

Encoding Explanatory Knowledge for Zero-shot Science Question Answering

Zili Zhou^{1,2}, Marco Valentino¹, Dónal Landers², André Freitas^{1,2,3}

¹Department of Computer Science, University of Manchester, United Kingdom

²Digital Experimental Cancer Medicine Team

Cancer Research UK Manchester Institute, United Kingdom

³Idiap Research Institute, Switzerland

{zili.zhou, marco.valentino}@manchester.ac.uk

donal.landiers@digitalecmt.org

andre.freitas@idiap.ch

Abstract

This paper describes N-XKT (Neural encoding based on eXplanatory Knowledge Transfer), a novel method for the automatic transfer of explanatory knowledge through neural encoding mechanisms. We demonstrate that N-XKT is able to improve accuracy and generalization on science Question Answering (QA). Specifically, by leveraging facts from background explanatory knowledge corpora, the N-XKT model shows a clear improvement on zero-shot QA. Furthermore, we show that N-XKT can be fine-tuned on a target QA dataset, enabling faster convergence and more accurate results. A systematic analysis is conducted to quantitatively analyze the performance of the N-XKT model and the impact of different categories of knowledge on the zero-shot generalization task.

1 Introduction

Contemporary Question Answering (QA) is evolving in the direction of addressing more abstract reasoning tasks (Thayaparan et al., 2020; Dua et al., 2019; Clark et al., 2018; Mihaylov et al., 2018), supported by multi-hop inference (Khot et al., 2020; Yang et al., 2018) and explanatory scientific facts (Jansen and Ustalov, 2019; Jansen et al., 2018, 2016).

This trend of aiming to address more complex, multi-evidence and chained inference is pushing the envelope for novel representation and architectural patterns (Ding et al., 2019; Qiu et al., 2019; Asai et al., 2020; Thayaparan et al., 2019; Kundu et al., 2019; Valentino et al., 2021), which are moving from modelling meaning from immediate distributional semantics patterns into deeper abstractive capabilities. This poses a paradigmatic challenge on the design of QA architectures, which need to operate over high-level semantic patterns and acquire the necessary knowledge to perform abstraction (Clark et al., 2018). At the same time, the

design of new strategies to incorporate explanatory knowledge into neural representation has the potential to address fundamental data efficiency problems and promote zero-shot generalisation on out-of-distribution examples.

Explanation-based Science QA (Jansen et al., 2018) provides a rich framework to evaluate these emerging requirements, as the task typically requires multi-hop reasoning through the composition of explanatory facts. While existing approaches in the field mainly focus on the construction of natural language explanations (Jansen et al., 2018; Jansen and Ustalov, 2019), this work aims to explore the impact of explanatory knowledge on zero-shot generalisation.

In this paper, we argue that explanation-centred corpora can serve as a resource to boost zero-shot capabilities on Question Answering tasks which demand deeper inference. To this end, we explore the adoption of latent knowledge representations for supporting generalisation on downstream QA tasks requiring multi-hop inference.

Our hypothesis is that explanatory scientific knowledge expressed in natural language can be transferred into neural network representations, and subsequently used to achieve knowledge based inference on scientific QA tasks. To validate this hypothesis, this paper proposes a *unified* approach that frames Question Answering as an explanatory knowledge reasoning problem. The unification between the two tasks allows us to explore the adoption of pre-training strategies over explanatory knowledge bases, and subsequently leverage the same paradigm to generalise on the Question Answering task.

An empirical evaluation is performed on Transformers-based architectures adopting the WorldTree corpus as a knowledge base (Xie et al., 2020; Jansen et al., 2018) and measuring generalisation on ARC (Clark et al., 2018) and OpenbookQA

(Mihaylov et al., 2018). The main contributions of this paper are as follows:

- We propose N-XKT, a neural mechanism for encoding and transferring explanatory knowledge for science QA. To the best of our knowledge, N-XKT is the first work tackling science QA tasks through the transfer of external explanatory knowledge via neural encoding mechanisms.
- We introduce the explanatory knowledge transfer task on explanation-centred knowledge bases, describing the methodology to implement N-XKT for knowledge acquisition and downstream Question Answering using Transformer-based models as neural encoders.
- We conduct a systematic empirical analysis to demonstrate the effectiveness of N-XKT on improving downstream QA accuracy and overall convergence speed in the training phase. An ablation analysis on different types of knowledge facts is performed to measure the impact of different knowledge categories.

2 Related Work

In this section we describe several works related to knowledge-based scientific QA.

Explanation Bank Explanation Bank¹ is a core component of the WorldTree corpus (Jansen et al., 2018; Xie et al., 2020). The dataset provides explanations for multiple-choice science questions in the form of graphs connecting questions and correct answers, where multiple sentences from a knowledge base (KB) are aggregated through lexical overlap between terms. The background knowledge used for the explanations is grouped in semi-structured tables, whose facts range from common-sense to core scientific statements. Explanation Bank has been proposed for the task of explanation regeneration (Jansen and Ustalov, 2019) – i.e. given a multiple-choice science question, regenerate the gold explanation supporting the correct answer. The explanation regeneration task has been framed as an Information Retrieval (IR) problem (Valentino et al., 2021). In this paper, we aim to leverage the knowledge expressed in the explanations to enhance generalisation and zero-shot capability on multiple-choice scientific question answering.

