

Cryptonite: A Cryptic Crossword Benchmark for Extreme Ambiguity in Language

Avia Efrat* Uri Shaham* Dan Kilman Omer Levy

Tel Aviv University

{avia.efrat, uri.shaham1}@gmail.com

Abstract

Current NLP datasets targeting ambiguity can be solved by a native speaker with relative ease. We present Cryptonite, a large-scale dataset based on *cryptic* crosswords, which is both linguistically complex and naturally sourced. Each example in Cryptonite is a *cryptic clue*, a short phrase or sentence with a misleading surface reading, whose solving requires disambiguating semantic, syntactic, and phonetic wordplays, as well as world knowledge. Cryptic clues pose a challenge even for experienced solvers, though top-tier experts can solve them with almost 100% accuracy. Cryptonite is a challenging task for current models; fine-tuning T5-Large on 470k cryptic clues achieves only 7.6% accuracy, on par with the accuracy of a rule-based clue solver (8.6%).

1 Introduction

The ambiguity of natural language is one of the most fundamental challenges in NLP research. While there are works and datasets specifically targeting ambiguity (Levesque et al., 2011; Raganato et al., 2017; Sakaguchi et al., 2020), these can be solved by a native speaker with relative ease. Can we design a dataset with ambiguities that pose a challenge even to competent native speakers?

We present Cryptonite, a large-scale dataset based on *cryptic* crosswords, which is both linguistically complex and naturally sourced. Cryptonite’s 523K examples are taken from professionally-authored cryptic crosswords, making them less prone to artifacts and biases than examples created by crowdsourcing (Gururangan et al., 2018; Geva et al., 2019). Each example in Cryptonite is a *cryptic clue*, a short phrase or sentence with a misleading surface reading, which poses a challenge even for humans experienced in cryptic crossword solving. A cryptic clue usually consists of two underlying parts: *wordplay* and *definition*. The




Figure 1: **How to solve the cryptic clue “One doesn’t like shifting earth (5)”**: Solving usually starts by figuring out which of the clue’s words belong to the *definition* (blue) and which to the *wordplay* (orange). Next, one needs to figure out the *type* of wordplay, which is often hinted by an *indicator* (purple). In our case, “shifting” hints that the answer is an anagram of some part of the wordplay. As the *enumeration* (gray) states the answer is a five-letter word, “earth” is a promising candidate for anagramming. Finally, given that “hater” is both an anagram of “earth” and a synonym of the definition, we conclude it to be the correct answer.

clue’s answer is both a disambiguation of the wordplay and, at the same time, directly answers the definition. While solving cryptic clues requires disambiguating semantic, syntactic, and phonetic wordplays, as well as world knowledge, clues are designed to have only one possible answer. See Section 1 for an example clue and its solution.

We provide a standard baseline by fine-tuning the generic T5-Large conditional language model (Raffel et al., 2020) on Cryptonite, achieving only 7.6% accuracy. For comparison, a rule-based cryptic clue solver (Deits, 2021) achieves 8.6% accuracy. These results highlight the challenge posed by Cryptonite, making it a candidate for assessing the disambiguation capabilities of future models.

Analyzing the results of both baselines, we find a correlation between performance on individual clues and a human assessment of the clue’s difficulty, and that the enumeration (answer length) is highly informative. Finally, we show that ensuring that the answers of train and test examples are mutually exclusive is critical for a candid estimation of T5’s ability to solve cryptic clues in general.

* Equal contribution.

2 Cryptic Crosswords

A cryptic crossword, just like a regular (non-cryptic) crossword, is a puzzle designed to entertain humans, comprised of *clues* whose answers are to be filled in a letter grid. Unlike regular crosswords, where answers are typically synonyms or hyponyms of the clues (Severyn et al., 2015), *cryptic* clues have a misleading surface reading, and solving them requires disambiguating wordplays. A cryptic clue has only one possible answer, even when taken outside the context of the letter grid.¹

Generally, a cryptic clue (henceforth, “clue”) consists of two parts: *wordplay* and *definition*. The wordplay-definition split is not given to the solver, and parsing it is usually the first step in solving a clue. Both the wordplay and the definition lead to the answer, but each in a different manner. While the definition is directly related to the answer (e.g. a synonym or a hypernym), the wordplay needs to be deciphered, usually with the help of an *indicator* that hints at the wordplay’s type.

