

Mind Your Outliers! Investigating the Negative Impact of Outliers on Active Learning for Visual Question Answering

Siddharth Karamcheti Ranjay Krishna Li Fei-Fei Christopher D. Manning

Department of Computer Science, Stanford University

{skaramcheti, ranjaykrishna, feifeili, manning}@cs.stanford.edu

Abstract

Active learning promises to alleviate the massive data needs of supervised machine learning: it has successfully improved sample efficiency by an order of magnitude on traditional tasks like topic classification and object recognition. However, we uncover a striking contrast to this promise: across 5 models and 4 datasets on the task of visual question answering, a wide variety of active learning approaches fail to outperform random selection. To understand this discrepancy, we profile 8 active learning methods on a per-example basis, and identify the problem as *collective outliers* – groups of examples that active learning methods prefer to acquire but models fail to learn (e.g., questions that ask about text in images or require external knowledge). Through systematic ablation experiments and qualitative visualizations, we verify that collective outliers are a general phenomenon responsible for degrading pool-based active learning. Notably, we show that active learning sample efficiency increases significantly as the number of collective outliers in the active learning pool decreases. We conclude with a discussion and prescriptive recommendations for mitigating the effects of these outliers in future work.

1 Introduction

Today, language-equipped vision systems such as VizWiz, TapTapSee, BeMyEyes, and CamFind are actively being deployed across a broad spectrum of users.¹ As underlying methods improve, these systems will be expected to operate over diverse visual environments and understand myriad language inputs (Bigham et al., 2010; Tellex et al., 2011; Mei et al., 2016; Zhu et al., 2017; Anderson et al., 2018b; Park et al., 2019). Visual Question Answering (VQA), the task of answering questions about

¹Applications can be found at <https://vizwiz.org/>, <https://taptapsee.com/>, <https://www.bemyeyes.com/>, and <https://camfindapp.com/>

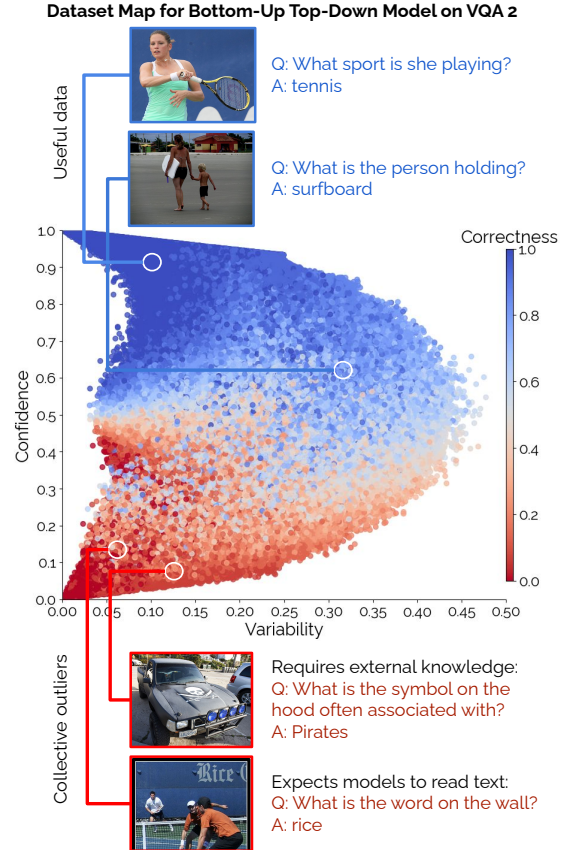


Figure 1: We systematically evaluate active learning on VQA datasets and isolate their inability to perform better than random sampling due to the presence of *collective outliers*. Active learning methods prefer to acquire these outliers, which are hard and often impossible for models to learn. We show that Dataset Maps, like the one shown here, can heuristically identify these collective outliers as examples assigned low model confidence and prediction variability during training.

visual inputs, is a popular benchmark used to evaluate progress towards such open-ended systems (Agrawal et al., 2015; Krishna et al., 2017; Gordon et al., 2018; Hudson and Manning, 2019). Unfortunately, today’s VQA models are data hungry: Their performance scales monotonically with more train-

ing data (Lu et al., 2016; Lin and Parikh, 2017), motivating the need for data acquisition mechanisms such as active learning, which maximize performance while minimizing expensive data labeling.

While active learning is often key to effective data acquisition when such labeled data is difficult to obtain (Lewis and Catlett, 1994; Tong and Koller, 2001; Culotta and McCallum, 2005; Settles, 2009), we find that 8 modern active learning methods (Gal et al., 2017; Siddhant and Lipton, 2018; Lowell et al., 2019) show little to no improvement in sample efficiency across 5 models on 4 VQA datasets – indeed, in some cases performing worse than randomly selecting data to label. This finding is in stark contrast to the successful application of active learning methods on a variety of traditional tasks, such as topic classification (Siddhant and Lipton, 2018; Lowell et al., 2019), object recognition (Deng et al., 2018), digit classification (Gal et al., 2017), and named entity recognition (Shen et al., 2017). Our negative results hold even when accounting for common active learning ailments: cold starts, correlated sampling, and uncalibrated uncertainty. We mitigate the cold start challenge of needing a representative initial dataset by varying the size of the seed set in our experiments. We account for sampling correlated data within a given batch by including Core-Set selection (Sener and Savarese, 2018) in the set of active learning methods we evaluate. Finally, we use deep Bayesian active learning to calibrate model uncertainty to high-dimensional data (Houlsby et al., 2011; Gal and Ghahramani, 2016; Gal et al., 2017).

