

Made of Steel?

Learning Plausible Materials for Components in the Vehicle Repair Domain

Annerose Eichel, Helena Schlipf, Sabine Schulte im Walde
Institute for Natural Language Processing, University of Stuttgart

{annerose.eichel,schulte}@ims.uni-stuttgart.de, helena.schlipf@gmail.com

Abstract

We propose a novel approach to learn domain-specific plausible materials for components in the vehicle repair domain by probing Pre-trained Language Models (PLMs) in a cloze task style setting to overcome the lack of annotated datasets. We devise a new method to aggregate salient predictions from a set of cloze query templates and show that domain-adaptation using either a small, high-quality or a customized Wikipedia corpus boosts performance. When exploring resource-lean alternatives, we find a distilled PLM clearly outperforming a classic pattern-based algorithm. Further, given that 98% of our domain-specific components are multiword expressions, we successfully exploit the compositionality assumption as a way to address data sparsity.

1 Introduction

Connecting a symptom to an underlying cause is a crucial building block for natural language understanding across domains. For example, as illustrated in Fig. 1, a standard approach that human mechanics in vehicle repair shops apply when tracing the cause of a symptom, is to exploit the link between a vehicle *component* and the component’s *materials*. This study tackles the task of automatically learning plausible materials for vehicle components. The information is crucial in vehicle repair shops, as mechanics are faced with constantly growing vehicle complexity, making it hard to manually identify the cause of a malfunction. While domain-specific information on vehicle components is often available, plausible domain-specific material information typically needs to be gathered from external data sources. For example, from “*Brake disk are usually manufactured from gray cast iron*”, one may retrieve the information that *brake disks* consist of *iron* (and possibly further materials not mentioned here).

We propose a novel approach to leverage state-of-the-art pretrained language models (PLMs) for

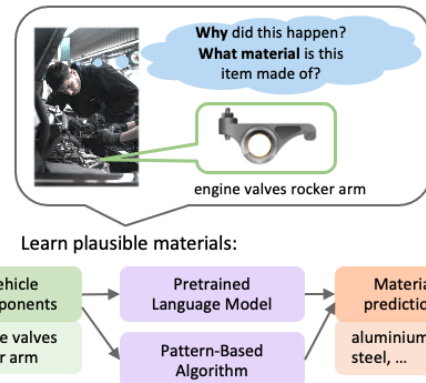


Figure 1: Conceptual overview for learning domain-specific plausible materials (e.g., *aluminium*) for vehicle components (e.g., *engine valves rocker arm*).

the task of learning plausible materials in highly domain-specific contexts, such as vehicle components in the vehicle repair domain. Our approach can also be applied to similarly specific domains, e.g., to learn which materials or substances are plausibly used in products in the health domain, textile industry, etc. As we tackle a task from material science, we make a contribution to a field that is not only very challenging, but heavily under-investigated from an NLP perspective. We focus on the PLM RoBERTa (Liu et al., 2019) and develop an approach that successfully learns plausible materials in English for components in the vehicle repair domain. We further exploit domain-adaptation variants and show (i) that harnessing a small high-quality domain-specific corpus boosts the performance over an out-of-the-box vanilla RoBERTa, and (ii) how similar performance can be reached by using a widely accessible data source such as Wikipedia if no domain-specific data is available.

To overcome the typical lack of annotated datasets for training and fine-tuning models in highly domain-specific tasks, we probe RoBERTa in a cloze task setting (Taylor, 1953; Petroni et al., 2019), compare variants for aggregating the results, and analyze the influence of a varied set of query templates on model predictions.

While PLMs come with great advantages, their hunger for storage, data, and computing power is even greater. We thus compare the performance of the vanilla and domain-adapted versions of RoBERTa with (i) a distilled version, i.e., the much smaller DistilRoBERTa (Sanh et al., 2019), and (ii) the seminal pattern-based bootstrapping algorithm Basilisk (Thelen and Riloff, 2002).

Finally, we tackle the challenge of handling a dataset with a vast majority of multiword expressions, i.e., 98% of our targets are noun compounds (such as *seat heating switch led*). We address the corresponding severe data sparsity by assuming that many domain-specific compounds are compositional (Clouet and Daille, 2014; Hätyy et al., 2021) such that our domain-specific model may fall back to information regarding the compound’s head (in the example above: the right-most simplex noun in the compound *led*), and thus improve on the data sparsity.

2 Related Work

Domain-specific IR In contrast to domains such as biomedicine and chemistry, “[I]everaging NLP tools in materials science remains in its infancy” (Olivetti et al., 2020, p.4), with main challenges including the development of task- and domain-specific tools to structure and harness knowledge for material synthesis and material discovery. Common approaches to elicit material mentions from text include unsupervised approaches such as topic modeling (Rani and Kumar, 2021) and LDA (Venugopal et al., 2021), as well as supervised methods for domain-specific NER, often focusing on a single material or material group (Mysore et al., 2017, 2019; Friedrich et al., 2020; Gupta et al., 2022; Nayak and Timmapathini, 2021; O’Gorman et al., 2021). Differently to previous work, we probe PLMs in a cloze-task style setting and compare the results to a similarly unsupervised pattern-based bootstrapping algorithm. While most unsupervised methods are not easily interpretable, our query templates, seed words, patterns and predicted materials are directly accessible.

Prompting and Cloze Query Engineering Recent work has leveraged cloze-task style settings to probe the knowledge that PLMs acquire during pretraining, thus targeting linguistic capabilities (Goldberg, 2019; Ettinger, 2020; Apidianaki and Garí Soler, 2021; Rogers et al., 2021), the understanding of rare words (Schick and Schütze,

2020) and conceptual abstractions (Ravichander et al., 2020), as well as factual and commonsense knowledge (Petroni et al., 2019; Jiang et al., 2021). An emerging strand of research highlights PLM sensitivity to the input in a cloze-task style setting. For example, Elazar et al. (2021) demonstrate that PLMs show deficiencies in prediction consistency when presented cloze-style query paraphrases. Others exploit this flaw by explicitly considering paraphrases of a cloze query in addressing a model’s sensitivity to a specific input query to elicit a desired output (Davison et al., 2019; Jiang et al., 2021). Pandia and Ettinger (2021) show that models lack robustness in their ability to harness relevant context information in the face of cloze tasks containing distracting contextual cues. In reverse, priming the model using *trigger tokens* (Shin et al., 2020) and *lexical cues* (Misra et al., 2020) might aid in guiding the model to predict the desired output. Specifically, Shin et al. (2020) propose AUTOPROMPT to develop automatically constructed prompts or patterns to elicit knowledge from pretrained PLMs for a variety of tasks. In contrast to previous work, we semi-automatically construct a variety of prompts to cover a salient set of paraphrases for eliciting plausible materials in a PLM; this allows for meaningful comparison and interpretation of the predictions’ quality and the influence of the cloze queries.

