

UKP-SQUARE: An Online Platform for Question Answering Research

Tim Baumgärtner,¹ Kexin Wang,¹ Rachneet Sachdeva,¹ Max Eichler,¹ Gregor Geigle,¹ Clifton Poth,¹ Hannah Sterz,¹ Haritz Puerto,¹ Leonardo F. R. Ribeiro,¹ Jonas Pfeiffer,¹ Nils Reimers,² Gözde Gül Şahin,³ Iryna Gurevych¹

¹Ubiquitous Knowledge Processing Lab, Department of Computer Science, Technical University of Darmstadt
²Hugging Face, ³Koç University Computer Science and Engineering Department

www.ukp.tu-darmstadt.de

Abstract

Recent advances in NLP and information retrieval have given rise to a diverse set of question answering tasks that are of different formats (e.g., extractive, abstractive), require different model architectures (e.g., generative, discriminative), and setups (e.g., with or without retrieval). Despite having a large number of powerful, specialized QA pipelines (which we refer to as Skills) that consider a single domain, model or setup, there exists no framework where users can easily explore and compare such pipelines and can extend them according to their needs. To address this issue, we present UKP-SQUARE, an extensible online QA platform for researchers which allows users to query and analyze a large collection of modern Skills via a user-friendly web interface and integrated behavioural tests. In addition, QA researchers can develop, manage, and share their custom Skills using our microservices that support a wide range of models (Transformers, Adapters, ONNX), datastores and retrieval techniques (e.g., sparse and dense). UKP-SQUARE is available on <https://square.ukp-lab.de>.¹

1 Introduction

Researchers in NLP have devoted significant resources to creating more powerful machine learning models for Question Answering (QA), and collecting high-quality QA datasets. Combined with the recent breakthroughs by large pretrained language models, we have witnessed rapid progress in the field across many different kinds of QA tasks (Rogers et al., 2021).

The great variety in QA tasks has led to specialized, domain-specific models trained on a single QA format such as *multiple choice* (Lai et al., 2017) (i.e., selecting the best answer out of multiple options), *extractive* (Rajpurkar et al., 2016)

¹The code is available on <https://github.com/UKP-SQUARE/square-core>

The screenshot displays the UKP-SQUARE web interface. At the top, there is a navigation bar with the logo and menu items: 'UKP-SQUARE', 'QA', 'Explainability', 'Docs', 'Feedback', and a user profile icon. The main interface is split into two columns. The left column, '1. Select a skill', features three dropdown menus: 'OpenSQuAD-DPR - Open-Dom', 'Compare up to three skills', and 'OpenBioASQ - Open-Domain Q', with a 'None' option at the bottom. The right column, '2. Enter you question', has a text input field containing 'What are the symptoms of MERS?' and three example questions: 'Where is Berlin?', 'When was the EU created?', and 'What are the symptoms of MERS?'. Below these is an 'Ask your question' button. The bottom section shows the results for two skills: 'OpenSQuAD-DPR' and 'OpenBioASQ'. The 'OpenSQuAD-DPR' result shows a score of 23.74% and a snippet of text about the 2005 Hurricane Katrina. The 'OpenBioASQ' result shows a score of 75.02% and a snippet of text about the Middle East respiratory syndrome coronavirus (MERS CoV).

Figure 1: QA page of UKP-SQUARE. The user selects a Skill (in this case, two open-domain Skills are selected), enters a question and then receives an answer.

(i.e., finding the answer span in a context) and *abstractive* (Kocisky et al., 2018) (i.e., generating an answer that is not a contiguous span in the context). The format may influence the model architecture (e.g., discriminative objective for *multiple choice*, generative objective for *abstractive*). Additionally, systems vary with how the context is provided. It can be given by the user, or retrieved from a Datastore which is commonly referred to as *open-domain* or *retriever-reader* setup (Chen et al., 2017). The retrieval mechanism can also be chosen from a set of sparse (e.g., BM25, Robertson et al., 1994) or dense (e.g., DPR, Karpukhin et al., 2020) techniques.

The speed of progress in the field makes it essential for researchers to explore, compare, and combine these different QA components as quickly as possible to identify the strengths and weaknesses

