

Dyna-bAbI: unlocking bAbI’s potential with dynamic synthetic benchmarking

Ronen Tamari^{†*} Kyle Richardson^{*} Noam Kahlon[†] Aviad Sar-Shalom[△]

Nelson F. Liu[◇] Reut Tsarfaty^{*‡} Dafna Shahaf[†]

[†]The Hebrew University of Jerusalem ^{*}Allen Institute for AI

[‡]Bar-Ilan University [△]Tel-Aviv University [◇]Stanford University

{ronent, dshahaf}@cs.huji.ac.il, {reutt, kyler}@allenai.org

Abstract

While neural language models often perform surprisingly well on natural language understanding (NLU) tasks, their strengths and limitations remain poorly understood. Controlled synthetic tasks are thus an increasingly important resource for diagnosing model behavior. In this work we focus on story understanding, a core competency for NLU systems. However, the main synthetic resource for story understanding, the bAbI benchmark, lacks such a systematic mechanism for controllable task generation. We develop Dyna-bAbI, a dynamic framework providing fine-grained control over task generation in bAbI. We demonstrate our ideas by constructing three new tasks requiring compositional generalization, an important evaluation setting absent from the original benchmark. We tested both special-purpose models developed for bAbI as well as state-of-the-art pre-trained methods, and found that while both approaches solve the original tasks (>99% accuracy), neither approach succeeded in the compositional generalization setting, indicating the limitations of the original training data. We explored ways to augment the original data, and found that though diversifying training data was far more useful than simply increasing dataset size, it was still insufficient for driving robust compositional generalization (with <70% accuracy for complex compositions). Our results underscore the importance of highly controllable task generators for creating robust NLU systems through a virtuous cycle of model and data development.¹

1 Introduction

Considerable progress has been made recently in natural language understanding (NLU), driven largely by advances in model pre-training (Devlin

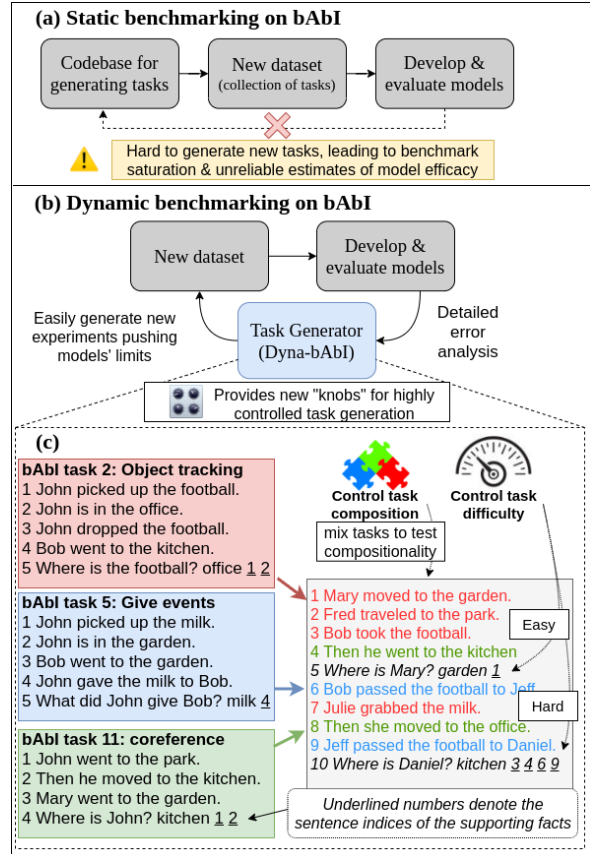


Figure 1: (a) Low task configurability leads to static datasets, benchmark saturation & unreliable model development. (b) We propose a dynamic benchmarking approach; developing models and tasks in a tight feedback loop using (c) Dyna-bAbI task generator. Dyna-bAbI provides fine-grained control over task structure, composition and difficulty, yielding challenging new test sets exposing limitations of state-of-the-art models.

et al., 2019; Raffel et al., 2020) and the development of large-scale NLU benchmarks across a wide range of tasks (Wang et al., 2018, 2019; Liang et al., 2020). Such successes, however, have coincided with the discovery of various shortcomings in existing human curated datasets, largely related to *annotation artifacts* (Gururangan

* Work begun during an internship at the Allen Institute.

¹ Data and code available at <https://dyna-babi-project.github.io/>.

et al., 2018), or systematic biases that create shortcuts that can inflate model performance and harm generalization.

In order to overcome these issues, two avenues of research have recently gained traction: 1) development of *dynamic benchmarks* (Potts et al., 2021; Kiela et al., 2021) where, in contrast to conventional *static* benchmarks, evaluation and data collection are conducted interactively with humans and models in a rapidly evolving feedback loop and; 2) renewed interest in *synthetic benchmarks* (Lake and Baroni, 2018; Sinha et al., 2019; Clark et al., 2020; Ruis et al., 2020) that allow for absolute control over the data creation process in order to help understand the strengths and weaknesses of existing models on targeted tasks and language phenomena.

Story understanding is a particularly important domain for research on dynamic and synthetic benchmarks; it is a core competency for NLU systems (McClelland et al., 2020; Dunietz et al., 2020), but the scale and annotation detail required make human data collection prohibitively costly. However, the main synthetic resource for story understanding remains the bAbI task suite (Weston et al., 2016), which is saturated by models reaching near-perfect performance (Liu et al., 2021), and further limited by exploitable biases in the data (Kaushik and Lipton, 2018). Despite its creators’ initial intentions, bAbI has largely remained a static benchmark limited to a small subset of the tasks potentially possible to generate within the bAbI “micro-world”. Accordingly, two natural questions arise: **(Q1)** *is near-perfect model performance on the original bAbI tasks a reliable indicator of story understanding competence?*; **(Q2)** *are there still interesting challenges to discover inside the broader bAbI task space that help identify weaknesses in current models and drive modelling innovation?*

To answer these questions, we employ a *dynamic synthetic benchmarking* approach on bAbI, combining the benefits of the agile approach of recent dynamic benchmarks with the scale and control provided by synthetic datasets. As illustrated in Figure 1, in dynamic synthetic benchmarks the data generator itself is designed for agile development, enabling experimentation with increasingly complex tasks and a wider range of linguistic phenomena.² Constructing

² While our framework does not enable *automatic* collection

challenging tasks is a challenge in and of itself, requiring precise control over the reasoning patterns underlying each question. To meet these requirements, we developed a new task generator for bAbI called Dyna-bAbI³.

Using Dyna-bAbI, we first devise new splits that systematically test *compositional generalization* across tasks; as shown in Fig. 1c, we test models on novel combinations (right side, line 10) of concepts seen at training, like co-reference and object tracking (left). We find that training on the original bAbI tasks (hereafter: bAbI 1.0) is not sufficient for models to attain good compositional generalization. Though general purpose pre-trained models far outperform special-purpose (non-pre-trained) architectures developed for bAbI, they still suffer a 20-50% drop in accuracy compared to the non-pre-trained models which suffer a 50-80% drop. Both types attain near perfect performance on the original tasks, suggesting that bAbI 1.0 is not challenging enough to differentiate between the two classes of models (**Q1**).

We next investigate how different enhancements of training data affect compositional generalization: (a) injecting more questions into bAbI 1.0, and (b) generating new, more diverse training samples. Compared to question injection, we find that diverse training data better facilitates compositional generalization, as well as being more data efficient. However, neither approach drives *reliable* compositional generalization; a representative state-of-the-art (SOTA) model, T5 (Raffel et al., 2020), demonstrates a lack of robustness to novel combinations and also exhibits knowledge inconsistency, for example, by correctly answering certain types of questions but systematically failing to answer equivalent paraphrases. These results suggest that there remain many important challenges within the broader bAbI task space (**Q2**) which can be discovered through more careful control of task generation.

To sanity-check the quality of our new tests compared with bAbI 1.0, we employ the notion of *concurrency* proposed by Liu et al. (2021);

of new data based on model errors as in other dynamic benchmarks, we still chose the term “dynamic” to highlight their important common function: data generation frameworks that enable easily “moving the goalposts” in meaningful directions (in our case, for probing models’ systematic generalization capacities).

³ Implemented in Python for improved accessibility compared with the original Lua implementation (<https://github.com/facebookarchive/bAbI-tasks>).

concurrency is a measure of correlation between models’ performance on a synthetic task and their performance on an existing, non-synthetic NLU benchmark. We find high concurrency between our new challenge tasks and the widely used SQuAD dataset (Rajpurkar et al., 2016), in contrast to bAbI 1.0, which achieved low concurrency.

Giving the continued interest in using bAbI 1.0 to evaluate new modelling approaches (Banino et al., 2020, 2021; Schlag et al., 2021), our new challenge splits and the Dyna-bAbI task generator contribute to more reliably guiding future efforts. While we focused on bAbI, our results apply more generally, telling a cautionary tale about the limits of static synthetic datasets, and motivating the development of controllable task generators for dynamic synthetic benchmarking.

