

Targeting the Benchmark: On Methodology in Current Natural Language Processing Research

David Schlangen

CoLabPotsdam / Computational Linguistics
Department of Linguistics, University of Potsdam, Germany
david.schlangen@uni-potsdam.de

Abstract

It has become a common pattern in our field: One group introduces a *language task*, exemplified by a *dataset*, which they argue is *challenging* enough to serve as a *benchmark*. They also provide a baseline *model* for it, which then soon is *improved* upon by other groups. Often, research efforts then move on, and the pattern repeats itself. What is typically left implicit is the argumentation for why this constitutes progress, and progress towards what. In this paper, I try to step back for a moment from this pattern and work out possible argumentations and their parts.

1 Introduction

The goal of any field of research is to make progress towards answering its foundational questions. To do so, a *methodology* is required that guides attempts at providing or improving answer proposals. In natural language processing, the object of study is human language, and any methodology for doing research in this field will need to have some contact with examples of this object. This contact has become more and more direct in the past decades, with samples of language becoming more directly the material from which proposals (in the form of statistical *models*) are derived. Recent years have seen an increase in the collection of samples specifically for the purpose of creating *benchmarks*, against which progress in devising models can be measured. It is this function of *benchmarking*, and its role in a progress-oriented methodology, that this paper aims to investigate.

Figure 1 illustrates the basic structure of a benchmarking methodology: A *language task* is devised that is a) restricted enough to be manageable with current methods, and b) deemed challenging for the *capabilities* that it involves.¹ For this task, a *dataset*

¹This figure is from (Schlangen, 2019), of which this is a shorter version developed in a somewhat different direction.

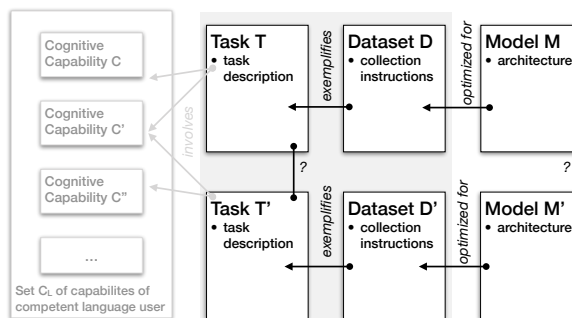


Figure 1: Relations between Research Objects in a Benchmark-Driven Methodology

is collected, often via crowd sourcing, on which in turn models are trained and compared, using evaluation metrics defined together with the task. What can we learn by following such a methodology? Let's look at the components first and then at ways in which this methodology is, might, and perhaps should be used.

2 What is a Language Benchmark?

2.1 What is a Benchmark?

In computing, a benchmark is “*a problem that has been designed to evaluate the performance of a system [which] is subjected to a known workload and the performance of the system against this workload is measured. Typically the purpose is to compare the measured performance with that of other systems that have been subject to the same benchmark test.*” (Butterfield et al., 2016).

The use of this term in NLP is related: here, benchmark tasks are also specifically designed for evaluation; however, an important difference is that what is being evaluated is not a full system that has a separate main purpose, but rather an *algorithm* that is instantiated on the benchmark itself. I will discuss the consequences of this below.

This kind of evaluation of learning algorithm has a long tradition in the field of machine learning research.² In this field, a new algorithm would

²For example, the UCI Machine Learning Repository has

normally be tested on a large collection of datasets, possibly ranging from classifications of flowers to classification of credit records, with no assumption of any internal connection between the datasets. Again, NLP is different here, as all datasets represent facets of the same underlying phenomenon, language use.

I will argue that these two differences (life outside of benchmarking, and internal connection between tasks) are important, but understudied. But first we look at the notion of a *language task* in some more detail.

2.2 What is a Language Task?

A *language task* is a mapping between an *input space* and an *output* or *action space*, at least one of which contains natural language expressions. The mapping has to conform to a *task description*, which is typically given only informally, making reference to theoretical or pre-theoretical constructs external to the definition, such as “translation” or “is true of”. I call this an *intensional description*. Typically, a task will also be specified *extensionally* through the provision of a *dataset* of examples of the mapping (that is, pairs of state and action). To collect such a dataset, the task description (e.g., “classification of entailment relations between sentence pairs”) must be operationalised into a collection instruction (“please mark whether the situation that is well described by sentence A could normally also be described by sentence B”).

