

LOCALRQA: From Generating Data to Locally Training, Testing, and Deploying Retrieval-Augmented QA Systems

Xiao Yu*, Yunan Lu*, Zhou Yu

Department of Computer Science, Columbia University, New York, NY
{xy2437, yl4021, zy2461}@columbia.edu

Abstract

Retrieval-augmented question-answering systems combine retrieval techniques with large language models to provide answers that are more accurate and informative. Many existing toolkits allow users to quickly build such systems using off-the-shelf models, but they fall short in supporting researchers and developers to customize the *model training, testing, and deployment process*. We propose LOCALRQA¹, an open-source toolkit that features a wide selection of model training algorithms, evaluation methods, and deployment tools curated from the latest research. As a showcase, we build QA systems using online documentation obtained from Databricks and Faire’s websites. We find 7B-models trained and deployed using LOCALRQA reach a similar performance compared to using OpenAI’s text-ada-002 and GPT-4-turbo.

1 Introduction

Retrieval-augmented question-answering (RQA) systems enhance large language models (LLMs) by enabling them to search through a large collection of documents before answering a user’s query. These systems have shown improved performance in providing more accurate, informative, and factually grounded answers compared to using LLMs alone (Guu et al., 2020; Izacard et al., 2022; Shi et al., 2023). Many existing toolkits, such as LlamaIndex (Liu, 2022) and LangChain (Chase, 2022), allow users to quickly build such an RQA system using off-the-shelf models such as text-ada-002 (OpenAI, 2022a) and GPT-4 (OpenAI, 2023). However, developers often find it costly to rely on these paid services, but also face difficulties to train/deploy smaller models with competitive performance. Researchers face even greater hurdles: they need to modify models/training algorithms, compare against prior work, and obtain

* denotes equal contribution.

¹GitHub: <https://github.com/jasonyux/LocalRQA>,
YouTube: <https://youtu.be/MEtFIcw7c1Y>.

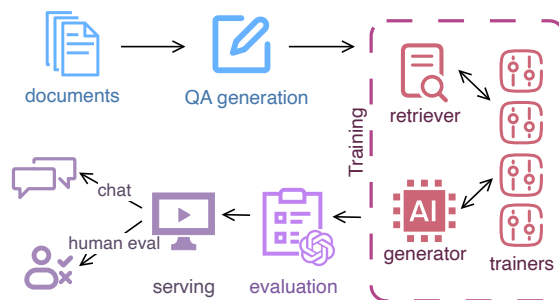


Figure 1: Given a collection of documents, LOCALRQA provides tools to generate RQA data, to train and test open-source models, and to deploy the RQA system for human evaluation or as an interactive chatbot.

human evaluation for their RQA system — all of which are largely neglected by existing toolkits.

We introduce LOCALRQA, an open-source toolkit that enables researchers and developers to easily train, test, and deploy RQA systems *using techniques from recent research*. Given a collection of documents, users can use pre-built pipelines in our framework to quickly assemble an RQA system using the best off-the-shelf models. *Alternatively*, users can create their own training data, train open-source models using algorithms from latest research, and deploy a local RQA system that achieves similar performance compared to using paid services such as OpenAI’s models.

To our knowledge, LOCALRQA is the first toolkit that provides a wide range of training algorithms and automatic evaluation metrics curated from the latest research (see Table 1 and Appendix A). This not only helps researchers to develop new RQA approaches and compare with prior work, but also helps developers to train and deploy more cost-effective models. Specifically, we provide many training algorithms for retrievers such as: distilling from an encoder-decoder’s cross-attention scores (Izacard and Grave, 2020a), distilling from a decoder’s language modeling probability (Shi et al., 2023), and using contrastive learning

	Easy Assembly	Flexible Training	Automatic Evaluation	Local Deployment
Haystack (Pietsch et al., 2019)	✓	✗	✓	✗
LangChain (Chase, 2022)	✓	✗	✗	✗
LLamaIndex (Liu, 2022)	✓	✗	✓	✗
FastRAG (Izsak et al., 2023)	✓	✗	✗	✗
LOCALRQA (Ours)	✓	✓	✓	✓

Table 1: Comparing LOCALRQA to related toolkits. *Easy Assembly* indicates that there are ready-made pipelines to allow users to easily assemble an RQA system; *Flexible Training* indicates if there is more than one training algorithm for retrievers/generators; *Automatic Evaluation* indicates if the toolkit provides automatic evaluation methods; and *Local Deployment* indicates if the toolkit supports methods to locally deploy their RQA system and allow external users to interact with them through a web interface.

approaches (Karpukhin et al., 2020; Wang et al., 2022, 2023). We also provide training algorithms for generative models such as: supervised fine-tuning using gold question-passage-answer pairs (Lewis et al., 2021), fine-tuning with a frozen retriever (Guu et al., 2020), and fusion-in-decoder training (Izacard and Grave, 2020b; Izacard et al., 2022). Then, to automatically evaluate the system’s performance, we implement metrics used in retrieval and question-answering domains, such as Recall@k, ROUGE (Lin, 2004), and GPT-4 Eval (Zheng et al., 2023a; Liu et al., 2023).

Furthermore, LOCALRQA provides two deployment methods to support researchers and developers to obtain human feedback for their RQA systems. First, we offer a static evaluation webpage where users can directly assess the system’s performance using a test dataset. This can be used to complement automatic evaluation. Next, we offer an interactive chat webpage where users can chat with the system and rate the helpfulness and correctness of each generated response. These ratings can be used to further improve models’ capability using techniques such as Reinforcement Learning from Human Feedback (RLHF, Ouyang et al. (2022)). To reduce latency and improve user experiences, our toolkit also integrates acceleration frameworks used to speed up document retrieval (Johnson et al., 2019) and LLM inference (Huggingface, 2023; Kwon et al., 2023; Zheng et al., 2023b). Together with our large collection of training algorithms and automatic metrics, LOCALRQA opens the possibility of future work to easily train, test, and deploy novel RQA approaches.

2 Background

RQA systems combine retrievers with powerful LLMs to provide answers that are more accurate and informative. Given a user query, a retriever first

selects k most relevant passages from a collection of documents. Then, a generative model produces an answer conditioned on the user’s query, selected passages, and a chat history. Popular methods to achieve this include concatenating all inputs into a single string and generating with decoder-only models (Chase, 2022; Ram et al., 2023), or processing the k passages in parallel and generating with fusion-in-decoder techniques (Izacard and Grave, 2020b; Izacard et al., 2022).