2 Related Work

Our work brings together two promising areas of current research: dynamic benchmarking such as Dynabench (Kiela et al., 2021) that address many existing issues with static benchmarks (Bowman and Dahl, 2021), and synthetic benchmarking, which is widely used for high-precision and data-intensive problems such as relational and logical reasoning (Sinha et al., 2019; Clark et al., 2020; Betz et al., 2021; Richardson and Sabharwal, 2022), robot planning (Banerjee et al., 2020), instruction following and language grounding (Long et al., 2016; Lake and Baroni, 2018) among many others (Richardson et al., 2020; Khot et al., 2021). Most approaches to synthetic benchmarking focus on model development on a static benchmark, and are not designed to facilitate agile and highly controlled task space exploration, which is our focus here.

The recent gSCAN dataset (Ruis et al., 2020) and later extensions (Qiu et al., 2021; Wu et al., 2021) can be seen as an example of a synthetic benchmark “going dynamic”. Our work differs in terms of target domain (story understanding as opposed to multi-modal language grounding), and we further focus attention on a more general research direction of intentional, a-priori design of NLU benchmarks for agile development. In this regard, our work can be seen as part of a trend towards *data-centric* research efforts in response to prevailing *model-centric* research, which generally focuses heavily on architectural design and novelty (Kaushik and Lipton, 2018), at

the expense of work on the data side which often receives less attention (Sambasivan et al., 2021; Rogers, 2021).

We address the domain of story understanding as a particularly core (and data-intensive) capacity underlying language use (McClelland et al., 2020), thought to require constructing and manipulating situation models of entities and their relations as they unfold throughout discourse (Zwaan, 2016; Tamari et al., 2020). Procedural text datasets (Dalvi et al., 2018; Tandon et al., 2020) are closely related in that they provide detailed annotation of entities and state changes, and have mostly focused on relatively small and static benchmarks using human collected data. Overall, recent works identify a lack of benchmarks which systematically probe the situation models constructed by systems processing discourse-level texts (Sugawara et al., 2021).

The bAbI benchmark (Weston et al., 2016) is seen as highly relevant in terms of objective (targeting situation modelling) (Dunietz et al., 2020), but has been viewed critically due to its constrained nature and exploitable artifacts (Kaushik and Lipton, 2018). Our work focuses on improving the evaluation in bAbI through compositional generalization, widely used across NLP to more rigorously probe model robustness (Finegan-Dollak et al., 2018; Keyzers et al., 2020; Gontier et al., 2020; Yanaka et al., 2021), but to our knowledge still not applied to story understanding or bAbI.

3 Synthetic Dynamic Benchmarking on bAbI

3.1 Dyna-bAbI

What makes a synthetic benchmark *dynamic*? We think of a dynamic synthetic benchmark as a highly controllable task generator, enabling rapid exploration of interesting areas of a task space. The original bAbI 1.0 simulator code does not readily facilitate such exploration; each of the bAbI 1.0 tasks is generated by a hard-coded script which does not enable parametric manipulation of interesting generation aspects such as question difficulty or compositionality.

Accordingly, we developed Dyna-bAbI, a Python-based version of the original simulator. Dyna-bAbI facilitates control of task generation through a configuration file, effectively abstracting away much of the underlying implementation complexity. The configuration file allows users

to specify high-level task parameters such as the set of target concepts, passage length, and filtering conditions to mine for harder/rarer examples. We also modularized the code to facilitate adding new questions and other concepts more easily.

In this next sections we describe the underlying structure of the bAbI 1.0 tasks, and how we combine them using Dyna-bAbI to create more complex compositional generalization tasks.

3.2 bAbI task structure

A task in bAbI 1.0 is a set of train, validation and test splits. Each split is a set of instances, where an instance is a tuple $(p, q, a) = (\textit{passage}, \textit{question}, \textit{answer})$. Passages are generated using a micro-world simulator by sampling a valid sequence of world events from an event set \mathcal{E} and generating a linguistic description of them. By default, linguistic descriptions are generated by a simple sentence-level mapping from an event to a natural language sentence. For example, the event `move(john, park)` could be translated to “John moved to the park.”

Some tasks also incorporate more complex linguistic mappings between events and sentences, such as co-reference: the event sequence $(\textit{move}(\textit{john}, \textit{park}), \textit{move}(\textit{john}, \textit{kitchen}))$ could be mapped to “John moved to the park. Then he went to the kitchen.” We denote the set of possible linguistic mappings by \mathcal{L} .

Finally, a valid question-answer pair (q, a) over p is sampled from question set \mathcal{Q} . In bAbI, each split is generated using some particular subset of all possible events, linguistic constructs and questions (§3.3); for a given split we can then define its *concept set*, $\mathcal{C} = \mathcal{E} \cup \mathcal{L} \cup \mathcal{Q}$. Instances also include a set of supporting facts (f), or the relevant lines from which a can be derived (see Fig. 1). The support composition (f_c) is the set of events and linguistic constructs contained in f (see examples in §4.2.1), and is useful for characterizing compositionality performance (§3.4).

3.3 Original bAbI 1.0 tasks

Our focus here is on a particular subset of 12 bAbI 1.0 tasks evaluating aspects of story understanding. Table 1 summarizes them, detailing $\mathcal{E}, \mathcal{L}, \mathcal{Q}$ for each task. For \mathcal{L} , we list only complex constructs beyond the default event-sentence mapping (which is present in every task). See appendix A.1 for additional details on task construction. Not all

Task	Events (\mathcal{E})	Linguistic Constructs (\mathcal{L})	Questions (\mathcal{Q})	Avg. sents. & supp. facts per story
1	MOVE	-	where-P	6, 1
2	MOVE, POSS	-	where-O	15.52, 2
3	MOVE, POSS	-	where-was-O	51.9, 3
5	GIVE, POSS	-	give-qs	20.1, 1
6	MOVE	-	yes-no	6.27, 1
7	GIVE, POSS	-	counting	8.67, 2.33
8	MOVE, POSS	-	list	8.75, 1.94
9	MOVE	NEGATE	yes-no	6, 1
10	MOVE	INDEF	yes-no	6, 1
11	MOVE	CO-REF	where-P	6, 2
12	MOVE	CONJ.	where-P	6, 1
13	MOVE	CONJ., CO-REF	where-P	6, 2

Table 1: Subset of 12 bAbI 1.0 tasks considered here. Each task is characterized by the possible events, linguistic constructs and questions that can occur in instances. POSS (possession) is short for GRAB and DROP events. Statistics based on training sets. A large space of task configurations remains unexplored.

of the story understanding tasks are considered. For example, tasks 14 and 20 address time reasoning and agent motivations, and we leave their integration for future work.

3.4 Compositional generalization on bAbI

As can be seen in Table 1, many possible task configurations are not covered by the original benchmark; which directions should be explored? We focus on out-of-distribution (OOD) robustness, which is increasingly seen as a vital evaluation criteria across AI/NLP research (Shanahan et al., 2020; Hendrycks et al., 2020). We target *compositional generalization*, a particularly important class of OOD problems (Lake et al., 2017; Lake and Baroni, 2018). Compositional generalization refers to the ability to systematically generalize to test inputs containing novel combinations of more basic elements seen at training time (Partee et al., 1995). For example, a model that has learned basic object tracking and co-reference *separately* (tasks 2 and 11, see Fig. 1c) could be expected to solve tasks requiring a *mixture* of both object tracking and co-reference (Fig. 1c, line 10 question on right side). Compositional tasks are absent from bAbI

1.0 which features only IID test sets (independent, identically distributed).⁴

Compositional task generation. To create compositional generalization tasks in practice, we create training (and validation) splits composed of M sub-tasks with concept sets $\{C_{\text{train}}^i\}_{i=1}^M$, and a test set C_{test} such that $C_{\text{test}} \neq C_{\text{train}}^i \forall i$, but $C_{\text{test}} = \bigcup_{i=1}^M C_{\text{train}}^i$. In other words, each training sub-task can be thought of focusing on a particular subset of test concepts, so models are exposed to all test concepts at training time, but not to all combinations of them (Yanaka et al., 2021).

Task difficulty. We hypothesize that support composition (f_c) and supporting fact set size ($|f|$) are main factors underlying a particular instance’s difficulty, and especially *novel* support compositions not seen at training time. Additionally, the difference between train and test splits results in potentially harder distractors, as test-time distractors appear in novel contexts.

Our notions of concepts and support composition resemble atoms and compounds in DBCA, a related study on compositionality (Keysers et al., 2020). While DBCA enables automatic creation of compositional train and test splits, we opt here for a more human-interpretable representation that allows more precise manual control of the combinations of concepts a model is exposed to at train and test time.

