

UKP-SQuARE v3: A Platform for Multi-Agent QA Research

Haritz Puerto, Tim Baumgärtner, Rachneet Sachdeva, Haishuo Fang, Hao Zhang,
Sewin Tariverdian, Kexin Wang, Iryna Gurevych

Ubiquitous Knowledge Processing Lab (UKP Lab),
Department of Computer Science and Hessian Center for AI (hessian.AI),
Technical University of Darmstadt
www.ukp.tu-darmstadt.de

Abstract

The continuous development of Question Answering (QA) datasets has drawn the research community’s attention toward multi-domain models. A popular approach is to use *multi-dataset models*, which are models trained on multiple datasets to learn their regularities and prevent overfitting to a single dataset. However, with the proliferation of QA models in online repositories such as GitHub or Hugging Face, an alternative is becoming viable. Recent works have demonstrated that combining expert agents can yield large performance gains over multi-dataset models. To ease research in *multi-agent models*, we extend UKP-SQuARE, an online platform for QA research, to support three families of multi-agent systems: i) agent selection, ii) early-fusion of agents, and iii) late-fusion of agents. We conduct experiments to evaluate their inference speed and discuss the performance vs. speed trade-off compared to multi-dataset models. UKP-SQuARE is open-source¹ and publicly available at square.ukp-lab.de.

1 Introduction

The current high-speed development of Artificial Intelligence yields thousands of datasets and trained models in repositories such as GitHub and Hugging Face (Rogers et al., 2023). These models are creating new research and application opportunities, such as high-performing Question Answering (QA) skills in chatbots (Burtsev et al., 2018; Miller et al., 2017). Comparing and analyzing these models usually requires learning libraries, writing code to run the models, and unifying their formats to compare them, which makes this process time-consuming and not scalable.

UKP-SQuARE (Baumgärtner et al., 2022b; Sachdeva et al., 2022) addresses this challenge, providing the first online platform that offers an ecosys-

tem for QA research enabling reproducibility, analysis, and comparison of QA models through a standardized interface and from multiple angles (i.e., general behavior, explainability, adversarial attacks, and behavioral tests).

The large variety of tasks and domains in QA datasets is pushing the research community towards creating models that generalize across domains (Fisch et al., 2019; Talmor and Berant, 2019; Khashabi et al., 2020). Currently, there are two main approaches to achieve this: i) multi-dataset models and ii) multi-agent models. While the former trains a model on multiple datasets (Talmor and Berant, 2019; Khashabi et al., 2020), the latter combines multiple expert agents (Geigle et al., 2021; Friedman et al., 2021; Puerto et al., 2023). Concurrently, large language models (LLM) such as GPT-3 (Brown et al., 2020) are emerging as new powerful systems for multi-task and multi-domain NLP applications. These LLM models are complementary to the focus of our work, multi-agent systems. While LLMs show impressive performance, they are extremely expensive to run and can usually only be accessed through APIs or deployed with great hardware resources. On the other hand, multi-agent systems offer a solution to create multi-domain models reusing available pretrained models that can be run on more modest hardware, which is an important requirement, e.g. where data cannot be sent to third parties.

Multi-agent models are particularly promising due to the thousands of models readily available on online model hubs and their current exponential growth.² This growth in the number of models is increasing the interest of the community in multi-agent model research (Wang et al., 2020; Matena and Raffel, 2021; Geigle et al., 2021; Friedman et al., 2021; Puerto et al., 2023; Wortsman et al., 2022; Jin et al., 2023). However, model hubs such as Hugging Face only allow inference on individ-