3 How Can It be Evaluated?

3.1 Relation Task / Dataset

Given a task and a dataset, the first question to ask is how well the latter exemplifies the former. Investigating this is relatively straightforward. First, the dataset should be *verified*, which is to check whether the provided input/output pairs can indeed be judged correct relative to the task (in its intensional description). If the examples are collected specifically for the purpose of exemplifying the task, this is the process of controlling annotation, and standard methodologies exist (Artstein and Poesio, 2008; Pustejovsky and Stubbs, 2013). Care needs to be taken that the task is actually well-defined enough to pose an unambiguous challenge to capable language users.³

been collecting and providing datasets for more than 20 years now (Dua and Graff, 2019).

³Pavlick and Kwiatkowski (2019), for example, show that the task of annotation textual entailments can lead to faultless

Validating a dataset is a less formalised process. It comprises arguing that the dataset indeed exemplifies the task intension well. For example, pairs only of images of giraffes and sentences describing them would arguably not exemplify the general task of *image description* very well (even if the descriptions are accurate), while perhaps exemplifying the task of *giraffe image description*.

Another way to evaluate a dataset is by trying to model it. If a model can “solve” the dataset even when deprived of information that for theoretical or pre-theoretical reasons is seen to be crucial, the dataset can be considered an unsatisfactory exemplification of the task. E.g., in a *visual (polar) question answering* setting (Antol et al., 2015), if in a dataset all and only the expressions that mention giraffes are true, a model could seize on this fact and perform well without needing the images, which would be evidence that the dataset is deficient relative to the task description.⁴

3.2 Relation Cognitive Capability / Task

While the dataset forms the visible surface of the task, it is the task itself that needs to provide value. We can categorise tasks by how they are embedded in further uses: a *product task* is one that can be argued to have direct value to consumers (such as translation, or search); an *annotation task* is one where the task description is theoretically motivated and the output a linguistically motivated object (which may be consumed in a pipeline that itself is motivated as a product task); finally, a *benchmark task* – which is the type that concerns us here – is one which gets its value from how well it tests a particular ability (and nothing else) and how well it discriminates learners based on this ability.⁵

For a language benchmark task, the argument roughly goes as follows (even if typically only made implicitly): To be good at task *T*, an agent

annotator disagreements.

⁴The task of visual question answering provides an interesting example case of such a development. After Antol et al. (2015) introduced the first large scale dataset for this task, it quickly became clear that this dataset could be handled competitively by models that were deprived of visual input (“language bias”, as noted e.g. by Jabri et al., 2016). This problem was then addressed by Goyal et al. (2017) with the construction of a less biased (and hence more valid) corpus for the same task.

⁵Martinez-Plumed and Hernandez-Orallo (2018), analysing AI benchmarks in general, distinguish between *difficulty* (which determines the ability level which must be reached to perform better than chance on a task) and *discrimination* (the slope of the graph plotting ability level vs. probability of correct response).

must possess a set C_T of capabilities (of representational or computational nature). If the $c \in C_T$ are capabilities that competent language users can be shown or argued to possess and make use of in using language—let’s call the set of these capabilities of a competent language user C_L , so that $C_T \subseteq C_L$ —then being able to model these capabilities (via modelling the task) results in progress towards the ultimate goal, which is to model competent language use. And hence, any task T that comes with an *interesting* set C_T is a good task.⁶

Under what conditions does this argument work? First of all, the assumed connection to the set of capabilities must indeed be there. We have already seen a way to challenge a claimed connection, through providing a model that can “solve” a given task (via a dataset) while not having access to information that, given our analysis of the task and interest in C_T , should be involved in the capability.⁷ (Although this challenge in the first instance only targets the dataset and not the task itself.)

Secondly, following usual scientific methodology (Popper, 1934), we can rank the value of an instantiation of this argument by how precisely the capability is specified, from the trivially correct “task T involves the capability to do task T ” to a statement that could be wrong, e.g. “task T involves the capability to compute the syntactic structure of a natural language sentence”. Such a statement must make reference to theoretical constructs belonging to the analysis of cognitive capabilities.

