

Database Reasoning Over Text

James Thorne
University of Cambridge
jt719@cam.ac.uk

Majid Yazdani
Facebook AI
myazdani@fb.com

Marzieh Saeidi
Facebook AI
marzieh@fb.com

Fabrizio Silvestri
Sapienza University, Rome
fsilvestri@diag.uniroma1.it

Sebastian Riedel
Facebook AI
University College London
sriedel@fb.com

Alon Halevy
Facebook AI
ayh@fb.com

Abstract

Neural models have shown impressive performance gains in answering queries from natural language text. However, existing works are unable to support database queries, such as “*List/Count all female athletes who were born in 20th century*”, which require reasoning over *sets* of relevant facts with operations such as join, filtering and aggregation. We show that while state-of-the-art transformer models perform very well for small databases, they exhibit limitations in processing noisy data, numerical operations, and queries that aggregate facts. We propose a modular architecture to answer these database-style queries over multiple spans from text and aggregating these at scale. We evaluate the architecture using WIKINLDB,¹ a novel dataset for exploring such queries. Our architecture scales to databases containing thousands of facts whereas contemporary models are limited by how many facts can be encoded. In direct comparison on small databases, our approach increases overall answer accuracy from 85% to 90%. On larger databases, our approach retains its accuracy whereas transformer baselines could not encode the context.

1 Introduction

Question answering (QA) over text has made significant strides in recent years owing to the availability of new datasets and models. Machines have surpassed human performance on the well-known SQuAD task (Rajpurkar et al., 2016) where models extract answer spans from a short passage of text. The subsequent body of work has further considered incorporating retrieval from large corpora such as Wikipedia (Dhingra et al., 2017; Joshi et al., 2017; Kwiatkowski et al., 2019) to identify relevant information, conditioning answer generation (Chen

¹<https://github.com/facebookresearch/NeuralDB>

Facts: (8 of 500 shown)

- Nicholas lives in Washington D.C. with his wife.
- Sheryl is Nicholas’s wife.
- Teuvo was born in 1912 in Ruskala.
- Sheryl’s mother gave birth to her in 1978.
- Nicholas is a doctor.
- Sarah was born in Chicago in 1982.
- Sarah married John in 2010.
- Sarah works in a hospital in NY as a doctor.

Queries:

List everyone born before 1980.
(Set) → Sheryl, Teuvo, ...

Whose spouse is a doctor?
(Join) → Sheryl, John, ...

Who is the oldest person?
(Max) → Teuvo

Who is Sheryl’s mother?
(Set) → NULL

Figure 1: Examples of set and aggregation queries over a natural language database: a database where facts are stored in free-form text without the need for a schema.

et al., 2017; Lewis et al., 2020b; Izacard and Grave, 2020). More sophisticated architectures have been proposed with incremental retrieval for multi-hop QA (Xiong et al., 2020; Das et al., 2019), where several passages are required, which may have low lexical or semantic similarity with the question.

This paper considers the problem of answering questions similar to database queries, such as those shown in Figure 1. For example, the query “*List all the female athletes in Wikipedia who were born in the 20th century*”, requires reasoning over hundreds or thousands of facts, retrieved from multiple Wikipedia pages, and applying set-based filters to them (e.g., gender, birth date). If our query further asked how many such athletes exist, we would

have to perform an aggregation function to count the result set. The ability to answer the aforementioned queries would enable a new kind of database (Thorne et al., 2021) where facts can be described in natural language and would therefore obviate the need for a pre-defined schema, which is a major limitation of current database systems. An example application for such flexible text databases exists in the area of storing knowledge for personal assistants where users store data about their habits and experiences, their friends and their preferences, for which designing a schema is impractical.

We introduce WIKINLDB, a benchmark dataset for exploring database reasoning over facts expressed in natural language. WIKINLDB contains a number of query types that require systems to return large set-based answers and aggregate over these (with operators such as `count`, `min`, and `max`). Our dataset is generated using publicly available knowledge graph data, enabling large volumes of instances to be generated with minimal effort. Most queries in WIKINLDB require reasoning over hundreds of facts to generate answers, exposing limitations in current neural models. In contrast to DROP (Dua et al., 2019) where queries are answered over *single* passages, and bAbI (Weston et al., 2015), where each query is based on a context of less than 20 facts, our dataset scales from databases of 25 instances to 1000, and could be extended further.

We also introduce a modular architecture to support database reasoning over text and characterize its behavior on our reference dataset. We find that even on small databases of 25 facts, naive application of transformers is insufficient. When provided with only the relevant facts, the baseline yields an answer accuracy of 85%, whereas applying our proposed architecture yields 90% by better answering queries, such as `count`, that require computation. It is well known that transformer models do not scale well to large inputs due to the use of self-attention. We found that mechanisms such as Fusion in Decoder (Izacard and Grave, 2020, FiD) and LongFormer (Beltagy et al., 2020), which mitigate the scaling issue, harm the model: combining more than 2 facts with FiD resulted in answer accuracies of 76% and 39%, respectively. These issues were mitigated by our approach which generates intermediate query-based derivations of small numbers of facts in the database, before using conventional computation to aggregate the results.

