

Towards Models that Can See and Read

Roy Ganz*
Technion, Israel

ganz@cs.technion.ac.il

Yair Kittenplon
AWS AI Labs

yairk@amazon.com

Oren Nuriel
AWS AI Labs

onurielamazon.com

Shai Mazor
AWS AI Labs

smazor@amazon.com

Aviad Aberdam
AWS AI Labs

aaberdam@amazon.com

Ron Litman
AWS AI Labs

litmanr@amazon.com

Abstract

Visual Question Answering (VQA) and Image Captioning (CAP), which are among the most popular vision-language tasks, have analogous scene-text versions that require reasoning from the text in the image. Despite their obvious resemblance, the two are treated independently and, as we show, yield task-specific methods that can either see or read, but not both. In this work, we conduct an in-depth analysis of this phenomenon and propose UniTNT, a Unified Text-Non-Text approach, which grants existing multimodal architectures scene-text understanding capabilities. Specifically, we treat scene-text information as an additional modality, fusing it with any pretrained encoder-decoder-based architecture via designated modules. Thorough experiments reveal that UniTNT leads to the first single model that successfully handles both task types. Moreover, we show that scene-text understanding capabilities can boost vision-language models' performance on general VQA and CAP by up to 2.69% and 0.6 CIDEr, respectively.

1. Introduction

In recent years, Vision-Language (VL) tasks, such as Visual Question Answering (VQA) [4, 18] and Image Captioning (CAP) [34, 2], have gained immense research interest [55, 39, 49, 30, 22, 46, 10, 47]. However, despite the remarkable success of VL models on these tasks, it was discovered a few years ago that such models are incapable of reasoning from the text in natural images [41, 8, 40]. This finding raised significant concerns, as understanding scene-text is crucial in almost any real-world application.

To address this issue, designated scene-text datasets were introduced for both VQA [41, 8] and CAP [40], aiming to highlight the importance of utilizing textual information in images. Following the introduction of the above datasets, a new line of research has arisen, focusing on scene-text-oriented tasks, evaluated individually and effectively dis-

*Work done during an Amazon internship.

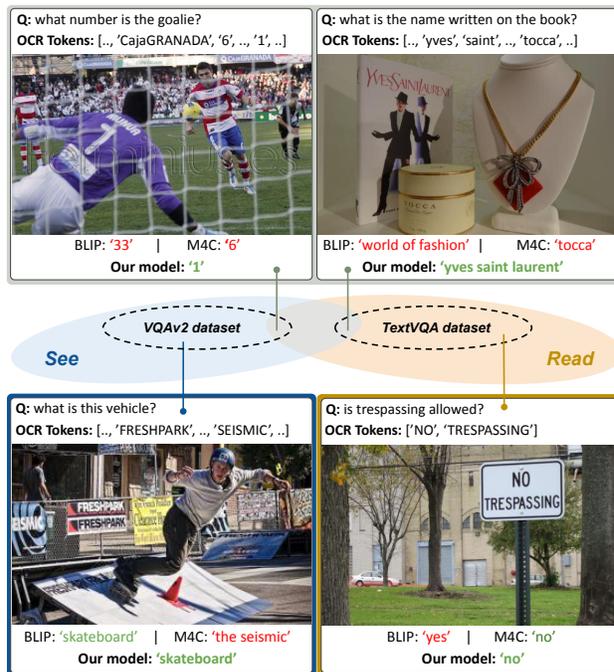


Figure 1: **See and read in VQA.** Illustration of the possible three types of reasoning required in VQA image-question pairs and representative datasets distributions (middle). Samples from the 'see' (bottom left), 'read' (bottom right), and 'see-read' (top) subsets are presented. Each sample includes an image, question, OCR, and model predictions.

sociated from the general one. From a user perspective, this separation is artificial and does not adequately reflect the objective of real-world VQA systems, and as we show, it encourages models to only excel on one task at a time. Therefore, we advocate that VL research should strive towards unified models, and thus, methods should be evaluated accordingly. To this end, we propose conducting *combined evaluation* for VL models on both general and scene-text benchmarks and treating the average results as expressing the "see" and "read" capabilities. We emphasize

that even the minority of works that evaluate both types of tasks [46, 11, 3] do it on separate models, which are fine-tuned per task, perpetuating the faulty tasks’ segregation.

Apart from being unjustified, this separation introduces biases [7, 48], providing the models with prior knowledge that implies which modality to focus on, which does not exist in real-world scenarios. Namely, it creates a shortcut that encourages models to excel solely on a specific benchmark by acquiring an understanding of either the visual or textual information in the image, but not both. In particular, Biten *et al.* [7] recently showed that SOTA performance on scene-text VQA can be achieved without using the visual modality, and Wang *et al.* [48] revealed that existing scene-text VQA models’ success stems from exploiting language priors. Our *combined evaluation* effectively addresses this problem by testing whether models can reason from both types of information, as exploiting such data biases and priors would yield low combined results.

From a more high-level view, three categories span the space of VL data; the first are examples that require reasoning over vision only (dominant in VQA [18] and CAP [9]), the second are instances in which using scene-text information solely is sufficient (dominant in scene-text VQA [8, 41] and scene-text CAP [40]), and the third are ones in which both are essential. We denote the three subsets as ‘see’, ‘read’, and ‘see-∩-read’, respectively. For completion, the whole space is denoted as ‘see-∪-read’, the union of all others. We illustrate this conceptual data distribution for VQA in Fig. 1. Examining the performance of existing VQA approaches over the three types of questions mentioned above, shown in Fig. 2, reveals that while some of the methods [30, 31, 47] perform well on the first subset and some [21, 52] on the second, none are optimal on the entire domain. Moreover, throughout our analysis, we reveal that the ‘see-∩-read’ subset, in which both visual and textual information are needed for answering, is very challenging and underrepresented, requiring a new dedicated benchmark.

In this work, while striving towards models that excel on the entire space of VL data, we propose UniTNT, a Unified Text-Non-Text model, which provides VL architectures with scene-text understanding capabilities. Specifically, we treat textual information in the image, *i.e.* tokens and positions, as a third modality and introduce it into the pretrained model. Adding a new modality to an already-trained model is challenging and might lead to suboptimal results [16, 43, 54, 3]. To overcome this, we encode such information using a designated encoder and inject it into the existing pretrained decoder via a novel fusing mechanism that gradually shifts between VL features to textual-enriched ones. Moreover, we propose scene-text-related intermediate supervision to encourage the already-trained model to leverage the newly added information. Being both task and model agnostic by its design, our method can be

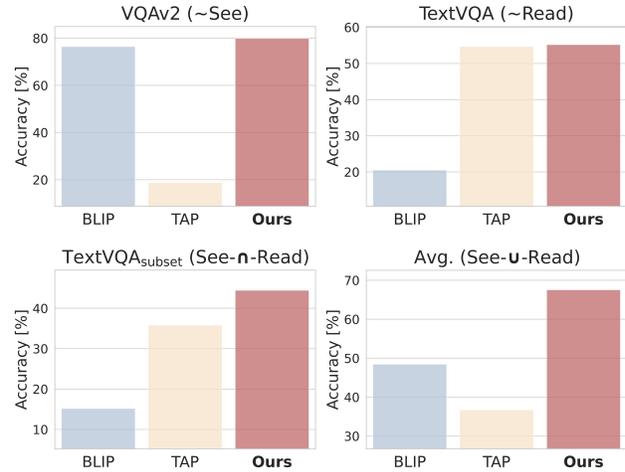


Figure 2: **Models’ accuracy on different types of VQA data.** Leading methods and UniTNT performance on different benchmarks. VQAv2 and TextVQA datasets mostly require reasoning from visual information only (‘See’) and textual information only (‘Read’), respectively. ‘See-∩-Read’ refers to a subset of the TextVQA dataset (Sec. 4.3), in which both modalities are essential for answering each question. ‘See-∪-Read’ represents the sets’ union.

applied to any VL encoder-decoder-based architecture.

We evaluate UniTNT on both general and scene-text benchmarks of VQA and CAP using *combined evaluation* and show that it leads to the first single model performing well on both tasks. We show that our method can be easily integrated into existing VL models, improving their scene-text understanding substantially by applying it to BLIP [30] and ALBEF [31]. Interestingly, such reasoning abilities boost the base model’s VQA results (*e.g.*, improves BLIP [30] by 2.69% on VQAv2 [18]), while achieving state-of-the-art competitive results on scene-text VQA benchmarks. A similar trend exists in captioning, where UniTNT enhances BLIP’s performance by 0.6 CIDEr points on COCO Captions [9] while substantially boosting its scene-text CAP performance. These improvements highlight the significance of scene-text comprehension in VL tasks, laying the foundation for future research on general multimodal architectures that can leverage scene-text.

To summarize:

- We thoroughly analyze current methods and reveal that the faulty text-non-text task separation leads to models that either reason from visual or textual information in images, but not both.
- We introduce UniTNT, a model-agnostic method to grant reading capabilities to pretrained VL models by fusing the scene-text information as an additional modality.
- Extensive experiments show that our method not only improves the scene-text benchmarks’ results but also significantly enhances the performance of VQA and CAP.