Quality comparison vs. bAbI 1.0 tasks. Intuitively, good synthetic datasets help drive the development of better modelling approaches. Our new compositional tasks might be harder than bAbI 1.0, but how do we know whether they are a more useful target? To provide a preliminary answer to this question, we adopt the notion of *concurrency* as a quality measure (Liu et al., 2021). Two benchmarks are said to have high concurrency when they rank a set of modelling approaches similarly. Concurrency offers a way to formalize the intuition above, as high concurrency between a synthetic and natural language benchmark suggests that the synthetic benchmark could have driven similar innovations. We follow the setup of Liu et al. (2021) using SQuAD for the natural language benchmark.⁵ Notably, bAbI 1.0 achieved very

⁴ Weston et al. (2016) noted that transfer learning was an important goal out of the original work’s scope.

⁵ Liu et al. (2021) consider a set of 20 modelling approaches used on SQuAD, including 10 pre-trained and 10 non-pre-trained methods.

Split	Type	Avg. length	Size	Avg. supp. fact set size
concat(T2)	Train	10.76	18,000	2
concat(T7)	Train	13.5	63,000	1.68
inject(T7)	Train	23.25	190,158	1.42
diverse(T7)	Train	20	17,000	2.17
concat(T12)	Train	10.8	108,000	1.42
inject(T12)	Train	15.97	368,831	1.28
diverse(T12)	Train	20	24,772	2.45
mix(T2)	Test	13.25	1,000	2.05
mix(T7)	Test	20	3,000	2.50
mix(T12)	Test	20	6,000	3.70

Table 2: Splits used for our experiments. All except the original data (*concat*) are created with Dyna-bAbI.

low concurrency with SQuAD; for example, pre-training consistently yields large gains on SQuAD, but on bAbI 1.0, both pre-trained and non-pre-trained models achieve perfect performance on many tasks. The low concurrency thus suggests that bAbI 1.0 may be an unreliable benchmark for model development, and highlights the importance of improving its quality.

4 Experiments

With the controllable task generation afforded by Dyna-bAbI, we can now create datasets probing deeper story understanding capabilities of models.

We present two main experiments targeting the following questions:

- Exp. 1: (q1.a) What role does model architecture play in the capacity for compositional generalization? (q1.b) What is the concurrency of our compositional tasks with real datasets, compared with bAbI 1.0?
- Exp. 2: (q2) How do training data quantity and diversity affect compositional generalization?

Data

For our experiments we created 4 kinds of splits over three subsets of bAbI 1.0 tasks, summarized in Table 2. We denote a subset of tasks T , and consider $T_2 = \{2, 11\}$, $T_7 = \{1, 2, 3, 5, 11, 12, 13\}$, and $T_{12} = \{1, 2, 3, 5, \dots, 13\}$.

- *concat* splits are simply concatenations of the official data for the tasks T . We considered the larger version where each task consists of 9,000/1,000 training/development examples; e.g., *concat*(T_2) consists of 18,000 training examples and 2,000 development examples.

- *inject* splits enrich the *concat* data as follows: for each question in the original data, we supplement it with all possible additional questions of the specified types. In this work, the supplement question types were *where-P* and *where-O* (to provide location information of objects and agents).
- *diverse* splits use rejection sampling to generate more diverse samples, such that the number of supporting facts per question is roughly uniform across all sub-task instances for a given question type. Without rejection sampling, most generated questions would be trivial (e.g., 1-2 supporting facts). Compositionality is retained by holding out certain combinations. In particular, at training time, complex linguistic constructs (e.g., co-reference) are only seen with MOVE events.
- *mix* are test splits generated using rejection sampling like *diverse*, and consist of instances which may feature elements from any of the considered tasks. As a result, questions in *mix* splits require novel/more complex reasoning patterns compared to those seen during training.

See appendix A.1 for examples and extended details on task generation.

4.1 Exp. 1: Can training on bAbI 1.0 facilitate compositional generalization?

For this experiment, we compared models on T_2 and T_7 , since they allow for a direct conversion to an extractive QA format,⁶ enabling us to use the same concurrence framework of Liu et al. (2021).

Models. We considered 3 classes of models:

- Non-pre-trained specialized architectures for bAbI 1.0 including EntNet (Henaff et al., 2017) and STM (Le et al., 2020), the latter being current SOTA on bAbI 1.0⁷.
- Non-pretrained general-purpose QA methods, such as BiDAF (Seo et al., 2017).
- General purpose pre-trained approaches including RoBERTa (Liu et al., 2020) and T5 (base) (Raffel et al., 2020).

The last two categories are comprised of the 20 models evaluated in Liu et al. (2021), with the addition of T5 to the last group. For implementation details, see appendix A.2.

⁶ Tasks 6-10 require generative QA, for answering *yes-no*, *count* and *list* questions.

⁷ As of March 10, 2022.

Results & Analysis

Experiment results are summarized in Table 3. All models perform well in IID settings, but performance drops considerably in OOD settings

Architecture alone is not a significant compositionality driver (q1.a). The large OOD performance gap between pre-trained and non-pre-trained models indicates that pre-training plays a much greater role than specialized architectures for QA performance, adding to similar findings in other NLP domains (Hendrycks et al., 2020). These results raise questions about special purpose relational reasoning architectures that continue to be developed today: the poor OOD performance suggests that such models may not be fulfilling their intended design. Either way, these results underscore the importance of rigorous evaluation to verify that modelling motivations are borne out in practice (Aina et al., 2019).

Compositionality increases concurrence (q1.b).

As can be seen in the Fig. 2 plots⁸, increasing compositionality is correlated with increased concurrence. In contrast to the original bAbI 1.0 tasks which exhibited virtually no correlation with SQuAD, our compositional task $mix(T_7)$ exhibits high concurrence of $r = 0.92, \tau = 0.78$ (Pearson and Kendall correlation functions, resp.). These results are comparable to other *natural* language as well as purpose-built synthetic datasets considered in Liu et al. (2021), which feature r, τ in the ranges $[0.87, 0.99]$ and $[0.77, 0.94]$, respectively. Our results thus extend the findings of Liu et al. (2021); they demonstrated the *existence* of high concurrence synthetic benchmarks, we additionally suggest a guiding principle for how to *create* them (incorporate compositionality evaluation).

4.2 Exp. 2: enriching bAbI 1.0 training data

The results above suggest that the bAbI data in their current form may not be rich enough to drive compositional generalization.⁹ In this experiment we probe this question, enriching the training data to better understand its impact on compositional generalization. In particular, we investigate two approaches to enriching the training data while maintaining the compositionality evaluation, corresponding to the *inject* and *diverse* splits.

⁸ See appendix A.4 for full numeric results.

⁹ An alternate hypothesis is that certain patterns may be too hard for models to learn; we confirm this is not the case by using the inoculation methodology of Liu et al. (2019), see details in Appendix A.3.

Name	Train	Test	Evaluation accuracy					SQuAD Concurrence	
			EntNet	STM	BiDAF	Roberta	T5	ρ	τ
2-task IID	concat(T2)	concat(T2)	98.95	99.85	100	100	99.85	[-0.35,0.08]	[-0.35,-0.19]
2-task OOD	concat(T2)	mix(T2)	72.0	67.6	97.2	98.7	98.1	0.48	0.51
7-task IID	concat(T7)	concat(T7)	96.8	99.4	99.98	99.98	99.8	[-0.4,0.08]	[-0.35,0.03]
7-task OOD	concat(T7)	mix(T7)	22.2	26.7	30.5	57.7	49.57	0.92	0.78
12-task IID	concat(T12)	concat(T12)	96.19	99.34	-	-	99.54	-	-
12-task OOD	concat(T12)	mix(T12)	31.97	35.65	-	-	67.4	-	-

Table 3: Experiment 1. OOD evaluation exposes large differences between pre-trained and non-pre-trained models, and also achieves high concurrence with the SQuAD benchmark. We report [min,max] concurrence for bAbi 1.0.

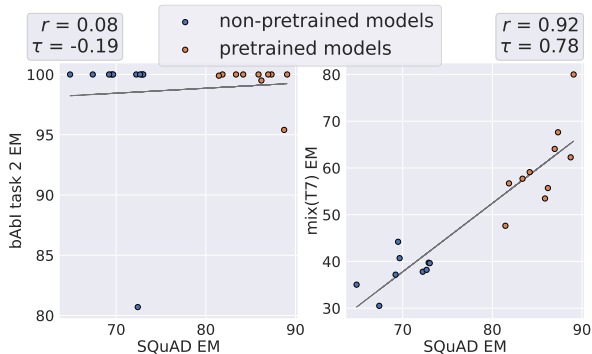


Figure 2: SQuAD concurrence plots for bAbi 1.0 task 2 (left; reproduced from Liu et al. (2021) with permission) and $mix(T_7)$ (right). bAbi task 2 has the highest concurrence of all T_7 tasks, yet exhibits virtually no correlation with SQuAD. $mix(T_7)$ exhibits high concurrence, highlighting the relevance of compositional evaluation.