2 Answering Database Queries over Text

2.1 Problem Definition

We refer to corpora that consist of unordered collections of facts expressed as short natural language sentences as Natural Language Databases (NLDBs). For example, a corpus may include all the utterances given to a personal assistant by its user, or all the claims uttered by a political figure. The texts in our corpora are similar to databases as they are sets of stand-alone facts. But unlike a database, they are not expressed as rows or triples in a pre-defined schema. For example, a sentence containing a single fact, “*Gustavo likes espresso*” or multiple facts, such as “*Robertson Howard, who attended the University of Virginia, is buried in the Congressional Cemetery*”.

A query Q over a database, D , produces a set of answers: $Q(D) = \{a_1, \dots, a_l\}$. We consider the following four query types (see examples in Table 5): (1) *Set queries* are extractive queries that return a list of spans, such as entities, from the facts. (2) *Boolean queries* return a True/False answer. (3) *Aggregation queries* require computation over answer sets with an operator, such as `count`, `min` and `max`. For example: “*How many people work for Yale Law School?*”). (4) *Join queries* require the combination of two (or more) facts to produce each answer. We combine join operations with set, Boolean and aggregation queries. For example, the query “*Who works in a company in France?*” considers both the relationship between people and employer as well as company locations.

2.2 Challenges

The NLP treatment of question answering, where systems encode the query and context (containing the background knowledge), forms a good starting point for NLDBs. Common model architectures are based on the transformer (Vaswani et al., 2017) in an encoder-decoder configuration. The encoder uses self-attention to conditionally encode the context with the query and the decoder allows conditional generation of outputs that are not necessarily present in the input. To scale question answering to reason over large knowledge-sources such as Wikipedia, task formulations typically retrieve text-spans from a corpus to condition answer generation (Chen et al., 2017; Dhingra et al., 2017). However, several challenges encountered in NLDBs preclude direct application of these techniques:

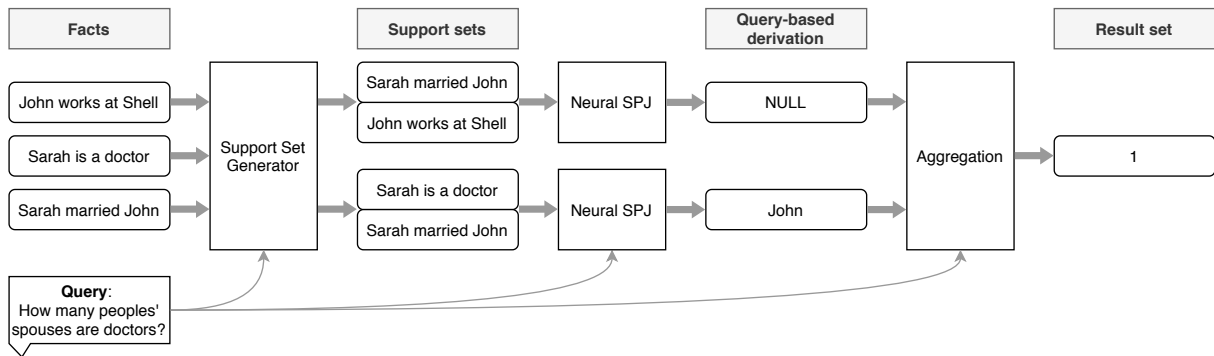


Figure 2: Overview of the proposed architecture. Consisting of a support set generator, SPJ and aggregation

Scale To scale neural reasoning to databases of non-trivial size, it would not be feasible to encode the entire database as input to the transformer. Question answering systems combine a retrieval mechanism to select relevant spans from knowledge sources as context. This task is usually referred to as open-domain QA (Lewis et al., 2020a; Izacard and Grave, 2020). It is common to use a maximum input size of 512 or 1024 tokens for context. While extensions such as Linformer (Wang et al., 2020), Longformer (Beltagy et al., 2020) and Fusion in Decoder (Izacard and Grave, 2020) enable larger contexts to be encoded, their application of self-attention varies and the number of tokens that may be encoded is limited by GPU memory.

Multiple answer spans The NLP formulation of question answering typically requires extracting a span from a single document or generating a short answer. Answering queries in a NLDB may require processing a large number of facts, generating a large number of items as answer, hundreds or thousands, and performing aggregations over large sets.

Locality and document structure NLDBs do not enjoy the locality properties that usually hold in open-domain QA. In NLDBs, a query may be dependent on multiple facts that can be anywhere in the database. In fact, by definition, the current facts in a database can be reordered and the query answers should not change. In contrast, in open-domain QA, the fact needed to answer a given question is typically located in a paragraph or document with multiple sentences about the same subject, in combination with a document title, where this additional context may help information recall.

Conditional retrieval Similar to open-domain question answering, NLDBs mandate an information retrieval component. When determining which

facts to input to the model, NLDBs may require conditional retrieval from the database. For example, to answer the query “*Whose spouse is a doctor?*” we’d first need to fetch spouses and then their professions. Recent work on multi-hop query answering (e.g., Asai et al. (2019)), has started considering this issue but is restricted to the case where we’re looking for a single answer. In NLDBs, we may need to perform multi-hops for sets of facts.