Figure 2 walks through several clues with different types of wordplays and their solution. These examples provide only a glimpse into the rich and diverse world of cryptic crosswords. For a deeper dive, see the guidebook by Moorey (2018).

3 Dataset

We introduce Cryptonite, a dataset of 523,114 cryptic clues from 17,375 English-language crosswords published in The Times² and The Telegraph³ between October 2000 and October 2020.⁴


For preprocessing, we remove clue-answer duplicates and examples whose answer and enumeration do not match. In addition, we remove any examples with the same clue but with a different answer. While this occurred in less than 0.1% of the data, these examples violate the principle that a clue must have a single solution once the wordplay is deciphered (see Section 2).

¹For example, a regular crossword can have a clue like “A Sesame Street character (4)”. Without existing letters from the grid, a solver cannot determine if the answer is “elmo” or “bert”. This cannot happen with a cryptic clue, which by design does not require any letter grid information for solving.


²<https://www.thetimes.co.uk/puzzleclub/crosswordclub/home/crossword-cryptic>

³<https://puzzles.telegraph.co.uk/crossword-puzzles/cryptic-crossword>

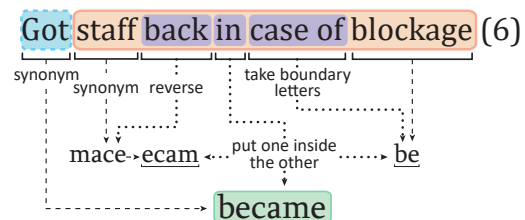
⁴We do not redistribute the dataset, but provide code for creating it with subscriptions to The Times and The Telegraph: <https://github.com/aviaefrat/cryptonite>



(a) A clue with a relatively simple *additive* wordplay, also requiring world knowledge. One can decipher the wordplay by identifying “and” as a concatenation indicator.



(b) Clues can also have *phonetic* wordplays. Here, “it’s said” implies that the wordplay is a homophone of the answer.



(c) Many clues combine more than one type of wordplay. This clue composes three: *reversing* the letters of a word (“back”), and *inserting* it (“in”) into the *boundary letters* of another (“case of”).

Getting fed up about midday (2,5)

A possible surface reading is “someone being mad about it being noon”.

A potential source of ambiguity is “fed up”. Try finding a different, perhaps more literal reading if the clue.

at lunch

(d) Although many wordplays can be roughly clustered into types and deciphered based on indicators, there is no silver bullet for solving cryptic crosswords. For this clue, even the standard wordplay-definition split does not apply; instead, the entire clue points to the answer.

Figure 2: Various examples of cryptic clues and how to solve them. Answers are in green, definitions in light blue (dashed frame), wordplays in orange, and indicators are in purple.

	Examples	Unique Answers	Clue Length	Answer Length
<i>train</i>	470,803	80,837	7.76	1.22
<i>valid</i>	26,156	4,534	7.77	1.20
<i>test</i>	26,157	4,538	7.77	1.24

Table 1: Overview of Cryptonite. Reported lengths are mean values. Words are delimited by spaces.

We follow the recent findings of Lewis et al. (2020), and split Cryptonite into train, validation, and test sets, where no answer is shared between them. Answer splitting creates a far more challenging benchmark for supervised models than naive random splits (see Section 4.3). Table 1 shows some basic statistics of the final Cryptonite dataset.

4 Experiments

We provide initial results on Cryptonite using two baselines: T5-Large (Raffel et al., 2020) and a rule-based cryptic clue solver (Deits, 2021). Despite training on half a million clues (T5) or being tailored to the task (rule-based solver), both approaches solve only a small portion of the test data, demonstrating that Cryptonite is indeed a challenging task. We further investigate two properties of the data: how difficulty (as perceived by humans) correlates with accuracy, and the informativeness of enumeration. In addition, we analyze how a naive data split affects the performance of T5, demonstrating that partitioning by answers is crucial for obtaining a candid estimate of the neural model’s ability to generalize to new cryptic clues.