¹<https://github.com/UKP-SQuARE/square-core>

²https://www.nazneenrajani.com/emnlp_keynote.pdf

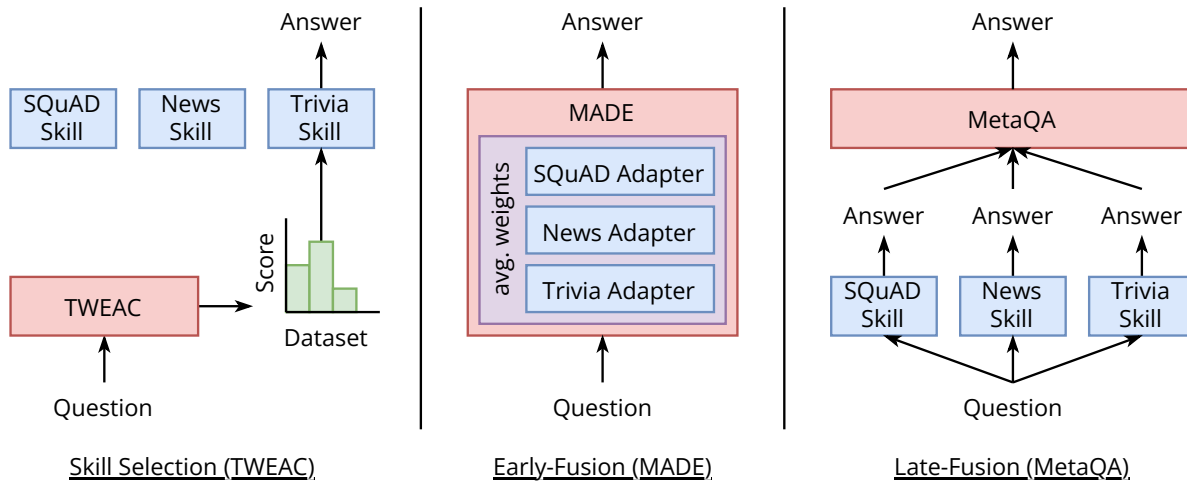


Figure 1: Overview of different multi-agent system architectures deployed in UKP-SQuARE. TWEAC (left) selects an agent (a *Skill* in UKP-SQuARE) based on which dataset it predicts the question is closest to and on which dataset a Skill was trained. MADE (center) fuses the weights of adapters trained on different datasets. MetaQA (right) predicts the final answer from a set of answers and their confidence scores. We illustrate the architectures with three different Skills. However, in practice, more Skills are used.

ual models, disregarding the possibility of combining them to make systems modular and multi-domain. This is a severe limitation as Puerto et al. (2023) showed that combining several QA models can yield performance gains of over 10 percentage points with respect to multi-dataset models (i.e., a single model trained on multiple datasets).

Therefore, we extend UKP-SQuARE to democratize access and research to multi-agent models. In particular, we add support to the three main methods to combine agents³: i) Skill selection, ii) early-fusion of Skills, and iii) late-fusion of Skills. The first consists of identifying the Skill with the highest likelihood of giving the correct answer and then routing the input to that Skill. We deploy TWEAC (Transformer With Extendable QA Agent Classifiers; Geigle et al., 2021) as an example of this method. The second one combines multiple models’ weights to obtain a new model with the distributional knowledge of the source weights. We deploy MADE (Multi-Adapter Dataset Experts; Friedman et al., 2021) as an example of this method. Lastly, the late-fusion of models consists of running multiple models to get their predictions and then combining them. This creates a system that can combine heterogeneous expert agents without reducing their performance in each domain. We provide MetaQA (Puerto et al., 2023) as an example of this method.

UKP-SQuARE facilitates research on multi-

agent QA systems by offering a platform equipped with dozens of agents and three methods to combine them. This upgrade holds paramount significance as the number of QA models created annually is increasing exponentially. UKP-SQuARE enables users to run, compare, and evaluate the strengths and weaknesses of multi-agent models, and compare them with multi-dataset models.

2 Related Work

The most famous types of multi-agent systems are Mixture of Experts (MoE) and ensemble methods. MoE consists of a gating mechanism that routes the input to a set of agents (Jacobs et al., 1991) while ensemble methods aggregate the outputs of multiple experts through a voting mechanism (Breiman, 1996; Freund and Schapire, 1996). Much work has been made to simplify the training of these multi-agent systems (Pedregosa et al., 2011; Chen and Guestrin, 2016; He et al., 2021; Hwang et al., 2022). However, as far as we know, there are no online platforms to run and compare them.

The most similar works to ours are the online model hubs such as Hugging Face’s Model Hub⁴ and AdapterHub (Pfeiffer et al., 2020a). They both offer a large number of models to download. In addition, Hugging Face’s Model Hub also allows running models through Spaces.⁵ However, this re-

³An agent is referred to as *Skill* in UKP-SQuARE.

⁴<https://huggingface.co/models>

⁵<https://huggingface.co/spaces>

quires implementing the Space, which can be non-trivial for complex scenarios such as ours (i.e., deploying and comparing multi-agent systems). UKP-SQuARE removes technical barriers and allows researchers to deploy multi-agent systems with a user-friendly interface.

Transformer (Vaswani et al., 2017) models using adapters (Houlsby et al., 2019) can also be seen as a type of multi-agent system. For this type of architecture, AdapterHub (Pfeiffer et al., 2020a) is a well-established library. In addition to simplifying the training of adapter-based models, it allows composing adapters (i.e., agents) with methods such as AdapterFusion (Pfeiffer et al., 2021) or stacking (Pfeiffer et al., 2020b). However, this library is not an online platform for analyzing models such as UKP-SQuARE. Their focus is to offer tools to create models based on adapters.

3 UKP-SQuARE

UKP-SQuARE (Baumgärtner et al., 2022b; Sachdeva et al., 2022) is the first online platform that offers an ecosystem for QA research. Its goal is to provide a common place to share, run, compare, and analyze QA models from multiple angles, such as explainability, adversarial attacks, behavioral tests, and I/O behaviors. The platform follows a flexible and scalable microservice architecture containing five main services:

- **Datstores:** Provide access to collections of unstructured text such as Wikipedia and Knowledge Graphs such as ConceptNet (Speer and Havasi, 2012).
- **Models:** Enable the dynamic deployment and inference of any Transformer model that implements a Hugging Face pipeline (Wolf et al., 2020) including models that use the adapter-transformers (Pfeiffer et al., 2020a) or sentence-transformers (Reimers and Gurevych, 2019) framework.
- **Skills:** central entity of the UKP-SQuARE. They specify a configurable QA pipeline (e.g., extractive, multiple-choice, and open-domain QA) leveraging Datstores and Models. Users interact with Skills since the platform’s goal is to remove technical barriers and focus on QA research (i.e., the QA pipeline). These Skills are equivalent to agents in the multi-agent system literature.