3 Architecture for querying NLDBs

To address the aforementioned challenges, we propose an instance of a Neural Database architecture (Thorne et al., 2021) that operates over textual facts with parallelizable non-blocking operators before aggregating the results. The three core components of the architecture, shown in Figure 2, are a Support Set Generator (SSG) which retrieves small sets of relevant facts called support sets, a parallelizable non-blocking Select-Project-Join (SPJ) operator which generates intermediate answers that can be unioned to produce the final answer, and an optional aggregation stage which uses conventional computation to perform numerical reasoning. The key insight underlying our architecture is to leverage neural models for what they excel at, namely, reasoning over a small set of facts.

Neural SPJ Operator Given a single support set and a query, the SPJ (Select-Project-Join) operator outputs a machine readable intermediate representation of the answer that can be generated from the support set. For example, given the query “*Who was born in Montevideo?*” and the support set {“*Mario Sagario was born in Montevideo, Uruguay, ...*”}, the Neural SPJ would output the entity literal `Mario Sagario`. Examples of outputs are provided in Figure 3.

The SPJ operator is performing three functions:

(1) for support sets that are insufficient to answer a question, the operator should return no output; (2) for queries that require short chains of reasoning over multiple facts, the SPJ operator joins the facts when generating the output; and (3) the SPJ generates a projection of the support set to a machine readable format dependent on the given query, and whether computation or aggregation is required.

Because the SPJ operator is run in parallel, it can scale independently of the limitations on the size of the input of a single transformer. In contrast, the use of self-attention when encoding all facts as one input precludes parallelization, has high latency, and is limited by the memory required to compute the self-attention. By using the SPJ operator to perform query-dependent information extraction, aggregations can be performed over the generated outputs using conventional computation, which trivially scales to thousands of operands. Furthermore, this allows large result sets to be generated by the model, whereas accurately decoding long sequences using an encoder-decoder architecture remains an open challenge (Hupkes et al., 2020).

Support Set Generator (SSG) A support set contains the minimal subset of sentences from the database needed to generate one single operand for the aggregation module by the SPJ operator. For example, for queries that are answered by a single sentence, e.g., “*Who is Sheryl’s husband?*”, the support set containing a single fact should be returned, e.g., {“*Sheryl is Nicholas’s spouse*”}. The output of the support set generator is a *set* of support sets, each of which is fed independently to a downstream SPJ module. Support sets may not be pairwise disjoint because some facts may be required for multiple answers.

The SSG output should satisfy the following two properties: (1) If multiple facts are needed to produce an intermediate answer, they should all be in the support set. For example, if we queried “*When was Sheryl’s husband born?*”, the support set should include a fact stating who the spouse is and a fact describing when they were born. (2) When performing aggregation, or outputting a set of answers, multiple support sets must be generated, each containing enough information to generate the intermediate results that are aggregated. For example, for the query “*Who is the oldest person?*”, each of the support sets would independently contain a fact that includes a person and indicates their age.

Aggregation The outputs of the SPJ modules are intermediate answers to the query. For some queries, e.g., “*who lives in London?*”, the final answer is simply the union of the intermediate answers. In other cases, e.g., “*how many countries grow coffee?*”, an aggregation operator needs to be applied to the union of intermediate answers. Because output of the SPJ operators are machine readable, we can hence guarantee accuracy and scalability by performing aggregation using conventional computation. In this paper, we consider the aggregation functions `min`, `max` and `count`.

4 The WIKINLDB dataset

In this section we introduce WIKINLDB, a novel dataset for training NLDBs which is generated by transforming structured data from Wikidata (Vrandečić and Krötzsch, 2014) into natural language facts and queries. Wikidata stores triples of the form (S, R, O) , where R is a relationship between the subject S and the object O , e.g., (Tim Cook, employedBy, Apple). The scale and breadth of Wikidata enables us to generate databases of many sizes and variety.

Facts To automate generation of questions and answers, sentences must be grounded in Wikidata identifiers. One approach to generate facts would be to use templates or collect them through grounded information extraction datasets such as T-REx (Elsahar et al., 2018). However, to ensure wider linguistic variety as well as accuracy of the mapping, we use verbalizations of knowledge graph triples that are synthesized through a sequence to sequence model. Concretely, we use generated sentences from KELM (Agarwal et al., 2020), which are not grounded with Wikidata IDs, and generate a post-hoc mapping back to Wikidata. For example, given the sentence: “*The Slice of Life manga series The Film Lives On was written by Osamu Tezuka.*” we map it to the Wikidata triple $(Q11332517, P50, Q193300)$. Our mapping is a two-step process: firstly, we look up entity names from Wikipedia, returning multiple matches for Osamu Tezuka, and secondly filter these based on which have an author relations to *The Slice of Life* in the Wikidata graph. While out of scope for this paper, this technique could be applied to generate training datasets for novel domains. WIKINLDB uses both atomic facts in KELM (about a single relation of an entity) or composite facts (about multiple relations).

Queries Following previous work on large-scale question answering (Hartmann et al., 2018; Talmor and Berant, 2018), queries are generated using templates. For each relation and operator, multiple templates were written by the authors where placeholders can be replaced with the subject and objects for each relation. While multiple templates are used to ensure variety, these are limited in diversity in comparison to the facts. Templates were generated for the first 25 relations on Wikidata with mapped data in KELM. To generate queries that require joins we apply the same technique, combining to combine two or more connected relations, chaining the entities. We further select the 15 most popular relations and generate additional templates which chain the two relations. For example, we chain $(Y, \text{locatedIn}, Z)$ and $(X, \text{employedBy}, Y)$ to create a template for the query “Does $\$X$ work at a company based in $\$Z$?”.

