

Read or Ignore? A Unified Benchmark for Typographic-Attack 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

TURING



東京大学
THE UNIVERSITY OF TOKYO



大学共同利用機関法人 情報・システム研究機構
国立情報学研究所
National Institute of Informatics



Institute of
SCIENCE
TOKYO



TOHOKU

TOHOKU
UNIVERSITY

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 defense** may reward models that **suppress all text (“Ignore-Text”)**.
- Real applications require **selective reading**: ignore distractors, but read text when relevant.

Read-or-Ignore VQA

TURING



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

Two question intents × 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.

We created a benchmark to evaluate read-or-ignore capability.

Obj-VQA

Attack word: Jug



(A) Tin can (B) Piano
(C) Camera (D) Jug

Recognize **object**,
while ignoring **text**

Text-VQA

Attack word: hoppy



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

Recognize **the key text**,
while ignoring **distractor text**



- 3,166 evaluation scenes
- 4 taxonomies:
(Obj/Text-VQA)
× (Clean/Attack)
- Overlay / image-
editing attacks

Key Finding

TURING



Object-centric Robustness Can Hide Lost Reading Ability.



IT-RT (Ignore-text defense) suppresses text; RIO-RT (ours) reads text selectively.

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**.