After concluding that negative results are consistent across all experimental conditions, we investigate active learning’s ineffectiveness on VQA as a data problem and identify the existence of *collective outliers* (Han and Kamber, 2000) as the source of the problem. Leveraging recent advances in model interpretability, we build *Dataset Maps* (Swayamdipta et al., 2020), which distinguish between collective outliers and useful data that improve validation set performance (see Figure 1). While global outliers deviate from the rest of the data and are often a consequence of labeling error, collective outliers cluster together; they may not individually be identifiable as outliers but collectively deviate from other examples in the dataset. For instance, VQA-2 (Goyal et al., 2017) is riddled with collections of hard questions that require external knowledge to answer (e.g., “What is the symbol

on the hood often associated with?”) or that ask the model to read text in the images (e.g., “What is the word on the wall?”). Similarly, GQA (Hudson and Manning, 2019) asks underspecified questions (e.g., “what is the person wearing?” which can have multiple correct answers). Collective outliers are not specific to VQA, but can similarly be found in many open-ended tasks, including visual navigation (Anderson et al., 2018b) (e.g., “Go to the grandfather clock” requires identifying rare grandfather clocks), and open-domain question answering (Kwiatkowski et al., 2019), amongst others.

Using Dataset Maps, we profile active learning methods and show that they prefer acquiring collective outliers that models are unable to learn, explaining their poor improvements in sample efficiency relative to random sampling. Building on this, we use these maps to perform ablations where we identify and remove outliers iteratively from the active learning pool, observing correlated improvements in sample efficiency. This allows us to conclude that collective outliers are, indeed, responsible for the ineffectiveness of active learning for VQA. We end with prescriptive suggestions for future work in building active learning methods robust to these types of outliers.

2 Related Work

Our work tests the utility of multiple recent active learning methods on the open-ended understanding task of VQA. We draw on the dataset analysis literature to identify collective outliers as the bottleneck hindering active learning methods in this setting.

Active Learning. Active learning strategies have been successfully applied to image recognition (Joshi et al., 2009; Sener and Savarese, 2018), information extraction (Scheffer et al., 2001; Finn and Kushmerick, 2003; Jones et al., 2003; Culotta and McCallum, 2005), named entity recognition (Hachey et al., 2005; Shen et al., 2017), semantic parsing (Dong et al., 2018), and text categorization (Lewis and Gale, 1994; Hoi et al., 2006). However, these same methods struggle to outperform a random baseline when applied to the task of VQA (Lin and Parikh, 2017; Jedoui et al., 2019). To study this discrepancy, we systematically apply 8 diverse active learning methods to VQA, including methods that use model uncertainty (Abramson and Freund, 2004; Collins et al., 2008; Joshi et al., 2009), Bayesian uncertainty (Gal and Ghahramani, 2016; Kendall and Gal, 2017), disagreement (Houlsby

et al., 2011; Gal et al., 2017), and Core-Set selection (Sener and Savarese, 2018).

Visual Question Answering. Progress on VQA has been heralded as a marker for progress on general open-ended understanding tasks, resulting in several benchmarks (Agrawal et al., 2015; Malinowski et al., 2015; Ren et al., 2015a; Johnson et al., 2017; Goyal et al., 2017; Krishna et al., 2017; Suhr et al., 2019; Hudson and Manning, 2019) and models (Zhou et al., 2015; Fukui et al., 2016; Lu et al., 2016; Yang et al., 2016; Zhu et al., 2016; Wu et al., 2016; Anderson et al., 2018a; Tan and Bansal, 2019; Chen et al., 2020). To ensure that our negative results are not dataset or model-specific, we sample 4 datasets and 5 representative models, each utilizing unique visual and linguistic features and employing different inductive biases.

Interpreting and Analyzing Datasets. Given the prevalence of large datasets in modern machine learning, it is critical to assess dataset properties to remove redundancies (Gururangan et al., 2018; Li and Vasconcelos, 2019) or biases (Torralba and Efros, 2011; Khosla et al., 2012; Bolukbasi et al., 2016), both of which negatively impact sample efficiency. Prior work has used training dynamics to find examples which are frequently forgotten (Krymolowski, 2002; Toneva et al., 2019) versus those that are easy to learn (Bras et al., 2020). This work suggests using two model-specific measures – *confidence* and *prediction variance* – as indicators of a training example’s “learnability” (Chang et al., 2017; Swayamdipta et al., 2020). Dataset Maps (Swayamdipta et al., 2020), a recently introduced framework uses these two measures to profile datasets to find learnable examples. Unlike prior datasets analyzed by Dataset Maps that have a small number of global outliers as hard examples, we discover that VQA datasets contain copious amounts of collective outliers, which are difficult or even impossible for models to learn.

3 Active Learning Experimental Setup

We adopt the standard pool-based active learning setup from prior work (Lewis and Gale, 1994; Settles, 2009; Gal et al., 2017; Lin and Parikh, 2017), consisting of a model \mathcal{M} , initial seed set of labeled examples $(x_i, y_i) \in \mathcal{D}_{\text{seed}}$ used to initialize \mathcal{M} , an unlabeled pool of data $\mathcal{D}_{\text{pool}}$, and an acquisition function $\mathcal{A}(x, \mathcal{M})$. We run active learning over a series of acquisition iterations

	Pool Size	# Answers
VQA-Sports	5,411 [5k]	20
VQA-Food	4,082 [4k]	20
VQA-2	411,272 [400k]	3130
GQA	943,000 [900k]	1842

Table 1: We evaluate active learning on 4 VQA datasets. We display the total available training examples, effective pool sizes we use [in brackets], and the total number of possible answers for each dataset.

T where at each iteration we acquire a batch of B new examples per: $\hat{x} \in \mathcal{D}_{\text{pool}}$ to label per $\hat{x} = \arg \max_{x \in \mathcal{D}_{\text{pool}}} \mathcal{A}(x, \mathcal{M})$.

Acquiring an example often refers to using an oracle or human expert to annotate a new example with a correct label. We follow prior work to simulate an oracle using existing datasets, forming $\mathcal{D}_{\text{seed}}$ from a fixed percentage of the full dataset, and using the remainder as $\mathcal{D}_{\text{pool}}$ (Gal et al., 2017; Lin and Parikh, 2017; Siddhant and Lipton, 2018). We re-train \mathcal{M} after each acquisition iteration.

Prior work has noted the impact of seed set size on active learning performance (Lin and Parikh, 2017; Misra et al., 2018; Jedoui et al., 2019). We run multiple active learning evaluations with varying seed set sizes (ranging from 5% to 50% of the full pool size). We keep the size of each acquisition batch B to a constant 10% of the overall pool size.

