

DART: A Lightweight Quality-Suggestive Data-to-Text Annotation Tool

†Ernie Chang, †Jeriah Caplinger, *Alex Marin, †Xiaoyu Shen, †Vera Demberg

†Dept. of Language Science and Technology, Saarland University

{cychang, jeriahc, xiaoyu, vera}@coli.uni-saarland.de

*Microsoft Corporation, Redmond, WA

{alemari}@microsoft.com

Abstract

We present a lightweight annotation tool, the *Data AnnotatoR Tool (DART)*, for the general task of labeling structured data with textual descriptions. The tool is implemented as an interactive application that reduces human efforts in annotating large quantities of structured data, *e.g.* in the format of a table or tree structure. By using a backend sequence-to-sequence model, our system iteratively analyzes the annotated labels in order to better sample unlabeled data. In a simulation experiment performed on annotating large quantities of structured data, *DART* has been shown to reduce the total number of annotations needed with active learning and automatically suggesting relevant labels.

1 Introduction

Neural data-to-text generation has been the subject of much research in recent years (Gkatzia, 2016). Traditionally, the task takes as input structured data which comes in the form of tables with attribute and value pairs, and generates free-form, human-readable text.

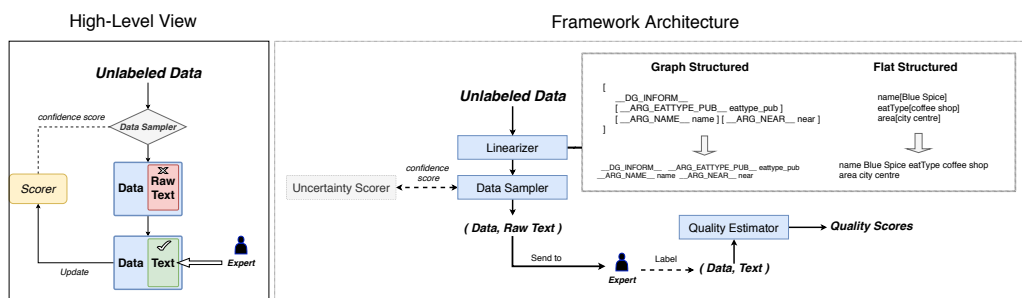


Figure 1: Left: Overview of the DART toolkit usage. Right: Diagram of the framework architecture.

Past example datasets include Restaurants (Wen et al., 2015) or graph-structure inputs (Balakrishnan et al., 2019). Analogously, most conversation systems (Williams et al., 2015; Crook et al., 2018) utilize intermediate meaning representation (data) as input to generate natural language sentences. In practice, however, these systems are highly reliant on the use of large-scale labeled data. Each new domain requires additional annotations to pair the new data with matching text. With the rise in development of natural language generation (NLG) systems from structured data, there is also an increased need for annotation tools that reduce labeling time and effort for constructing complex sentence labels. Unlike other labeling tasks, such as sequence tagging (Lin et al., 2019), where the labels are non-complex and correspond to fixed sets of classes, data-to-text generation entails providing complete sentence labels for each data instance. To construct textual description is time-consuming and therefore it can be beneficial for the system to automatically suggest texts and allow the annotators to accept or *partially correct* them. To this end, we propose to create an interactive annotation tool: Data AnnotatoR Tool (DART¹) that

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

¹Demo is available at <https://youtu.be/onPUgQ2ixpI>

System	Label Types	Programming Language	Has Label Recommendation ?	User Interface	Use Active Learning ?
DART	Text	Python	✓	GUI	✓
AlpacaTag (Lin et al., 2019)	Tag	Python	✓	GUI	✓
YEDDA (Yang et al., 2018)	Tag	Python	✓	GUI	✗

Table 1: A general comparison of relevant GUI-based annotation tools.

reduces structured data-to-text annotation efforts by incorporating *automatic label suggestion* and the *uncertainty-based active learning algorithm* (Lewis and Catlett, 1994; Culotta and McCallum, 2004). *DART* serves as a natural complement to downstream data-to-text systems, rather than an end-to-end NLG system. As such, it can assist in the development of both traditional rule-based systems (e.g. (Reiter, 2007)), and the recent neural systems (e.g. (Balakrishnan et al., 2019; Chang et al., 2020; Shen et al., 2020; Hong et al., 2019)).