Semantic Plausibility While classical distributional models tend to model selectional preferences and thematic fit instead of capturing semantic plausibility (Erk et al., 2010), we have recently seen advances to model plausibility across various dimensions, including *physical* and *abstract* semantic plausibility (Wang et al., 2018; Porada et al., 2019). SOTA models for event plausibility however still rely on straightforward conditional probabilities of co-occurrences as estimated by distributional models (Emami et al., 2021; Porada et al., 2021). Considering our task of learning plausible materials for vehicle components, we go beyond selectionally preferred component materials as predicted with high probabilities, aiming also for less frequently observed cases that are nevertheless plausible.

(Domain-specific) MWEs and Compositionality Multiword expressions (MWEs) are challenging for any natural understanding system, given that MWE meanings are idiosyncratic to some degree, i.e., the meaning of an MWE is not entirely (or

even not at all) predictable from the meanings of the constituents (Sag et al., 2002; Reddy et al., 2011; Salehi et al., 2014; Schulte im Walde et al., 2016; Cordeiro et al., 2019; Schulte im Walde and Smolka, 2020). Even though MWEs are ubiquitous not only in general- but also in domain-specific language (Clouet and Daille, 2014; Häty et al., 2021), up to date only few NLP systems have exploited MWE meaning modules, as in machine translation (Cholakov and Kordoni, 2014; Weller et al., 2014). This study is faced with 98% noun compounds among our domain-specific targets, and we test the compound-head compositionality assumption (e.g., a *seat heating switch led* "is a type of" *led* but an *engine valves rocker arm* "is not a type of" *arm*) to fight the severe MWE-triggered data sparsity.

3 Data

Vehicle Component Dataset As targets for our components, we rely on a set of 7,069 unique component names curated by experts from the vehicle repair domain.¹ A component name may denote a tangible physical component such as *cooling blower*, as well as intangible functional and software components such as *ABS warning lamp function* and *road test*. The dataset comprises 155 single-word components and 6,914 multiword components with up to eight constituents (see Fig. 6 in App. A for the distribution). When applying the assumption of MWE compositionality, we fall back to the right-most constituent word as the MWE head (AltaKhaine, 2019), resulting in a total of 725 different heads across the 6,914 multiword components.

Domain-Specific Corpora We utilize two domain-specific English corpora for extracting materials for vehicle components from text: (i) a domain-specific vehicle repair manual written by domain-experts, containing approx. 800K English tokens (henceforth: DOMAIN) and (ii) a portion of the English Wikidata (henceforth: WIKI). For this, we download a Wikidata dump², from which we only keep those 118,154 articles which contain one or more of the 6,914 multiword components in the titles or body (Häty et al., 2020). For domain-adaptation, we delete sentences that do not contain any of the multiword components, resulting in approx. 225K sentences.

¹The Vehicle Component Dataset and the DOMAIN corpus are provided by Bosch.

²<https://dumps.wikimedia.org/enwiki/>, 2.3.2022

4 PLM Experiments

Our PLM experiments investigate to what extent and how PLMs encode domain-specific component materials when prompted with cloze query variants, as illustrated in Fig. 2: We use the vehicle component dataset to construct a set of 504 cloze query templates for each individual component. We first probe a RoBERTa model without any modifications and gather the top-5 predictions for each component, resulting in a total of >2K predictions. To aggregate these predictions, we experiment with three aggregation methods and determine the top-5 most plausible predictions per component and aggregation method for expert evaluation. In a follow-up experiment, we perform domain-adaptation of RoBERTa and DistilRoBERTa for the vehicle repair domain using our two domain-specific corpora, and compare the probing results of the vanilla and the domain-adapted models. The following subsections describe our PLM experiments in detail.

4.1 Cloze Query Prediction and Processing

Our cloze query handling involves three steps (again, cf. Fig. 2): we first construct a set of cloze query templates, then we aggregate model predictions, and finally we analyze the effect of individual cloze queries to select the most meaningful templates for the final top-5 predictions. As Vanilla RoBERTa model (Liu et al., 2019) (henceforth: Vanilla RB), we draw on the `roberta-base` implementation by `huggingface` (Wolf et al., 2020) without any modifications.

Step 1: Cloze Query Template Construction

We develop a set of templates and generate cloze statements for each vehicle component to probe our PLM for plausible component materials: First of all, we define a set of 18 paraphrases typically used to express that a component is made from one or more materials (Davison et al., 2019; Jiang et al., 2021). For this, we start with a set of highly frequent verb relations in our corpus (such as *contain*) and use WordNet synsets (Fellbaum, 1998) to determine additional relevant verbs. We then apply both plural and singular forms of the component nouns with corresponding indefinite articles to define queries (see Table 1). To include prompts that refer to prototypical (and thus presumably highly plausible) materials, we follow Apidianaki and Garí Soler (2021) and integrate the quantifiers {MOST} and {ALL}; we also add the quantifier {MANY}. Furthermore, we define a query




Figure 2: Overview of prediction generation and aggregation, with the numbers in the orange boxes referring to queries/predictions for *one* example component, e.g., *battery*.

Cloze Query Templates	
[context] + [noun_quant.] + indef. article + component + verb_relation + [adverb] + <mask>.	
SINGULAR	a battery consists of <mask>.
PLURAL	batteries consist of <mask>.
noun_quantifier + PLURAL	most batteries consists of <mask>.
adverb + SINGULAR	a battery usually consists of <mask>.
adverb + PLURAL	batteries usually consists of <mask>.
context + SINGULAR	when used in a vehicle, a battery consists of <mask>.
context + PLURAL	when used in a vehicle, batteries consists of <mask>.

Table 1: Cloze query template and filled-in queries for the vehicle component *battery*. SINGULAR and PLURAL are placeholders that are filled in with vehicle component nouns in singular and plural form, respectively.

template element for adverbs referring to typicality, by including the adverbs {USUALLY, GENERALLY, NORMALLY}. Fostering materials that are not necessarily prototypical but still potentially plausible, we use the noun quantifier {SOME} (Apidianaki and Garí Soler, 2021) and the adverbs {POSSIBLY, PLAUSIBLY}. In our domain-specific scenario, humans tend to intuitively restrict the answer space of plausible materials by providing contexts such as “WHEN USED IN A VEHICLE, ...”. Accordingly, we follow recent work investigating the impact of so-called *trigger tokens* (Shin et al., 2020) and *lexical cues* (Misra et al., 2020) to predict a related token, and define six limiting context phrases to leverage intuitive human behaviour.

The above procedure constructs 504 cloze queries for each vehicle component such as the input query “*a battery contains <mask>*.”, where the plausible material candidate is masked and to be predicted by a PLM. A full overview of query variants is presented in Table 1.

Step 2: Cloze Query Prediction Aggregation

For each component, we prompt our PLM with the corresponding 504 cloze queries and obtain the top-5 predicted tokens ranked by probability for each of the queries, resulting in 2,520 predictions per component, see Fig. 2.³ Given that we want only a small list of highly plausible materials, we experiment with three approaches to aggregate the

2,520 predictions for each component such that the most plausible material candidates are ranked at the top of a component’s material list.

