-specific labels compared to baselines. However, the improvement in performance is not as drastic as seen for chart-specific answers. This may be due to the fact that many of the chart-specific questions include precise measurement *e.g.* ‘How many people prefer object **O**?’ which could be beyond the capacity of the SAN architecture.

6. Discussion

In this paper we presented the DVQA dataset and explored models for DVQA. Our experiments show that VQA algorithms are only capable of answering simple structure questions. They perform much more poorly on data retrieval and reasoning questions, whereas our approaches, SANDY and MOM, are able to better answer these questions. Moreover, SANDY and MOM can both produce answers that are novel to the test set, which is impossible for traditional VQA algorithms. Finally, SANDY can also encode questions with novel words from the bar chart.

We studied SANDY’s performance using a real OCR and a perfect oracle OCR system. While the performance dropped when real OCR was used, it still surpassed other



Q: How many objects are preferred by less than 7 people in at least one category?
 SAN: two ✓ MOM: two ✓ SANDY: two ✓

Q: What category does the medium purple color represent?
 SAN: closet ✗ MOM: lisit ✗ SANDY: list ✓

Q: Which item sold the most number of units summed across all the stores?
 SAN: closet ✗ MOM: branch ✓ SANDY: branch ✓

Q: How many units of the item branch were sold in the store sister?
 SAN: 9 ✓ MOM: 9 ✓ SANDY: 9 ✓

Q: Are the values in the chart presented in a percentage scale?
 SAN: no ✓ MOM: no ✓ SANDY: no ✓

Q: How many units of items lead and pure were sold?
 SAN: 8 ✗ MOM: 8 ✗ SANDY: 7 ✓

Figure 5: Example results for different models on DVQA. Outputs of oracle version of SANDY model are shown. SAN completely fails to predict chart-specific answers whereas MOM model often makes small OCR errors (left). Both MOM and SAN are also incapable of properly encoding chart-specific labels in questions (right).

algorithms across all categories. Despite its success, the proposed dynamic encoding used in SANDY is simple and offers a lot of room for expansion. Currently, the dynamic encoding is inferred based on the position of previously detected words. Any error in the OCR system in detecting a single word will propagate throughout the chain and render the encoding for the whole image useless. While this is not a problem for a perfect OCR, removing the cascaded reliance on correctness of each OCR results can help improve performance for an imperfect real-world OCR system.

Recently, multiple compositional models for VQA, such as neural module networks (NMN) [3, 4, 11], have been developed. These recursive neural network systems consist of stacked sub-networks that are executed to answer questions, and they work well on the compositional reasoning questions in CLEVR [13]. However, current NMN formulations are unable to produce chart-specific answers, so they cannot be used for DVQA without suitable modifications.

SANDY and MOM are both built on top of SAN, and they try to solve chart-specific answer generation in two distinct ways that are agnostic to SAN’s actual architecture. SANDY uses DEM and OCR to encode an image’s text, whereas MOM attempts to predict the location of the text it needs to generate an answer. As VQA systems continue to evolve, upgrading SAN with a better VQA algorithm could improve the performance of our systems.

Our dataset currently only contains bar charts. We are developing a follow-up version that will contain pie charts, plots, and other visualizations in addition to bar-charts. Since neither MOM nor SANDY are designed specifically

for bar-charts, they can operate on these alternative diagrams with only minor modifications.

We conducted an additional study to assess how well these models work on real bar charts. We manually annotated over 500 structure understanding questions for real bar charts scraped from the Internet. Without any fine-tuning, all of the SAN-based models achieved about 59% accuracy on these questions, a 15% absolute improvement over the image blind (QUES) baseline. This shows a positive transfer from synthetic to real-world bar charts. Training on entirely real charts would be ideal, but even then there would likely be a benefit to using synthetic datasets as a form of data augmentation [17].

7. Conclusion

Here, we described DVQA, a dataset for understanding bar charts. We demonstrated that VQA algorithms are incapable of answering simple DVQA questions. We proposed two DVQA algorithms that can handle chart-specific words in questions and answers. Solving DVQA will enable systems that can be utilized to intelligently query massive repositories of human-generated data, which would be an enormous aid to scientists and businesses. We hope that the DVQA dataset, which will be made publicly available, will promote the study of issues that are generally ignored with VQA with natural images, *e.g.*, out-of-vocabulary words and dynamic question encoding. We also hope that DVQA will serve as an important proxy task for studying visual attention, memory, and reasoning capabilities.