- **Explainability:** Provides saliency maps, behavioral tests, and graph visualizations⁶ that explains the outputs of a Skill.
- **Adversarial Attacks:** Create modified versions of the input to create adversarial attacks to expose vulnerabilities of the Skills.

All these services allow UKP-SQuARE to offer an ecosystem of tools to analyze Skills through a user-friendly interface without writing any code or complex configurations. UKP-SQuARE helps researchers identify the models’ strengths and weaknesses to push the boundaries of QA research.

3.1 Target Users and Scenarios

This new update of UKP-SQuARE targets researchers working on multi-agent and multi-dataset systems. These users can use the platform as a showcase of their systems. The dozens of Skills already available in UKP-SQuARE simplify the deployment of multi-agent systems since users can employ our user-friendly interface to select the Skills they want to combine using the three families of methods we deploy. Furthermore, researchers can deploy their new multi-skill methods through a pull request in our repository. The platform can also be used to analyze and compare multiple multi-agent systems from efficiency (i.e., inference time) and effectiveness (i.e., performance) points of view. Furthermore, it can also be used to compare multi-agent with multi-dataset systems. Lastly, UKP-SQuARE can also be used for teaching QA. The ecosystem of QA tools can be used to help students understand explainability, adversarial attacks, multi-dataset, and multi-agent models through interactive explanations with examples. Our platform can also be used to design homework where students train QA models and analyze them with the aforementioned QA tools.

4 Multi-Agent Systems

Multi-Agent systems are a type of multi-domain system that aggregate multiple expert agents from different domains to create a unified system. i.e., their focus is on the agents (*Skills* in UKP-SQuARE). On the other hand, multi-dataset systems aim to learn a unified model from multiple data distributions to create a single, general agent. For example, UnifiedQA (Khashabi et al., 2020) is

⁶For graph-based models.

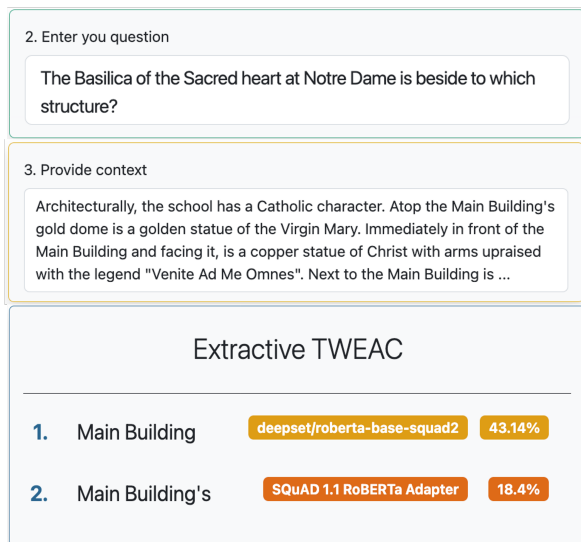


Figure 2: TWEAC predicts that the question is *SQuAD-like* and routes it to Skills trained on this dataset.

a QA model trained on multiple datasets using a generative model to overcome format boundaries.

However, Raffel et al. (2020) show that a model trained on multiple datasets may underperform the same architecture trained on a single dataset, i.e., multi-dataset models may underfit certain distributions. Based on this observation, Puerto et al. (2023) show that multi-agent models can avoid this limitation while being data-efficient to train and even outperform multi-dataset models by large margins in both in-domain and out-of-domain scenarios. This is possible because instead of using a very general architecture to solve multiple tasks, it uses a list of expert agents with specific architectures designed to solve those tasks (i.e., SOTA agents) and establishes a collaboration between these agents. However, this performance comes at a cost. The inference time is higher because it needs to run more than one model (at least one expert agent and one answer aggregator).

Therefore, we extend UKP-SQuARE to add support to the three main approaches for multi-agent systems, which we refer to as Meta-Skills on the platform: i) Skill Selection (§4.1), ii) Early-Fusion of Skills (§4.2), and iii) Late-Fusion of Skills (§4.3). An overview of the different architectures is illustrated in Figure 1.

4.1 Skill Selection

Skill selection is the simplest method of the three. It aims to identify the Skill with the highest likelihood of returning the correct answer to the input question and then route the input to that Skill. More

formally, it defines a function $f : Q \rightarrow S$ that maps any question Q to an available Skill S . Geigle et al. (2021) follow this approach and propose TWEAC (Transformer with Extendable QA Agent Classifiers), a Transformer model with a classification head for each Skill that maps questions to Skills. However, instead of predicting Skills, they predict *datasets*, i.e., they identify the dataset from which the input question comes. Then, they select a Skill trained on that dataset. Using this method, they report a Skill prediction accuracy higher than 90% across ten different QA types.