¹<http://cognitiveai.org/explanationbank/>

Bidirectional Encoder Representations from Transformers BERT represents the foundation which defines the state-of-the-art in several NLP tasks (Devlin et al., 2019). This model adopts a Transformer-based architecture composed of several layers of attention (Vaswani et al., 2017) that are used to learn a deep bidirectional representation of language. BERT-based models have demonstrated remarkable results in Question Answering when directly fine-tuned on the answer prediction task or additionally pre-trained using domain specific knowledge (Clark et al., 2020; Beltagy et al., 2019). A recent line of research attempts to enrich the input of BERT with background knowledge in the form of explanations in order to boost generalisation and accuracy for challenging QA settings. Here, the explanations are explicitly constructed through the adoption of language models (Rajani et al., 2019) or information retrieval (IR) approaches (Valentino et al., 2021; Yadav et al., 2019). Conversely, this paper explores mechanisms to implicitly encode explanatory knowledge in the neural representation to improve the capability of performing downstream inference. Specifically, in this work, we adopt Transformers as text neural encoders.

Leveraging External Knowledge for Scientific QA Recently, many solutions have been proposed for science QA that leverage either external reference corpora (Khot et al., 2017; Khashabi et al., 2018; Zhang et al., 2018) or existing knowledge graphs (Li and Clark, 2015; Sachan et al., 2016; Wang et al., 2018; Musa et al., 2019; Zhong et al., 2019). Generally, previous works rely on Information Retrieval models or on structural embeddings for Knowledge Bases, while our work focuses on directly encoding explanatory knowledge, evaluating it in a downstream scientific QA setting.

3 Methodology

Scientific Question Answering has the distinctive property of requiring the articulation of multi-hop and explanatory reasoning. This can be contrasted with the lexical-retrieval style of factoid Question Answering. Additionally, the explanatory chains required to arrive at the correct answer typically operate at an abstract level, through the combination of definitions and scientific laws (Thayaparan et al., 2020). This characteristic makes the generalisation process more challenging, as the answer

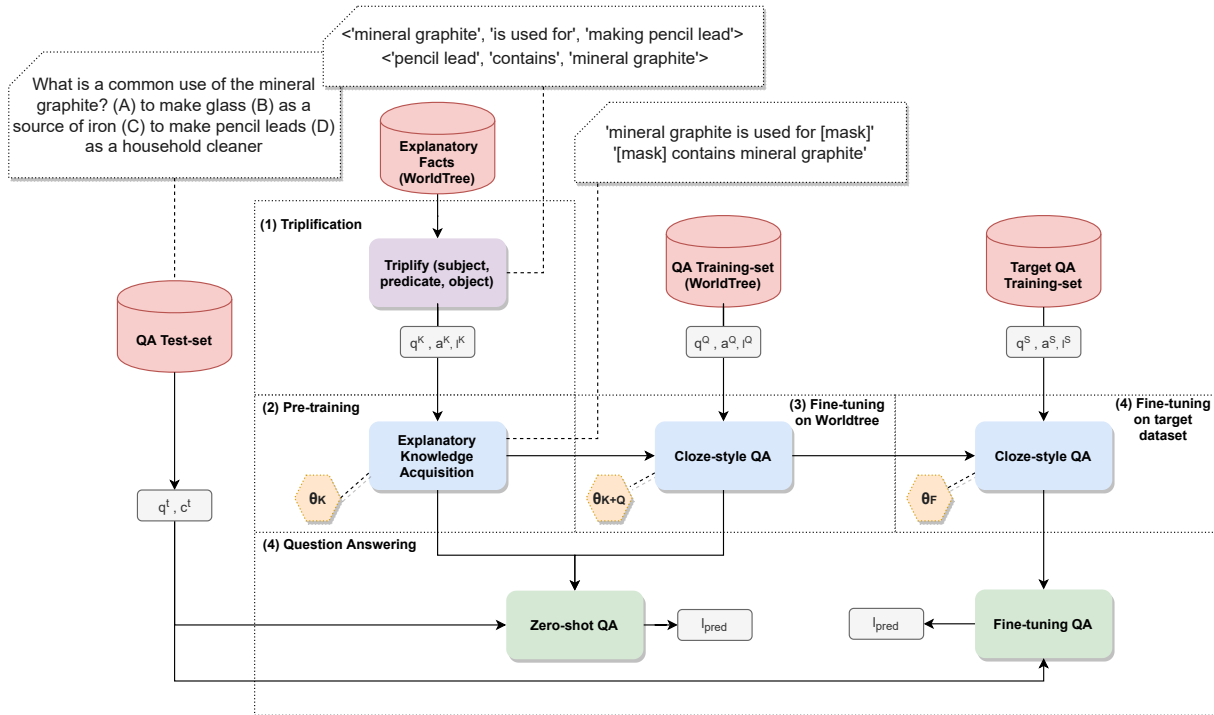


Figure 1: Outline of the proposed approach.

prediction model needs to acquire the ability to perform abstraction from the specific context in the question.

This paper hypothesises that it is possible to automatically transfer abstractive knowledge from explanatory facts into neural encoding representation for more accurate scientific QA, and for enabling zero-shot generalization. To this end, we propose N-XKT (Neural encoding based on eXplanatory Knowledge Transfer) which encodes abstractive knowledge into neural representation to improve the effectiveness in both zero-shot QA task and fine-tuning based QA task. The general neural encoding mechanism is evaluated adopting the following training tasks:

1. **Explanatory Knowledge Acquisition:** In this pre-training task, the N-XKT model encodes the explanatory textual knowledge from a set of explanatory facts into supporting embeddings. This process aims to acquire the necessary explanatory knowledge to test generalization on downstream science QA. We frame this problem as a knowledge base completion task. Specifically, after casting each explanatory fact in the knowledge base into a tuple composed of subject, object, and predicate, the model is trained on completing each fact by alternatively masking each element in

the tuple (additional details can be found in section 3.1).