3.1 Models

Visual Question Answering (VQA) requires reasoning over two modalities: images and text. Most models use feature “backbones” (e.g., features from object recognition models pretrained on ImageNet, and pretrained word vectors for text). For image features we use grid-based features from ResNet-101 (He et al., 2016), or object-based features from Faster R-CNN (Ren et al., 2015b) fine-tuned on Visual Genome (Anderson et al., 2018a). We evaluate with a representative sample of existing VQA models, including the following:²

LogReg is a logistic regression model that uses either ResNet-101 or Faster R-CNN image features with mean-pooled GloVe question embeddings (Pennington et al., 2014). Although these models

²Key implementation details can be found in the appendix. In the interest of full reproducibility and further work in active learning and VQA, we release our code and results here: <https://github.com/siddk/vqa-outliers>.

are not as performant as the subsequent models, logistic regression has been effective on VQA (Suhr et al., 2019), and is pervasive in the active learning literature (Schein and Ungar, 2007; Yang and Loog, 2018; Musmann and Liang, 2018).

LSTM-CNN is a standard model introduced with VQA-1 (Agrawal et al., 2015). We use more performant ResNet-101 features instead of the original VGGNet features as our visual backbone.

BUTD (Bottom-Up Top-Down Attention) uses object-based features in tandem with attention over objects (Anderson et al., 2018a). BUTD won the 2017 VQA Challenge (Teney et al., 2018), and has been a consistent baseline for recent work in VQA.

LXMERT is a large multi-modal transformer model that uses BUTD’s object features and contextualized BERT (Devlin et al., 2019) language features (Tan and Bansal, 2019). LXMERT is pre-trained on a corpus of aligned image-and-textual data spanning MS COCO, Visual Genome, VQA-2, NLVR-2, and GQA (Lin et al., 2014; Krishna et al., 2017; Goyal et al., 2017; Suhr et al., 2019; Hudson and Manning, 2019), initializing a cross-modal representation space conducive to fine-tuning.³

3.2 Acquisition Functions

Several active learning methods have been developed to account for different aspects of the machine learning training pipeline: while some acquire examples with high aleatoric uncertainty (Settles, 2009) (having to do with the natural uncertainty in the data) or epistemic uncertainty (Gal et al., 2017) (having to do with the uncertainty in the modeling/learning process), others attempt to acquire examples that reflect the distribution of data in the pool (Sener and Savarese, 2018). We sample a diverse set of these methods:

Random Sampling serves as our baseline passive approach for acquiring examples.

Least Confidence acquires examples with lowest model prediction probability (Settles, 2009).

³Results for LXMERT in Tan and Bansal (2019) are reported *after* pretraining on training and validation examples from the VQA datasets we use. While this is fair if the goal is optimizing for test performance, this exposure to training and validation examples leaks important information; to remedy this, we obtained a model checkpoint from the LXMERT authors trained *without* VQA data. This is also why our LXMERT results are lower than the numbers reported in the original paper – however, the general boost provided by cross-modal pretraining holds.

Entropy acquires examples with the highest entropy in the model’s output (Settles, 2009).

MC-Dropout Entropy (Monte-Carlo Dropout with Entropy acquisition) acquires examples with high entropy in the model’s output averaged over multiple passes through a neural network with different dropout masks (Gal and Ghahramani, 2016). This process is a consequence of a theoretical casting of dropout as approximate Bayesian inference in deep Gaussian processes.

BALD (Bayesian Active Learning by Disagreement) builds upon Monte-Carlo Dropout by proposing a decision theoretic objective; it acquires examples that maximise the decrease in expected posterior entropy (Houlsby et al., 2011; Gal et al., 2017; Siddhant and Lipton, 2018) – capturing “disagreement” across different dropout masks.

Core-Set Selection samples examples that capture the diversity of the data pool (Sener and Savarese, 2018; Coleman et al., 2020). It acquires examples to minimize the distance between an example in the unlabeled pool to its closest labeled example. Since Core-Set selection operates over a representation space (and not an output distribution, like prior strategies) and VQA models operate over two modalities, we employ three Core-Set variants: **Core-Set (Language)** and **Core-Set (Vision)** operate over their respective representation spaces while **Core-Set (Fused)** operates over the “fused” vision and language representation space.

4 Experimental Results

We evaluate the 8 active learning strategies across the 5 models described in the previous section. Figures 2–5 show a representative sample of active learning results across datasets. Due to space constraints, we only visualize 4 active learning strategies – Least-Confidence, BALD, CoreSet-Fused, and the Random Baseline – using 3 models (LSTM-CNN, BUTD, LXMERT).⁴ Results and trends are consistent across the different acquisition functions, models and seed set sizes (see the appendix for results with other models, acquisition functions, and seed set sizes). We now go on to provide descriptions of the datasets we evaluate against, and the corresponding results.

⁴For LXMERT, running Core-Set selection is prohibitive, so we omit these results; please see Appendix B for more details.




Figure 2: Results for varied active learning methods on VQA-Sports, a simplified VQA dataset. Strategies perform on par with or worse than the random baseline, when using 10% of the full dataset as the seed set.




Figure 3: Results for the full VQA-2 dataset, also using 10% of the full dataset as a seed set. Similar to the plot above, all active learning methods perform similar to a random baseline.



Figure 4: Results on VQA-2 using 50% of the dataset as a seed set. While methods are *relatively* better when using a larger seed set—confirming results from (Lin and Parikh, 2017)—no methods outperform random.




Figure 5: Results on GQA using 10% of the dataset for the seed set. Even with different question structures, the above trends hold, with strategies performing worse than or equivalent to random.




Figure 6: We visualize the difference in acquisition preferences between random and active learning acquisitions (least confidence and BALD) across multiple iterations. Active learning methods prefer to sample impossible examples which models are unable to learn, hurting sample efficiency relative to the random baseline.