3 LOCALRQA

We introduce LOCALRQA, a Python-based toolkit designed to help users flexibly train, test, and deploy RQA systems. As shown in Figure 2, our toolkits employ a modular design to allow users to: generate and prepare RQA data (Section 3.1), train retrieval and generative models (Section 3.2 and Section 3.3), build an RQA system (Section 3.4), evaluate the system (Section 3.5), and finally deploy the system (Section 3.6).

3.1 Prepare Data

A prerequisite for training and evaluating RQA systems is a dataset of (question, answer, passage) pairs, denoted as $\langle q, a, p \rangle$. However, full $\langle q, a, p \rangle$ pairs may not always be available in practice. To cater to various scenarios, our toolkit provides: 1) scripts to generate $\langle q, a, p \rangle$ pairs from a collection of documents, and 2) scripts to convert existing QA datasets into $\langle q, a, p \rangle$ pairs. These scripts can be useful for researchers to create RQA datasets for new domains, or for developers to prepare training/testing data for specific applications.

Generate RQA Data Given a collection of documents, our scripts first use a sampling algorithm to select a set of gold (and hard negative) documents, and then use LLMs to generate questions and answers from each gold document (see Appendix C

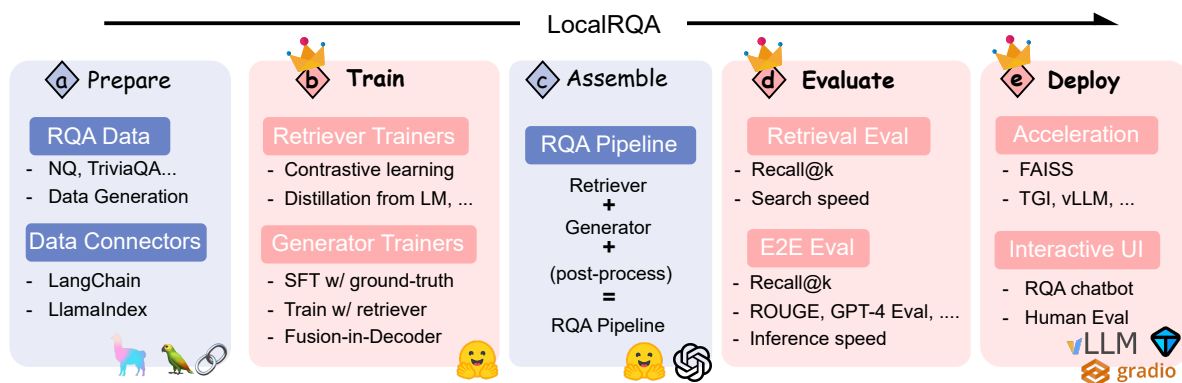


Figure 2: An overview of the LOCALRQA toolkit, which supports the entire pipeline of developing an RQA system: from data processing to training, testing, and serving an RQA system. Different from many existing toolkits, we feature a wide selection of training, testing, and serving methods curated from the latest RQA research.

for more details). These scripts can be used to create $\langle q, a, p \rangle$ pairs not only from a collection of documents, but also from a collection of $\langle q, p \rangle$ pairs (e.g., from information retrieval datasets).

Convert from Existing Datasets Many existing QA datasets include supporting passages for each gold question-answer pair. We provide scripts to download and reformat these datasets into $\langle q, a, p \rangle$ pairs compatible with the rest of our toolkit. This includes popular datasets such as Natural Questions (Kwiatkowski et al., 2019), TriviaQA (Joshi et al., 2017), and MS-Marco (Bajaj et al., 2018). These scripts allow researchers to easily compare against prior work that also uses these datasets.

3.2 Train Retrievers

Given a dataset of $\langle q, a, p \rangle$ pairs, users can train a retriever to select the most relevant passages for a given query. Prior work shows that using better retrievers often leads to more performant RQA systems (Karpukhin et al., 2020), and that fine-tuning them with task-specific data can greatly improve their performance (Izacard et al., 2021; Wang et al., 2022). To this end, LOCALRQA implements: 1) lexical-based and embedding-based methods, and 2) various trainers that finetune open-source embedding models to achieve a better performance.

Supported Models For lexical-based methods, we support BM25 (Trotman et al., 2014). For embedding-based methods, we support all huggingface (Wolf et al., 2020) encoder models such as Contriever (Izacard et al., 2021), E5 (Wang et al., 2022), and BGE (Xiao et al., 2023).

Trainers We implement trainers for encoders that distill from a down-stream LM, and trainers

that perform contrastive learning using a dataset of $\langle q, p \rangle$ pairs (and optionally hard negative examples). This includes trainers that: (1) distill from cross-attention scores of an encoder-decoder model (Izacard and Grave, 2020a); (2) distill from a decoder model’s LM probability (Shi et al., 2023); and (3) train using contrastive learning (Izacard et al., 2021; Wang et al., 2022, 2023). We provide easy-to-use Python scripts for each trainer, where all training hyperparameters can be specified in a single command line.

3.3 Train Generative Models

Besides improving retrievers, using better generative models can more effectively incorporate retrieved passages. To this end, our toolkit provides: 1) direct support for many open-source generative models, and 2) various training algorithms to finetune these models to improve their task-specific performance.

Supported Models We support all huggingface (Wolf et al., 2020) decoder-only models such as LLaMA-2 (Touvron et al., 2023), and all T5 based encoder-decoder models such as FLAN-T5 (Chung et al., 2022). The former is compatible with our supervised trainers, and the latter is compatible with our fusion-in-decoder trainers.

Trainers We implement supervised fine-tuning trainers that concatenate input queries with ground-truth or retrieved passages, and fusion-in-decoder trainers that process retrieved passages in parallel. This includes trainers that: (1) supervised finetune a decoder using ground-truth $\langle q, a, p \rangle$ pairs (Lewis et al., 2021); (2) supervised finetune a decoder with a frozen retriever (Guu et al., 2020); and (3) train

an encoder-decoder with fusion-in-decoder training (Izacard and Grave, 2020b; Izacard et al., 2022). We provide easy-to-use Python scripts for each trainer, where all training hyperparameters can be specified in a single command line.