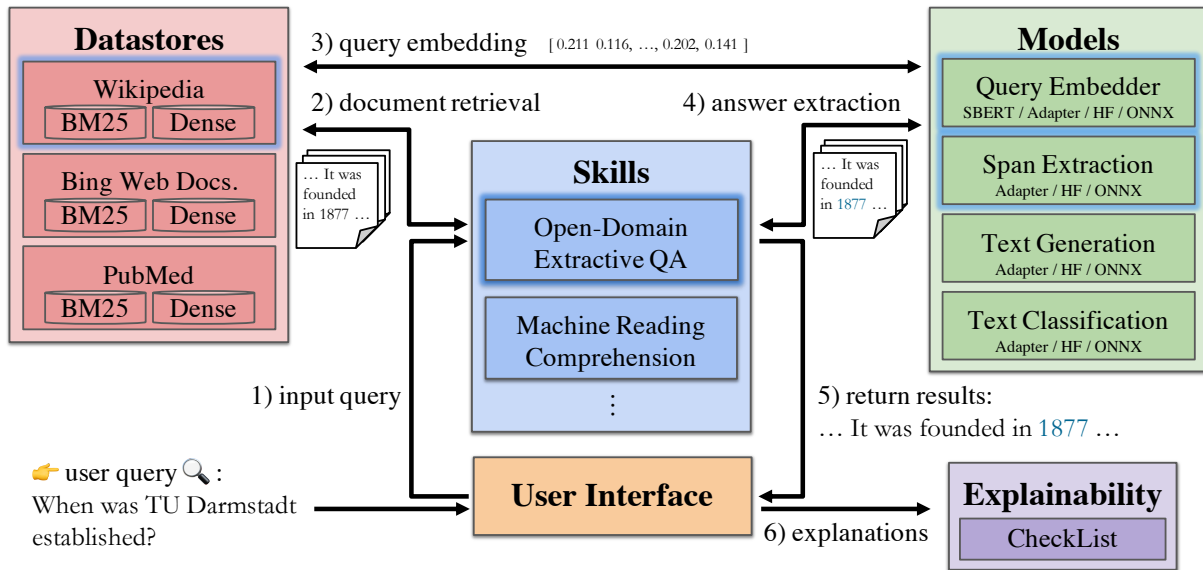


Figure 2: Overall architecture of UKP-SQUARE, illustrating an open-domain, extractive QA Skill. (1) First a user selects a **Skill** and issues a query via the **User Interface**. (2) The selected QA Skill forwards the query to the respective **Datstore** for document retrieval. (3) The **Datstore** gets the query embedding from the **Models**, uses it for semantic document retrieval and returns the top documents to the **Skill**. (4) The **Skill** sends the query and retrieved documents to the reader model for answer extraction. (5) Finally, the answers are shown to the user. (6) Optionally, the user can view the results of the predefined behavioural tests for the **Skill**.

of the current state of the art. Even though there exists a number of powerful QA systems (Dibia, 2020; Khashabi et al., 2020) and frameworks such as Haystack,² those approaches focus only on one component (e.g., retrieval, QA format, domain), hence do not allow plug-and-play of different Datastores, domains, model architectures or retrieval techniques. This considerably limits their applicability and reusability across the diverse, rapidly progressing area of QA research, making it infeasible for researchers to quickly integrate novel models and QA pipelines.

To address this gap, we introduce UKP-SQUARE, a flexible and extensible QA platform to enable users to easily implement, manage and share their custom QA pipelines, which we call **Skills**, using our microservices. As shown in Fig. 1, UKP-SQUARE also allows users to query and compare different **Skills** via an easy-to-use user interface and systematically analyze their strengths and weaknesses through integrated behavioural tests.³

2 UKP-SQUARE

The system is implemented as a modern microservice architecture using Docker containers.⁴ The

²<https://deepset.ai/haystack>

³Screenshots for adding Skills, the outputs of different QA formats and behavioural tests are shown in Appendix D.

⁴<https://docker.com>

major components are **Skills**, **Datastores**, **Models**, **Explainability** and the **User Interface**. The process flow across the components is illustrated in Fig. 2 on an open-domain, extractive QA Skill. The central component of the system is the **Skill** that specifies how a user query is processed (e.g., which QA type, retrieval mechanism, model or adapter to be used in which order). The **Skill** leverages the other services for query execution. **Datastores** hold multiple collections of documents with sparse indices, e.g., BM25 (Robertson et al., 1994) and dense indices, e.g., DPR (Karpukhin et al., 2020), allowing fast and efficient retrieval of background knowledge in an extensible way. The **Model** service hosts numerous models, combined with Adapters (Houlsby et al., 2019; Pfeiffer et al., 2020), to support a wide range of tasks such as text embedding (for queries in open-domain QA), sequence and token classification (for multiple-choice and extractive QA) and sequence-to-sequence generation (for abstractive QA). The **Explainability** component performs behavioural tests on the deployed Skills for better understanding of the models. Details of each service are provided in the following sections.

Furthermore, while we host UKP-SQUARE on our infrastructure and make it available for the community, we also provide the option to set up the system locally. Additionally, the **Datastores** and

Models services are exposed via an API⁵.

2.1 Skills

Skills define how the user query should be processed by the Datastores and Models components and how the respective answers are obtained. For question answering, this might involve retrieving background knowledge, extracting spans from context or selecting an answer from multiple choices.

Skills are not necessarily equivalent to a model trained on a dataset. Instead a Skill is more general and can use multiple models to arrive at an output. A Skill might work on a specialized domain (e.g. biomedical, movies, etc.) or a specific format (e.g. extractive, abstractive, etc.), but also combinations are possible. For example, a Skill could combine Wikipedia and a news based extractive reader model to answer factoid and news questions. The degree of specialization or generalization of a Skill is up to its developer. In UKP-SQUARE the Skill only defines the pipeline, i.e., pre-processing, information retrieval or answer extraction/generation/classification. These steps are facilitated and executed by the usage of the other components: Models (§2.2) and Datastores (§2.3).

Importantly, Skills can be added to the system by the community. They can be added privately, thereby only giving a specific user access to it, or made public, allowing everyone to use it (§3.1). This allows great flexibility in the design of question answering pipelines, keeping implementation effort and required compute low, thereby democratizing the usage of question answering models.

2.2 Models

The Models component is responsible for hosting NLP models required for document retrieval and answer extraction/generation tasks. Our platform supports a wide variety of models comprising HuggingFace (HF) Transformers (Wolf et al., 2020), Adapters, Sentence-Transformers (Reimers and Gurevych, 2019), and a limited selection of ONNX (Open Neural Network Exchange) (Bai et al., 2019) models. Specifically, the inclusion of memory-efficient adapters in our platform allows having a variety of task-specific models while maintaining storage efficiency. Moreover, for faster inference, the high performance inference engine, ONNX Runtime⁶ can be used for the ONNX mod-

⁵<https://square.ukp-lab.de/docs/>

⁶<https://github.com/microsoft/onnxruntime>

els provided on our platform.

The Models component comprises of two main services: **inference** and **management**. The inference service is responsible for loading models and getting predictions for the input queries. The management service allows the user to list, deploy, update and remove models (available on HF, Adapterhub and Sentence-Transformers) on the UKP-SQUARE platform. This allows to deploy and query models beyond the ones we already provide, for example multilingual models. To maintain a scalable architecture, we host every deployed model in its separate Docker container and use Traefik⁷ to route the user query to the specific model instance for inference. The inference service of the model API can be queried using the Skills (§2.1) as per the end-user’s requirements.