Furthermore, we can rank the motivation given for a task by how explicit it is in delineating the set of capabilities it involves. For a given $c \in C_T$, is “ c as required by T ” fully *separable* from any

⁶To give some examples of informal versions of this argument, and choosing papers more or less randomly, here are some quotes:

From the paper that introduced the *visual question answering* task (Antol et al., 2015): “What makes for a compelling AI-complete task? [...] Open-ended questions require a potentially vast set of AI capabilities to answer – fine-grained recognition [...], object detection [...], activity recognition [...], knowledge based reasoning [...], and commonsense reasoning [...]”

Williams et al. (2018), on computing entailments: “The task of natural language inference (NLI) is well positioned to serve as a benchmark task for research on NLU. [...] In particular, a model must handle phenomena like lexical entailment, quantification, coreference, tense, belief, modality, and lexical and syntactic ambiguity.”

⁷Such an attack challenges the claim of there being a *necessary* connection between handling T and possessing capability c . It might still very well be that humans can only perform this task if they possess capability c (and all the knowledge involved in it), because they wouldn’t be able to pick up the statistical correlations that could be exploited.

other tasks involving c ? Or is “ c as required by T ” perhaps all that there is to know about c , that is, is c *exhaustively* represented by T ?

Finally, underlying the benchmarking methodology — where the benchmark is not just a measuring tool, but also a modelling target — there has to be the assumption that some sort of transferable knowledge is generated by modelling T , so that what the model (and not just the modeller!) has learned about (a sufficiently generally specified) c can be used in other tasks that involve c . (Let’s call this *transferability*; which strictly speaking is a property of models, not of tasks.) More on this below.

To sum up, a benchmark task must point beyond itself and get its value from its connection to a particular facet of language, a particular capability of language users; this in turn seems to be difficult to specify without access to terms from theories of the domain, which allow us to name these capabilities.^{8,9}

4 How are Language Benchmarks Used?

In the way that these tasks are set up, as single-step tasks that humans can quickly do (“describe this image”, “is the elephant [in this image] sleeping?”, “does sentence A follow from sentence B?”), it is tempting to see a similarity to tasks used in (human) intelligence testing (see e.g. Borsboom (2005) for an introduction). There is a crucial difference, however: Where intelligence testing works more in the way standard computing benchmarking works (subjecting the *otherwise functioning* learner to a standardised workload), in NLP, benchmarks are both the testing instrument as well as the training material.¹⁰ The question then cannot be “to what extent does system Σ possess capability c ”, it has to be “to what extent can algorithm A learn c from dataset D ?” — and what does that tell us?

4.1 Single-Task Models

Let’s assume we have defined a task T that we are sufficiently convinced is well represented by

⁸And one will indeed find that papers introducing such tasks make mention of terms like *syntax*, *semantics*, *compositionality*, *quantifiers*, etc.

⁹We can also note that with this focus on benchmarking normally comes a certain top-down approach, where the collected data is not investigated for how exactly the human participants went about solving their task. (But see (van Miltenburg, 2019) for a detailed study along those lines, for the task of image description.)

¹⁰For a recent paper also discussing the relation between AI benchmarking and intelligence testing, see (Chollet, 2019).

dataset D . We have trained a model M that performs well on this dataset. What have we learned? We have learned that a learning algorithm of the type of M can model D . Further, we have learned that the information to do task T (as exemplified in D), is contained in D , and M can pick it up.

Under what conditions can we now say that we have modelled T , rather than just D ? If we have convinced ourselves that D represents T faithfully, then we might be willing to make this leap, and with it, claim that we have modelled C_T . We can get further support by collecting more data D' that also exemplifies T , but perhaps operationalises it differently. The prediction should be at least that the learning algorithm can also learn to model D' ; but more significantly, we'd also want the model M learned from D to perform well on D' . Similarly, if we have another task T' of which we think that it involves similar capabilities, we should expect it to be amenable to being modelled with a learning algorithm of similar type to M .

What do we learn from a model M' (introducing architectural innovation κ over M) performing better on T (via D)? We can take this as indication that κ may be what is responsible for increasing performance, and hence what is leading to a more adequate model of C_T .

4.2 Multi-Task Models

With the advent of pre-training in NLP (Peters et al., 2018; Devlin et al., 2018), where a model is trained on (a typically large amount of) data under a specific task-regime (typically language modelling, i.e. the task of predicting the next word in a running text) and then becomes part of the model for a target task, it has become common to test on a collection of tasks (Wang et al., 2019b,a). What do we learn from such a setup? In our Figure 1, if we find a task on which we can pre-train a model M_P that becomes a part of models M and M' , and which makes them more powerful than models that do not have access to the pre-trained model, then we can infer that whatever M_P models is a shared part of M and M' as well (and hence *involves* the hypothesised joint capability C'). This then provides an instrument to study the tasks: if the pre-trained model works well on some but not all, there must be something that those groups have in common. To make this intelligible, however, recourse to theoretical terms must again be taken. (E.g., assuming that these tasks involve the use of

certain types of representation, or certain actions over representations.)

