

# Automatic Generation of Contrast Sets from Scene Graphs: Probing the Compositional Consistency of GQA

Yonatan Bitton, Gabriel Stanovsky, Roy Schwartz, Michael Elhadad

## Overview

Models often exploit data artifacts to achieve good test scores.



<https://thegradient.pub/shortcuts-neural-networks-love-to-cheat/>


McCoy, R. Thomas, et al.  
"Right for the Wrong Reasons:  
Diagnosing Syntactic Heuristics in  
Natural Language Inference.",  
ACL 2019.

Gururangan, Suchin, et al.  
"Annotation artifacts in natural  
language inference data.",  
NAACL 2018.

Jia, Robin, et al. "Adversarial  
examples for evaluating reading  
comprehension systems.",  
EMNLP 2017.

Contrast sets quantify this phenomenon. Used as a more accurate  
evaluation the for models true capabilities 🔍.



**Contrast sets**  
Gardner, Matt, et al.  
"Evaluating models'  
local decision  
boundaries via  
contrast sets",  
Findings of EMNLP  
2020



In many cases, contrast sets have been built manually, requiring  
extensive human effort and expertise 🧑.

Original Instance	Contrastive Instance (color = edit)
Hardly one to be faulted for his ambition or his vision, it is genuinely unexpected, then, to see all Park's effort add up to so very little. ... The premise is promising, gags are copious and offbeat humour abounds but it all fails miserably to create any meaningful connection with the audience. (Label: Negative)	Hardly one to be faulted for his ambition or his vision, <b>here we see all Park's effort come to fruition.</b> ... The premise is <b>perfect</b> , gags are <b>hilarious</b> and offbeat humour abounds, <b>and it creates a deep</b> connection with the audience. (Label: Positive)



We propose a method for automatic construction of large contrast sets for the Visual Question Answering task, by leveraging scene-graphs input representations.



We demonstrate the effectiveness of our method on the GQA dataset.



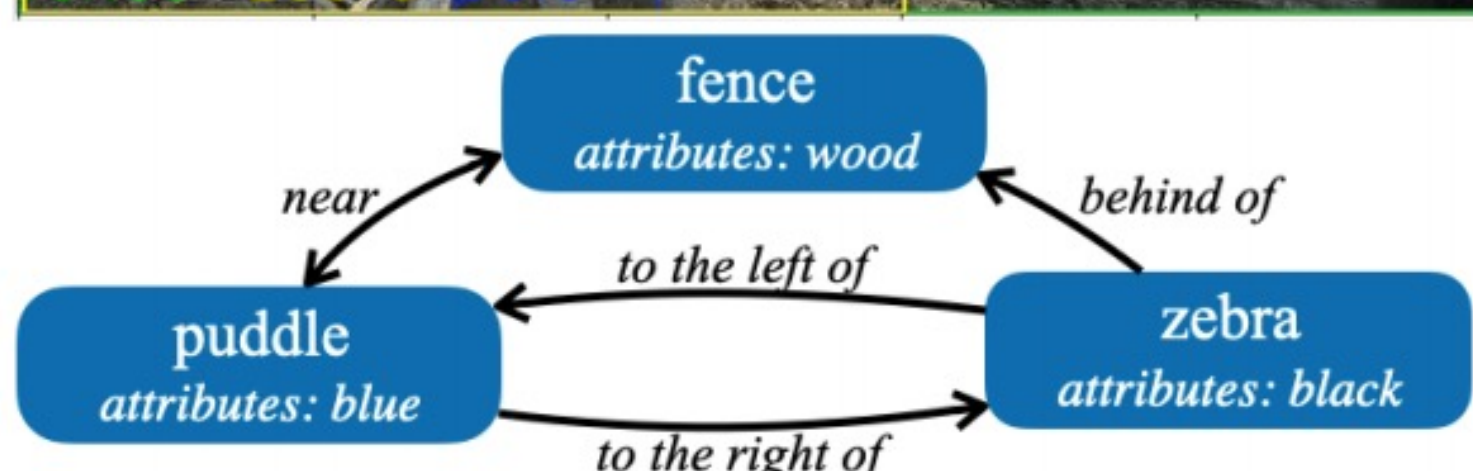

Hudson, Drew A, et al.  
"GQA: A new dataset for real-world  
visual reasoning and compositional  
question answering."  
CVPR 2019.



SEMANTIC REPRESENTATIONS COMPOSITIONAL BALANCED

Starting from  $(image, scene\ graph, Q, A)$  we generate a set of  
variants  $\{(image, scene\ graph, Q_i', A_i')\}$  s.t  $Q_i'$  is a minimal change of  
 $Q$ , and  $A \neq A_i'$ .

Original Q Is there **a fence** near the puddle? Label: Yes Pred: Yes  
Aug. Q #1 Is there **a wall** near the puddle? Label: No Pred: Yes  
Aug. Q #2 **Are** there **men** near the puddle? Label: No Pred: Yes  
Aug. Q #3 Is there **an elephant** near the puddle? Label: No Pred: No



## Automatic Contrast Set Construction



### Identifying Recurring Patterns in GQA

Question template	Tested attributes	Example
On which side is the $X$ ?	Relational (left vs. right)	On which side is the <i>dishwasher</i> ? → On which side are the <i>dishes</i> ?
What color is the $X$ ?	Color identification	What color is the <i>cat</i> ? → What color is the <i>jacket</i> ?
Do you see $X$ or $Y$ ?	Compositions	Do you see <i>laptops</i> or cameras? → Do you see <i>headphones</i> or cameras?
Are there $X$ near the $Y$ ?		<i>Are</i> there any <i>cats</i> near the boat? → <i>Is</i> there any <i>bush</i> near the boat?
Is the $X$ <i>Rel</i> the $Y$ ?	Spatial, relational	Is the boy to the <i>right</i> of the man? → Is the boy to the <i>left</i> of the man?
Is the $X$ <i>Rel</i> the $Y$ ?		Is the boy to the right of the <i>man</i> ? → Is the boy to the right of the <i>zebra</i> ?

### Illustrating the perturbation process

Is the **teddy bear** to the left of a **suitcase**? No → Is the *teddy bear* to the left of a **blanket**? Yes

Is the  $X$  *Rel* the  $Y$ ?



## Main Findings

### Models struggle with our contrast sets


	MAC		LXMERT	
	Original	Aug.	Original	Aug.
On which side is the $X$ ?	68%	57%	94%	81%
What color is the $X$ ?	49%	49%	69%	62%
Are there $X$ near the $Y$ ?	85%	66%	98%	79%
Do you see $X$ or $Y$ ?	88%	53%	95%	65%
Is the $X$ <i>Rel</i> the $Y$ ?	85%	44%	96%	69%
Is the $X$ <i>Rel</i> the $Y$ ?	71%	38%	93%	55%
<b>Overall</b>	<b>65%</b>	<b>52%</b>	<b>84%</b>	<b>67%</b>

### Training on perturbed set leads to more robust models

Model	Training set	Original	Augmented
MAC	Baseline	64.9%	51.5%
	Augmented	64.4%	<b>68.4%</b>
LXMERT	Baseline	83.9%	67.2%
	Augmented	82.6%	<b>77.2%</b>

Since our method is automatic, we can augment  
the training set as well

### Consistency drops as the number of augmentations grow



Augmentations per instance	Contrast sets	Acc.	Consistency
1	11,263	66%	63.4%
3	23,236	67%	51.1%
5	28,968	67%	46.1%