2. **Cloze-style Question Answering:** To keep the encoding mechanism consistent with the pre-training explanatory knowledge acquisition task, we cast Multiple-choice Question Answering into a cloze-style QA problem. Specifically, we train the N-XKT model to complete the question with the expected candidate answer. This task aims to acquire additional knowledge for addressing downstream science QA since the patterns in the questions are typically more complex than the background explanatory facts (additional details can be found in section 3.2).

The training tasks defined above can be used to encode different types and levels of knowledge into the N-XKT model, allowing us to perform a detailed evaluation on both zero-shot and fine-tuning-based Question Answering tasks.

Figure 1 shows a schematic representation of the proposed approach.

3.1 Explanatory Knowledge Acquisition

The WorldTree corpus (Jansen et al., 2018) contains natural language explanatory facts, which are stored in semi-structured tables whose columns correspond to semantic roles. The knowledge base

contains a total of 82 tables, where each table represents a different knowledge type, with different arity and argument types. N-XKT can be used as a unified approach for transferring knowledge from heterogeneous explanatory facts via a neural encoding mechanism.

To acquire the explanatory knowledge in a unified way for subsequent transfer learning, we normalize the semi-structured facts using a binary predicate-argument structure as typical practice in standard knowledge-base completion tasks (Bordes et al., 2013; Wang et al., 2014; Lin et al., 2015). Specifically, for each table, we map the columns into three main components: subject, predicate, and object. After performing the mapping for each table in the knowledge base, we generate triples for all the facts in the knowledge base.

By framing the explanatory knowledge acquisition task as a knowledge base completion problem, we alternatively mask subjects and objects from the triples and train the model to predict the missing component in the triple by giving in input the remaining ones. Specifically, we simulate a question answering problem adopting either subject or object as an answer, and the other two components in the triple as a question.

The neural encoder of N-XKT learns an embedding representation for each pair in input. A softmax layer is added on top of the embedding to predict the probability of the missing component in the triple. The configuration adopted for the N-XKT model is described in equation 1;

$$\theta_K \leftarrow \operatorname{argmin}_{\theta} \mathcal{L}(\mathbf{N-XKT}_{\theta}(q^K, a^K), l^K) \quad (1)$$

Here, q^K and a^K represent the simulated question-answer pair generated from a generic explanatory fact triple, while l^K represents the target labels (i.e. 1 if a is the correct component for completing the triple, 0 otherwise). θ_K is the set of parameters optimised during the explanatory knowledge acquisition stage. The negative samples are generated by replacing each correct answer with a random component extracted from different explanatory facts in the knowledge base.

The transformer neural network is used as a textual neural encoder component of N-XKT, where each question-answer pair is compiled into the input token sequence:

$$[CLS][question][SEP][answer][SEP] \quad (2)$$

The final hidden vector $C \in \mathbb{R}^H$ of the Transformer neural network that corresponds to the first input

token ([CLS]) is used as an embedding to perform the final classification.

3.2 Cloze-style Question Answering

Normally, the explanatory knowledge patterns do not contain the complete information to address downstream Question Answering. However, the questions in WorldTree can be used as additional knowledge to deal with complex structured science questions, allowing N-XKT to learn to recognize more complex patterns.

To acquire additional knowledge while keeping the encoding mechanism consistent with the pre-training explanatory knowledge acquisition task, we cast Multiple-choice Question Answering into a cloze-style QA problem. The particular encoding configuration of the N-XKT model can be used in fact to address this type of question answering problems, where the model is trained to complete the question with the expected candidate answer. The detailed parameters and inputs adopted for cloze-style QA are described in equation 3:

$$\theta_{K+Q} \leftarrow \operatorname{argmin}_{\theta} \mathcal{L}(\mathbf{N-XKT}_{\theta_K}(q^Q, a^Q), l^Q) \quad (3)$$

The setting adopted for cloze-style QA is similar to the one adopted for explanatory knowledge acquisition, but with two main differences: 1) In this case, the question q^Q , the answer a^Q , and the target label l^K are generated from the WorldTree multiple-choice question answering set, where the right candidate answer of each question acts as a positive sample, and the incorrect candidate answers act as the negative samples. 2) The initial parameters are initially set with θ_K , that is, we adopt the parameters that have been optimised during the explanatory knowledge acquisition stage.

3.3 Zero-shot and Fine-tuning Settings

Given a multiple-choice science question, N-XKT can perform question answering by framing it as a sequence classification problem, where the question is paired with each candidate answer to compute a probability score. The candidate choice with highest score can then be selected as the predicted answer. We evaluate N-XKT in two different settings: zero-shot and fine-tuning-based QA.