As a lightweight, standalone desktop application, *DART* can be easily distributed to domain experts and installed on local devices. *DART* consists of a user-friendly interface that allows experts to iteratively improve the overall corpus quality with partial corrections. Overall, the toolkit provides three advantages: (1) It reduces labeling difficulty by automatically providing natural language label recommendations; (2) It efficiently solicits data for which it has low confidence (or high uncertainty) in its generated text to be annotated, so that overall annotation efforts can be reduced; (3) Lastly, it provides real-time in-progress updates with statistics about the labeled corpus as to help direct the overall annotation process. This is achieved with a myriad of quality estimators that assess corpus diversity and the overall text quality.

2 Annotation Framework

DART is a desktop application built with PyQt5². It is compiled into a single executable with PyInstaller³, a tool that supports both Mac OS and Windows environments. It contains an intuitive interface as described in section 3. Annotation experts interact with *DART* in the following way: (1) A file containing unlabeled data is uploaded. (2) The system samples some data instances from the file, with a selection strategy based on signals from the sequence-to-sequence *uncertainty scorer* (section 2.1) and performed with the *data sampler* (section 2.2). (3) Experts then annotate the provided data by correcting the suggested labels (available after the first iteration of (1)-(2)). (4) During the process of annotation, the labeled corpus quality is indicated by the annotation *quality estimators* (section 2.3) for experts to determine if the process were to be terminated. We discuss each component in more detail below.

2.1 Uncertainty Scorer

We represent the structured unannotated corpus as $D = d_{i=1}^N$ where each data sample d_i comprises of a token sequence linearized from underlying structured data samples x_i , as motivated by past works in the multilingual surface realization tasks (Mille et al., 2018). We employ the Transformer-based (Vaswani et al., 2017) encoder-decoder architecture as the sequence-to-sequence model. The sequences d_i are fed into the model in order to generate a text sequence $t_i = w_1, w_2, \dots, w_{M_i}$ of length M_i .

Since the model is given only the input data d , we compute reconstruction scores for this data, and use the cross-entropy loss as the *uncertainty score*. To do so, we perform round-trip training⁴ where the source data is reconstructed to achieve *cycle consistency*. In this setup, the same encoder and decoder are used in both forward and backward training *i.e.* a forward model $M_{forward}$ goes from data to text and the backward model $M_{backward}$ converts the text back into data. We define the round-trip training log-loss as the uncertainty score $S_{uncertainty}$ (as Eq. 1a) where $\mathcal{L}(\cdot)$ is the cross-entropy loss. d' is the

²<https://riverbankcomputing.com/software/pyqt/intro>

³<https://www.pyinstaller.org/>

⁴Defined in (Lample et al., 2017) as the back-translation technique where input data is reconstructed with $M_{forward}$ and $M_{backward}$ to compute consistency loss.

generated data using $M_{forward}$ and $M_{backward}$ given input data d .

$$\mathcal{S}_{\text{uncertainty}} = \mathbb{E}_{d \sim p_{\text{data}}(D)} \mathcal{L}(d', d) \quad (1a)$$

Next, we discuss how $\mathcal{S}_{\text{uncertainty}}$ is used in *uncertainty sampling* during data selection.

2.2 Data Sampler

The process of data selection identifies N data instances to be labeled such that the overall generation quality is improved. This can be achieved by learning the structure over data D (Tosh and Dasgupta, 2018). We adopt a simple technique to represent each data instance d as a bag-of-word (BOW) vector and further divide each *attribute-type* (first layer) into k clusters (*sub-type*, second layer) with the K-means algorithm⁵ (Alsabti et al., 1997), which splits data instances into k clusters based on selected centroids.

We first rank the order of batch-size samples within each *sub-type* using *uncertainty scores* (in section 2.1) so that experts can annotate the ones with the least confident scores first. At sampling time, we obtain unlabeled data instances from all *sub-types* across all *attribute-types* iteratively. One sample is obtained from each *sub-type* before moving on to the next *sub-type*.

For data presented to be labeled, the system also suggests labels in order to reduce annotation efforts. *DART* employs a simple retrieval-based technique to obtain a text label t for each data instance d . Using the BOW representation of d , we simply find the most similar d' (cosine similarity) in the labeled pool of (d_i, t_i) pairs, and use its text label t_i as the suggestion⁶. The sampling process continues until either all data instances are labeled or a satisfactory threshold value is reached for the quality metric on the labeled corpus (as defined in section 2.3).

