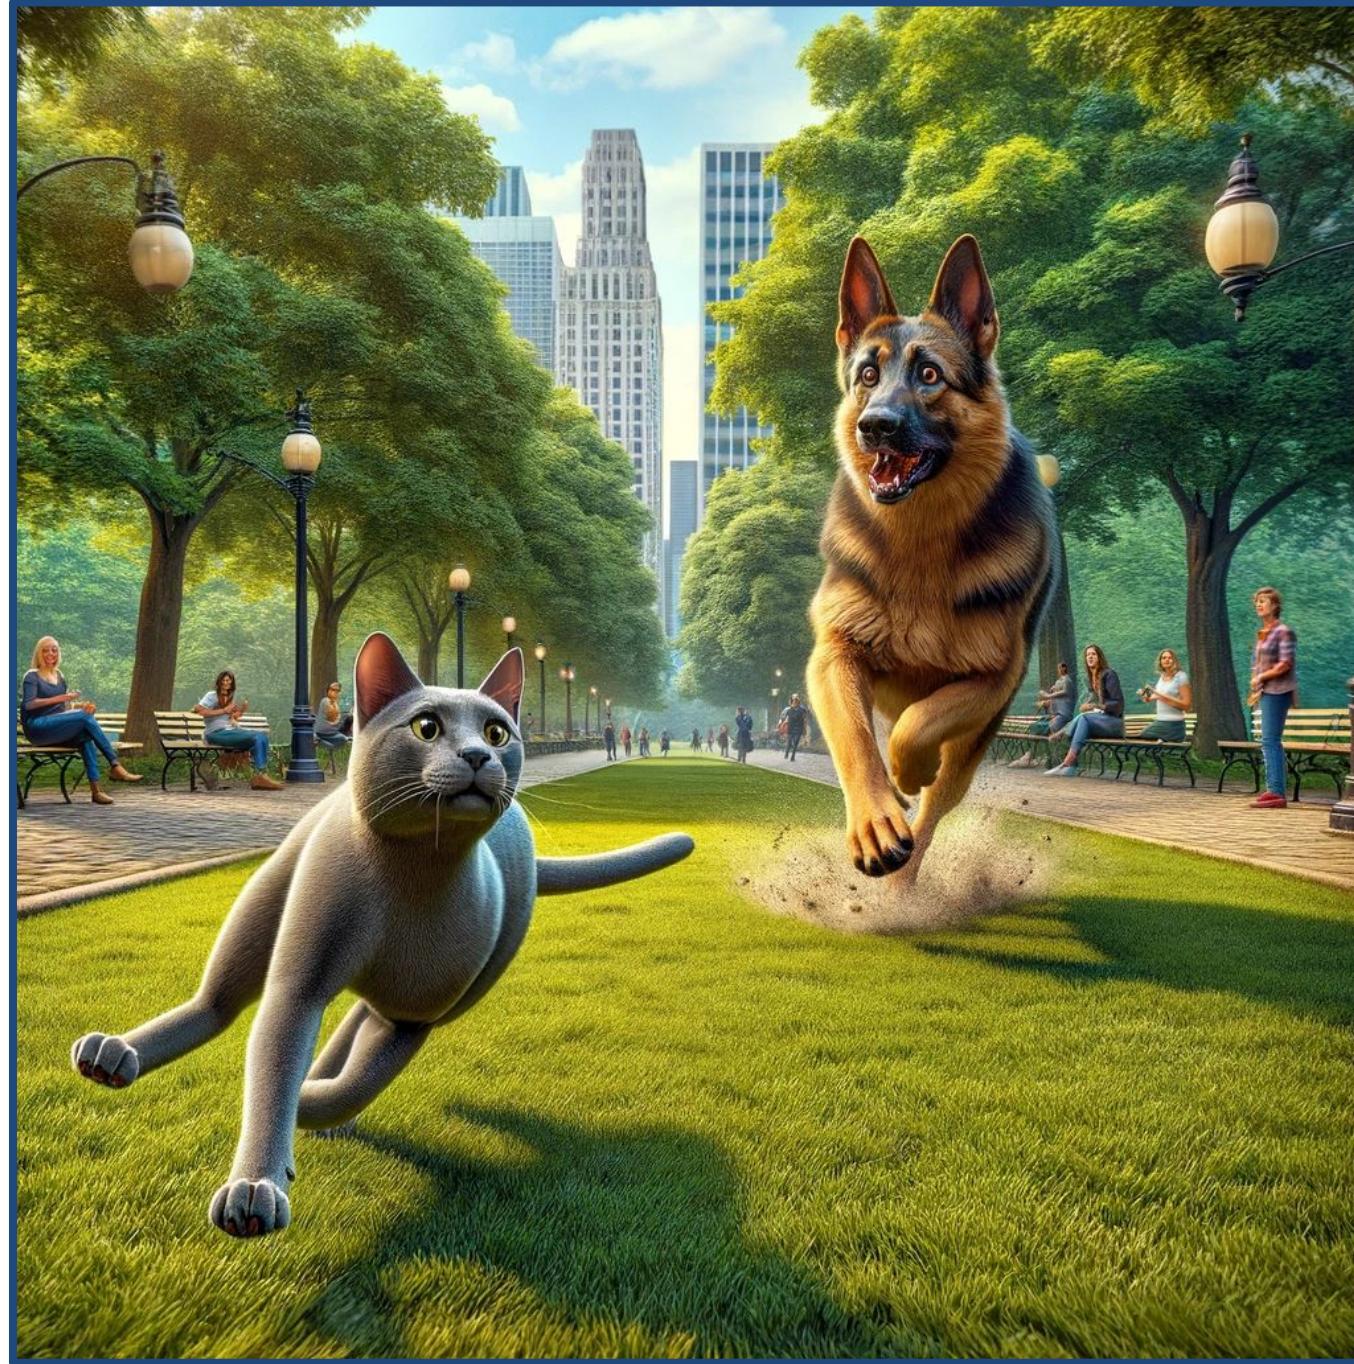
