

UKP-SQUARE v2

Explainability and Adversarial Attacks for Trustworthy QA

Rachneet Sachdeva*, Haritz Puerto*, Tim Baumgärtner, Sewin Tariverdian, Hao Zhang, Kexin Wang, Hossain Shaikh Saadi, Leonardo F. R. Ribeiro, 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

Question Answering (QA) systems are increasingly deployed in applications where they support real-world decisions. However, state-of-the-art models rely on deep neural networks, which are difficult to interpret by humans. Inherently interpretable models or post hoc explainability methods can help users to comprehend how a model arrives at its prediction and, if successful, increase their trust in the system. Furthermore, researchers can leverage these insights to develop new methods that are more accurate and less biased. In this paper, we introduce SQUARE v2, the new version of SQUARE, to provide an explainability infrastructure for comparing models based on methods such as saliency maps and graph-based explanations. While saliency maps are useful to inspect the importance of each input token for the model’s prediction, graph-based explanations from external Knowledge Graphs enable the users to verify the reasoning behind the model prediction. In addition, we provide multiple adversarial attacks to compare the robustness of QA models. With these explainability methods and adversarial attacks, we aim to ease the research on trustworthy QA models. SQUARE is available at <https://square.ukp-lab.de>.¹

1 Introduction

The recent explosion of Question Answering datasets and models is pushing the boundaries of QA systems and making them widely used by the general public in virtual assistants or chatbots (Rogers et al., 2021). This ubiquitous adoption is making regulators start preparing policies for artificial intelligence with special emphasis on explainability and robustness to adversarial attacks.²

*Equal Contribution.

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

²<https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence>

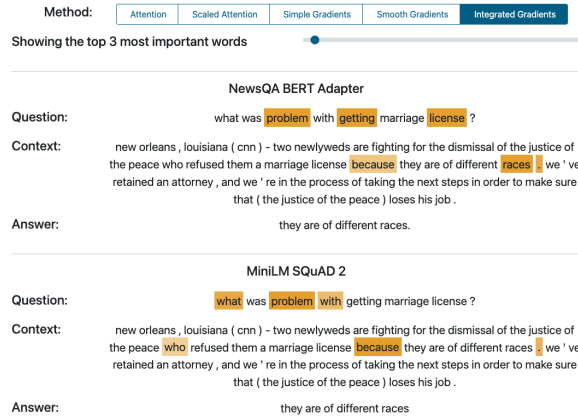


Figure 1: Visualization of two saliency maps computed using integrated gradients. The darker the highlighting color, the higher its importance to get the prediction. Hovering on a word shows its importance value.

There are multiple methods to explain the predictions of AI models (Danilevsky et al., 2020) and analyze their robustness (Zhang et al., 2020). Some explainability methods focus on specific input attributions such as attention- and gradient-based saliency maps (Simonyan et al., 2014). Others design interpretable models instead of using post hoc methods (Yasunaga et al., 2021). Lastly, most approaches that analyze the robustness of AI systems are based on *adversarial attacks*, i.e., the use of inputs such as questions with minor modifications that change the system’s output.

However, exploring and comparing these methods is not straightforward for most models. Researchers usually need to manipulate libraries and create interfaces to compare them in a satisfactory manner, which is a slow and complicated process that hinders the research in trustworthy QA.

The SQUARE platform (Baumgärtner et al., 2022) simplifies the process of comparing QA models by empowering NLP researchers with an online platform to deploy, run, and compare the most common QA pipelines while removing technical barriers such as model and infrastructure configu-

rations. It includes dozens of models of multiple types, namely open-domain, extractive, multiple choice, and abstractive QA. However, the only explainability method currently implemented is behavioral test (Appendix 3), limiting the comparison between QA models based solely on the models' final predictions.