2.3 Datastores

The Datastores are responsible for storing document collections as knowledge bases of QA Skills, supporting retrieval on these collections. Each Datastore contains a collection of documents and several indices of them for retrieval. The document collections are stored by an Elasticsearch⁸ instance. Within one Datastore, the document collection is indexed by sparse or dense retrieval models.

For sparse retrieval, we use BM25 provided by the Elasticsearch instance; for dense retrieval, we use dual-encoder neural networks (Karpukhin et al., 2020; Xiong et al., 2021) with Approximate Nearest Neighbor (ANN) indexing provided by Faiss (Johnson et al., 2021). The Datastores are agnostic to the ANN methods. Among them, we use *IndexIVFScalarQuantizer* (Jégou et al., 2011) from Faiss as the default choice. For scalability, we maintain each dense-retrieval index within one Docker container and use Traefik to route the queries to the specific index. For each query using dense retrieval, the Datastores forward the query to the Models to get the query embedding (e.g., via the Query Embedder in Fig. 2 and Table 2) and then input this embedding to the ANN search for retrieving relevant documents.

As the built-in Datastores, Wikipedia⁹ with the DPR encoder (Karpukhin et al., 2020), PubMed¹⁰

⁷<https://traefik.io>

⁸<https://elastic.co>

⁹The English Wikipedia dump preprocessed by Karpukhin et al. (2020).

¹⁰From the BioASQ8 edition (Nentidis et al., 2020).

Minimum Functionality Test (MFT)- <i>Taxonomy</i>
C: There is a tiny purple box in the room.
Q: What size is the box?
Test: Check if the prediction is tiny
INVariance- <i>Robustness</i>
C: ...Newcomen designs had a duty of about 7 million, but most were closer to 5 million....
Q: What was the ideal [duty -> udty] of a Newcomen engine?
Test: Check whether the prediction changes or not.

Table 1: Examples for two most common test types. **Top:** Minimum Functionality Test (MFT), **Bottom:** Invariance Test (INV). **C:** refers to context and **Q:** is the question.

and Bing web documents¹¹ with the TAS-B encoder (Hofstätter et al., 2021) are supported. We plan to add more Datastores in the future.

2.4 Explainability

Recently, many interpretability techniques to understand black-box neural models such as influence functions and input/token attribution methods (Madsen et al., 2021) have been introduced. Most of these techniques provide only local explanations and require access to the back-propagation function. One exception is *CheckList* (Ribeiro et al., 2020), which is a type of behavioural testing that treats models—in our case Skills—as black-boxes and compares their behaviour against the expected one. This is achieved by *unit tests* designed by the end-users or the system experts. Two most common test types are **Minimum Functionality Test (MFT)** and **INVariance (INV)** as shown in Table 1. MFTs are designed to measure a capability (e.g., *Taxonomy* capacity of matching object properties to categories) via specifying the expected behaviour (e.g., “tiny” in Table 1). INVs tests are similarly refined for capabilities (e.g., robustness under spelling errors in question), however the expected behaviour is already known, i.e., the answer should remain the same. We adapt the machine comprehension tests from Ribeiro et al. (2020) for behavioural testing of our Skills. In our current setup, the tests for all the deployed Skills are curated manually, saved in as JSON file and made available via the UI. The test results are shown on demand via a separate tab (§3.3).

2.5 User Interface

We host UKP-SQUARE as a web application built with VueJS¹² to make it easily accessible to re-

¹¹From the MS MARCO dataset (Nguyen et al., 2016).

¹²<https://vuejs.org>

Training Dataset for Models	Domain
Text Generation (Abstractive QA)	
NarrativeQA (Kocisky et al., 2018)	Stories
Span Extraction (Extractive QA)	
BioASQ (Tsatsaronis et al., 2015)	Biomedical
DROP (Dua et al., 2019)	Wikipedia
DuoRC (Saha et al., 2018)	Movies
Natural Questions (Kwiatkowski et al., 2019)	Wikipedia
NewsQA (Trischler et al., 2017)	News
Quoref (Dasigi et al., 2019)	Wikipedia
SQuAD 1.1 (Rajpurkar et al., 2016)	Wikipedia
SQuAD 2.0 (Rajpurkar et al., 2018)	Wikipedia
TriviaQA (Joshi et al., 2017)	Wikipedia, Web
Text Classification (Multiple-Choice QA)	
BioASQ (Tsatsaronis et al., 2015)	Biomedical
BoolQ (Clark et al., 2019)	Wikipedia
CommonsenseQA (Talmor et al., 2019)	-
CosmosQA (Huang et al., 2019)	Personal Narratives
MultiRC (Khashabi et al., 2018)	Fiction, Textbook, Wikipedia, News, etc.
Quail (Rogers et al., 2020)	Fiction, News, Blogs, User Stories
Quartz (Tafjord et al., 2019)	Relationships
RACE (Lai et al., 2017)	News, Stories, Ads, Biography, Philosophy
SocialQA (Sap et al., 2019)	Social Interactions
Query Embedder (Retrieval)	
Natural Questions (Kwiatkowski et al., 2019)	Wikipedia
MS MARCO (Nguyen et al., 2016)	Bing Web Docs.

Table 2: Available Models fine-tuned on various datasets upon the release of UKP-SQUARE.

searchers. Once a Skill has been created by a user (§3.1) it can be added, edited, and deleted in the Skill management section of the application in the “My Skills” menu. For each Skill, its URL, metadata, requirements for context, and visibility can be adjusted (see Appendix Fig. 3). The functionality of the user interface is split into QA and explainability.