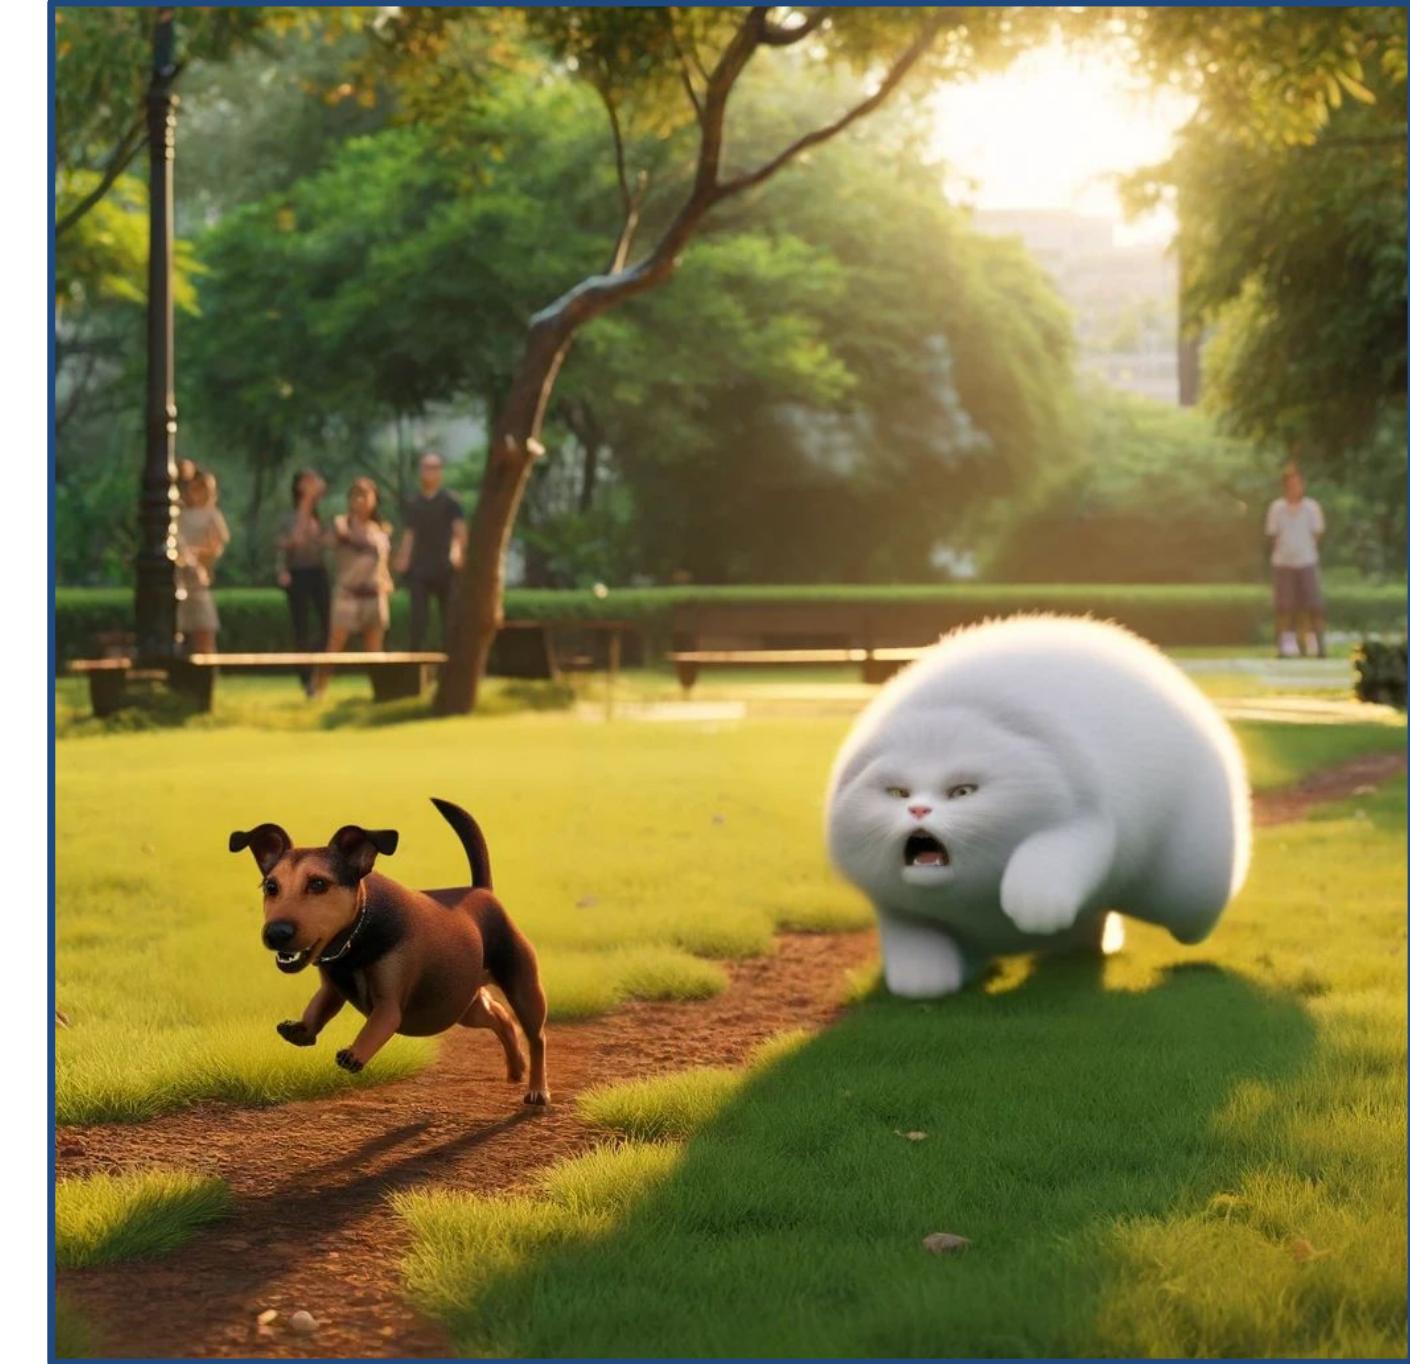