In this work, we propose SQUARE v2, a new online platform for trustworthy QA research implementing various explainability, interpretability, and robustness methods and interfaces to facilitate research in trustworthy QA models. Specifically, we make the following contributions: 1) SQUARE v2 supports the comparison of models based on different post hoc explainability methods. We create interactive saliency maps that illustrate the importance of each input token for the model's prediction (Simonyan et al., 2014). 2) We extend the Datastores to include support for knowledge graphs (KG), deploy QA-GNN (Yasunaga et al., 2021), an interpretable graph-based model, and create an interactive visualization graph. 3) SQUARE v2 further provides various adversarial attacks, which change the prediction by modifying the input but keeping its semantics in order to evaluate the robustness of QA models (Ebrahimi et al., 2018).

2 Related Work

AllenNLP demo³ is the closest system to SQUARE v2. They provide a web interface to interact with their library, where users can explore explainability functionalities (Gardner et al., 2018; Wallace et al., 2019). However, only two non-Transformer models include saliency maps and attack methods. In addition, users cannot deploy their models on this web demo, and instead, they would need to install their library and create their own interface.

Among the explainability libraries, Captum (Kokhlikyan et al., 2020) is of special relevance. It is a model interpretability library for PyTorch that includes multiple saliency maps and provides built-in visualizations. However, it does not provide a user interface to run all their methods and compare them at a glance. On the other hand, it provides an adversarial attack method, Fast Gradient Sign Method (Goodfellow et al., 2015), and some variants; however, these are not designed for NLP.

Lastly, there are some efforts to ease the study of adversarial attacks on NLP models. Textattack

³<https://demo.allennlp.org>

(Morris et al., 2020) is a library that supports several attacks and is model agnostic. However, they do not provide a web interface, so users must therefore create their own visualizations in order to be able to easily compare attacks on multiple models.

In summary, SQUARE is a single entry-point for NLP practitioners to analyze, compare, and teach QA through models' outputs, explainability, and robustness with a user-friendly interface.

3 UKP-SQuARE

SQUARE (Baumgärtner et al., 2022) is an open-source, online platform for NLP researchers to share, run, compare, and analyze their QA models. The platform implements a flexible and scalable microservice architecture containing four high-level services:

1. **Datastores:** Provides efficient access to large-scale background knowledge such as Wikipedia.
2. **Models:** Allows the dynamic deployment and inference of a wide variety of models implemented in the Hugging Face transformers library (Wolf et al., 2020) or adapters (Pfeiffer et al., 2020).
3. **Skills:** Implements a configurable QA pipeline (e.g. multiple-choice, open-domain, or extractive QA) leveraging the Datastores and Models service. They can be added dynamically by the users to the system.
4. **Explainability:** provides a set of unit tests (questions and answers in our case) (Ribeiro et al., 2020) to compare the predictions with the expected answers and, in this way, analyze the biases and weaknesses of the Skills.

SQUARE is designed to ease the comparison and analysis of models. Users can deploy their models using a simple interface without the need of any code and then, they can compare outputs of different models side-by-side. This paper describes a new major update of SQuARE.

4 Trustworthy Methods for QA

Modern neural networks have significantly improved in performance in recent years; however, their explainability have not followed the same improvement (Rogers et al., 2020). Additionally, despite their impressive performance, the models

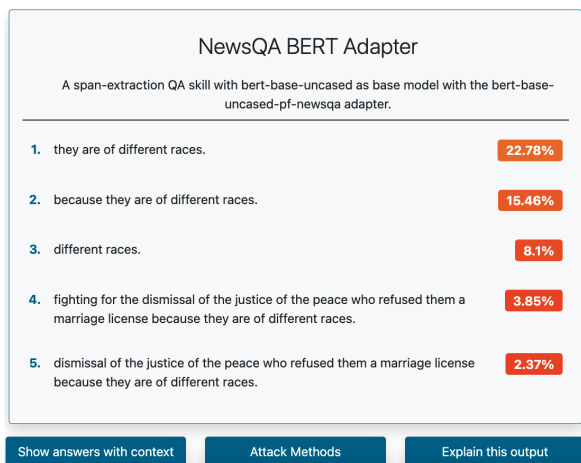


Figure 2: The "Explain this output" and "Attack Methods" are shown under the predictions of the Skill.

are vulnerable to adversarial attacks. The goal of SQUARE v2 is to provide the research community with a set of tools to facilitate the research on trustworthy QA. SQUARE simplifies and provides visualizations for saliency maps, graph-based interpretable models, and adversarial attacks. The following sections briefly describe the methods provided in SQUARE.