	Method	OCR System	Visual Question Answering					Image Captioning					
			VQA		TextVQA		ST-VQA	Avg.	COCO		TextCaps		Avg.
			test-dev	test-std	val	test	test-ANLS		Karpathy-test	val	test		
Separate	M4C [21, 40]	✓	27.47	27.70	46.53	47.42	0.43	37.56	4.7	95.5	90.1	47.4	
	TAP [52]	✓	18.76	18.81	54.71	53.97	0.60	36.39	4.6	109.2	103.2	53.9	
	ALBEF [31]	✗	75.22	75.38	11.67	13.88	0.19	44.63	-	-	-	-	
	BLIP [30]	✗	76.39	76.59	20.50	23.74	0.34	50.16	133.3	59.4	61.9	97.6	
	OFA _{Large} [47]	✗	79.70	79.85	22.10	21.47	0.27	50.66	150.7	64.5	66.8	108.8	
Comb.	M4C [21]	✓	59.11	59.04	47.22	48.61	0.50	53.83	109.8	102.7	98.0	103.9	
	ALBEF [31]	✗	75.61	75.87	16.15	17.04	0.22	46.46	-	-	-	-	
	BLIP [30]	✗	77.40	77.39	32.43	31.48	0.44	54.44	133.4	101.4	91.8	112.6	

Table 1: **Current status of VQA and CAP models.** The results of leading methods on both scene-text and general VQA and CAP benchmarks reveal that currently, no method performs well on both scene-text and general benchmarks, even when applying combined training. *Separate* and *Comb.* summarize the results described in Sections 2.1, 2.2 and 2.3, respectively.

2. See and Read: Analyzing Methods and Data

In this paper, contrary to the common practice in VL research, we highlight the importance of models to “see” and “read” altogether and start by comprehensively analyzing such capability via a “see- \cup -read”-oriented *combined evaluation*. Our analysis reveals that existing models’ reasoning abilities over both types of information are lacking, prompting the question of whether this limitation is due to inherent method constraints or biased data. Our evaluation focuses on the performance of leading general and scene-text-oriented models on VQAv2 [18], TextVQA [41], and ST-VQA [8] for VQA, and COCO Captions [9] and TextCaps [40] for captioning.

2.1. Visual Question Answering

General VQA Methods: During the vision-language revolution, numerous methods [33, 39, 31, 30, 47, 49, 46, 3, 10, 55, 53, 13, 22] have been proposed for various multimodal tasks, including VQA, which have advanced the state-of-the-art. These methods can leverage vast online image-caption pairs via vision-language pretraining [32, 12, 39], followed by task-specific fine-tuning. However, a few years ago, such models were shown to be ineffective in reasoning from textual information in the scene, as they primarily focus on the images’ visual content [41, 8].

Nevertheless, such models have advanced significantly in the past few years. Thus, to reveal the current status of such models in scene-text understanding, we examine the performance of three leading VQA models, ALBEF [31], BLIP [30], and OFA [47], using unconstrained open-vocabulary generation, on scene-text VQA tasks. As seen in Tab. 1, although such methods perform well on VQA, as expected, their results on the analogous scene-text VQA datasets are unsatisfactory, testifying their incompetence in scene-text understanding. Interestingly, their inability to utilize scene-text information hinders its performance even on VQA, as we later show in Sec. 4.

Scene-Text VQA Methods: Several methods have been

proposed to improve the scene-text understanding of VQA models [21, 19, 17, 52, 23, 36, 7]. These models utilize an off-the-shelf OCR system’s output alongside the image and question as input to a multimodal transformer. However, some recent studies [48, 7] have indicated that scene-text VQA datasets may have biases discouraging models from relying on the visual modality. To properly test such claims, we evaluate M4C [21] and TAP [52] on general VQA, which requires strong visual understanding and report the results in Tab. 1. As can be seen, M4C and TAP obtain only 27.70% and 18.81%, respectively. When compared to, for example, BLIP’s 76.59%, it testifies that, indeed, such methods disregard the visual information. Interestingly, although TAP consistently outperforms M4C on the scene-text benchmarks, it achieves lower results on the general one, implying the data biases in the former datasets.

2.2. Image Captioning

Similar to VQA’s analysis, we conduct a captioning *combined evaluation* using TextCaps and COCO Captions for both types of models and report the average CIDEr scores. Our empirical results in Tab. 1 demonstrate that while general models (BLIP and OFA) and scene-text ones (M4C-Captioner and TAP) perform well on their designated benchmarks, they fail to obtain satisfactory results on the analogous one. In particular, BLIP obtains a CIDEr score of only 61.9 on TextCaps, compared to 90.1 of M4C-Captioner. On the other hand, the latter achieves 4.7 on COCO captions, compared to BLIP’s 133.3. In addition, like in VQA, while TAP outperforms M4C in TextCaps, it does not occur on COCO Captions. These findings suggest that existing methods exhibit unsatisfactory performance when evaluated on both captioning benchmarks.

2.3. The Role of the Datasets in ‘See and Read’

We now examine whether this limitation stems from a lack of representative training data rather than method limitations. Specifically, we test if the inferior performance of

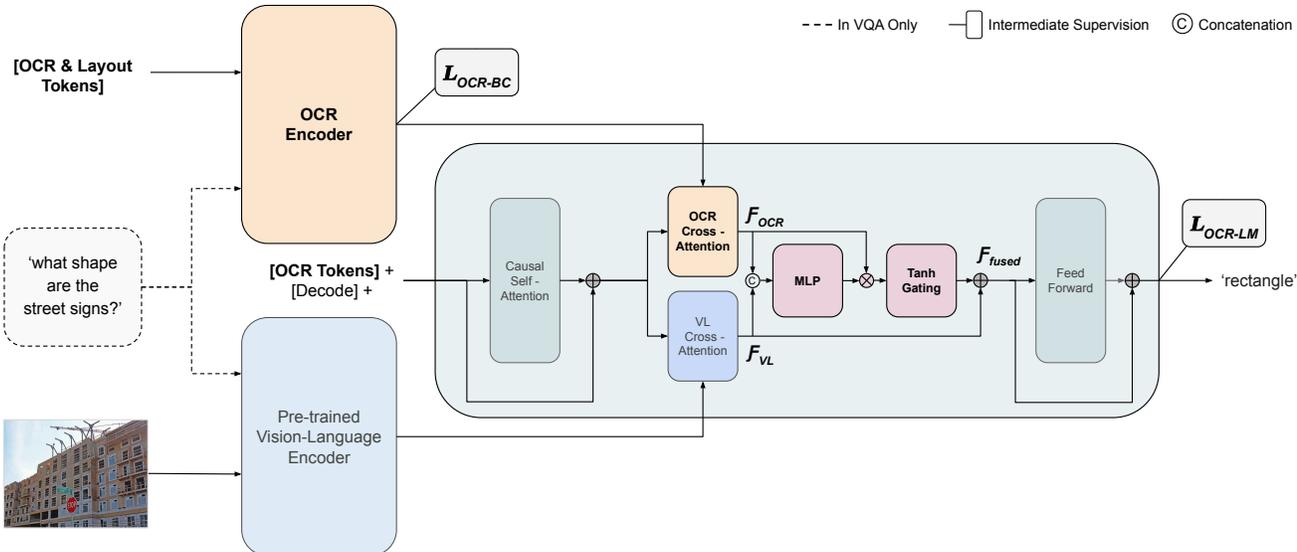


Figure 3: **An overview of UniTNT.** Our method endows existing general VL models with scene-text understanding capability. The OCR information is encoded separately and injected into the decoder via a gated cross-attention-based fusing mechanism as complementary information. \mathcal{L}_{OCR-BC} and \mathcal{L}_{OCR-LM} are auxiliary losses, enforcing the model to utilize the scene-text information. UniTNT newly introduced components are presented in bold. ‘See’, ‘Read’, and ‘Fusing’ related modules are in blue, orange, and red, respectively.

scene-text-oriented models on visual tasks and vice versa is solely due to the training data’s bias towards reasoning over solely one type of information. To test this claim, we merge two datasets, conduct combined training for both general and scene-text-oriented methods, and report the results in Tab. 1. As can be seen, while unified training leads to improved performance on both types of VL benchmarks, there is a substantial performance gap – scene-text models lag behind general ones on the general benchmarks and vice-versa. Nevertheless, these results indicate that reasoning from text and vision are not at odds and suggest a symbiotic relationship between the two tasks. Furthermore, they provide further motivation for avoiding the common practice of separating the tasks, as done in previous work [3, 11, 46]. To conclude, while joint training is a step forward, it is not enough to achieve our ultimate goal.

3. Method

In this section, we describe UniTNT, a method aimed to obtain our titular goal by granting pretrained general VL models the ability to reason over scene-text information during finetuning while retaining their original reasoning capabilities, depicted in Fig. 3. By doing so, we propose a change of perspective compared to top-performing ST methods, such as [7, 52], that harnesses an OCR-oriented pretrained model but fails to enrich it with visual understanding during finetuning. Adapting pretrained models to consider additional inputs, absent during pretraining, is a non-trivial task tackled by recent literature [43, 3]. On the

one hand, we wish to encourage the model to utilize the new stream of information and, on the other hand, to prevent it from neglecting the original stream. To address this, we encode the OCR information via a designated OCR encoder and fuse it residually, retaining the former stream of information and gradually shifting towards an OCR-enriched one. Moreover, we propose auxiliary losses, encouraging the pretrained decoder to utilize this information. Similarly to previous works [21, 52, 36, 7], we utilize an off-the-shelf OCR system to extract the scene-text information.

3.1. Architecture

We design our architecture in a task-agnostic way – enabling compatibility with both visual question answering and image captioning tasks. In addition, UniTNT is model agnostic and can be applied to any encoder-decoder-based VL model. In this work, we integrate our approach into two top-performing open-source methods – ALBEF [31], and BLIP [30] as a case study, denoted as UniTNT_{ALBEF/BLIP}.

OCR Encoder Rather than utilizing the pre-existing encoder to process the OCR alongside the visual modality, as in [7, 21, 52], we introduce a dedicated OCR encoder, which maps the scene-text information into features fed into the existing system’s decoder. This encoder receives the question alongside OCR information, namely tokens and 2-dimensional (2D) positional information, both extracted by the OCR system. The positional information was proven to be valuable for documents and scene-text understanding

tasks [50, 51, 5, 7]. Not only that our approach outperforms the one that utilizes the pre-existing text encoder to process the OCR tokens (demonstrated in Sec. 5), but it also provides flexibility to address tasks that do not utilize a text encoder, such as image captioning.