QA Interface. The QA section of the user interface provides access to the Skill by allowing the user to enter their question and optionally a context. Public Skills are accessible to everyone while private Skills require the user to be signed in. The UI provides distinct visualizations depending on the selected Skill type. For extractive Skills, e.g., SQuAD (Rajpurkar et al., 2016), a document and multiple spans are returned and ranked by the model’s confidence. In this setup, we also provide the option to show the span highlighted in its position in the document (see Fig. 1). Categorical Skills, e.g., BoolQ (Clark et al., 2019), show an interface with boolean output scores (see Appendix Fig. 5). A multiple-choice Skill requires multiple options separated by newlines in the context field. These are then ranked and returned with their

respective scores (see Appendix Fig. 6). When multiple Skills are selected, the user can see and compare their outputs side-by-side and better understand their behavioural differences.

Explainability Interface. A Skill selector is provided at the top which allows users to visualize and compare the results of the CheckList machine reading tests for the selected Skills. A list of tests with their name, type, capability, and failure rate is shown. The list can be expanded for a detailed description along with a small number of failed examples with their questions, context, and predictions.

3 Use Cases

3.1 Skill Publishing

A major contribution of our platform is to support developers creating their own Skills. This allows practitioners to easily make their research publicly available, without having to take care of engineering heavy topics such as infrastructure, web development and security. To publish a new Skill, developers need to implement a single function that defines the question answering pipeline. They are provided with utility functions that facilitate interacting with other components such as the Datastores, Models and the UI. A code snippet implementing a Skill is given in Appendix A.

Allowing developers to implement their own Skills enables us to greatly extend the system to have stronger models. For instance, multiple Datastores with potentially different retrieval methods can be combined to find complementary background knowledge, e.g., from Wikipedia and biomedical articles. Similarly, different models could be used to precisely answer a diverse set of questions that might require different capabilities, such as answerability (Rajpurkar et al., 2018), numerical (Dua et al., 2019) or multi-hop (Yang et al., 2018) reasoning. Once a developer creates their Skill, it can be added to UKP-SQUARE via the UI. The Skill developer can further make the Skill publicly available.

Allowing the community to implement Skills comes with a technical challenge such as deploying unreliable code on our servers. We therefore allow three different ways of hosting Skills. (1) First, Skills can be hosted directly on UKP-SQUARE. For this, a pull request for the new Skill should be submitted to our public repository, which can

then be added to the system upon a code review. While processing the submitted Skill requires a human in the loop, this option simplifies the hosting process for the Skill developer. (2) Second, in order to provide an option to make Skills instantly and independently available, we also allow Skills to be hosted on third party cloud platforms such as Amazon Web Services, Google Cloud and Microsoft Azure. All these cloud providers allow to easily host a lightweight function that can be used by UKP-SQUARE. (3) Lastly, we allow developers to host Skills on their own hardware. The only requirement is that the Skill needs to be publicly accessible. In the latter two cases, developers will still have access to UKP-SQUARE’s components (e.g., Datastores and Models), but the Skill itself will run on the cloud or on other hardware. For quick development of Skills we recommend using options (2) and (3). For long-term availability and usage of a Skill, adding it via the public github repository is recommended. We provide extensive documentation for all possibilities to host Skills.¹³

3.2 Skill Querying

Once a developer makes their Skill public in UKP-SQUARE, other users can obtain answers from it. Upon release of the system, we make a wide range of question answering Skills available. These span over different QA formats (extractive, multiple-choice, abstractive), setups (open-domain, machine reading comprehension) and to different domains (wikipedia, web, biomedical, etc.). The list of available models for different formats is given in Table 2. This allows the public to test current state-of-the-art question answering models. Moreover, researchers can use it for qualitative analysis, for example to discover potential biases, strengths or weaknesses in models by behavioural testing. Furthermore, we support querying multiple Skills at the same time. This is particularly useful to compare capabilities of different models. For example see Fig. 1, where two open domain, extractive Skills can be compared.

3.3 Behavioural Testing of Skills

The users can choose the Skill they want to investigate from the drop-down menu. The selected Skill can be analyzed standalone or alongside two different compatible Skills.

The tests are displayed showing the Skill fail-

¹³<https://square.ukp-lab.de/docs/>

	Supported Models	Retrieval	QA Types	Expl.	Ext.
Haystack	HF Transformers	sparse, dense	EX, AB	×	×
Dibia (2020)	HF Transformers	sparse	EX	✓	×
Karpukhin et al. (2020)	DPR	dense	EX	×	×
Khashabi et al. (2020)	T5	×	EX, AB, MC, YN	×	×
UKP-SQUARE	HF Transformers, ONNX, adapters	sparse, dense	EX, AB, MC, YN	✓	✓

Table 3: Qualitative comparison of UKP-SQUARE to previous works. HF: HuggingFace, Expl.: Explainability component, Ext.: Extensible by the end-user, EX: Extractive, AB: Abstractive, MC: Multiple-choice, YN: Yes/No

ure rate and the failed examples can be viewed by clicking on the ‘Expand’ button. An exemplary visualization for negation and coreference testing of SQuAD Skills is given in Appendix Fig. 4. For replacement tests, e.g., where names are perturbed, colored markers are used to highlight how the input was modified for the test. This allows the user to quickly identify changes the Skill could not handle. To analyze or process a Skill’s test performance in more detail, a full JSON report of all test examples can be downloaded.

4 User Study

We evaluate the usability of our system by conducting a pilot attitudinal user study with five participants. We recruited graduate students, our main target user group, and instructed them to compare and analyze several Skills. We provided them with a list of predefined questions to input into the system to help them use it. After the students used the system we asked them several questions to discover whether they understood every element of the interface effortlessly (i.e., the input and the output of the Skills, the list of behavioral cards of the Skills, and their specific contents). All users understood the input and output of the Skills and stated that the interface allows them to compare the Skills effortlessly. They also stated that the behavioral cards of the explainability component are useful to analyze the strong and weak points of the models and could help develop new Skills. However, most of them could not understand them in a glimpse. Hence, we will improve the presentation of these cards in a future update. Appendix C provides the list of questions and responses. To finish the study, we employed the System Usability Scale (SUS) questionnaire ([Brooke, 1996](#)) to quantitatively assess the global usability of the system. The average score is 70 out of 100, which refers to a “good usability” ([UIUX-Trend, 2021](#)).