4.1 Saliency Maps

Saliency Maps assign an attribution weight to the input tokens to assess their importance in the model prediction, as illustrated in Fig. 1. To obtain this visualization, a user needs to click on the button "Explain this output" located after the predictions of any Skill, as shown in Fig. 2.

In SQUARE, we use two families of attribution methods to construct saliency maps: i) Gradient-based methods and ii) Attention-based methods.

4.1.1 Gradient-based Methods

A common approach to obtaining an importance score for the input tokens is to compute the gradients on the embedding layer against the model prediction. The magnitude of the gradient corresponds to the change of the prediction when updating the embedding. Therefore, a large gradient has a large effect on the prediction, indicating the importance of the input.

Vanilla Gradient (Simonyan et al., 2014) utilizes the plain gradients of the embedding layer of the model as importance weights of the inputs.

Integrated Gradient (Sundararajan et al., 2017) integrates the straight line path from the vector of zeros to the input token embedding. The value of

this integral is the weight of this token to make the prediction since it represents the amount of information given with respect to the zero vector (i.e., no information).

SmoothGrad (Smilkov et al., 2017) adds gaussian noise to the input to create multiple versions and then average their saliency scores. In this way, this method can smooth the saliency scores and alleviate noise from local variations in the partial derivatives.

4.1.2 Attention Methods

Neural NLP models have broadly incorporated attention mechanisms, which are frequently recognized for enhancing transparency and increasing performance (Vaswani et al., 2017). These methods compute a distribution over the input tokens that can be considered to reflect what the model believes to be important. Following (Jain et al., 2020), we build a saliency map using the average **attention weights** of the heads from the CLS token to the other tokens of the input. However, Serrano and Smith (2019) argue that attention weights are inconsistent and may not always correlate with the human notion of importance. Thus, they propose an alternative, **Scaled Attention**, which we also integrate in SQUARE, that multiplies the attention weights by their corresponding gradients to make it more stable.

4.2 Interpretable Graph-based Models

Knowledge graphs store knowledge in the form of relations (edges) between entities (nodes). In addition to the explicit facts they represent, they enable explainable predictions by providing reasoning paths (Yasunaga et al., 2021). In SQUARE v2, we deploy QA-GNN (Yasunaga et al., 2021), a graph-based QA model, as a Skill (more details on Appendix B) and ConceptNet (Speer et al., 2017) as a graph Datastore. Since QA-GNN uses a KG (i.e., ConceptNet) for QA reasoning, it is possible to analyze its working graph to identify the most important entities and relations for the answer prediction. As shown in Fig. 3 and later discussed in §6.2, we provide an interface that enables the visualization of the graph-based reasoning process executed by the model.

User Interface. In order to plot the graphs, SQUARE provides users with a "Show graph" button after the predictions of the QA-GNN Skill at the bottom of the page. Clicking on this button

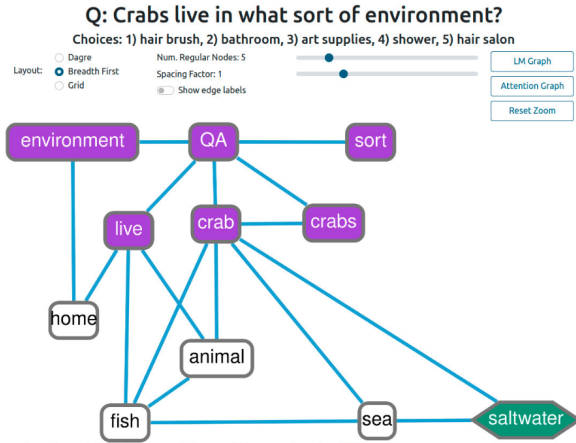


Figure 3: Visualization of the graph used by QA-GNN Skill to answer the question. Question nodes in purple, answer nodes in green.