We train TWEAC on 16 datasets (shown in Appendix 5) with an accuracy of 79% and deploy it in UKP-SQuARE. The cause of the accuracy difference is the selection of the datasets. While the authors experiment on widely different QA tasks such as SQuAD, CommunityQA, and Weather Report, we use the most popular QA datasets, including the 2019 MRQA Shared Task (Fisch et al., 2019), which are more similar and thus, the task becomes more challenging since it is more difficult to distinguish the type of questions. We deploy two TWEAC Skills on UKP-SQuARE: one for extractive QA and another for multiple-choice. Figure 2 shows an extractive QA TWEAC that identifies the question as *SQuAD-like* and routes it to two Skills trained on SQuAD.

4.2 Early-Fusion of Skills

This method combines the weights of multiple models to create a new model that generalizes across all the input models.

Friedman et al. (2021) propose to train adapter weights for individual datasets while sharing the weights of a common Transformer that is also trained with those adapters. Later, in a second training phase, they freeze the Transformer weights and fine-tune each adapter on its corresponding dataset. The intuition behind this is that the shared parameters encode the regularities of the QA task while the adapters model the sub-distributions. This training schema yields a model that performs robustly on new domains by averaging its adapter weights.

Following this work, we extend UKP-SQuARE to allow the creation of Skills that average the weights of a series of adapters. To do this, on the Skill creation page (Figure 3), users are prompted to select whether they wish to combine adapters and, if affirmative, which ones to average.

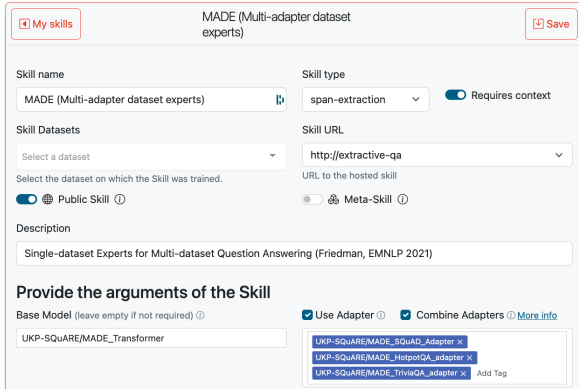


Figure 3: UKP-SQuARE allows combining adapters by simply writing the list of adapters.

4.3 Late-Fusion of Skills

Lastly, Puerto et al. (2023) propose MetaQA, a system that combines 18 heterogeneous expert agents across multiple formats. This system yields significant gains over multi-dataset models because some tasks require particular architectures to solve them, such as DROP (Dua et al., 2019), which requires numerical reasoning. Thus, while a *one-size-fits-all* architecture cannot learn such a wide variety of distributions, a multi-agent system that combines predictions can use expert agents to solve these datasets and yield a higher-performing model in general. Figure 4 shows how MetaQA answers a question from the *DuoRC* dataset but selects an out-of-domain (OOD) agent instead of the in-domain agent to answer, which gives a wrong answer. Thanks to the interface provided by UKP-SQuARE, it is easier to analyze the collaboration between the Skills established by MetaQA.

One limitation of this type of system is its need to run multiple models, which makes it more expensive than the previous two approaches. To alleviate this limitation, we run the expert agents in parallel. In this way, the inference time of MetaQA remains close to the other multi-agent systems, as shown in Table 1.

4.4 Comparison of Multi-Skill Models

In this section, we compare the inference time of the deployed multi-skill systems (i.e., MetaQA, TWEAC, and MADE) and UnifiedQA as a representative of the multi-dataset models. We extract 20 random questions from the six datasets from the MRQA 2019 Shared Task (Fisch et al., 2019) yielding a total of 120 questions and measure the time needed by each Skill to solve them. We repeat this process with five different random seeds and

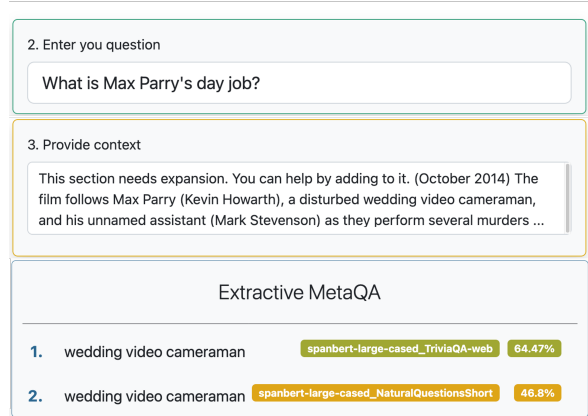


Figure 4: UKP-SQuARE simplifies the analysis of the collaboration between the agents. The question comes from the DuoRC dataset. However, while the in-domain agent gives a wrong answer (not shown), MetaQA selects an out-of-domain agent that gives a correct answer. Only an excerpt of the context is shown.

show the means and standard deviations in Table 1. Each model has 8 CPUs⁷ assigned to it and runs behind an asynchronous API.