4.1 Baselines

T5-Large Following current NLP methodology, we fine-tune the 770M parameter T5-Large (Raffel et al., 2020) on Cryptonite. The model’s encoder takes the clue as input, and uses the decoder to predict the answer using teacher forcing during training and beam search ($b = 5$) during inference.

We use HuggingFace (Wolf et al., 2020) with the recommended settings (Raffel et al., 2020), optimizing with AdaFactor (Shazeer and Stern, 2018) at a constant learning rate of 0.001. We train until convergence with a patience of 10 epochs and a batch size of 7000 tokens, selecting the best model checkpoint using validation set accuracy.

T5 uses SentencePiece tokenization (Kudo and Richardson, 2018), which might incur some information loss, as many clues require character-level manipulations.

Baseline	Accuracy	
	Validation	Test
T5-Large	7.44%	7.64%
Rule-based Solver	8.26%	8.58%

Table 2: Baseline performance on Cryptonite. Accuracy is measured using exact string match.

Rule-based Solver We also gauge the abilities of a rule-based solver with a manually-crafted probabilistic grammar (Deits, 2021). Building on the assumption that a clue can usually be split into a wordplay and a definition (Section 2), the solver tries to find the most probable parse such that the wordplay yields a semantically-similar result to the definition. The similarity between the definition and the parsed wordplay is calculated using expert-authored resources such as WordNet (Miller, 1995). Some less frequent wordplay types, such as *homophones* (Figure 2b) and *hidden-at-intervals* (Moorey, 2018, Chapter 3), are not implemented in the solver’s grammar.

4.2 Main Benchmark

We first evaluate our baselines on the main dataset. Table 2 shows that both approaches are able to solve a small portion of the clues. Even though the T5 model is trained on roughly half a million examples, it does not exceed the performance of the rule-based solver. For comparison, top-tier human experts are able to solve even very hard clues with almost 100% accuracy (Friedlander and Fine, 2016, 2018), though this expertise is acquired through significant training. Appendix A shows a selection of examples and the respective predictions of T5.

4.3 Analysis

Correlation with human perception of difficulty *Quick* cryptic crosswords is a subgenre of cryptic crosswords aimed at beginners, with clues designed to be easier to solve. Cryptonite’s test set contains 2,081 such clues. Examining the results of our main benchmark, Table 3 shows that both baselines perform better on quick clues, suggesting a correlation between human assessment of linguistic difficulty and the models’ performance on clues.⁵

⁵All quick clues are taken from *Times Quick Cryptic* (TQC). For a fair comparison to the quick clues, we consider a clue as non-quick only if it was published in The Times after March 10th 2014 (when TQC was introduced), and not as a part of TQC. Cryptonite’s test set contains 4,653 such non-quick clues.

	Quick Clues	Non-Quick Clues
T5-Large	12.83%	3.40%
Rule-based Solver	13.50%	5.78%

Table 3: Accuracy on quick cryptic clues vs non-quick cryptic clues. Both baselines perform better on clues that were deemed easier by human experts.

	With Enum.	Without Enum.
T5-Large	7.64%	4.90%
Rule-based Solver	8.58%	3.58%

Table 4: Comparison of baseline accuracy when enumeration is provided and when it is not provided.

The effect of enumeration The enumeration is the number (or numbers) in parentheses at the end of a clue indicating the number of letters in its answer, e.g. (7) or (5,4). To measure the informativeness of enumeration, we run our main experiment again, this time without providing the enumeration. Table 4 shows an accuracy drop in both baselines when the enumeration is not provided.⁶ While it is to be expected that enumeration helps the rule-based solver, we see that T5 is able to leverage this information as well.

Why do we split the data by answer? Many clues that share the same answer are paraphrases of each other (Appendix B). A neural model such as T5 might exploit this information and by copying answers from memorized training examples. Therefore, to test whether a model has learnt a general process for solving cryptic clues, we follow Lewis et al. (2020) and make Cryptonite’s default split the *answer split*, in which the answers of the train, validation, and test sets are mutually exclusive.