- BEST-SCORE probabilities aggregate the most probable PLM-predicted material types from the 2,520 top-5 predictions across all queries. In this way, predictions that the model considers highly probable in a specific query constellation are considered highly plausible, no matter how often they have been predicted.
- AVG probabilities are obtained by summing up the probabilities for each query variant on the material type level and then averaging by the number of queries that predicted that material. This way, probability mass is taken away from materials that were predicted with high probabilities by only one or few queries. The aggregation is thus directed towards materials that are less prominent for individual queries but more pervasive across query variants.
- PREVALENCE refers to a method where the most salient PLM-predicted materials are obtained by ranking only according to how often the material was predicted across all top-5 suggestions of all query variants. Completely ignoring the probabilities in this step accounts for the amount of used cloze queries and provides insight into whether the model leverages semantic information across specific queries.

³We apply basic post-processing as described in App. B.2.

	AVG	BEST-SCORE	PREVALENCE
Total@5	0.08	0.18	0.42
Acc@5	0.36	0.58	0.86

Table 2: Results of aggregation methods on top-5 predicted materials for 100 components each, i.e., a total of 500 predictions per method, reporting total@5 and accuracy in %.

After aggregating the final top-5 predictions by using each of the three presented methods, we analyze⁴ for 100 components (see App. C for sampling details) whether the top-5 material predictions are plausible. We follow previous work (Ettinger, 2020; Apidianaki and Garí Soler, 2021) and evaluate by accuracy, i.e. the proportion of vehicle components for which a suggested material is among the model’s top-5 predictions.

Results are reported in Table 2. The suggested materials ranked by PREVALENCE with over 40% of correct predictions clearly outperform BEST-SCORE and AVG. Especially AVG is observed to promote rather implausible predictions such as *dna*, *food*, or *death*. Based on this preliminary assessment, we use the PREVALENCE method for aggregating the top-5 predictions in our final experiments.

Step 3: Cloze Query Template Selection Using the top-5 material predictions by PREVALENCE, we now investigate the productiveness of query templates for 100 components (see App. C for sampling details) in order to select the most salient ones for our final experiments. Productiveness investigates which queries are *productive*, i.e., contribute a material to the final top-5 materials.

In our first productiveness analysis, we analyze the 68,651 queries that trigger any of the top-5 predictions across all 100 sampled components and all query template variants. Fig. 3 shows the number of queries actively triggering a top-5 prediction in comparison to the overall number of queries that could have potentially been activated (in absolute numbers). Overall, plural variants exhibit a slightly higher productiveness than singular variants.

In our second productiveness analysis, we examine whether some verb relations "do all the work" in comparison to others that are not activated and

⁴The analysis is performed by one of the authors of this paper, who is familiar with the vehicle repair domain. Note that this preliminary assessment is only for comparing aggregation methods, and therefore independent of the final evaluation involving three experts, see Fig. 2.




Figure 3: Query template productiveness across the respective defined query template variants.




Figure 4: Relation productiveness across singular and plural query template variants.

should therefore be excluded for the final experiments. Fig. 4 shows the relations triggering a top-5 prediction across all query variants. The most productive relations for singular and plural variants are [*comprise, formed of, build with*] and [*made up of, composed of, consist of*], respectively. For all but two cases, query productiveness of plural variants is higher than for the singular equivalents.

Our two productiveness analyses suggest that our set of domain-specific query templates represents a well-defined set which is *successful as a whole* in cooperation with a suitable method for aggregating a large number of resulting predictions, and does not include strongly over- or underperforming variants. For our final PLM experiments, we thus leverage the full set of query templates.

4.2 Domain Adaptation of RoBERTa

We experiment with our two domain-specific corpora DOMAIN and WIKI (see §3) to adapt RoBERTa to the vehicle repair domain (henceforth: DOMAIN RB and WIKI RB, respectively). Additionally, we compare RoBERTa to a domain-specific resource-lean model DOMAIN DistilRB, where we adapt DistilRoBERTa (Sanh et al., 2019) to the vehicle repair domain using DOMAIN.

The domain adaptations draw on the `roberta-base` and `distilroberta-base` implementations from `huggingface` (Wolf et al., 2020); we split our input texts into train (90%) and validation (10%) sets. Following common choices, we set the argument specifying the fraction of tokens to mask to 15%. We train with a learning rate of $2e-5$, a weight decay of 0.01, and a batch size of 1,024 (cf. App. B.3 for further details). We train for three epochs and save the best model according to validation set performance. We train one model for each of the aforementioned variants (DOMAIN RB, WIKI RB, and DOMAIN DistilRB). Model results will be presented in §7, in comparison to Basilisk results.

5 Basilisk Experiments: A Classic

In comparison to the PLM experiments, we explore the seminal pattern-based bootstrapped learner Basilisk (Thelen and Riloff, 2002) for learning domain-specific plausible materials for vehicle components. Basilisk is designed to learn high-quality semantic lexicons for one or more categories, provided an unstructured natural language text corpus plus seed words for each semantic category to be learnt. The resource-lean learner leverages extraction patterns representing linguistic contexts, thus exploiting lexico-syntactic structures to capture word meaning.

As the starting step for Basilisk, we select ten seed nouns from the most frequent words in DOMAIN⁵ that belong to the semantic class *material*. Since Basilisk depends on extraction patterns to supply contextual support for additional words belonging to the same semantic class, we rely on dependency parsing to create domain-specific syntactic contextual patterns by using the seed words obtained from DOMAIN and only the DOMAIN corpus. We then apply Basilisk as bootstrapping algorithm to select the best-performing patterns for a pattern pool and subsequently fill a candidate word pool with the extractions of the highest-scoring patterns. We retrieve, for example, `SEED*<GDep>:<compound>:<dependent>:alloy` on the basis of which the word `alloy` gets added to the candidate word pool. We stop Basilisk after n bootstrapping rounds and store the retrieved

⁵Set of seeds: {*water, steel, metal, glass, rubber, plastic, aluminum, copper, polyester, quartz*} In the beginning, we created individual seed word lists for DOMAIN and WIKI, where the latter lacked sufficient quality. Thus, we harness the DOMAIN seed words for both datasets.

semantic lexicon of n plausible materials for a given vehicle component. To then connect these possible materials with the components, we first process the underlying unstructured data sources and only keep sentences containing at least one component name mention. Second, we filter the obtained sentences with the semantic lexicon and store all component-material candidate matches.

In this way, we generate 396,887 extraction patterns, and use the Basilisk implementation (Thelen and Riloff, 2002) with all corresponding patterns and seed words on the DOMAIN dataset. We limit the candidate word pool size to $n = 200$, and generate possible material candidates for both the DOMAIN and the WIKI datasets.

6 Evaluation

To our knowledge, no gold standard is available for the task of plausible material extraction in the vehicle repair domain. We thus perform an expert evaluation of the quality of predictions for each slot obtained from the vanilla and the domain-adapted RoBERTA models. To compare PLM-based results to resource-leaner methods, we also evaluate results from the pattern-based Basilisk algorithm and DistilRoBERTa.