## Can Multimodal LLMs distinguish between



a dog chasing a cat



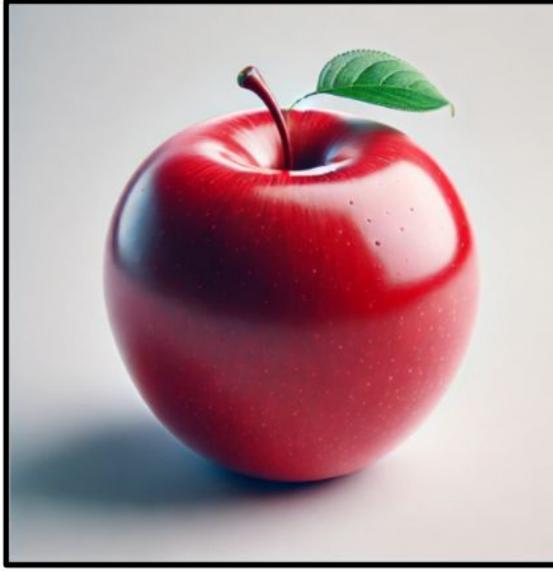
a cat chasing a dog

On Winoground, a benchmark of such questions,  
MLLM (GPT-4V) performance = 33%  
Human performance = 85%

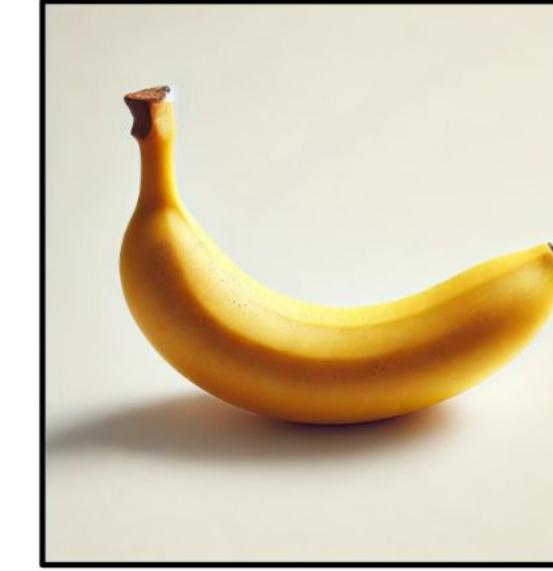
## Compositionality:

Identifying parts/concepts and how they are composed to give rise to a given scene