Regarding the **zero-shot setting**, the N-XKT is trained only on the explanatory knowledge acquisition task and then directly tested on downstream Question Answering. We also evaluate the model

trained jointly on explanatory knowledge and science questions in WorldTree, evaluating its generalization capabilities on different multiple-choice Question Answering datasets, such as ARC² (Clark et al., 2018) and OpenBook QA³ (Mihaylov et al., 2018). For each pair of question and candidate answer, the scores are computed as described in equation 4. Here, (q^T, c^T) represent the test question and a candidate answer, while l_{pred}^T is the score predicted by the model.

$$l_{pred}^T = \mathbf{N}\text{-}\mathbf{XKT}_{\theta_{K+Q}}(q^T, c^T) \quad (4)$$

In the **fine-tuning setting**, the N-XKT model is additionally fine-tuned on each target QA dataset as in equation 6. Here, (q^S, a^S) represents a question-answer pair from the target QA training set, while l^S is the label indicating whether the answer is correct or not.

$$\theta_F \leftarrow \operatorname{argmin}_{\theta} \mathcal{L}(\mathbf{N}\text{-}\mathbf{XKT}_{\theta_{K+Q}}(q^S, a^S), l^S) \quad (5)$$

As shown in equation 6, we adopt the same configuration as in the zero-shot setting, where the only difference is represented by the fine-tuned parameters set θ_F :

$$l_{pred}^T = \mathbf{N}\text{-}\mathbf{XKT}_{\theta_F}(q^T, c^T) \quad (6)$$

4 Empirical Evaluation

We conduct our experiments on four widely used science QA datasets, WorldTree V2.0 (Xie et al., 2020), ARC Easy and Challenge (Clark et al., 2018), and Openbook QA (Mihaylov et al., 2018). The results tend to confirm our research hypothesis that explanatory knowledge encoding can improve generalization in downstream science Question Answering (QA) tasks. Furthermore, we systematically analyze several factors which may have an impact on the final results, including the use of Transformer-based models with a larger number of parameters (BERT-large), testing the model on QA tasks using different types of explanatory background knowledge, and measuring training and test performance by further fine-tuning the model on other datasets.

4.1 Experimental Setup

QA dataset size. In order to conduct a thorough quantitative analysis, we use four science QA

²<https://allenai.org/data/arc>

³<https://allenai.org/data/open-book-qa>

Table 1: QA datasets size.

Dataset	#Train	#Dev	#Test
WorldTree V2.0	3,947	1,019	4,165
ARC Easy	2,251	570	2,376
ARC Challenge	1,119	299	1,172
Openbook QA	4,957	500	500

Table 2: Number of instances in each explanatory knowledge category.

Type	Size
All	9,701
Retrieval	7,006
Inference-supporting	1,670
Complex Inference	1,025

datasets, WorldTree V2.0 (Xie et al., 2020), ARC Easy and Challenge (Clark et al., 2018), and Openbook QA (Mihaylov et al., 2018). The number of question-answer pairs in each dataset is listed in Table 1.

Explanatory knowledge dataset size. We encode different types of explanatory knowledge in the WorldTree corpus into Transformer neural networks. The statistics of the adopted explanatory facts are reported in Table 2. Because we further analyze the impact of different types of knowledge, the number of each knowledge type is also given in the table.

Hyperparameters configuration. We adjust two major hyperparameters for the training of the model, namely batch size and learning rate. We optimize the parameters considering the following combinations: we adopt training batch sizes in $\{16, 32\}$, and learning rate in $\{1e-5, 3e-5, 5e-5\}$. The best results are obtained with batch size 32 and learning rate $3e-5$ for the BERT-base model, and batch size 16 and learning rate $1e-5$ for BERT-large (Devlin et al., 2019).

Information Retrieval baseline. We adopt an Information Retrieval (IR) baseline similar to the one described in Clark et al. (2018). Given a question q , for each candidate answer $c_i \in \mathcal{C} = \{c_1, \dots, c_n\}$, the IR solver uses BM25 vectors and cosine similarity to retrieve the top K sentences in the WorldTree corpus that are most similar to the concatenation of q and c_i . The score of a candidate answer c_i is then obtained by considering the sum of the BM25 relevance scores associated to

Table 3: N-XKT Question Answering accuracy results.

Config	Explanation Bank		ARC Easy		ARC Challenge		Openbook QA	
	Dev	Test	Dev	Test	Dev	Test	Dev	Test
IR BM25 (K = 5)	50.29%	44.55%	54.56%	50.00%	37.46%	31.14%	24.80%	26.80%
K base	49.30%	44.74%	50.18%	50.89%	34.38%	33.17%	30.96%	32.72%
Q base	44.86%	40.34%	50.81%	47.43%	24.41%	26.86%	27.92%	33.12%
K+Q base	58.14%	50.42%	58.53%	57.98%	37.46%	35.87%	35.32%	37.60%
K large	51.62%	45.85%	52.81%	52.58%	37.53%	33.07%	31.72%	34.12%
Q large	47.54%	43.47%	53.61%	51.41%	27.09%	28.63%	28.24%	36.04%
K+Q large	60.16%	50.98%	61.19%	58.24%	39.00%	37.63%	35.64%	38.20%
base FT	-	-	53.61%	53.82%	36.72%	32.71%	53.64%	53.16%
K base FT	-	-	53.61%	52.81%	35.79%	34.90%	53.60%	54.60%
Q base FT	-	-	59.05%	58.44%	33.65%	35.09%	56.04%	57.08%
K+Q base FT	-	-	59.33%	58.79%	38.13%	38.09%	56.12%	56.56%

Table 4: Question Answering accuracy results using different explanatory knowledge categories.