4.1 Simplified VQA Datasets

One complexity of VQA is the size of the output space and the number of examples present (Agrawal et al., 2015; Goyal et al., 2017); VQA-2 has 400k training examples, and in excess of 3k possible answers (see Table 1). However, prior work in active learning focuses on smaller datasets like the 10-class MNIST dataset (Gal et al., 2017), binary classification (Siddhant and Lipton, 2018), or small-cardinality (≤ 20 classes) text categorization (Lowell et al., 2019). To ensure our results and conclusions are not due to the size of the output space, we build two meaningful, but narrow-domain VQA datasets from subsets of VQA-2. These simplified datasets reduce the complexity of the underlying learning problem and provide a fair comparison to existing active learning literature.

VQA-Sports. We generate VQA-Sports by compiling a list of 20 popular sports (e.g., soccer, football, tennis, etc.) in VQA-2, and restricting the set of questions to those with answers in this list. We picked the sports categories by ranking the GloVe vector similarity between the word “sports” to answers in VQA-2, and selected the 20 most commonly occurring answers.

VQA-Food. We generate the VQA-Food dataset similarly, compiling a list of the 20 commonly occurring food categories by GloVe vector similarity to the word “food.”

Results. Figure 2 presents results for VQA-Sports, with an initial seed set restricted to 10% of the total pool (500 examples). The appendix reports similar results on VQA-Food. For LSTM-CNN, *Least-Confidence* appears to be slightly more sample efficient, while all other strategies perform

on par with or worse than random. For BUTD, all methods are on par with random; for LXMERT, they perform worse than random. Generally on VQA-Sports, active learning performance varies, but fails to outperform random acquisition.

4.2 VQA-2

VQA-2 is the canonical dataset for evaluating VQA models (Goyal et al., 2017). In keeping with prior work (Anderson et al., 2018a; Tan and Bansal, 2019), we filter the training set to only include answers that appear at least 9 times, resulting in 3130 unique answers. Unlike traditional VQA-2 evaluation, which treats the task as a *multi-label* binary classification problem, we follow prior active learning work on VQA (Lin and Parikh, 2017), which formulates it as a *multi-class* classification problem, enabling the use of acquisition functions such as uncertainty sampling and BALD.

Results. Figures 3 and 4 show results on VQA-2 with different seed set sizes – 10% (40k examples) and 50% (200k examples). Active learning performs relatively better with larger seed sets but still underperforms random. Surprisingly, when initialized with 50% of the pool as the seed set, the gain in validation accuracy after acquiring the entire pool of examples (400k examples total) is only 2%. This is an indication that the lack of sample efficiency might be a result of the underlying data, a problem we explore in the next section.

4.3 GQA

GQA was introduced as a means for evaluating compositional reasoning (Hudson and Manning, 2019). Unlike VQA’s natural human-written questions, GQA contains synthetic questions of the form “what is inside the bottle the glasses are to




Figure 7: Example groups of collective outliers in the VQA-2 and GQA datasets.

the right of?”. We use the standard GQA training set of 943k questions, 900k of which we use for the active learning pool.

Results. Figure 5 shows results on GQA using a seed set of 10% of the full pool (90k examples). Despite its notable differences in question structure to VQA-2, active learning still performs on par with or slightly worse than random.

5 Analysis via Dataset Maps

The previous section shows that active learning fails to improve over random acquisition on VQA across models and datasets. A simple question remains – *why*? One hypothesis is that sample inefficiency stems from the data itself: there is only a 2% gain in validation accuracy when training on half versus the whole dataset. Working from this, we characterize the underlying datasets using Dataset Maps (Swayamdipta et al., 2020) and discover that active learning methods prefer sampling “hard-to-learn” examples, leading to poor performance.

Mapping VQA Datasets. A Dataset Map (Swayamdipta et al., 2020) is a model-specific graph for profiling the learnability of individual training examples. Dataset Maps present holistic pictures of classification datasets relative to the training dynamics of a given model; as a model trains for multiple epochs and sees the same examples repeatedly, the mapping process logs statistics about the confidence assigned to individual predictions. Maps then visualize these statistics against two axes: the y-axis plots the average model confidence assigned to the correct answer over training epochs, while the x-axis plots the spread, or variability, of these values. This introduces a 2D representation of a dataset (viewed through its relationship with individual model) where examples are placed on the map by coarse statistics describing their “learnability”. We show the Dataset Map for BUTD trained on VQA-2 in Figure 1. For our work, we build this map post-hoc, training on the

entire pool as a means for analyzing what active learning is doing – treating it as a diagnostic tool for identifying the root cause why active learning seems to fail for VQA.

In an ideal setting, the majority of examples in the training set should lie in the upper half of the graph – i.e., the mean confidence assigned to the correct answer should be relatively high. Examples towards the upper-left side represent the “easy-to-learn” examples, as the variability in the confidence assigned by the model over time is fairly low.

A curious feature of VQA-2 and other VQA datasets is the presence of the 25-30% of examples in the bottom-left of the map (shown in red in Figure 1) – examples that have low confidence and variability. In other words, models are unable to learn a large proportion of training examples. While prior work attributes examples in this quadrant to “labeling errors” (Swayamdipta et al., 2020), labeling errors in VQA are sparse, and cannot account for the density of such examples in these maps.

Interpreting Acquisitions. We profile the acquisitions made by each active learning method, contextualizing the acquired examples via their placement on the associated Dataset Map. We segregate training examples into four buckets using the map’s y-axis: easy (≥ 0.75), medium (≥ 0.50), hard (≥ 0.25), and impossible (≥ 0.00). Ideally, active learning should be robust to “hard-to-learn” examples, focusing instead on learnable, high uncertainty examples towards the upper-right portion of the Dataset Map. Instead, we find that active learning methods acquire a large proportion of impossible examples early on and concentrate on the easier examples only after the impossible examples dwindle (see Figure 6). In contrast, the random baseline acquires examples proportional to each bucket’s density in the underlying map; acquiring easier examples earlier and performing on par with or better than all others.




Figure 8: Using Dataset Maps, we remove hard-to-learn examples, which we identify as collective outliers. With the outliers removed, active learning methods demonstrate up to 2–3x sample efficiency versus random sampling.