5 Related Work

A qualitative comparison with similar frameworks is given in Table 3. The closest work to ours is Haystack, which is an open-source and scalable framework for building search systems over large document collections. Although it supports both sparse and dense retrieval techniques, models from the Huggingface (HF), and different QA types (abstractive and extractive) it lacks support for faster ONNX or memory efficient adapter models. Furthermore, it has to be set up by the users on their own infrastructure which requires technical expertise and sufficient hardware resources. [Dibia \(2020\)](#) introduce NeuralQA, an interactive tool for QA that leverages the benefits of sparse retrieval along with the HF reader models. However, NeuralQA is limited to extractive QA. [Karpukhin et al. \(2020\)](#) provide a simple user interface that employs efficient dense retrieval but only support models for open-domain QA. Finally, UnifiedQA ([Khashabi et al., 2020](#)) provides a demo page¹⁴ that employs a custom T5 based model trained on a wide range of QA datasets, hence supports a variety of QA formats. However, (1) it lacks the retrieval component, (2) is not scalable (to include different model formats), and (3) is not flexible (not possible to use models with different retrieval techniques). Unlike other previous systems, UKP-SQUARE is dynamically extendable allowing users to easily contribute with new Skills. Finally, except from gradient-based explanations in [Dibia \(2020\)](#), none of the systems have an explainability component.

6 Conclusion and Future Work

We introduce the UKP-SQUARE platform that enables researchers and developers to study and compare QA pipelines, i.e., Skills, that comprises a selection of Datastores, retrieval mechanisms and reader models. The platform enables querying ex-

¹⁴<https://unifiedqa.apps.allenai.org/>

isting public Skills, as well as implementing custom ones using UKP-SQUARE’s microservices and utility functions that support a large collection of model types and Datastores. Furthermore, users can simultaneously query multiple Skills, and analyze them through integrated behavioural tests.

Our architecture is scalable and flexible to incorporate most of the latest developments in the QA domain. Future versions will include automated deployment of custom models and Datastores, automated Skill selection by incorporating previous works (Puerto et al., 2021; Geigle et al., 2021) and increasing the number of supported Datastores (e.g., wikidata, Vrandečić and Krötzsch, 2014). We also plan to incorporate specialized models (e.g., using graph encoders, Ribeiro et al., 2021), structured reasoning approaches (Yasunaga et al., 2021) and interpretability techniques such as saliency maps (Li et al., 2016).

Ethics and Broader Impact Statement

Data This work does not generate new data. All datasets employed in used to construct Skills as described in §2.2, §2.3, and Table 2. The datasets are well-known to be safe for research purposes and do not contain any personal information or offensive content. We comply with the licenses and intended uses of each dataset. The licenses of each dataset can be seen in Appendix B.

Intended Use. The intended use of UKP-SQUARE is i) bringing different QA components together to share them as a skill with the rest of the world and ii) the analysis of these Skills. Our platform allows NLP practitioners to share their Skills with the community removing technical barriers such as configuration and infrastructure so that any person can reuse these models. In addition, users can analyze the available Skills through behavioral tests and compare them thanks to a user-friendly UI. This has a straightforward benefit for the research community (i.e., reproducible research and analysis of prior works), but also to the general public because UKP-SQUARE allows them to run state-of-the-art models without requiring them any special hardware and hiding complex settings such as virtual environments and package management.

Potential Misuse. Our platform makes use of Skills uploaded by the community. However, this current version does not incorporate any mechanism to ensure that these models are fair and with-

out bias. Nonetheless, UKP-SQUARE includes a module for explainability that uses CheckLists (Ribeiro et al., 2020) to analyze the strong and weak points of the Skills and to detect their biases and unfair content. Thus, we currently delegate the fairness checks to the authors of the models. We are not held responsible for errors, false, or offensive content generated by the Skills. Users should use them at their discretion.

Environmental Impact. Since UKP-SQUARE empowers the community to run publicly available Skills on the cloud, it has the potential to reduce CO₂ emissions from retraining previous models to make the comparisons needed when developing new models.

User Study. The participants are junior graduate students recruited on a voluntary basis. They are not part of this work, and never saw the user interface before the study. Before starting the study, they were given detailed instructions on the goals and scope of the study, and how the data was going to be used. Only non-personal data was recorded.

Acknowledgements

We thank Jan-Christoph Klie for his insightful feedback and suggestions on a draft of the paper and the project. We thank Richard Eckart de Castilho for advice on the general infrastructure and Nandan Thakur and Hossain Shaikh Saadi for their preliminary work on the project. We also thank Serkan Bayraktaroğlu for designing our logo.

This work has been funded by: (i) the German Research Foundation (DFG) as part of the UKP-SQUARE project (grant GU 798/29-1), (ii) the DFG as part of the QASciInf project (GU 798/18-3), (iii) the DFG within the project “Open Argument Mining” (GU 798/25-1), associated with the Priority Program “Robust Argumentation Machines (RA-TIO)” (SPP-1999), (iv) the DFG-funded research training group “Adaptive Preparation of Information from Heterogeneous Sources” (AIPHES, GRK 1994/1), (v) the European Regional Development Fund (ERDF) and the Hessian State Chancellery – Hessian Minister of Digital Strategy and Development under the promotional reference 20005482 (TexPrax), and (vi) the LOEWE initiative (Hesse, Germany) within the emergenCITY center.