3.4 Assemble an RQA System

Given a retriever and a generative model, users can now assemble an end-to-end RQA system. Similar to frameworks such as LlamaIndex, LOCALRQA uses a modular design to support arbitrary combinations of retrievers, generative models, as well as user-defined modules (see ?? for more details), such as safety filters and decision planners (Kim et al., 2023; Peng et al., 2023). For a quick start, users can use ready-made RQA pipelines to assemble a system within five lines of code (Listing 1). These built-in pipelines support: **retrievers** available on huggingface, retrievers trained from Section 3.2, BM25 (Robertson and Zaragoza, 2009), and OpenAI embedding models (OpenAI, 2022a); **generative models** available on huggingface, models trained from Section 3.3, and OpenAI models such as ChatGPT (OpenAI, 2022b).

Alternatively, a user can also customize an RQA pipeline by implementing new/modifying existing modules. As an example, we provide an implementation of a (dummy) safety filter added to the SimpleRQA pipeline in Listing 2 in Appendix. In Listing 2, the DontKnowSafetyFilter module will ignore the answers generated by the previous components of the SimpleRQA modules, and always return “I don’t know.” as an answer.

In general, users can easily add new modules to an existing pipeline by: 1) implementing a class that inherits from Component, which requires defining a run method and run_input_keys, and 2) append the module to the components field. Alternatively, researchers can create a fully customized pipeline by inheriting from the RQAPipeline class. For more documentation and examples, please refer to our GitHub pages.

3.5 Evaluate an RQA System

Given an RQA system, LOCALRQA implements many automatic evaluation metrics to help users measure their system’s performance. This can be used by researchers to compare their system’s performance against prior work, or by developers to find the most cost-effective models/training methods suitable for their applications. We provide scripts to automatically evaluate the perfor-

Listing 1 Assembling an RQA system.

```

1 from local_rqa import ...
2 ### pre-built RQA Pipeline
3 rqa = SimpleRQA.from_scratch(
4     database_path="db_path/",
5     embedding_model_name_or_path="...",
6     qa_model_name_or_path="...",
7 )
8 response = rqa.qa(
9     batch_questions=['What is ...?'],
10    batch_dialogue_session=[
11        DialogueSession()
12    ],
13 )

```

mance of *any RQA system* that inherits from the RQAPipeline class. These scripts will also save the evaluation results in a JSONL file, which can be used to further obtain human evaluation using our serving methods (see Section 3.6). We describe the supported automatic metrics below.

Retrieval To test the performance of a retriever, we provide an evaluation script that measures: (1) Recall@k and nDCG@k score, and (2) runtime. Recall and nDCG scores are often used in information retrieval benchmarks such as BEIR (Thakur et al., 2021) and MTEB (Muennighoff et al., 2022). Runtime is important for real-world applications.

End-to-End To test the end-to-end performance of an RQA system, we provide an automatic evaluation script that measures: (1) retrieval performance such as Recall@k; (2) generation performance such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004) and GPT-4 Eval (Zheng et al., 2023a; Liu et al., 2023); and (3) end-to-end metrics such as runtime. BLEU and ROUGE scores are often used in open-ended generation tasks such as machine translation and summarization. GPT-4 Eval is a recent method using GPT-4 (OpenAI, 2023) to evaluate the quality of model-generated responses (Liu et al., 2023; Zheng et al., 2023a).

3.6 Deploy an RQA System

Finally, researchers and developers may want to showcase their RQA systems to the public, or to collect human feedback to further improve their systems using techniques such as RLHF (Ouyang et al., 2022). We provide: (1) support for efficient retrieval and LLM inference acceleration methods to reduce latency during interactive chats, and (2)

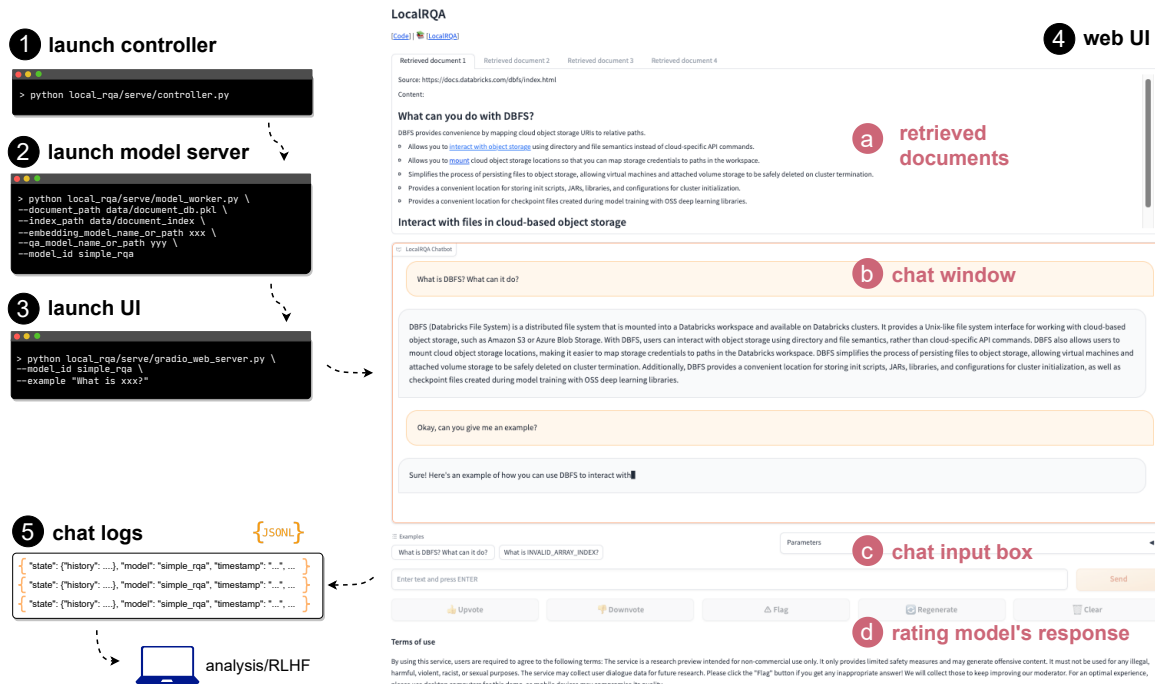


Figure 3: Researchers can launch an interactive chat page with LOCALRQA using three commands. LOCALRQA uses a model controller back-end (Zheng et al., 2023a) to handle load-balancing. Chat histories are automatically saved for researchers to conduct further analysis or model training.

implementations to easily launch an interactive chat webpage or a static evaluation webpage, given a user-built RQA system.

Acceleration Frameworks To speed up document retrieval, we support FAISS (Johnson et al., 2019), a library for efficient similarity search across billion-scale document datasets. To speed up LLM inference, we support Text Generation Inference (TGI, Huggingface (2023)), vLLM (Kwon et al., 2023), and SGLang (Zheng et al., 2023b). These inference acceleration frameworks support many decoder architectures such as LLaMA-2 and encoder-decoder architectures such as FLAN-T5.