2.3 Quality Estimator

To better manage the annotation process, we include the diversity metrics used in (Balakrishnan et al., 2019): *number of unique tokens*, *number of unique trigrams*, *Shannon token entropy*, *conditional bigram entropy*. Following (Novikova et al., 2016), we also measure various types of lexical richness including *type-token ratio (TTR)* and *Mean Segmental TTR (MSTTR)* (Lu, 2012), where higher values of TTR and MSTTR correspond to more diverse corpus. *DART* displays these scores on the *Status Display* as shown in Figure 2. These scores serve as on-the-fly quality estimates that help experts decide when sufficient labels have been collected.

3 User Interface

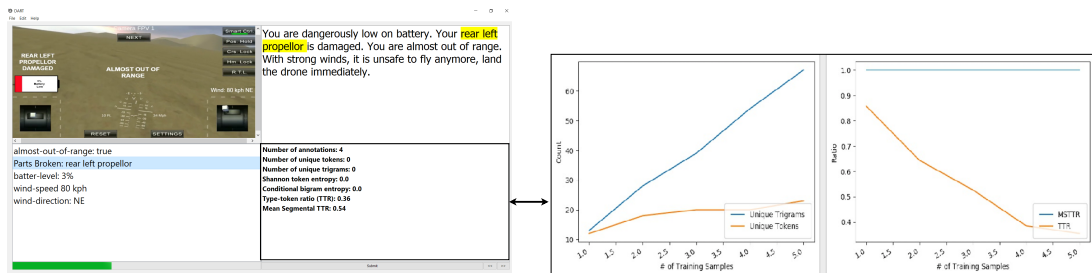


Figure 2: Screenshot of sample annotation interface showing the provided data, and the text box for annotation. On the bottom right, user can select either the statistics or plots indicating the corpus diversity and annotation count.

The interactive interface divides the interactive application window into a few compartments: *DART* includes a configuration editor interface that allows the experts to modify the delimiter (e.g. “;”) between *attribute:value* pairs. For graph structured input, the delimiter (e.g. “_”) is used to identify the attribute

⁵Using the implemented version from <https://scikit-learn.org>

⁶There are no suggestions for the initial batch of annotations.




Figure 3: Performance comparison between DART’s **data sampler**, **random sampling**, and retrieving labels from the full dataset (**ALL**) on E2E (Left) and the Weather (Right) datasets (42k data instances for E2E and 32k for weather), using the same retrieval method.

tags instead. Note that the system supports three granularity of tokenization: (1) word, (2) character, and (3) byte-pair encoding (BPE) (Sennrich et al., 2016). The top half of Figure 2 shows the main annotation page where experts can input constructed sentences into the text boxes based on suggested texts and the provided image (or short clip) ⁷. As the expert annotates, the progress bar below the text box indicates when the background *uncertainty scorer* training session will begin. The bottom half of Figure 2 shows the annotation progress statistics, including the percentage of data types that have been annotated and the quality of overall templates. When a specified number of annotations has been created, experts can download both the annotated data samples along with data with predicted labels. In general, a high-quality corpus maintain a high corpus diversity (e.g. a MSTTR score of 0.75 or TTR of 0.01 in the E2E dataset (Novikova et al., 2016)) even as the number of annotation increases.

4 Experiments

Data. We use two different types of structured data: (A) Attribute-value pairs as used in the crowd-sourced *E2E* dataset (Novikova et al., 2017), and (B) the graph-structured data as defined in (Balakrishnan et al., 2019) on the *weather* domain. To simulate the annotation process, we employ the given training, development and test sets of each datasets for annotation tool evaluation, with the test set kept fixed. This amounts to roughly 42k samples for E2E and 32k for the weather training sets.

Simulation Study. To evaluate the effectiveness of *DART*, we perform a simulated experiment for each of our two datasets, E2E and Weather. We simulate the labeling process using both the retrieval-based method (**Sampler**, as discussed in section 2.2) and the baseline approach using random selection of data (**Random**) and compare the performance of the two methods relative to using the full dataset (**All**).

Results for the two datasets, E2E and Weather, are presented in Figure 3. On both datasets, the data sampler allows the retriever to obtain the same performance (i.e. with similar BLEU score as *ALL*) using only 10k labeled data instances, which is significantly less than that of the original dataset (i.e. 42k for *E2E* and 32k for *Weather*). As such, the number of required annotations to arrive at the same performance using all labeled data is significantly reduced for both datasets, to one-fifth of the original dataset size. In contrast, performance obtained using the *Random* selection is significantly worse (on the order of 6 to 10 BLEU points lower than the baseline) compared to using the *Sampler* selection while matching the number of training samples.