Data Quality We manually inspect randomly selected queries and facts and score them using the categories introduced in this section. For queries, we sample 70 instances, 10 for each query type. We score each query for fluency and intelligibility. Out of 70 queries, only one question was marked as non-fluent due to a typo which was corrected for the final dataset. All 70 queries were intelligible. We observed that the clarity of some queries depended on the facts in the database to provide context (e.g. “Who is male?”), but otherwise met the task requirements.

To assess the quality of mapped facts from KELM, a sample of 50 was evaluated based on 6 categories: intelligibility, fluency, inclusivity (conveying information from all the mapped relations), faithfulness to these relations, and whether extraneous information (not in the mapped relations) is present. 49/50 facts were intelligible and 45/50 facts were fluent. The remaining 5 had redundant information or missing conjunctions. 50/50 facts contained all mapped relations and 48/50 were faithful to these relations. 8/50 facts had extraneous information for relations that could not be mapped. The relations that could not be mapped are not used for query generation and did not affect how answers were automatically generated.

WIKINLDB Statistics We create databases over 25 common relationships from Wikidata, and create 643 templates from which queries are phrased. For *join*-type queries, we chain a fur-

DB Size (up to)	Avg #Q/DB	#DBs		
		Train	Valid	Test
25	8	4000	631	621
50	7	4986	498	499
100	13	2500	250	250
250	53	1000	100	100
500	66	500	50	50
1000	70	250	25	25

Table 1: The statistics for datasets with varying size of DBs (i.e. number of facts). Average number of queries per each DB instance and also the number of DB instances per split is displayed.

Example Input
Which place has the highest yearly number of visitors? [SEP] The Ibaraki Prefectural Museum of History has a visitor count of 93976 per year.
Example Output
[ARGMAX] Ibaraki Prefectural Museum of History [SEP] +93976

Figure 3: Example input and output of the Neural SPJ operator (blue: query, brown: support set sentences)

ther 15 relations with a further 86 template fragments. The relations we chose were selected from a weighted sample of the most common entity types in KELM. In total, we generate five variants of the dataset containing databases of size 25 to 1000 facts where each fact has between 30-50 tokens. Dataset statistics are reported in Table 1.

5 Models

5.1 Neural Select-Project-Join

The SPJ operator is trained as a sequence-to-sequence model to generate intermediate results from a support set and a given query. All facts in the support set are concatenated with the query before being input to a transformer model.

The model is trained to output different derivations depending on the query type. For the *min*, *max* operators, the projection is a machine-readable key-value pair, illustrated in Figure 3. For example “*which place has the highest yearly number of visitors?*” has the projection of the form: $(\text{place}, \text{number of visitors})$ allowing an *argmax* operation by the downstream aggregation module. For queries with Boolean answers, the output is a token indicating whether the answer is true or false. And for all other queries where a set of results is returned or counted, the output is

simply a span, such as an entity or numerical value, extracted from the support set.

Even though we use intermediary annotation for the SPJ operator, we believe that collecting such annotation is a simpler labeling task compare to collecting the answers to the queries. For example, given the fact “*Serena Jameka Williams (born September 26, 1981) is an American professional tennis player and former world No.*” and the query “*List all the female athletes who were born in 20th century.*”, it seems relatively simple to provide the label “*Serena Jameka Williams*”. However, it is non-trivial to produce a list of potentially hundreds of entities as answer (e.g. [*Serena Jameka Williams, Simona Halep, Mary Lou Retton, Megan Rapinoe, Kim Simmone, Mary Abichi, . . .*]). The training of the components in our proposed architecture does not depend on the final answer and instead, on the simpler intermediary labels.

Predicting Aggregation Operator Rather than using a separate classifier to predict the question type, we encode the choice of operator as a special token that is predicted by the SPJ operator prepended to the model output (Figure 3). The aggregation operator is chosen using a majority vote over all generated derivations from all support sets.

Negative Example Generation It is important for the SPJ to be resilient to extraneous facts that might be returned by a low-precision high-recall SSG. Negative instances for training are generated in two ways: (1) queries are paired with randomly sampled facts and the model is trained to generate a NULL projection (indicating the support set does not contribute to the answer). For example, a fact about someone’s date of birth isn’t useful when answering a query about the visitor count of an attraction. (2) for a portion of the training instances, we additionally sample extraneous unrelated facts and append these to the support sets simulating false-positive facts from the SSG.

5.2 Support Set Generator

For simple queries over single facts, conventional information retrieval, such as TF-IDF could be considered a primitive SSG. However, this would not scale for joins, aggregation queries or for queries outputting a set of answers as generating relevant sets requires incremental decoding, conditioning on already retrieved facts.