Formally, each OCR instance is represented by $(t, x_0, y_0, x_1, y_1, w, h)$, namely, its word token, bounding box’s top-left, bottom-right, width, and height values, respectively. We embed each value separately using designated embedding layers E (*i.e.*, `torch.nn.Embedding`). Next, we sum the 2D representations, pass them via a 2-layer MLP and add it to the token’s representation, yielding the OCR representation:

$$\mathbf{e}_{\text{OCR}} = E_{\text{OCR}}(t) + \alpha * \text{MLP}(E_x(x_0) + E_y(y_0) + E_x(x_1) + E_y(y_1) + E_w(w) + E_h(h)) \quad (1)$$

where α is a predefined hyperparameter. As for the question, we embed its tokens using the same embedding layer. Since both the OCR and the question representations are fed into the same model, we equip the question representations with pseudo-2D information corresponding to the size of the entire image, yielding the final question representation \mathbf{e}_q . Finally, we concatenate them to obtain the OCR encoder’s input, $\{\mathbf{e}_q^1 \dots \mathbf{e}_q^M, \mathbf{e}_{\text{OCR}}^1 \dots \mathbf{e}_{\text{OCR}}^N\}$, where M and N are the lengths of the question and OCR, respectively.

VL-OCR Decoder To integrate the OCR information into the decoder, we add a dedicated OCR Cross Attention (CA) and a fusing mechanism, as visualized in Fig. 3. We place the OCR CA block parallel to the pre-existing VL CA module to enrich the decoded features with textual information in the image. This architectural design yields two data streams (visual and scene-text-oriented ones) that need to be merged adequately into a single VL-OCR representation. To this end, we introduce a fusing mechanism composed of a gated cross-attention mechanism, which gradually shifts from VL features to fused, OCR-enriched ones.

Formally, our fusing mechanism merges the output of our new OCR CA with the one of the VL CA, denoted as \mathcal{F}_{OCR} and \mathcal{F}_{VL} respectively. Specifically, this module receives two features sequences, $\mathcal{F}_{\text{OCR}}, \mathcal{F}_{\text{VL}} \in \mathbb{R}^{B \times L \times C}$, and outputs $\mathcal{F}_{\text{fused}} \in \mathbb{R}^{B \times L \times C}$, where B, L, C are the batch size, sequence length and the number of channels, respectively. First, we concatenate \mathcal{F}_{OCR} and \mathcal{F}_{VL} across the channel dimension and insert them into a simple 2-layer MLP to obtain an attention map $\mathcal{F}_{\text{attn}} \in \mathbb{R}^{B \times L \times C}$. Next, we pass the element-wise product of \mathcal{F}_{OCR} and $\mathcal{F}_{\text{attn}}$ in a tanh gating mechanism [20, 3]. The goal of the tanh gating is to enable gradual OCR blending with the VL one by multiplying its inputs with $\tanh(\beta)$, where β is a learnable parameter initialized to zero. At initialization, it ensures that the added modules are skipped, preserving the model pre-training’s data flow. Finally, we sum the output of the tanh

gating with \mathcal{F}_{VL} to obtain the fused features:

$$\mathcal{F}_{\text{attn}} = \text{MLP}(\text{concat}(\mathcal{F}_{\text{VL}}, \mathcal{F}_{\text{OCR}})), \quad (2)$$

$$\mathcal{F}_{\text{fused}} = \mathcal{F}_{\text{VL}} + \tanh(\beta)(\mathcal{F}_{\text{OCR}} \odot \mathcal{F}_{\text{attn}}), \quad (3)$$

where \odot is the Hadamard product.

3.2. Scene-text Auxiliary Losses

We propose two auxiliary losses, encouraging the model to utilize the scene-text signal rather than ignoring it - OCR Causal Language Modeling (OCR-LM) and OCR Binary Classification (OCR-BC).

OCR Causal Language Modeling To better fuse the scene-text information, we add a causal language modeling supervision over the OCR tokens. Specifically, we prepend the shifted OCR tokens (according to the OCR system reading order) to the inputs of the decoder and train the system to predict the next OCR token based on previous ones,

$$\mathcal{L}_{\text{OCR-LM}} = - \sum_{i=1}^N \log(\mathbb{P}(t^i | t^{<i})) \quad (4)$$

where t^i is the i^{th} OCR token. Minimizing such loss enforces the system to account for the scene-text signal, as desired. While variants of such a loss were previously used during pretraining [52, 7], we are the first to utilize it during finetuning. Moreover, inserting the OCR into the decoder at inference has another significant advantage, as it serves as a prefix and enables the model to condition its answers on the OCR. Such behavior is desirable since the OCR can provide meaningful information for general and scene-text VL tasks, as we experimentally demonstrate in Sec. 4.1.

OCR Binary Classification To obtain more meaningful and task-beneficial OCR encodings, we propose a binary classification objective of predicting whether each OCR token is a part of the ground-truth answer. We build a binary linear classifier on top of the outputs of the OCR encoder and train it using a binary cross-entropy loss. More specifically, since most of the OCR tokens are not part of the answer, we employ a weighted version, as such classification task is highly imbalanced. We denote this loss as $\mathcal{L}_{\text{OCR-BC}}$.

3.3. Training Procedure

So far, we have described the main building blocks in our method, and now, as illustrated in Fig. 3, we put it all together. First, we harness a trained general encoder-decoder VL model and modify it as described above in Sec. 3.1. Next, we freeze the VL model’s pre-existing image encoder, similarly to [54, 3], and train UniTNT on a unified dataset (*i.e.*, general and scene-text VQA datasets

	Method	OCR System	VQA		TextVQA		ST-VQA	Avg.
			test-dev	test-std	val	test	test-ANLS	
VQA	SimVLM _{large} [49]	✗	79.32	79.56	-	-	-	-
	GIT _{large} ^{VQA} [46]	✗	75.51	-	-	-	-	-
	ALBEF [31]	✗	75.22	75.38	11.67	13.88	0.19	44.63
	OFA _{large} [47]	✗	79.70	79.85	22.10	21.47	0.27	50.66
	mPLUG _{VTE-B} [29]	✗	79.79	79.81	-	-	-	-
	BLIP [30]	✗	76.39	76.59	20.50	23.74	0.34	50.17
TextVQA	UniTNT _{BLIP}	✓	79.68	79.78	36.33	35.90	0.50	57.84
	Δ		↑3.28	↑3.19	↑15.83	↑12.16	↑0.16	↑7.67
	GIT _{large} ^{TextVQA} [46]	✗	-	-	37.47	-	-	-
	SA-M4C [23]	✓	-	-	45.4	44.6	0.50	-
	LOGOS [36]	✓	-	-	51.53	51.08	0.58	-
	M4C [21]	✓	27.47	27.70	46.53	47.42	0.43	37.56
Comb.	TAP [52]	✓	18.76	18.81	54.71	53.97	0.60	36.39
	LaTr [7]	✓	-	-	59.53	59.55	0.68	-
	BLIP [30]	✗	40.16	40.39	30.12	27.72	0.36	34.06
	UniTNT _{BLIP}	✓	37.01	37.24	50.19	47.39	0.59	42.32
	Δ		↓3.15	↓3.15	↑20.07	↑19.67	↑0.23	↑8.26
	M4C [21]	✓	59.11	59.04	47.22	48.61	0.50	53.83
Comb.	ALBEF [31]	✗	75.61	75.87	16.15	17.04	0.22	46.46
	UniTNT _{ALBEF}	✓	77.60	77.80	43.73	44.13	0.58	60.97
	Δ		↑1.99	↑1.93	↑27.58	↑27.09	↑0.36	↑14.51
	BLIP [30]	✗	77.40	77.39	32.43	31.48	0.44	54.44
	UniTNT _{BLIP}	✓	79.90	80.08	55.21	55.35	0.66	67.72
	Δ		↑2.50	↑2.69	↑22.77	↑23.87	↑0.22	↑13.28

Table 2: **VQA results.** Accuracy of general, scene-text oriented VQA methods and UniTNT using three training regimes – separate VQA and TextVQA and combined training, where non-open vocabulary methods results are in gray. Δ indicates improvement over the base architecture in the same regime. These results highlight our method’s effectiveness, significantly improving the general VQA results by enriching VL models with scene-text understanding.

or general and scene-text captioning datasets). Specifically, $\mathcal{L}_{\text{UniTNT}} = \mathcal{L}_{\text{base}} + \alpha_1 \mathcal{L}_{\text{OCR-LM}} + \alpha_2 \mathcal{L}_{\text{OCR-BC}}$ is minimized, where $\mathcal{L}_{\text{base}}$ is the base task-dependent loss term used in our base architecture, and α_1, α_2 are tunable hyperparameters.