Knowledge	Config	Explanation Bank		ARC Easy		ARC Challenge		Openbook QA	
		Dev	Test	Dev	Test	Dev	Test	Dev	Test
None	Q base	44.86%	40.34%	50.81%	47.43%	24.41%	26.86%	27.92%	33.12%
Retrieval	K base	39.05%	38.72%	44.42%	45.25%	23.75%	26.25%	27.12%	29.96%
	K+Q base	51.00%	46.08%	51.79%	53.22%	34.65%	33.00%	31.96%	32.96%
Inference-supporting	K base	41.60%	38.24%	45.96%	44.77%	26.09%	26.02%	27.40%	30.88%
	K+Q base	52.72%	47.33%	54.35%	54.32%	34.85%	34.40%	33.64%	37.16%
Complex inference	K base	41.01%	38.58%	46.32%	45.98%	24.95%	23.75%	26.96%	29.76%
	K+Q base	52.99%	46.12%	55.30%	52.74%	34.78%	34.51%	32.08%	35.08%
All	K base	49.30%	44.74%	50.18%	50.89%	34.38%	33.17%	30.96%	32.72%
	K+Q base	58.14%	50.42%	58.53%	57.98%	37.46%	35.87%	35.32%	37.60%

the retrieved sentences. The predicted answer corresponds to the candidate choice with the highest score. To test the generalisation of this approach on ARC and OpenbookQA, we keep the same background knowledge throughout the experiments.

Configuration Setting. We adopt different configurations in the experiments to control for training data, Transformer model, and target QA test dataset fine-tuning. We report the different configurations in the ‘‘Config’’ column of Table 6 and Table 7. The label ‘‘K’’ indicates that the model is trained only on the explanatory knowledge acquisition task, ‘‘Q’’ means that the model is trained only on the cloze-style QA task using WorldTree as reference dataset, ‘‘K+Q’’ means that the model is pre-trained for explanatory knowledge acquisition and then further fine-tuned on cloze-style QA (again using only WorldTree as training dataset). Moreover, ‘‘base’’ means using BERT-base as Transformer model, while ‘‘large’’ means using BERT-large. Finally, ‘‘FT’’ means that the model is additionally fine-tuned on the target QA dataset’s training data.

4.2 Overall Results on Zero-shot Science Question Answering

In Table 6, we report the performance of N-XKT under different configurations along with the accuracy of the BM25 baseline with $K = 5$ number of facts. The models are tested across multiple QA datasets including WorldTree, ARC, and OpenbookQA.

From the results, we derive the following conclusions. First, the proposed N-XKT model can clearly achieve better accuracy than the BM25 baseline since N-XKT uses Transformer-based neural mechanisms to acquire and encode external knowledge. Second, using BERT-large instead of BERT-base as initial Transformer can improve the performance since BERT-large contains more parameters than BERT-base. However, we found that the advantage of using BERT-large is not significant since more parameters implies more resources needed for training. Third, we observe that N-XKT obtains better performance than pre-trained BERT when fine-tuning on the target datasets.

Table 5: Accuracy comparison between N-XKT and other approaches. External KB adopted by the models: 1.ARC-corpus (Clark et al., 2018), 2.ConceptNet (Speer et al., 2017), 3.Wikipedia (<https://www.wikipedia.org/>), 4.SciTail (Khot et al., 2018), 5.SNLI (Bowman et al., 2015), 6.MultiNLI (Williams et al., 2018), 7.RACE (Lai et al., 2017), 8.MCScript (Ostermann et al., 2018), 9.WorldTree (Jansen et al., 2018).

	ARC Easy	ARC Challenge	Openbook QA	External KB	IR-based	Fine-tuned
IR BM25 (K = 5)	50.00%	31.14%	26.80%	9	yes	no
Clark et al. (2018)	62.60%	20.30%	-	1	yes	yes
Mihaylov et al. (2018)	-	-	50.20%	2, 3	yes	yes
Khot et al. (2018)	59.00%	27.10%	24.40%	4	yes	yes
Zhang et al. (2018)	-	31.70%	-	1	no	yes
Yadav et al. (2018)	58.40%	26.60%	-	none	no	yes
Musa et al. (2019)	52.20%	33.20%	-	1	yes	yes
Zhong et al. (2019)	-	33.40%	-	2	no	yes
Pirtoacă et al. (2019)	61.10%	26.90%	-	4, 5, 6	no	yes
Ni et al. (2019)	-	36.60%	-	7, 8	no	yes
<i>GPT^{II}</i> (Radford, 2018)	57.00%	38.20%	52.00%	7	no	yes
<i>RS^{II}</i> (Sun et al., 2019)	66.60%	40.70%	55.20%	7	no	yes
N-XKT K+Q base (ours)	57.98%	35.87%	37.60%	9	no	no

4.3 Ablation Analysis on Impact of Different Explanatory Knowledge Types

To understand the impact of different types of explanation on the final accuracy, we breakdown the facts stored in the knowledge base using three different categories (i.e., retrieval, inference-supporting and complex inference) and rerun the training of the N-XKT model using only one category per time.

The adopted categories are provided in the WorldTree corpus and can be described as follows:

- *Retrieval*: facts expressing knowledge about taxonomic relations and/or properties.
- *Inference-Supporting*: Facts expressing knowledge about actions, affordances, requirements.
- *Complex Inference*: Facts expressing knowledge about causality, processes, and if/then relationships.

The obtained accuracy is showed in Table 7. The results highlight the importance of using all the explanation categories to achieve the final accuracy for the combined approach. However, the retrieval category seems to have a higher impact on the generalisation. We believe that this result is due to the taxonomic knowledge encoded in the retrieval category (i.e. “*x is a kind of y*”), which facilitates the acquisition of the implicit explanatory capabilities necessary for answering science questions.