Naively generating the set of all relevant support sets, $SSG_Q(D) \subset \mathcal{P}(D)$, would be intractable as

Algorithm 1: SSG modeled as multi-label classification: using maximum inner product search (MIPS) over vector encodings of facts U and state V

Input: Bi-encoders C : C_U (for actions), C_V (for state), Database D , Query Q , Threshold τ
Output: Set of support sets $(\hat{D}_1, \dots, \hat{D}_b) \subset \mathcal{P}(D)$
 $open := \{\{\}\}$ $closed := \{\}$;
 $U := [C_U(u_1); \dots; C_U(u_n); C_U(STOP)]$ for $u_i \in D$;
while $open \neq \{\}$ **do**
 $next := \{\}$;
 for \hat{D}_k **in** $open$ **do**
 $V := [C_V(Q, u_1 \dots u_m)]$, for $u_i \in \hat{D}_k$;
 $A := MIPS(U, V, \tau)$;
 for a_j **in** A **do**
 if $a_j == STOP$ **then**
 $closed := closed \cup \{\hat{D}_k\}$;
 else
 $next := next \cup \{a_j \cup \hat{D}_k\}$;
 $open := next$;
return $closed$;

it is akin to enumerating the powerset. We construct support sets efficiently by taking an incremental approach, starting from the empty set (see Algorithm 1). At each step, the classifier considers the partially generated support set \hat{D}_k and the query and predicts which candidate facts $u_i \in D$ from the database should be added, or whether to stop the iteration, these choices being modeled as a multi-label classification task. If STOP is predicted, the partial result set \hat{D}_k is closed (i.e., it forms part of the output); otherwise, for each fact added, a new intermediate (*open*) support set is generated which is explored in the next iteration. For efficiency, we use a bi-encoder architecture that independently encodes the facts in the database and the state (query and a partial support set) and computes the inner product between the encoded representations to generate a score: $C_U(u_i)^T C_V(Q, \hat{D}_k)$. The encoders are pre-trained transformers fine-tuned to yield a high inner product between the state’s encodings and relevant facts to be added. At prediction time, the vectors encoding the facts are static and are pre-computed offline. At each step, t , we encode the state using a transformer by concatenating the query tokens and the facts in the partially generated support set D_k . The SSG is trained with full supervision of all partial support sets from the dataset and trained to predict which facts to add to the support set using a contrastive loss.

Complexity of SSG The inner loop of Algorithm 1 involves a Maximum Inner Product Search (MIPS) between the encoded state and the encodings of the facts, which is linear in the number of facts. Approximate search, such as FAISS (Johnson et al., 2019), accelerate retrieval to $O(\log^2 n)$. If we assume a query needs a maximum of b support sets, and the average size of a support set is m , then the complexity of the SSG algorithm is $O(bm \log^2 n)$. Both b and m are bounded by the number of facts in the database n , but in practice we’d expect only one of b or m factors to be large. However, there is fertile ground for developing methods for indexing (and/or clustering) the facts in the database so that only few facts need to be considered in each iteration of the inner loop of the algorithm, leading to significant speedups.

5.3 Baselines

We compare our proposed architecture to transformer-based models that explore the effect of three attention mechanisms representative of the state-of-the-art. Self-attention in transformers captures both *inter*-fact as well as *intra*-fact interactions between tokens. However, computing self-attention is quadratic with respect to memory and scaling beyond 1024 tokens is non-trivial. In our baselines, the task formulation is a sequence to sequence model, similar to that used in question answering. All (relevant) facts are encoded with the query and the transformer is trained to predict the answer without using any intermediate representations. We compare *full* self-attention against independently encoding the facts (in the context of the query) and fusing the embeddings in the decoder (Izacard and Grave, 2020, *Fusion in Decoder (FiD)*). Because FiD independently encodes contexts, run-time complexity is reduced to be linear with respect to the number of facts at the expense of not having inter-fact attention. We additionally compare to using windowed attention over facts with global attention to the query using *Longformer* (Beltagy et al., 2020). Inter-fact attention is captured only within the window.

6 Implementation

We use the HuggingFace (Wolf et al., 2020) transformers library and its implementations of T5 and Longformer. For SSG, we use BERT to generate encodings, which has a comparable architecture to T5. The learning-rate for fine-tuning and number

Model	Answer Accuracy (%)	
	PerfectIR	WholeDB
NeuralSPJ + Aggr (ours)	90.10 ± 0.3	-
T5	85.59 ± 0.2	65.96 ± 0.5
Longformer	76.43 ± 3	58.58 ± 0.4
Fusion in Decoder	39.61 ± 0.2	23.18 ± 0.6

Table 2: T5 and Longformer both capture inter-fact attention whereas Fusion in Decoder does not. Regardless of how attention is used, using all facts in the database harms the model.

of epochs were selected through maximizing the Exact-Match (EM) accuracy on a held-out validation set for the tasks. For each experiment, we train 3 separate models with different seeds and report mean accuracy. The SPJ models are only trained on the small database of 25 facts and applied to larger databases at test time.

For most queries, we measure correctness using Exact Match (EM), which is 1 if the answer string generated by the model is exactly equal to the reference answer and 0 otherwise. This metric is used to score outputs where either a Boolean, null answer, string or numeric answer is expected. When a set of results is returned, we compute the F_1 score considering exact matches of set elements. When comparing models and reporting results, we report macro-averages over all instances in the test set. We collectively refer to this as *Answer Accuracy*.

7 Experiments & Results

We first consider the suitability of transformer models over small databases of 25 facts comparing two information retrieval settings: *PerfectIR*, which is representative of other question answering approaches that combine an information retrieval system to select only the facts needed to answer a query, and *WholeDB*, where the entire database is encoded by the model, assessing resilience to unrelated information and noise.