4. Experiments

In this section, we experimentally examine UniTNT, comparing its performance with state-of-the-art methods with a similar capacity on both VQA and CAP tasks, using separate and combined training. In particular, to better study the effects of our method, we test it and the baselines in three distinct training regimes; (i) separate training on the general datasets, (ii) separate training on the scene-text ones, and (iii) combined training approach, denoted as Comb. As we focus on models’ see and read capabilities, we emphasize the combined training regime and view it as the most crucial one. However, the separate training regimes can provide insights into the impact of scene-text understanding on the general benchmarks and the biases within the scene-text datasets. As in Sec. 2, for each of the regimes, we consider three standard benchmarks for VQA: VQAv2 [18], TextVQA [41] and ST-VQA [8], and two for CAP: COCO Captions [9] and TextCaps [40]. We report the performance on each benchmark and the non-weighted averaged one (*combined evaluation*) to quantify the models’ reasoning capabilities from both visual and textual information as a single number. For VQA, we calculate this

	Method	OCR System	COCO		TextCaps		Avg.
			Karpathy-test	val	val	test	
Caps	VinVL [55]	✗	129.3	-	-	-	-
	LEMON _{base} [22]	✗	133.3	-	-	-	-
	GIT _{large} ^{Cap} [46]	✗	138.5	-	-	-	-
	SimVLM _{large} [49]	✗	142.6	-	-	-	-
	OFA _{large} [47]	✗	150.7	64.5	66.8	108.8	-
	BLIP [30]	✗	133.3	59.4	61.9	97.6	-
TextCaps	UniTNT _{BLIP}	✓	133.7	59.6	62.8	98.3	-
	Δ		↑0.4	↑0.2	↑0.9	↑0.7	-
	GIT _{large} ^{TextCap} [46]	✗	-	106.3	-	-	-
	MMA-SR [45]	✓	-	98.0	88.0	-	-
	CNMT [42]	✓	-	-	93.0	-	-
	M4C-Captioner [40]	✓	4.7	95.5	90.1	47.4	-
Comb.	TAP [52]	✓	4.6	109.2	103.2	53.9	-
	BLIP [30]	✗	84.8	112.7	103.7	94.3	-
	UniTNT _{BLIP}	✓	70.4	130.5	123.1	96.8	-
	Δ		↓14.4	↑17.8	↑19.4	↑2.5	-
	M4C-Captioner [40]	✓	109.8	102.7	98.0	103.9	-
	BLIP [30]	✗	133.4	101.4	91.8	112.6	-
Comb.	UniTNT _{BLIP}	✓	134.0	119.1	109.4	121.7	-
	Δ		↑0.6	↑17.7	↑17.6	↑9.1	-

Table 3: **CAP results.** CIDEr scores of general, scene-text oriented CAP methods and UniTNT using three training regimes – separate Caps and TextCaps and combined training. Δ indicates improvement over the base architecture in the same regime. These results highlight our method’s effectiveness, significantly improving the general CAP results by enriching VL models with scene-text understanding.

score only on VQAv2 and TextVQA test sets. Lastly, in Sec. 4.3, we present a new subset evaluation setting for scene-text VQA to measure the model’s ability to answer questions requiring reasoning over all modalities simultaneously. For all datasets, we extract OCR information using Amazon Text-in-Image¹ [35, 38, 1, 26]. The supplementary materials list the implementation details, additional dataset information and for completeness, a comparison with other methods, disregarding the models’ size.

4.1. Visual Question Answering Experiments

We integrate our approach to two models, ALBEF and BLIP, denoted as UniTNT_{ALBEF} and UniTNT_{BLIP}, respectively, and report their performance using three training regimes: (i) VQA, (ii) TextVQA, and (iii) Comb., as shown in Tab. 2. In the first regime, training UniTNT_{BLIP} exclusively on VQAv2 results in performance improvements of +3.19% and +12.16% on VQA, and TextVQA, respectively, leading to a significant boost of +7.67% in the average score. Even though VQAv2 mainly focuses on reasoning from visual information, these results stress the importance of scene-text understanding in this benchmark and the effectiveness of our method. Interestingly, despite the marginal presence of OCR in VQAv2, UniTNT_{BLIP} manages to effectively harness it and obtain 35.90% on

¹<https://docs.aws.amazon.com/rekognition/latest/dg/text-detection.html>

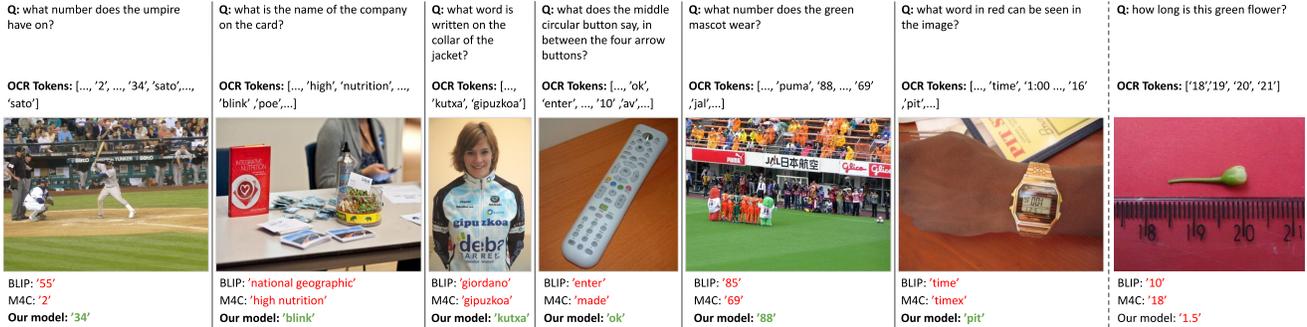


Figure 4: **Reasoning over all modalities.** We curate a subset out of TextVQA [41] validation set, containing only the samples which require reasoning over both vision and scene-text in the same question. Presented are representative examples from this subset, each includes an image, question, OCR input tokens, and model predictions. Green and red stand for correct and wrong predictions, respectively.

Method	TextVQA	TextVQA _{Read}	TextVQA _{See\capRead}	Gap \downarrow
M4C [21]	46.53	47.94	35.69	12.25
TAP [52]	54.71	56.24	35.83	20.41
UniTNT _{BLIP}	55.21	56.32	44.44	11.88

Table 4: **TextVQA splits.** Accuracy of leading scene-text VQA methods on the two non-overlapping subsets of TextVQA validation data, and the gap between them. 'See \cap Read' refers to our subset, in which reasoning over all modalities is needed for each sample. 'Read' stands for the rest of the TextVQA validation set.

TextVQA, outperforming BLIP that trained solely on TextVQA itself (27.72%). In the scene-text configuration, performance improves by +19.67% on TextVQA; however, it decreases by -3.15% on VQA. This reinforces previous findings [48, 7], suggesting that scene-text VQA datasets contain biases encouraging models to over-rely on the OCR and disregard the visual information. As BLIP's scene-text understanding is very restricted, it cannot fully exploit such biases and retains its visual understanding better, expressed via better VQAv2 results. In the final combined training configuration, we showcase the versatility of our approach by presenting results for both UniTNT_{ALBEF} and UniTNT_{BLIP}, highlighting its model-agnostic nature. When trained on both types of datasets, UniTNT_{BLIP} improves BLIP by +2.5%, +22.77%, and +13.28% on VQA, TextVQA, and on average, respectively, achieving the highest average score. The results indicate that despite the biases in scene-text VQA datasets, UniTNT can harness them without sacrificing the visual reasoning capability. Moreover, UniTNT_{BLIP} trained on the combined dataset outperforms models trained on each task separately, attesting to the tasks' mutually beneficial relationship, motivating the community to strive towards models that can see and read.

To gain a better understanding of the enhancements achieved by UniTNT across both "see" and "read" datasets, we conducted a qualitative analysis of our method, com-

paring it to BLIP and M4C in Fig. 1, Fig. 4 and in the supplementary materials. Our analysis indicates that the improvements observed in VQA are due to questions that necessitate reading and those that become easier to answer with OCR information. Regarding scene-text VQA, M4C struggles to reason from visual information, while UniTNT excels in this regard, resulting in significant performance improvements on both TextVQA and ST-VQA.

4.2. Image Captioning Experiments

Similar to our VQA experiments, we evaluate the performance of UniTNT on CAP by comparing it to top-performing methods using the same three training regimes (Caps, TextCaps, and Comb.). We only integrate our approach to BLIP (UniTNT_{BLIP}), as ALBEF was not applied to captioning. As the results in Tab. 3 indicate, in the Caps regime, our approach slightly improves both the per-task and the average results. Like in VQA, in the TextCaps training regime, UniTNT results in significant gains over TextCaps (+19.4 CIDEr points) but a decline in COCO (-14.4 CIDEr points), compared to BLIP. Moreover, while combined training leads to the best COCO results, the best TextCaps performance is achieved via designated TextCaps finetuning. This phenomenon aligns with the earlier findings by [40], attributing it to the different nature of ground truth captions in the scene-text and general benchmarks (additional analysis appears in the supplementary materials). Nevertheless, the combined trained UniTNT leads to the best average score across all methods and regimes.

4.3. A subset for Reasoning Over All Modalities

As illustrated in Fig. 1, VQA data is composed of three categories. Some questions can be answered using just vision ('see'), some by reasoning over the scene-text information only ('read'), and some require reasoning over both modalities at once ('see \cap read'). Since most of the questions in current benchmarks fall either in the 'see' or

OCR-Sys	OCR-Enc	Fuse	\mathcal{L}_{OCR-LM}	\mathcal{L}_{OCR-BC}	2-D	VQA		CAP	
						VQA _{v2}	TextVQA	COCO	TextCap
✗	✗	✗	✗	✗	✗	77.40	32.43	133.4	101.4
✓	✗	✗	✗	✗	✗	77.66	43.02	133.5	109.7
✓	✓	✗	✗	✗	✗	78.41	46.13	133.5	110.4
✓	✓	✓	✗	✗	✗	78.65	47.38	133.5	118.3
✓	✓	✓	✓	✗	✗	79.86	52.66	133.8	118.9
✓	✓	✓	✓	✓	✗	79.81	52.66	-	-
✓	✓	✓	✓	✓	✗	79.74	52.93	134.0	119.1
✓	✓	✓	✓	✓	✓	79.90	55.21	-	-

Table 5: **UniTNT design choices.** UniTNT_{BLIP} results on VQA and CAP w.r.t its different building blocks.

‘read’ category, unifying them is beneficial for testing methods’ performance on the whole space, denoted by ‘see- \cup -read’, eliminating the model’s prior on whether a question is of type ‘see’ or ‘read’. However, the more challenging and intriguing questions are the ones that require reasoning over scene-text and visual information altogether, denoted as ‘see- \cap -read’. To provide a more reliable way to evaluate VQA models on this questions’ category, we manually curate all such image-question pairs from the TextVQA [41] validation set, producing an evaluation subset of 480 image-question pairs out of the total 5000 ($\pm 10\%$).