In Table 7, we compare the impact of different explanatory knowledge types and get the following conclusion. 1) All three types of explanatory

knowledge are helpful for further science QA task. The results using all three types of knowledge are significantly better than the results obtained when using no explanatory knowledge at all (first row in Table 7). 2) The model trained on all explanatory knowledge outperforms the models using each individual type of knowledge alone, confirming that different types of knowledge are complementary for achieving the final performance.

4.4 Evaluating Zero-shot N-XKT with Start-of-the-art baselines

In Table 5, we evaluate several start-of-the-art methods as baselines along with N-XKT trained only on the WorldTree. The table reports the accuracy results on ARC and OpenbookQA. In the “External KB” column, we list the external Knowledge Bases (KB) adopted by different models. The “IR-based” column indicates whether the model adopts Information Retrieval (IR) techniques, and the “Fine-tuned” column indicates whether the approach is fine-tuned on the target dataset.

Table 5 is intended to provide a general comparative analysis between N-XKT and the baseline models, most of them fine-tuned on the target datasets. N-XKT is able to achieve comparable performance under a transfer learning setting. The generalization performance of the proposed model is more noticeable for the ARC Challenge dataset, which requires the implicit encoding of more complex explanatory knowledge.




Figure 2: Convergence curve when fine-tuning different version of N-XTK on the target QA datasets.

4.5 Improvement on Fine-tuning Convergence

In Figure 2, we visualize the convergence curve for the fine-tuning over three science QA tasks (ARC Easy, ARC Challenge and OpenBookQA), comparing a pure BERT-based N-XKT model with a pre-trained N-XKT models using different configurations, AFK (pre-trained on explanatory knowledge acquisition), QAP (pre-trained on WorldTree cloze-style QA), AFK+QAP (pre-trained on both). It is noticeable that the encoding of explanatory knowledge impacts the convergence of the model for all three datasets, with a particular emphasis on the two ARC variants.

5 Conclusion

In this paper, we proposed a neural encoding mechanism for explanatory knowledge acquisition and transfer, N-XKT. We evaluated the impact of the encoding mechanism on downstream science QA. The proposed model delivers better generalisation and accuracy for QA tasks that require multi-hop and explanatory inference. The proposed encoding mechanism can be used to deliver zero-shot inference capabilities, providing comparable performance when compared to supervised models on QA. These results supports the hypothesis that pre-training tasks targeting abstract and explanatory knowledge acquisition can constitute and impor-

tant direction to improve inference capabilities and generalization of state-of-the-art neural models.

References

- Akari Asai, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caiming Xiong. 2020. [Learning to retrieve reasoning paths over wikipedia graph for question answering](#). In *International Conference on Learning Representations*.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). 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 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#).
- Peter Clark, Oren Etzioni, Tushar Khot, Daniel Khashabi, Bhavana Mishra, Kyle Richardson, Ashish Sabharwal, Carissa Schoenick, Carissa Schoenick, Oyvind Tafjord, Niket Tandon, Sumithra Bhakthavatsalam, Dirk Groeneveld, Michal Guerquin, and Michael Schmitz. 2020. [From ‘f’ to ‘a’ on the n.y. regents science exams: An overview of the aristo project](#). *AI Magazine*, 41(4):39–53.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ming Ding, Chang Zhou, Qibin Chen, Hongxia Yang, and Jie Tang. 2019. [Cognitive graph for multi-hop reading comprehension at scale](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2694–2703, Florence, Italy. 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.
- Peter Jansen, Niranjana Balasubramanian, Mihai Surdeanu, and Peter Clark. 2016. [What’s in an explanation? characterizing knowledge and inference requirements for elementary science exams](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2956–2965, Osaka, Japan. The COLING 2016 Organizing Committee.
- Peter Jansen and Dmitry Ustalov. 2019. [TextGraphs 2019 shared task on multi-hop inference for explanation regeneration](#). In *Proceedings of the Thirteenth Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-13)*, pages 63–77, Hong Kong. Association for Computational Linguistics.
- Peter Jansen, Elizabeth Wainwright, Steven Marmorstein, and Clayton Morrison. 2018. [WorldTree: A corpus of explanation graphs for elementary science questions supporting multi-hop inference](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Daniel Khashabi, Tushar Khot, Ashish Sabharwal, and D. Roth. 2018. [Question answering as global reasoning over semantic abstractions](#). In *AAAI*.
- Tushar Khot, Peter Clark, Michal Guerquin, Peter Jansen, and Ashish Sabharwal. 2020. [Qasc: A dataset for question answering via sentence composition](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8082–8090.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2017. [Answering complex questions using open information extraction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 311–316, Vancouver, Canada. Association for Computational Linguistics.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2018. [Scitail: A textual entailment dataset from science question answering](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5189–5197. AAAI Press.