Interactive UIs We provide (1) a static evaluation webpage where users directly evaluate the quality of pre-generated responses (e.g., computed from a test set); and (2) an interactive chat webpage where users can chat with a system and rate the correctness and helpfulness of each response. Both web interfaces can be easily launched with our toolkits, which not only support a variety of models (see Section 3.4) but also integrate with acceleration frameworks mentioned in the previous paragraph. See Figure A1 for an example of the human evaluation page, and Figure 3 for the interactive chat page.

4 Applications

To showcase our toolkit, we built two RQA systems using data scraped from Databricks and Faire’s online documentations (under consent). Databricks provides the world’s first data intelligence platform powered by generative AI, providing products that facilitate building, sharing, and maintaining data at scale. Faire is an online wholesale marketplace that connects independent retailers and brands around the world. Since the documents we obtained include many company/product-specific details, we believe this is an ideal use case for RQA systems.

First, we describe the documentation datasets we collected in Section 4.1. Then, we describe our model training, baselines, and evaluation procedures in Section 4.2, Section 4.3, and Section 4.4. Finally, we present our main results in Section 4.5.

4.1 Datasets

Databricks We use data provided by Databricks’ technical team, which includes documentations such as API references and technical tutorials from docs.databricks.com and kb.databricks.com. After applying our data processing scripts, we obtain a dataset of 11,136 passages with a maximum length of 400 tokens. See Appendix E for examples of preprocessed documents.

Retriever	Generator	Databricks				Faire			
		Retrieval		Generation		Retrieval		Generation	
		Recall@1	Recall@4	ROUGE-L	GPT4-Acc	Recall@1	Recall@4	ROUGE-L	GPT4-Acc
text-ada-002	GPT-3.5-turbo	47.36	<i>67.11</i>	46.47	<i>86.84</i>	44.00	<i>76.00</i>	35.59	81.33
text-ada-002	GPT-4-turbo	47.36	<i>67.11</i>	36.62	89.47	44.00	<i>76.00</i>	32.55	86.67
Contriever (DCA)	FastChat-T5-3B (FiD)	34.21	50.00	20.20	19.73	26.67	65.33	23.64	44.00
Contriever (RPG)	StableLM-3B (SFT)	28.94	53.94	43.30	52.63	34.66	68.00	47.53	68.00
Contriever (CTL)	StableLM-3B (SFT)	28.94	60.52	45.86	61.33	<i>42.66</i>	69.33	46.27	72.00
Contriever (CTL)	Vicuna-7B (SFT)	28.94	60.52	42.95	72.37	<i>42.66</i>	69.33	45.72	75.68
E5 (CTL)	Vicuna-7B (SFT)	34.21	77.63	45.35	73.68	40.00	<i>76.00</i>	46.16	<i>76.00</i>
Contriever (CTL)	Vicuna-7B (SwR)	28.94	60.52	44.10	60.53	<i>42.66</i>	69.33	48.72	<i>76.00</i>
E5 (CTL)	Vicuna-7B (SwR)	34.21	77.63	50.02	69.33	40.00	<i>76.00</i>	47.98	72.00
E5 (CTL)	Starling-7B (SFT)	34.21	77.63	42.06	72.36	40.00	<i>76.00</i>	46.32	78.67
E5 (CTL)	Mistral-7B (SFT)	34.21	77.63	<i>51.56</i>	80.26	40.00	<i>76.00</i>	49.63	77.33
BGE (CTL)	Starling-7B (SFT)	<i>39.47</i>	77.63	51.67	76.32	37.33	77.33	<i>50.64</i>	86.67
BGE (CTL)	Mistral-7B (SFT)	<i>39.47</i>	77.63	49.49	77.63	37.33	77.33	51.18	<i>84.00</i>

Table 2: Retrieval-augmented QA systems locally trained and tested using the LOCALRQA framework. Training algorithm used is denoted as “model(*trainer name*)”. All generation results use the top-4 passages retrieved. *GPT4-Acc* is GPT-4 evaluation of whether the generated answer is correct. Best is highlighted in **bold**, and runner-up is highlighted in *gray*.

Faire We first crawled guides and FAQ documents from faire.com/support, and then processed the data to only keep raw texts (e.g., removing image hyperlinks). Similar to Databricks, we then apply the data processing scripts and obtain a dataset of 1,758 passages. See Appendix F for examples of preprocessed documents.

Since both datasets only contain document passages p , we use LOCALRQA to generate $\langle q, a, p \rangle$ pairs for training and testing. See Appendix H for more details.

4.2 Models and Training Algorithms

LOCALRQA supports a large variety of models and training algorithms. To demonstrate the flexibility of our toolkit, we experiment with all available trainers and the most capable open-source models.

Retrievers We consider the best open-source encoder models according to the MTEB benchmark (Muennighoff et al., 2022) as of Jan 24, 2024. However, as these models vary *greatly* in capability and size, for simplicity we use the best models of similar sizes. This includes E5-base (Wang et al., 2022), Contriever-base (Izacard et al., 2021), and BGE-base (Xiao et al., 2023). We also consider all trainers in our toolkit including: (1) distilling from cross-attention scores, denoted as *DCA*; (2) distilling LM probability, denoted as *RPG*; and (3) training with contrastive learning, denoted as *CTL*.

Generators We consider the best generator models according to the Chatbot Arena leaderboard (Zheng et al., 2023a) as of Jan 24, 2024. Since training LLMs is time and resource intensive, we

focus on the best open-source models up to 7B parameters. This includes encoder-decoder models such as FastChat-T5-3B (Zheng et al., 2023a; Chung et al., 2022), and decoder-only models such as: StableLM-3B (Tow et al., 2023), Vicuna-7B (Chiang et al., 2023), Starling-7B (Zhu et al., 2023), and Mistral-7B (Jiang et al., 2023). We also consider all trainers in our toolkit including: (1) supervised fine-tuning, denoted as *SFT*; (2) supervised fine-tuning with a frozen retriever, denoted as *SwR*; and (3) fusion-in-decoder training, denoted as *FiD*.

4.3 Baselines

Since LOCALRQA features developing new RQA systems *locally*, we compare against the most powerful models accessible *remotely*. This include using text-ada-002 (OpenAI, 2022a) as the retriever, and prompting GPT-3.5-turbo (ChatGPT) and GPT-4-turbo as the generative models.