We compare the answer split with a *naive* (random) partition of the data. Table 5 shows that a naive split of Cryptonite will grossly overestimate the performance of T5; while the rule-based solver’s performance barely changes, T5-Large is now able to solve an additional 50% of the entire test set. Further analyzing the naive test set (Figure 3), we observe that the probability of T5 solving a clue is highly correlated with the number of times its answer appeared in the training set. This result indicates that that a significant part of the performance difference is due to the paraphrasing artifact,

⁶Cryptonite’s metadata contains additional information that could help a solver, such as *orientation* (whether the clue is across or down in the grid). Knowing the orientation can help in finding the clue’s wordplay-definition split.




Figure 3: The naive split exhibits a strong correlation between T5’s accuracy on clues from the test set (vertical axis) and the number of times their answer appears in the train set (horizontal). Each dot’s size represents the number of clues from the test set whose answer appears in the train set n times. Trend line is logarithmic.

Data Split	T5-Large	Rule-based Solver
Answer	7.64%	8.58%
Naive	56.16%	8.43%

Table 5: Comparing test set accuracy of our *answer* split (where the answers of test examples do not appear during training) and the *naive* random split.

and that ensuring unseen test answers is critical for establishing a true estimate of a model’s ability to solve cryptic clues.

5 Related Work

Cryptic crosswords Williams and Woodhead (1979) attempt to devise a formal language for describing cryptic clues. Hart and Davis (1992) define four stages of rule-based solving, and implement the second stage – “syntactic identification”. In our work we focus on creating a large-scale dataset of a cryptic clues and apply neural and rule-based methods to establish a strong baseline. Hardcastle (2001, 2007) focuses on rule-based approaches for *creating* cryptic clues given a word as an answer. Although in our work we test solving abilities, the reverse direction of creating a clue from an answer is also challenging, and the Cryptonite dataset could prove useful in this direction as well.

Language disambiguation In addition to works and datasets specifically targeting disambiguation on the word level (Levesque et al., 2011; Raganato et al., 2017; Sakaguchi et al., 2020), there are other domains strongly related to language disambiguation. Among them are pun disambiguation (Miller and Gurevych, 2015; Miller et al., 2017),

and sarcasm detection (Joshi et al., 2017; Oprea and Magdy, 2020). However, to the best of our knowledge Cryptonite is the first dataset both large in scale (unlike pun disambiguation), and containing a variety of wordplays (unlike sarcasm detection).

Non-cryptic crosswords As described in Section 2, non-cryptic (“regular”) crosswords are the common crosswords found in most newspapers. There are works introducing regular crossword datasets, some even containing a small percentage of more “tricky” clues⁷ (Littman et al., 2002). However, identifying this small portion of clues requires human effort, whereas Cryptonite is already guaranteed to consist entirely of cryptic clues. In addition, works on solving regular crosswords typically rely on an external database of clues (Ernandes et al., 2005; Barlacchi et al., 2014; Severyn et al., 2015). When given a clue as an input, these systems search the database for the most similar clues, in hope they share the answer with the input clue. In Cryptonite, the answers of the train, validation, and test sets are mutually exclusive (Section 3). In doing so, we hope to shift the focus of solving from memorization to reasoning, which is especially interesting in the setting of *cryptic* clues.

6 Conclusion

We present Cryptonite, a large-scale dataset based on cryptic crosswords, whose solving requires disambiguating a variety of wordplays. We saw that the standard approach of fine-tuning T5-Large on Cryptonite does not outperform an existing rule-based model, achieving 7.6% and 8.6% accuracy respectively, while human experts achieve close to 100% accuracy. These results highlight the challenge posed by Cryptonite, and will hopefully encourage further research on disambiguation tasks that are not easily solved by a native speaker.