References

- [1] A. Agrawal, D. Batra, D. Parikh, and A. Kembhavi. Don't just assume; look and answer: Overcoming priors for visual question answering. In *CVPR*, 2018. 4
- [2] R. A. Al-Zaidy and C. L. Giles. Automatic extraction of data from bar charts. In *Proceedings of the 8th International Conference on Knowledge Capture*, page 30. ACM, 2015. 2
- [3] J. Andreas, M. Rohrbach, T. Darrell, and D. Klein. Deep compositional question answering with neural module networks. In *CVPR*, 2016. 2, 8
- [4] J. Andreas, M. Rohrbach, T. Darrell, and D. Klein. Learning to compose neural networks for question answering. In *NAACL*, 2016. 2, 8
- [5] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. VQA: Visual question answering. In *ICCV*, 2015. 1, 2
- [6] S. Elzer, S. Carberry, and I. Zukerman. The automated understanding of simple bar charts. *Artificial Intelligence*, 175(2):526–555, 2011. 2
- [7] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016. 2
- [8] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *CVPR*, 2017. 2
- [9] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *ICML*, 2006. 5
- [10] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [11] R. Hu, J. Andreas, M. Rohrbach, T. Darrell, and K. Saenko. Learning to reason: End-to-end module networks for visual question answering. In *ICCV*, 2017. 8
- [12] A. Jabri, A. Joulin, and L. van der Maaten. Revisiting visual question answering baselines. In *ECCV*, 2016. 3, 4
- [13] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2017. 2, 8
- [14] K. Kafle and C. Kanan. Answer-type prediction for visual question answering. In *CVPR*, 2016. 2
- [15] K. Kafle and C. Kanan. An analysis of visual question answering algorithms. In *ICCV*, 2017. 2, 4
- [16] K. Kafle and C. Kanan. Visual question answering: Datasets, algorithms, and future challenges. *Computer Vision and Image Understanding*, 2017. 2
- [17] K. Kafle, M. Yousefhusien, and C. Kanan. Data augmentation for visual question answering. In *INLG*, 2017. 8
- [18] S. E. Kahou, A. Atkinson, V. Michalski, A. Kadar, A. Trischler, and Y. Bengio. Figureqa: An annotated figure dataset for visual reasoning. *arXiv preprint arXiv:1710.07300*, 2017. 3
- [19] J. S. Kallimani, K. Srinivasa, and R. B. Eswara. Extraction and interpretation of charts in technical documents. In *Advances in Computing, Communications and Informatics (ICACCI), 2013 International Conference on*, pages 382–387. IEEE, 2013. 2
- [20] V. Kazemi and A. Elqursh. Show, ask, attend, and answer: A strong baseline for visual question answering. *arXiv preprint arXiv:1704.03162*, 2017. 5
- [21] A. Kembhavi, M. Salvato, E. Kolve, M. Seo, H. Hajishirzi, and A. Farhadi. A diagram is worth a dozen images. In *ECCV*, 2016. 2, 3
- [22] A. Kembhavi, M. Seo, D. Schwenk, J. Choi, A. Farhadi, and H. Hajishirzi. Are you smarter than a sixth grader? textbook question answering for multimodal machine comprehension. In *CVPR*, 2017. 2, 3
- [23] J.-H. Kim, K.-W. On, J. Kim, J.-W. Ha, and B.-T. Zhang. Hadamard product for low-rank bilinear pooling. In *ICLR*, 2017. 2
- [24] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [25] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. Bernstein, and L. Fei-Fei. Visual Genome: Connecting language and vision using crowdsourced dense image annotations. 2016. 2
- [26] J. Lu, J. Yang, D. Batra, and D. Parikh. Hierarchical question-image co-attention for visual question answering. In *NIPS*, 2016. 2
- [27] M. Malinowski and M. Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *NIPS*, 2014. 1, 2
- [28] H. Noh and B. Han. Training recurrent answering units with joint loss minimization for VQA. *arXiv preprint arXiv:1606.03647*, 2016. 2
- [29] J. Poco and J. Heer. Reverse-engineering visualizations: Recovering visual encodings from chart images. In *Computer Graphics Forum*, volume 36, pages 353–363. Wiley Online Library, 2017. 2
- [30] M. Ren, R. Kiros, and R. Zemel. Exploring models and data for image question answering. In *NIPS*, 2015. 2
- [31] M. Savva, N. Kong, A. Chhajta, L. Fei-Fei, M. Agrawala, and J. Heer. Revision: Automated classification, analysis and redesign of chart images. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 393–402. ACM, 2011. 2
- [32] Q. Wu, D. Teney, P. Wang, C. Shen, A. Dick, and A. v. d. Hengel. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*, 2017. 2
- [33] K. Xu, J. Ba, R. Kiros, A. Courville, R. Salakhutdinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. 2
- [34] Z. Yang, X. He, J. Gao, L. Deng, and A. J. Smola. Stacked attention networks for image question answering. In *CVPR*, 2016. 2, 5