4.4 Metrics

We present a subset of automatic evaluation metrics from LOCALRQA, and also include human evaluations on the best-performing models using UIs from Section 3.6. To measure retrievers’ performance, we report Recall@1 and Recall@4 which are commonly used in information retrieval (Thakur et al., 2021; Muennighoff et al., 2022). To measure the final generation performance, we report ROUGE-L (Lin, 2004) and GPT-4 Eval (Zheng et al., 2023a; Liu et al., 2023), which are used in open-domain generation tasks (Zheng et al., 2023a). For the best models, we additionally perform human evaluation and report the accuracy of the generated answers.

Retriever	Generator	Databricks	Faire
		Human-Acc	Human-Acc
text-ada-002	GPT-3.5-turbo	78.00	86.00
text-ada-002	GPT-4-turbo	84.00	88.00
E5 (CTL)	Mistral-7B (SFT)	78.00	90.00
BGE (CTL)	Starling-7B (SFT)	80.00	88.00

Table 3: Comparing the best LOCALRQA-trained models in Table 2 against ChatGPT and GPT-4. *Human-Acc* is authors’ judgement of whether the final answer is correct. We use the first 50 test samples for evaluation.

4.5 Main Results

Table 2 presents our non-exhaustive combination of retrievers and generators trained and tested using LOCALRQA. First, we find contrastive learning (CTL) most effective for training retrievers. We believe this is because CTL was also used to pre-train all the encoders we investigated (Izacard et al., 2021; Wang et al., 2022; Xiao et al., 2023). We also find that simple supervised fine-tuning (SFT) with gold $\langle q, a, p \rangle$ pairs is suitable for generators, given the answers in the training data are generated only using the gold passage.

Next, we find using more powerful retriever models (BGE-base and E5-base) and generator models (Mistral-7B and Starling-7B) improves Recall@4 and GPT4-Acc score. This is understandable since these models have a better base performance. We also find that ROUGE-L does not correlate well with GPT4-Acc (or our human evaluation). This is consistent with Cohan and Goharian (2016); Nekvinda and Dušek (2021), since open-ended generations are inherently difficult to evaluate using automatic metrics.

Lastly, we use the best models from Table 2 according to GPT4-Acc scores, and further validate their performance with our human evaluation UI (see Section 3.6). In Table 3, we find the best local models reach a similar performance as the OpenAI’s baselines, despite being only 7B in size (see Appendix G for some examples). These results underscores the effectiveness of our toolkit in training and developing cost-effective RQA systems.

5 Conclusion

We present LOCALRQA, a Python-based toolkit designed to help users develop novel retrieval-augmented QA systems. Different from existing frameworks such as LlamaIndex and LangChain, our toolkit features a wide collection of training algorithms, evaluation metrics, and deploy-

ment methods to help users quickly develop cost-effective RQA systems. Strong results using models and training algorithms from recent research pave the way for future work to explore RQA methods in both practical and academic settings.

6 Limitations and Future Work

Model Size We performed all of our experiments using a single A100 80G GPU, and investigated a large combination of model choices and training methods. Therefore, we considered the best-performing models up to 7B parameters due to time and resource concerns. We believe experimenting with larger, more capable models could further improve the systems’ performance, and we leave this for future work.

More Training Algorithms Besides providing tools to help users easily build an RQA system, LOCALRQA features a collection of training algorithms and evaluation methods curated from latest research. However, this collection is non-exhaustive (Zhong et al., 2022; Asai et al., 2022; Min et al., 2023; Asai et al., 2023; Ram et al., 2023). We commit to add support for more models, training algorithms, and testing methods to reflect ongoing advancements in RQA research.

Compute Requirement LOCALRQA features methods to help users develop novel RQA systems *locally*. Compared with using paid services such as OpenAI’s text-ada-002 and GPT-4, this approach is less expensive but requires access to compute resources (e.g., GPUs). To make our toolkit more accessible, we not only support open-source models from huggingface of various sizes, but also support using “remote” models such as OpenAI’s ChatGPT and GPT-4.

7 Ethical Considerations

Our work describes a toolkit that can be used to help researchers develop new RQA systems. LOCALRQA offers a suite of tools, starting from data generation to locally training, testing, and serving an RQA system. While most toolkits are not designed for unethical usage, there is often potential for abuse in their applications. In our demo (Section 4), we apply our toolkit to train RQA systems based on documentations obtained from two companies’ website, Databricks and Faire. However, since our toolkit can be used with any kind of data, it is possible to use it for unethical tasks, such

as scamming and generating harmful responses (Gehman et al., 2020; Welbl et al., 2021). We do not condone the use of LOCALRQA for any unlawful or morally unjust purposes.

8 Acknowledgement

We thank Xiangrui Meng and Quinn Leng from Databricks, and Wenhao Liu from Faire for their valuable support and discussions.

References

- Abubakar Abid, Ali Abdalla, Ali Abid, Dawood Khan, Abdulrahman Alfozan, and James Zou. 2019. [Gradio: Hassle-free sharing and testing of ML models in the wild.](#)
- Akari Asai, Timo Schick, Patrick Lewis, Xilun Chen, Gautier Izacard, Sebastian Riedel, Hannaneh Hajishirzi, and Wen tau Yih. 2022. [Task-aware retrieval with instructions.](#)
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection.](#)
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. 2018. [Ms marco: A human generated machine reading comprehension dataset.](#)
- Harrison Chase. 2022. [LangChain.](#)
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality.](#)
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models.](#)
- Arman Cohan and Nazli Goharian. 2016. [Revisiting summarization evaluation for scientific articles.](#) In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 806–813, Portorož, Slovenia. European Language Resources Association (ELRA).
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. [RealToxicityPrompts: Evaluating neural toxic degeneration in language models.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, Online. Association for Computational Linguistics.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [Realm: Retrieval-augmented language model pre-training.](#)
- Huggingface. 2023. [Large language model text generation inference.](#)
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. [Unsupervised dense information retrieval with contrastive learning.](#)
- Gautier Izacard and Edouard Grave. 2020a. [Distilling knowledge from reader to retriever for question answering.](#)
- Gautier Izacard and Edouard Grave. 2020b. [Leveraging passage retrieval with generative models for open domain question answering.](#)
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. [Atlas: Few-shot learning with retrieval augmented language models.](#)
- Peter Izsak, Moshe Berchansky, Daniel Fleischer, and Ronen Laperdon. 2023. [fastRAG: Efficient Retrieval Augmentation and Generation Framework.](#)
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. [Mistral 7b.](#)
- Jeff Johnson, Matthijs Douze, and Herv  J gou. 2019. [Billion-scale similarity search with GPUs.](#) *IEEE Transactions on Big Data*, 7(3):535–547.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension.](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas O uz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen tau Yih. 2020. [Dense passage retrieval for open-domain question answering.](#)
- Jinhwa Kim, Ali Derakhshan, and Ian G. Harris. 2023. [Robust safety classifier for large language models: Adversarial prompt shield.](#) *ArXiv*, abs/2311.00172.

- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#).
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Jerry Liu. 2022. [LlamaIndex](#).
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Sewon Min, Weijia Shi, Mike Lewis, Xilun Chen, Wen tau Yih, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2023. [Nonparametric masked language modeling](#).
- Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2022. [Mteb: Massive text embedding benchmark](#). *arXiv preprint arXiv:2210.07316*.
- Tomáš Nekvinda and Ondřej Dušek. 2021. [Shades of bleu, flavours of success: The case of multiwoz](#).
- OpenAI. 2022a. [New and improved embedding model](#).
- OpenAI. 2022b. [OpenAI: Introducing ChatGPT](#).
- OpenAI. 2023. [GPT-4 technical report](#).
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars Liden, Zhou Yu, Weizhu Chen, and Jianfeng Gao. 2023. [Check your facts and try again: Improving large language models with external knowledge and automated feedback](#).
- Malte Pietsch, Timo Möller, Bogdan Kostic, Julian Risch, Massimiliano Pippi, Mayank Jobanputra, Sara Zanzottera, Silvano Cerza, Vladimir Blagojevic, Thomas Stadelmann, Tanay Soni, and Sebastian Lee. 2019. [Haystack: the end-to-end NLP framework for pragmatic builders](#).
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [In-context retrieval-augmented language models](#).
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: Bm25 and beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023. [Replug: Retrieval-augmented black-box language models](#).
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. [BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).

- Jonathan Tow, Marco Bellagente, Dakota Mahan, and Carlos Riquelme. 2023. [Stablelm 3b 4e1t](#).
- Andrew Trotman, Antti Puurula, and Blake Burgess. 2014. [Improvements to bm25 and language models examined](#). In *Proceedings of the 19th Australasian Document Computing Symposium, ADCS '14*, page 58–65, New York, NY, USA. Association for Computing Machinery.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. [Text embeddings by weakly-supervised contrastive pre-training](#).
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2023. [SimLM: Pre-training with representation bottleneck for dense passage retrieval](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2244–2258, Toronto, Canada. Association for Computational Linguistics.
- Johannes Welbl, Amelia Glaese, Jonathan Uesato, Sumanth Dathathri, John Mellor, Lisa Anne Hendricks, Kirsty Anderson, Pushmeet Kohli, Ben Coppin, and Po-Sen Huang. 2021. [Challenges in detoxifying language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2447–2469, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Huggingface’s transformers: State-of-the-art natural language processing](#).
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2023. [C-pack: Packaged resources to advance general chinese embedding](#).
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023a. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).
- Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Jeff Huang, Chuyue Sun, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E. Gonzalez, Clark Barrett, and Ying Sheng. 2023b. [Efficiently programming large language models using sglang](#).
- Zexuan Zhong, Tao Lei, and Danqi Chen. 2022. [Training language models with memory augmentation](#).
- Banghua Zhu, Evan Frick, Tianhao Wu, Hanlin Zhu, and Jiantao Jiao. 2023. [Starling-7b: Improving llm helpfulness and harmlessness with rlaif](#).

A Comparison against Existing Toolkits

Many existing toolkits, such as Haystack, LangChain, and LlamaIndex help *users* quickly build an RQA system (Pietsch et al., 2019; Chase, 2022; Liu, 2022). However, these frameworks provide very little support for *researchers* to train, test, and serve their RQA systems using recent advances in retrieval-augmented QA research. For instance, LlamaIndex only includes basic supervised finetuning methods to train “LLaMA-2 for better text-to-SQL”, or finetune “GPT-3.5-turbo to distill GPT-4”². We provide three different retriever training algorithms (Section 3.2) and three different generator training algorithms (Section 3.3). We highlight our main contributions compared to other existing toolkits in Table 1.

B Supported Data Formats

LOCALRQA support data coming from many different sources, by providing integration with frameworks such as LangChain (Chase, 2022) and LlamaIndex (Liu, 2022). This not only includes loading data of different formats (e.g., JSON, HTML, PDF files), but also data from different locations (e.g., Google Drive, S3 bucket, websites and more). For more details on data loading, please refer to our project website.

C Details on Data Generation

LOCALRQA provides data generation scripts that can be used to create questions q from a set of documents p , and answers from a set of $\langle q, p \rangle$ pairs. These scripts can also be easily modified to use: 1) custom prompts to generate a question or answer, and 2) custom filtering functions to use a subset of the documents for question/answer generation.

Question Generation Given a set of documents, LOCALRQA first creates a set of gold passages by sampling. Since contrastive learning (Section 3.2) benefits from using hard negative passages (related passages but does not contain the answer), we also sample nearby passages as hard negatives. This is achieved by first organizing all passages according to their source s_i (e.g., URL or title):

$$\{p_0^{s_0}, p_1^{s_0}, \dots, p_n^{s_0}, p_0^{s_1}, p_1^{s_1} \dots\}$$

and then sample from $\{p_j^s\}_{j \neq i}$ as hard negatives for p_i^s . Next, an LLM of choice (e.g., ChatGPT) is

²<https://docs.llamaindex.ai/en/stable/optimizing/fine-tuning/fine-tuning.html>, visited on Feb 12, 2024.

prompted to generate k questions given a sampled gold passage. To filter duplicate questions, LOCALRQA uses ROUGE-L score (Lin, 2004) to remove questions with high word overlap with others.

Answer Generation Given a set of $\langle q, p \rangle$ pairs, LOCALRQA prompts an LLM of choice (e.g., GPT-4) to generate answers conditioned on the question q and the gold passage p .

See Appendix E for examples on how to customize the data generation scripts and Appendices E and F for examples commands.

D More Details on Serving RQA Pipelines

LOCALRQA offers two serving methods: 1) an interactive chat page where users can chat with an RQA system while also providing ratings for each generated response, and 2) a static evaluation page where users directly evaluate the quality (e.g., accuracy, helpfulness, harmlessness) of the pre-generated response. The front-end UIs are created using Gradio (Abid et al., 2019), and the model back-end (for interactive chat) is modified from Zheng et al. (2023a). We provide an example of using each serving method in Figure 3 and Figure A1, respectively.