5 Conclusions

While a wide range of annotation tools for NLP tasks exists, most of these tools are targeted at non-textual labels. *DART* is designed to enable the ease of annotation where the labels are textual descriptions and the inputs are structured data. This is the initial version of the tool, and we hope to extend it to include a web-based version and to expand its functionality in the following ways: (1) support different types of encoders, and (2) improve upon the data sampling process.

⁷The use of pictures is shown to elicit better data (Novikova et al., 2016)

Acknowledgements

This research was funded in part by the German Research Foundation (DFG) as part of SFB 248 “Foundations of Perspicuous Software Systems”. We sincerely thank the anonymous reviewers for their insightful comments that helped us to improve this paper.

References

- Khaled Alsabti, Sanjay Ranka, and Vineet Singh. 1997. An efficient k-means clustering algorithm.
- Anusha Balakrishnan, Jinfeng Rao, Kartikeya Upasani, Michael White, and Rajen Subba. 2019. Constrained decoding for neural NLG from compositional representations in task-oriented dialogue. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 831–844, Florence, Italy, July. Association for Computational Linguistics.
- Ernie Chang, David Ifeoluwa Adelani, Xiaoyu Shen, and Vera Demberg. 2020. Unsupervised pidgin text generation by pivoting english data and self-training. *arXiv preprint arXiv:2003.08272*.
- Paul A Crook, Alex Marin, Vipul Agarwal, Samantha Anderson, Ohyoung Jang, Aliasgar Lanewala, Karthik Tangirala, and Imed Zitouni. 2018. Conversational semantic search: Looking beyond web search, q&a and dialog systems. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 763–766. ACM.
- Aron Culotta and Andrew McCallum. 2004. Confidence estimation for information extraction. In *Proceedings of HLT-NAACL 2004: Short Papers*, pages 109–112. Association for Computational Linguistics.
- Dimitra Gkatzia. 2016. Content selection in data-to-text systems: A survey. *arXiv preprint arXiv:1610.08375*.
- Xudong Hong, Ernie Chang, and Vera Demberg. 2019. Improving language generation from feature-rich tree-structured data with relational graph convolutional encoders. In *Proceedings of the 2nd Workshop on Multilingual Surface Realisation (MSR 2019)*, pages 75–80.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2017. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*.
- David D Lewis and Jason Catlett. 1994. Heterogeneous uncertainty sampling for supervised learning. In *Machine learning proceedings 1994*, pages 148–156. Elsevier.
- Bill Yuchen Lin, Dong-Ho Lee, Frank F. Xu, Ouyu Lan, and Xiang Ren. 2019. AlpacaTag: An active learning-based crowd annotation framework for sequence tagging. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 58–63, Florence, Italy, July. Association for Computational Linguistics.
- Xiaofei Lu. 2012. The relationship of lexical richness to the quality of esl learners’ oral narratives. *The Modern Language Journal*, 96(2):190–208.
- Simon Mille, Anja Belz, Bernd Bohnet, Emily Pitler, and Leo Wanner, editors. 2018. *Proceedings of the First Workshop on Multilingual Surface Realisation*, Melbourne, Australia, July. Association for Computational Linguistics.
- Jekaterina Novikova, Oliver Lemon, and Verena Rieser. 2016. Crowd-sourcing nlg data: Pictures elicit better data. In *Proceedings of the 9th International Natural Language Generation conference*, pages 265–273.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2017. The e2e dataset: New challenges for end-to-end generation. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 201–206.
- Ehud Reiter. 2007. An architecture for data-to-text systems. In *Proceedings of the Eleventh European Workshop on Natural Language Generation*, pages 97–104. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725.
- Xiaoyu Shen, Ernie Chang, Hui Su, Jie Zhou, and Dietrich Klakow. 2020. Neural data-to-text generation via jointly learning the segmentation and correspondence. *arXiv preprint arXiv:2005.01096*.

- Christopher Tosh and Sanjoy Dasgupta. 2018. Interactive structure learning with structural query-by-committee. In *Advances in Neural Information Processing Systems*, pages 1121–1131.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721, Lisbon, Portugal, September. Association for Computational Linguistics.
- Jason D. Williams, Eslam Kamal, Mokhtar Ashour, Hani Amr, Jessica Miller, and Geoff Zweig. 2015. Fast and easy language understanding for dialog systems with Microsoft language understanding intelligent service (LUIS). In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 159–161, Prague, Czech Republic, September. Association for Computational Linguistics.
- Jie Yang, Yue Zhang, Linwei Li, and Xingxuan Li. 2018. Yedda: A lightweight collaborative text span annotation tool. In *Proceedings of ACL 2018, System Demonstrations*, pages 31–36.