5 But Are We Making Progress?

Within the logic of this methodology, we are clearly making enormous progress at two links in the chain illustrated in Figure 1: For many of the established tasks, models have been and continue to be proposed that perform better, according to the metrics defined for the tasks. In addition, for many of the tasks, better datasets have been collected, avoiding exploitable biases. Where there is less activity is in systematically studying the implications of success at one task for success at others. The presentation above was largely idealised (or normative): In reality, there is very little explicitness about the assumed connection between tasks and capabilities, and no theory of how (or whether) language competence decomposes into capabilities that could be learned separately and then be assembled into a whole, and there is very little explicit knowledge about the vertical links in the Figure, from one task / model to the next.¹¹

6 Conclusions

In this short paper, I have discussed the methodology of using *language tasks* to drive research on models of language competence. I have argued that the success of this approach hinges on how well progress on one task can be translated into progress on other tasks. While some steps have been taken in this direction, current work still appears to mostly focus on isolated tasks (or groups of tasks). Overcoming this, in my opinion, will require more explicit considerations about how tasks and capabilities are connected, and how the set of capabilities is structured—to ensure movement is not only uphill, but rather up the *right* hill (Bender and Koller, 2020), and it indeed is a single hill. For this, a (re-)connection with the fields that study the composition of language competence—linguistics and cognitive and developmental psychology—seems advisable (if only to disagree *explicitly*). As a positive proposal, I suggest that a focus should be put on assembling a *curriculum* of tasks, organised in a complexity and inclusion hierarchy, and that the benchmarking target should be the developmental trajectory on this. Working this out in detail I must leave for future work.

¹¹In the neighbouring field of Computer Vision, there recently have been attempts to “disentangle task transfer learning” (Zamir et al., 2018).

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *International Conference on Computer Vision (ICCV)*.
- Ron Artstein and Massimo Poesio. 2008. *Inter-Coder Agreement for Computational Linguistics*. *Computational Linguistic*, 34(4):555–596.
- Emily M Bender and Alexander Koller. 2020. Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2, pages 5185–5198.
- Denny Borsboom. 2005. *Measuring the Mind: Conceptual Issues in Contemporary Psychometrics*. Cambridge University Press, Cambridge, UK.
- Andrew Butterfield, Gerard Ekembe Ngondi, and Anne Kerr, editors. 2016. *A Dictionary of Computer Science*, 7th edition. Oxford University Press, Oxford, UK.
- François Chollet. 2019. *On the Measure of Intelligence*. *arXiv e-prints*, page arXiv:1911.01547.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.
- Dheeru Dua and Casey Graff. 2019. *UCI machine learning repository*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. In *CVPR 2017*.
- Allan Jabri, Armand Joulin, and Laurens van der Maaten. 2016. Revisiting Visual Question Answering Baselines. In *European Conference on Computer Vision (ECCV)*.
- Fernando Martinez-Plumed and José Hernandez-Orallo. 2018. Dual indicators to analyse ai benchmarks: Difficulty, discrimination, ability and generality. *IEEE Transactions on Games*, pages 1–1.
- Emiel van Miltenburg. 2019. *Pragmatic factors in (automatic) image description*. Ph.D. thesis, Vrije Universiteit Amsterdam.
- Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent Disagreements in Human Textual Inferences. *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Karl Popper. 1934. *Logik der Forschung*. Mohr Siebeck, Tübingen, Germany.
- James Pustejovsky and Amber Stubbs. 2013. *Natural Language Annotation for Machine Learning*. O’Reilly, Sebastopol, CA, USA.
- David Schlangen. 2019. Language tasks and language games: On methodology in current natural language processing research. *CoRR*, abs/1908.10747.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems. In *NeurIPS*, July, pages 1–30.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019b. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *ICLR 2019*, pages 1–20.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Amir Zamir, Alexander Sax, William Shen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. 2018. Taskonomy: Disentangling Task Transfer Learning. In *CVPR 2018*.