displays a modal window with multiple options to render the graph, as shown in Fig. 3. Controls include a switch to show or hide edge labels, a slider to show the top k nodes, another slider to select the spacing factor between nodes, and a group of radio buttons to select the layout (Dagre⁴, Breadth First, and Grid). In addition, we offer two types of visualizations: i) a graph where the nodes are sorted by the relevance scores generated by the model and ii) a graph with the nodes sorted by the sum of the attention scores of their incoming edges.

4.3 Adversarial Attacks

Adversarial attacks make use of inputs that expose vulnerabilities of machine learning models to understand their robustness and identify how to improve them (Ebrahimi et al., 2018). To simplify the exploration of adversarial attacks on a wide range of Skills, we implement the following four methods in SQUARE for span-extraction Skills and leave the other Skills for future updates.

Figure 4: HotFlip Attack. Changing one word changes the prediction.

⁴<https://github.com/cytoscape/cytoscape.js-dagre>

HotFlip (Ebrahimi et al., 2018) uses a saliency method (§4.1) to score input words and subsequently replaces the top words with semantically similar words to alter the prediction of the model. An example of the interface is shown in Fig. 4. The words highlighted in green are replacements, and when hovering over them, the original word is shown in a tooltip.

Input Reduction (Feng et al., 2018) iteratively removes unimportant words from the question based on their saliency scores (§4.1), without changing the model’s prediction. An example is shown in Fig. 5 (Appendix C).

Sub-Span Jain et al. (2020) computes the saliency scores of the input words to select a contiguous span that maximizes the accumulative saliency score and uses this span as an explainability method. Instead, we leverage this method to create an adversarial attack. We identify a sub-span of the context that explains the output and use it as the whole context. In this way, the model has the key information, such as a phrase containing the answer, but not the whole context. Therefore, it is possible to identify a sub-span that lacks the nuance to answer the question properly, but since the answer occurs in the sub-span, the model may retrieve it due to spurious correlations. Fig. 6 (Appendix C) and § 6.3 show an use case of this attack.

Top K. Similarly as in the previous case, Jain et al. (2020) compute the saliency scores of the input words to identify the top k words from the context that explains the output answer. We leverage this method to create an adversarial attack. While the top k words are key to obtaining the answer, they are usually not contiguous. Therefore, creating a new context by concatenating these words yields a grammatically and semantically incorrect text. If the model still identifies the correct answer using this new context, it would be due to spurious correlations. An example of this attack is shown in Fig. 7 (Appendix C).

User Interface After the user queries any Skill, the button “Attack Method” is shown under the predictions, as shown in Fig. 2. After clicking on it, a modal page is shown where users can conduct adversarial attacks.

5 Datastores for Knowledge Graphs

To best re-use the existing Datastore while being efficient and robust, we rely on an Elasticsearch instance to store KGs. In particular, we represent

nodes and edges as documents and include information to recreate the graph structure, such as their connectivity. We show the schema of these documents in Appendix A.

In addition, we implement two main functionalities: Firstly, users can dynamically add and update new KGs as long as the structure of the KG can be converted to the schema shown in Appendix A. This allows supporting any KG that is requested by the community. For demonstration purposes, we provide ConceptNet (Speer et al., 2017) as a built-in KG. More information is available on the Datastores documentation⁵. Secondly, we implement a subgraph extraction method. Given a list of root nodes (e.g., the entities in a question), it extracts all the nodes and edges in the vicinity of k hops to the roots. Since ConceptNet is densely connected, we limit the maximum number of hops to 3. However, this is a parameter that can be adjusted for any KG. Lastly, after the extraction, we prune the disconnected nodes.