- Souvik Kundu, Tushar Khot, Ashish Sabharwal, and Peter Clark. 2019. [Exploiting explicit paths for multi-hop reading comprehension](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2737–2747, Florence, Italy. Association for Computational Linguistics.
- 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.
- Yang Li and Peter Clark. 2015. [Answering elementary science questions by constructing coherent scenes using background knowledge](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2007–2012, Lisbon, Portugal. Association for Computational Linguistics.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI’15, page 2181–2187. AAAI Press.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*.
- Ryan Musa, Xiaoyan Wang, Achille Fokoue, Nicholas Mattei, Maria Chang, Pavan Kapanipathi, Bassem Makni, Kartik Talamadupula, and Michael Witbrock. 2019. [Answering science exam questions using query reformulation with background knowledge](#). In *Automated Knowledge Base Construction (AKBC)*.
- Jianmo Ni, Chenguang Zhu, Weizhu Chen, and Julian McAuley. 2019. [Learning to attend on essential terms: An enhanced retriever-reader model for open-domain question answering](#). 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 335–344, Minneapolis, Minnesota. Association for Computational Linguistics.
- Simon Ostermann, Ashutosh Modi, Michael Roth, Stefan Thater, and Manfred Pinkal. 2018. [MCScript: A novel dataset for assessing machine comprehension using script knowledge](#).
- George-Sebastian Pîrtoacă, Traian Rebedea, and Ștefan Rușeți. 2019. [Improving retrieval-based question answering with deep inference models](#). In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.
- Lin Qiu, Yunxuan Xiao, Yanru Qu, Hao Zhou, Lei Li, Weinan Zhang, and Yong Yu. 2019. [Dynamically fused graph network for multi-hop reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6140–6150, Florence, Italy. Association for Computational Linguistics.
- A. Radford. 2018. Improving language understanding by generative pre-training.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Explain yourself! leveraging language models for commonsense reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4932–4942, Florence, Italy. Association for Computational Linguistics.
- Mrimmaya Sachan, Kumar Dubey, and Eric Xing. 2016. [Science question answering using instructional materials](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 467–473, Berlin, Germany. Association for Computational Linguistics.
- 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*, AAAI’17, page 4444–4451. AAAI Press.
- Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. 2019. [Improving machine reading comprehension with general reading strategies](#). 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 2633–2643, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mokanarangan Thayaparan, Marco Valentino, and André Freitas. 2020. [A survey on explainability in machine reading comprehension](#).
- Mokanarangan Thayaparan, Marco Valentino, Viktor Schlegel, and André Freitas. 2019. [Identifying supporting facts for multi-hop question answering with document graph networks](#). In *Proceedings of the Thirteenth Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-13)*, pages 42–51, Hong Kong. Association for Computational Linguistics.
- Marco Valentino, Mokanarangan Thayaparan, and André Freitas. 2021. [Unification-based reconstruction of multi-hop explanations for science questions](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 200–211, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz

- Kaiser, and Illia Polosukhin. 2017. *Attention is all you need*. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Liang Wang, Meng Sun, Wei Zhao, Kewei Shen, and Jingming Liu. 2018. *Yuanfudao at SemEval-2018 task 11: Three-way attention and relational knowledge for commonsense machine comprehension*. pages 758–762.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, AAAI’14, page 1112–1119. AAAI Press.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. *A broad-coverage challenge corpus for sentence understanding through inference*. 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 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Zhengnan Xie, Sebastian Thiem, Jaycie Martin, Elizabeth Wainwright, Steven Marmorstein, and Peter Jansen. 2020. *WorldTree v2: A corpus of science-domain structured explanations and inference patterns supporting multi-hop inference*. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5456–5473, Marseille, France. European Language Resources Association.
- Vikas Yadav, Steven Bethard, and Mihai Surdeanu. 2019. *Quick and (not so) dirty: Unsupervised selection of justification sentences for multi-hop question answering*. 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 2578–2589, Hong Kong, China. Association for Computational Linguistics.
- Vikas Yadav, Rebecca Sharp, and M. Surdeanu. 2018. Sanity check: A strong alignment and information retrieval baseline for question answering. *The 41st International ACM SIGIR Conference on Research Development in Information Retrieval*.
- 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.
- Y. Zhang, H. Dai, Kamil Toraman, and L. Song. 2018. Kg2: Learning to reason science exam questions with contextual knowledge graph embeddings. *ArXiv*, abs/1805.12393.
- Wanjun Zhong, Duyu Tang, Nan Duan, M. Zhou, Jiahai Wang, and J. Yin. 2019. Improving question answering by commonsense-based pre-training. In *NLPCC*.

A Hyperparameters tuning

The N-XKT mainly use a transformer network as natural language encoder component, the hyperparameters of transformer network training have been tuned manually for the optimisation is the maximisation of the accuracy in answer prediction. Specifically, 3 parameters should be set for training, train batch size β , learning rate α , and train epoch \mathcal{N} . The values used in pre-training on explanation knowledge base are as follows:

- $\beta = 32$
- $\alpha = 5e-5$
- $\mathcal{N} = 5$

The values used in fine-tuning on Question Answer are as follows:

- $\beta = 32$
- $\alpha = 5e-5$
- $\mathcal{N} = 40$

B Data

We use two versions of Explanation Bank Scientific Question Answer datasets in this paper. The version 1 of Explanation Bank dataset can be downloaded at the following URL: http://cognitiveai.org/dist/worldtree_corpus_textgraphs2019sharedtask_withgraphvis.zip. The version 2 of Explanation Bank dataset is available at the following URL: <https://github.com/cognitiveailab/tg2020task>.

C Computing Infrastructure

To accelerate the training process of the experiments, we adopt a NVIDIA Tesla V100 GPU.

D Accuracy Results Including Standard Deviation

We repeat the N-XKT model Question Answering training process on all the dataset for 5 times, each time with random parameters initialization. Addition to the tables provided in paper, we report the detailed results with standard deviation in following tables.