References

- Gianni Barlacchi, Massimo Nicosia, and Alessandro Moschitti. 2014. [Learning to rank answer candidates for automatic resolution of crossword puzzles](#). In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 39–48, Ann Arbor, Michigan. Association for Computational Linguistics.
- Robin Deits. 2021. [rdeits/crypticcrosswords.jl: v0.1.1](#).
- ⁷Compare Florida fruit (6) → orange to where to get a date (4) → palm.
- Marco Ernandes, Giovanni Angelini, and Marco Gori. 2005. [Webcrow: A web-based system for crossword solving](#). In *AAAI*, pages 1412–1417.
- Kathryn J. Friedlander and Philip A. Fine. 2016. [The grounded expertise components approach in the novel area of cryptic crossword solving](#). *Frontiers in Psychology*, 7:567.
- Kathryn J. Friedlander and Philip A. Fine. 2018. [“the penny drops”: Investigating insight through the medium of cryptic crosswords](#). *Frontiers in Psychology*, 9:904.
- Mor Geva, Yoav Goldberg, and Jonathan Berant. 2019. [Are we modeling the task or the annotator? an investigation of annotator bias in natural language understanding datasets](#). 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 1161–1166, Hong Kong, China. Association for Computational Linguistics.
- 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.
- D Hardcastle. 2001. [Using the bnc to produce dialectic cryptic crossword clues](#). In *Corpus Linguistics 2001*, pages 256–265.
- David Hardcastle. 2007. *Riddle posed by computer (6): the computer generation of cryptic crossword clues*. Ph.D. thesis, Citeseer.
- M Hart and RH Davis. 1992. [Cryptic crossword clue interpreter](#). *Information and Software Technology*, 34(1):16 – 27.
- Aditya Joshi, Pushpak Bhattacharyya, and Mark J. Carman. 2017. [Automatic sarcasm detection: A survey](#). *ACM Comput. Surv.*, 50(5).
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- H. Levesque, E. Davis, and L. Morgenstern. 2011. [The winograd schema challenge](#). In *KR*.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2020. [Question and answer test-train overlap in open-domain question answering datasets](#).

- Michael L Littman, Greg A Keim, and Noam Shazeer. 2002. A probabilistic approach to solving crossword puzzles. *Artificial Intelligence*, 134(1-2):23–55.
- George A. Miller. 1995. [Wordnet: A lexical database for english](#). *Commun. ACM*, 38(11):39–41.
- Tristan Miller and Iryna Gurevych. 2015. [Automatic disambiguation of English puns](#). 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 719–729, Beijing, China. Association for Computational Linguistics.
- Tristan Miller, Christian Hempelmann, and Iryna Gurevych. 2017. [SemEval-2017 task 7: Detection and interpretation of English puns](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 58–68, Vancouver, Canada. Association for Computational Linguistics.
- Tim Moorey. 2018. *How to Crack Cryptic Crosswords*. Collins.
- Silviu Oprea and Walid Magdy. 2020. [iSarcasm: A dataset of intended sarcasm](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1279–1289, Online. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Alessandro Raganato, Jose Camacho-Collados, and Roberto Navigli. 2017. [Word sense disambiguation: A unified evaluation framework and empirical comparison](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 99–110, Valencia, Spain. Association for Computational Linguistics.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Winogrande: An adversarial winograd schema challenge at scale. In *AAAI*.
- Aliaksei Severyn, Massimo Nicosia, Gianni Barlacchi, and Alessandro Moschitti. 2015. [Distributional neural networks for automatic resolution of crossword puzzles](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 199–204, Beijing, China. Association for Computational Linguistics.
- Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive learning rates with sublinear memory cost](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4596–4604, Stockholmsmässan, Stockholm Sweden. PMLR.
- P. W. Williams and D. Woodhead. 1979. [Computer assisted analysis of cryptic crosswords](#). *The Computer Journal*, 22(1):67–70.
- 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.

A Example Predictions

Clue	Answer
Act like tragic heroine with cold extremity	mimic
Group of musicians prohibited on the radio	band
Assumed diamonds to be shelved	put on ice
Is in control of distant armies abroad	administrates
Second parasite	tick

Table 6: Examples of correct predictions of T5-Large.

Clue	Answer	Prediction
Suffer death at riverside	endure	strand
Travel free heading for eastbourne	ride	trip
Toast gordon!	brown	gobbi
Pair of braces	four	pair
Performing insect begged for food	eggs	beef
	benedict	wellington

Table 7: Examples incorrect predictions of T5-Large.

B Similar Clues

Clue	Answer
greek upper chamber greek for 'upper room'	attic
flat race maybe failing to finish flat race possibly unfinished	even
one playing minor part in run one with small part in film run?	extra
think to make changes in partner contemplate change in partner	meditate
beginning assault start of an attack	onset
what we learn by accepting established award what we learn by accepting fixed award	rosette
animal in forest, a grizzly animal in forest, a gazelle	stag

Table 8: Examples of pairs of similar clues from the naive split, one from the train set and one from the test set.