

IEEE/ACM
2025 INTERNATIONAL
CONFERENCE ON
COMPUTER-AIDED
DESIGN

44th Edition

MM-GRADE: A Multi-Modal EDA Tool Documentation QA Framework Leveraging Retrieval Augmented Generation

Yuan Pu ♠♥

Hairuo Han ♠

David Pan ◇

Zhuolun He ♠♥

Haisheng Zheng ♥

Bei Yu ♠

Shutong Lin ♠

Yuqi Jiang ♣

Jiajun Qin ♠

Cheng Zhuo ♣

Xinyun Zhang ♠

Qi Sun ♣

♠ The Chinese University of Hong Kong

♣ Zhejiang University

♥ ChatEDA Technology

◇ University of Texas at Austin

Oct. 29, 2025



Outline

- 1 Introduction
- 2 Proposed Algorithm
- 3 Experimental Results

Introduction

EDA Tool Documentation QA

- EDA tool involves powerful but complex functionalities/documents → steep learning curve.
- Multi-Modal Automatic Question-Answering (QA) system:
 - Given the **user query**, the QA system generates **answer** referring to the relevant **documents**.
 - User query/document contain useful visual information:
 - GUI screenshots.
 - Design layouts.
 - Screenshot of error traceback.

In OpenROAD, I don't want pins to be placed in the region highlighted by the green rectangle, what should I do?



Retrieved Document:

The 'place_pins' command ... use the '-exclude' option to specify a region where pins cannot be placed ...

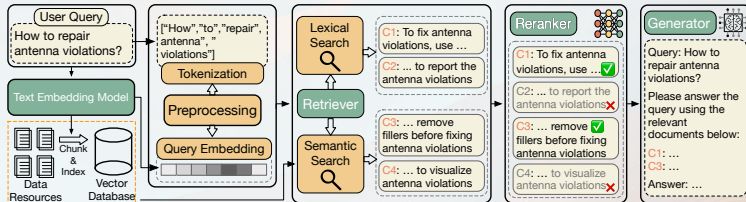


The green rectangle is at the bottom of the layout.
You can use the command: **place_pins —exclude bottom:***
To avoid pins placed at the bottom region.

A Typical Multi-Modal Query in OpenROAD.

RAG-based EDA QA Systems and Its Limitations

- Trend: Use LLMs and Retrieval Augmented Generation for EDA Tool QA.
- Existing works: RAG-EDA^[1], etc.
- Stages:
 - Retriever: retrieve the relevant documents to the user query.
 - Generator: usually LLM, generates the answer.
- Limitations: They are designed for **text-only** query, ignore **visual information** → fundamental gap in understanding the user's true intent.



Overall Flow of RAG-EDA.

^[1] Yuan Pu et al. (2024). "Customized Retrieval Augmented Generation and Benchmarking for EDA Tool Documentation QA". In: *Proc. ICCAD*.

General-Purpose Multi-Modal RAG Methods: Limitations

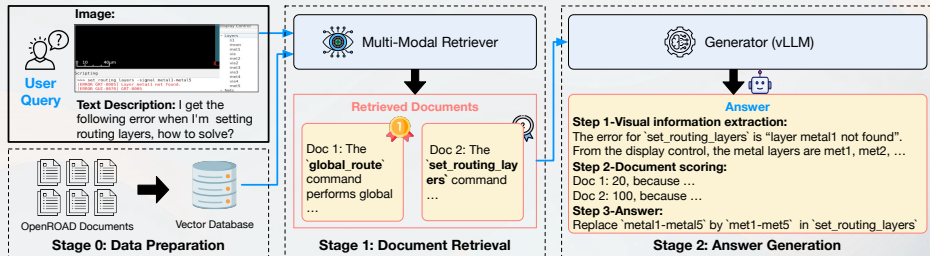
- Existing Multi-Modal RAG Methods: VisRAG^[2], EchoSight^[3], etc.
- Limitations of being applied to EDA tool documentation QA:
 - Retriever
 - General MM-RAG retrievers are trained on natural scenes, charts → struggle to comprehend the complex, specialized information in EDA visuals.
 - Generator
 - General Vision-Language Models (VLMs) lack EDA domain knowledge → *spurious understanding* of the user's query and poor quality answers

^[2] Shi Yu et al. (2025). "VisRAG: Vision-Based Retrieval-Augmented Generation on Multi-Modality Documents". In: *Proc. ICLR*.

^[3] Yibin Yan et al. (2024). "EchoSight: Advancing Visual-Language Models with Wiki Knowledge". In: *Proc. EMNLP*.

Proposed Algorithm

Overall Flow



Overall Flow of MM-GRADE.