Evaluation Task For each model, we select and evaluate 100 randomly sampled components,⁶ (i) balancing the number of constituents per target component, (ii) providing the full component vs. providing only its head, and (iii) varying the underlying corpus (for Basilisk only). For both 1- and 2-constituent words,⁷ we sample 30 components with 20 full/10 head samples with equal amounts from DOMAIN/WIKI. For 3-constituent words, we sample 20 instances with equal amounts of full/head and DOMAIN/WIKI samples. For both 4- and 5-constituent words, we sample 10 head instances with equal amounts of DOMAIN/WIKI samples (cf. App. C for further details).

We obtain the aggregated top-5 predictions for vehicle components from each model. For Basilisk, we sample up to five answer options whenever more than five entries are extracted for a given component. To minimize annotator bias, we compile a set of up to 18 materials for each vehicle component while keeping track of prediction origins. This way, annotators see each evaluated component only once and rate all predictions from all models for

⁶We make the list of evaluated components available [here](#).

⁷A 1-constituent word is a simplex word.

Model	IAA	A1::A2	A2::A3	A3::A1
Vanilla RB	0.85	0.89	0.81	0.86
DOM. RB	0.79	0.85	0.74	0.78
WIKI RB	0.81	0.86	0.74	0.83
Basilisk	0.79	0.79	0.74	0.83
DOM. DistilRB	0.82	0.88	0.78	0.81

Table 3: Averaged (left) and pair-wise (right) IAA on material predictions.

each vehicle component without potential unconscious comparison to previously seen predictions from another model. In the majority of cases, the models suggested materials either as a singular or as a plural form (e.g., *metal* and *metals*); for the few cases where we are faced with both versions, we decide to not merge the model predictions and ask the annotators to tick both options if both are considered plausible.

Evaluation Setup We present each component instance with corresponding materials to three experts⁸ from the vehicle repair domain. The annotators are asked to rate the plausibility of material predictions for each vehicle component, by ticking the correctly identified materials. They are also provided the options “none of these” if no material is plausible, and “I do not know the answer” if this is the case. Inter-annotator agreement (IAA) is shown in Table 3, with overall average IAA on the left and IAA calculated for each annotator pair on the right. All IAA scores indicate substantial agreement. Further details on the annotation setup and on inter-annotator analyses on both material and component level are provided in App. C.2.

7 Results and Discussion

Table 4 presents results on model performances. On the left, we show accuracy on component level across n annotators. Accuracy is defined as the proportion of components for which at least one of a model’s top-5 material predictions is rated plausible by the annotators. On the right, we display accuracy on material level where accuracy is defined as the proportion of plausible material predictions among a model’s top-5 predictions, as rated by n expert annotators.⁹

⁸The participating annotators are not affiliated with the company providing the Vehicle Component Dataset.

⁹Material accuracy is naturally lower than component accuracy, as we evaluate five material predictions per component. Hence, achieving perfect agreement for material predictions is more difficult than perfect agreement for components.

Model	COMPONENTS		MATERIALS	
	$\geq 1A$	3A	$\geq 1A$	3A
Vanilla RB	0.87	0.68	0.49	0.24
DOMAIN RB	0.93	0.73	0.62	0.28
WIKI RB	0.91	0.66	0.56	0.24
Basilisk	0.73	0.40	0.45	0.14
DOMAIN DistilRB	0.87	0.69	0.53	0.23

Table 4: Model performance. COMPONENT accuracy: proportion of components for which at least one material among a model’s top-5 predictions is rated plausible by n expert annotators nA . MATERIAL accuracy: proportion of material predictions among a model’s top-5 predictions rated plausible by n expert annotators nA .

Domain-Adaption Approaches: Vanilla, DOMAIN, and WIKI RoBERTa RoBERTa adapted to the domain using a small high-quality corpus (DOMAIN RB) beats all other models with an accuracy of 0.93, and RoBERTa adapted to the domain using a customized portion of the English Wikipedia (WIKI RB) is similarly successful (0.91); in comparison, Vanilla RoBERTa reaches an accuracy of 0.87. With perfect agreement among annotators (i.e., 3A), DOMAIN RB is still able to reach an accuracy of 0.73, in comparison to 0.66 for WIKI RB and 0.68 for Vanilla RB. The top results on material level resemble those for component accuracy: DOMAIN RB (0.62) outperforms all other models, followed by WIKI RB (0.56). In contrast to the component level, Vanilla RB is no longer on par with the smaller but domain-adapted DOMAIN DistilRB, while still outperforming Basilisk. The differences between the three resource-intensive approaches (Vanilla RB, DOMAIN RB and WIKI RB) are not significant¹⁰ on component level for nA . On material level, however, differences are significant between PLMs (incl. DOMAIN DistilRB) for $\geq 1A$, but not for 3A. All approaches significantly outperform the resource-lean Basilisk on both levels for all nA .

As a particular example to illustrate the positive impact of domain adaptation, we inspect results for the vehicle component *engine valves rocker arm* from our introductory example (cf. Fig 1). Vanilla RB results include [*wood, metal, steel, bones, legs*], where the predictions *bones* and *legs* refer to (parts of) an extremity of an animate being instead of a vehicle component. The materials [*wood, metal, steel, joints, aluminium*] as obtained from WIKI RB instead encompass the material *aluminium* as

¹⁰Significance tests apply χ^2 , $p < 0.05$.

well as *joints* which refer to connections between body parts as well as vehicle parts. Finally, DOMAIN RB results comprise all plausible materials predicted by the other models, namely [*steel, metal, parts, aluminium, plastic*]. Moreover, the results feature the additional in-domain material *plastic*, while not including out-of-domain predictions such as *wood* anymore. We note that annotators are quite likely to rate one of the top-5 predictions for the same component as plausible, however, material prediction accuracy values indicate that the likelihood to also agree on this specific material drops with an increasing number of annotators.

The findings from these results are twofold. Domain adaptation of a given model results in a clear increase of plausible materials for items within a given domain. Second, if no highly domain-specific and well-curated corpus is available, domain adaptation using WIKI leads to very viable results and therefore represents a strong alternative.

Resource-Lean Approaches Results for the pattern-based algorithm Basilisk significantly underperform all other models for both component (0.73) and material (0.45) accuracy. In contrast, the domain-adapted resource-lean DistilRoBERTa (DOMAIN DistilRB) is on par with Vanilla RB and WIKI RB. Manually analyzing predictions for our example discussed earlier, i.e., the component *engine valves rocker arm*, we find that the Basilisk results [*structures, glass, core, steel, plating*] are potentially within the domain, but often neither a material (*structures, core*) nor rated plausible for the respective target component (*glass, plating*). Results for DOMAIN DistilRB, in contrast, include [*steel, metal, parts, plastic, aluminium*] and basically equal DOMAIN RB predictions.

The comparison of larger vs. leaner and smaller approaches by the examples of the well-established classic Basilisk and the transformer-based DistilRoBERTa model thus clearly demonstrates a preference for the latter.