6 Case Study

6.1 Saliency Maps

Our new saliency map interface allows users to compare the explanation of the outputs of up to three Skills. As shown in Fig. 1, thanks to this visualization, we can easily observe that the first Skill, *NewsQA BERT Adapter*, gives the correct answer for the right reasons since it identifies "races" as a keyword. Even though the second Skill, *MiniLM SQuAD 2*, also returns the correct answer, the Skill does not seem to understand the context properly. In particular, the most important words for the predictions are not related to the answer. We argue that this interface can provide insights into whether the model understands the task and thus make the Skills more trustworthy.

6.2 Interpretable Models

ConceptNet provides background knowledge that can boost the commonsense abilities of NLP models. As shown in Fig. 3, the QA-GNN Skill makes use of the KG to connect the entities *crab* with *sea* and with *saltwater*, the answer. This explicit path helps to identify why the model returns its answers. However, it still requires some human effort to interpret the graph. For example, ConceptNet

does not include the triple (*sea, is a, environment*), which could be seen as counter-intuitive.

On the other hand, other non-graph-based Skills need post hoc explainability methods such as saliency maps (§4.1) to explain their output. However, post hoc methods have raised concerns about the possibility of not being faithful to the actual computations performed by the model or giving incomplete explanations as in saliency maps (Liu et al., 2021). In particular, saliency maps identify what parts of the input are relevant for the prediction, but they do not explain how or why the model obtains the output.

6.3 Adversarial Attacks

Using the Sub-span attack method shown in Fig. 6 (Appendix C), we can observe that the Skill gives the correct answer even though it does not have information about *Super Bowl 50*, which is needed. A robust Skill should instead return "not enough information." This example suggests that the Skill is conducting a superficial question-context overlap matching without understanding the nuances of the question, a phenomenon previously identified by Lim et al. (2020). Similarly, the input reduction attack shown in Fig. 5 (Appendix C) shows the same phenomenon. After removing most words from the question, the resulting question is not semantically complete, yet the Skill gives the correct answer.

7 Conclusion and Future Work

We present SQUARE v2, a web platform that unifies three families of methods for analyzing QA models: saliency maps, adversarial attacks, and interpretable models. Firstly, we offer an interactive interface that allows users to compare multiple saliency map methods for all the Skills deployed in SQUARE. Secondly, we provide an interface to conduct adversarial attacks. This interface allows the community to study the robustness of QA models. Lastly, we deploy an interpretable graph-based model and provide an interface to visualize the reasoning paths that the model may conduct. To deploy this Skill, we extend the Datastores module to support both text documents and KGs. These contributions give SQUARE a set of tools to compare, analyze, and explain the behavior of QA models. Since SQUARE allows the deployment of almost any Transformer-based model effortlessly, our new explainability interface empowers the community with tools for trustworthy QA research. SQUARE

⁵<https://square.ukp-lab.de/docs/api/datastores/>

is actively under development. Future updates will include new KGs such as WikiData (Vrandečić and Krötzsch, 2014), automated Skill selection (Geigle et al., 2021), and Skill collaboration (Puerto et al., 2021).

Limitations

Although saliency maps attempt to explain the output of the models, they should be analyzed with skepticism. As discussed in §4.1.2, attention-based saliency maps may not correlate with the human interpretation of importance, and in general, they do not explain how and why the model creates the outputs. Instead, saliency maps only aim to identify regions of the input that upon removal, changes the output.

Currently, we only deploy one graph-based model (QA-GNN) and one knowledge graph (ConceptNet). However, our Datastores §5 and graph visualization interface §4.2 are flexible enough to accommodate any other model, and thus, we invite the community to create pull requests and deploy their graph-based models on SQUARE.