### Atomic Concepts

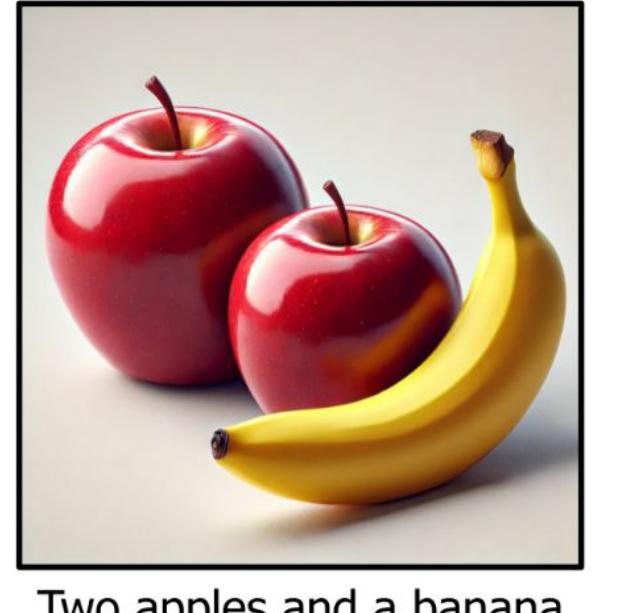


Apple



Banana

### Infinite Compositions - Nothing New