6 Collective Outliers

This leaves two questions: 1) can we characterize these “hard” examples, and 2) are these examples responsible for the ineffectiveness of active learning on VQA? We first identify hard-to-learn examples as collective outliers and explain why active learning methods prefer to acquire them. Next, we perform ablation experiments, removing these outliers from the active learning pool iteratively, and demonstrate a corresponding boost in sample efficiency relative to random acquisition.

Hard Examples are Collective Outliers. Collective outliers are groups of examples that deviate from the rest of the examples but cluster together (Han and Kamber, 2000) – they often present as fundamental subproblems of a broader task. For instance (Figure 7), in VQA-2, we identify clusters of hard-to-learn examples that require optical character recognition (OCR) for reasoning about text (e.g., “What is the first word on the black car?”); another cluster requires external knowledge to answer (“What is the symbol on the hood often associated with?”). In GQA, we identify different clusters of collective outliers; one cluster stems from innate underspecification (e.g., “what is on the shelf?” with multiple objects present on the shelf); another cluster requires multiple reasoning hops difficult for current models (e.g., “What is the vehicle that is driving down the road the box is on the side of?”).

We sample 100 random “hard-to-learn” examples from both VQA-2 and GQA and find that 100% of the examples belong to one of the two aforementioned collectives. Since hard-to-learn examples constitute 25–30% of the data pool, active learning methods cannot avoid them. Uncertainty-

based methods (e.g., Least-Confidence, Entropy, Monte-Carlo Dropout) identify them as valid acquisition targets because models lack the capacity to correctly answer these examples, assigning low confidence and high uncertainty. Disagreement-based methods (e.g., BALD) are similar; model confidence is generally low but high variance (lower middle/lower right of the Dataset Maps). Finally, diversity methods (e.g., Core-Set selection) identify these examples as different enough from the existing pool to warrant acquisition, but fail to learn meaningful representations, fueling a vicious cycle wherein they continue to pick these examples.

Ablating Outliers. To verify that collective outliers are responsible for the degradation of active learning performance, we re-run our experiments using active learning pools with varying numbers of outliers removed. To remove these outliers, we sort and remove all examples in the data pool using the product of their model confidence and prediction variability (x and y-axis values of the Dataset Maps). We systematically remove examples with a low product value and observe how active learning performance changes (see Figure 8).

We observe a 2–3x improvement in sample efficiency when removing 50% of the entire data pool, consisting mainly of collective outliers (Figure 8c). This improvement decreases if we only remove 25% of the full pool (Figure 8b), and further degrades if we remove only 10% (Figure 8a). This ablation demonstrates that active learning methods are more sample efficient than the random baseline when collective outliers are absent from the unlabelled pool.

7 Discussion and Future Work

This paper asks a simple question – why does the modern neural active learning toolkit fail when applied to complex, open ended tasks? While we focus on VQA, collective outliers are abundant in tasks such as natural language inference (Bowman et al., 2015; Williams et al., 2018) and open-domain question answering (Kwiatkowski et al., 2019), amongst others. More insidious is their nature; collective outliers can take multiple forms, requiring external domain knowledge or “common-sense” reasoning, containing underspecification, or requiring capabilities beyond the scope of a given model (e.g., requiring OCR ability). While we perform ablations in this work removing collective outliers, demonstrating that active learning fails as collective outliers take up larger portions of the dataset, this is only an analytical tool; these outliers are, and will continue to be, pervasive in open-ended datasets – and as such, we will need to develop better tools for learning (and performing active learning) in their presence.

Selective Classification. One potential direction for future work is to develop systems that abstain when they encounter collective outliers. Historical artificial intelligence systems, such as SHRDLU (Winograd, 1972) and QUALM (Lehnert, 1977), were designed to flag input sequences that they were not designed to parse. Ideas from those methods can and should be resurrected using modern techniques; for example, recent work suggests that a simple classifier can be trained to identify out-of-domain data inputs, provided a seed out-of-domain dataset (Kamath et al., 2020). Active learning methods can be augmented with a similar classifier, which re-calibrates active learning uncertainty scores with this classifier’s predictions. Other work learns to identify novel utterances by learning to intelligently set thresholds in representation space (Karamcheti et al., 2020), a powerful idea especially if combined with other representation-centric active learning methods like Core-Set Sampling (Sener and Savarese, 2018).

Active Learning with Global Reasoning. Another direction for future work to explore is to leverage Dataset Maps to perform more global, holistic reasoning over datasets, to intelligently identify promising examples – in a sense, baking part of the analysis done in this work directly into the active learning algorithms. A possible instantiation

of this idea would be in training a discriminator to differentiate between “learnable” examples (upper half of each Dataset Map) from the “unlearnable”, collective outliers with low confidence and low variability. Between each active learning acquisition iteration, one can generate an updated Dataset Map, thereby reflecting what models are learning as they obtain new labeled examples.

Machine learning systems deployed in real-world settings will inevitably encounter open-world datasets, ones that contain a mixture of learnable and unlearnable inputs. Our work provides a framework to study when models encounter such inputs. Overall, we hope that our experiments serve as a catalyst for future work on evaluating active learning methods with inputs drawn from open-world datasets.

Reproducibility

All code for data preprocessing, model implementation, and active learning algorithms is made available at <https://github.com/siddk/vqa-outliers>. Additionally, this repository also contains the full set of results and dataset maps as well.

The authors are fully committed to maintaining this repository, in terms of both functionality and ease of use, and will actively monitor both email and Github Issues should there be problems.

Acknowledgements

We thank Kaylee Burns, Eric Mitchell, Stephen Mussman, Dorsa Sadigh, and our anonymous ACL reviewers for their useful feedback on earlier versions of this paper. We are also grateful to Hao Tan for providing us with the LXMERT checkpoint trained without access to VQA datasets, as well as for general LXMERT fine-tuning pointers.

Siddharth Karamcheti is graciously supported by the Open Philanthropy Project AI Fellowship. Christopher D. Manning is a CIFAR Fellow.

References

- Yotam Abramson and Yoav Freund. 2004. Active learning for visual object recognition. Technical report, University of California, San Diego.
- Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C. Lawrence Zitnick, Devi Parikh, and Dhruv Batra. 2015. VQA: Visual question answering. *International Journal of Computer Vision*, 123:4–31.