As shown in Table 1, MetaQA is the slowest model. This is expected since it needs to run all the expert agents to get the predictions. However, its inference time is remarkably close to both MADE and TWEAC. TWEAC is surprisingly as fast as MADE, considering that TWEAC has to run at least two models (router and expert agent), while MADE only runs one. We conjecture that MADE is not faster because the adapter layers increase the depth of the transformer stack. UnifiedQA is the fastest model, as expected, since it is a multi-dataset model and hence, does not need to combine multiple agents.

Beyond inference, training time and cost are also interesting factors to consider. TWEAC and MetaQA are considered cheap to train assuming the existence of pretrained agents on online model hubs such as the Hugging Face Model Hub.⁸ Hence, the part that they train is a small router or answer aggregator. On the other hand, MADE and UnifiedQA require training a neural network from scratch in the task of question answering, which is much more challenging than simply routing questions or aggregating answers. Therefore, MADE and UnifiedQA need more training data than TWEAC and MetaQA, making them more expensive.

⁷AMD EPYC 7543 with 2.8GHz.

⁸3.6K models on https://huggingface.co/models?pipeline_tag=question-answering&sort=downloads. Accessed on Feb 2023

Model	F ₁	Inference Time [s]	Training
TWEAC	77.65	5.38 ± 0.06	cheap
MADE	82.20	5.45 ± 0.18	expensive
MetaQA	81.13	7.08 ± 0.16	cheap
UnifiedQA	77.30	2.15 ± 0.02	expensive

Table 1: Comparison of inference time on UKP-SQuARE averaged over 600 predictions. Performance from their respective papers.

Table 1 shows the trade-off between performance, training, and inference efficiency. Although MetaQA is the slowest Skill to run, its inference time is very close to the other models’ thanks to the parallel inference of the expert agents offered by UKP-SQuARE (cf. Figure 1). Furthermore, it is cheap to train, has almost the highest performance, and is compatible with any QA format. This makes it interesting for scenarios where model updating, performance, and flexibility are vital. TWEAC is also cheap and as flexible as MetaQA. Although, it is significantly worse than MetaQA on extractive QA datasets. This makes TWEAC ideal in the same scenarios as MetaQA but where running the expert agents in parallel is difficult (i.e., when MetaQA cannot be used). MADE has the highest performance and is as fast as TWEAC. However, it is more expensive to train than MetaQA and TWEAC, and it is not as flexible as MetaQA and TWEAC since it cannot be used for multiple formats simultaneously. Therefore, it should be used when inference, performance, and simple deployment are vital, while the model is not expected to need re-training (i.e., updates) often and is not required to be compatible with multiple QA formats at the same time. Lastly, UnifiedQA is compatible with any text-based QA format but has lower (although competitive) results. Although it is the fastest to run, it is more expensive to train than TWEAC and MetaQA. Thus, its ideal use case is a scenario where a simple deployment is needed while being flexible to process any text-based QA format. Therefore, this small study suggests that in scenarios where new domains are introduced often, router-based systems such as MetaQA might be more suitable, whereas, in scenarios where inference speed or simple deployment are needed, MADE and UnifiedQA might be more appropriate.

5 Conclusions and Discussions

In this work, we have extended UKP-SQuARE to support multi-agent models. In particular, we deployed a routing system, TWEAC (Geigle et al., 2021), a method to combine adapter weights, MADE (Friedman et al., 2021), and a model that combines the prediction of multiple Skills, MetaQA (Puerto et al., 2023). We have conducted experiments on these three models and UnifiedQA (Khashabi et al., 2020), a multi-dataset system, to analyze the trade-off between the performance, efficiency, and flexibility of these systems. We showed that in scenarios where new domains or expertise are often needed, MetaQA provides the best trade-off since its performance is close to the best model, it is compatible with any QA format, cheap to train, and its inference runtime is close to TWEAC and MADE using the parallel engine provided by UKP-SQuARE. However, when simple deployment is needed or the model is not expected to be updated, MADE and UnifiedQA might be more appropriate.

This update of UKP-SQuARE is of utmost importance due to the current speed of development of QA models that creates thousands of models per year. Our platform eases the deployment, running, comparison, and analysis of QA Skills. With this update, we also facilitated the aggregation of these Skills into Multi-Skills simplifying research on multi-agent systems. We leave as future work the comparison of these modular systems with prompting-based QA in large language models (Brown et al., 2020; Zhong et al., 2022).