Compositionality of Domain-specific MWEs

Results regarding the assumption of compositionality on domain-specific MWEs are presented in Fig. 5. While accuracy values are lower when using the right-most head instead of the full multiword component (exception: Basilisk), the average difference is only 0.07 and 0.01 for component and material prediction accuracy, respectively. Note, however, that IAA differs especially on material




Figure 5: Comparison of full and head component performance on component (full: $x/59$, head: $x/40$) and material (full: $x/395$, head: $x/200$) accuracy level.

level with an average agreement of 0.44 using full and 0.32 drawing on head components. This indicates that predictions for full component names tend to be more clearly connected to this very item than predictions retrieved for a head.

We further analyse model performance regarding the influence of MWE constituent number. The results indicate that specificity beats brevity when learning plausible materials for vehicle components. However, our analysis also points towards a threshold separating a beneficial number of constituents (2) from a number pointing towards a detrimental degree of specificity (3+, cf. App. D.2 for further details).

Our findings suggest that domain-specific multiword components are largely compositional and therefore justify a simplified meaning representation via head constituents.

The Influence of Cloze Queries To shed more light on the impact of using cloze query templates for material prediction, we examine whether the number of queries triggering predictions¹¹ for a given component correlates with component accuracy. We find moderate-to-strong correlation values for all our PLMs, thus suggesting that a high number of queries yielding the same prediction indicates the prediction’s plausibility. We also test whether this correlation is directed and fit a corresponding linear regression model. However, the resulting significant R^2 values, ranging around 0.25, are relatively weak and leave us with a mixed picture regarding a directed relationship between the number of queries and component accuracy (cf. App. D.1 for further details).

When zooming into each template part, we find the adverb *GENERALLY* outperforming the other ad-

¹¹An average number of 718 queries is activated for each component across models, with numbers of queries ranging from a minimum of <250 and a maximum of >1.200 queries activated per component.

verbs. Concerning quantifiers, we observe MANY yielding the most activated queries across models, followed by MOST and ALL. Among contexts, we identify WHEN USED IN VEHICLES,... demonstrating best performance. When inspecting the impact of our domain-specific verb relations, we find mostly distinctive groups of verbs accumulating at the extremes of the curve, e.g., *build with*, *consist of*, *made up of* and *involve, made of, include* yielding maximum and minimum numbers of plausible predictions across models and singular/plural variants, respectively. While overall observations suggest that templates in the passive voice produce more plausible predictions overall, examples such as *consist of* underline the importance of a well-rounded template set that is *successful as a whole*.

To further explore the impact of designing high-quality query templates for automatic query construction, we perform the following experiment. We take the five most frequent main verbs *make, say, go, use, take* from ENCOW16AX (Schäfer, 2015) that are considered to be *not* directly related to describing what material an item is made of. We construct a set of cloze queries to probe our best-performing domain-adapted PLM DOMAIN RB and aggregate predictions as described in §4.1. We analyze the final top-5 material predictions and find only 9 correct materials. Instead, we observe a small number of types being predicted for many of the sampled components, as well as a small number of component-specific predictions that might be semantically related, but are no material, e.g., top-5 predictions for the component *fuel tank* include [*sense, noise, hydrogen, oil, cold*]. These findings emphasize that using a set that is (i) comprehensive, (ii) domain-specific, and (iii) syntactically diverse is crucial in order to avoid distortion of results.

Exploring Prediction Pool Size Finally, we explore whether aggregating results from a larger pool of predictions obtained from the model benefits our goal of learning all plausible materials for a component. For this, we use our original set of 504 cloze queries per component to obtain the top-10 instead of previously top-5 predicted tokens from the PLM Vanilla RB, resulting in a set of 5040K predictions per component. We aggregate this doubled amount of predictions as explained in §4.1 and analyze the final top-5 predictions as outlined in §4.1. While component accuracy differs by only 1% (in favor of a smaller pool), drawing from a larger pool seems to benefit learning plausible materials

with an 5% increase of correct predictions when using top-10 model predictions for aggregation (cf. App. D.3 for details). We attribute these gains to the increase in relevant material predictions that are appropriately aggregated using PREVALENCE, thus displacing, for example, non-relevant material predictions such as *wood* from top-k positions.¹²

8 Conclusion

We tackled the task of learning domain-specific plausible materials for components in the vehicle repair domain from a novel perspective by probing SOTA language models in a cloze task style setting to overcome the lack of annotated datasets. Based on a diverse set of semi-automatically constructed cloze queries combined with a suitable aggregation method, we presented a new method to efficiently extract knowledge regarding vehicle components and their materials as acquired by a PLM. While showing that domain-adaptation using either a small, high-quality or a customized Wikipedia corpus boosts performance, we also demonstrated the power of resource-lean alternatives such as the PLM DistilRoBERTa, and found that the bottleneck for domain-adaptation with respect to our task might not be model size but rather corpus quality and suitability. Finally, we successfully exploited the compositionality assumption for our domain-specific multiword expressions as a way to address data sparsity.

Acknowledgements

The authors would like to thank the Robert Bosch GmbH for providing access to the Vehicle Component Dataset and the DOMAIN corpus. We are grateful to Frank Hartmann, Andreas Körber, Patrick Reiser, and Boris Haselbach for valuable discussions and their support with the Bosch Vehicle Component Dataset.

We also thank Prisca Piccirilli, Neele Falk, and the SemRel group for constructive suggestions and feedback on versions of this work. Helena Schlipf was funded by the Robert Bosch GmbH. Annerose Eichel received funding by the Robert Bosch GmbH and the Hanns Seidel Foundation’s Talent Program.

¹²Statistics w.r.t. *wood* as a top-1 material prediction with pool size 5: 0.36% (correct: 0.1%); with pool size 10: 0.23% (correct: 0.1%).

Limitations

When learning plausible materials for components in the vehicle repair domain, we build a varied set of query templates to probe PLMs and seek to aggregate obtained predictions in an optimal way. We are, however, fully aware that a set of queries that is optimal for machines is not necessarily the set that also makes perfect sense to humans.

As far as the transfer of the suggested approach to languages other than English is concerned, we call attention to the potential need to adapt the query templates, e.g., when working with languages that allow for more flexibility in word order such as Polish, Turkish, German, Hindi, or Finnish. Further, while it might be difficult to find well-curated domain-specific corpora in some languages, we show that using a customized version of Wikipedia of moderate size (approx. 225K sentences) in a given language represents a very viable alternative. Additionally, researchers could use a multilingual model or an adapter-based approach to navigate in-/output in other or multiple languages.

Our work focuses on RoBERTa as a backbone model, which has been shown to perform well in cloze task style settings. We also carried out initial experiments with newer models such as ELECTRA, not yielding desired results. We also conducted experiments with generative models such as GPT3. However, the output contained a lot of noise that did not aid our goals. Moreover, the focus in this paper is devoted to the development of a novel contribution regarding the prompting and aggregation techniques as well as analyzing results rather than benchmarking a wide variety of models using an already existing method. While experiments with a wider variety of models represent an interesting future task, high-quality evaluation might be a bottleneck due to the availability and cost factor of domain experts.