- Solution: MM-GRADE, a multi-modal, RAG-based EDA tool QA framework.
- Techniques:
 - Multi-modal Retriever: Bi-level negative mining strategy → contrastive learning.
 - Generator (vision large language model): customized 3-stage reasoning pipeline: Extract-Score-Answer (ESA) pipeline.
- Benchmark: ORD-MMBench, consisting of 120 QAs (from OpenROAD) with screenshots

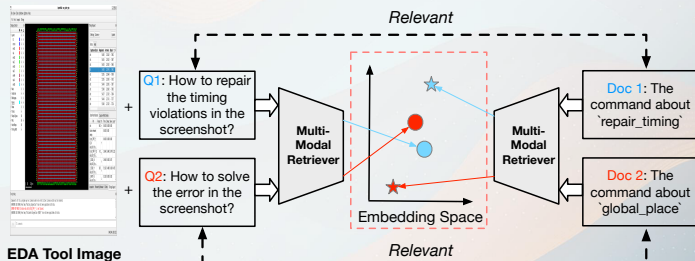
Existing Training Strategy of Multi-Modal Retriever

- Mechanism of multi-modal retriever for document retrieval:
 - Step 1: project the user query and all document chunks (multi-modality) into a unified embedding space.
 - Step 2: retrieve the most relevant document chunks by similarity matching.
- Training strategy: Contrastive learning with **document-level hard negative mining**:
 - Given a user query q (consists of text part q_i^{txt} and image part q_i^{img}).
 - positive samples: its relevant document d_i^+
 - negative samples: its n irrelevant documents $d_i^- = \{d_{i,1}^-, d_{i,2}^-, \dots, d_{i,n}^-\}$
- Training loss:

$$L_d^i = -\log \frac{e^{f(q_i^{\text{img}} \| q_i^{\text{txt}})^\top f(d_i^+) / \tau}}{e^{f(q_i^{\text{img}} \| q_i^{\text{txt}})^\top f(d_i^+) / \tau} + \sum_{j=1}^n e^{f(q_i^{\text{img}} \| q_i^{\text{txt}})^\top f(d_{i,j}^-) / \tau}}, \quad (1)$$

Limitations of Multi-Modal Retriever

- Observation
 - One EDA tool screenshot contains complex and diverse information → triggers multiple, distinct user questions.
- Consequence
 - Many different EDA Tool user questions share one common image (screenshot).
- Limitation
 - Exiting general-purpose Multi-Modal retrievers focus on visual information → distinct user queries (with the same image) are projected closely in the embedding space → Low document retrieval accuracy.
- Example:



Solution of EDA-Customized Multi-Modal Retriever: Bi-Level Hard Negative Mining

- Solution: contrastive learning with **query-level hard negative mining**:
 - Preliminary: For one screenshot image q_i^{img} , m different user queries with corresponding text descriptions $\{q_{i,1}^{\text{txt}}, q_{i,2}^{\text{txt}}, \dots, q_{i,m}^{\text{txt}}\}$, can be asked. These m queries correspond to different relevant documents, denoted as $\{d_{i,1}^+, d_{i,2}^+, \dots, d_{i,m}^+\}$.
 - Given a user query $q = \{q_i^{\text{img}}, q_{i,j}^{\text{txt}}\}$ ($j \in [1, m]$).
 - Positive Samples: its relevant documents $d_{i,j}^+$
 - Negative Samples: distinct queries with **the same image** $\{(q_i^{\text{img}}, q_{i,k}^{\text{txt}}) | k \in [1, m] \wedge k \neq j\}$
- Training loss:

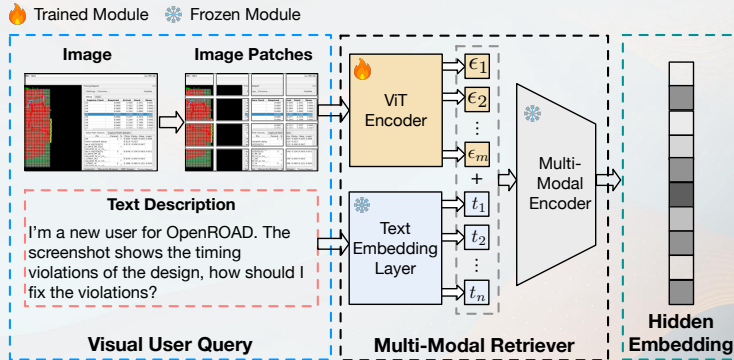
$$L_q^i = - \sum_{j=1}^m \log \frac{e^{f(q_i^{\text{img}} \| q_{i,j}^{\text{txt}}) T f(d_{i,j}^+) / \tau}}{e^{f(q_i^{\text{img}} \| q_{i,j}^{\text{txt}}) T f(d_{i,j}^+) / \tau} + \sum_{k \neq j} e^{f(q_i^{\text{img}} \| q_{i,j}^{\text{txt}}) T f(q_i^{\text{img}} \| q_{i,k}^{\text{txt}}) / \tau}}. \quad (2)$$