Notably, Exp. 2 can be seen as a first iteration of the dynamic benchmarking loop depicted in Fig. 1: based on the error analysis of Exp. 1, we leverage Dyna-bAbi for targeted creation of new tasks, which allow us to systematically test our hypotheses.

In this experiment we focus on pre-trained models, as they significantly out-performed non-pre-trained methods. We use T5 as a representative since its generative abilities make it straightforward to apply also to T_{12} (unlike the extractive methods which were applicable only to T_7).

Injecting supplementary questions. One hypothesis for the poor performance of models on the mix splits could be that the original bAbi tasks do not provide enough supervision for models to learn the basic event semantics. For example, tasks 5 and 7 are the only bAbi 1.0 tasks featuring the GIVE event, and neither includes any questions about the location of participants. However, test-time compositional questions may require models to infer that the participants in a GIVE event

Train	Test	Evaluation accuracy / # supporting facts			
		1	2	3+	Total
inject(T7)	concat(T7)	99.83	100	93.35	99.05
inject(T7)	mix(T7)	89.82	80.55	64.16	71.57
diverse(T7)	concat(T7)	99.58	100	78.36	96.94
diverse(T7)	mix(T7)	100	98.44	93.84	95.8
inject(T12)	concat(T12)	99.94	99.97	91.91	99.35
inject(T12)	mix(T12)	92.45	85.29	67.67	72.2
diverse(T12)	concat(T12)	99.75	98.73	76.81	97.73
diverse(T12)	mix(T12)	99.01	96.29	81.24	84.82

Table 4: Enriching the training data. Injecting knowledge to the original bAbi tasks doesn’t substantially improve compositionality. Sampling more structurally diverse instances yields more significant improvements, though is still limited, especially for more complex compositions.

share the same location (e.g., line 10 question in Fig. 1c). Error analysis shows that such implicit inferences are indeed challenging for models trained on the $concat$ splits (see details in appendix A.5). Perhaps the $inject$ splits supplementing the original tasks with relevant information will improve compositionality performance? Table 4 displays the result of this experiment; performance on mix is improved only marginally, despite a 3-fold increase in training data (Table 2).

Sampling structurally diverse training data. As shown in Table 2, though $inject$ splits significantly increase dataset size, their diversity remains low: most questions require only one or two supporting facts. Therefore, we next enrich training data through sampling more structurally diverse samples. This method is known to improve data efficiency for both compositional generalization as well as IID settings (Oren et al., 2021). As can be seen in Table 4, training on the $diverse$ splits yields a more significant improvement; similar to the findings of Oren et al. (2021), sampling more diverse training data leads to greater

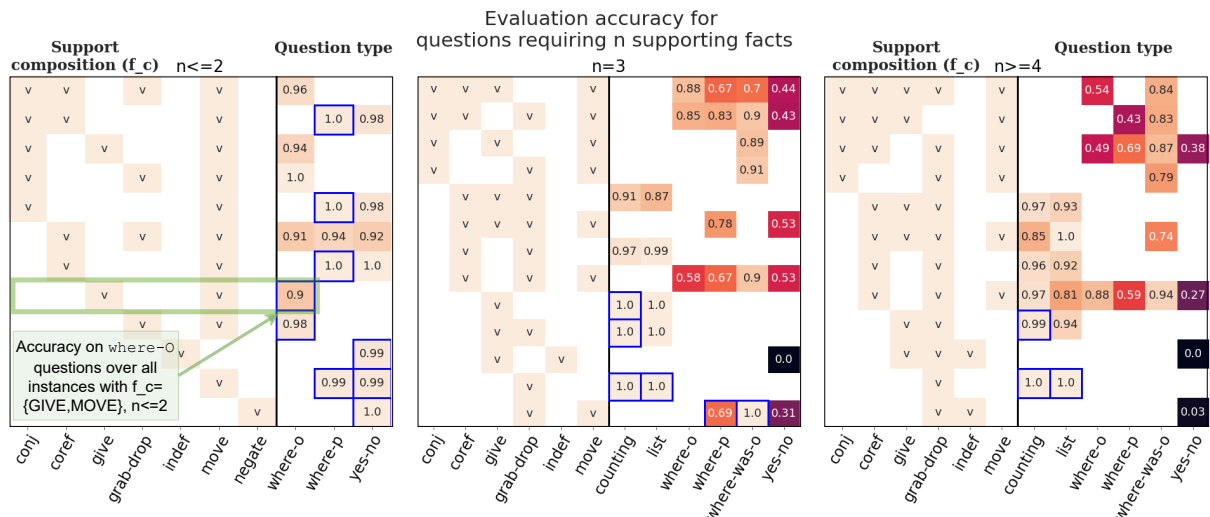


Figure 3: Error analysis on $mix(T_{12})$ for T5 trained on $diverse(T_{12})$ data. The sub-plots break down performance on questions requiring $\{\leq 2, 3, \geq 4\}$ supporting facts. For each sub-plot, the left side of each row corresponds to a particular support composition (f_c), and the right side displays accuracy over inputs sharing f_c , across various question types. Performance on f_c seen at training time (blue frames) is generally high, but overall generalization is not systematic, as evidenced by high variance across different f_c , especially for higher complexity ($n = 3, n \geq 4$) and more novel compositions.

generalization as well as much improved data efficiency.¹⁰ However, as the error analysis of the next section shows, performance on compositional generalization is still fundamentally limited.

4.2.1 Discussion and error analysis

Figure 3 breaks down the performance of T5 on $mix(T_{12})$ after training on $diverse(T_{12})$. The heatmaps plot performance across various support compositions (f_c) occurring in the test data, subdivided by the number of required supporting facts n per question. Performance on support compositions seen at training time (blue frames) is generally high, indicating the importance of training pattern diversity for better generalization. The plots indicate that T5 shows some ability to generalize to new support compositions, especially for lower n . Furthermore, certain question types appear to be more learned more robustly; for *list* and *count* questions, performance remains relatively high even for larger n and across novel f_c . We hypothesize that such questions may be easier as simple counting rules suffice to reach an answer, and these are “close to the surface”; unlike other events that may implicitly convey

```

1 Bill and Jeff moved to the park.
2 Following that they journeyed to the bathroom.
3 Bill is either in the hallway or the office.
4 Jeff picked up the apple.
5 Following that he dropped the apple.
6 John is in the school.
7 Fred is either in the garden or the office.
8 Mary is in the bedroom.
9 Bill grabbed the milk.
10 Afterwards he grabbed the football.
11 Julie and John travelled to the bedroom.
12 Bill is either in the kitchen or the bathroom.
13 Daniel is in the hallway.
14 Sandra is in the school.
15 Bill got the apple.
16 Jeff travelled to the garden.
17 After that he travelled to the bathroom.
18 Daniel is either in the garden or the school.
19 Bill dropped the apple.
20 Bill handed the milk to Jeff.
21 Is Bill in the bathroom? yes 1 2 4 5 15 T5: maybe
22 Where is Bill?bathroom 1 2 4 5 15 T5: bathroom

```

Figure 4: Example $mix(T_{12})$ instance demonstrating the question phrasing sensitivity failure mode in T5: the model correctly answers the question in *where-P* form (line 22), and incorrectly in *yes-no* form (line 21).

information, in our stories, changes of possession are always explicit in the text.

In general however, the plots indicate that T5 is far from robust compositional generalization:

Performance deteriorates with increased complexity. Performance is near perfect for simple compositions ($n \leq 2$) but deteriorates significantly for more complex cases ($n \geq 3$).

¹⁰The relatively low performance of *diverse* trained models in the “3+” column for *concat* splits is predominantly due to length discrepancies at train and test time: *concat* contains some very long stories which are challenging for the model trained on the uniform length and shorter *diverse* stories.

Question phrasing sensitivity. The discrepancy between the relatively high performance on *where-P* questions compared with very low performance on *yes-no* questions suggests that models are learning highly question-dependent story representations. E.g., if a model answers y correctly to some “Where is p ?” question, we would expect it to answer “yes” correctly for the same question in *yes-no* format, “Is p at y ?”. Figure 4 shows a characteristic example: T5 answers correctly in the *where-P* format, but incorrectly answers “maybe” for the *yes-no* format, likely thrown off by the distractor indefinite phrase in sentence 3. The pattern of answering “maybe” to questions about the location of an actor mentioned in an indefinite is commonly observed in training.

We present further empirical support for this finding in appendix A.6. These results suggest models may be learning shortcuts that work well for the story/question pairs seen at training time, but not more robust rules that also generalize to novel test instances (Geirhos et al., 2020). Such highly question-dependent story representation stands in contrast to more human-like narrative comprehension, which is thought to involve the construction of *situation models*, or structured representations of entities and their relations as depicted by the text. Situation models are less dependent on a-priori knowledge of a question (or its phrasing), and are often generated on-line during the course of comprehension, supporting a wide range of inferences (Graesser et al., 1994; Cayol and Nazir, 2020).