Ethics and Broader Impact Statement

Intended Use. The intended use of SQUARE is to facilitate the comparison of QA models through multiple angles such as performance, explainability, interpretability, and robustness. Our platform allows NLP practitioners to share their models with the community removing technical barriers such as configuration and infrastructure so that any person can reuse these models. 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 SQUARE allows them to run state-of-the-art models without requiring any special hardware and hiding complex settings such as virtual environments and package management.

Potential Misuse. Our platform makes use of models uploaded by the community. However, this current version does not incorporate any mechanism to ensure that these models are fair and without bias. We hope that the new tools we provide in this work can help the community understand the outputs of QA models and identify potential biases or unfair behaviors. Thus, we currently delegate the fairness checks to the authors of the models. We are not held responsible for errors, false or of-

fensive content generated by the models. Users should use them at their discretion.

Environmental Impact. Since 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.

Acknowledgements

We thank Irina Bigoulaeva and Haishuo Fang for their insightful comments on a previous draft of this paper. We also thank the anonymous reviewers for their insightful feedback.

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, 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. 2022. [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.
- Marina Danilevsky, Kun Qian, Ranit Aharonov, Yanis Katsis, Ban Kawas, and Prithviraj Sen. 2020. [A survey of the state of explainable AI for natural language processing](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 447–459, Suzhou, China. Association for Computational Linguistics.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. [HotFlip: White-box adversarial examples for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- Shi Feng, Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, and Jordan Boyd-Graber. 2018.

- Pathologies of neural models make interpretations difficult. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3719–3728, Brussels, Belgium. Association for Computational Linguistics.
- Max Franz, Christian Tannus Lopes, Gerardo Huck, Yue Dong, Selçuk Onur Sümer, and Gary D. Bader. 2016. Cytoscape.js: a graph theory library for visualisation and analysis. *Bioinform.*, 32(2):309–311.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. AllenNLP: A deep semantic natural language processing platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. 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.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Sarthak Jain, Sarah Wiegrefe, Yuval Pinter, and Byron C. Wallace. 2020. Learning to faithfully rationalize by construction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4459–4473, Online. Association for Computational Linguistics.
- Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, and Orion Reblitz-Richardson. 2020. Captum: A unified and generic model interpretability library for pytorch. *CoRR*, abs/2009.07896.
- Doyeon Lim, Haritz Puerto San Roman, and Sung-Hyon Myaeng. 2020. Analysis of the semantic answer types to understand the limitations of mrqa models. *Journal of KIISE : Software and Applications*, 47(3):298–309.
- Zixuan Liu, Ehsan Adeli, Kilian M Pohl, and Qingyu Zhao. 2021. Going beyond saliency maps: Training deep models to interpret deep models. In *International Conference on Information Processing in Medical Imaging*, pages 71–82. Springer.
- John Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. TextAttack: A framework for adversarial attacks, data augmentation, and adversarial training in NLP. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 119–126, Online. Association for Computational Linguistics.
- 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. *CoRR*, abs/2112.01922.
- Marco Tulio 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*, pages 4902–4912, Online. Association for Computational Linguistics.
- 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, and Anna Rumshisky. 2020. A primer in BERTology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Sofia Serrano and Noah A. Smith. 2019. Is attention interpretable? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2931–2951, Florence, Italy. Association for Computational Linguistics.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*.
- Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda B. Viégas, and Martin Wattenberg. 2017. Smoothgrad: removing noise by adding noise. *CoRR*, abs/1706.03825.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4444–4451. AAAI Press.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 3319–3328. PMLR.
- 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.
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: A free collaborative knowledgebase](#). *Commun. ACM*, 57(10):78–85.
- Eric Wallace, Jens Tuyls, Junlin Wang, Sanjay Subramanian, Matt Gardner, and Sameer Singh. 2019. [AllenNLP interpret: A framework for explaining predictions of NLP models](#). 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): System Demonstrations*, pages 7–12, Hong Kong, China. 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.
- 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.
- Wei Emma Zhang, Quan Z Sheng, Ahoud Alhazmi, and Chenliang Li. 2020. [Adversarial attacks on deep-learning models in natural language processing: A survey](#). *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(3):1–41.