This subset can serve as a foundation for measuring models’ capabilities on what we believe are the more challenging questions that the research community should tackle. In Fig. 4, we depict examples from this subset alongside the prediction of M4C [21], BLIP [30], and UniTNT. This qualitative analysis confirms that both scene-text and general VQA models struggle to cope with this type of questions, while UniTNT is substantially better. Moreover, in Tab. 4 we report the quantitative results of leading scene-text-oriented methods and UniTNT on the non-overlapping subsets of TextVQA validation set, *i.e.*, the ‘TextVQA_{See \cap Read}’ subset and its complementary set, ‘TextVQA_{Read}’, exposing the performance degradation that occurred on the former. As these findings suggest, our method leads to the best performance, affirming that it is indeed better at reasoning on scene-text and visual information simultaneously. Nevertheless, as can be seen, while UniTNT is a step forward, there is still a big room to improve on these types of challenging questions.

5. Ablation Studies

In this section, we study the effect of our key contributions and test the impact of freezing the vision encoder.

Design Choices: We ablate UniTNT’s components on both the general and scene-text-oriented datasets in Tab. 5, where all numbers are reported under the *Comb.* settings. Since the trends in CAP results are similar to VQA, we will focus on analyzing the latter. First, we report the added performance of a naive approach – simply inserting the OCR tokens as an additional input to BLIP’s existing text encoder, similar to [21, 52, 7]. As seen in Tab. 5, the accu-

Method	Freeze VE	VQA test-dev	TextVQA val	Avg.
UniTNT _{ALBEF}	✗	75.55	40.60	58.08
UniTNT _{ALBEF}	✓	77.60	43.73	60.67
UniTNT _{BLIP}	✗	78.77	52.45	65.61
UniTNT _{BLIP}	✓	79.90	55.21	67.56

Table 6: **Visual encoder freezing.** VQA accuracy of UniTNT_{BLIP} and UniTNT_{ALBEF}, with and without freezing the visual encoder, attesting to the freezing’s importance.

racy on TextVQA improves by +10.59% (from 32.43% to 43.02%) while improving VQA results by +0.22%. Our designated OCR encoder increases TextVQA performance to 46.13% (+3.11%) while obtaining an additional +0.75% gain in VQA. Introducing our VL-OCR decoding scheme (denoted as ‘Fuse’) boosts us to 47.37% on TextVQA and an extra +0.24% on VQA. Furthermore, using \mathcal{L}_{OCR-LM} significantly improves TextVQA performance by +5.22% (from 47.38% to 52.66%) while gaining an extra +1.21% on VQA. Finally, the combination of \mathcal{L}_{OCR-BC} with the 2-D information gets us to 55.21% and 79.9% on TextVQA and VQA. Overall, UniTNT leads to significant +22.78% and +2.50% improvements on TextVQA and VQA over the combined trained BLIP.

The Effect of Freezing the Visual Encoder: Recently, a few works [43, 3, 54] have examined different freezing configurations to avoid knowledge forgetness when combining pretrained models. Inspired by these works, we examine the effect of freezing the Visual Encoder (VE) weights while applying UniTNT, preserving its valuable knowledge acquired in pretraining, and summarize the results in Tab. 6. As our findings suggest, freezing the VE significantly improves the results on VQA for both UniTNT_{ALBEF} and UniTNT_{BLIP} by +2.05% and +1.13%, and on TextVQA by +3.13% and +2.75%, respectively.

6. Discussion and Conclusions

We wish to convey a few take-home messages to the VL research community. First, current SOTA methods cannot adequately reason over both scene-text and vision information. Our experiments demonstrate that this occurs even when combining training datasets, suggesting a fundamental limitation of existing methods. Second, our findings discover the symbiotic nature of these two types of reasoning capabilities, as performance on both tasks can be improved jointly. Moreover, by proposing UniTNT, we present the first single model that successfully handles both task types. Finally, we argue that the VL research community should strive to develop models that can simultaneously reason over vision, language, and scene-text. To facilitate this, we curate a suitable subset to serve as a benchmark foundation.

References

- [1] Aviad Aberdam, Roy Ganz, Shai Mazor, and Ron Litman. Multimodal semi-supervised learning for text recognition. *arXiv preprint arXiv:2205.03873*, 2022. [6](#)
- [2] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8948–8957, 2019. [1](#)
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022. [2](#), [3](#), [4](#), [5](#), [8](#)
- [4] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015. [1](#), [12](#)
- [5] Srikanth Appalaraju, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R Manmatha. Docformer: End-to-end transformer for document understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 993–1003, 2021. [5](#)
- [6] Jeffrey P Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, et al. Vizwiz: nearly real-time answers to visual questions. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, pages 333–342, 2010. [13](#)
- [7] Ali Furkan Biten, Ron Litman, Yusheng Xie, Srikanth Appalaraju, and R Manmatha. Latr: Layout-aware transformer for scene-text vqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16548–16558, 2022. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [8] Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluís Gomez, Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. Scene text visual question answering. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4291–4301, 2019. [1](#), [2](#), [3](#), [6](#), [12](#)
- [9] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. [2](#), [3](#), [6](#), [12](#)
- [10] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022. [1](#), [3](#)
- [11] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022. [2](#), [4](#)
- [12] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: universal image-text representation learning. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXX*, volume 12375 of *Lecture Notes in Computer Science*, pages 104–120. Springer, 2020. [3](#)
- [13] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022. [3](#)
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [13](#)
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019. [12](#)
- [16] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999. [2](#)
- [17] Chenyu Gao, Qi Zhu, Peng Wang, Hui Li, Yuliang Liu, Anton Van den Hengel, and Qi Wu. Structured multimodal attentions for textvqa. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [3](#)
- [18] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. [1](#), [2](#), [3](#), [6](#)
- [19] Wei Han, Hantao Huang, and Tao Han. Finding the evidence: Localization-aware answer prediction for text visual question answering. *arXiv preprint arXiv:2010.02582*, 2020. [3](#)
- [20] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 1997. [5](#)
- [21] Ronghang Hu, Amanpreet Singh, Trevor Darrell, and Marcus Rohrbach. Iterative answer prediction with pointer-augmented multimodal transformers for textvqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9992–10002, 2020. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [22] Xiaowei Hu, Zhe Gan, Jianfeng Wang, Zhengyuan Yang, Zicheng Liu, Yumao Lu, and Lijuan Wang. Scaling up vision-language pre-training for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17980–17989, 2022. [1](#), [3](#), [6](#)
- [23] Yash Kant, Dhruv Batra, Peter Anderson, Alexander Schwing, Devi Parikh, Jiasen Lu, and Harsh Agrawal. Spatially aware multimodal transformers for textvqa. In *European Conference on Computer Vision*, pages 715–732. Springer, 2020. [3](#), [6](#)

- [24] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. Icdar 2015 competition on robust reading. In *2015 13th international conference on document analysis and recognition (ICDAR)*, pages 1156–1160. IEEE, 2015. [13](#)
- [25] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluís Pere De Las Heras. Icdar 2013 robust reading competition. In *2013 12th international conference on document analysis and recognition*, pages 1484–1493. IEEE, 2013. [13](#)
- [26] Yair Kittenplon, Inbal Lavi, Sharon Fogel, Yarin Bar, R. Manmatha, and Pietro Perona. Towards weakly-supervised text spotting using a multi-task transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4604–4613, June 2022. [6](#)
- [27] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. [13](#)
- [28] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4. *International Journal of Computer Vision*, 128(7):1956–1981, 2020. [13](#)
- [29] Chenliang Li, Haiyang Xu, Junfeng Tian, Wei Wang, Ming Yan, Bin Bi, Jiabo Ye, Hehong Chen, Guohai Xu, Zheng Cao, et al. mplug: Effective and efficient vision-language learning by cross-modal skip-connections. *arXiv preprint arXiv:2205.12005*, 2022. [6](#)
- [30] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086*, 2022. [1](#), [2](#), [3](#), [4](#), [6](#), [8](#), [12](#)
- [31] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. [2](#), [3](#), [4](#), [6](#), [12](#)
- [32] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 9694–9705, 2021. [3](#)
- [33] Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, and Haifeng Wang. Unimo: Towards unified-modal understanding and generation via cross-modal contrastive learning. *arXiv preprint arXiv:2012.15409*, 2020. [3](#)
- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. [1](#), [13](#)
- [35] Ron Litman, Oron Anshel, Shahar Tsiper, Roei Litman, Shai Mazor, and R Manmatha. Scatter: selective context attentional scene text recognizer. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11962–11972, 2020. [6](#)
- [36] Xiaopeng Lu, Zhen Fan, Yansen Wang, Jean Oh, and Carolyn P Rosé. Localize, group, and select: Boosting text-vqa by scene text modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2631–2639, 2021. [3](#), [4](#), [6](#)
- [37] A. Mishra, K. Alahari, and C. V. Jawahar. Image retrieval using textual cues. In *ICCV*, 2013. [13](#)
- [38] Oren Nuriel, Sharon Fogel, and Ron Litman. Textadain: Fine-grained adain for robust text recognition. *arXiv preprint arXiv:2105.03906*, 2021. [6](#)
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [1](#), [3](#)
- [40] Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. Textcaps: a dataset for image captioning with reading comprehension. In *European conference on computer vision*, pages 742–758. Springer, 2020. [1](#), [2](#), [3](#), [6](#), [7](#), [12](#), [13](#)
- [41] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#), [12](#)
- [42] Anonymous submission. Cnmt, textcaps challenge 2020 top entry. <https://evalai.cloudcv.org/web/challenges/challenge-page/573/overview>, 2020. [6](#)
- [43] Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *Advances in Neural Information Processing Systems*, 34:200–212, 2021. [2](#), [4](#), [8](#)
- [44] Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*, 2016. [13](#)
- [45] Jing Wang, Jinhui Tang, and Jiebo Luo. Multimodal attention with image text spatial relationship for ocr-based image captioning. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4337–4345, 2020. [6](#)
- [46] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022. [1](#), [2](#), [3](#), [4](#), [6](#)