E More Details on Databricks Demo

Collected Documents We use documents provided by Databrick’s technical team, which are already cleaned and parsed into markdown format. We present an example in Table A1.

QA Generation We generate questions and answers using the data generation scripts in LOCALRQA. We first customize the prompts and filtering functions in order to obtain high-quality questions based mostly on technical tutorials rather than version release notes³. This only requires:

1. creating a new python script with from `scripts.data.doc_to_q` import *
2. defining a custom prompt and filter function
3. assigning `filter_fn=your_filter_fn` and `doc2q_prompt=YOUR_PROMPT` in the imported main function

For a complete example, please refer to our `scripts/data/doc_to_q_databricks.py`. Finally, we run the above question generation script

³<https://docs.databricks.com/en/release-notes/runtime/index.html>

	Databricks Example Document	Faire Example Document
content	<pre> ## Step 3: Configure access to the 'default.people10m' table Enable the user you created in [Step 1](#add-a-user) to access the 'default.people10m' table you created in [Step 3](#create- a-table). You can configure access using [Data Explorer](#data-explorer) or [SQL editor](#sql-editor). ### Data explorer - Click the <Data Icon> Data in the sidebar. - In the drop-down list at the top right ...</pre>	<pre>*Please note these settings are applicable to all products, if you'd like to change these settings on a product level, you can do so in the 'Catalog Synchronization' tab. # API Token This is where you will enter the API token provided by Faire. Once entered, you will also receive a message to confirm that the connection is confirmed between Faire and Prestashop. # Catalog Import for Wholesale This is ...</pre>
metadata	<pre>"source": "https://docs.databricks.com/.../admin-set-up-user -to-query-table.html", "seq_num": 574, "description": "", "text": "..."</pre>	<pre>"source": "https://www.faire.com/support/artic- les/8726114634779", "seq_num": 426, "subtitle": "Prestashop Integration with Faire", "title": "Prestashop Integration with Faire"</pre>

Table A1: Example documents collected from Databrick’s documentation pages and Faire’s support pages. Omitted details are indicated as “...”.

	Databricks	Faire
Train	1,185	575
Validation	74	74
Test	76	76

Table A2: Number of $\langle q, a, p \rangle$ pairs used in training, validation, and testing. During testing, we use all available documents to measure an RQA system’s retrieval and generation performance.

followed by our answer generation script to obtain a collection of $\langle q, a, p \rangle$ pairs. We used ChatGPT (OpenAI, 2022b) and GPT-4-turbo (OpenAI, 2023) to generate questions and answers, respectively.

Model Training To show the flexibility of LOCALRQA training, we present at least one run of using trainer in our main experiments Table 2. All trainings are performed on a single A100 80G GPU. Please refer to our GitHub for more details on training hyperparameters and other command-line arguments used. We note that we did not rigorously hyperparameter-tune each model and trainer combination due to the large number of experiments to perform. We believe results in Table 2 may be further improved if a hyperparameter search is ran for each method.

F More Details on Faire Demo

Collected Documents We contacted Faire’s Sales team and crawled documents from [faire.com/support](https://www.faire.com/support) according to their suggestions. We only kept raw texts by removing all hyperlinks for images and other websites. We present an example in Table A1.

QA Generation Since document data from Faire include simpler guides and QAs compared to Databricks, we find using the default generation script in LOCALRQA sufficient to obtain high-quality questions and answers. Therefore, we simply ran `scripts/data/doc_to_q.py` to generate questions, and `scripts/data/docq_to_a.py` to generate answers. Similar to Databricks, we used ChatGPT (OpenAI, 2022b) and GPT-4-turbo (OpenAI, 2023) to generate questions and answers, respectively.

Model Training Similar to the training process implemented in Databricks, we conduct experiments across various trainers and model choices. See Table 2 for an example. We note that we did not rigorously hyperparameter-tune each model and trainer combination due to the large number of experiments to perform. We believe results in Table 2 may be further improved if a hyperparameter search is ran for each method.

G Example Model Answers

We provide example generated answers and retrieved documents in [Table A3](#) and [Table A4](#) for Databricks and Faire, respectively. In [Table A3](#), the weaker retrievers fail to find the relevant document given the user’s query. As a result, the generative models also fail to provide the correct answer. This shows the benefit of training and using more accurate retrievers. In [Table A4](#), despite all methods retrieved the correct document, the weaker generative models failed to incorporate all relevant information in its response. This underscores the importance to fine-tune capable generative models to improve the RQA system’s performance.

H Dataset Statistics

We follow the procedure outlined in [Appendix E](#) and [Appendix F](#) to obtain $\langle q, a, p \rangle$ pairs from documentation data. For Databricks, we first use our sampling algorithm to obtain 600 passages for training, and 75 for validation and for testing. For Faire, we sampled 300 passages for training as Faire has fewer documents, and 75 for validation and for testing. We then use ChatGPT ([OpenAI, 2022b](#)) as the LLM to generate two questions for each sampled passage, and GPT-4-turbo ([OpenAI, 2023](#)) to generate an answer for each $\langle q, p \rangle$ pair (see [Appendix C](#) for more details). This results in a training, evaluation, and testing dataset with statistics reported in [Table A2](#).

Databricks RQA Example	
User: How many catalogs can a metastore have in Databricks?	
text-ada-002+ GPT-3.5-turbo	In Databricks, a metastore can have multiple catalogs. There is no specific limit on the number of catalogs that can be associated with a metastore. Catalogs are used to organize [...] Users can see all catalogs on which they have been assigned the ‘USE CATALOG’ data permission.
text-ada-002+ GPT-4-turbo	A metastore in Databricks can contain multiple catalogs. [...] There isn’t a specified limit to the number of catalogs a metastore can have mentioned in the documentation, which suggests that users can create multiple catalogs as needed to manage their data in an organized manner [...].
E5(CTL)+ Mistral-7B(SFT)	As of August 25, 2022 , a metastore can have up to 1000 catalogs .
BGE(CTL)+ Starling-7B(SFT)	As of August 25, 2022 , a metastore in Databricks can have up to 1000 catalogs .
reference document:	<p># <UC> GA release note</p> <p>**August 25, 2022**</p> <p><UC> is now generally available on <Databricks>.</p> <p>**This article describes <UC> as of the date of its GA release.** It focuses primarily on the features and updates [...]</p> <p>As of August 25, 2022</p> <ul style="list-style-type: none"> - Your <Databricks> account can have only one metastore per region - A metastore can have up to 1000 catalogs. - A catalog can have up to 10,000 schemas. - A schema can have up to 10,000 tables. <p>For current <UC> quotas, see [...].</p>

Table A3: Example answers generated by different RQA systems. Questions come from the test dataset for Databricks, which is described in [Appendix H](#). Due to space limitations, we display the most relevant document amongst all documents retrieved. In this example, text-ada-002 failed to retrieve the most relevant documents, causing GPT-3.5-turbo and GPT-4-turbo to produce incorrect answers.