Performance below chance for certain question types. The heatmaps expose a particularly challenging class of *yes-no* questions involving disjunctions over indefinites (center and right plots, bottom right); accuracy for such questions is close to zero. See appendix A.7 for an example instance.

5 Future work & conclusions

Our work opens up multiple new directions for future research. Our new tool, Dyna-bAbI is readily extendable for systematic probing of more diverse linguistic phenomena. A beneficial first step could include integration of additional bAbI tasks. That said, our experience suggests that the design of truly scalable synthetic and dynamic benchmarks poses significant theoretical and engineering challenges, warranting deeper research on their own right.

Our results raise new questions about the viability of learning robust situation models using standard question-answering training, and our datasets present new challenges for future efforts.

Additionally, Dyna-bAbI can naturally complement parallel work probing the the situation representations constructed by neural language models (Li et al., 2021) by facilitating tailored data generation for specific questions, thus broadening and deepening the scope of possible research.

In conclusion, we introduced Dyna-bAbI, a new framework for highly controllable bAbI task generation. We used it to create compositional generalization datasets providing new modelling challenges for state-of-the-art neural language models. More broadly, our results underscore the importance in development of benchmarks themselves, beyond only the models solving them.

Broader Impact

While large, neural language models are increasingly seen as foundations for a wide array of NLP tasks, we still lack a clear understanding of their capabilities and failure modes. Our work joins many recent efforts using carefully controlled synthetic tasks to more rigorously evaluate models’ language comprehension abilities.

While our choice of a synthetic language benchmark allows more precise control over evaluation, the synthetic nature of the data is an obvious limitation. Similar to the original bAbI benchmark, our tasks are not a substitute for real natural language datasets, but should rather complement them. Even if a method works well on our data, it should be shown to perform well on real data as well. Rather, our tasks are better thought of as comprehension “unit-tests”, where poor performance on our tasks serves as a warning sign suggesting the model may exhibit limited systematicity and robustness on more difficult, naturalistic inputs.

Acknowledgements

We thank the Aristo team at the Allen Institute for AI for valuable support and feedback at various stages of this work. Ronen Tamari was supported by the Center for Interdisciplinary Data-science Research (CIDR) at HUJI. This work was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant no.

852686, SIAM, Shahaf). Part of this research is also supported by the European Research Council, ERC-StG grant no. 677352 (Tsarfaty), which we gratefully acknowledge.

References

- Laura Aina, Carina Silberer, Ionut-Teodor Sorodoc, Matthijs Westera, and Gemma Boleda. 2019. [What do entity-centric models learn? insights from entity linking in multi-party dialogue](#). 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 3772–3783, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pratyay Banerjee, Chitta Baral, Man Luo, Arindam Mitra, Kuntal Pal, Tran C. Son, and Neeraj Varshney. 2020. [Can transformers reason about effects of actions?](#) *Computing Research Repository*, arXiv:2012.09938.
- Andrea Banino, Adrià Puigdomènech Badia, Raphael Köster, Martin J. Chadwick, Vinicius Zambaldi, Demis Hassabis, Caswell Barry, Matthew Botvinick, Dharshan Kumaran, and Charles Blundell. 2020. [Memo: A deep network for flexible combination of episodic memories](#). In *International Conference on Learning Representations*.
- Andrea Banino, Jan Balaguer, and Charles Blundell. 2021. [Pondernet: Learning to ponder](#). In *8th ICML Workshop on Automated Machine Learning (AutoML)*.
- Gregor Betz, Christian Voigt, and Kyle Richardson. 2021. Critical thinking for language models. *Proceedings of IWCS*.
- Lukas Biewald. 2020. [Experiment tracking with weights and biases](#). Software available from wandb.com.
- Samuel R. Bowman and George Dahl. 2021. [What will it take to fix benchmarking in natural language understanding?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4843–4855, Online. Association for Computational Linguistics.
- Zoé Cayol and Tatjana A Nazir. 2020. [Why Language Processing Recruits Modality Specific Brain Regions: It Is Not About Understanding Words, but About Modelling Situations](#). *Journal of Cognition*, 3(1):1–23.
- Peter Clark, Oyvind Tafjord, and Kyle Richardson. 2020. [Transformers as soft reasoners over language](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3882–3890. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Bhavana Dalvi, Lifu Huang, Niket Tandon, Wen-tau Yih, and Peter Clark. 2018. [Tracking state changes in procedural text: a challenge dataset and models for process paragraph comprehension](#). 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 1595–1604, New Orleans, Louisiana. Association for Computational Linguistics.
- 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.
- Jesse Dunietz, Greg Burnham, Akash Bharadwaj, Owen Rambow, Jennifer Chu-Carroll, and Dave Ferrucci. 2020. [To test machine comprehension, start by defining comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7839–7859, Online. Association for Computational Linguistics.
- William Falcon et al. 2019. Pytorch lightning. *GitHub*. Note: <https://github.com/PyTorchLightning/pytorch-lightning>, 3.
- Catherine Finegan-Dollak, Jonathan K. Kummerfeld, Li Zhang, Karthik Ramanathan, Sesh Sadasivam, Rui Zhang, and Dragomir Radev. 2018. [Improving text-to-SQL evaluation methodology](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 351–360, Melbourne, Australia. Association for Computational Linguistics.
- Robert Geirhos, Jörn Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. [Shortcut learning in deep neural networks](#). *Nature Machine Intelligence*, 2(11):665–673.
- Nicolas Gontier, Koustuv Sinha, Siva Reddy, and Christopher Pal. 2020. [Measuring systematic generalization in neural proof generation with transformers](#). In *Advances in Neural Information Processing Systems 33*. Curran Associates, Inc.
- Arthur C. Graesser, Murray Singer, and Tom Trabasso. 1994. [Constructing Inferences During Narrative Text Comprehension](#). *Psychological Review*, 101(3):371–395.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*,

- pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Mikael Henaff, Jason Weston, Arthur Szlam, Antoine Bordes, and Yann LeCun. 2017. Tracking the world state with recurrent entity networks. 5th International Conference on Learning Representations, ICLR 2017 ; Conference date: 24-04-2017 Through 26-04-2017.
- Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzić, Rishabh Krishnan, and Dawn Song. 2020. [Pretrained transformers improve out-of-distribution robustness](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2744–2751, Online. Association for Computational Linguistics.
- Divyansh Kaushik and Zachary C. Lipton. 2018. [How much reading does reading comprehension require? a critical investigation of popular benchmarks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5010–5015, Brussels, Belgium. Association for Computational Linguistics.
- Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. 2020. [Measuring compositional generalization: A comprehensive method on realistic data](#). In *International Conference on Learning Representations*.
- Tushar Khot, Kyle Richardson, Daniel Khoshdel, and Ashish Sabharwal. 2021. [Learning to Solve Complex Tasks by Talking to Agents](#). *arXiv preprint arXiv:2110.08542*.
- Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. [Dynabench: Rethinking benchmarking in nlp](#).
- Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A method for stochastic optimization](#).
- Brenden Lake and Marco Baroni. 2018. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International conference on machine learning*, pages 2873–2882. PMLR.
- Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. 2017. Building machines that learn and think like people. *Behavioral and brain sciences*, 40.
- Hung Le, Truyen Tran, and Svetha Venkatesh. 2020. Self-attentive associative memory. In *International Conference on Machine Learning*, pages 5682–5691. PMLR.
- Belinda Z. Li, Maxwell Nye, and Jacob Andreas. 2021. [Implicit representations of meaning in neural language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1813–1827, Online. Association for Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Nelson F. Liu, Tony Lee, Robin Jia, and Percy Liang. 2021. [Can small and synthetic benchmarks drive modeling innovation? a retrospective study of question answering modeling approaches](#). *Computing Research Repository*, arXiv:2102.01065.
- Nelson F. Liu, Roy Schwartz, and Noah A. Smith. 2019. [Inoculation by fine-tuning: A method for analyzing challenge datasets](#). 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 2171–2179, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Ro{bert}a: A robustly optimized {bert} pretraining approach](#).
- Reginald Long, Panupong Pasupat, and Percy Liang. 2016. [Simpler context-dependent logical forms via model projections](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1456–1465, Berlin, Germany. Association for Computational Linguistics.
- James L. McClelland, Felix Hill, Maja Rudolph, Jason Baldridge, and Hinrich Schütze. 2020. [Placing language in an integrated understanding system: Next steps toward human-level performance in neural language models](#). *Proceedings of the National Academy of Sciences*, arXiv:1707(Xx):201910416.
- Inbar Oren, Jonathan Herzig, and Jonathan Berant. 2021. [Finding needles in a haystack: Sampling structurally-diverse training sets from synthetic data for compositional generalization](#).
- Barbara Partee et al. 1995. Lexical semantics and compositionality. *An invitation to cognitive science: Language*, 1:311–360.