Limitations

UKP-SQuARE v3 does not aim to provide all existing multi-skill systems off the shelf. Instead, we deploy three different approaches and encourage the community to share, deploy and compare their multi-skill systems. Using the modular Skill system of UKP-SQuARE and the reference implementations, users can reconfigure the existing multi-skill pipelines or implement and deploy their own through a streamlined pull request.⁹

Another limitation is that the multi-skill systems deployed in this paper have been shown to work effectively with no more than 20 Skills. Hence, the effectiveness of multi-skill systems remains unknown for a larger number of Skills. We hope

⁹For details, we refer to the documentation at <https://square.ukp-lab.de/docs/home/components/skills>

that UKP-SQuARE v3 can help shed light on this topic.

Lastly, since multi-skill systems combine several models, it is feasible that the resulting system can inherit biases and unfair behaviors. Although the Skills we used are not intended to exhibit any bias or unfairness, users should use them at their own discretion.

Ethics Statement

Intended Use The intended use of UKP-SQuARE v3 is deploying, running, comparing, analyzing, and combining Skills. Our platform provides dozens of Skills readily available to be combined using the implemented multi-agent systems or new systems to be created by the community. This simplifies the analysis of these systems and thus fosters multi-agent QA research.

Potential Misuse A malicious user could train multiple Skills with biased and unfair behaviors, such as a QA system that responds harmful answers, and combine them with the deployed methods available in UKP-SQuARE. UKP-SQuARE does not provide any Skill with such an intended behavior, but the community is free to upload any model to our platform. Therefore, we encourage the community not to publicly upload such models unless there is a clear research intention with a discussion of the ethics of such research, and in this case, make the Skills private, so that nobody can use them in an unintended way. We are not liable for errors, false, biased, offensive, or any other unintended behavior of the Skills. Users should use them at their own discretion.

Environmental Impact The use of UKP-SQuARE can reduce the computational cost of reproducing prior research since it prevents the community from training models that are already trained.

Acknowledgements

We thank Haris Jabbar, Martin Tutek, and Imbesat Hassan Rizvi for their insightful comments on a previous draft of this paper.

This work has been funded by the German Research Foundation (DFG) as part of the UKP-SQuARE project (grant GU 798/29-1), the QASci-Inf project (GU 798/18-3), and by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research,

Science and the Arts (HMWK) within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

References

- Tim Baumgärtner, Leonardo F. R. Ribeiro, Nils Reimers, and Iryna Gurevych. 2022a. [Incorporating relevance feedback for information-seeking retrieval using few-shot document re-ranking](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8988–9005, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tim Baumgärtner, Kexin Wang, Rachneet Sachdeva, Gregor Geigle, Max Eichler, Clifton Poth, Hannah Sterz, Haritz Puerto, Leonardo F. R. Ribeiro, Jonas Pfeiffer, Nils Reimers, Gözde Şahin, and Iryna Gurevych. 2022b. [UKP-SQUARE: An online platform for question answering research](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 9–22, Dublin, Ireland. Association for Computational Linguistics.
- Leo Breiman. 1996. [Bagging predictors](#). *Mach. Learn.*, 24(2):123–140.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Mikhail Burtsev, Alexander Seliverstov, Rafael Airapetyan, Mikhail Arkhipov, Dilyara Baymurzina, Nickolay Bushkov, Olga Gureenkova, Taras Khakhulin, Yuri Kuratov, Denis Kuznetsov, Alexey Litinsky, Varvara Logacheva, Alexey Lymar, Valentin Malykh, Maxim Petrov, Vadim Polulyakh, Leonid Pugachev, Alexey Sorokin, Maria Vikhрева, and Marat Zaynutdinov. 2018. [DeepPavlov: Open-source library for dialogue systems](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 122–127, Melbourne, Australia. Association for Computational Linguistics.
- Shulin Cao, Jiaxin Shi, Liangming Pan, Lunyiu Nie, Yutong Xiang, Lei Hou, Juanzi Li, Bin He, and Hanwang Zhang. 2022. [KQA pro: A dataset with explicit compositional programs for complex question answering over knowledge base](#). In *Proceedings of the*