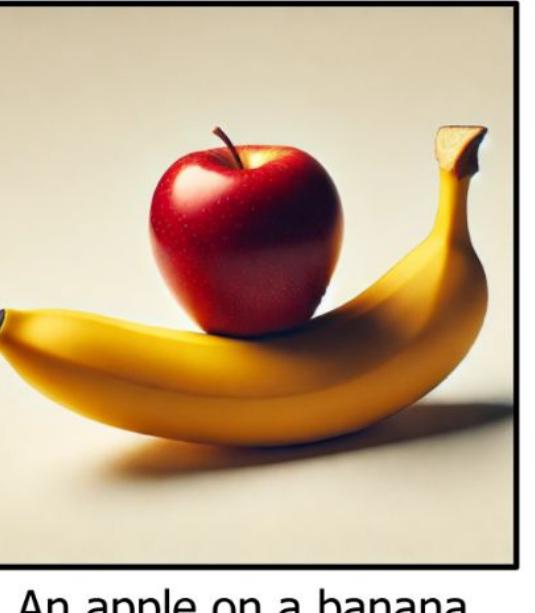
Two apples and a banana



An apple and two bananas



A banana on an apple

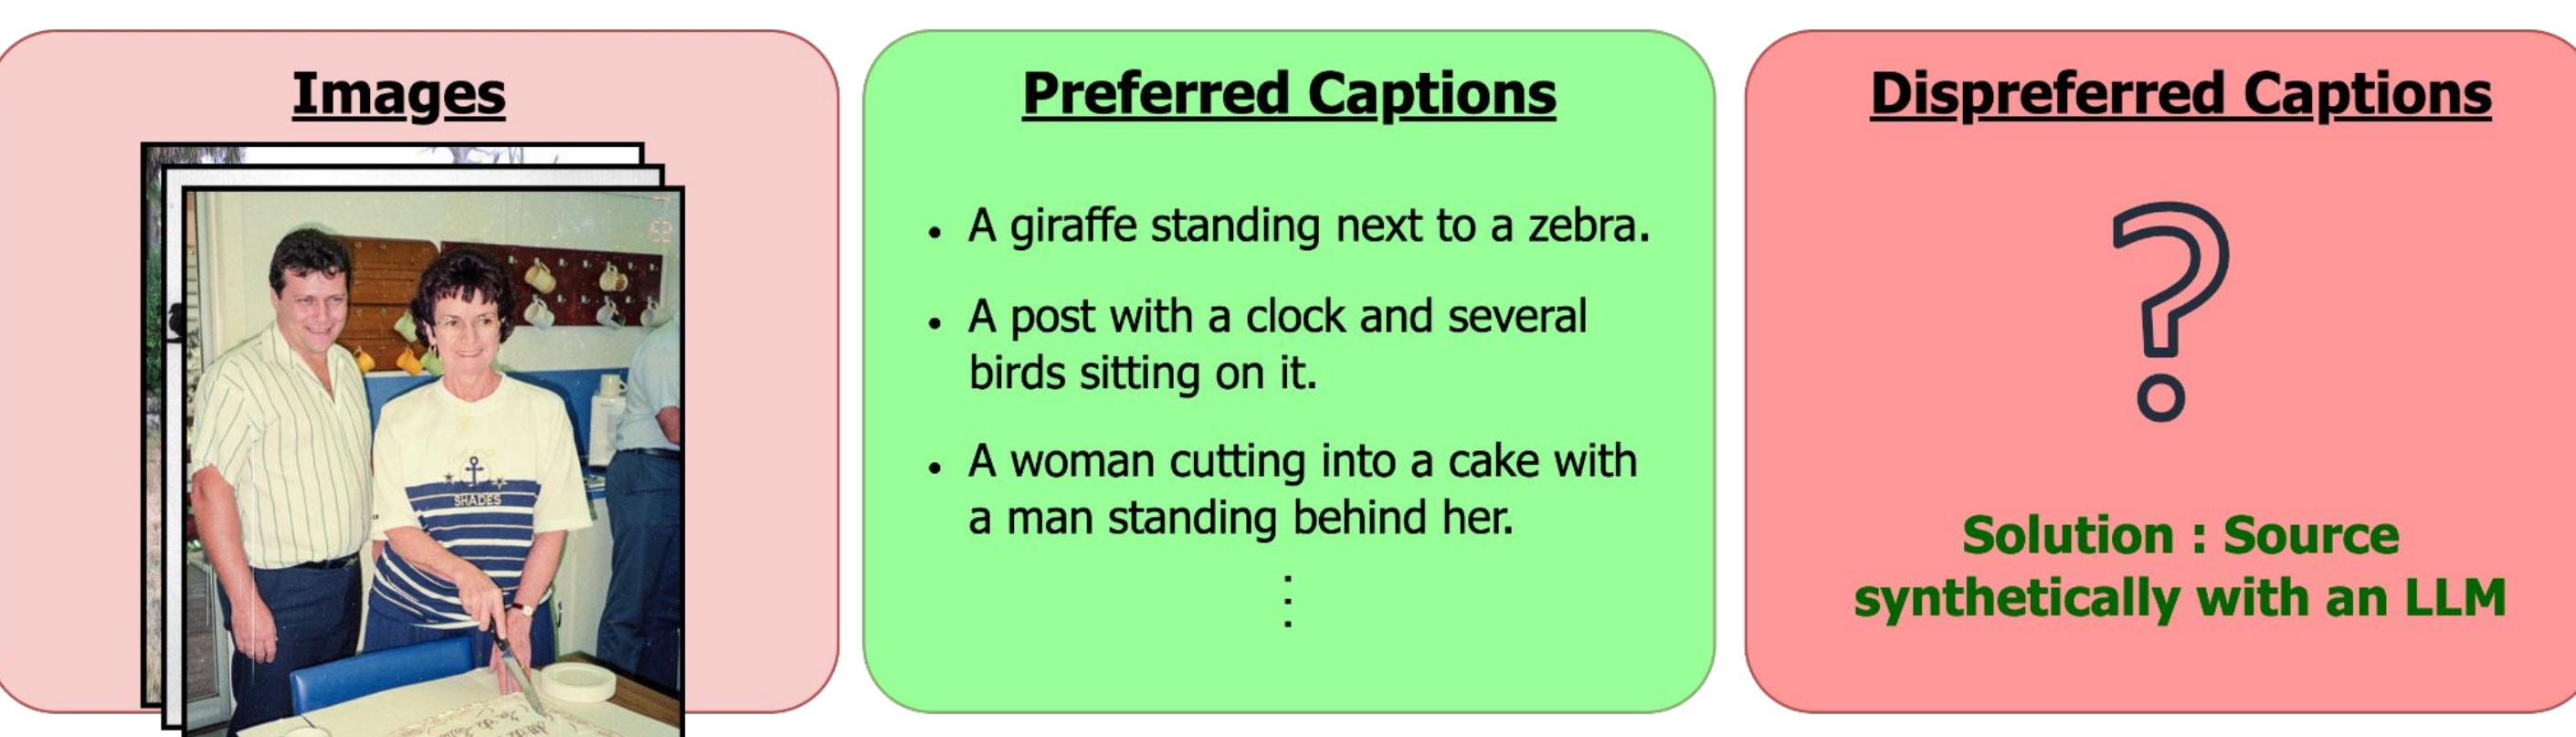


