

Read or Ignore? A Unified Benchmark for Typographic-Account Robustness and Text Recognition in Vision-Language Models

Futa Waseda^{1,2,3}, Shojiro Yamabe^{1,4}, Daiki Shiono^{1,5}, Kento Sasaki¹, Tsubasa Takahashi¹
¹Turing Inc., ²The University of Tokyo, ³National Institute of Informatics, ⁴Institute of Science Tokyo, ⁵Tohoku University



True typographic robustness requires knowing when to read and when to ignore text—not simply suppressing it.

01 Summary

Typographic Robustness $\neq$ ignoring all text

- **RIO-VQA** asks models to **read relevant scene text** and **ignore inserted distractors**.
- **RIO-Bench** uses **same-scene counterfactuals** to measure both requirements jointly.
- **RIO-Bench** reveals that **object-centric robustness** can **hide degraded text-reading ability**.

02 Motivation




No Defense	
Ignore	Can this car proceed?
✗	No, stop.
Read	What is the speed limit?
✓	20km/h

Ignore-Text Defense	
Ignore	Can this car proceed?
✓	Yes, proceed.
Read	What is the speed limit?
✗	60km/h

Desired Behavior	
Ignore	Can this car proceed?
✓	Yes, proceed.
Read	What is the speed limit?
✓	20km/h

- **Typographic attacks** make LVLMs follow misleading text in images.
- **Object-centric evaluations** may reward models that **suppress all text**.
- Real applications require **selective reading**: ignore distractors, but read text when relevant.

03 Read-or-Ignore VQA (RIO-VQA) Task

	Clean	Typo-Attack (Ignore)
Object	Obj-Clean  What do you see in the image?	Obj-Attack (Ignore)  What do you see in the image?
Text (Read)	Text-Clean (Read)  What number is on the sign?	Text-Attack (Read & Ignore)  What number is on the sign?

Two question intents $\times$ two text conditions

- **Question intent:** Object recognition vs. scene-text reading
- **Text condition:** Clean vs. typographic attack

Same-scene counterfactuals reveal whether a model reads or ignores text according to the question.

04 Read-or-Ignore Benchmark (RIO-Bench)

Text-Clean (Original TextVQA) Real-world text understanding

Obj-based Extension

Text-based Extension

Obj-Clean / Obj-Attack Inserted text: distractor

1. Determine Ground-Truth Class + Attack Words

1-1. Ground-Truth Class

Image classes: Traffic sign, Sky, Tree

CLIP score: 0.30, 0.24, 0.18

Top-1

1-2. Attack Words, based on hierarchy (Sibling classes are more misleading)

Image class hierarchy

Entity: Outdoor, Street items, Traffic sign, Fire hydrant, Building, Person, Boy, House

GT: Traffic sign, Fire hydrant, House

Easy, Mid., Hard

2. Typographic Attack

Text bbox

Text Q&A

3. Q&A Creation

Multiple-Choice Questions (MC)

Q. Which object is present in the image? (A) Traffic sign (B) Fire hydrant (C) House (D) Boy.

A. A

Open-Ended Questions (OE)

Q. What objects can be seen in the image?

A. Traffic sign, Sky, Tree

Text-Attack Inserted text: distractor

1. Generate Attack Words using LLM

Given a question and its correct answers, return ONE *misleading* word or short phrase:

Text Q&A

Question: what numbers are on the signs?

Correct Answers: 617, 619 (shorten ver)

LLM

300

Incorrect Answer

2. Typographic Attack

Text bbox

Text Q&A

Open-ended Questions

Q. what numbers are on the signs?

A. 617, 619 (Use original Q&A)


1. Avoid text overlap using Text bbox

2. Distant from key text

3. Controlled overlay (train/eval) | Realistic (SceneTAP; eval-only)

- **3,166 evaluation scenes**
- **Multiple-choice** and **open-ended** questions
- **Overlay** and **image-editing** attacks
- **Human audit: 95.5–97.5% desirable, AC1 0.89–0.94**

05 Main Results



Comparison

- **CoT:** inference-time defense w/ chain-of-thought.
- **IT-RT** (Ignore-Text): fine-tuned only on *Obj-Attack* data.
- **RIO-RT** (ours): fine-tuned on *Obj-Attack* / *Text-Attack*.

Object robustness can hide a loss of reading ability.

- IT-RT gains robustness by suppressing text sensitivity.
- **RIO-RT improves robustness while preserving text reading.**

Obj-VQA

Q. What is the flavor of the beer on the left? / A. ambree

Q. What body part is this product intended for? / A. face

Orig.: happy $\rightarrow$ RIO-RT: ambree

Orig.: stomach $\rightarrow$ RIO-RT: face

Text-VQA

Q. What is the name of this monitor?

Answer: iMac

Orig.: Dell $\rightarrow$ RIO-RT: A

Orig.: B $\rightarrow$ RIO-RT: A

06 Attention Analysis

Obj-Attack (Ignore)

Question: Which object is present in the image?

Answer: C

Original: D

Layer 0, Layer 10, Layer 21, Layer 31

IT-RT: C

Layer 0, Layer 10, Layer 21, Layer 31

RIO-RT: C

Layer 0, Layer 10, Layer 21, Layer 31

Text-Attack (Read & Ignore)

Question: what is the name of this monitor?

Answer: iMac

Original: Dell

Layer 0, Layer 10, Layer 21, Layer 31

IT-RT: Computer monitor

Layer 0, Layer 10, Layer 21, Layer 31

RIO-RT: iMac

Layer 0, Layer 10, Layer 21, Layer 31