- 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 6101–6119, Dublin, Ireland. Association for Computational Linguistics.
- Tianqi Chen and Carlos Guestrin. 2016. [Xgboost: A scalable tree boosting system](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, pages 785–794. ACM.
- 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.
- Matthew Dunn, Levent Sagun, Mike Higgins, V. Ugur Güney, Volkan Cirik, and Kyunghyun Cho. 2017. [Searchqa: A new q&a dataset augmented with context from a search engine](#). *CoRR*, abs/1704.05179.
- Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. [MRQA 2019 shared task: Evaluating generalization in reading comprehension](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 1–13, Hong Kong, China. Association for Computational Linguistics.
- Yoav Freund and Robert E. Schapire. 1996. Experiments with a new boosting algorithm. In *Machine Learning, Proceedings of the Thirteenth International Conference (ICML '96), Bari, Italy, July 3-6, 1996*, pages 148–156. Morgan Kaufmann.
- Dan Friedman, Ben Dodge, and Danqi Chen. 2021. [Single-dataset experts for multi-dataset question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6128–6137, Online and Punta Cana, Dominican Republic. 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.
- Jiaao He, Jiezhong Qiu, Aohan Zeng, Zhilin Yang, Jidong Zhai, and Jie Tang. 2021. [Fastmoe: A fast mixture-of-expert training system](#). *CoRR*, abs/2103.13262.
- 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, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Changho Hwang, Wei Cui, Yifan Xiong, Ziyue Yang, Ze Liu, Han Hu, Zilong Wang, Rafael Salas, Jithin Jose, Prabhat Ram, Joe Chau, Peng Cheng, Fan Yang, Mao Yang, and Yongqiang Xiong. 2022. [Tutel: Adaptive mixture-of-experts at scale](#). *CoRR*, abs/2206.03382.
- Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. 1991. [Adaptive mixtures of local experts](#). *Neural Comput.*, 3(1):79–87.
- Xisen Jin, Xiang Ren, Daniel Preotiuc-Pietro, and Pengxiang Cheng. 2023. [Dataless knowledge fusion by merging weights of language models](#). In *The Eleventh International Conference on Learning Representations*.
- 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.
- Jungo Kasai, Keisuke Sakaguchi, Yoichi Takahashi, Ronan Le Bras, Akari Asai, Xinyan Yu, Dragomir R. Radev, Noah A. Smith, Yejin Choi, and Kentaro Inui. 2022. [Realtime QA: what’s the answer right now?](#) *CoRR*, abs/2207.13332.
- 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.
- 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.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Michael Matena and Colin Raffel. 2021. [Merging models with fisher-weighted averaging](#). *CoRR*, abs/2111.09832.

- Alexander Miller, Will Feng, Dhruv Batra, Antoine Bordes, Adam Fisch, Jiaseen Lu, Devi Parikh, and Jason Weston. 2017. [ParIAI: A dialog research software platform](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 79–84, Copenhagen, Denmark. Association for Computational Linguistics.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake VanderPlas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Edouard Duchesnay. 2011. [Scikit-learn: Machine learning in python](#). *J. Mach. Learn. Res.*, 12:2825–2830.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. [AdapterFusion: Non-destructive task composition for transfer learning](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020a. [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.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020b. [MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Haritz Puerto, Gözde Şahin, and Iryna Gurevych. 2023. [MetaQA: Combining expert agents for multi-skill question answering](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3566–3580, Dubrovnik, Croatia. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- 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.
- Anna Rogers, Matt Gardner, and Isabelle Augenstein. 2023. [QA dataset explosion: A taxonomy of NLP resources for question answering and reading comprehension](#). *ACM Comput. Surv.*, 55(10):197:1–197:45.
- Rachneet Sachdeva, Haritz Puerto, Tim Baumgärtner, Sewin Tariverdian, Hao Zhang, Kexin Wang, Hosain Shaikh Saadi, Leonardo F. R. Ribeiro, and Iryna Gurevych. 2022. [UKP-SQuARE v2: Explainability and adversarial attacks for trustworthy QA](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 28–38, Taipei, Taiwan. Association for Computational Linguistics.
- Robyn Speer and Catherine Havasi. 2012. [Representing general relational knowledge in ConceptNet 5](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 3679–3686, Istanbul, Turkey. European Language Resources Association (ELRA).
- Alon Talmor and Jonathan Berant. 2019. [MultiQA: An empirical investigation of generalization and transfer in reading comprehension](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4911–4921, Florence, Italy. 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.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Jesse Vig. 2019. [A multiscale visualization of attention in the transformer model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 37–42, Florence, Italy. Association for Computational Linguistics.
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: a free collaborative knowledgebase](#). *Commun. ACM*, 57(10):78–85.

- Jing Wang, Mayank Kulkarni, and Daniel Preotiuc-Pietro. 2020. [Multi-domain named entity recognition with genre-aware and agnostic inference](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8476–8488, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 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.
- Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. 2022. [Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 23965–23998. PMLR.
- 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.
- Wanjun Zhong, Yifan Gao, Ning Ding, Yujia Qin, Zhiyuan Liu, Ming Zhou, Jiahai Wang, Jian Yin, and Nan Duan. 2022. [ProQA: Structural prompt-based pre-training for unified question answering](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4230–4243, Seattle, United States. Association for Computational Linguistics.

A Further Updates

Although our focus is the development of multi-domain systems, we further extended UKP-SQuARE with other minor but important features.

A.1 Knowledge Graph Question Answering

To maximize the advantage of multi-agent systems, UKP-SQuARE needs to be compatible with most QA systems available. UKP-SQuARE is already compatible with most QA formats (i.e., extractive, multiple-choice, boolean, and abstractive), and in this version, we add support for Knowledge Graph Question Answering (KGQA) systems. Sachdeva et al. (2022) include support for neuro-symbolic systems that combine language models with ConceptNet but lack support for KGQA models. Thus, this paper implements a generic KGQA Skill compatible with any knowledge graph and any generation model that generates SPARQL queries. As a demo, we deploy a BART-based (Lewis et al., 2020) semantic parser on KQA Pro (Cao et al., 2022), which is a complex KGQA dataset based on Wikidata (Vrandečić and Krötzsch, 2014) with nine different question types. The KGQA Skill parses a question into an executable SPARQL query, which is then executed against a KG to get the final answer. For this purpose, a dataset-centric subgraph is deployed using virtuoso.¹⁰ Thanks to the modularity of UKP-SQuARE, we can flexibly combine different semantic parsers with different KGs to get the final answer. Therefore, all aforementioned multi-Skill methods can be easily adapted for multi-Skill KGQA, which we leave as future work.