- [47] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. OFA: unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 23318–23340. PMLR, 2022. [1](#), [2](#), [3](#), [6](#)
- [48] Qingqing Wang, Liqiang Xiao, Yue Lu, Yaohui Jin, and Hao He. Towards reasoning ability in scene text visual question answering. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 2281–2289, 2021. [2](#), [3](#), [7](#)
- [49] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*, 2021. [1](#), [3](#), [6](#)
- [50] Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. Layoutlm: Pre-training of text and layout for document image understanding. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1192–1200, 2020. [5](#)
- [51] Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, et al. Layoutlmv2: Multi-modal pre-training for visually-rich document understanding. *arXiv preprint arXiv:2012.14740*, 2020. [5](#)
- [52] Zhengyuan Yang, Yijuan Lu, Jianfeng Wang, Xi Yin, Dinei Florencio, Lijuan Wang, Cha Zhang, Lei Zhang, and Jiebo Luo. Tap: Text-aware pre-training for text-vqa and text-caption. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8751–8761, 2021. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [53] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. [3](#)
- [54] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18123–18133, 2022. [2](#), [5](#), [8](#)
- [55] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Making visual representations matter in vision-language models. 2021. [1](#), [3](#), [6](#)

A. Implementation Details

In this section, we provide full implementation specifics of UniTNT and divide it into three parts – (1) architecture; (2) training procedure; (3) Scene-text information.

A.1. Architecture

We harness the model agnosticism of UniTNT and apply it to two top-performing VL models. Specifically, we utilize the publicly-available code bases of ALBEF [31]² and BLIP³ [30] and apply our method to them. We design our approach in a modular way enabling simple integration into existing models. Below we list the architectural specifics for both UniTNT_{ALBEF} and UniTNT_{BLIP}.

OCR Encoder We use a pretrained BERT-base⁴ [15] as our encoder and introduce it with 2-dimensional information, as can be seen in Equation 1. Specifically, we use three separate embedding layers (*i.e.*, `torch.nn.Embedding`)– for the word token and its x and y axis positions for both the OCR and the question. In particular, we define the minimal and the maximal spatial position as 0 and 1000, respectively, and set these values for the question tokens (referred to as “pseudo-2D information” in the main paper). We restrict the number of OCR and question token lengths to 128 and 35, respectively. Next, we sum the 2D-related embeddings and pass them in a 2-layer MLP with a hidden dimension of 768 for additional processing. Finally, we multiply it by α (set to 0.1) and sum it with the token representation to obtain the final one fed into the encoder.

VL-OCR Decoder In order to introduce the pretrained decoder with scene-text information, we create new OCR Cross Attention (OCR-CA) blocks and place them in parallel to the existing VL ones. Such newly added components are identical to the existing ones and initialized with the pretrained weights of the latters⁷. To fuse the outputs of the OCR CA and the VL CA, \mathcal{F}_{OCR} and \mathcal{F}_{VL} , we concatenate them along the channel dimension and pass them via attention based 2-layers MLP with a hidden size of 768 to obtain \mathcal{F}_{attn} , an attention map that multiplies \mathcal{F}_{OCR} ($\mathcal{F}_{OCR} \odot \mathcal{F}_{attn}$). Namely, this mechanism highlights the important and meaningful features in \mathcal{F}_{OCR} and masks the less relevant ones. Then, we pass the multiplication output via a learnable gating module (by multiplying it by $\tanh(\beta)$, where β is learnable and initialized to 0), aimed to gradually blend the OCR features into the existing VL one.

²<https://github.com/salesforce/ALBEF>

³<https://github.com/salesforce/BLIP>

⁴https://huggingface.co/docs/transformers/model_doc/bert

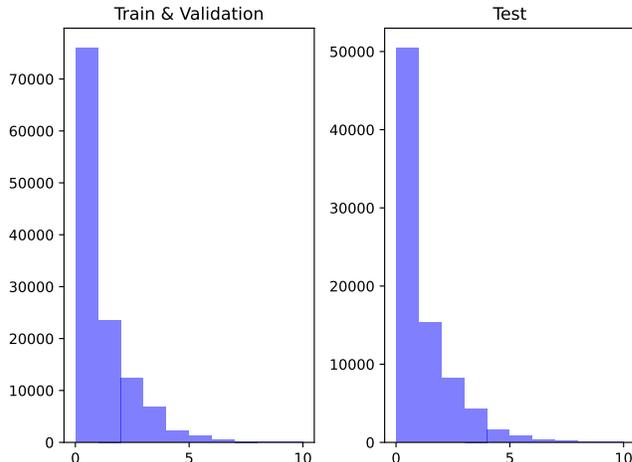


Figure 5: **OCR prevalence in VQAv2.** Histogram of the number of OCR instances per-image in VQAv2 dataset.

A.2. Training Procedure

We train all of our models to minimize $\mathcal{L}_{\text{UniTNT}} = \mathcal{L}_{\text{base}} + \alpha_1 \mathcal{L}_{\text{OCR-LM}} + \alpha_2 \mathcal{L}_{\text{OCR-BC}}$ using 8 A100 GPUs, where α_1 and α_2 are hyperparameters.

Visual Question Answering We train both UniTNT_{ALBEF} and UniTNT_{BLIP} on a unified Text-Non-Text VQA dataset, containing VQAv2 [4], TextVQA [41] and ST-VQA [8] for 10 epochs using a batch size of 8 and 16 for ALBEF and BLIP, respectively. Moreover, we set $\alpha_1 = \alpha_2 = 1$ and keep the other training-related hyperparameters as in the original papers.

Image Captioning We train UniTNT_{BLIP} on a the unified Text-Non-Text CAP dataset, comprised of COCO Captions [9] and TextCaps [40], for 5 epochs with batch size of 32. We set $\alpha_1 = 0.05$ and $\alpha_2 = 0$ since contrary to VQA, CAP does not contain textual information available both in training and inference time, making it infeasible to implement OCR-BC. Moreover, we keep the rest of the hyperparameters as in BLIP.

A.3. Scene-text information

As specified in the paper, we extract the scene-text information (word tokens and 2-dimensional position) for all the VQA and CAP datasets (both the general and scene-text counterparts) using Amazon Text-in-Image. To better understand the prevalence of OCR in the non-scene-text datasets, we plot the statistics of OCR in VQAv2 in Fig. 5 (same images are in COCO Captions as well). While a large portion of the images does not contain text in them, there is a large amount of such with OCR (38.36% and 38.03% of

Method	Vision-oriented dataset	OCR-oriented dataset	VQA test-dev	TextVQA val	Avg.	COCO Caps val	TextCaps val	Avg.
BLIP	✗	✓	40.16	30.12	35.14	84.8	112.7	98.8
UniTNT _{BLIP}			37.01	50.19	43.60	70.4	130.5	100.5
BLIP	✓	✗	76.39	20.50	48.45	133.3	59.4	96.4
UniTNT _{BLIP}			79.68	36.33	58.01	133.7	59.6	96.7
BLIP	✓	✓	77.40	32.43	54.92	133.4	101.4	117.4
UniTNT _{BLIP}			79.90	55.21	67.56	134.0	119.1	126.6

Table 7: **The impact of training data.** We show the effect of each dataset configuration for training UniTNT and BLIP.

train and test images contain OCR). Since OCR conveys meaningful information, it sheds light on the significant improvement of UniTNT up his baselines (ALBEF and BLIP).

B. Datasets

B.1. Visual Question Answering

VQA_{v2} contains 204,721 images (82,783, 40,504, and 81,434) from COCO [34], 1,105,904 questions (443,757, 214,354, and 447,793), and 6,581,110 answers (4,437,570, 2,143,540, and the test answers are held-out). Answering the questions requires vision-language understanding and commonsense knowledge. Each question has ten ground-truth answers.

TextVQA contains 28,408 images from OpenImages [28], 45,336 questions and 453,360 ground-truth answers. The annotators were instructed to formulate questions that require reasoning from the text in the image. As in VQA_{v2}, each question has 10 ground-truth answers.

ST-VQA is a fusion of computer-vision datasets – ImageNet [14], VizWiz [6], Visual Genome [27], IIIT Scene Text Retrieval [37], ICDAR 2013 [25], ICDAR 2015 [24] and COCO Text [44]. It contains 31K questions, split into training (26K) and testing (5K), requiring scene-text understanding.

B.2. Image Captioning

COCO Captions contains over one and a half million captions describing over 330,000 images from the COCO dataset. Each image has five human-generated captions.

TextCaps is composed of 28,408 images and 142,040 captions (5 captions per image). The images are from the TextVQA dataset, and the captions are based on the text in the image. Specifically, models have to reason over the scene-text information to generate correct captions.

C. The Impact of Training Data

In this section, we study the effect of the different combinations of training datasets and report our findings in Tab. 7. In particular, we experiment with UniTNT and BLIP in Visual Question Answering and Image Captioning using separate training on vision-oriented and OCR-oriented datasets and combined training. In VQA, using both dataset types leads to the best standalone and average performance in the tested benchmarks. This attests to the symbiosis between general and scene-text-oriented VQA, encouraging avoidance of the common practice of separate finetuning.