Regarding our model predictions, we have not yet attempted to detect and organize potential semantic relations between predictions, such as hypernym/hyponym relations between *metal* and *aluminium*. We leave organizing and utilizing such relations as potentially relevant milestone for connecting and tracing symptoms back to a cause for future work. Furthermore, material predictions are not yet categorized along dimensions such as main and auxiliary materials or surface and inner materials, which might be of interest in the vehicle repair and material science domains.

Finally, we would like to point out that all tested models and algorithms struggle with predicting plausible materials for intangible items which might, however, be more prevalent in other domains. Even if a model is able to predict correct material predictions, rating the plausibility of predictions that are not tangible but rather abstract materials such as *data*, *parameters* and *functions* cannot be considered a trivial task for expert annotators and needs corresponding guidance.

Ethics Statement

Two of the resources used in this work, the Vehicle Component Dataset and one of the two domain-specific corpora (DOMAIN) have been kindly provided to us. To provide insight into the nature of the data and foster reproducibility using comparable data, we make a sample of the Vehicle Component Dataset available (cf. §6). We also investigated comparable publicly available alternatives for DOMAIN and showed that leveraging a portion of the English Wikipedia customized to the domain of interest (WIKI) represents a viable substitute in case no custom corpus is available. Note that Wikipedia text content including Wikipedia dumps is licensed under both the Creative Commons Attribution-ShareAlike 3.0 License and the GNU Free Documentation License.

In the context of our evaluation task, we collected plausibility ratings from human participants. For this, the participants were provided an Informed Consent Letter with the name and the contact of the principal investigators; the title, purpose and procedure of the study; risks, benefits and compensation for participating in the study; confirmation of confidential anonymous data handling; and confirmation that participation in the study is voluntary. The Informed Consent Letter was signed by both the participant and the investigators before the participants took part in the study.

We use and adapt PLMs as provided and licensed under the Apache License 2.0 by `huggingface` (Wolf et al., 2020). We acknowledge that material predictions learnt using the outlined approach are a product of unsupervised learning methods which might be prone to error. We point out that predictions should be approved by an expert or flagged otherwise in case they are used in a downstream application to avoid potential risks including harm of objects or safety risks in case of incorrect repair procedures.

References

- Abdel Rahman Mitib Altakhaineh. 2019. [A cross-linguistic perspective on the Right-Hand Head Rule: The rule and the exceptions](#). *Linguistics Vanguard*, 5(1):20180033.
- Marianna Apidianaki and Aina Garí Soler. 2021. [ALL dolphins are intelligent and SOME are friendly: Probing BERT for nouns' semantic properties and their prototypicality](#). In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 79–94, Punta Cana, Dominican Republic.
- Kostadin Cholakov and Valia Kordoni. 2014. [Better Statistical Machine Translation through Linguistic Treatment of Phrasal Verbs](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 196–201, Doha, Qatar.
- Elizaveta Loginova Clouet and Béatrice Daille. 2014. [Splitting of compound terms in non-prototypical compounding languages](#). In *Proceedings of the First Workshop on Computational Approaches to Compound Analysis*, pages 11–19, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.
- Silvio Cordeiro, Aline Villavicencio, Marco Idiart, and Carlos Ramisch. 2019. [Unsupervised Compositionality Prediction of Nominal Compounds](#). *Computational Linguistics*, 45(1):1–57.
- Joe Davison, Joshua Feldman, and Alexander Rush. 2019. [Commonsense Knowledge Mining from Pretrained Models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 1173–1178, Hong Kong, China.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhिलाsha Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. [Measuring and Improving Consistency in Pretrained Language Models](#). *Transactions of the Association for Computational Linguistics*, 9:1012–1031.
- Ali Emami, Ian Porada, Alexandra Olteanu, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2021. [ADEPT: An Adjective-Dependent Plausibility Task](#). 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 7117–7128, Online.
- Katrin Erk, Sebastian Padó, and Ulrike Padó. 2010. [A Flexible, Corpus-Driven Model of Regular and Inverse Selectional Preferences](#). *Computational Linguistics*, 36(4):723–763.
- Allyson Ettinger. 2020. [What BERT Is Not: Lessons from a New Suite of Psycholinguistic Diagnostics for Language Models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Christiane Fellbaum, editor. 1998. *WordNet – An Electronic Lexical Database*. Language, Speech, and Communication. MIT Press, Cambridge, MA, USA.
- Annemarie Friedrich, Heike Adel, Federico Tomazic, Johannes Hingerl, Renou Benteau, Anika Marusczyk, and Lukas Lange. 2020. [The SOFC-Exp Corpus and Neural Approaches to Information Extraction in the Materials Science Domain](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1255–1268, Online.
- Yoav Goldberg. 2019. [Assessing BERT's Syntactic Abilities](#). *CoRR*, abs/1901.05287.
- Tanishq Gupta, Mohd Zaki, N. M. Anoop Krishnan, and Mausam. 2022. [MatSciBERT: A Materials Domain Language Model for Text Mining and Information Extraction](#). *npj Computational Materials*, 8(1):102.
- Anna HäTTY, Julia Bettinger, Michael Dorna, Jonas Kuhn, and Sabine Schulte im Walde. 2021. [Compound or Term Features? Analyzing Saliency in Predicting the Difficulty of German Noun Compounds across Domains](#). In *Proceedings of the 10th Joint Conference on Lexical and Computational Semantics*, pages 252–262, Bangkok, Thailand (online).
- Anna HäTTY, Dominik Schlechtweg, Michael Dorna, and Sabine Schulte im Walde. 2020. [Predicting Degrees of Technicality in Automatic Terminology Extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2883–2889, Online.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. [How Can We Know When Language Models Know? On the Calibration of Language Models for Question Answering](#). *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). *CoRR*, abs/1907.11692.
- Kanishka Misra, Allyson Ettinger, and Julia Rayz. 2020. [Exploring BERT's Sensitivity to Lexical Cues using Tests from Semantic Priming](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4625–4635, Online.
- Sheshera Mysore, Zachary Jensen, Edward Kim, Kevin Huang, Haw-Shiuan Chang, Emma Strubell, Jeffrey Flanagan, Andrew McCallum, and Elsa Olivetti. 2019. [The Materials Science Procedural Text Corpus: Annotating Materials Synthesis Procedures with Shallow Semantic Structures](#). In *Proceedings of the 13th Linguistic Annotation Workshop*, pages 56–64, Florence, Italy.