- Christopher Potts, Zhengxuan Wu, Atticus Geiger, and Douwe Kiela. 2021. [DynaSent: A dynamic benchmark for sentiment analysis](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2388–2404, Online. Association for Computational Linguistics.
- Linlu Qiu, Hexiang Hu, Bowen Zhang, Peter Shaw, and Fei Sha. 2021. [Systematic generalization on gscan: What is nearly solved and what is next?](#) *Computing Research Repository*, arXiv:2109.12243.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Kyle Richardson, Hai Hu, Lawrence Moss, and Ashish Sabharwal. 2020. [Probing natural language inference models through semantic fragments](#). In *Proceedings of AAAI*.
- Kyle Richardson and Ashish Sabharwal. 2022. Pushing the limits of rule reasoning in transformers through natural language satisfiability. *Proceedings of AAAI*.
- Anna Rogers. 2021. [Changing the world by changing the data](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2182–2194, Online. Association for Computational Linguistics.
- Laura Ruis, Jacob Andreas, Marco Baroni, Diane Bouchacourt, and Brenden M Lake. 2020. [A benchmark for systematic generalization in grounded language understanding](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 19861–19872. Curran Associates, Inc.
- Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. “everyone wants to do the model work, not the data work”: Data cascades in high-stakes ai. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, New York, NY, USA. Association for Computing Machinery.
- Imanol Schlag, Tsendsuren Munkhdalai, and Jürgen Schmidhuber. 2021. [Learning associative inference using fast weight memory](#). In *International Conference on Learning Representations*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. [Bidirectional attention flow for machine comprehension](#). In *International Conference on Learning Representations*.
- Murray Shanahan, Matthew Crosby, Benjamin Beyret, and Lucy Cheke. 2020. Artificial intelligence and the common sense of animals. *Trends in Cognitive Sciences*, 24(11):862–872.
- Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. 2019. [CLUTRR: A diagnostic benchmark for inductive reasoning from 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 4506–4515, Hong Kong, China. Association for Computational Linguistics.
- Saku Sugawara, Pontus Stenetorp, and Akiko Aizawa. 2021. [Benchmarking machine reading comprehension: A psychological perspective](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1592–1612, Online. Association for Computational Linguistics.
- Ronen Tamari, Chen Shani, Tom Hope, Miriam R L Petruck, Omri Abend, and Dafna Shahaf. 2020. [Language \(re\)modelling: Towards embodied language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6268–6281, Online. Association for Computational Linguistics.
- Niket Tandon, Keisuke Sakaguchi, Bhavana Dalvi, Dheeraj Rajagopal, Peter Clark, Michal Guerquin, Kyle Richardson, and Eduard Hovy. 2020. [A dataset for tracking entities in open domain procedural text](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6408–6417, Online. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomáš Mikolov. 2016. [Towards ai-complete question answering: A set of prerequisite toy](#)

tasks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Zhengxuan Wu, Elisa Kreiss, Desmond C. Ong, and Christopher Potts. 2021. [ReaSCAN: Compositional reasoning in language grounding](#). *NeurIPS 2021 Datasets and Benchmarks Track*.

Hitomi Yanaka, Koji Mineshima, and Kentaro Inui. 2021. [SyGNS: A systematic generalization testbed based on natural language semantics](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 103–119, Online. Association for Computational Linguistics.

Rolf A. Zwaan. 2016. [Situation models, mental simulations, and abstract concepts in discourse comprehension](#). *Psychonomic Bulletin and Review*, 23(4):1028–1034.

A Appendix

A.1 Extended task construction details

This section provides further details of the training and test splits used for our experiments.

Table 5 enumerates the basic “building blocks”, or concepts underlying the tasks, as presented in §3.2.

Tables 6 and 7 detail the concept sets for each of the sub-tasks comprising the training and test sets, for the T_2 , T_7 and T_{12} groups of tasks.

As can be seen from the tables, the main sources of compositionality are:

- Following the bAbI 1.0 task structure, at training time, all of the more complex linguistic constructs are seen only with MOVE events (and none of the other event types).
- Similarly, at training time, *yes-no* questions are always seen only with MOVE events (and none of the other event types), and with the INDEF or NEGATE linguistic constructs (but not others, such as COREF).
- *where-was-O* questions are never seen in stories with GIVE events.

Language templates. For our new generated tasks we use the same language templates as used in the original bAbI 1.0 benchmark (e.g., the same entity names, verb synonyms). The only modification to the language generation engine was that we completely omit the use of “there”; in the original benchmark, “there” could be used in confusing contexts, as shown in Fig. 5.

A.1.1 Example instances

Figure 6 shows examples from each of the 4 types of splits used in our experiments. The *concat* instance is from the original bAbI 1.0 task 5. The *inject* data contains the same passages as *concat*, but adds supplementary questions on agent and object locations. *diverse* instances

1 Mary journeyed to the bathroom. 2 Sandra went to the garden. 3 Daniel went back to the garden. 4 Daniel went to the office. 5 Sandra grabbed the milk there. 6 Sandra put down the milk there. 7 Where is the milk? garden 6 2

Figure 5: Example from original bAbI 1.0 benchmark with confusing usage of “there”. In Dyna-bAbI we do not include “there”, to avoid this confusion.

contain more diverse support compositions (f_c), but certain combinations are held out. In particular, *diverse* instances only feature non-default linguistic mappings with MOVE events, never with POSS (GRAB or DROP) or GIVE. In the *mix* instances, all combinations of support compositions are possible, as shown in the example which features possession (POSS) events along with co-reference.

A.1.2 Long instances in the bAbI 1.0 tasks

For the T5 experiments, we used a slightly modified version of the bAbI 1.0 tasks, where we trimmed all training and validation examples that didn’t fit into the 512-token input window. This resulted in trimming 1,585 training instances and 175 validation instances from T_7 and T_{12} (common to both sets). These data points are not consequential as our analysis focuses on the effects of compositionality and not story length; all instances in *diverse* and *mix* are substantially shorter than the 512-token maximum input window size.

A.2 Implementation details

T5. We use the publicly available HuggingFace pre-trained T5-base implementation (Wolf et al., 2020) which has 220M parameters. We similarly use the HuggingFace tokenization pipeline. We fine-tune T5 for 12 epochs on our bAbI data, using the Adam optimizer (Kingma and Ba, 2017), an initial learning rate of $5 * 10^{-5}$ and training batch size of 8.

STM. We used the official STM implementation¹¹, with the only change being a batch size of 32 instead of 128, due to technical constraints.

EntNet. We re-implemented the model in PyTorch, similarly using a batch-size of 32. Following the official Lua reference implementation¹², we used 20 memory units each with dimension 100. We used the SGD optimizer.

For both the EntNet and STM, we trained models for 200 epochs, and took the best of 10 tries, following Henaff et al. (2017).

For the 20-model concurrence benchmark, refer to Liu et al. (2021) for model details, as we used the same experimental setup.

¹¹<https://github.com/thaihungle/SAM>

¹²<https://github.com/facebookarchive/MemNN/tree/master/EntNet-babi>

Events	Template	Example	Notes
MOVE	P {moved} to the L.	John traveled to the park.	
GRAB	P {grabbed} the O.	Mary picked up the apple.	
DROP	P {dropped} the O.	Daniel dropped the milk.	
GIVE	P1 {gave} P2 the O.	John handed Mary the apple.	
Linguistic Constructs			
COREF	P (MOVE GRAB DROP) Following that, {he} (MOVE GRAB DROP).	John went to the garden. Following that, he moved to the store	Co-reference
CONJ	P1 and P2 {moved} to the L1.	Jeff and Fred went to the cinema.	Conjunction
COMPOUND	P1 and P2 {moved} to the L1. Then they {moved} to the L2.	Jeff and Fred went to the cinema. Then they traveled to the school.	Compound co-reference
NEGATE	P is not at the L.	Julie is not in the park.	Negation
INDEF	P is either at the L1 or the L2.	John is either in the park or the school.	Indefinite expression
Questions			
where-P	Where is P?	Where is John?	
where-O	Where is the O?	Where is the football?	
where-was-O	Where was the O before the L?	Where was the football before the hallway?	
yes-no	Is P at the L?	Is John at the park?	
list	What is P carrying?	What is John carrying?	
counting	How many objects is P carrying?	How many objects is John carrying?	
give-qs	Who gave the O to P2?	Who gave the football to John?	Constitutes multiple question types over GIVE events.
	Who gave the O?	...	
	Who received the O?		
	Who did P1 give the P2 to?		
	What did P1 give to P2?		

Table 5: Details of the events, linguistic constructs and questions constituting the bAbI tasks covered in this work. Words in {brackets} are drawn from a small set of synonyms.