However, using a unified training set in CAP leads to the best COCO Captions and average results, but not in TextCaps. Specifically, separate finetuning on TextCaps achieves a CIDEr score of 130.5, compared to 119.1 in the combined training. It corresponds with [40], which shows that combining COCO Captions with an upsampled version of TextCaps reduces the model’s performance on the former. It is because while training on TextCaps encourages the model to insert OCR into the caption, training on COCO Captions which barely contains OCR in its captions, penalizes such behavior, leading to an intrinsic tradeoff. To better understand the effects of training models solely on TextCaps, we qualitatively test them on COCO Captions. Notably, we finetune both BLIP and UniTNT of TextCaps and demonstrate their performance on COCO Captions in Fig. 6. Our analysis shows that as TextCaps contains OCR in all its captions, separate finetuning causes models to fixate on OCR, regardless of their importance. Moreover, in images without an OCR signal, the models sometimes hallucinate text in the image. While both models showcase similar behavior, since UniTNT has better scene-text understanding, it is more prone to such phenomena. It is also expressed in Tab. 7, where BLIP and UniTNT trained on TextCaps obtain 84.8 and 70.4 on COCO Captions, respectively. Despite the improved performance on TextCaps when performing separate finetuning on it, our findings highlight its drawbacks. Thus, we claim that also in CAP, combined training should be applied.

From a general view, we hypothesize that since numerous valid captions exist for a given image, both with and

without OCR, the model struggles to decide whether to use the OCR in its caption. Due to the datasets' sizes in combined training that favors the vision-oriented ones, the model opts to reduce its use of OCR, not fully maximizing its performance on TextCaps. It is contrary to VQA, where the conditioning over the question makes it easier for the model to decide whether to use OCR or not (*e.g.*, "What is written in the sign?" versus "What color is this shirt?").

D. Qualitative Analysis

Visual Question Answering We provide an additional qualitative demonstration of UniTNT and compare it to BLIP and M4C on both TextVQA validation set (Fig. 7) and VQAv2 test set (Fig. 8). We depict in the four leftmost columns success-cases and the rightmost, fail cases, and color in green the correct answers and red, incorrect ones. Moreover, we divide the figures such that the upper part corresponds with the benchmark's goal (VQAv2 – see, TextVQA – read) and the lower one with the counterpart goal (VQAv2 – read, TextVQA – see). These results further demonstrate that UniTNT is capable of reasoning over both visual and scene-text information, while other competing methods perform unsatisfactorily on at least one of the benchmarks. Moreover, as the visualizations in Fig. 8 testify, granting scene-text understanding also benefit VQAv2, corresponding with the quantitative evidence in the main paper. It is demonstrated in the bottom part of the figure, where the OCR is crucial for answering the questions or providing meaningful information that facilitates answering them.

Image Captioning Similar to the VQA demonstration, we present a visualization of UniTNT performance on TextCaps (Fig. 9) and COCO Captions (Fig. 10) and compare the performance to M4C and BLIP. On the left columns, we show images where our method outperforms the other methods, and on the right, its failure cases. Moreover, we list the CIDEr scores of each prediction and color in green the highest one. These findings attest that BLIP is incapable of incorporating scene-text information, which results in relatively low CIDEr results. Interestingly, M4C is too overfitted for TextCaps, causing it to fail completely on COCO Captions where OCR is scarce. Specifically, it focuses on the OCR regardless of their importance (*e.g.*, the third example in the last row of Fig. 10) and thus provides an irrelevant caption. Despite the intrinsic tradeoff described in the paper between TextCaps and COCO Captions, UniTNT is capable of providing adequate captions for both benchmarks. Specifically, our method is the only one to cope satisfactorily on both benchmarks altogether and is capable of harnessing both scene-text and visual information.

Hallucinating OCR



BLIP: a young boy is eating a piece of cake with a yellow frosting on it (54.5)

Ours: a young boy is eating a cake with the word cake on it (47.9)



BLIP: a traffic light has a red light on it (42.5)

Ours: a traffic light has a red light that says red on it (28.1)



BLIP: two women are decorating a cake with a pepsi logo on it (96.9)

Ours: two women are decorating a cake on a counter (197.4)

Over-fixation on OCR



BLIP: a man is surfing in the ocean and is wearing a swim suit (18.9)

Ours: a man is surfing in the ocean with the name jimmy bravo (8.8)



BLIP: a display of donuts with a coca cola can in the background (24.3)

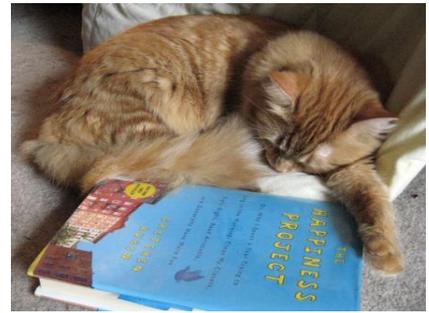
Ours: a coca cola box is behind some donuts (14.6)



BLIP: a seagull is flying over a body of water with the words mr nicholas (26.5)

Ours: a seagull is flying over the water with the words sharklady adventures (71.4)

OCR is useful



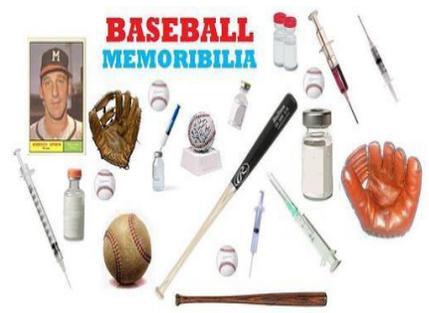
BLIP: a cat sleeping on top of a book that has the word paris on it (103.1)

Ours: a cat sleeping on a book titled happiness project (184.0)



BLIP: a boy wearing a green and yellow jersey with the word fell on it (82.8)

Ours: a boy in a jerlin baseball uniform holds a bat (127.2)



BLIP: a poster with a baseball player and the words baseball memories (34.5)

Ours: a picture of baseball items and the words baseball memorabilia (93.2)

Figure 6: **Qualitative demonstration of the effects of finetuning on TextCaps.** BLIP and UniTNT results of COCO Captions when finetuned solely on TextCaps. In some cases, scene-text understanding helps the models, but it also leads to over-reliance on the OCR signal and even to the hallucination of OCR. While such phenomena occur in both models, it is more prevalent in UniTNT due to its better scene-text understanding.



What is the name of this gateway?

M4C: motel
BLIP: golf course
Ours: grand canyon
GT: grand canyon



What is it nice to be?

M4C: the nice
BLIP: to be funny
Ours: important
GT: important



What company makes these cakes?

M4C: sweet
BLIP: unknown
Ours: dq
GT: dq



What is the last name on the right?

M4C: london
BLIP: clocktower
Ours: peckham
GT: peckham



What brand liquor is on the right?

M4C: morangier islay
BLIP: wilddorf
Ours: scotch whisky
GT: bowmore



What are the first 3 letters of the left boxer's name?

M4C: m
BLIP: mkt
Ours: moa
GT: moa



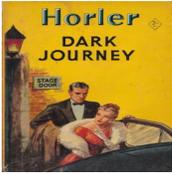
What is the name of the food truck?

M4C: real estate
BLIP: dominio's
Ours: dominic's
GT: dominic's



What is the most this cup can measure?

M4C: cups
BLIP: 12 ounces
Ours: 16 oz
GT: 16 oz



What is the title of the book?

M4C: dark horse
BLIP: unknown
Ours: dark journey
GT: dark journey



What brand is the bottle with red label?

M4C: jack daniels
BLIP: jagermeister
Ours: jim beam
GT: red label



What is the highest number on the players short?

M4C: 4
BLIP: 6
Ours: 8
GT: 8



What does the sign below the stop one mean?

M4C: only
BLIP: curve in road
Ours: right turn only
GT: right turn only



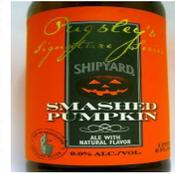
Whose revenge ale is this?

M4C: foster's
BLIP: peroni's
Ours: perry's
GT: perry's



What is the name on the label?

M4C: the hungry
BLIP: softcover
Ours: sprecher
GT: sprecher



What is the name of this ale?

M4C: smashed ale
BLIP: pumpkin ale
Ours: shipley's
GT: smashed pumpkin



Continuing straight take you where?

M4C: dubai
BLIP: al musalla road
Ours: Hatta, oman
GT: Hatta, oman



What does it say on the man's stripped shirt?

M4C: Bwin foundation
BLIP: qantas
Ours: qatar foundation
GT: qatar foundation



What number is on the black and white sign?

M4C: 15
BLIP: 20
Ours: 201
GT: 201



What do the red signs say behind the men in field?

M4C: Red sox
BLIP: cups
Ours: clips
GT: clips



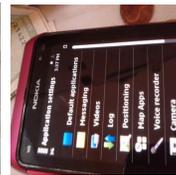
What is that in the orange bottle?

M4C: rock
BLIP: orange juice
Ours: sprite
GT: sunkist



What does the graffiti say?

M4C: simple casual way
BLIP: my best guess yes
Ours: sym
GT: sym



What does it say next to the microphone?

M4C: settings
BLIP: no microphone
Ours: voice recorder
GT: voice recorder



What number is on the tail of the plane?

M4C: 004
BLIP: 055
Ours: j-5004
GT: j-5004



What number is the pitcher?

M4C: 11
BLIP: 0
Ours: 23
GT: 23



How many grams of flour are in the measuring cup?

M4C: 60
BLIP: two
Ours: 10
GT: 201

Figure 7: Qualitative demonstration on TextVQA validation. UniTNT, M4C, and BLIP answers, containing both success (left) and fail (right) cases of our method on image-question pairs that require mainly reading (top) and ones that require also visual reasoning (bottom).