References

- Junjie Bai, Fang Lu, Ke Zhang, et al. 2019. ONNX: Open Neural Network Exchange. <https://github.com/onnx/onnx>.
- John Brooke. 1996. SUS-a quick and dirty usability scale. *Usability evaluation in industry*, 189.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pradeep Dasigi, Nelson F. Liu, Ana Marasović, Noah A. Smith, and Matt Gardner. 2019. [Quoref: A reading comprehension dataset with questions requiring coreferential reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5925–5932, Hong Kong, China. Association for Computational Linguistics.
- Victor Dibia. 2020. [NeuralQA: A usable library for question answering \(contextual query expansion + BERT\) on large datasets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 15–22, Online. Association for Computational Linguistics.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. [DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2368–2378, Minneapolis, Minnesota. Association for Computational Linguistics.
- Gregor Geigle, Nils Reimers, Andreas Rücklé, and Iryna Gurevych. 2021. [TWEAC: transformer with extendable QA agent classifiers](#). *CoRR*, abs/2104.07081.
- Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy Lin, and Allan Hanbury. 2021. [Efficiently teaching an effective dense retriever with balanced topic aware sampling](#). In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 113–122. ACM.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [Cosmos QA: Machine reading comprehension with contextual commonsense reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2391–2401, Hong Kong, China. Association for Computational Linguistics.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. [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.
- Herve Jégou, Matthijs Douze, and Cordelia Schmid. 2011. [Product quantization for nearest neighbor search](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(1):117–128.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. [Looking beyond the surface: A challenge set for reading comprehension over multiple sentences](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 252–262, New Orleans, Louisiana. Association for Computational Linguistics.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hananeh Hajishirzi. 2020. [UNIFIEDQA: Crossing format boundaries with a single QA system](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online. Association for Computational Linguistics.

- Tomas Kocisky, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gabor Melis, and Edward Grefenstette. 2018. [The narrativeqa reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, 6(0):317–328.
- 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.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale ReAding comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Jiwei Li, Xinlei Chen, Eduard H. Hovy, and Dan Jurafsky. 2016. [Visualizing and understanding neural models in NLP](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 681–691.
- Andreas Madsen, Siva Reddy, and Sarath Chandar. 2021. [Post-hoc interpretability for neural NLP: A survey](#). *arXiv*, abs/2108.04840.
- Anastasios Nentidis, Anastasia Krithara, Konstantinos Bougiatiotis, Martin Krallinger, Carlos Rodríguez Penagos, Marta Villegas, and Georgios Paliouras. 2020. [Overview of bioasq 2020: The eighth bioasq challenge on large-scale biomedical semantic indexing and question answering](#). In *Experimental IR Meets Multilinguality, Multimodality, and Interaction - 11th International Conference of the CLEF Association, CLEF 2020, Thessaloniki, Greece, September 22-25, 2020, Proceedings*, volume 12260 of *Lecture Notes in Computer Science*, pages 194–214. Springer.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. [MS MARCO: A Human Generated Machine Reading Comprehension Dataset](#). In *Proceedings of the Workshop on Cognitive Computation, NIPS*.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020. [AdapterHub: A framework for adapting transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 46–54, Online. Association for Computational Linguistics.
- Haritz Puerto, Gözde Gül Sahin, and Iryna Gurevych. 2021. [MetaQA: Combining expert agents for multi-skill question answering](#). *arXiv*, abs/2112.01922.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Leonardo F. R. Ribeiro, Yue Zhang, and Iryna Gurevych. 2021. [Structural adapters in pretrained language models for AMR-to-Text generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4269–4282, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Marco Túlio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with checklist](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4902–4912.
- Stephen E. Robertson, Steve Walker, Susan Jones, Micheline Hancock-Beaulieu, and Mike Gatford. 1994. [Okapi at TREC-3](#). In *Proceedings of The Third Text REtrieval Conference, TREC 1994, Gaithersburg, Maryland, USA, November 2-4, 1994*, volume 500-225 of *NIST Special Publication*, pages 109–126. National Institute of Standards and Technology (NIST).
- Anna Rogers, Matt Gardner, and Isabelle Augenstein. 2021. [QA dataset explosion: A taxonomy of NLP resources for question answering and reading comprehension](#). *arXiv*, abs/2107.12708.
- Anna Rogers, Olga Kovaleva, Matthew Downey, and Anna Rumshisky. 2020. [Getting closer to ai complete question answering: A set of prerequisite real tasks](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8722–8731.
- Amrita Saha, Rahul Aralikkatte, Mitesh M. Khapra, and Karthik Sankaranarayanan. 2018. [DuoRC: Towards complex language understanding with paraphrased reading comprehension](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational*

- Linguistics (Volume 1: Long Papers)*, pages 1683–1693, Melbourne, Australia. Association for Computational Linguistics.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. **Social IQa: Commonsense reasoning about social interactions**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Oyvind Tafjord, Matt Gardner, Kevin Lin, and Peter Clark. 2019. **QuaRTz: An open-domain dataset of qualitative relationship questions**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5941–5946, Hong Kong, China. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. **CommonsenseQA: A question answering challenge targeting commonsense knowledge**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. **NewsQA: A machine comprehension dataset**. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 191–200, Vancouver, Canada. Association for Computational Linguistics.
- George Tsatsaronis, Georgios Balikas, Prodromos Malakasiotis, Ioannis Partalas, Matthias Zschunke, Michael R. Alvers, Dirk Weissenborn, Anastasia Krithara, Sergios Petridis, Dimitris Polychronopoulos, Yannis Almirantis, John Pavlopoulos, Nicolas Baskiotis, Patrick Gallinari, Thierry Artières, Axel-Cyrille Ngonga Ngomo, Norman Heino, Eric Gaussier, Liliana Barrio-Alvers, Michael Schroeder, Ion Androutsopoulos, and Georgios Paliouras. 2015. **An overview of the bioasq large-scale biomedical semantic indexing and question answering competition**. *BMC Bioinformatics*, 16(1):138.
- UIUX-Trend. 2021. Measuring and interpreting system usability scale. <https://uiuxtrend.com/measuring-system-usability-scale-sus>. Accessed: 2022-01-21.
- Denny Vrandečić and Markus Krötzsch. 2014. **Wiki-data: A free collaborative knowledgebase**. *Commun. ACM*, 57(10):78–85.
- 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. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. **Approximate nearest neighbor negative contrastive learning for dense text retrieval**. In *International Conference on Learning Representations*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. **HotpotQA: A dataset for diverse, explainable multi-hop question answering**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, and Jure Leskovec. 2021. **QA-GNN: Reasoning with language models and knowledge graphs for question answering**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 535–546, Online. Association for Computational Linguistics.