The overall scores, in Table 2, indicate that without a retrieval mechanism (i.e., WholeDB), all models were susceptible to distractor facts. Furthermore, encoding all facts in a single model is not a viable solution to answer queries posed to NLDBs as this approach does not accurately answer queries that combine multiple support sets, illustrated in Figure 4, and cannot easily scale to thousands of facts. Using a transformer yields errors when the query requires computation, such as counting, highlighted when comparing rows 1 and 3 of Table 3.




Figure 4: (PerfectIR) Even when provided with the correct contexts, baseline scores decrease for queries requiring the combination of multiple support sets.

Inter-fact attention Applying FiD, which does not capture inter-fact attention, to scale to larger databases would not be successful because answer accuracy further decreases with with support set size. Applying Longformer, which captures inter-fact attention within a window could yield outcomes similar to the T5 transformer baseline where relevant facts are encoded with similar locality. However, in the limit, where context falls between different attention windows, the model could degrade to be similar to FiD.

7.1 Evaluating the SSG+SPJ architecture

Our architecture consists of a support set generator (SSG), a select-project-join (SPJ) operator that generates derivations over the support sets and an aggregation function over the results of the SPJ operators. Assuming a perfect SSG, the SPJ accurately answers more queries than the T5 transformer baseline (Table 2) because of the computation within the aggregation function that yields higher scores for min/max and count queries, displayed in Table 3. In combination with SSG, the overall score decreases to 67% due to retrieval errors. However, SSG+SPJ still exceeds the WholeDB baselines.

It is tricky to evaluate the SSG in isolation because errors here not necessarily translate into errors in query answers. For example, the SSG may return a superset of a support set, but the SPJ may still generate the correct answer. Table 4 shows the performance of the SSG for a database of 25 facts. An output is considered an exact match if it is exactly the same as a support set in the reference data and soft match if it is a superset thereof.

Decoding machine-readable outputs The aggregation operator was selected by predicting a

Method	Answer Accuracy (%)			
	Min/Max	Bool	Count	Set
SPJ PerfectIR	89.72	99.10	94.68	85.25
SSG + SPJ	74.03	77.79	50.75	65.32
<i>T5 PerfectIR</i>	78.23	99.34	87.33	89.19

Table 3: Using retrieved evidence achieves results competitive to the PerfectIR on a DB of 25 facts.

Query Type	Exact Match (%)		Soft Match (%)	
	Precision	Recall	Precision	Recall
Boolean	64.00	80.28	66.15	80.68
Set	63.28	80.77	65.23	81.30
Count	60.21	83.11	61.58	83.41
Min/Max	70.88	93.25	71.80	93.41
<i>Average</i>	65.96	86.51	67.36	86.82

Table 4: Precision and recall of supervised SSG w.r.t. the reference set. Note that errors in retrieval do not necessarily translate to wrong query answers because the SPJ operator is trained to be robust to noise.

special token decoded by the SPJ. For 1.4% of instances, an incorrect choice of aggregation function was made or the machine-readable outputs from the SPJ could not be parsed.

7.2 Scaling to larger databases

We scale the baseline transformers to larger databases using TF-IDF and DPR to retrieve appropriate facts. However, these models are still limited by the encoder size of the transformer. In contrast, the SPJ operates over support sets of 1-2 facts and, in combination with the SSG, can scale to arbitrarily large databases, illustrated in Figure 5. For Boolean queries, the combination of T5 and TF-IDF scored 89%, exceeding the accuracy of the SSG+SPJ. This is because TF-IDF exploits token matching between the query and facts. For larger databases, the retrieval errors resulted in lower answer accuracy. While, with a perfect SSG, the the SPJ accurately answers most query types, as database size increases, the propagation of errors from the SSG resulted in erroneous answers.

8 Related Work

Database queries require reasoning over a large set of relevant and non-redundant facts and performing aggregation. While in-roads have been made to perform discrete reasoning and computation over passages (Dua et al., 2019), with explicit computation (Andor et al., 2019) or differentiable modules




Figure 5: Scaling to larger databases with a model trained using 25 facts and tested on larger databases.

(Gupta et al., 2020), these use only a single passage rather than requiring aggregation over large numbers of facts from different texts.

Multi-hop question answering requires finding supporting evidence in multiple documents (see (Welbl et al., 2018; Talmor and Berant, 2018; Wolfson et al., 2020) for datasets facilitating this research). In answering multi-hop questions, the works decompose the question into simpler sub questions (Min et al., 2019; Wolfson et al., 2020), or condition each hop on the previously retrieved documents (Asai et al., 2019; Xiong et al., 2020).

While tasks such as ComplexWebQuestions (Talmor and Berant, 2018) and BREAK (Wolfson et al., 2020) focus on complex queries that can be broken down into simpler ones, our focus is on set-based and aggregation queries where the complexity comes from the need to retrieve and process a large number of non-redundant relevant facts. In contrast to the *set* and *count* tasks in bAbI (Weston et al., 2015), where each query is based on a small context (less than 20 facts), our dataset scales from databases of 25 facts to 1000.