<p>concat(T12) + inject(T12)</p> <ol style="list-style-type: none"> 1 Bill travelled to the office. 2 Bill picked up the football there. 3 Bill went to the bedroom. 4 Bill gave the football to Fred. 5 What did Bill give to Fred? football {4} 6 Where is the football? bedroom {3, 4} 7 Where is Bill? bedroom {3} 8 Where is Fred? bedroom {3, 4} 	<p>diverse(T12)</p> <ol style="list-style-type: none"> 1 Fred went back to the garden. 2 Sandra travelled to the cinema. 3 Fred went to the bathroom. 4 Fred got the football. 5 Fred travelled to the garden. 6 Bill journeyed to the garden. 7 Fred passed the football to Bill. 8 Bill discarded the football. 9 Jeff got the football. 10 Jeff discarded the football. 11 Sandra journeyed to the office. 12 Fred journeyed to the kitchen. 13 Bill got the football. 14 Bill travelled to the office. 15 Bill passed the football to Julie. 16 Julie passed the football to Daniel. 17 Daniel left the football. 18 Mary journeyed to the bedroom. 19 Bill picked up the football. 20 Bill left the football. 21 Where is Jeff? garden <p>f = {6, 8, 9} f_c = {MOVE, POSS}</p>	<p>mix(T12)</p> <ol style="list-style-type: none"> 1 John is no longer in the bedroom. 2 Bill is in the bedroom. 3 Bill took the apple. 4 Afterwards he discarded the apple. 5 Bill is no longer in the bedroom. 6 Daniel is either in the kitchen or the bathroom. 7 Fred and Bill journeyed to the kitchen. 8 Jeff is either in the park or the office. 9 Daniel is either in the garden or the kitchen. 10 Sandra is in the school. 11 Bill is either in the bathroom or the school. 12 Mary is not in the office. 13 Sandra journeyed to the hallway. 14 After that she grabbed the milk. 15 Julie is either in the bedroom or the office. 16 Daniel is no longer in the garden. 17 Jeff moved to the bathroom. 18 Julie picked up the apple. 19 Following that she got the football. 20 Jeff is in the hallway. 21 Where is the football? bedroom <p>f = {2, 3, 4, 18, 19} f_c = {MOVE, POSS, COREF}</p>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 6: Example instances from each of the 4 types of splits used in our experiments.

		Events				Linguistic Constructs			Questions			
Sub-task	Type	Move	Grab	Drop	Give	Co-reference	Conjunction	Compound co-ref.	where-P	where-O	where-was-O	give
1	Train	✓							✓			
2	Train	✓	✓	✓					I/D	✓		
3	Train	✓	✓	✓					I	I	✓	
5	Train	✓	✓	✓	✓				I/D	I/D		✓
11	Train	✓				✓			✓			
12	Train	✓					✓		✓			
13	Train	✓						✓	✓			
mix(T_2)	Test	✓	✓	✓		✓			✓	✓		
mix(T_7)	Test	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	

Table 6: Concept sets for the T_2 and T_7 sub-set of the original bAbI tasks, and the new tasks generated with Dyna-bAbI. Train sub-task numbering follows the original bAbI numbering. The *inject* and *diverse* tasks inherit the same concept set from the original tasks, and additionally “I”, “D” denote question types included only in the *inject* or *diverse* tasks, respectively. “I/D” denotes question types included in both.

		Events				Linguistic Constructs				Questions							
Task	Type	Move	Grab	Drop	Give	Co-reference	Conjunction	Compound co-ref.	Negation	Indefinite	where-P	where-O	where-was-O	yes-no	counting	list	give
1	Train	✓								✓							
2	Train	✓	✓	✓						I/D	✓						
3	Train	✓	✓	✓						I	I	✓					
5	Train	✓	✓	✓	✓					I/D	I/D						✓
6	Train	✓								I/D			✓				
7	Train	✓	✓	✓	✓					I	I			✓			
8	Train	✓	✓	✓						I	I				✓		
9	Train	✓						✓		I/D			✓				
10	Train	✓							✓	I/D			✓				
11	Train	✓				✓				✓							
12	Train	✓					✓			✓							
13	Train	✓						✓		✓							
mix(T_{12})	Test	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 7: Concept sets for the T_{12} sub-set of the original bAbI tasks, and the new tasks generated with Dyna-bAbI. Train sub-task numbering follows the original bAbI numbering. The *inject* and *diverse* tasks inherit the same concept set from the original tasks, and additionally “I”, “D” denote question types included only in the *inject* or *diverse* tasks, respectively. “I/D” denotes question types included in both.

For the T5 experiments, we used the PyTorch Lightning (Falcon et al., 2019) trainer implementation, and Weights & Biases (Biewald, 2020) for experiment tracking and artifacts management.

We used standard hyper-parameter settings for all models, with slight changes in the case of memory issues as described above.

Experimental infrastructure details. Our experiments were performed using an RTX-8000 GPU, with a total computational budget of roughly 1,000 GPU hours.

A.3 Inoculation experiment results

To rule out the hypothesis that certain patterns may be too hard for models to learn, we follow the inoculation methodology presented in Liu et al. (2019): after training on the original tasks, we fine-tune the T5 on small amounts of OOD data (disjoint from the test data), and evaluate performance as a function of “inoculation dose”. As can be seen in Fig. 7, we find that performance quickly (with only 500 additional inoculation samples per question type) reaches over 90% accuracy on both the $mix(T_7)$ and $mix(T_{12})$ challenge sets. These results support the hypothesis that the training data is not rich enough, indicating clearly that the model is capable of quickly learning to solve the challenge tasks, given exposure to training samples with similar enough patterns.

A.4 Concurrence experiments

Table 8 presents the full results for the concurrence experiments of §4.1. SQuAD and bAbI task 2 results are reproduced from Liu et al. (2021), see there also for implementation details of the models used.

A.5 Extended error analysis: GIVE events

We analyze the performance of models on the $mix(T_7)$ split after being trained on $concat(T_7)$, and in particular we focus on GIVE events. As noted in §4.2, compositions involving GIVE are intuitively challenging as they entail multiple inferences which are not explicit in the text: the actors share the same location, and the possession of the object being given is transferred from the giver to the recipient. The only task in $concat(T_7)$ featuring GIVE events is task 5, which never asks about the locations of actors or objects, but only about the participant roles in the event (e.g., who was the giver or recipient; see Fig. 1 example from task 5).

Model	Evaluation accuracy			
	SQuAD	mix(T2)	mix(T7)	babi task 2
rasor	64.86	88.20	35.03	100.00
bidaf	67.39	97.20	30.50	100.00
documentreader	69.66	90.20	40.70	100.00
documentreader (no_features)	69.21	82.50	37.17	100.00
bidafplusplus	69.49	99.50	44.20	80.70
mnemonicreader	73.02	98.20	39.63	100.00
mnemonicreader (no_features)	72.67	97.50	38.20	100.00
qanet	72.41	67.70	-	100.00
fusionnet	72.90	99.50	39.73	100.00
fusionnet (no_features)	72.24	88.10	37.80	100.00
bert	81.46	95.50	47.63	100.00
bert_large	84.17	98.30	59.10	100.00
bert_large_wwm	87.33	98.70	67.63	99.90
albert	81.86	98.20	56.70	100.00
albert_xxlarge	89.07	99.80	80.00	100.00
roberta	83.37	98.70	57.70	100.00
roberta_large	86.96	99.80	64.07	100.00
electra	85.88	98.70	53.47	100.00
spanbert	86.20	98.40	55.70	99.50
spanbert_large	88.74	98.60	62.27	95.40

Table 8: Full results of concurrence experiments presented in §4.1.




Figure 7: Inoculation experiment results.

# supporting facts	# samples	Evaluation accuracy		
		BiDAF	RoBERTa	T5
1	334	53.3	93.4	86.8
2 (without GIVE)	734	51.50	82.3	71.8
2 (with GIVE)	99	3.03	7.07	5.05
3 (without GIVE)	1365	24.6	47.2	44.3
3 (with GIVE)	468	4.27	7.05	15.2

Table 9: Breakdown of model performance on $mix(T_7)$ for questions including GIVE events in the supporting fact set. The poor performance indicates that training on the bAbI 1.0 data does not facilitate generalization to novel compositions of GIVE.

<i>where-P</i> (\rightarrow)	correct	incorrect
<i>yes-no</i> (\downarrow)		
correct	209	4
incorrect	145	88

Table 10: Confusion matrix displaying question phrasing sensitivities in T5. We pose a question in two formats: (1) *yes-no*: “Is X at L ? yes” vs (2) *where-P*: “Where is X ? L ”. We find performance is considerably higher for questions posed in the *where-P* format, indicating the model isn’t learning the equivalence of both forms.