A Skill Implementation

The code below implements an open-domain, extractive QA Skill. First, a set of utility classes are loaded and initialized for facilitating interaction with UKP-SQUARE's Models and Datastore components (lines 1-5). Next, in the `predict` function, the Datastores are queried for retrieval. The Datastores component takes the user query, the datastore (Wikipedia snapshot from Natural Questions) and what index to use (dense, based on DPR) as input and returns the top documents. From these results, the document text and respective scores are extracted (lines 11-17). Subsequently, the query and the top documents are passed to the Models component for span extraction. In this implementation, a BERT base model with an adapter trained on SQuAD V2.0 is used (lines 21-30). Finally, the top answers are returned (lines 32-36).

```
1 from square_skill_api.models import QueryOutput, QueryRequest
2 from square_skill_helpers import ModelAPI, DataAPI
3
4 model_api = ModelAPI()
5 data_api = DataAPI()
6
7 async def predict(request: QueryRequest) -> QueryOutput:
8
9     # Dense document retrieval using the Datastores
10    # on a Wikipedia snapshot with DPR embeddings
11    data_api_output = await data_api(
12        datastore="nq",
13        index_name="dpr",
14        query=request.query,
15    )
16    context = [d["document"]["text"] for d in data_api_output]
17    context_score = [d["score"] for d in data_api_output]
18
19    # Answer extraction from the top document using the Model API
20    # using bert-base-uncased base model with SQuAD2.0 adapter
21    model_api_request = {
22        "input": [[request.query, c] for c in context],
23        "task_kwargs": {"topk": 1},
24        "adapter_name": "qa/squad2@ukp",
25    }
26    model_api_output = await model_api(
27        model_name="bert-base-uncased",
28        pipeline="question-answering",
29        model_request=model_api_request,
30    )
31
32    return QueryOutput.from_question_answering(
33        model_api_output=model_api_output,
34        context=context,
35        context_score=context_score
36    )
```

Listing 1: Example Implementation of an open-domain, span extraction Skill.

B Dataset Licences

Table 4 shows the license of each dataset. In the case of RACE, the authors did not provide any license but specified that it can only be used for non-commercial research purposes. In the case of the other datasets without any specified license, the authors did not provide any license, but the datasets are freely available to download and use in a research context. BioASQ is available by Courtesy of the U.S. National Library of Medicine.

Dataset	License
NarrativeQA	Apache 2.0
BioASQ	National Library of Medicine Terms and Conditions
DROP	CC BY-SA 4.0
DuoRC	MIT
Natural Questions	MIT
NewsQA	MIT
Quoref	CC BY 4.0
SQuAD 1.1	CC BY-SA 4.0
SQuAD 2.0	CC BY-SA 4.0
TriviaQA	Apache 2.0
BoolQ	CC BY-SA 3.0
CommonSenseQA	NA
CosmosQA	NA
MultiRC	NA
Quail	NA
Quartz	NA
RACE	NA
SocialIQA	NA
MS MARCO	CC BY 4.0

Table 4: License of each dataset.

C Questions of the User Study

Table 5 contains the answers of the participants of the user study (§4) to each question we asked to evaluate their understanding of the interface.

Question	Avg. Ans.
SQuARE provides a user interface that allows me to tell the difference between both Skills	4.4
I understand in a glimpse each card.	2.6
I can get a quick overall view of the weak points of the skill.	3.8
The examples of each CheckList item are useful.	4.4

Table 5: List of questions to understand the usefulness of the system. 1 represents "strongly disagree" and 5 represents "strongly agree."

D User Interface

UI screenshots for visualizing categorical and multiple choice Skill results are given in Fig. 5 and 6 respectively. In Fig. 3 the UI for managing a Skill is shown. Navigating through behavioural test results is given in Fig. 4.

My skills
SQuAD 1.1 BERT Adapter
Save

Skill name

Skill type

Description

Minimum multiple choice options

Requires context

Public

Skill URL

URL to the hosted skill (scheme://host:port/base_path)

Available

Provide example questions

These examples will be featured alongside your skill.

Example 1

<p>Question</p> <p>To whom did the Virgin Mary allegedly appear in 1858 in Lourdes France?</p>	<p>Context</p> <p>Architecturally, the school has a Catholic character. Atop the Main Building's gold dome is a golden statue of the Virgin Mary. Immediately in front of the Main Building and facing it, is a copper statue of Christ with arms upraised with the legend "Venite Ad Me Omnes". Next to the Main Building is the</p>
---	--

Figure 3: User interface for managing a Skill.