- Bi-level hard negative mining (HNM): integrate document-level and query-level HNM:

$$L = \sum_{i=1}^N L_d^i + L_q^i.$$

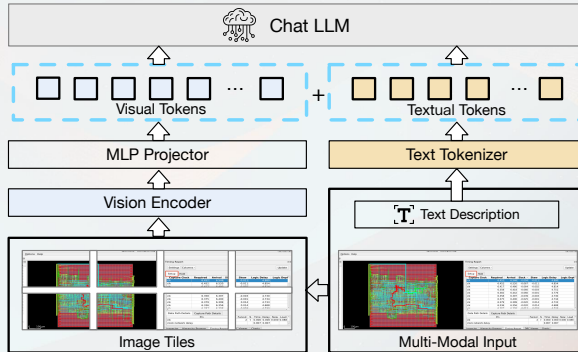
Multi-Modal Retriever: Proposed Model Architecture

- Components:
 - ViT Encoder: Encodes the visual input.
 - Text Embedding Layer: Encodes the textual input
 - Multi-Modal Encoder: Text-visual fusion for embedding generation.



Multi-Modal Generator: Model Architecture and Limitations

- Model Architecture: Based on InternVL-2^[4].
- Limitations when applied to EDA Tool QA: Lacks domain-specific inference logic.



^[4] Zhe Chen et al. (2024). "InternVL: Scaling Up Vision Foundation Models and Aligning for Generic Visual-Linguistic Tasks". In: *Proc. CVPR*.

Solution of Multi-Modal Generator: ESA Pipeline


Extract-Score-Answer Pipeline:

- Pipeline the answer into 3 stages:
 - Stage 1: Visual information extraction.
 - Stage 2: Document relevance scoring.
 - Stage 3: Answer
- Purpose: used for both training and inference.
- Scoring Criteria of Stage 2, between query q and document d :

Score	Description
100	d is perfectly relevant to q and can be used to directly answer q .
80	d is relevant to q and can partially answer q .
60	d is relevant to q but can not be directly used to answer q .
40	It is possible that d is relevant with q , but more information is required to verify the relevance.
20	d is basically not relevant with q .
0	d is totally irrelevant with q .

Solution of Multi-Modal Generator: ESA Pipeline

User Visual Query-Image:



User visual query text description:

I don't want the pins to be placed on the region of the green rectangle in the screenshot, what should I do?

Retrieved Documents:

Doc 1: `place_pin`: the command to place a specific pin ...

Doc 2: use the option `exclude` of the `place_pins` command to ...

Extraction-Scoring-Answer Strategy

Stage 1: Visual information extraction

Prompt:

Extract the information from the image which can be used to understand and answer the question.

Answer:

The pins highlighted in the green rectangle is the bottom region of the layout.

Stage 2: Document Relevance Scoring

Prompt:

Now given the retrieved documents, analyse and score the relevance between each document and the visual query. The score ranges from 0 to 100, higher scores represent higher relevance.

Answer:

Doc 1: 20, because the query is about the command to place all pins, but `place_pin` command only places one pin.

Doc 2: 100, because the document introduces the usage of `place_pins` and the `exclude` flag.

Stage 3: Answer

Prompt:

Answer this user query according to the extracted visual information in stage 1 and the document relevance scoring results in stage 2.

Answer:


You can use the command: `place_pins -hor_layers <hor layers> -ver_layers <ver layers> -exclude bottom:*` to disable pins placed on the bottom region.

Prompt Example of the ESA Pipeline

Experimental Results

Benchmark: ORD-MMBench

- ORD-MMBench^[5]: 120 high-quality multi-modal question-document-answer triplets based on the OpenROAD^[6] document.
- Four categories of questions:



^[5] <https://github.com/lesliepy99/MM-GRADE-Benchmarks-ICCAD>

^[6] Tutu Ajayi et al. (2019). "Toward an Open-Source Digital Flow: First Learnings from the OpenROAD Project". In: *Proc. DAC*.

Experimental Setting

- Multi-Modal Retriever:
 - Base model: BGE-Visualized-M3^[7].
 - Training data: 1075 query-document corpus items generated by gpt-4o.
 - Finetuned by 3000 steps (batch size = 8) on 4 A100 80G GPUs.
- Multi-Modal Generator:
 - Base model: InternVL-2-26B^[8].
 - Training data: 2307 high-quality query-document-answer triplets based on OpenROAD^[9].
 - Trained by 2 epochs (batch size = 1) on 8 A100 80G GPUs.