Bridging the gap between unstructured natural language data and database-style querying has been a long-standing theme in database research (Halevy et al., 2003). The work on information extraction has developed techniques for translating segments of natural language text into triples that can be further processed by a database system. There has been significant work on translating queries posed in natural language into SQL queries on a database whose schema is known (Androustopoulos et al., 1995; Li and Jagadish, 2014; Zeng et al., 2020), with extensions to semi-structured data and knowledge bases (Pasupat and Liang, 2015; Berant et al., 2013). More recently, systems such as

BREAK (Wolfson et al., 2020) and ShARC (Saeidi et al., 2018) have trained models to translate a natural language query into a sequence of relational operators (or variants thereof).

9 Conclusions

Database systems are the workhorse of data analysis but they require a pre-defined schema. Part of their power stems from the fact that a data analyst can explore the data by easily posing a wide variety of queries. Given the rise in the amount of data that is becoming available in text, images and other modalities, we would like to build systems that enable the flexibility of posing complex queries against such data, but without the need for a pre-defined schema.

This paper proposed an architecture for neural databases and the associated WIKINLDB dataset, as first steps towards realizing a system for querying multi-modal data. Our architecture is capable of overcoming the limitations of transformer models because it runs multiple transformers in parallel, each taking a small set of facts. Consequently, NLDBs can scale to large databases.

Additional research is required in order to scale NLDBs to larger datasets, more complex queries, and to multi-modal data. In particular, one of the key components of the architecture is the SSG module that retrieves the relevant facts to feed to each instance of the neural SPJ. We believe that in practice, the semantics of the application will provide a strong hint on which facts may be relevant. For example, when querying a large corpus of social-media posts, each post is a candidate support set as long as the query does not require joining data from multiple posts. In addition, we assumed that our databases describe a snapshot of the world. In practice, we may have facts that override previous ones (e.g., ‘Samantha works for Apple’, followed by ‘Samantha works for Twitter’) and we would need to reason about which facts should be ignored.

Acknowledgments

We would like to thank Yann LeCun and Antoine Bordes for the initial discussion that sparked the idea of neural databases. This work was performed while James Thorne and Fabrizio Silvestri were at Facebook.

Broader Impact Statement

Ethical Concerns A NL database is very similar to a traditional database in terms of applications with a difference that it extends the use of databases on unstructured text. For example, NL databases can be used to produce analytics on data expressed in natural language. For an NL database to be applicable in the context of a virtual assistance, they will likely need to be trained on real-world conversations. Privacy preserving ML methods should be considered for such applications.

Environmental Concerns Large transformer-based models take a lot of computational resources and energy for pre-training and fine-tuning. As a result such models raise environmental concerns. In our proposed architecture, we only fine-tune transformer models on small support sets. We then use several instances of such models in parallel for inference, instead of a single large model, even on large datasets. Therefore, the model is relatively efficient, both during the fine-tuning and during the inference.

References

- Oshin Agarwal, Heming Ge, Siamak Shakeri, and Rami Al-Rfou. 2020. [Large scale knowledge graph based synthetic corpus generation for knowledge-enhanced language model pre-training](#).
- Daniel Andor, Luheng He, Kenton Lee, and Emily Pitler. 2019. [Giving BERT a calculator: Finding operations and arguments with reading comprehension](#). 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)*, volume 2, pages 5947–5952. Association for Computational Linguistics.
- I Androustopoulos, G D Ritchie, and P Thanisch. 1995. [Natural Language Interfaces to Databases - an Introduction](#). *Natural Language Engineering*, 1(1):29–81.
- Akari Asai, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caiming Xiong. 2019. [Learning to retrieve reasoning paths over wikipedia graph for question answering](#). *arXiv preprint arXiv:1911.10470*.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#).
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, October, pages 1533–1544. Association for Computational Linguistics.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879. Association for Computational Linguistics.
- Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, and Andrew McCallum. 2019. [Multi-step retriever-reader interaction for scalable open-domain question answering](#). *arXiv preprint arXiv:1905.05733*.
- Bhuvan Dhingra, Kathryn Mazaitis, and William W Cohen. 2017. [Quasar: Datasets for question answering by search and reading](#). *arXiv preprint arXiv:1707.03904*.
- 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.
- Hady Elsahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon Hare, Frederique Laforest, and Elena Simperl. 2018. [T-REx: A large scale alignment of natural language with knowledge base triples](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Nitish Gupta, Kevin Lin, Dan Roth, Sameer Singh, and Matt Gardner. 2020. [Neural module networks for reasoning over text](#). In *International Conference on Learning Representations (ICLR)*.
- Alon Y. Halevy, Oren Etzioni, AnHai Doan, Zachary G. Ives, Jayant Madhavan, Luke K. McDowell, and Igor Tatarinov. 2003. [Crossing the structure chasm](#). In *CIDR 2003, First Biennial Conference on Innovative Data Systems Research, Asilomar, CA, USA, January 5-8, 2003, Online Proceedings*. www.cidrdb.org.
- Ann-Kathrin Hartmann, Edgard Marx, and Tommaso Soru. 2018. [Generating a large dataset for neural question answering over the DBpedia knowledge base](#).
- Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. 2020. [Compositionality Decomposed: How do Neural Networks Generalise?](#) *Journal of Artificial Intelligence Research*, 67:757–795.
- Gautier Izacard and Edouard Grave. 2020. [Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering](#).