- Peter Anderson, X. He, C. Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018a. Bottom-up and top-down attention for image captioning and visual question answering. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6077–6086.
- Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. 2018b. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Computer Vision and Pattern Recognition (CVPR)*.
- Jeffrey P Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, and Tom Yeh. 2010. VizWiz: nearly real-time answers to visual questions. In *User Interface Software and Technology (UIST)*, pages 333–342.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4349–4357.
- Samuel Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew Peters, Ashish Sabharwal, and Yejin Choi. 2020. Adversarial filters of dataset biases. In *International Conference on Machine Learning (ICML)*, pages 1078–1088.
- Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. 2017. Active bias: Training more accurate neural networks by emphasizing high variance samples. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1002–1012.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *European Conference on Computer Vision (ECCV)*, pages 104–120.
- Cody Coleman, Christopher Yeh, Stephen Mussmann, Baharan Mirzasoleiman, Peter Bailis, Percy Liang, Jure Leskovec, and Matei Zaharia. 2020. Selection via proxy: Efficient data selection for deep learning. In *International Conference on Learning Representations (ICLR)*.
- Brendan Collins, Jia Deng, Kai Li, and Li Fei-Fei. 2008. Towards scalable dataset construction: An active learning approach. In *European Conference on Computer Vision (ECCV)*, pages 86–98.
- Aron Culotta and Andrew McCallum. 2005. Reducing labeling effort for structured prediction tasks. In *Association for the Advancement of Artificial Intelligence (AAAI)*, pages 746–751.
- Yue Deng, KaWai Chen, Yilin Shen, and Hongxia Jin. 2018. Adversarial active learning for sequences labeling and generation. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4012–4018.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Association for Computational Linguistics (ACL)*, pages 4171–4186.
- Li Dong, Chris Quirk, and Mirella Lapata. 2018. Confidence modeling for neural semantic parsing. In *Association for Computational Linguistics (ACL)*.
- Aidan Finn and Nicolas Kushmerick. 2003. Active learning selection strategies for information extraction. In *Proceedings of the International Workshop on Adaptive Text Extraction and Mining (ATEM-03)*, pages 18–25.
- 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. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning (ICML)*.
- Yarin Gal, R. Islam, and Zoubin Ghahramani. 2017. Deep Bayesian active learning with image data. In *International Conference on Machine Learning (ICML)*.
- Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. 2018. IQA: Visual question answering in interactive environments. In *Computer Vision and Pattern Recognition (CVPR)*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *Computer Vision and Pattern Recognition (CVPR)*.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A Smith. 2018. Annotation artifacts in natural language inference data. In *Association for Computational Linguistics (ACL)*, pages 107–112.
- Ben Hachey, Beatrice Alex, and Markus Becker. 2005. Investigating the effects of selective sampling on the annotation task. In *Computational Natural Language Learning (CoNLL)*, pages 144–151.

- Jiawei Han and Micheline Kamber. 2000. *Data Mining: Concepts and Techniques*. Morgan Kaufmann.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*.
- Steven CH Hoi, Rong Jin, Jianke Zhu, and Michael R Lyu. 2006. Batch mode active learning and its application to medical image classification. In *Proceedings of the 23rd international conference on Machine learning*, pages 417–424.
- Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. 2011. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*.
- Drew A. Hudson and Christopher D. Manning. 2019. GQA: A new dataset for real-world visual reasoning and compositional question answering. In *Computer Vision and Pattern Recognition (CVPR)*.
- Khaled Jedoui, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. 2019. Deep Bayesian active learning for multiple correct outputs. *arXiv preprint arXiv:1912.01119*.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Computer Vision and Pattern Recognition (CVPR)*.
- Rosie Jones, Rayid Ghani, Tom Mitchell, and Ellen Riloff. 2003. Active learning for information extraction with multiple view feature sets. In *International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 26–34.
- Ajay J Joshi, Fatih Porikli, and Nikolaos Panikolopoulos. 2009. Multi-class active learning for image classification. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2372–2379.
- Amita Kamath, Robin Jia, and Percy Liang. 2020. Selective question answering under domain shift. In *Association for Computational Linguistics (ACL)*.
- Siddharth Karamcheti, Dorsa Sadigh, and Percy Liang. 2020. Learning adaptive language interfaces through decomposition. In *EMNLP Workshop for Interactive and Executable Semantic Parsing (IntEx-SemPar)*.
- Alex Kendall and Yarin Gal. 2017. What uncertainties do we need in Bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5574–5584.
- Aditya Khosla, Tinghui Zhou, Tomasz Malisiewicz, Alexei A Efros, and Antonio Torralba. 2012. Undoing the damage of dataset bias. In *European Conference on Computer Vision (ECCV)*, pages 158–171.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidi, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Fei-Fei Li. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123:32–73.
- Yuval Krymolowski. 2002. Distinguishing easy and hard instances. In *International Conference on Computational Linguistics (COLING)*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: A benchmark for question answering research. In *Association for Computational Linguistics (ACL)*.
- Wendy Lehnert. 1977. *The Process of Question Answering*. Ph.D. thesis, Yale University.
- David D Lewis and Jason Catlett. 1994. Heterogeneous uncertainty sampling for supervised learning. In *International Conference on Machine Learning (ICML)*, pages 148–156.
- David D Lewis and William A Gale. 1994. A sequential algorithm for training text classifiers. In *ACM Special Interest Group on Information Retrieval (SIGIR)*.
- Yi Li and Nuno Vasconcelos. 2019. Repair: Removing representation bias by dataset resampling. In *Computer Vision and Pattern Recognition (CVPR)*, pages 9572–9581.
- 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 (ECCV)*, pages 740–755.
- Xiao Lin and Devi Parikh. 2017. Active learning for visual question answering: An empirical study. *arXiv preprint arXiv:1711.01732*.
- David Lowell, Zachary C. Lipton, and Byron C. Wallace. 2019. Practical obstacles to deploying active learning. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. Hierarchical question-image co-attention for visual question answering. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. 2015. Ask your neurons: A neural-based approach to answering questions about images. In *International Conference on Computer Vision (ICCV)*, pages 1–9.