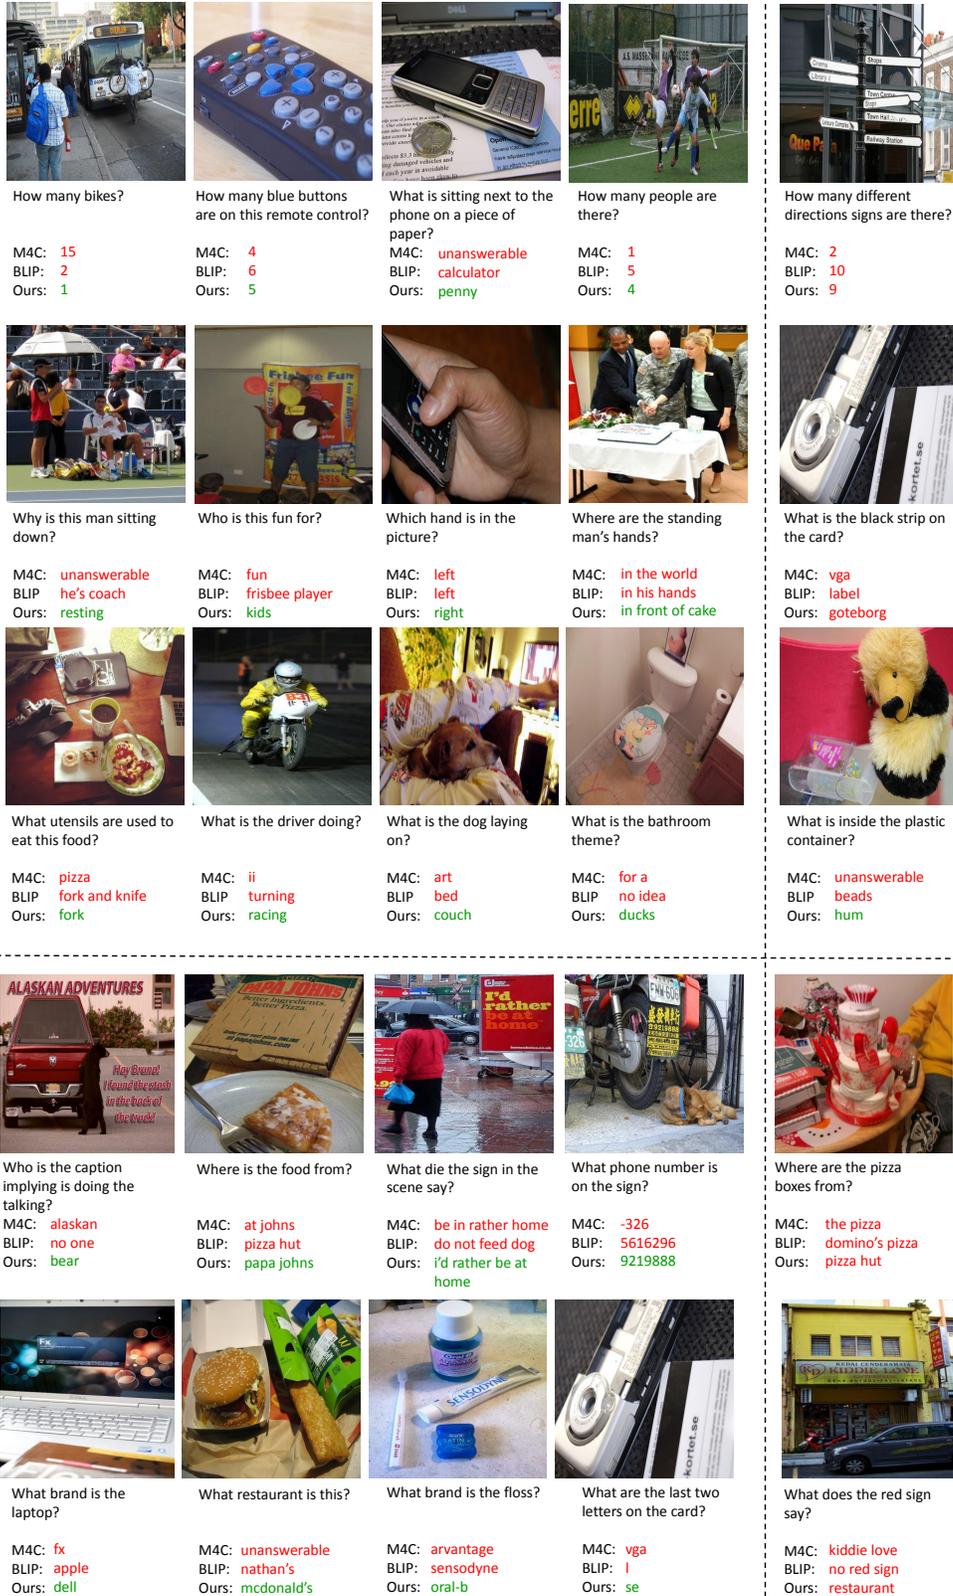


Figure 8: **Qualitative demonstration on VQAv2 test.** UniTNT, M4C, and BLIP answers, containing both success (left) and fail (right) cases of our method on image-question pairs that require mainly vision (top) and ones that require also scene-text understanding (bottom).

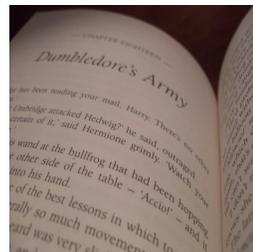
			
M4C: a red stop sign that is outside in the daytime (94.5)	M4C: a bottle of virgin wine is on a white surface (84.3)	M4C: a glass of big omaha beer is sitting on a table (184.7)	M4C: a poster that says ' city traveler ' on it (9.6)
BLIP: a stop sign in front of a building (125.9)	BLIP: a close up of a bottle of wine on a table (19.7)	BLIP: a large metal bucket sitting on top of a table (49.2)	BLIP: a black background with the words air port in different languages (22.2)
Ours: stop sign with arabic writing on it (192.5)	Ours: a bottle of extra virgin extra virgin olive oil (258.5)	Ours: a bucket that says big omaha 2009 on it (330.4)	Ours: the word pdx is on a black background (14.2)
			
M4C: a bottle of holmes point marlborough from marlborough (247.3)	M4C: a road sign for the giessen of winchester (135.7)	M4C: a green lenovo phone with the time of 11:00 (104.4)	M4C: a woman stands in front of a large screen that says flood on it (2.9)
BLIP: a bottle of wine sitting on top of a table (41.8)	BLIP: a street sign sitting on the side of a road (28.2)	BLIP: a close up of a cell phone on a table (53.0)	BLIP: a group of people sitting at a table in front of a screen (15.9)
Ours: a bottle of holmes point sauvignon blanc wine (443.6)	Ours: a sign for the city of winchester in england (226.8)	Ours: a black lenovo cell phone on a white surface (251.yuc4)	Ours: a screen shows a woman speaking at a conference (9.6)
			
M4C: a united states navy plane is flying in the sky (290.6)	M4C: a book is open to a page that says 'a dumb army' (137.7)	M4C: a car with a yellow license plate that says m6 tal (274.0)	M4C: several coins on a table including one that says 'united states of america' (34.2)
BLIP: a small propeller plane flying through a blue sky (36.6)	BLIP: a close up of an open book on a table (51.1)	BLIP: a silver car with a yellow license plate (150.9)	BLIP: a bunch of different types of coins (18.7)
Ours: a man flying through the air while riding a skateboard (180.4)	Ours: a book is open to a page titled dumbledore's army (399.8)	Ours: a silver car with the license plate m6 tal (332.1)	Ours: a collection of united states quarters (9.5)

Figure 9: **Qualitative demonstration on TextCaps.** UniTNT, M4C-Captioner, and BLIP answers, containing both success (left) and fail (right) cases of our method alongside the per-caption CIDEr score.



M4C: a white car with the number 3 on it (0.5)

BLIP: a group of people sitting on top of a sandy beach (59.2)

Ours: a group of people on the beach under an umbrella (162.6)



M4C: a large white and red sign that says 'd' on it (0.4)

BLIP: a man sitting at a table with a plate of food (68.5)

Ours: a man in a tie is smiling for the camera (116.7)



M4C: a white car with the word "no" on it (2.5)

BLIP: a man riding a wave on top of a surfboard (123.4)

Ours: a man riding a surfboard on top of a river (181.8)



M4C: a small white sign that says "w" on it (1.5)

BLIP: a group of people standing in a room (27.7)

Ours: a man and a woman standing next to each other (17.3)



M4C: a sign that says vote & laduke on it (0.1)

BLIP: a man sitting at a table with a plate of food (42.8)

Ours: a man in a green shirt holding a glass of wine (186.2)



M4C: a large white sign that says "no parking" on it (10.9)

BLIP: a man riding a skateboard up the side of a ramp (69.0)

Ours: a man flying through the air while riding a skateboard (180.4)



M4C: a picture of a woman in a suit with a sign that says "say say cheese!" (0)

BLIP: two stuffed animals sitting next to each other on a chair (52.3)

Ours: two stuffed animals are sitting next to a book (99.3)



M4C: a picture of a man with the number 3 on it (1.7)

BLIP: a bathroom with a washer and a window (23.9)

Ours: a bathroom with a washer and dryer in it (20.9)



M4C: a picture of a woman and a yellow and white dress with the word "middle middle" on it (0.2)

BLIP: a couple of people that are holding a skateboard (7.8)

Ours: a man holding a snowboard next to another man (196.4)



M4C: a large number of a train is on the ground with a red and white sign that says "sample" (0.4)

BLIP: a person holding a piece of fabric in their hand (41.7)

Ours: a person holding a tie in their hand (205.5)



M4C: a book called warcraft is on a table with a picture of a person in the background (5.0)

BLIP: a brown teddy bear sitting on top of a desk (156.7)

Ours: teddy bear wearing headphones sitting on a desk (189.5)



M4C: a green sign that says 'no parking' on it (6.7)

BLIP: a group of people walking across a street next to a tall building (16.4)

Ours: a group of people walking across a street (14.7)

Figure 10: **Qualitative demonstration on COCO Captions.** UniTNT, M4C-Captioner, and BLIP answers, containing both success (left) and fail (right) cases of our method alongside the per-caption CIDEr score.