- Sheshera Mysore, Edward Kim, Emma Strubell, Ao Liu, Haw-Shiuan Chang, Srikrishna Kompella, Kevin Huang, Andrew McCallum, and Elsa Olivetti. 2017. [Automatically Extracting Action Graphs from Materials Science Synthesis Procedures](#). *CoRR*, abs/1711.06872.
- Anmol Nayak and Hariprasad Timmapathini. 2021. [Wiki to automotive: Understanding the distribution shift and its impact on named entity recognition](#). *CoRR*, abs/2112.00283.
- Tim O’Gorman, Zach Jensen, Sheshera Mysore, Kevin Huang, Rubayyat Mahbub, Elsa Olivetti, and Andrew McCallum. 2021. [MS-Mentions: Consistently Annotating Entity Mentions in Materials Science Procedural Text](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1337–1352, Online and Punta Cana, Dominican Republic.
- Elsa A. Olivetti, Jacqueline M. Cole, Edward Kim, Olga Kononova, Gerbrand Ceder, Thomas Yong-Jin Han, and Anna M. Hiszpanski. 2020. [Data-driven materials research enabled by natural language processing and information extraction](#). *Applied Physics Reviews*, 7(4):041317.
- Lalchand Pandia and Allyson Ettinger. 2021. [Sorting through the noise: Testing robustness of information processing in pre-trained language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1583–1596, Online and Punta Cana, Dominican Republic.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language Models as Knowledge Bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 2463–2473, Hong Kong, China.
- Ian Porada, Kaheer Suleman, and Jackie Chi Kit Cheung. 2019. [Can a Gorilla Ride a Camel? Learning Semantic Plausibility from Text](#). In *Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing*, pages 123–129, Hong Kong, China.
- Ian Porada, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2021. [Modeling Event Plausibility with Consistent Conceptual Abstraction](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1732–1743, Online.
- Seema Rani and Mukesh Kumar. 2021. [Topic modeling and its applications in materials science and engineering](#). *Materials Today: Proceedings*, 45:5591–5596. Second International Conference on Aspects of Materials Science and Engineering.
- Abhilasha Ravichander, Eduard Hovy, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2020. [On the Systematicity of Probing Contextualized Word Representations: The Case of Hypernymy in BERT](#). In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 88–102, Barcelona, Spain (Online).
- Siva Reddy, Diana McCarthy, and Suresh Manandhar. 2011. [An Empirical Study on Compositionality in Compound Nouns](#). In *Proceedings of the 5th International Joint Conference on Natural Language Processing*, pages 210–218, Chiang Mai, Thailand.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2021. [A Primer in BERTology: What We Know About How BERT Works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Ivan A. Sag, Timothy Baldwin, Francis Bond, Ann Copestake, and Dan Flickinger. 2002. [Multiword Expressions: A Pain in the Neck for NLP](#). In *Proceedings of the Conference on Intelligent Text Processing and Computational Linguistics*, Mexico City, Mexico.
- Bahar Salehi, Paul Cook, and Timothy Baldwin. 2014. [Using Distributional Similarity of Multi-way Translations to Predict Multiword Expression Compositionality](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 472–481, Gothenburg, Sweden.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter](#). *CoRR*, abs/1910.01108.
- Roland Schäfer. 2015. [Processing and querying large web corpora with the COW14 architecture](#). Proceedings of the 3rd Workshop on Challenges in the Management of Large Corpora, Lancaster, 20 July 2015, pages 28 – 34, Mannheim.
- Timo Schick and Hinrich Schütze. 2020. [Rare Words: A Major Problem for Contextualized Embeddings and How to Fix it by Attentive Mimicking](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8766–8774.
- Sabine Schulte im Walde, Anna Häddy, and Stefan Bott. 2016. [The Role of Modifier and Head Properties in Predicting the Compositionality of English and German Noun-Noun Compounds: A Vector-Space Perspective](#). In *Proceedings of the 5th Joint Conference on Lexical and Computational Semantics*, pages 148–158, Berlin, Germany.
- Sabine Schulte im Walde and Eva Smolka, editors. 2020. [The Role of Constituents in Multi-Word Expressions: An Interdisciplinary, Cross-Lingual Perspective](#). Number 4 in *Phraseology and Multiword Expressions*. Language Science Press, Berlin, Germany.

- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235, Online.
- Wilson L. Taylor. 1953. “Cloze Procedure”: A New Tool for Measuring Readability. *Journalism Quarterly*, 30(4):415–433.
- Michael Thelen and Ellen Riloff. 2002. [A Bootstrapping Method for Learning Semantic Lexicons using Extraction Pattern Contexts](#). In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*, pages 214–221.
- Vineeth Venugopal, Sourav Sahoo, Mohd Zaki, Manish Agarwal, Nitya Nand Gosvami, and N. M. Anoop Krishnan. 2021. [Looking through glass: Knowledge discovery from materials science literature using natural language processing](#). *Patterns*, 2(7):100290.
- Su Wang, Greg Durrett, and Katrin Erk. 2018. [Modeling Semantic Plausibility by Injecting World Knowledge](#). 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 303–308, New Orleans, Louisiana.
- Marion Weller, Fabienne Cap, Stefan Müller, Sabine Schulte im Walde, and Alexander Fraser. 2014. [Distinguishing Degrees of Compositionality in Compound Splitting for Statistical Machine Translation](#). In *Proceedings of the 1st Workshop on Computational Approaches to Compound Analysis*, pages 81–90, Dublin, Ireland.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [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.

A Data

Dataset Distribution Fig. 6 shows the distribution of multiword component names in the Vehicle Component Dataset with the absolute number of vehicle components per number of constituent, e.g., the dataset comprises 1,563 multiword components with 2 constituents.




Figure 6: Distribution of multiword components in the Vehicle Component Dataset (in absolute numbers).

B Cloze Query Prediction and Processing

B.1 Cloze Query Templates

Verb Relation Variables consist of, comprise, contain, be formed by, be formed of, be made up of, be made up from, be made of, be made out of, be composed of, be manufactured from, encompass, entail, include, involve, incorporate, be build of, be build with

Cloze Query Statements Table 1 presents an overview of the defined cloze query templates and provides an example of filled-in prompts for the vehicle component *lamp*.

B.2 Post-Processing

PLM-based Predictions Having created a given set of cloze queries, we feed the queries into a PLM and obtain the top-5 predictions. Based on our goals, we apply the following post-processing steps. We lower-case the predictions and filter out numbers, punctuation, 1-character tokens, as well as a list of standard stopwords obtained from `spacy` with nine customized stopwords including non-informative tokens such as *material* and *component*. Additionally, spelling variants such as *aluminum* and *aluminium* are merged. We make sure that a predicted token does not equal the singular or plural form of the full input component.

Basilisk-based Predictions We obtain n predictions for each vehicle component using DOMAIN or WIKI for the Basilisk algorithm, and apply post-processing as described for the PLM predictions.

B.3 PLM and Basilisk Experiments: Computing Infrastructure and Model Information

As Vanilla RoBERTa model, we use the `roberta-base` implementation from `huggingface` (Wolf et al., 2020) that comes with 125M parameters. We adapt this model to the vehicle repair domain. As DistilRoBERTa implementation, we use the `distilroberta-base` implementation from `huggingface` (Wolf et al., 2020) which has 82M parameters. See the [official huggingface documentation](#) for further details. For all final experiments, including obtaining predictions from the various models, we use a single NVIDIA RTX A600 GPU. We adapted our models once (DOMAIN RoBERTa, WIKI RoBERTa, and DOMAIN DistilRoBERTa) and used the best model according to validation set performance, i.e. our domain-adapted models are the product of a single fine-tuning run.