An apple on a banana

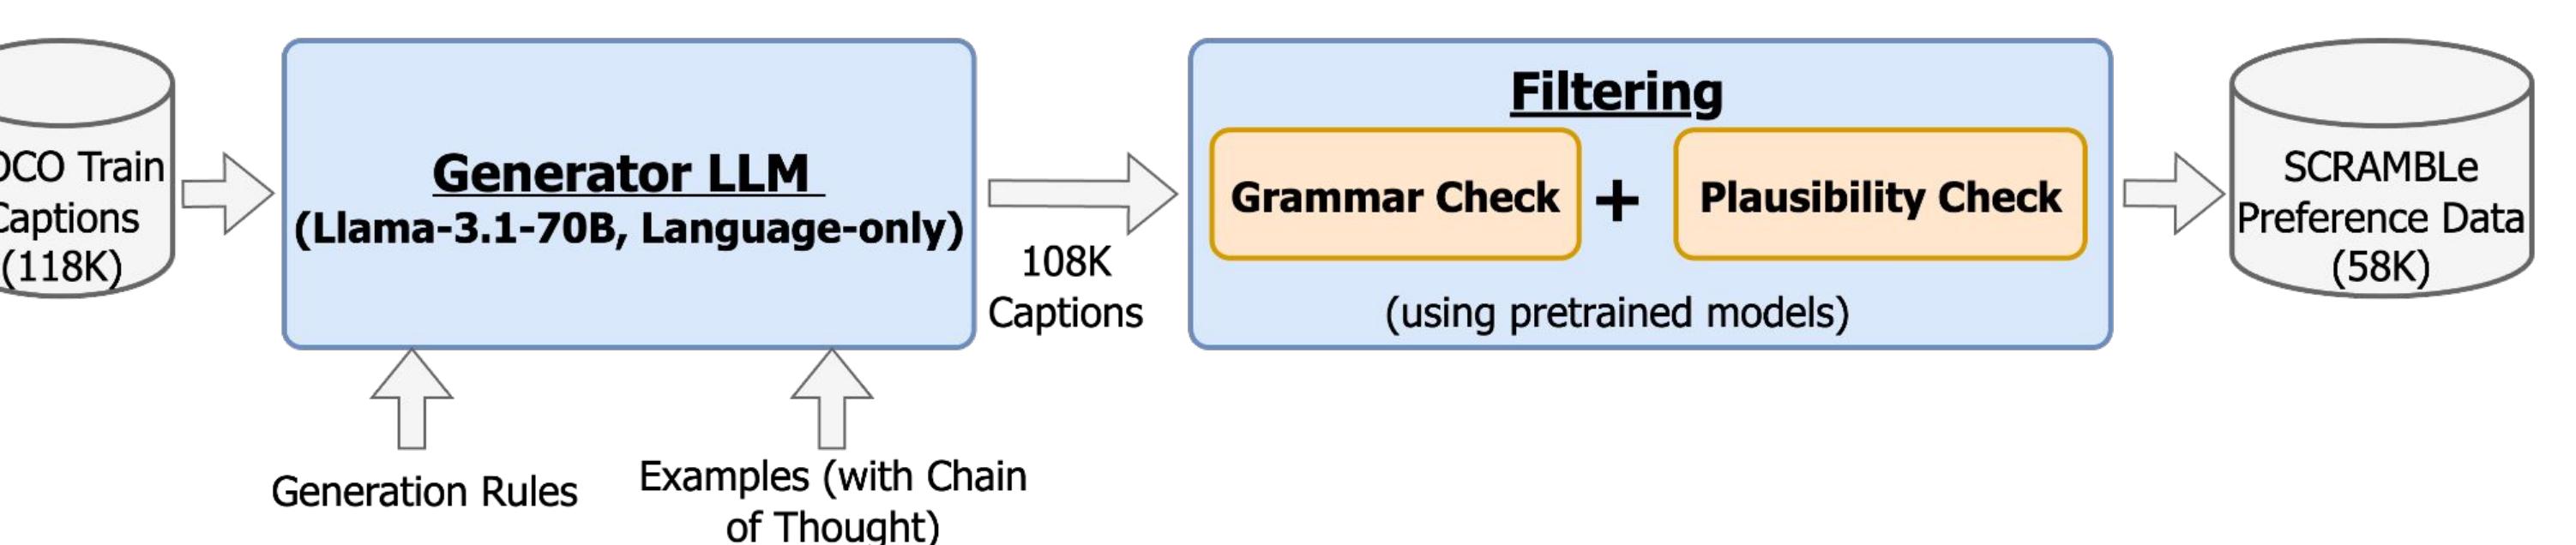
## Improving Compositionality with Preference Tuning