Table 6: N-XKT Question Answering accuracy result comparison

Config	Explanation Bank		ARC Easy		ARC Challenge		Openbook QA	
	Dev	Test	Dev	Test	Dev	Test	Dev	Test
IR BM25 (K = 5)	50.29%	44.55%	54.56%	50.00%	37.46%	31.14%	24.80%	26.80%
K base	49.30%	44.74%	50.18%	50.89%	34.38%	33.17%	30.96%	32.72%
	± 0.0238	± 0.0166	± 0.0167	± 0.0198	± 0.0255	± 0.0165	± 0.0359	± 0.0273
Q base	44.86%	40.34%	50.81%	47.43%	24.41%	26.86%	27.92%	33.12%
	± 0.0229	± 0.0087	± 0.0258	± 0.0136	± 0.0101	± 0.0049	± 0.0342	± 0.0176
K+Q base	58.14%	50.42%	58.53%	57.98%	37.46%	35.87%	35.32%	37.60%
	± 0.0119	± 0.0039	± 0.0047	± 0.0014	± 0.0135	± 0.0149	± 0.0124	± 0.0085
K large	51.62%	45.85%	52.81%	52.58%	37.53%	33.07%	31.72%	34.12%
	± 0.0159	± 0.0089	± 0.004	± 0.0136	± 0.0109	± 0.0129	± 0.0199	± 0.0232
Q large	47.54%	43.47%	53.61%	51.41%	27.09%	28.63%	28.24%	36.04%
	± 0.0131	± 0.0061	± 0.0176	± 0.0073	± 0.012	± 0.0125	± 0.0118	± 0.0167
K+Q large	60.16%	50.98%	61.19%	58.24%	39.00%	37.63%	35.64%	38.20%
	± 0.0168	± 0.0102	± 0.0108	± 0.0076	± 0.0268	± 0.0155	± 0.0076	± 0.0161
base FT	-	-	53.61%	53.82%	36.72%	32.71%	53.64%	53.16%
	-	-	± 0.0168	± 0.0093	± 0.0104	± 0.0086	± 0.0182	± 0.0223
K base FT	-	-	53.61%	52.81%	35.79%	34.90%	53.60%	54.60%
	-	-	± 0.0159	± 0.0241	± 0.0218	± 0.0239	± 0.0248	± 0.0281
Q base FT	-	-	59.05%	58.44%	33.65%	35.09%	56.04%	57.08%
	-	-	± 0.0177	± 0.0070	± 0.0280	± 0.0065	± 0.0126	± 0.0178
K+Q base FT	-	-	59.33%	58.79%	38.13%	38.09%	56.12%	56.56%
	-	-	± 0.0187	± 0.0087	± 0.0224	± 0.0124	± 0.0186	± 0.0111

Table 7: Question Answering accuracy result in different abstractive knowledge categories

Knowledge	Config	Explanation Bank		ARC Easy		ARC Challenge		Openbook QA	
		Dev	Test	Dev	Test	Dev	Test	Dev	Test
None	Q base	44.86%	40.34%	50.81%	47.43%	24.41%	26.86%	27.92%	33.12%
		± 0.0229	± 0.0087	± 0.0258	± 0.0136	± 0.0101	± 0.0049	± 0.0342	± 0.0176
RET	K base	39.05%	38.72%	44.42%	45.25%	23.75%	26.25%	27.12%	29.96%
		± 0.0258	± 0.0106	± 0.011	± 0.0139	± 0.0165	± 0.0141	± 0.0099	± 0.0202
	K+Q base	51.00%	46.08%	51.79%	53.22%	34.65%	33.00%	31.96%	32.96%
		± 0.0173	± 0.0135	± 0.0178	± 0.0141	± 0.0321	± 0.0128	± 0.0192	± 0.0182
INSUPP	K base	41.60%	38.24%	45.96%	44.77%	26.09%	26.02%	27.40%	30.88%
		± 0.0149	± 0.0075	± 0.0127	± 0.0118	± 0.0164	± 0.0099	± 0.0168	± 0.0122
	K+Q base	52.72%	47.33%	54.35%	54.32%	34.85%	34.40%	33.64%	37.16%
		± 0.0247	± 0.0062	± 0.0206	± 0.0092	± 0.031	± 0.0128	± 0.0279	± 0.0306
COMPLEX	K base	41.01%	38.58%	46.32%	45.98%	24.95%	23.75%	26.96%	29.76%
		± 0.0132	± 0.0035	± 0.0134	± 0.0091	± 0.0263	± 0.0066	± 0.012	± 0.0163
	K+Q base	52.99%	46.12%	55.30%	52.74%	34.78%	34.51%	32.08%	35.08%
		± 0.0098	± 0.0131	± 0.0081	± 0.0087	± 0.0112	± 0.0194	± 0.018	± 0.0153
All	K base	49.30%	44.74%	50.18%	50.89%	34.38%	33.17%	30.96%	32.72%
		± 0.0238	± 0.0166	± 0.0167	± 0.0198	± 0.0255	± 0.0165	± 0.0359	± 0.0273
	K+Q base	58.14%	50.42%	58.53%	57.98%	37.46%	35.87%	35.32%	37.60%
		± 0.0119	± 0.0039	± 0.0047	± 0.0014	± 0.0135	± 0.0149	± 0.0124	± 0.0085

Tab. 6 is for overall accuracy of N-XKT model on QA tasks, and Tab. 7 is for ablation analysis results, only use part of explanations in training process.