- Hongyuan Mei, Mohit Bansal, and Matthew R Walter. 2016. Listen, attend, and walk: Neural mapping of navigational instructions to action sequences. In *Association for the Advancement of Artificial Intelligence (AAAI)*.
- Ishan Misra, Ross Girshick, Rob Fergus, Martial Hebert, Abhinav Gupta, and Laurens Van Der Maaten. 2018. Learning by asking questions. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11–20.
- Stephen Mussmann and Percy Liang. 2018. On the relationship between data efficiency and error in active learning. In *International Conference on Machine Learning (ICML)*.
- Junwon Park, Ranjay Krishna, Pranav Khadpe, Li Fei-Fei, and Michael Bernstein. 2019. AI-based request augmentation to increase crowdsourcing participation. In *Association for the Advancement of Artificial Intelligence (AAAI)*, volume 7, pages 115–124.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Mengye Ren, Ryan Kiros, and Richard Zemel. 2015a. Exploring models and data for image question answering. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2953–2961.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015b. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 39:1137–1149.
- Tobias Scheffer, Christian Decomain, and Stefan Wrobel. 2001. Active hidden Markov models for information extraction. In *International Symposium on Intelligent Data Analysis*, pages 309–318.
- A. Schein and Lyle H. Ungar. 2007. Active learning for logistic regression: An evaluation. *Machine Learning*, 68:235–265.
- Ozan Sener and Silvio Savarese. 2018. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations (ICLR)*.
- Burr Settles. 2009. Active learning literature survey. Technical report, University of Wisconsin, Madison.
- Yanyao Shen, Hyokun Yun, Zachary C Lipton, Yakov Kronrod, and Animashree Anandkumar. 2017. Deep active learning for named entity recognition. In *Proceedings of the Second Workshop on Representation Learning for NLP (Repl4NLP)*.
- Aditya Siddhant and Zachary C Lipton. 2018. Deep Bayesian active learning for natural language processing: Results of a large-scale empirical study. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. A corpus for reasoning about natural language grounded in photographs. In *Association for Computational Linguistics (ACL)*.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. Dataset cartography: Mapping and diagnosing datasets with training dynamics. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Hao Hao Tan and Mohit Bansal. 2019. LXMERT: Learning cross-modality encoder representations from transformers. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew R Walter, Ashis Gopal Banerjee, Seth J Teller, and Nicholas Roy. 2011. Understanding natural language commands for robotic navigation and mobile manipulation. In *Association for the Advancement of Artificial Intelligence (AAAI)*.
- Damien Teney, Peter Anderson, Xiaodong He, and Anton V. D. Hengel. 2018. Tips and tricks for visual question answering: Learnings from the 2017 challenge. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4223–4232.
- Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. 2019. An empirical study of example forgetting during deep neural network learning. In *International Conference on Learning Representations (ICLR)*.
- Simon Tong and Daphne Koller. 2001. Support vector machine active learning with applications to text classification. *Journal of machine learning research*, 2(0):45–66.
- Antonio Torralba and Alexei A Efros. 2011. Unbiased look at dataset bias. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1521–1528.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Association for Computational Linguistics (ACL)*, pages 1112–1122.
- Terry Winograd. 1972. *Understanding Natural Language*. Academic Press.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Qi Wu, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. 2016. Ask me anything: Free-form visual question answering based on

knowledge from external sources. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4622–4630.

Yazhou Yang and Marco Loog. 2018. A benchmark and comparison of active learning for logistic regression. *Pattern Recognition*, 83.

Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. 2016. Stacked attention networks for image question answering. In *Computer Vision and Pattern Recognition (CVPR)*.

Bolei Zhou, Yuandong Tian, Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. 2015. Simple baseline for visual question answering. *arXiv preprint arXiv:1512.02167*.

Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. 2016. Visual7W: Grounded question answering in images. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4995–5004.

Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. 2017. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *International Conference on Robotics and Automation (ICRA)*, pages 3357–3364.

A Overview

Due to the broad scope of our experiments and analysis, we were unable to fit all our results in the main body of the paper. Furthermore, given the limited length provided by the appendix, we provide only salient implementation details and other representative results here; however, we make all code, models, data, results, active learning implementations available at this link: <https://github.com/siddk/vqa-outliers>.

Generally, any combination of $\{\text{active learning strategy} \times \text{model} \times \text{seed set size} \times \text{analysis/acquisition plot}\}$ is present in this paper, and is available in the public code repository.

B Implementation Details

B.1 Models & Training

Where applicable, we implement our models based on publicly available PyTorch implementations. For the LSTM-CNN model, we base our implementation off of this repository: <https://github.com/Shivanshu-Gupta/Visual-Question-Answering>, while for the Bottom-Up Top-Down Attention Model, we use this repository: <https://github.com/hengyuan-hu/bottom-up-attention-vqa>, keeping default hyperparameters the same.

Logistic Regression. When implementing Logistic Regression, we base our PyTorch implementation on the broadly used Scikit-Learn (<https://scikit-learn.org>) implementation, using the default parameters (including L2 weight decay). We optimize our models via stochastic gradient descent.

LXMERT. As mentioned in Section 3, the default LXMERT checkpoint and fine-tuning code made publicly available in Tan and Bansal (2019) (associated code repository: <https://github.com/airsplay/lxmert>) is pretrained on data from VQA-2 and GQA, leaking information that could substantially affect our active learning results. To mitigate this, we contacted the authors, who kindly provided us with a checkpoint of the model without VQA pretraining.

However, in addition to this model obtaining different results from those reported in the original work, the provided pretrained checkpoint behaves slightly differently during fine-tuning, requiring different hyperparameters than provided in the original repository. We perform a coarse grid search

over hyperparameters, using the LXMERT implementation provided by HuggingFace Transformers (Wolf et al., 2019), and find that using an AdamW optimizer rather than the BERT-Adam Optimizer used in the original work *without any special learning rate scheduling* results in the best fine-tuning performance.