## Data Source : Existing Image Caption Dataset (COCO)



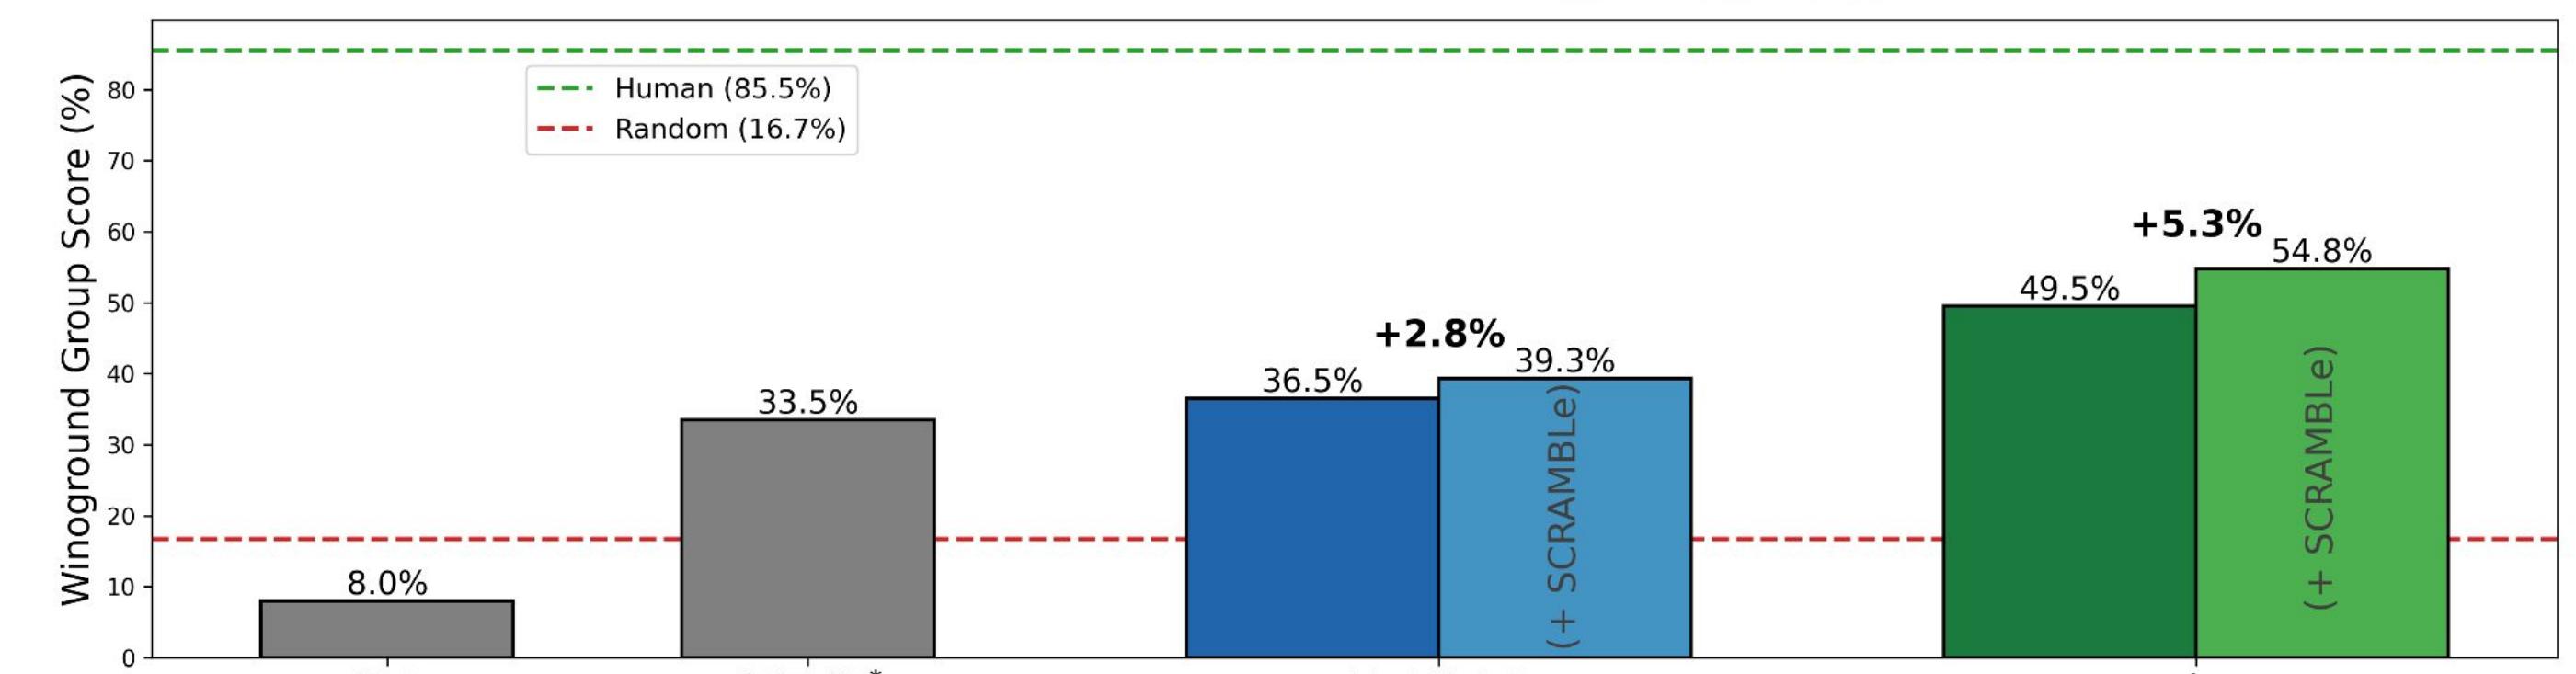
## Synthetic Data Generation



## References

[1] Mitra, Chanchik, et al. "Compositional chain-of-thought prompting for large multimodal models." Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2024.  
 [2] Lin, Zhiqiu, et al. "Evaluating text-to-visual generation with image-to-text generation." European Conference on Computer Vision. Cham: Springer Nature Switzerland, 2024.

## Winoground Accuracy



\*GPT-4V eval based on verbose responses and LLM as a judge [1]. In contrast LLaVA and Molmo use VQAScore [2] based on model output logprobs.

## Hmm, could this be overfitting?

Seemingly not. Models retain or improve general QA performance (on Control Benchmarks)

Model Name	Compositionality Benchmarks			Control Benchmarks	
	Winoground	COLA	ConMe	SEED-Bench	MM-Vet
LLaVA-1.5-13B	36.5	49.5	62.3	<b>68.23</b>	$36.2 \pm 0.3$
+ SCRAMBLe	<b>39.3</b>	<b>55.7</b>	<b>64.5</b>	68.19	<b>38.6 \pm 0.1</b>
MoLMO-7B	49.5	57.1	72.2	74.04	$59.3 \pm 0.2$
+ SCRAMBLe	<b>54.8</b>	<b>60.5</b>	<b>74.6</b>	<b>74.61</b>	<b>60.9 \pm 0.4</b>

## Chats with SCRAMBLe-Molmo