As a Basilisk implementation, we use the Java source code as provided by Thelen and Riloff (2002). We use `spacy` (Honnibal et al., 2020) for text preprocessing and dependency parsing.

C Evaluation

C.1 Sampling Vehicle Components

To test whether the head of a MWE preserves the meaning of the full MWE in our domain-specific data, we compare extracted material predictions using only the head of the MWE vs. all constituents (full). To evaluate results, we apply the following procedure when sampling vehicle components consisting of 2+ words. Full components are sampled as is along with the extracted results. Whenever only the head of a component is used, the results for only the head component are obtained. For the evaluation task, the head component is mapped back to the original full component from the vehicle information dataset. For example, during sampling the head component *sensor* gets mapped to the full 2-constituent component *pressure sensor*. In this way, we can evaluate whether results extracted based on only the head *sensor* are rated as plausible also for the more domain-specific full multiword expression *pressure sensor*.

For the Basilisk-based approach, no results are extracted for one item of the 2-constituent full components and five items of the 3-constituent full components. These results were thus expanded with head component extraction results. In one case, stopword removal leads to an empty prediction list (*alternator start charge current reduction, []*, using DOMAIN). The prediction is back-filled with a randomly sampled token from WIKI-based results for the head component (*reduction, [lithium]*, using WIKI).

C.2 Expert Evaluation

Evaluation Task Setup The evaluation task was carried out in a remote setting using Google Forms. Annotators were provided detailed written guidelines including example questions and borderline decisions. In case of questions, annotators had the option to contact the authors of the paper. The evaluation could be completed flexibly in the course of a week. Annotators could take as much as time as they needed for completing the evaluation (average time effort: ~1.5 hours). Our three recruited annotators are based in Germany, male, and speak German as their first language. All annotators have 10+ years knowledge and continued education and training in English in general, as well as profound work experience with vehicle repair domain data in German and English. Annotators received a representation allowance for their effort. Each annotator submitted one unique set of answers.

The collected data does not include any information that names or uniquely identifies individual people or offensive content. We check for this by (i) not giving out data to annotators containing such information, and (ii) not asking for any of this information when collecting ratings regarding plausible material candidates for vehicle components, and (iii) manually inspecting collected data to confirm anonymity of annotators and potential other entities. Letters of Consent (cf. §8) are signed before participation and stored separately from the collected ratings.

Evaluation Results Analysis We evaluate a total of 2,342 material predictions from five different models and algorithms on a set of 100 vehicle components from the vehicle repair domain. For 99 components, we collect an average of 5.6 plausible material predictions valid for one or more models. For exactly one component, all annotators agree on *I do not know the answer*.

Pair-wise inter-annotator agreement (IAA) on the level of vehicle components is shown in Table 5. We calculate the percentage of components where two given annotators agree on *at least one* material being plausible (i.e. accuracy on inter-annotator level) or where two given annotators agree that at least one material prediction is **not** plausible. This includes instances where up to four predictions are not plausible as well as instances where all five predicted material options are not rated as plausible (*none of these*). We exclude instances where one or both annotators agree on the answer option *I do not know the answer*.

Model	IAA	A1::A2	A2::A3	A3::A1
Vanilla RB	1.00	1.00	1.00	1.00
DOMAIN RB	1.00	0.99	1.00	1.00
WIKI RB	1.00	1.00	1.00	1.00
Basilisk	0.93	0.97	0.91	0.91
DOM. DistilRB	1.00	1.00	1.00	1.00

Table 5: Proportion of components where two *Ans* agree on at least one plausible material or where *Ans* agree on at least one material prediction being *not* plausible.

D Results and Discussion

D.1 Statistical Analysis: Number of Cloze Queries and Component Accuracy

We statistically analyze the relationship between the number of cloze queries triggering predictions for a given component and the actual accuracy values on component level.¹³ For this, we first perform a correlation analysis to see whether the number of activated queries is correlated to component accuracy, i.e., whether 1+ of the predictions for this component is rated plausible by 1+ expert annotator. Results are reported in Table 6 with the investigated models on the left and Pearson’s correlation coefficients in the first column. The obtained results indicate that the relationship between the number of activated queries and observed component accuracy is strongly correlated (coefficient >0.5) for the Vanilla and WIKI RB models, and moderately correlated (coefficient >0.3) for DOMAIN RB and DOMAIN DistilRB.

To see whether these correlations are in fact directed, we employ linear regression modeling with the number of queries as predictor and component accuracy as outcome variable. Results are presented in Table 6 with investigated models on

¹³As no cloze queries are generated for the Basilisk algorithm, we only investigate PLM results.

Model	Pearson’s Coefficients	R^2
Vanilla RB	0.51***	0.26***
DOMAIN RB	0.49***	0.24***
WIKI RB	0.58***	0.34***
DOM. DistilRB	0.40***	0.16***

Table 6: Pearson’s correlation coefficients for the number of activated queries and component accuracy. All coefficients are significant with a p-value < 0.01 (***)

the left and R^2 values in the second column. The relatively low R^2 values suggest that the relationship between the number of activated queries and observed *component* (\neq material) accuracy might more strongly rely on other factors than the number of queries activated for a given component.

D.2 n -Constituent Component Analysis

Results regarding MWEs with n constituents are shown in Fig. 7.




Figure 7: Results regarding n -constituent performance on component accuracy level (in %).

Model performance is presented on component accuracy level. It can be observed that 2-constituent components receive the highest number of plausible material annotations, followed by simplex components, consisting of only one constituent. 4-constituent components outperform the shorter 3-constituent and the longest investigated components, namely 5-constituent components. The results indicate that (domain-)specificity beats brevity when learning plausible materials for vehicle components. However, our findings also suggest that there might be a threshold separating a beneficial number of constituents (2) from a number pointing towards a detrimental degree of specificity (3+).

D.3 Prediction Pool Size Analysis

We present results regarding the size of the pool from which predictions are drawn for aggregation in Table 7. The percentage of correct predictions per top-k predictions are depicted as Hits@1-5. Here, using a larger pool leads to a gain of 18% more plausible materials as the top-1 prediction.

	Top-10 Pool	Top-5 Pool
Hits@1	0.54	0.37
Hits@2	0.45	0.56
Hits@3	0.45	0.42
Hits@4	0.50	0.42
Hits@5	0.41	0.31
Comp. Acc@5	0.85	0.86
Total Hits@5	0.47	0.42

Table 7: Correct predictions (Hits@k) per top-k predictions ($\times/100$), component accuracy ($\times/100$), and total correct predictions per top-5 predictions (Total Hits@5).

While accuracy on component level stays basically the same (top-10 pool: 0.85%, top-5 pool: 0.86%), an increase of 5% can be observed when building the final top-5 predictions from a larger pool of predictions.