1. Select a skill

SQuAD 1.1 BERT Adapter — Extractive QA, bert-base-uncased

Compare up to three skills

SQuAD 1.1 RoBERTa Adapter — Extractive QA, roberta-base

None

Show Checklist

SQuAD 1.1 BERT Adapter	SQuAD 1.1 RoBERTa Adapter
<p><small>Download all examples</small></p> <p>Extractive QA, bert-base-uncased</p>	<p><small>Download all examples</small></p> <p>Extractive QA, roberta-base</p>
<p>20</p> <p>Negation in question only.</p> <p>Min Func Test test on Negation</p> <p><small>Expand</small></p>	<p>17</p> <p>Negation in question only.</p> <p>Min Func Test test on Negation</p> <p><small>Expand</small></p>
<p>40</p> <p>Basic coref, he / she</p> <p>Min Func Test test on Coref</p> <p><small>Expand</small></p>	<p>24</p> <p>Basic coref, he / she</p> <p>Min Func Test test on Coref</p> <p><small>Expand</small></p>

Figure 4: User interface for behavioural tests from Checklist.

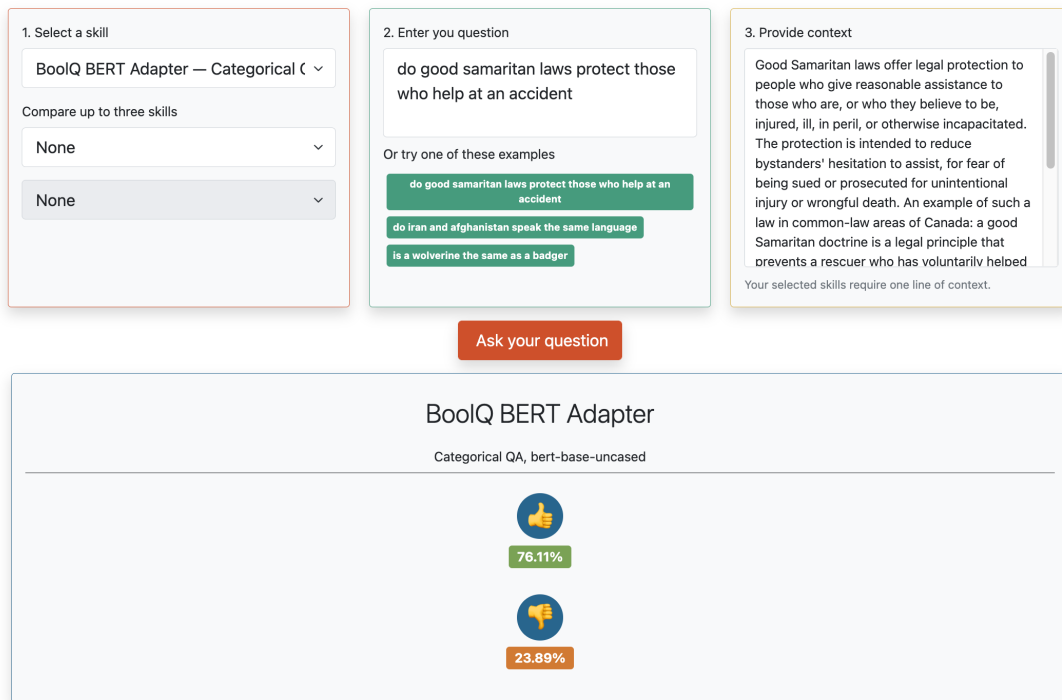


Figure 5: User interface for visualizing categorical Skill results.

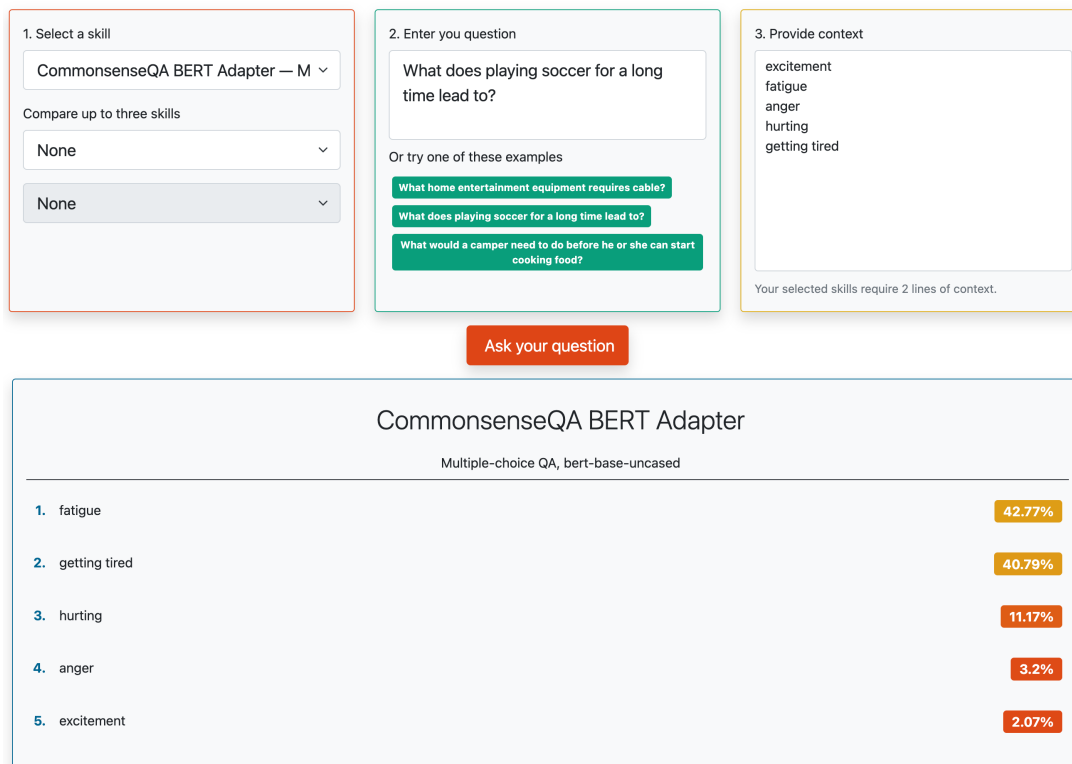


Figure 6: User interface for visualizing multiple choice Skill results.