- Jeff Johnson, Matthijs Douze, and Herve Jegou. 2019. [Billion-scale similarity search with GPUs](#). *IEEE Transactions on Big Data*, pages 1–1.
- 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. 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.
- Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020a. [Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks](#).
- Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020b. Retrieval-augmented generation for knowledge-intensive nlp tasks. *arXiv preprint arXiv:2005.11401*.
- Fei Li and H V Jagadish. 2014. Constructing an Interactive Natural Language Interface for Relational Databases. *Proceedings of the VLDB Endowment* 2, 8(1):73–84.
- Sewon Min, Victor Zhong, Luke Zettlemoyer, and Hananeh Hajishirzi. 2019. [Multi-hop reading comprehension through question decomposition and rescoring](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6097–6109. Association for Computational Linguistics.
- Panupong Pasupat and Percy Liang. 2015. [Compositional semantic parsing on semi-structured tables](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1470–1480. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392. Association for Computational Linguistics.
- Marzieh Saeidi, Max Bartolo, Patrick Lewis, Sameer Singh, Tim Rocktäschel, Mike Sheldon, Guillaume Bouchard, and Sebastian Riedel. 2018. [Interpretation of natural language rules in conversational machine reading](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2087–2097. Association for Computational Linguistics.
- Alon Talmor and Jonathan Berant. 2018. [The web as a knowledge-base for answering complex questions](#). 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 641–651. Association for Computational Linguistics.
- James Thorne, Majid Yazdani, Marzieh Saeidi, Fabrizio Silvestri, Sebastian Riedel, and Alon Y. Halevy. 2021. From natural language processing to neural databases. *Proc. VLDB Endow.*, 14.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *31st Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, CA, USA.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. [Linformer: Self-Attention with Linear Complexity](#). 2048(2019).
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. [Constructing datasets for multi-hop reading comprehension across documents](#). *Transactions of the Association for Computational Linguistics*, 6:287–302.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Huggingface’s transformers: State-of-the-art natural language processing](#).
- Tomer Wolfson, Mor Geva, Ankit Gupta, Matt Gardner, Yoav Goldberg, Daniel Deutch, and Jonathan Berant. 2020. [Break it down: A question understanding benchmark](#). *Transactions of the Association for Computational Linguistics*, 8:183–198.

Wenhan Xiong, Xiang Lorraine Li, Srini Iyer, Jingfei Du, Patrick Lewis, William Yang Wang, Yashar Mehdad, Wen-tau Yih, Sebastian Riedel, Douwe Kiela, et al. 2020. Answering complex open-domain questions with multi-hop dense retrieval. *arXiv preprint arXiv:2009.12756*.

Jichuan Zeng, Xi Victoria Lin, Steven C.H. Hoi, Richard Socher, Caiming Xiong, Michael Lyu, and Irwin King. 2020. [Photon: A robust cross-domain text-to-SQL system](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, volume abs/2007.15280, pages 204–214. Association for Computational Linguistics.

A Appendix

A.1 Sample Data and Dataset Statistics

Example: Set	
Question	Who studied at University of Minnesota?
Supporting Facts	<ol style="list-style-type: none">1. [John B Totushek was born on 7 September 1944 in Minneapolis. He attended the University of Minnesota and became a US Naval Aviator. Mr. Totushek was also a human being.]2. [Melvin Maas graduated from the University of Minnesota and is buried at Arlington National Cemetery. He is a native of Minnesota and his language is English.]3. [Clarence Larson graduated from the University of Minnesota and is a member of the National Academy of Engineering.]4. [Ted Mann, who is the surname of Ted Mann, attended Duke University and the University of Minnesota. He is a human being.]
Answer	[John B. Totushek, Ted Mann, Clarence Larson, Melvin Maas]
Example: count	
Question	How many people work for Yale Law School?
Supporting Facts	<ol style="list-style-type: none">1. [Michael Ponsor, born in Oxford, graduated from Pembroke College in Oxford. He was awarded the Rhodes Scholarship and is an employee at Yale Law School. He is an expert in the field of human rights.]2. [Stephen Wizner is an American legal scholar who graduated from Dartmouth College and is a graduate of the University of Chicago Law School. He works at Yale Law School.]
Answer	2
Example: Min/Max	
Question	What is the largest yearly attendance?
Supporting Facts	<ol style="list-style-type: none">1. [The musee en herbe has a visitor per year of] 70000.2. [The total number of visitors to the Hirschsprung Collection is 71779 per year.]...24. [The Tate Modern has a visitor count of 5839197 visitors per year.]25. [Catocin Mountain Park attracts 221750 visitors per year.]
Answer	5839197
Example: Bool	
Question	Is North Carolina State University the employer of Wes Moore?
Supporting Facts	<ol style="list-style-type: none">1. [Wes Moore is a human being who is employed at Francis Marion University and is a basketball player for North Carolina State University.]
Answer	TRUE
Example: Join	
Question	Who plays for a team in Ligue 1?
Supporting Facts	<ol style="list-style-type: none">1. [Thomas Allofs started his career in 1989 with RC Strasbourg Alsace. He finished his career in 1990., RC Strasbourg Alsace is an association football club in the Ligue 1 league. It was founded in 1906 and is located in Strasbourg, France.]
Answer	[Thomas Allofs]

Table 5: Examples of different types of queries, their supporting facts and answers. These examples are based on databases of size 25.




Figure 6: Dataset statistics for DBs of varying sizes provided with WIKINLDB