[7] https://huggingface.co/BAAI/bge-visualized/blob/main/Visualized_m3.pth

[8] <https://huggingface.co/OpenGVLab/InternVL2-26B>

[9] Tutu Ajayi et al. (2019). "Toward an Open-Source Digital Flow: First Learnings from the OpenROAD Project". In: *Proc. DAC*.

Experimental Results: Retrieval Baselines

- Metric: Recall@ k , defined as the proportion of relevant documents that are retrieved among the Top- k results returned by the retriever/reranker.
- Baselines:
 - Text-Only Retriever:
 - BGE-M3^[10], RAG-EDA-Retriever
 - Multi-Modal Retriever:
 - DEDR^[11], BLIP^[12], VisRAG-Retriever^[13], EchoSight-Retriever^[14], BGE-Visualized-M3

^[10] Jianlv Chen et al. (2024). “BGE M3-Embedding: Multi-Lingual, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation”. In: *Proc. ACL*.

^[11] Alireza Salemi et al. (2023). “A Symmetric Dual Encoding Dense Retrieval Framework for Knowledge-Intensive Visual Question Answering”. In: *Proc. SIGIR*.

^[12] Junnan Li et al. (2022). “Blip: Bootstrapping Language-Image Pre-Training for Unified Vision-Language Understanding and Generation”. In: *Proc. ICML*.

^[13] Shi Yu et al. (2025). “VisRAG: Vision-Based Retrieval-Augmented Generation on Multi-Modality Documents”. In: *Proc. ICLR*.

^[14] Yibin Yan et al. (2024). “EchoSight: Advancing Visual-Language Models with Wiki Knowledge”. In: *Proc. EMNLP*.

Experimental Result: Multi-Modal Retrieval

- Ablation Study
 - Train MM-GRADE Retriever (Ours) with only query-level/document-level hard negative mining (HNM).
- MM-GRADE retriever with bilevel HNM achieves the best performance.

Model Type	Recall@1	Recall@2	Recall@3	Recall@4	Recall@5
BGE-M3	0.590	0.694	0.731	0.761	0.776
RAG-EDA-Retriever	0.530	0.612	0.679	0.724	0.739
DEDR	0.112	0.187	0.231	0.284	0.299
BLIP	0.127	0.209	0.276	0.299	0.313
VisRAG-Retriever	0.025	0.042	0.050	0.083	0.108
EchoSight-Retriever	0.383	0.492	0.583	0.625	0.667
VISTA: BGE-Visualized-M3	0.582	0.716	0.739	0.761	0.799
MM-GRADE-Retriever w. d-level HNM	0.664	0.784	0.806	0.828	0.851
MM-GRADE-Retriever w. q-level HNM	0.582	0.694	0.746	0.761	0.828
MM-GRADE Retriever (Ours)	0.679	0.828	0.843	0.866	0.896

Table: Performance of Document Retrieval on ORD-MMBench.

Experimental Result: Overall RAG Flow

- Metrics:
 - LLM-Score: GPT-4o to score each answer (from 0 to 100) referring to a scoring criteria [5].
 - BLEU and Rouge-L.
- RAG Flow Baselines:
 - Text-Only: RAG-EDA.
 - Multi-Modal: VisRAG, EchoSight, VISTA-RAG.

RAG Flow	ORD-MMBench▶error traceback			ORD-MMBench▶layout			ORD-MMBench▶functionality			ORD-MMBench▶GUI			ORD-MMBench▶all		
	LLM-Score	BLEU	ROUGE-L	LLM-Score	BLEU	ROUGE-L	LLM-Score	BLEU	ROUGE-L	LLM-Score	BLEU	ROUGE-L	LLM-Score	BLEU	ROUGE-L
VisRAG	41.129	0.080	0.223	12.500	0.098	0.242	18.750	0.111	0.247	4.762	0.024	0.162	20.417	0.084	0.225
EchoSight	51.613	0.087	0.241	54.688	0.165	0.300	62.500	0.191	0.322	64.286	0.091	0.214	57.917	0.139	0.276
RAG-EDA	22.419	0.064	0.187	61.719	0.156	0.276	59.722	0.190	0.332	51.190	0.107	0.267	49.125	0.134	0.268
VISTA-RAG	65.323	0.126	0.266	68.750	0.216	0.348	71.528	0.232	0.368	73.810	0.116	0.261	69.583	0.180	0.317
MM-GRADE w/o ESA	71.129	0.159	0.303	78.125	0.273	0.418	75.833	0.291	0.442	90.476	0.160	0.312	77.792	0.229	0.377
MM-GRADE (ours)	82.258	0.244	0.391	80.469	0.307	0.468	80.000	0.294	0.450	83.333	0.178	0.323	81.292	0.264	0.417

Table: Performance of the Multi-Modal RAG Flows on ORD-MMBench.

[5] <https://github.com/lesliepy99/MM-GRADE-Benchmarks-ICCAD>

THANK YOU!