Faire RQA Example	
User: How should I crop my feature image for it to be approved?	
text-ada-002+ GPT-3.5-turbo	To ensure that your feature image is approved, you should crop it to a square shape. The image must fill a 1:1 ratio , meaning that the width and height should be equal. By cropping your image to a square, you can meet this requirement and increase the chances of approval.
text-ada-002+ GPT-4-turbo	To ensure your feature image gets approved, you should follow these cropping and content guidelines: <ol style="list-style-type: none"> **Square Crop Ratio**: Crop your image to a perfect square. The aspect ratio should be 1:1, which means [...]. **No Overlaid Text or Logos**: Avoid adding any text or logos [...]. **Simplicity is Key**: Steer clear of using collages or images that are too busy. A simple, singular focus on the product [...]. **Visual Context**: Provide context by showing the product in use or in a setting where it is naturally found [...]. **Avoid Direct Eye Contact**: When including models, aim to avoid direct eye contact with the camera. [...]
E5(CTL)+ Mistral-7B(SFT)	To have your feature image approved, you should ensure that it is cropped to fit a square aspect ratio, which must fill a 1:1 ratio . Additionally, it should not include any words or logos . The recommended approach is to choose a high-resolution lifestyle image that showcases your best-selling product(s) in use, such as a cheese knife with cheese and grapes on a cutting board. Also, you should provide some visual context in [...]
BGE(CTL)+ Starling-7B(SFT)	To have your feature image approved, you should crop it to a square shape that fills a 1:1 ratio .
reference document:	<p>## Feature image guidelines</p> <p>Crop your image to a square. It must fill a 1:1 ratio.</p> <p>Don't include overlaid text or logos.</p> <p>Avoid collages and keep it simple. Collages cause clutter in a space already crowded with many other brands. Choose a single image with a clear product focus. Simple, uncluttered often stand out more, as they provide rest to the shopper's eye that may be weary from browsing many pages of content.</p> <p>Provide visual context (avoid plain white backgrounds). While it's true that great packaging design attracts people all by itself, showing your product in use helps get people excited about your brand and helps tell a more visual story about what you have to offer.</p> <p>Avoid showing models with direct eye contact. [...]</p>

Table A4: Example answers generated by different RQA systems. Questions come from the test dataset for Faire, which is described in [Appendix H](#). Due to space limitations, we display the most relevant document amongst all documents retrieved. In this example, all methods retrieved the correct document, but GPT-3.5-turbo and Starling-7B failed to include details other than "fill a 1:1 ratio" (c.f. [Table 3](#)).

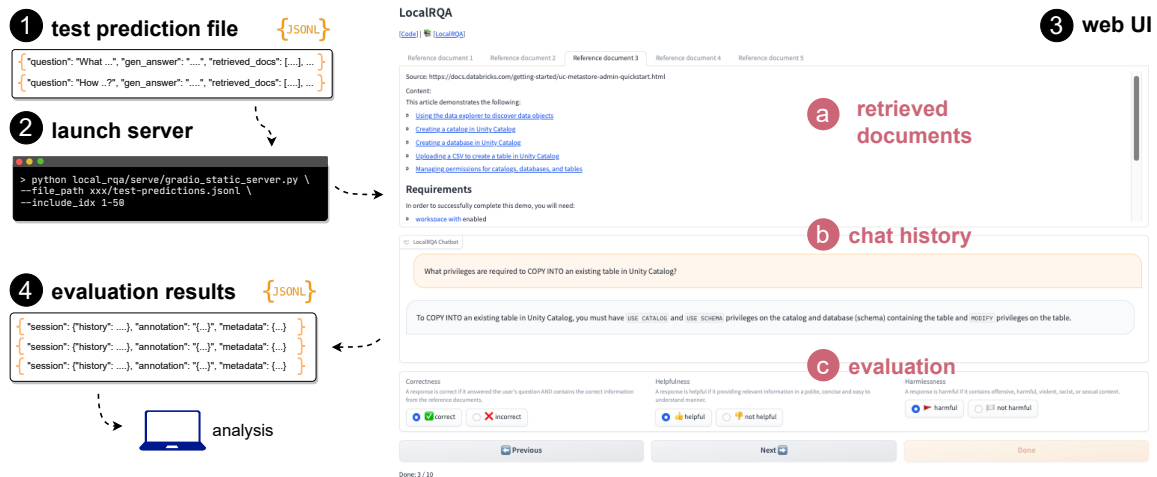


Figure A1: Researchers can launch a human evaluation page using LOCALRQA in a single command line. Given a prediction file (see Section 3.5), LOCALRQA launches a web server that allows other users to evaluate the quality of pre-generated responses. Evaluation results are automatically saved for researchers to conduct further analysis.

Listing 2 Adding a custom component to the SimpleRQA pipeline. In general, researchers can easily extend an existing pipeline by editing the `.components` field, or create new pipelines by inheriting from the `RQAPipeline` class.

```

1 from local_rqa import ...
2 ### pre-built RQA pipeline
3 rqa = SimpleRQA.from_scratch(
4     database_path="db_path/",
5     embedding_model_name_or_path="...",
6     qa_model_name_or_path="...",
7 )
8
9 ### custom module: safety filter
10 class DontKnowSafetyFilter(Component):
11     run_input_keys = [
12         "batch_questions",
13         "batch_source_documents",
14         "batch_dialogue_session",
15         "batch_answers",
16     ]
17
18     def run(self, *args, **kwargs):
19         return RQAOutput(
20             batch_answers="I don't know.",
21             batch_source_documents=kwargs["batch_source_documents"],
22             batch_dialogue_session=kwargs["batch_dialogue_session"],
23         )
24
25 rqa.components.append(DontKnowSafetyFilter())
26
27 ### run QA!
28 rqa.qa(...) # output: "I don't know."

```