To measure this intuition empirically, we analyze a subset of 567 questions including GIVE events in the supporting facts set. As shown in Table 9, performance for all models on questions including GIVE is extremely low, far below performance for questions without it. Qualitative analysis indicates many failure cases follow the pattern shown in the right-side example of Fig. 1c, question on line 10: the location of an entity (e.g., Daniel) must be inferred via the known (co-)location of a second participant in the GIVE event (e.g., Jeff). These results strengthen the hypothesis that standard QA training on the original bAbI data does not drive strong event comprehension in models.

A.6 Extended error analysis: question phrasing sensitivity

This section presents further empirical analysis of the question phrasing sensitivities discussed in §4.2.1, relating to the performance of the T5 model trained on the *diverse*(T_{12}) data and evaluated on the challenge set *mix*(T_{12}).

We collected all *yes-no* questions from *mix*(T_{12}) for which the answer was “yes”, yielding 446 questions in total. For each such (question, answer) pair, of the form (“Is *person* at the *location*?”, “yes”), we created an equivalent pair in the format of a *where-P* question, (“Where is *person*?”, *location*). Figure 4 shows a characteristic example. Ideally, we would expect a model to be agnostic to equivalent phrasings of a question. However, as displayed in Table 10, we find that T5 is considerably more accurate for questions posed in the *where-P* format, likely due to exposure to a larger variety of such questions at training time.

1 Bill grabbed the milk.
2 Bill put down the milk.
3 John is either in the bedroom or the kitchen.
4 Fred journeyed to the kitchen.
5 John grabbed the football.
6 Following that he put down the football.
7 Bill picked up the milk.
8 Following that he went to the bedroom.
9 Bill is in the office.
10 Bill is in the cinema.
11 Bill passed the milk to Julie.
12 Julie handed the milk to Bill.
13 Jeff is not in the school.
14 John took the football.
15 Fred and Jeff moved to the school.
16 Afterwards they journeyed to the bathroom.
17 Bill handed the milk to Julie.
18 John dropped the football.
19 Daniel is either in the school or the bedroom.
20 Daniel took the football.
21 Is John in the bedroom? yes 3 18 19 20

Figure 8: Double disjunction example from *mix*(T_{12}).

A.7 Extended error analysis: double disjunctions

As the shown in the §4.2.1 error analysis, a particularly difficult class of questions are double disjunctions over indefinite expressions. Figure 8 displays a typical example from *mix*(T_{12}), where the locations of two actors are given in indefinite form (sentences 3 and 19), and are also known to be co-located, since they share the location of the object “football”, as inferred from sentences 18 and 20. Hence it is possible to infer their location as the intersection of the two indefinite expressions (here “bedroom”). Rather than answering “yes” to the question “Is John in the bedroom?”, T5 invariably answers “maybe” for such cases. This pattern is likely due to the fact that in the training data “maybe” is a typical answer for *yes-no* questions about actors mentioned by indefinite expressions (task 10 in bAbI 1.0).

B Datasheet for datasets

Motivation

For what purpose was the dataset created?

Was there a specific task in mind? Was there a specific gap that needed to be filled? Please provide a description.

Few synthetic resources for probing NLP models' performance on discourse-level narrative understanding texts. Existing resources lack customizability (control over data created + amenable to extension).

Who created this dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?

Joint team of researchers at Hebrew University of Jerusalem (Israel) and the Allen Institute for Artificial Intelligence.

Who funded the creation of the dataset? If there is an associated grant, please provide the name of the grantor and the grant name and number.

Work was supported by the Center for Interdisciplinary Data-science Research (CIDR) at HUJI. This work was also supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant no. 852686, SIAM) and NSF-BSF grant no. 2017741 (Shahaf). Part of this research is also supported by the European Research Council, ERC-StG grant no. 677352 (Tsarfaty).

Any other comments?

Composition

What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)? Are there multiple types of instances (e.g., movies, users, and ratings; people and interactions between them; nodes and edges)? Please provide a description.

Instances represent variable length stories.

How many instances are there in total (of each type, if appropriate)?

Any size dataset can be created (programmatic generation).

Does the dataset contain all possible instances or is it a sample (not necessarily

random) of instances from a larger set?

If the dataset is a sample, then what is the larger set? Is the sample representative of the larger set (e.g., geographic coverage)? If so, please describe how this representativeness was validated/verified. If it is not representative of the larger set, please describe why not (e.g., to cover a more diverse range of instances, because instances were withheld or unavailable).

Used rejection sampling for some datasets to cover more diverse instances.

What data does each instance consist of? "Raw" data (e.g., unprocessed text or images) or features? In either case, please provide a description.

Simple textual stories generated using templates ("John went to the kitchen. He grabbed the apple.").

Is there a label or target associated with each instance? If so, please provide a description.

Each instance is accompanied by a (question, answer) pair, both in natural language.

Is any information missing from individual instances? If so, please provide a description, explaining why this information is missing (e.g., because it was unavailable). This does not include intentionally removed information, but might include, e.g., redacted text.

N/A

Are there recommended data splits (e.g., training, development/validation, testing)? If so, please provide a description of these splits, explaining the rationale behind them.

The data is organized in splits, which are explained in section 4 of the paper.

Are there any errors, sources of noise, or redundancies in the dataset? If so, please provide a description.

Template based language generation may result in somewhat unnatural texts.

Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)? If it links to or relies on external resources, a) are there guarantees that they will exist, and remain constant, over time; b) are there official archival versions of the complete dataset

(i.e., including the external resources as they existed at the time the dataset was created); c) are there any restrictions (e.g., licenses, fees) associated with any of the external resources that might apply to a future user? Please provide descriptions of all external resources and any restrictions associated with them, as well as links or other access points, as appropriate.

Self contained.

Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety? If so, please describe why.

No.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

No.

Any other comments?

Collection Process

How was the data associated with each instance acquired? Was the data directly observable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or language)? If data was reported by subjects or indirectly inferred/derived from other data, was the data validated/verified? If so, please describe how.

Programmatically generated using logical rules and templates.

If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?

Rejection sampling was used in some cases, described in Section 4.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

No.

Any other comments?

Preprocessing/cleaning/labeling

Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)?

If so, please provide a description. If not, you may skip the remainder of the questions in this section.

No.

Was the “raw” data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)? If so, please provide a link or other access point to the “raw” data.

N/A

Is the software used to preprocess/clean/label the instances available? If so, please provide a link or other access point.

N/A

Any other comments?

Uses

Has the dataset been used for any tasks already? If so, please provide a description.

Benchmark to guide model development for reading comprehension and textual reasoning tasks.

Is there a repository that links to any or all papers or systems that use the dataset? If so, please provide a link or other access point.

Not currently, we will use the <https://paperswithcode.com/> integration to track results.

What (other) tasks could the dataset be used for?

N/A

Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses? For example, is there anything that a future user might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other undesirable harms (e.g., financial harms,

legal risks) If so, please provide a description. Is there anything a future user could do to mitigate these undesirable harms?

N/A

Are there tasks for which the dataset should not be used? If so, please provide a description.

Similar to the original bAbI benchmark, our tasks are not a substitute for real natural language datasets, but should rather complement them. Even if a method works well on our data, it should be shown to perform well on real data as well. Rather, our tasks are better thought of as comprehension “unit-tests”, where poor performance on our tasks serves as a warning sign suggesting the model may exhibit limited systematicity and robustness on more difficult, naturalistic inputs.

Any other comments?

Distribution

Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created? If so, please provide a description.

N/A

How will the dataset will be distributed (e.g., tarball on website, API, GitHub) Does the dataset have a digital object identifier (DOI)?

GitHub + Weights and Biases. No DOI currently.

When will the dataset be distributed?

Data and code-base for task generation to be uploaded upon publication.

Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)? If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.

Will be available with standard MIT license.

Have any third parties imposed IP-based or other restrictions on the data associated with the instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any

relevant licensing terms, as well as any fees associated with these restrictions.

N/A

Do any export controls or other regulatory restrictions apply to the dataset or to individual instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any supporting documentation.

N/A

Any other comments?

Maintenance

Who will be supporting/hosting/maintaining the dataset?

Corresponding author of paper.

How can the owner/curator/manager of the dataset be contacted (e.g., email address)?

Via email with corresponding author, and through dedicated GitHub repository.

Is there an erratum? If so, please provide a link or other access point.

N/A

Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)? If so, please describe how often, by whom, and how updates will be communicated to users (e.g., mailing list, GitHub)?

Extensions will be maintained via GitHub.

If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were individuals in question told that their data would be retained for a fixed period of time and then deleted)? If so, please describe these limits and explain how they will be enforced.

N/A

Will older versions of the dataset continue to be supported/hosted/maintained? If so, please describe how. If not, please describe how its obsolescence will be communicated to users.

Data versioning supported natively through Weights and Biases.

If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so? If so, please provide a description. Will these contributions be validated/verified? If so, please describe how. If not, why not? Is there a process for communicating/distributing these contributions to other users? If so, please provide a description.

The codebase can be freely extended, we will only be responsible of course for changes to the main branch.

Any other comments?