A Knowledge Graph Document Schema

The nodes of a knowledge graph are stored in the Datastore as a json document using the following schema:

```
{
  "node_id": {
    "_id": "keyword",
    "name": "keyword",
    "description": "text",
    "type": "keyword"
  }
}
```

The edges of a knowledge graph are stored in the Datastore as a json document using the following schema:

```
{
  "edge_id": {
    "_id": "keyword",
    "name": "keyword",
    "description": "text",
    "type": "keyword",
    "in_id": "keyword",
    "out_id": "keyword",
    "weight": "double"
  }
}
```

B QA-GNN Implementation


We implement the QA-GNN inference pipeline on SQUARE based on the official implementation of QA-GNN model.⁶ We disregard the training code since training QA models is not in the scope of SQUARE and connect the model with the Datastore service holding the KG. This makes it more flexible for future updates of ConceptNet. Lastly, the retrieved nodes with corresponding attention weights and relevance scores are accessible along with the answer prediction. With this information, we plot the graph using the JavaScript library Cytoscape.js (Franz et al., 2016).

⁶<https://github.com/michiyasunaga/qagnn>

C Adversarial Attack Figures

Attack Methods

Method: HotFlip **Input Reduction** Sub-Span Top K

Reductions = 10 

SQuAD 1.1 BERT Adapter

Question: to whom did the virgin mary allegedly appear in 1958 in lourdes france ?

Context: 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 Basilica of the Sacred Heart. Immediately behind the basilica is the Grotto, a Marian place of prayer and reflection. It is a replica of the grotto at Lourdes, France where the Virgin Mary reputedly appeared to Saint Bernadette Soubirous in 1858. At the end of the main drive (and in a direct line that connects through 3 statues and the Gold Dome), is a simple, modern stone statue of Mary.

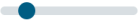
New Answer: saint bernadette soubirous **Old Answer:** Saint Bernadette Soubirous

Figure 5: Input Reduction. After removing tokens from the question, the new question is not specific enough to be answerable. Yet, the model still gives the same answer evidencing a spurious correlation.

Attack Methods

Method: HotFlip Input Reduction **Sub-Span** Top K

Saliency Method: Simple Gradients Smooth Gradients IntegratedGrad Attention

Length of sub-span = 10 

SQuAD 1.1 BERT Adapter

Question: Which NFL team represented the AFC at Super Bowl 50?

Context: super bowl 50 was an american football game to determine the champion of the national football league { nfl } for the 2015 season : the american football conference (afc) champion denver broncos defeated the national football conference { nfc } champion carolina panthers 24 – 10 to earn their third super bowl title :

New Answer: denver broncos **Old Answer:** Denver Broncos

Figure 6: Sub-Span Attack. Removing part of the context leaves a new context without the nuances needed to properly respond to the question (i.e., *at Super Bowl 50*).

Attack Methods

Method: HotFlip Input Reduction Sub-Span Top K

Saliency Method: Simple Gradients Smooth Gradients IntegratedGrad Attention

Top k = 19

SQuAD 1.1 BERT Adapter

Question: Which NFL team represented the NFC at Super Bowl 50?

Context: super bowl 50 was an american football game to determine the champion of the national football league (nfl) for the 2015 season . the american football conference (afe) champion denver broncos defeated the national football conference (nfc) champion carolina panthers 24 - 10 to earn their third super bowl title .

New Answer: carolina panthers **Old Answer:** Carolina Panthers

Figure 7: Top K Attack. Using as context the highlighted words, the Skill still gives the same answer even though the context is semantically and grammatically incomplete and does not include *Super Bowl 50*.