A.2 BERTViz

We also extended our explainability ecosystem by adding BERTViz (Vig, 2019), a method that allows the exploration of the attention weights as shown in Figure 5. While UKP-SQuARE v2 focuses on high-level explanations through saliency maps, BERTViz offers a low-level explanation utilizing the attention weights across all layers of the transformer models.

A.3 Datastores

Lastly, regarding Datastores, while UKP-SQuARE v1 focuses on document collections such as Wikipedia or PubMed, and UKP-SQuARE v2 focuses on Knowledge Graphs, UKP-SQuARE v3 offers a datastore that is updated in real-time. This

type of information is vital for real-time questions (i.e., questions whose answers may change over time; Kasai et al. (2022)). For instance, some facts, such as the president of a country, can change quickly. Therefore, we deploy a real-time datastore by using the Bing Search API. This datastore does not store documents and, instead, relies on the Bing Search engine to retrieve online documents (i.e., websites) that are more likely to be updated.

Furthermore, we create an information-retrieval Skill that allows the inference of only an IR model (instead of combining them with a reader model). We allow providing relevance-feedback for sparse retrieval, as it has shown to perform well in information-seeking and interactive scenarios (Baumgärtner et al., 2022a).

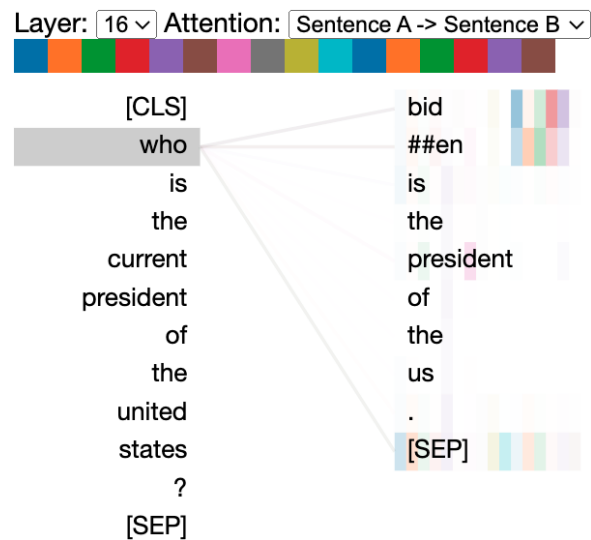


Figure 5: Screenshot of the BERTViz attention visualization in UKP-SQuARE. The token *who* attends to *Biden* (the answer) more than to other tokens.

¹⁰<https://virtuoso.openlinksw.com/>

B Datasets

Dataset	Characteristics	Train	Dev	Test	License
SQuAD (Rajpurkar et al., 2016)	Crowdsourced questions on Wikipedia	6573	5253	5254	MIT
NewsQA (Trischler et al., 2017)	Crowdsourced questions about News	74160	2106	2106	MIT
HotpotQA (Yang et al., 2018)	Crowdsourced multi-hop questions on Wikipedia	72928	2950	2951	MIT
SearchQA (Dunn et al., 2017)	Web Snippets, Trivia questions from J! Archive	117384	8490	8490	MIT
NQ (Kwiatkowski et al., 2019)	Wikipedia, real user queries on Google Search	104071	6418	6418	MIT
TriviaQA-web (Joshi et al., 2017)	Web Snippets, crowdsourced trivia questions	61688	3892	3893	MIT

Table 2: Summary of the datasets used in §4.4.

C Expert Agents

Expert Agents	Link
Span-BERT Large for SQuAD	https://huggingface.co/haritzpuerto/spanbert-large-cased_SQuAD
Span-BERT Large for NewsQA	https://huggingface.co/haritzpuerto/spanbert-large-cased_NewsQA
Span-BERT Large for HotpotQA	https://huggingface.co/haritzpuerto/spanbert-large-cased_HotpotQA
Span-BERT Large for SearchQA	https://huggingface.co/haritzpuerto/spanbert-large-cased_SearchQA
Span-BERT Large for NQ	https://huggingface.co/haritzpuerto/spanbert-large-cased_NaturalQuestionsShort
Span-BERT Large for TriviaQA-web	https://huggingface.co/haritzpuerto/spanbert-large-cased_TriviaQA-web
Span-BERT Large for QAMR	https://huggingface.co/haritzpuerto/spanbert-large-cased_QAMR
Span-BERT Large for DuoRC	https://huggingface.co/haritzpuerto/spanbert-large-cased_DuoRC

Table 3: List of the expert agents used for TWEAC and MetaQA.