B.2 Acquisition Functions

We use standard implementations of the 8 active learning strategies described, borrowing from prior implementations (Mussmann and Liang, 2018) and existing code repositories (<https://github.com/google/active-learning>). We provide additional details below.

Monte-Carlo Dropout. For our implementations of the deep Bayesian active learning methods (Monte-Carlo Dropout w/ Entropy, BALD), we follow Gal and Ghahramani (2016) and estimate a Dropout distribution via test-time dropout, running multiple forward passes through our neural networks, with different, randomly sampled Dropout masks. We use a value of $k = 10$ forward passes to form our Dropout distribution.

Amortized Core-Set Selection. In the original Core-Set selection active learning work introduced by Sener and Savarese (2018), it is shown that Core-Set selection for active learning can be reduced to a version of the k -centers problem, which can be solved approximately (2-OPT) with a greedy algorithm. However, running this algorithm on high-dimensional representations, across large pools can be prohibitive; Core-Set selection is *batch-aware*, requiring recomputing distances from each “cluster-center” (points in the set of acquired examples) to all points in the active learning pool *after each acquisition in a batch*. While we can run this out completely for smaller datasets (and indeed, this is what we do for our small datasets VQA-Sports and VQA-Food), a single acquisition iteration for a large dataset for the full VQA-2 dataset takes approximately 20 GPU-hours on the resources we have available, or up to 9 days for a single Core-Set selection run. For GQA, performing exact Core-Set selection takes at least twice as long.

To still capture the spirit of Core-Set diversity-based selection in our evaluation, we instead introduce an *amortized implementation of Core-Set selection*, which is comprised of two steps. We first downsample the high-dimensional representations

(of either the fused language and text, or either uni-modal representations) via Principal Component Analysis (PCA) to make the distance computation faster by an order of magnitude. Then, rather than updating distances from examples in our acquired set to points in our pool *after each acquisition* \hat{x} , we delay updates, instead only refreshing the distance computation every 2000 acquisitions (roughly 5% of an acquisition batch for VQA-2). This allows us to report results for Core-Set selection with the three different proposed representations (Fused, Language-Only, Vision-Only) for VQA-2; unfortunately, for GQA and LXMERT (due to the high cost of training), even running this amortized version of Core-Set selection is prohibitive, so we report a subset of results, and omit the rest.

C Active Learning Results

We include further results from our study of active learning applied to VQA, including results on VQA-Food (not included in the main body), active learning results for the two logistic regression models – Log-Reg (ResNet-101) and Log-Reg (Faster R-CNN), as well as with the 4 acquisition strategies not included in the main body of the paper – Entropy, Monte-Carlo Dropout w/ Entropy, Core-Set (Language), and Core-Set (Vision).

C.1 VQA-Food

Figure 9 shows results on VQA-Food with the LSTM-CNN, BUTD, and LXMERT models, with a seed set comprised of 10% of the total pool. The results are mostly similar to those reported in the paper; strategies track or underperform random sampling, with the exception of Least-Confidence for the LSTM-CNN model – however, this is the sole exception, and the LSTM-CNN has the highest training variance of all the models we try.

C.2 Logistic Regression (ResNet-101)

Figure 10 shows active learning results for the LogReg (ResNet-101) model on VQA-Sports (seed set = 10%), and VQA-2 (seed set = 10%, 50%). Results are similar to those reported in the paper, with active learning failing to outperform random acquisition.

C.3 Logistic Regression (Faster R-CNN)

Figure 11 presents the same set of experiments as the prior section, except with the LogReg (Faster R-CNN) model. While the object-based Faster

R-CNN representation enables much higher performance than the ResNet-101 representation, active learning results are consistent with those reported in the paper.

C.4 Other Acquisition Strategies

Figure 12 presents results for the four other active learning strategies we implement – Entropy, Monte Carlo Dropout w/ Entropy, Core-Set (Language), and Core-Set (Vision) – for the BUTD model. Results are across VQA-Sports (seed set = 10%), and VQA-2 (seed set = 10%, 50%) – despite the unique features of each strategy, the trends remain consistent with those in the paper.

D Dataset Maps & Acquisitions

To provide further context around active learning acquisitions across datasets, Figures 13–16 present Dataset Maps and acquisitions for the BUTD Model across VQA-Sports, VQA-Food, and GQA respectively. Interesting to note is that while VQA-Sports and VQA-Food are generally easier, with fewer “hard-to-learn” examples, active learning still has a bias for picking those examples. For GQA, our earlier analysis is confirmed; active learning is picking the collective outliers populating the bottom half of the Dataset Map.




Figure 9: Results for the representative active learning methods on VQA-Food, a simplified VQA dataset similar to VQA-Food, across LSTM-CNN, BUTD, and LXMERT.




Figure 10: Active learning results using the Logistic Regression (ResNet-101) model on VQA-Sports (10% seed set), and VQA-2 (10% and 50% seed set). Most strategies either track or underperform random acquisition.




Figure 11: Active learning results using the Logistic Regression (Faster R-CNN) model on VQA-Sports (10% seed set), and VQA-2 (10% and 50% seed set). While the Faster R-CNN representation leads to better validation accuracies, active learning performance remains consistent.




Figure 12: Results with the BUTD on VQA-Sports, VQA-2 and GQA using the alternative 4 acquisition strategies not included in the main body of the paper. Unsurprisingly, results are consistent with those reported in the paper.




Figure 13: Dataset Maps for the Bottom-Up Top-Down Attention model on VQA-Sports, VQA-Food, and GQA respectively. Note that VQA-Sports and VQA-Food have fewer “hard-to-learn” examples.



Figure 14: Acquisitions with the BUTD Model on VQA-Sports. The dataset has fewer “hard-to-learn” examples, but active learning strategies pick the medium–hard examples, which still negatively impact performance.



Figure 15: Acquisitions with the BUTD Model on VQA-Food. Despite the sparsity of hard examples, active learning strategies still tend towards them. BALD is high-variance, selecting examples all over the map.



Figure 16: Acquisitions with the BUTD Model on the full GQA dataset. Given that the map for GQA is similar to the map for VQA-2, it is not surprising that the active learning acquisitions follow a similar trend, preferring to select “hard-to-learn” examples.